



HPE ATP – Server Solutions V4

Official Certification Study Guide (Exams HPE0-S52
and HPE0-S53)

HPE Partner Ready Certification and Learning

Master all HPE0-S52 exam questions

Assess your knowledge with chapter quizzes

Understand key concepts and objectives



Hewlett Packard
Enterprise

HPE ATP – Server Solutions V4

OFFICIAL CERTIFICATION STUDY GUIDE

(EXAMS HPE0-S52 AND HPE0-S53)

First Edition

Ken Radford

HPE Press
660 4th Street, #802
San Francisco, CA 94107

HPE ATP – Server Solutions V4

Official Certification Study Guide (Exams HPE0-S52 and HPE0-S53)

Ken Radford

© 2018 Hewlett Packard Enterprise Development LP,

Published by:

Hewlett Packard Enterprise Press

660 4th Street, #802

San Francisco, CA 94107

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

ISBN: **978-1-942741-84-8**

WARNING AND DISCLAIMER

This book provides information about the topics covered in the HPE ATP – Building HPE Server Solutions V4 (HPE0-S52) certification exam. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an “as is” basis. The author, and Hewlett Packard Enterprise Press, shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the author and are not necessarily those of Hewlett Packard Enterprise Press.

TRADEMARK ACKNOWLEDGEMENTS

All third-party trademarks contained herein are the property of their respective owner(s).

GOVERNMENT AND EDUCATION SALES

This publisher offers discounts on this book when ordered in quantity for bulk purchases, which may include electronic versions. For more information, please contact U.S. Government and Education Sales 1-855-447-2665 or email sales@hpepressbooks.com.

Feedback Information

At HPE Press, our goal is to create in-depth reference books of the best quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the expertise of members from the professional technical community.

Readers' feedback is a continuation of the process. If you have any comments regarding how we could improve the quality of this book, or otherwise alter it to better suit your needs, you can contact us through email at hpepress@epac.com. Please make sure to include the book title and ISBN in your message.

We appreciate your feedback.

Publisher: Hewlett Packard Enterprise Press

HPE Contributors: Wilfred Brown, Jeff Holderfield, Chris Powell, Sam Morgan, and Chris Bradley.

HPE Press Program Manager: Michael Bishop

About the Author

Ken Radford is a certified HPE Master ASE – Advanced Server Solutions Architect V3, HPE ASE - Synergy Solutions Integrator, HPE ASE - Hybrid Infrastructure and Cloud Architect, HPE ASE – Converged Infrastructure Architect, HPE ATP – Cloud Administrator, and HPE Certified Instructor. He was author of the Integrating HPE Synergy Solutions official certification study guide, and courseware developer for the Designing Hybrid Infrastructure and Cloud Solutions Rev. 17.21 and Integrating HPE Synergy Solutions, Rev. 17.31 courses.

Introduction

This study guide helps you prepare for the HP ATP – Building HPE Server Solutions V4 certification exam (HPE0-S52). Organized along the lines of exam topics and objectives, chapters can be studied independently when preparing for certification. If you already hold the Building HPE Server Solutions V3 certification and want to acquire the Building HPE Server Solutions V4 certification, this guide also covers the topics in the Building HPE Server Solutions (HPE0-S53) delta exam.

This guide provides a solid foundation for recommending HPE server products and solutions. The exam covers such topics as functions, features, and benefits of these products and solutions, the HPE server solution marketplace, and basic planning and design concepts.

Certification and Learning

Hewlett Packard Enterprise Partner Ready Certification and Learning provides end-to-end continuous learning programs and professional certifications that can help you open doors and succeed in your career as an IT professional. HPE provides continuous learning activities and job-role based learning plans to help you keep pace with the demands of the dynamic, fast paced IT industry; professional sales and technical training and certifications to give you the critical skills needed to design, manage and implement the most sought-after IT disciplines; and training to help you navigate and seize opportunities that enable business advantage today.

As a Partner Ready Certification and Learning certified member, your skills, knowledge, and real-world experience are recognized and valued in the marketplace. To continue your professional and career growth, you have access to our large HPE community of world-class IT professionals, trend-makers and decision-makers. Share

ideas, best practices, business insights, and challenges as you gain professional connections globally.

To learn more about HPE Partner Ready Certification and Learning certifications and continuous learning programs, please visit

<http://certification-learning.hpe.com>

Audience

This guide is suitable for anyone seeking an understanding of foundational level HPE server technology.

Assumed Knowledge

It is assumed that you have a basic understanding of servers, storage, networking, and server management technologies.

Minimum Qualifications

There are no prerequisites for the HP ATP – Building HPE Server Solutions exam (HPE0-S52). However, it is strongly recommended that you have foundational knowledge of HPE server solutions and underlying technologies, and that you have taken the recommended trainings.

The specific requirements for each certification are available on the HPE Partner Ready Certification and Learning website.

Relevant Certifications

After you pass these exams, your achievement may be applicable toward more than one certification. To determine which certifications can be credited with this achievement, log in to The Learning Center and view the certifications listed on the exam's More Details tab. You might be on your way to achieving additional certifications.

Preparing for the Certification Exam (HPE0-S52)

This self-study guide does not guarantee that you will have all the knowledge you need to pass the HP ATP – Building HPE Server Solutions V4 (HPE0-S52) exam. It is expected that you will also need to draw on real-world experience and would benefit

from completing the hands-on activities delivered in the instructor-led training.

Preparing for the Delta Exam (HPE0-S53)

The HP ATP – Building HPE Server Solutions V4 delta exam (HPE0-S53) is available for those candidates who already hold the HP ATP – Building HPE Server Solutions V3 certification and who wish to upgrade their certification to V4. It is recommended that candidates read this study guide in its entirety to prepare for the delta exam.

Recommended HPE Training

Recommended training to prepare for each exam is accessible from the exam's page in The Learning Center. See the exam attachment, "Supporting courses," to view and register for the courses.

Obtain Hands-on Experience

You are not required to take the recommended, supported courses, and completion of training does not guarantee that you will pass the exams. Hewlett Packard Enterprise strongly recommends a combination of training, thorough review of courseware and additional study references, and sufficient on-the-job experience prior to taking an exam.

Exam Registration

To register for an exam, go to http://certification-learning.hpe.com/tr/certification/learn_more_about_exams.html

CONTENTS

1 **Transforming IT Strategy**

Prelearning check

The Idea Economy is here

Creating disruptive waves of new demands and opportunities

Faster application development enables accelerated innovation

Long time to value is costly

IT must become a value creator that bridges the old and the new

HPE Transformation Solutions

Applications and data fuel the digital enterprise

HPE Transformation Solutions

Accelerating time to value

Hybrid IT

Learning check

HPE Composable Infrastructure

The problem with bimodal computing

A new category of infrastructure is needed to power the Idea Economy

Moving to Composable Infrastructure

HPE Composable Infrastructure vision

Composable Infrastructure core principles

Learning check

Customer requirements

Introducing a new generation compute experience

HPE Gen10 servers—Relevant and differentiating innovations

Learning check

Summary

2 **Security**

Prelearning check

Gen10 servers—Security

Can you trust your hardware vendor?

HPE Security focus

HPE Secure Compute Lifecycle

HPE secure server management from anywhere, anytime

[Secure Start—Featuring Silicon Root of Trust](#)

[Learning check](#)

[Secure recovery](#)

[Security built into every level](#)

[Modes for HPE Secure Compute Lifecycle](#)

[New secure server options for Gen10](#)

[Server network interface card adapter security](#)

[Learning check](#)

[HPE Gen10 server physical security features](#)

[Gen10 Trusted Platform Module](#)

[Server intrusion detection](#)

[HPE rack and power infrastructure](#)

[Learning check](#)

[Summary](#)

3 [HPE Server Technologies and Tools](#)

[Prelearning check](#)

[HPE ProLiant Gen10 server innovation](#)

[Features of ProLiant servers](#)

[Intel® Xeon® Processor Scalable Family](#)

[Learning check](#)

[HPE IST](#)

[Learning check](#)

[HPE Gen10 DDR4 memory technologies](#)

[HPE Smart Array controllers and storage](#)

[Activity: Features of HPE Smart Array controller](#)

[HPE Server Storage portfolio](#)

[HPE PCIe Workload Accelerators for ProLiant servers](#)

[HPE ProLiant server solution technologies](#)

[Gen10 power strategy](#)

[Learning check](#)

[HPE server networking](#)

[HPE configuration tools](#)

[Activity: Features of HPE Power Advisor](#)

[Learning check](#)

[HPE server management technologies](#)

[On-system management](#)

HPE iLO 5

Activity: Exploring the iLO 5 interface

Activity: HPE iLO 5 license installation

Intelligent Provisioning

Activity: HPE IP

On-system management with UEFI

Gen 10 UEFI enhancements

Activity: HPE Management Tools

Activity: UEFI Shell

On-premises management with HPE OneView

On-cloud management with HPE Insight Online

Learning check

Summary

4 HPE ProLiant Tower Servers

Prelearning check

HPE ProLiant tower servers

HPE ProLiant ML servers

HPE ProLiant ML350 Gen10

Storage options

ProLiant ML110 Gen10 Server

HPE ProLiant ML350 and ML110 Gen10 comparison

HPE ProLiant ML350 Gen10 to Gen9 comparison

HPE ProLiant MicroServer Gen10

ClearOS on HPE ProLiant

HPE ProLiant Gen10 workload-based positioning

Learning check

ProLiant ML10 Gen9 Server

ProLiant ML10 Gen9 front view

ProLiant ML10 Gen9 rear view

ProLiant ML30 Gen9 Server

ProLiant ML110 Gen9 Server

ProLiant ML150 Gen9 Server

Learning check

Activity: Matching ProLiant tower servers to customer needs

Activity debrief

Learning check

[Activity: Matching ProLiant tower servers to customer needs](#)

[Summary](#)

5 [HPE Rack Server Solutions](#)

[Prelearning check](#)

[HPE ProLiant rack servers](#)

[HPE ProLiant Gen10 servers](#)

[ProLiant Gen10 rack server portfolio](#)

[ProLiant server product numbers](#)

[HPE ProLiant DL360 Gen10 Server](#)

[Learning check](#)

[HPE ProLiant DL380 Gen10 Server](#)

[Activity: Using HPE QuickSpecs](#)

[HPE ProLiant DL560 Gen10 Server](#)

[HPE ProLiant DL560 Gen10—Front system detail](#)

[HPE ProLiant DL580 Gen10 Server](#)

[HPE ProLiant Gen10 workload-based positioning](#)

[Learning check](#)

[HPE ProLiant DL385 Gen10](#)

[HPE ProLiant Gen9 rack server portfolio](#)

[HPE ProLiant DL20 Gen9 Server](#)

[HPE ProLiant DL60 Gen9 Server: Internal view](#)

[HPE ProLiant DL80 Gen9 Server](#)

[HPE SimpliVity 380 Server](#)

[Hyperconvergence solves common business problems](#)

[HPE Intelligent Series Racks](#)

[Learning check](#)

[Summary](#)

6 [High Performance Computing Solutions](#)

[Prelearning check](#)

[Supercomputing and HPC](#)

[HPE purpose-built portfolio for HPC](#)

[HPE Apollo 2000—Enterprise bridge to HPC](#)

[HPE Apollo 2000—r2800 Chassis—24 SFF](#)

[Apollo 2000 system offerings](#)

[HPE Apollo 2000—r2800 Chassis with four 1U nodes](#)

[HPE Apollo 2000—r2800 Chassis with two 2U nodes](#)

[HPE Apollo 2000—r2800 chassis fans](#)

[HPE Apollo 2000—r2800 RCM Module](#)

[HPE ProLiant XL170r and XL190r Server—Gen 10](#)

[HPE Apollo 2000—XL170r Gen10](#)

[HPE Apollo 2000—XL190r Gen10](#)

[Learning check](#)

[HPE Apollo 10 Series](#)

[HPE Apollo 20 Family—k120](#)

[HPE Apollo 40 Family—sx40 Gen10](#)

[HPE Apollo 40 Family—pc40 Gen10](#)

[Learning check](#)

[HPC use cases](#)

[Why GPU and coprocessor acceleration?](#)

[DL and GPU market—Projections](#)

[GPU accelerators—HPC and DL](#)

[Using HPE Apollo 2000 for storage](#)

[HPE Apollo 2000 networking choices](#)

[Learning check](#)

[Summary](#)

7 [HPE BladeSystem Solutions](#)

[Prelearning check](#)

[HPE BladeSystem](#)

[BladeSystem meets data center challenges](#)

[HPE BladeSystem—The ultimate converged infrastructure](#)

[BladeSystem Gen10 and Gen9 solutions](#)

[BladeSystem with HPE OneView](#)

[ProLiant BL server blades](#)

[Scalable blade form factors and device bays](#)

[The ProLiant BL server portfolio](#)

[ProLiant BL460 Gen10](#)

[Learning check](#)

[BladeSystem enclosure](#)

[BladeSystem c7000 enclosure](#)

[OA module](#)

[Insight Display](#)

[HPE Thermal Logic technologies](#)

[Learning check](#)

[BladeSystem interconnect options](#)

[HPE BladeSystem signal midplane](#)

[Mezzanine and FlexLOM cards](#)

[HPE Dynamic Smart Array P204i SR Gen10 Controller](#)

[HPE Dynamic Smart Array P408e-m SR Gen10 Controller](#)

[Learning check](#)

[BladeSystem Virtual Connect modules](#)

[Activity: Assessing HPE BladeSystem Innovations: 20 Gb adapter](#)

[VC network configuration](#)

[VC firmware](#)

[Learning check](#)

[Summary](#)

8 **[HPE Synergy](#)**

[Prelearning check](#)

[HPE Synergy](#)

[Emergence of bi-modal IT](#)

[The industry's first platform architected for composability](#)

[Composable Infrastructure](#)

[Learning check](#)

[Synergy components](#)

[Synergy Composer](#)

[HPE Synergy Image Streamer](#)

[Synergy compute module portfolio](#)

[HPE Synergy 660 Gen10 Compute Module](#)

[HPE Synergy 620 and HPE Synergy 680 Gen9 Compute Modules](#)

[Run any workload better with HPE Synergy Gen10](#)

[Learning check](#)

[HPE Synergy 12000](#)

[Synergy 12000 frame front view](#)

[Synergy 12000 frame rear view](#)

[HPE Synergy fabric configuration](#)

[Managing HPE Synergy](#)

[Synergy management subsystem](#)

[Management architecture](#)

[Reliable, efficient multi-frame management](#)

[Learning check](#)

[HPE Synergy fabric and ICMs](#)

[Synergy frictionless scaling](#)

[HPE Synergy 25/50 Gb Converged Network Adapter](#)

[HPE Synergy 10/20 Gb Converged Network Adapters](#)

[Synergy fabric portfolio choices](#)

[Mellanox SH2200 Switch Module for HPE Synergy](#)

[Synergy storage](#)

[HPE Smart Array P416ie-m SR Gen10 controller](#)

[Remote support in Synergy via embedded HPE OneView](#)

[Learning check](#)

[Summary](#)

9 **[HPE On-system Management](#)**

[Prelearning check](#)

[Preparing a ProLiant system for deployment](#)

[Installing hardware](#)

[Installing DIMMs](#)

[Gen10 DIMM installation guidelines for a single processor](#)

[Configuring or upgrading other components](#)

[Environmental standards](#)

[Learning check](#)

[At server management](#)

[iLO Service Port](#)

[Learning check](#)

[HPE iLO 5 management technologies](#)

[Setting up and connecting to iLO](#)

[ProLiant boot process and server configuration](#)

[Server boot process](#)

[BIOS/Platform Configuration](#)

[Performing shell operations in UEFI](#)

[BIOS RESTful API moves to industry-standard Redfish](#)

[Learning check](#)

[Viewing and updating firmware from the iLO 5 web interface](#)

[iLO manual reboot](#)

[HPE on-system management tools](#)

[HPE Intelligent Provisioning](#)
[Intelligent Provisioning startup menu](#)
[Home screen](#)
[Intelligent Provisioning Perform Maintenance options](#)
[Intelligent Provisioning OS installation](#)

[Learning check](#)

[HPE Smart Update solution stack](#)

[HPE Smart Update Tools](#)

[HPE SUM](#)

[HPE SPP](#)

[HPE Firmware update technology](#)

[HPE RESTful Interface Tool](#)

[Activity: Updating firmware with SUM](#)

[Activity debrief](#)

[Agentless Management](#)

[Preboot System Health Summary](#)

[Active Health System](#)

[Server storage configuration utilities](#)

[SSA](#)

[SSA diagnostics](#)

[Performing a diagnostic task](#)

[Activity: Creating logical volumes using the SSA GUI](#)

[Activity debrief](#)

[Activity: Creating logical volumes using the RESTful Interface Tool](#)

[Activity debrief](#)

[Learning check](#)

[Summary](#)

10 [HPE On-premises and On-cloud Management](#)

[Prelearning check](#)

[HPE server management](#)

[On-premises management with HPE OneView](#)

[Designed for simplicity](#)

[HPE OneView: Your infrastructure automation engine](#)

[HPE OneView architecture](#)

[HPE OneView resource managers](#)

[REST API](#)

[SCMB](#)

[Introduction to server profiles](#)

[Adding enclosures to HPE OneView](#)

[Learning check](#)

[Activity: Creating server profiles in HPE OneView](#)

[Activity debrief](#)

[Easy access to reporting of inventory and health status](#)

[Near real-time search feature](#)

[Learning check](#)

[HPE OneView partner integrations](#)

[HPE OneView for Microsoft System Center](#)

[HPE OneView for VMware vCenter](#)

[HPE OneView for RHEV](#)

[Remote management](#)

[Remote management with HPE iLO Advanced](#)

[HPE Insight Online](#)

[Insight Online and Remote Support integration](#)

[Target markets of direct connect for remote support](#)

[Learning check](#)

[Summary](#)

[**Learning Check Answers**](#)

11 [Practice Test](#)

[Introduction](#)

[Ideal candidate for this exam](#)

[Exam details](#)

[HPE0-S52 testing objectives](#)

[Test preparation questions and answers](#)

[Questions](#)

[Answers](#)

1 Transforming IT Strategy

WHAT IS IN THIS CHAPTER FOR YOU?

After completing this chapter, you should be able to:

Identify how the approach to IT strategy is changing in the Idea Economy.

Describe the HPE Composable Infrastructure strategy.

Explain how the HPE Transformation Solutions reflect what customers consider most important.

Describe the HPE ProLiant Gen10 Compute Experience values.

Prelearning check

A customer is considering migrating their data center to a Cloud Service Provider. They are concerned about the risks involved and would rather have the ability to host some of their applications on-premises and some in the cloud.

How should you respond?

The Idea Economy is here

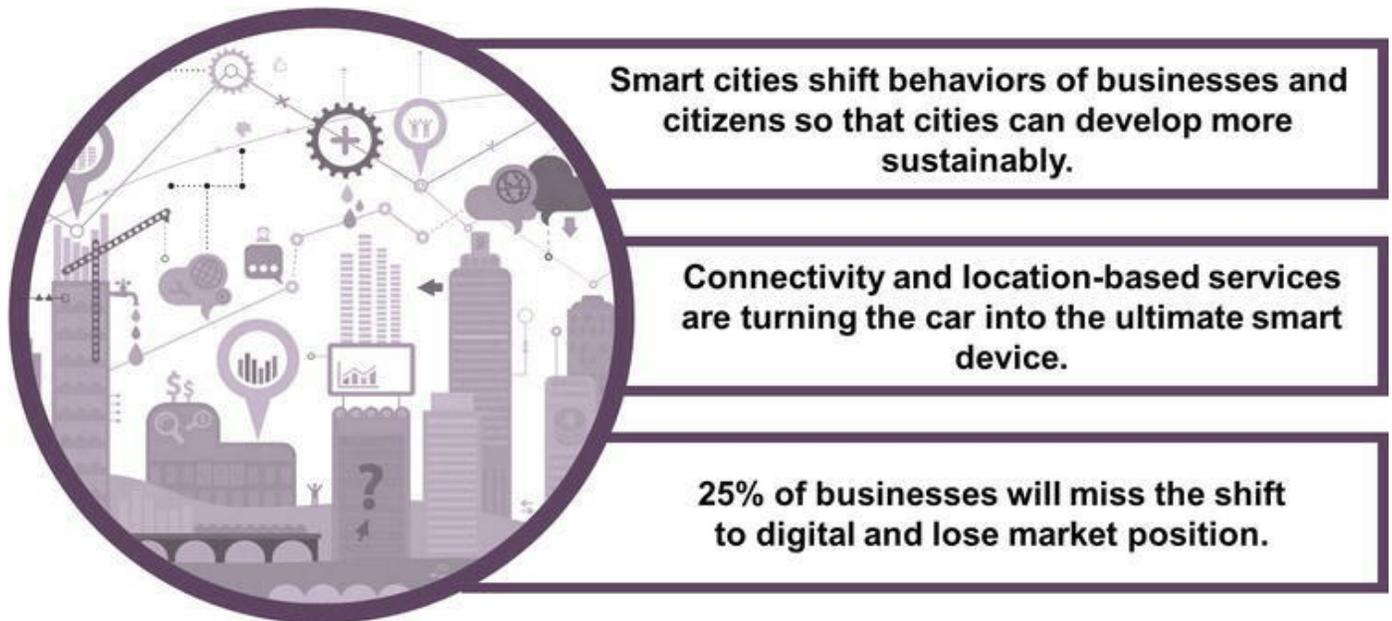


Figure 1-1 Exponential access, data, and connections drive exponential innovation

Ideas have always fueled business success. Ideas have built companies, markets, and industries. However, there is a difference today. [Figure 1-1](#) shows how businesses operate in the Idea Economy, which is also called the digital, application, or mobile economy. Doing business in the Idea Economy means turning an idea into a new product, capability, business, or industry.

Today, an entrepreneur with a good idea has access to the infrastructure and resources that a traditional Fortune 1000 company would have. That entrepreneur can rent compute capacity on demand, implement a Software-as-a-Service (SaaS) enterprise resource planning (ERP) system, use PayPal or Square for transactions, market products and services using Facebook or Google, and have FedEx or UPS run the supply chain.

Companies such as Vimeo, One Kings Lane, Dock to Dish, Uber, Pandora, Salesforce, and Airbnb used their ideas to change the world with very little start-up capital. Uber had a dramatic impact after launching its application connecting riders and drivers in 2009. Three years after its founding, the company expanded internationally. Without owning a single car, Uber now serves more than 758 cities worldwide (as of December 11, 2017). The company has disrupted the taxi industry; San Francisco Municipal Transportation Agency reported that cab use in San Francisco dropped 65% in two years.

In a technology-driven world, it takes more than just ideas to be successful. Success is defined by how quickly ideas can be turned into value. To succeed in the Idea Economy,

businesses need to:

Deliver new services faster—be a service provider

Spend more time developing and improving applications

Spend less time managing and maintaining infrastructure

Implement a flexible infrastructure that allows a “build-once, run-many” operations model

Creating disruptive waves of new demands and opportunities



Figure 1-2 The Idea Economy presents new opportunities and challenges

The Idea Economy presents an opportunity and a challenge for most enterprises. On one hand, cloud, mobile, Big Data, and analytics give businesses the tools to accelerate time to value. This increased speed allows organizations to combine applications and data to create dramatically new experiences, even new markets.

On the other hand, most organizations were built with rigid IT infrastructures that are costly to maintain. This rigidity makes it difficult, if not impossible, to implement new ideas quickly.

[Figure 1-2](#) shows how creating and delivering new business models, solutions, and

experiences require harnessing new types of applications, data, and risks. It also requires implementing new ways to build, operate, and consume technology. This new way of doing business no longer just supports the company—it becomes the core of the company.

Faster application development enables accelerated innovation

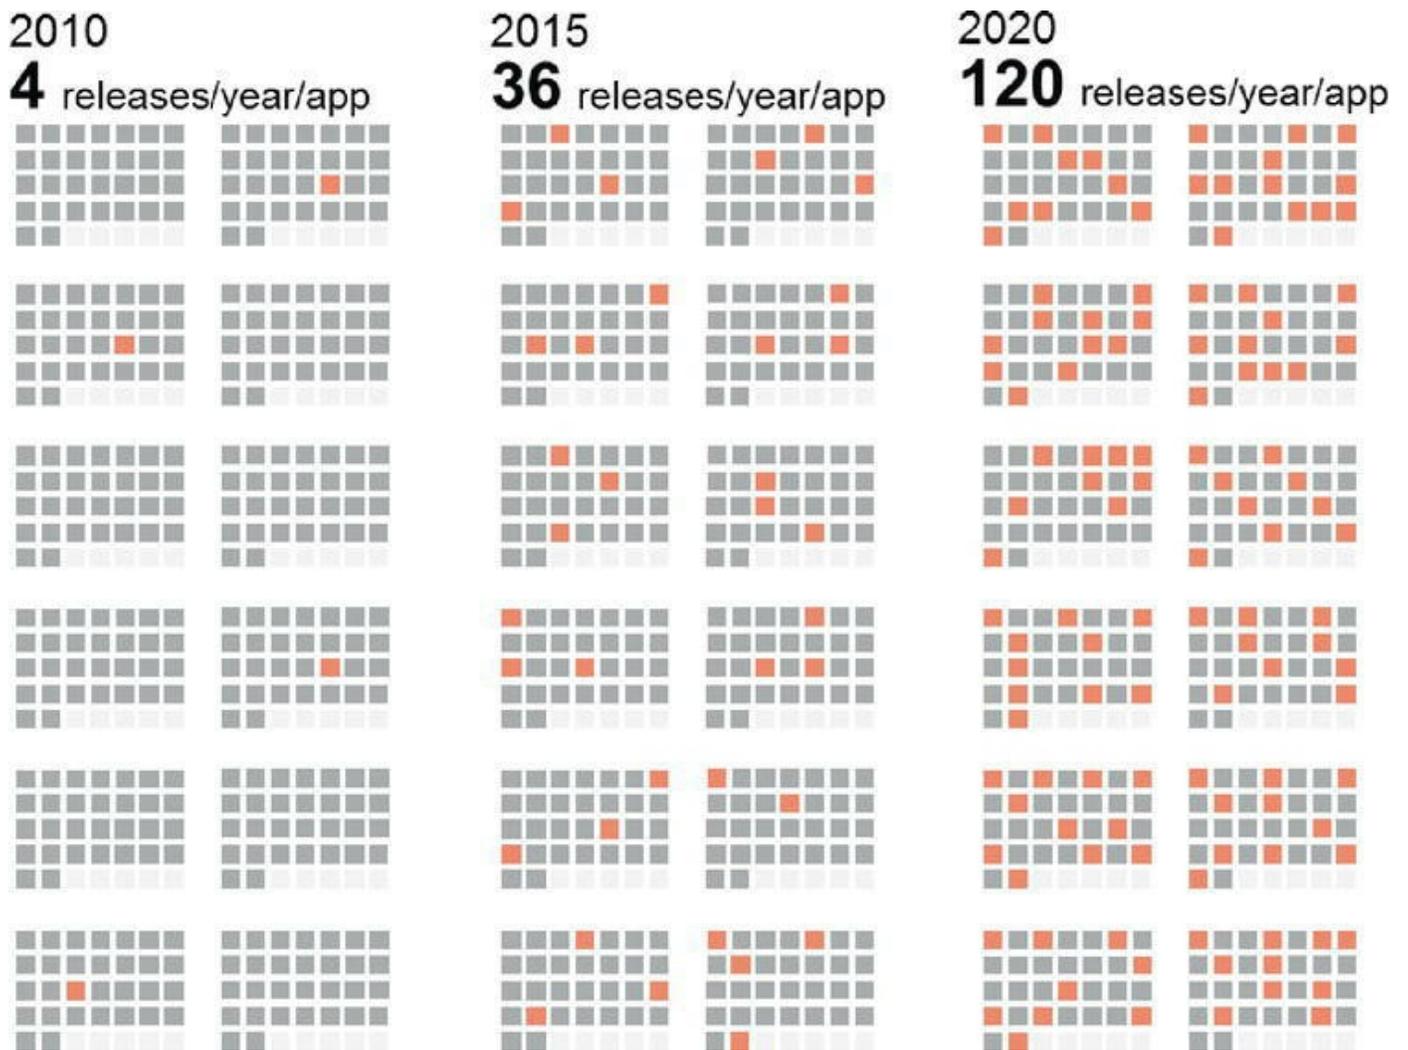


Figure 1-3 Application release cycles have vastly increased and continue to increase

From 2010 to 2015, much changed from an application development perspective. In 2010, the average application release cycle was four releases per year, per application. In 2015, this number went up to 36 releases per year, per application. It is projected that by 2020, there will be 120 releases per year, per application (30 times more releases

than in 2010) as shown in [Figure 1-3](#).

Considerations from the Forrester Thought Leader Paper commissioned by HPE, “Better Outcomes, Faster Results: Continuous Delivery and the Race for Better Business Performance,” helps summarize this trend.

Agility is paramount.

“Even when delivering at cadences of less than a week, 20% of organizations want to go even faster.”

Developers need flexibility.

Companies want open, flexible architectures for application portability and lock-in prevention.

Long time to value is costly

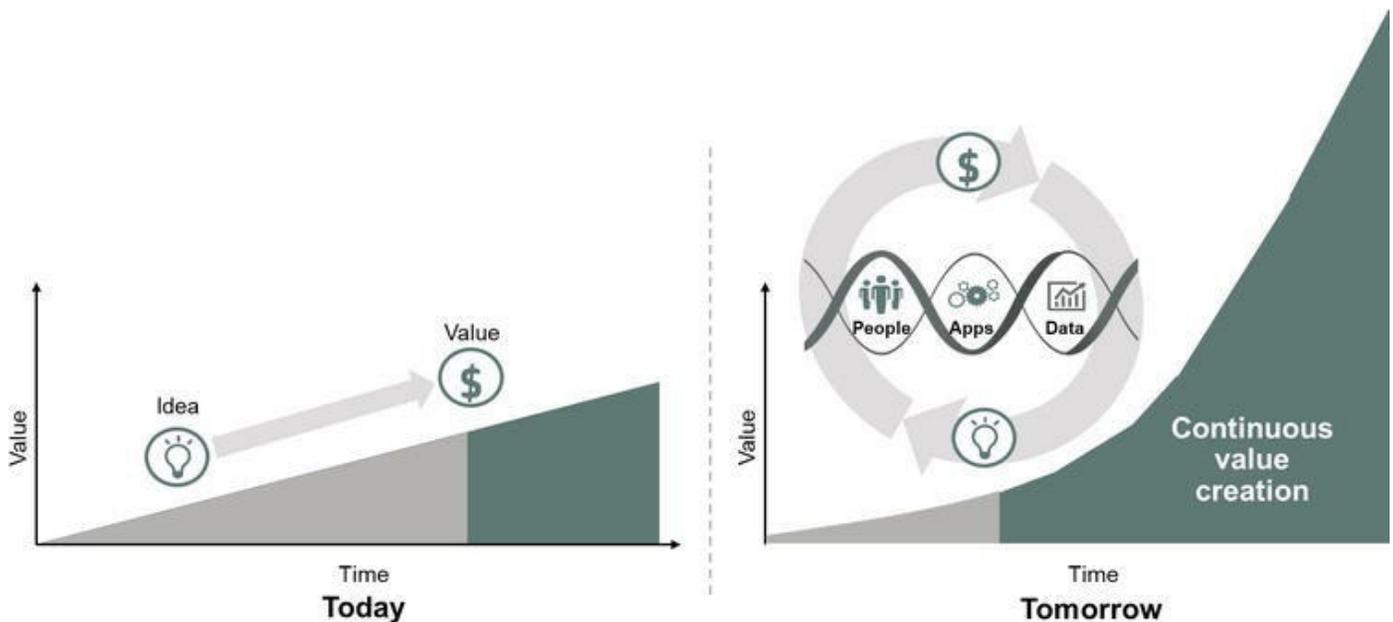


Figure 1-4 An organization must be able to alter its digital DNA continuously

As shown in [Figure 1-4](#), success today is defined by how quickly an enterprise can turn ideas into value—how quickly a business can experiment, learn, test, tune, and make things better. Speed is a key differentiator in all industries.

Uber did not invent a new technology. Instead, the company took advantage of the explosion of smartphones and mobile applications to design a compelling customer experience, ultimately creating a new way of doing business.

This example is not only about Uber executing a good idea; it is also about the taxicab industry's inability to act quickly to transform its business models to compete. Examples such as Uber serve as a warning. Every Fortune 1000 company is at risk of missing a market opportunity, failing to secure its enterprise, and being disrupted by a new idea or business model.

Timelines for IT projects and new applications used to be planned over years and months. Today, these projects take weeks or days. Increasingly, it is shrinking to hours. Now is the time for a company to ask:

How quickly can the company capitalize on a new idea?

How rapidly can the company seize a new opportunity?

How fast can the company respond to a new competitor that threatens the business?

The good news is that any company can use these technologies to adapt quickly to changing business models and achieve faster time to value. However, many established companies are working with a rigid IT infrastructure that can present significant challenges. Changing a data center strategy is a costly, lengthy, and complex process.

IT must become a value creator that bridges the old and the new

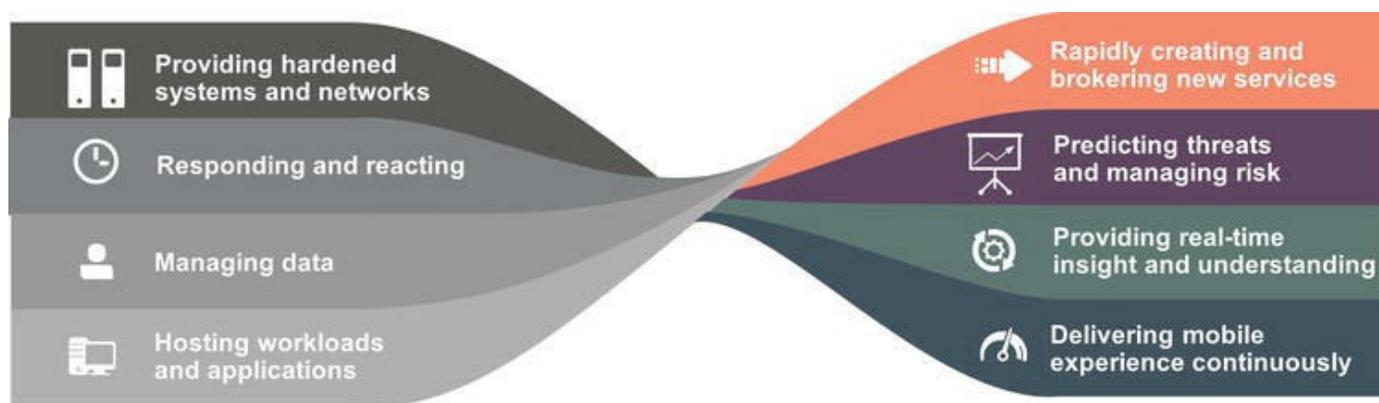


Figure 1-5 IT must transform from a cost center to a value creator

To respond to the disruptions created by the Idea Economy, IT must transform from a cost center to a value creator as shown in [Figure 1-5](#). In order to evolve, IT must shift focus:

From efficiently hosting workloads and services to continuously creating and delivering

new services

From simply providing hardened systems and networks to proactively managing and mitigating risks

From just storing and managing data to providing real-time insight and understanding

From using software to automate business systems to differentiating products and services

Customers need to make IT environments more efficient, productive, and secure. They need to enable their organizations to act rapidly on ideas by creating, consuming, and reconfiguring new solutions, experiences, and business models.

One of the first steps in achieving this kind of agility is to break down the old infrastructure silos that make enterprises resistant to new ideas internally and vulnerable to new ideas externally. Designing compelling new experiences and services does not work if the infrastructure cannot support them.

The right compute platform can make a significant impact on business outcomes and performance. Examples include storage that “thinks” as much as it stores; networking that moves information faster and more securely than ever before; and orchestration and management software that provides predictive capabilities.

HPE Transformation Solutions

Applications and data fuel the digital enterprise



Figure 1-6 Apps are the digital products driving innovation and causing market disruption

We are living in an Idea Economy, where the ability to turn an idea into a new product or service has never been easier. In today's hyper-connected world, applications are the mechanism for delivering innovative ideas, and we use apps for everything from our day-to-day life to business, as shown in [Figure 1-6](#).

Apps make life easy, but they also make it easy to switch suppliers in an instant, so maintaining customer loyalty requires providing a better customer experience through continuously delivering new app features. Every business is now a digital business, and apps are the digital products driving innovation and causing market disruption.

HPE Transformation Solutions

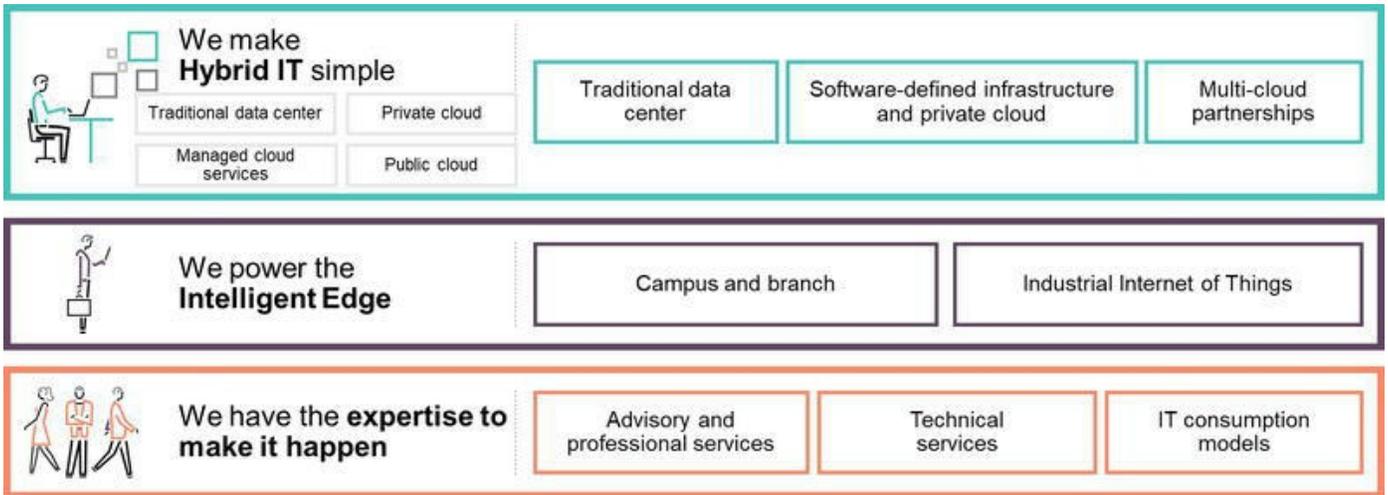


Figure 1-7 HPE Transformation Solutions

The HPE strategy represents a logical transition from the Transformation Areas: Transform, Protect, Empower, and Enable (the what) to Transformation Solutions (the how). As shown in [Figure 1-7](#), the strategy is designed to help customers harness the right mix of hybrid IT to accelerate service delivery, build next-generation data foundations to accelerate time to insight, and power a new generation of digital experiences to create business value.

First, the world is going to be hybrid and HPE makes hybrid IT simple. This is achieved through offerings in the traditional data center, software-defined infrastructure, systems software, private cloud, and through public cloud partnerships.

Second, HPE powers the Intelligent Edge through offerings in campus and branch networking via Aruba products and the Industrial Internet of Things (IoT) with products like Edgeline Converged Systems.

Third, services are going to be more critical than ever. HPE has the expertise to make it happen and can accelerate customers' journey through 22,000 Technology Services professionals. HPE can build customer solutions from the ground up, and deliver world-class support, advisory, and consumption models.

HPE also has a Financial Services organization. This is strategically important to customers as they increasingly look for financial flexibility and consumption models.



Note

HPE servers play a key role in allowing customers transform to a hybrid environment. Consequently, the material in this book focuses on the first Transformation Solution.

Accelerating time to value

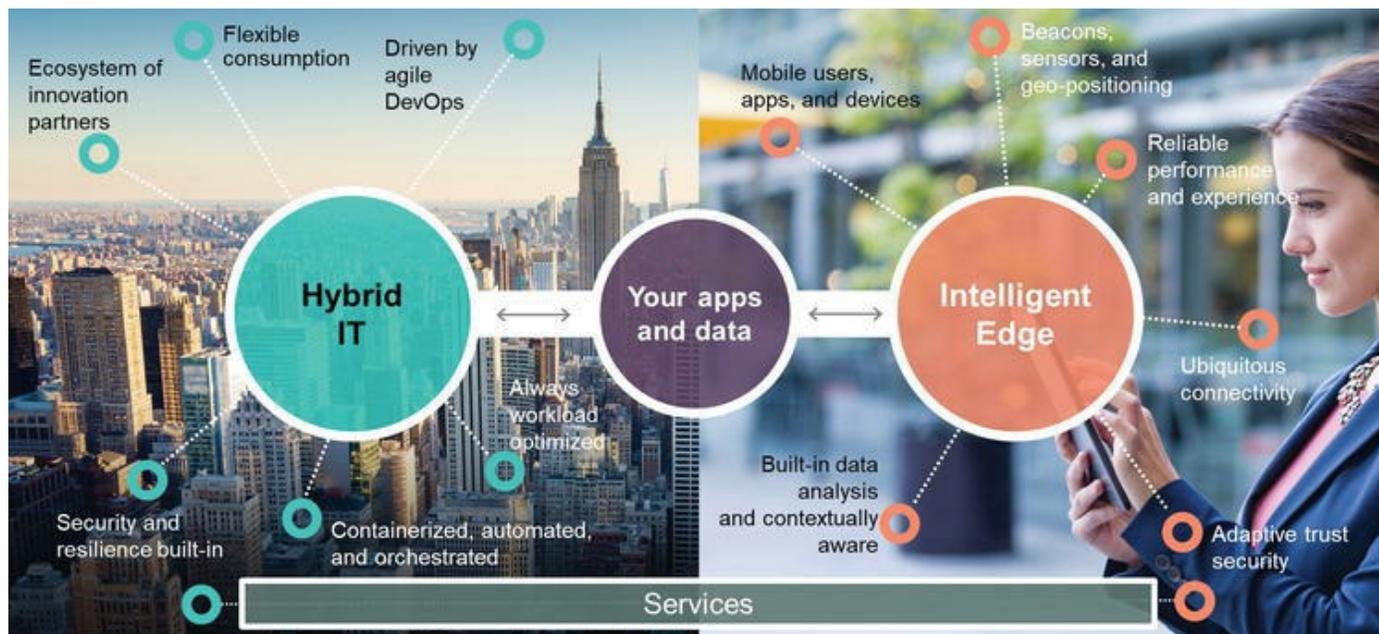


Figure 1-8 Speeding time to value across a hybrid world

HPE believes speeding time to value across a hybrid world is the number one priority, and to do this, the new generation of apps and data demand a vision that provides the flexibility to operate seamlessly across these environments.

As shown in [Figure 1-8](#), the vision is based on three core beliefs.

The world will be hybrid.

The Intelligent Edge is going to unleash an industrial IoT revolution.

Data and a new breed of apps is the digital fuel that will drive the future.

Whether your apps and data are in your data center, private and public clouds, multi-clouds, Hybrid IT, or at the edge of the network—the Intelligent Edge—where enterprises touch customers, employees, and the outside world, and where new experiences, new value, and new data is being created, it is critical that data collection and analytics be built-in to provide context and insights that drive value.

HPE believes that enterprises must unify apps old and new with data—wherever it lives in this hybrid IT context.

HPE has broken up infrastructure, software, and services into separate businesses to enable each to focus on excelling in this digital reality. HPE is focused on infrastructure and sells complete solutions to solve customers' business problems.

Hybrid IT

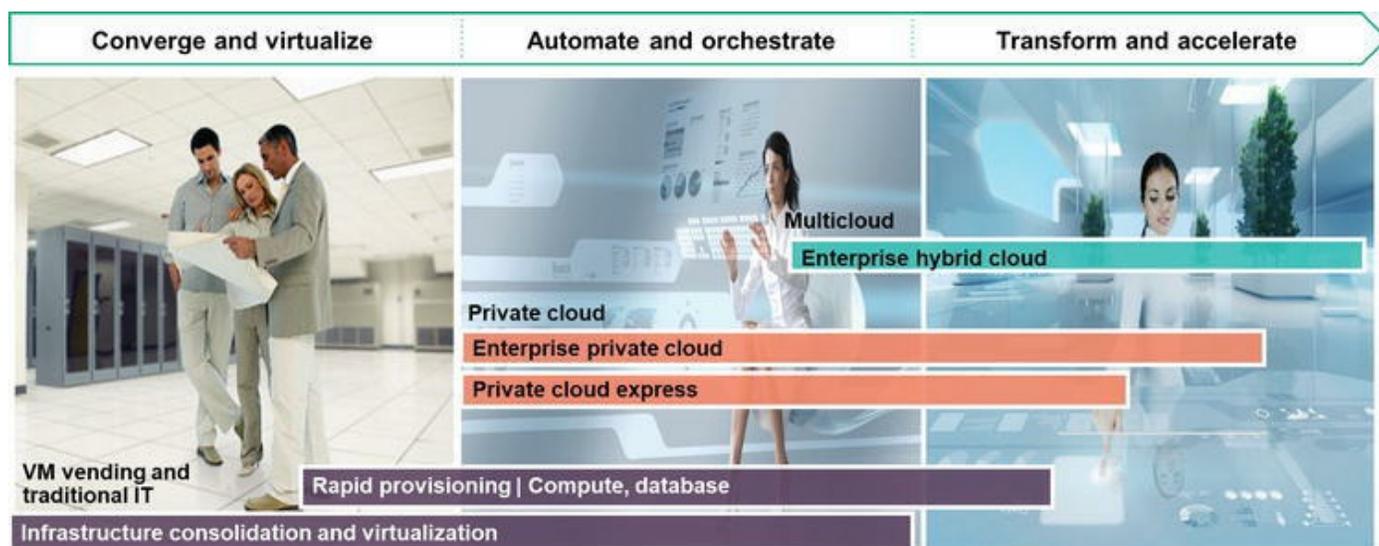


Figure 1-9 The solution continuum through digital transformation

Digital enterprises must manage apps and services using a single as-a-service IT model and optimize workload placement across Hybrid IT to balance workload service level agreement (SLA) needs.

As shown in [Figure 1-9](#), digital transformation requires:

Standardization, consolidation, and virtualization across the data center to reduce cost and improve utilization

Automation to orchestrate IT processes and tasks to improve operational efficiencies

Transformation to a DevOps delivery model to accelerate application and service delivery for business agility

The HPE portfolio of solutions and services enables organizations to meet this challenge and deliver business outcomes. HPE makes Hybrid IT simple.

Learning check

A customer is considering migrating their data center to a Cloud Service Provider. They are concerned about the risks involved and would rather have the ability to host some of their application on-premises and some in the cloud. How should you respond?

HPE Composable Infrastructure

HPE introduced the concept of a converged infrastructure because businesses that were trying to compete in the applications-driven economy found their traditional business infrastructure struggling. Traditional infrastructure is stable and allows organizations to slowly (over longer, planned periods of time) roll out applications that support the business, such as online transaction processing (OLTP) applications and databases. However, maintaining a traditional infrastructure in an application-driven economy is a challenge for which there is a solution.

The problem with bimodal computing

Many CIOs and data center system administrators feel challenged to deliver traditional business applications while at the same time standing up new applications such as mobile and cloud-native apps that drive revenue. Operations-driven and cost-focused, traditional IT environments make it difficult to deliver value to the business—IT cannot move fast enough for today’s application delivery goals.

Traditional applications are designed to support and automate existing business processes such as collaboration, data processing and analytics, supply chain, and web infrastructure. They include applications such as ERP and other large databases that have been prepackaged and pretested. These applications and services typically go through one or two release cycles per year. IT has been built around these for the last 20–30 years.

New apps-driven and agility-focused IT environments deliver apps and services that drive revenue and enhanced customer experiences through mobility, big data, and cloud-native technologies. These apps challenge IT to maintain a digital enterprise in a digital economy alongside traditional applications. But maintaining two different sets of infrastructure, one designed for traditional apps and another designed for cloud-native apps, increases costs and complexity. This approach is not sustainable.

Gartner gives the name “bi-modal” computing to the strategy of maintaining an existing infrastructure for traditional applications while transitioning to infrastructure and tools for emerging applications. The HPE vision is to pull both together with one infrastructure that provides the agility of on-premise cloud infrastructure.

A new category of infrastructure is needed to power the

Idea Economy

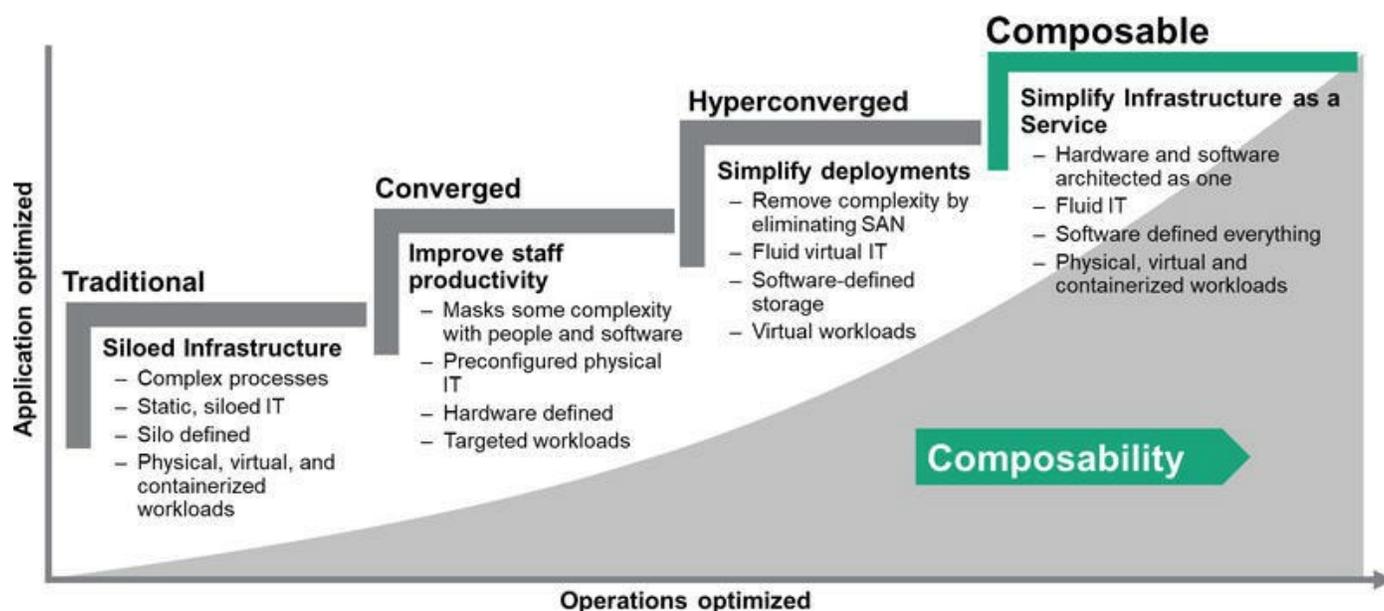


Figure 1-10 Composable Infrastructure must be agile and flexible

[Figure 1-10](#) shows the different categories of infrastructure. In addition to its focus on stability, reliability, and scalability, the traditional infrastructure is often siloed and burdened with complex processes. It can take IT months to stand up environments for traditional applications—regardless of whether these environments are meant to support physical, virtual, or containerized workloads.

HPE followed its converged infrastructure with hyper-converged solutions that only require minutes to deploy. These workload-specific offerings deliver fluid virtual IT and software-defined storage (SDS) in a single appliance.

HPE converged infrastructure and hyper-converged appliances help organizations succeed in the Idea Economy. Doing business in the Idea Economy—which is also called the digital, application, or mobile economy—means turning an idea into a new product, capability, business, or industry. To stay competitive, organizations need to develop new IT capabilities, create new outcomes, proactively manage risk, be predictive, and create a hyper-connected workplace. HPE converged infrastructure, software-defined management, and hyper-converged systems reduce costs and increase operational efficiency by eliminating silos and freeing available compute, storage, and networking resources.

However, most organizations must still maintain their traditional infrastructure, which can involve using a different set of tools and application programming interfaces (APIs). Businesses need a single, fluid infrastructure that is optimized to deliver

infrastructure for both traditional and new applications in seconds. This infrastructure must be agile and flexible enough to change personalities dynamically so that it can meet traditional and Idea Economy workload demands without overprovisioning resources for either. This architectural approach is known as Composable Infrastructure.

Moving to Composable Infrastructure

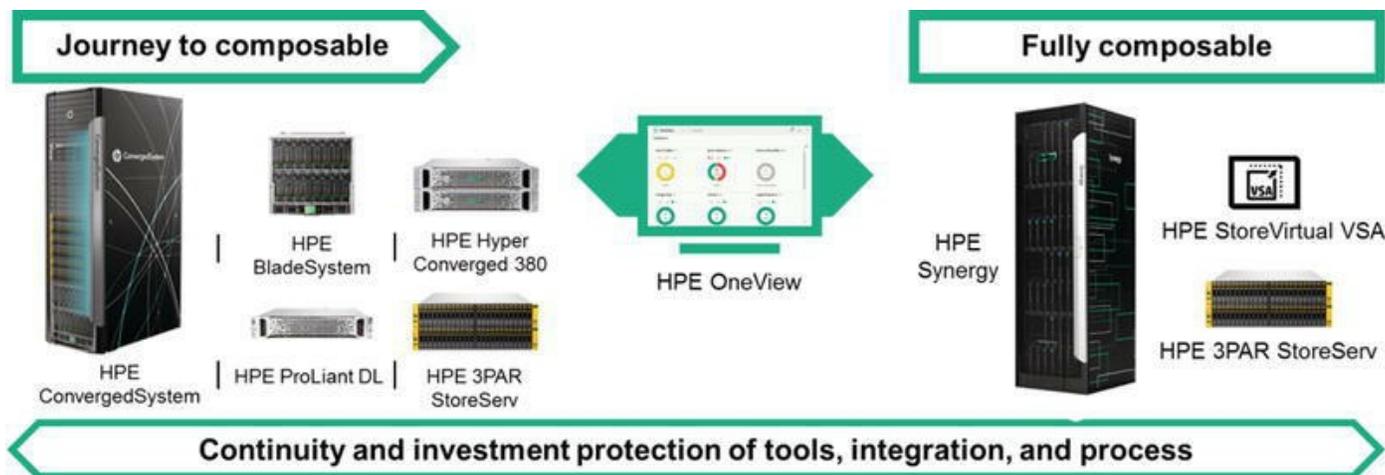


Figure 1-11 Journey to Composable Infrastructure

Building on a converged infrastructure, HPE has designed the Composable Infrastructure around three core principles.

Fluid resource pools meet each application's changing needs by allowing for the composition and recomposition of single blocks of disaggregated compute, storage, and fabric resources needed to run applications in sizes at the appropriate scale. This allows new applications to be quickly provisioned.

Software-defined intelligence provides a single-management interface to integrate operational silos and eliminate complexity. Workload templates speed deployment and frictionless change eliminates unnecessary downtime.

The unified API provides a single interface to discover, search, inventory, configure, provision, update, and diagnose the Composable Infrastructure. A single line of code enables full infrastructure programmability and can provision the infrastructure required for an application.

A Composable Infrastructure facilitates the move to a continuous services and application delivery model and enables applications to be updated as needed, rather than just once or twice a year. IT operates in a manner similar to how a cloud provider

handles lines of business and the extended enterprise. This type of framework consistently meets SLAs and provides the predictable performance needed to support core workloads.

HPE offers several solutions that provide customers with a bridge to Composable Infrastructure, as shown in [Figure 1-11](#). For example, HPE developed ConvergedSystem and hyper-converged solutions to help customers respond more quickly to their demands.

With the release of Synergy, customers can move to a fully Composable Infrastructure. Synergy greatly reduces the operational complexity of managing infrastructure and enables customers to accelerate application deployment. It is designed for today and architected to anticipate the evolution of technologies in the future.

HPE Composable Infrastructure vision



Figure 1-12 Create and deliver new value instantly and continuously in the Idea Economy

HPE Composable Infrastructure offers an experience that empowers IT to create and deliver new value instantly and continuously. The HPE Composable Infrastructure is a multi-year vision intended to deliver value to IT in four ways as shown in [Figure 1-12](#).

Run anything—Optimize all apps and service levels by storing data on a single infrastructure with fluid pools of physical and virtual compute, storage, and fabric resources.

Move faster—Accelerate app and service delivery through a single interface that precisely composes logical infrastructures at near-instant speeds.

Work efficiently—Reduce operational efforts and costs through internal software-defined intelligence that provides template-driven, frictionless operations.

Unlock value—Increase productivity and control across the data center by integrating and automating infrastructure operations.

The ability to provision an infrastructure as code means that you can provision a bare-metal infrastructure with one line of code in the same way that you would provision virtual machines and cloud resources. In addition, a single unified API allows ops and developers to access the resources as code.

Composable Infrastructure core principles

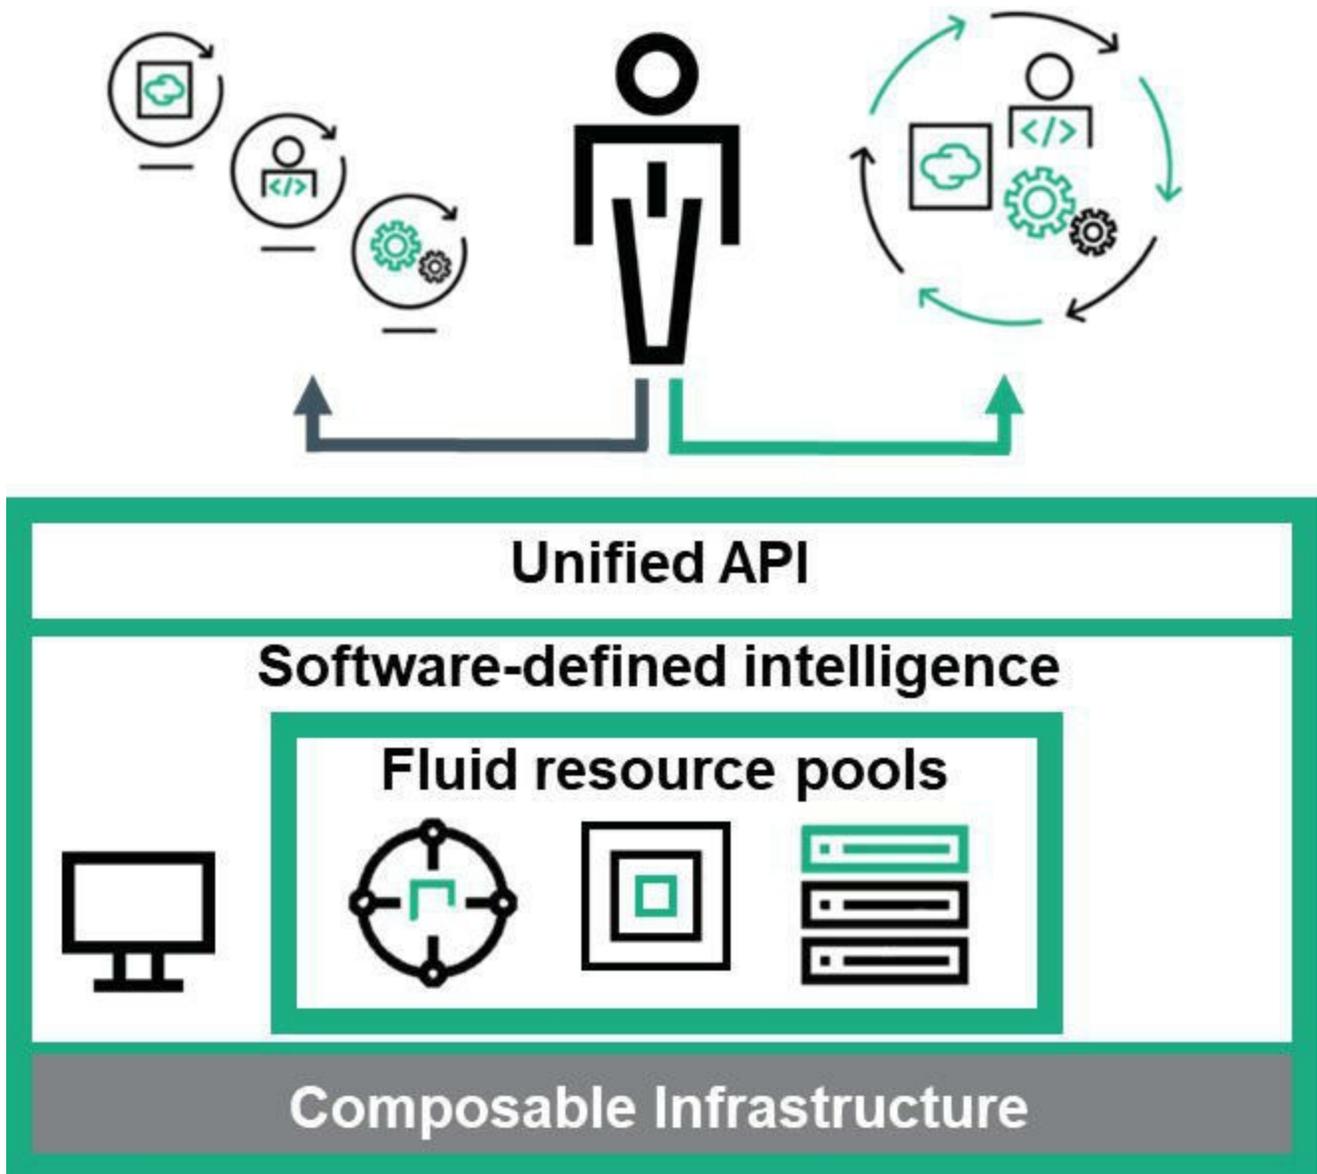


Figure 1-13 A two-mode approach designed around three core principles

HPE Synergy enables IT to accelerate application and service delivery through a single interface that composes and recomposes logical infrastructures into any combination. Composable resources are provisioned together with their state (BIOS settings, firmware, drivers, protocols, and so forth) and the operating system image using repeatable templates. This is ideal for traditional IT as well as a DevOps approach because it eliminates the time-consuming provisioning processes across operational silos.

Synergy is a next-generation approach designed around three core principles as shown in [Figure 1-13](#).

Fluid resource pools are combined into a single structure that boots up ready for any

workload with fluid pools of compute, storage, and fabric that can be instantly turned on and flexed. They effortlessly meet each application's changing needs by allowing for the composition and recomposition of single blocks of disaggregated compute, storage, and fabric infrastructure.

Software-defined intelligence embeds intelligence into the infrastructure. It uses workload templates to compose, recompose, and update in a repeatable, frictionless manner to speed deployment and eliminate unnecessary downtime. It provides a single-management interface to integrate operational silos and eliminate complexity.

Unified API provides a single interface to discover, search, inventory, configure, provision, update, and diagnose the Composable Infrastructure. A unified API allows the infrastructure to be programmed like code so it can become Infrastructure as a Service (IaaS). A single line of code enables full infrastructure programmability and can provision the infrastructure required for an application.

A Composable Infrastructure facilitates the move to a continuous services and application delivery model and enables applications to be updated as needed, rather than just once or twice a year. It enables IT to operate like a cloud provider to lines of business and the extended enterprise, consistently meet SLAs, and provide the predictable performance needed to support core workloads.

The three foundational components of an HPE Composable Infrastructure are:

HPE Synergy

HPE BladeSystem or HPE ConvergedSystem with HPE 3PAR

HPE ProLiant with HPE StoreVirtual VSA

HPE Composable Infrastructure solutions and HPE OneView

The journey to Composable Infrastructure begins with converged blocks that have composable attributes such as software-defined automation and the unified API from HPE OneView.

HPE Composable Infrastructure lets administrators and developers use infrastructure as code to control their internal environments. The unified API integrates dozens of popular management tools. For example, Chef needs an infrastructure to direct, such as a virtual machine, public cloud, or virtual cloud instance. The Composable Infrastructure API powered by HPE OneView creates, aggregates, and hosts internal IT resources that Chef can provision on demand, without needing a detailed understanding of the underlying physical elements. By connecting Chef tools with HPE OneView, bare-metal infrastructure can be directed in the same way as virtual and public cloud

resources.

Provisioning with Chef is just one example of how a Composable Infrastructure offers infrastructure as code. The same concepts allow:

Other automation tools such as Ansible and Puppet to provision bare-metal infrastructure

VMware and Microsoft technologies to create and flex virtualization clusters

HPE Helion to provision bare-metal hosts or virtualization clusters for multi-tiered applications through the OpenStack Ironic project

Docker to deploy infrastructure to host containers

Developers and independent software vendors (ISVs) can programmatically control a Composable Infrastructure through a single, open API that is native in HPE OneView. Continuous delivery of applications and services requires fast, policy-based automation of applications and infrastructure. The HPE Composable Infrastructure API enables developers to integrate with development, testing, and production automation toolchains and drive a more aligned and responsive delivery of IT services. HPE offers the hardware, software, services, and partner ecosystem needed to deliver an end-to-end offering with a single point of accountability.

These products enable continuity and provide investment protection of tools, integration, and processes.

Learning check

Which concept is central to the operation of HPE Composable Infrastructure?

- A. A single unified API allows operators and developers to access IT resources.
- B. Multiple differentiated APIs allow operators and developers to access IT resources.
- C. A single portal allows operators and developers to have seamless access to systems management functions such as Integrated Lights-Out 5 (iLO 5).
- D. Multiple differentiated portals allow operators and developers to have seamless access to systems management functions such as iLO 5.

Customer requirements

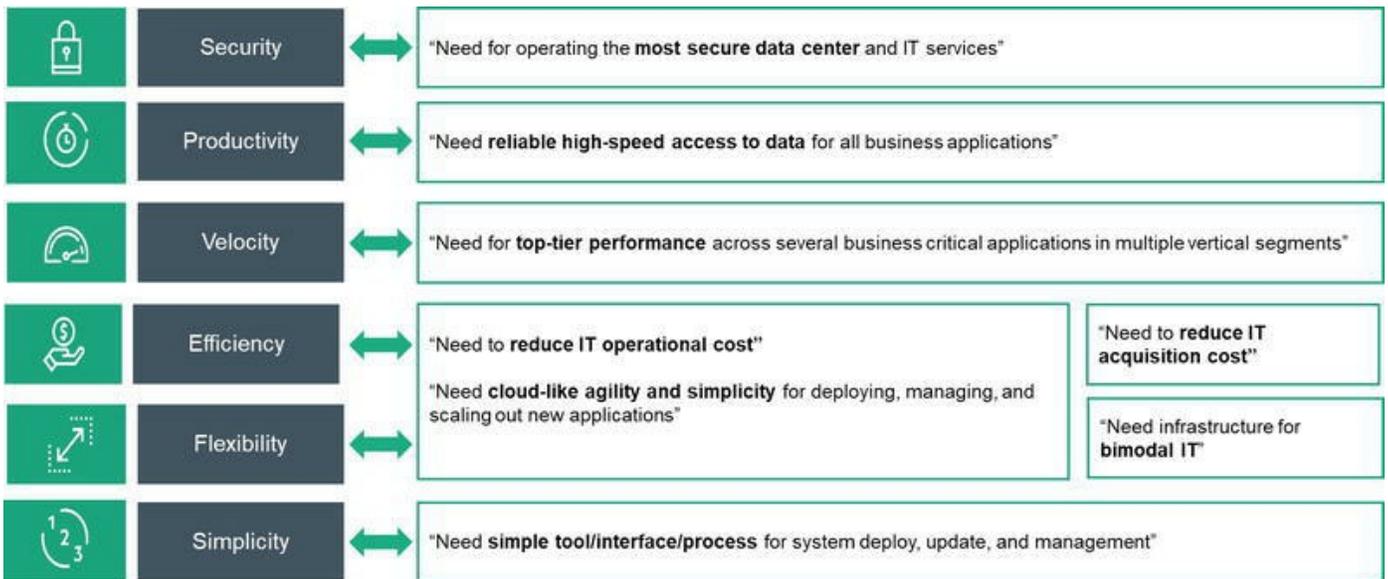


Figure 1-14 Gen10 innovation solves customer requirements

As shown in [Figure 1-14](#), customer concerns can be categorized around security, productivity, velocity, efficiency, flexibility, and simplicity. These concerns are discussed in the next few pages.

The HPE Gen10 Compute Experience has been shaped to accommodate new trends in IT and solve current and emerging customer requirements. There are three main pillars recognizing the need to provide a better way to deliver business results, a better way to protect customer business and data, and a better way to consume and pay for what you use.

Introducing a new generation compute experience



Figure 1-15 Gen10 servers offer customers agility, security, and economic control

Only HPE is uniquely able to deliver on this new IT experience not only because of the incredible innovation already brought to market such as Composable Infrastructure but now enhancing the whole experience with a new generation of capabilities powered by Gen 10 that offer customers the ability to accelerate applications and business insights, as shown in [Figure 1-15](#).

Provide a new experience in business and IT agility

HPE iLO is an industry-leading server management solution, differentiating itself by enabling three crucial requirements of server management: configuration, monitoring, and remote management. The introduction of HPE iLO 5 incorporates several improvements to simplify operations such as reducing maintenance windows with updates performed efficiently at scale with the least possible impact to production and the ability to easily roll back changes in the event of a firmware being compromised.

Simplified management enabled by iLO 5 with integrated Smart Update to schedule, stage, and rollback updates, Intelligent Provisioning integration with HPE smart Array Controller and faster configuration, and HPE iLO Amplifier Pack with the power to discover, inventory, and update HPE servers at unmatched speed and scale (rapidly discover and inventory up to 100,000 servers).

Data-centric applications performance is boosted with the introduction of in-memory compute with HPE Scalable Persistent Memory. These performance gains benefit from much larger in-memory compute with persistence, expanding to TBs instead of low 100s of GB with non-volatile dual in-line memory modules (NVDIMMs). The new higher capacity HPE 16GB NVDIMM delivers twice the capacity of the first generation

HPE NVDIMM and provides an ideal solution for any workload constrained by storage bottlenecks.

New to HPE Gen10 is Intelligent System Tuning (IST) making it easier to determine the best performance to match a given workload by dynamically tuning the performance on the server. These technologies, Jitter Smoothing, Core Boosting, and Workload Matching, are unique to HPE ProLiant Servers.

Superior business resilience and protection

All current HPE solutions are secure. HPE Gen9 offers FIPS Level 1, Measured Boot, Unified Extensible Firmware Interface (UEFI) Secure Boot, Trusted eXecution Technology, and secure options like TPM 2.0. However, as security threats have increased significantly, recently, HPE has introduced, in Gen10 servers, new security measures to provide customers with even higher levels of protection. Gen10 servers with iLO 5 include Silicon Root of Trust, Runtime Firmware Verification, and secure recovery of essential firmware and security options such as chassis intrusion detection. At the edge of the IT infrastructure, HPE Aruba ClearPass creates a strong networking security clearance protocol for clearing anyone requesting access to the network. Recently acquired Niara will offer the ability to monitor the activity of users inside the network, identifying, and reporting abnormal activity that may resemble potential malicious behavior.

Gain a new level of control regarding how you consume IT

Economic control means providing a better way to consume and pay only for what is used. HPE provides active capacity management, ensuring that customers have sufficient supply to meet demands available in minutes either on-premise or cloud. This will be of great benefit to meet unpredictable customer demands without the need to overprovision on-premise IT. As their business grows, the infrastructure can scale without incurring exponential costs, as well as having the flexibility to pay for IT as it is used, across the entire mix of hybrid infrastructure.

HPE is delivering a cloud-like experience with the security and economic control of customer's data center providing the best of both worlds for the future of hybrid IT.

HPE Gen10 servers—Relevant and differentiating innovations

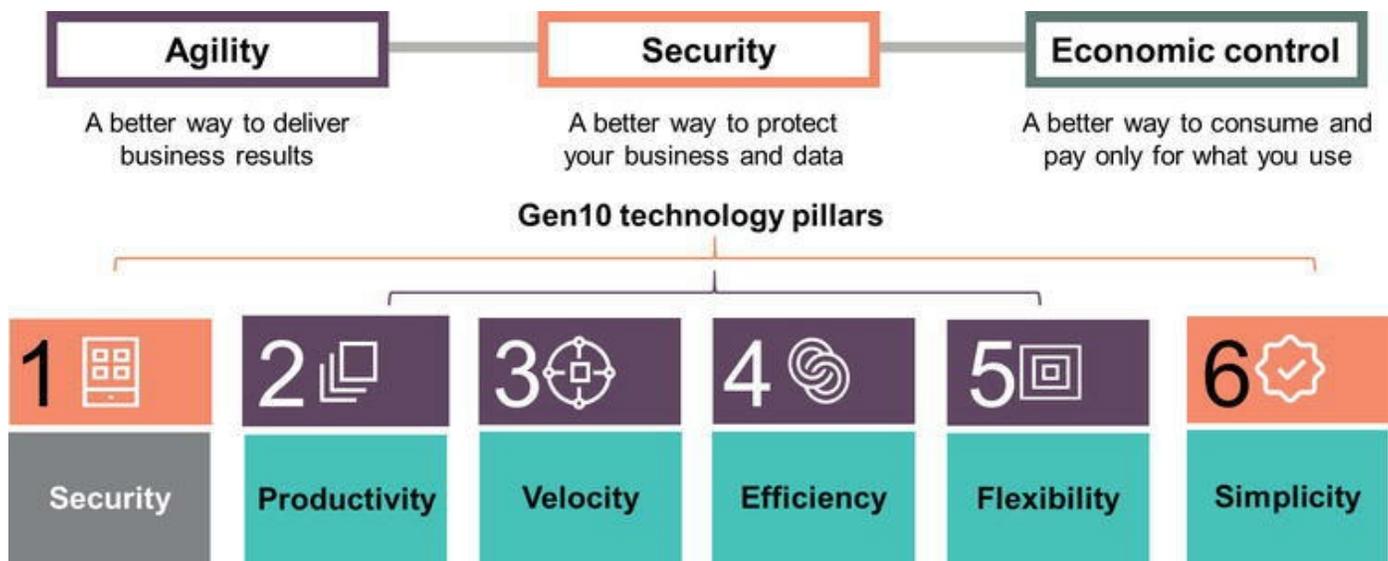


Figure 1-16 HPE Gen10 technology pillars

The six Gen10 technology pillars shown in [Figure 1-16](#) are summarized below:

Security

- Unmatched threat protection
- Supply chain attack detection
- Extensive standards compliance
- Unparalleled ability to recover firmware and operating system after attack

Productivity

- High-capacity
- Data acceleration
- Non-volatile memory
- Second generation of memory-driven compute

Velocity

- Enhanced GPU, PCIe SSD, and NVME drive performance and choice
- Higher speed memory access
- Next-generation industry-standard CPUs

Efficiency

- Processor performance tuning

- Predictable latency reduction
- Balanced workload optimization

Flexibility

- Substantially greater NVMe capacity
- Enhanced in-server storage density and drive count

Simplicity

- Large-scale firmware deployment
- Improved GUI and industry standard APIs
- Easy system debug access
- Convenient warranty entitlement validation

Learning check

Which of the following represent customer requirements? (Select four.)

- A. Velocity
- B. Legacy
- C. Efficiency
- D. Flexibility
- E. Dynamic
- F. Portability
- G. Security

Write a summary of the key concepts presented in this chapter.

Summary

Today's Idea Economy means enhanced access, data, and connections are driving exponential innovation that creates disruptive new challenges and opportunities for IT. IT organizations are being challenged to keep the business running with traditional applications while also developing new applications to drive new business opportunities.

Building on a converged infrastructure, HPE has designed the Composable Infrastructure, which is a next-generation approach designed around three core principles.

- Fluid resource pools
- Software-defined intelligence
- Unified API

The HPE Gen10 Compute Experience has been shaped to accommodate new trends in IT and solve current and emerging customer requirements.

2 Security

WHAT IS IN THIS CHAPTER FOR YOU?

After completing this chapter, you should be able to:

Explain how security risks impact customers.

Identify the security technologies embedded in HPE ProLiant servers.

Describe how the new Gen10 security features make HPE unique.

Describe HPE chassis and rack-scale security.

Prelearning check

You are attending a meeting with a customer to discuss their future data center modernization plans. The subject of security is high on their list of concerns. While they are comfortable with their progress in protecting against application-layer attack, they are aware of the possibility of attacks that are directed at the hardware level. They raise the question about what HPE offers in the Gen10 platform around protecting against firmware-level attack. How should you respond?

A customer is hosting solutions on behalf of other organizations and is interested in providing rack and server physical access security. For cost reasons, they would rather not have to build secure cages to provide this level of security. How should you respond?

Gen10 servers—Security

As previously discussed the three main areas HPE has focused on for the new generation of compute experience are agility, economic control, and security. In this chapter, we concentrate on security.

To support the claim that Gen10 is the world's most secure compute platform HPE has engaged a third-party company to test the security capability. They performed their penetration tests on the Gen10 platform and compared the result against some competitor offerings. The outcome of these tests places the HPE Gen10 platform ahead of the competition.

There are many examples of security breaches that can be found described on the web. Some notable attacks over the last couple of years include eBay, Target, JP Morgan Chase, Adobe, and European Central Bank. Each of these companies has invested huge amounts of time and money in their security architectures, but they still have had their security breached. For example, recently Yahoo revealed 500 million records have been stolen then later admitted a further 500 million.

Further investigation reveals that the average time that cyber criminals may be inside an organization before detection is 99 days, as of 2017, which is a significant improvement over the previous two years. In 2016, it was 146 days before detection, and in 2015, it was 229 days. Clearly, there is an improvement which can be attributed to the ways that organizations are investing in security programs and the fact that security companies are introducing security technologies into their products such as machine learning to improve detection capabilities. However, there is still the fact that criminals are inside the organization for over three months causing considerable damage to the organization. Further evidence shows that many breaches happen at the application layer. The source of these breaches could be from the development of bespoke applications instead of buying in more secure, robust, and tested solutions to achieve a more competitive edge. Independent software vendors such as Adobe and Microsoft invest heavily in their security teams and are quick to release patches to resolve any security holes. Bespoke applications may be released without so much concern for their security vulnerabilities.

While security breaches may be detected at the application layer, other breaches are occurring at the physical layer, a fact recently reported by the National Institute of Standards and Technology (NIST), part of the U.S. Department of Commerce.

It is the physical layer where the HPE ProLiant Gen10 development is focused to deal with growing concern for potential firmware or BIOS attacks. ProLiant Gen10 has been designed to provide a secure computing platform and protect the BIOS and firmware against attacks. This is a major concern for future solutions where unprotected BIOS or firmware may be the source of security breaches. This is of concern because the BIOS

is responsible for launching the system. If the BIOS is already compromised, then anything could happen further up in the stack of an operating system or application. Rootkit and Keyloggers that could potentially be added to firmware are distinct possibilities in the future.

The mean annualized cost of cybercrime incident could be as high as \$7.7 million, peaking at several hundred million. The Sony breach is reported to be as high as \$200 million. Not only is the financial cost high, the impact on a company's reputation may be incalculable.

The number of threats that companies are experiencing is increasing—something like a new piece of malware is detected every 4.2 seconds. However, the amount of money made available by organizations to harden security is dropping. The ideal would be to be able to detect cybercrime sooner. New regulations coming into effect in Europe in 2018, for example, the General Data Protection Regulation (GDPR), will update the data protection directive that has been in place since 1995. Basic GDPR compliance requires organizations to implement technologies like encryption and pseudo anonymization to protect customer data and state-of-the-art security controls. It also states that companies must not only implement policies to protect against breaches but also to inform of a breach discovery within 72 hours. GDPR has the ability to fine an organization up to 4% of global revenue for data theft or other security breaches.

Hybrid cloud and hybrid IT is the catalyst for business transformation. This transformation fuels innovation, but brings new risks, new exposures, and attack surfaces, thereby increasing cost and complexities of regulatory pressures to combat increasingly sophisticated cyberattacks.

In addition to data theft, there are other forms of attack that need to be considered.

DoS—Denial-of-Service attack (DoS attack) is a cyberattack where the perpetrator seeks to make a machine or network resource unavailable to its intended users by temporarily or indefinitely disrupting services of a host connected to the internet.

DDoS—A Distributed DoS attack (DDoS attack) occurs when multiple systems flood the bandwidth or resources of a targeted system, usually one or more web servers. Such an attack is often the result of multiple compromised systems (for example, a botnet) flooding the targeted system with traffic.

PDoS—Permanent DoS attack (PDoS attack), rather than collecting data or providing some ongoing nefarious function, its only aim is to completely prevent its target device from functioning.

Can you trust your hardware vendor?

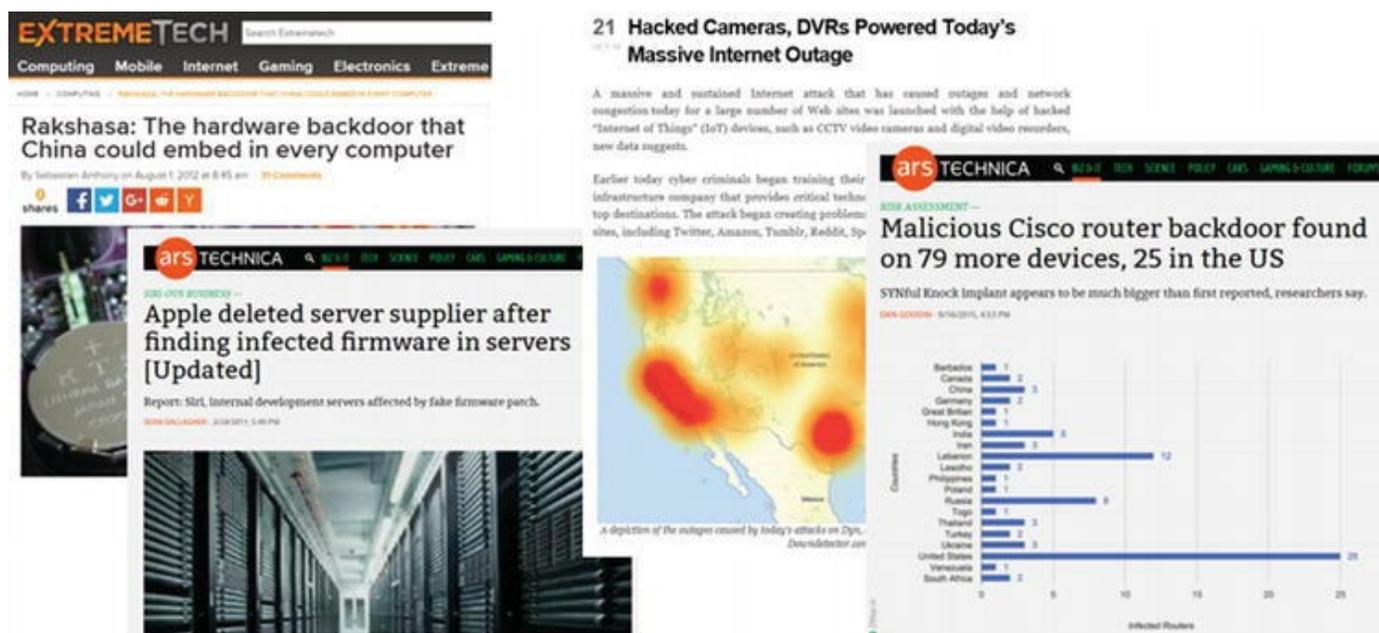


Figure 2-1 Some examples of firmware-level security breaches

The ProLiant Gen10 security features have developed around the concept of trust in hardware vendors.

There are a few of examples reported over the past few years that really demonstrate the requirement for firmware-level security, as shown in [Figure 2-1](#).

The first is a recent situation with Apple. Apple made the decision to move away from a tier one server vendor because they wanted to save money in their data centers. Apple chose Supermicro as the vendor, resulting in lower cost with less-sophisticated management features and security capabilities. Apple had deployed these servers in their Siri and internal development server data centers.

Someone hacked into the Supermicro update server and placed a malicious firmware patch on this server which was subsequently downloaded to the servers at Apple. Luckily for Apple, they were able to detect the malicious update before it could do harm and they were able to isolate the servers.

When working with a vendor that does not invest in security, with things like protecting firmware with signed certificates, the potential for mishap is huge. Apple learned their lesson by going with cheaper hardware and suffered the problems that it caused.

A problem was encountered by Cisco several years ago where there was a vulnerability called the SYNful Knock implant. This involved the ability to send a specially

constructed sequence of TCP packets to a Cisco router then use that to install malicious firmware, which provided an opportunity to take the device over remotely.

Another case involved a malware approach called Rakshasha. This is a proof of concept that looked at Chinese motherboards and discussed the possibility of installing backdoors in the hardware underneath the firmware level so they would never be identified during the boot sequence on the laptops that were being used.

While examples of malicious breaches at the server firmware level are not massive today there are enough examples out there that can be used to show customers what they should be thinking about when they are investing in new infrastructure.

PDoS is the main focus for the HPE Silicon Root of Trust security technology, and this is discussed in more detail on the next page.

Sources:

<https://arstechnica.com/information-technology/2017/02/apple-axed-supermicro-servers-from-datacenters-because-of-bad-firmware-update/>

<https://www.extremetech.com/computing/133773-rakshasha-the-hardware-backdoor-that-china-could-embed-in-every-computer>

<https://arstechnica.com/security/2015/09/malicious-cisco-router-backdoor-found-on-79-more-devices-25-in-the-us/>

HPE Security focus

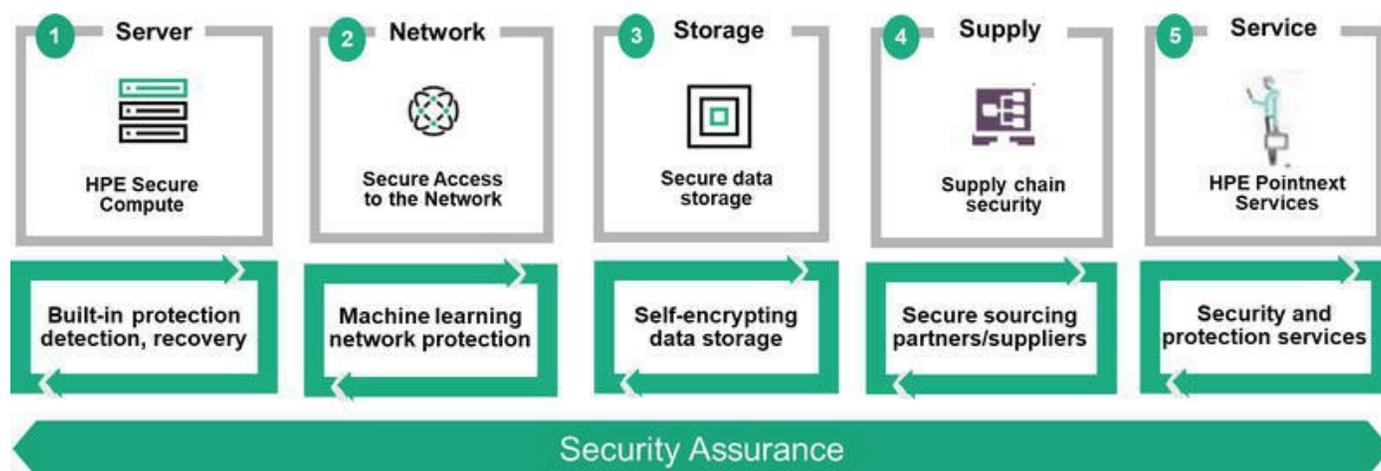


Figure 2-2 The areas of the HPE Security focus

The HPE Security focus is concerned with ensuring the best protection for customers by embedding security technologies into hardware. [Figure 2-2](#) shows all the areas of this focus. This chapter is largely concerned with the HPE Server Secure Compute Lifecycle with built-in protection, detection, and recovery capability.

Other HPE portfolios, such as networking, offer machine learning on the network and admission control. From a storage perspective, there is self-encrypting data storage. Security capability is extended to the supply chain to ensure that third-party and partner suppliers consider security when manufacturing components, in line with the HPE Security requirements. Finally, the services capability provided by HPE Pointnext allows HPE to offer additional security and protection services to customers.

From the HPE ProLiant server perspective, there are three main areas that the inbuilt security is aimed at providing protection for customers. While DDoS attacks are commonly reported, it is PDoS that is of concern here. PDoS refers to an attack vector that could be used to kill a server remotely by damaging firmware and affecting the server's ability to boot. In ProLiant Gen10, an immutable silicon-based fingerprint is installed in the server, thus ensuring the server will never boot with compromised firmware.

If an attack does occur, there is a built-in ability to automatically detect and recover from compromised firmware, and a known good state is restored to the server.

Finally, there is an option to perform the highest-level security protection inside the platform by complying with Commercial National Security Algorithms (CNSA) Suite, a U.S. government standard that allows these servers to be deployed in top-level security

environments in the U.S. government. This will have global relevance for other countries with other government customers.

HPE Secure Compute Lifecycle

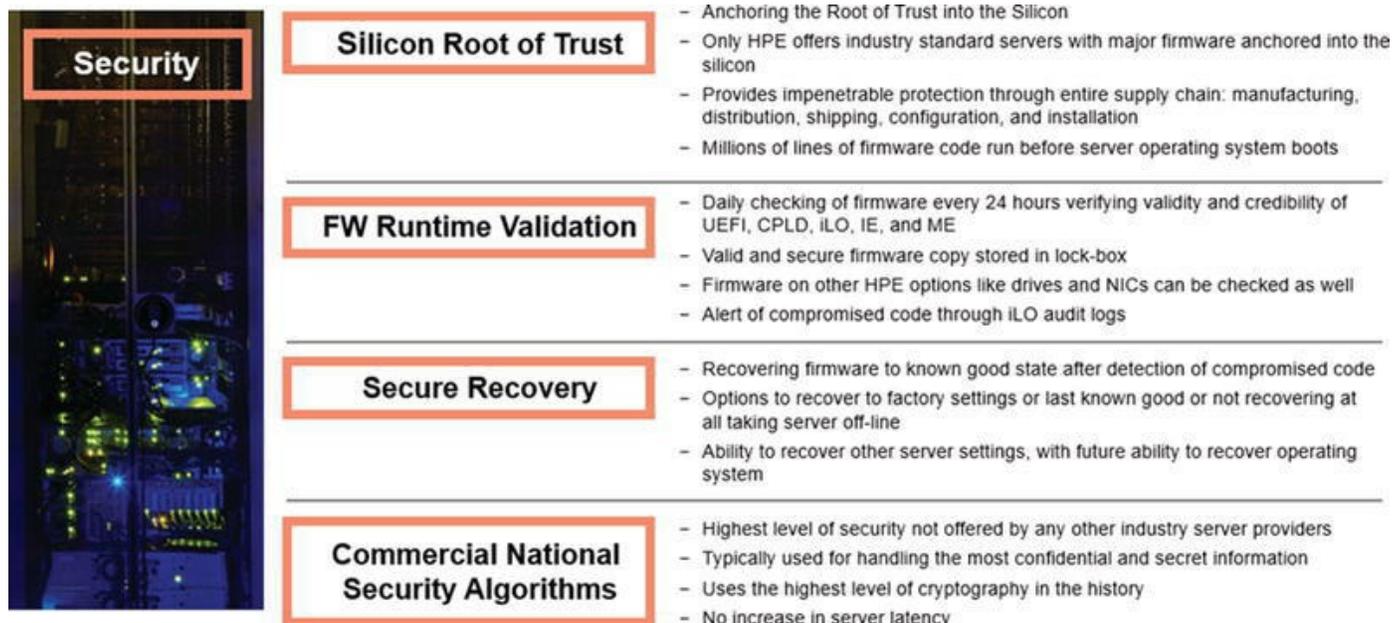


Figure 2-3 The HPE Secure Compute Lifecycle

[Figure 2-3](#) gives an overview of the HPE Secure Compute Lifecycle.

Silicon Root of Trust and HPE Gen10 servers:

Removes vulnerabilities that expose infrastructure firmware to malicious attacks

Enables the detection of the previously undetectable with runtime firmware verification to monitor for attacks and generate audit logs for security information and event management (SIEM) software

Avoids lasting damage to the business by quickly restoring the firmware and operating system to the factory settings or the last known authenticated safe setting, in the unlikely event of a breach

The HPE Silicon Root of Trust is anchored in the silicon and is unique in the industry, giving HPE an impenetrable protection right through the entire supply chain. Therefore, at all stages of the manufacturing process, authenticity and integrity is assured for the server platform.

While the job of CIO has always been split between introducing new technologies to improve the business while keeping it all running and secure at the right cost, never before have we seen such tension between the desire for organizations to, as seamlessly as possible, reinvent themselves digitally while maintaining strict business continuity with the highest degrees of security against threats, both within and without.

Worries about cybersecurity, including a potentially career-ending public hack, has become a top concern to the extent that some IT management surveys put it at the very top of the spending list. Security is pulling focus away from other digital priorities, with not much to show for in terms of new value creation, making today's relatively flat IT budgets an even greater challenge. CIOs are seeking creative ways to manage security and looking for IT security breakthroughs to help them spend more time leading other top-level priorities.

HPE is resolute to stay one step ahead. Products are engineered based on the belief that infrastructure should be the strongest defense, armed with the latest innovations to prevent, detect, and recover from security attacks. Just as customers expect and deserve high-quality and reliable products, HPE also believes that customers should expect the most secure infrastructure in the industry, which is why security is embedded into all HPE offerings. Poor quality and reliability can slow down your business, but security vulnerabilities can harm your business and your brand.

HPE views the supply chain as an essential element of cybersecurity because of the possibility that products could be compromised at their source. HPE reduces the risk of exposing the supply chain to threats such as counterfeit materials, malicious software embedded in products, and other untrustworthy components by vetting component vendors and sourcing from Trade Agreements Act (TAA) designated countries. Because of its unique and privileged position within the IT infrastructure, HPE chose to reduce security concerns and threats to BIOS firmware by developing BIOS firmware and Application-Specific Integrated Circuits (ASICs) in-house.

Firmware security risks are nothing new but are rarely considered as part of a risk assessment. NIST has published two firmware protection guidance documents in the public domain, highlighting growing concerns. In a recent study, more than 50% of companies reported at least one incident of malware-infected firmware with 17% indicating the incident resulted in a material impact.

Firmware attacks can be introduced in the supply chain or appear as vulnerabilities in previously trusted firmware that could result in a targeted attack on an organization, either a PDoS or a persistent malware presence. Compromised firmware is both hard to detect and difficult to recover from, and it is mostly invisible to the software layer of a system, including most antivirus and spyware tools. Malicious people can plant rootkits (compact and dormant malicious hooks) in the firmware whose primary function is to provide an Application Programming Interface (API) to other viruses and worms on an infected system.

The potential damage from a malicious firmware attack is far worse than a software attack due to the persistent nature of the firmware. For example, a DoS attack on the system software stack may result in a corrupted operating system stack that would need

to be reinstalled. In contrast, a compromised firmware DoS attack may result in a completely unbootable and unusable platform.

Protection requires making the firmware image part of a Root of Trust. To achieve a hardware Root of Trust is a major advancement in security protection as it originates from the hardware passing to the software through the firmware, making it a critical part of the Chain of Trust.

Design ownership and control over the hardware and firmware is required to achieve this level of security. HPE is the first manufacturer to implement a Silicon Root of Trust in a commercial server, a boot process that depends on a series of trusted handshakes starting from a signature key embedded in the hardware at the factory with state-of-the-art encryption and breach detection technologies that set them apart from the competition.

The responsibility for firmware verification is rooted in the Integrated Lights-Out (iLO) management system. Further protection provided by the CNSA feature really locks down iLO access, providing stricter logon processes, stronger cryptography on the data going over the network, and more locked down interfaces.

HPE secure server management from anywhere, anytime

<p>Maintain complete control of your secure server, proactively managing it with ease and minimal manual intervention</p>  <p>Available on most Gen10 ProLiant, Apollo, and Synergy servers</p>	<p>Customer needs:</p> <ul style="list-style-type: none">- Uncompromising security- Hassle-free server management and integration into infrastructure management ecosystem using industry standards- Intuitive, user-friendly server maintenance <p>Key new features:</p> <ul style="list-style-type: none">- Immutable Silicon Root of Trust for Secure Start with ability to automatically rollback to known-good firmware- Common Access Card (CAC) 2-factor authentication support- OpenLDAP support- Additional iLO security modes- Granular control of all iLO interfaces- Run-time Firmware Validation to verify integrity of iLO and BIOS- With 2x the CPU MHz in the iLO 5, virtual media performance is twice as fast vs iLO 4- Open IPMI mode for increased interoperability with industry IPMI tools
--	--

Figure 2-4 HPE iLO 5 chipset

Many of the new security capabilities shown in [Figure 2-4](#) were introduced in the ProLiant Gen10 platform and are centered in the iLO 5 chipset embedded in the server

motherboard. This lights-out management chipset provides easy server management, with both direct connect capability via a web interface as well as API capabilities.

iLO supports integrated security features. HPE designs, writes, and develops the code that goes into the iLO chipset in-house. HPE is unlike competitors who, when they need firmware written, will engage with third parties to do it for them; therefore, introducing concerns over trust and how the third-party deals with security. With iLO 5, HPE has introduced a Root of Trust embedded in the chipset, which means that all the way from the start of the manufacturing process HPE ensures that the platform is secure.

Features:

Immutable Silicon Root of Trust for Secure Start with ability to automatically rollback to known good firmware

Common Access Card (CAC) two-factor authentication support

OpenLDAP support

Additional iLO security modes

Granular control of all iLO interfaces

Run-time Firmware Validation to verify integrity of iLO and BIOS

Secure Start—Featuring Silicon Root of Trust

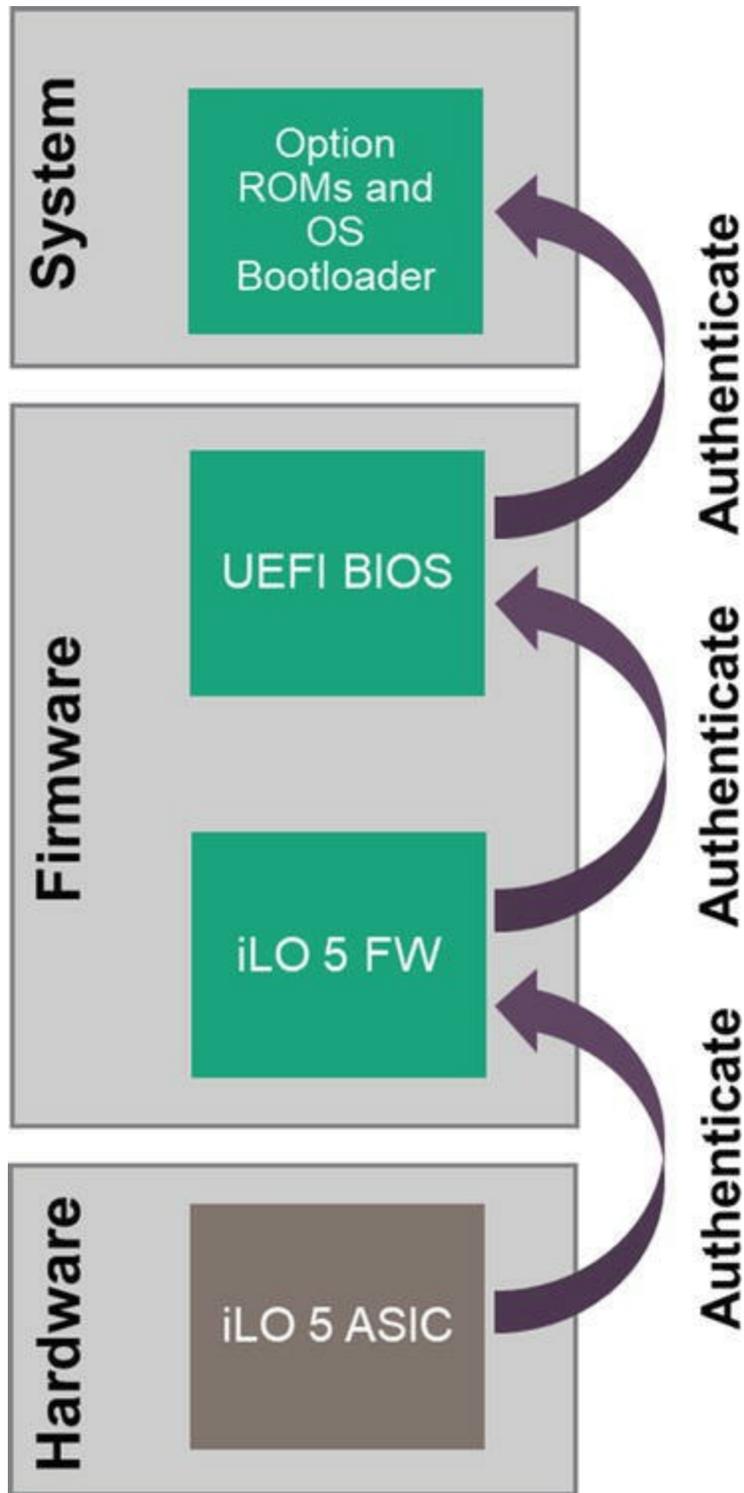


Figure 2-5 iLO and BIOS team up to provide enhanced security features

[Figure 2-5](#) shows how iLO and BIOS team up to provide enhanced security features.

Silicon Root of Trust begins when iLO powers up

- HPE-designed logic in iLO chip validates the iLO firmware

- After successful validation iLO waits for server power on
 - If iLO firmware validation fails, iLO performs recovery to last known good image
- iLO firmware then validates the system ROM
- Digital signature must match or the ROM is not executed
 - iLO firmware is trusted, then the ROM is trusted (Chain of Trust)

ROM then validates Option ROMs and the operating system bootloader via Unified Extensible Firmware Interface (UEFI) Secure Boot

- Option ROMs and operating system bootloader are not executed if they fail authentication.

When the platform is manufactured, a fingerprint is burned into the chipset in the form of a digital signature created using private keys to which only HPE has access. During the iLO boot process, the digital signature is used to verify the authenticity of the firmware. The firmware is booted only if the signatures match, thus indicating that the firmware is authentic and has not been compromised. The iLO ASIC progresses to the next level, where the iLO firmware verifies the UEFI BIOS. The UEFI BIOS checks that the Option ROMs and the OS Bootloader are all authenticated, effectively following a Chain of Trust from the iLO all the way up the stack through the firmware, through the BIOS and eventually up to the Option ROMs and operating system bootloader.

HPE calls this feature Secure Start, which effectively is the first two and a half layers in the stack. The UEFI Secure Boot process takes over, which provides the BIOS and bootloader checking process.

UEFI Secure Boot

UEFI Secure Boot is a platform feature in UEFI that replaces the traditional BIOS. Secure Boot provides:

UEFI standard. Works with HPE and third-party cards and major operating systems including:

- o Windows 8/Server 2012+SLES 11SP3+, RHEL7+, Ubuntu 12.10+, Fedora 18+
- o vSphere 6.5+

Functionality

All UEFI drivers, OS bootloaders, and UEFI applications are digitally signed.

Binaries are verified using a set of embedded trusted keys.

Only validated and authorized components are executed.

Creates a Chain of Trust. Improved solution over TCG Trusted Boot.

Learning check

You are attending a meeting with a customer to discuss their future data center modernization plans. The subject of security is high on their list of concerns. While they are comfortable with their progress in protecting against application-layer attack, they are aware of the possibility of attacks that are directed at the hardware level. They raise the question about what HPE offers in the Gen10 platform around protecting against firmware-level attack. How should you respond?

What is the name that HPE has given to the process that includes the “Silicon Root of Trust”?

- A. Fast Start
- B. New Start
- C. Secure Start
- D. Up Start
- E. Valid Start

Secure recovery

Many of the iLO 5 security features are available with the iLO Standard License. The iLO secure recovery feature is initiated automatically by iLO, recovering from a known good version of the firmware in the event of the iLO firmware validation failure. The iLO Advanced Premium Security Edition license is required to achieve automatic recovery of the ROM. Run-time scan takes place periodically to provide continuous firmware authentication and to detect and report compromised firmware. Features include:

Redundancy

- ROM and iLO have built-in redundancy

Recovery

- Factory-installed recovery set on non-volatile storage.
- Recovery administrator can setup new recovery set.
- iLO can automatically recover iLO (iLO Standard feature).
- iLO can automatically recover ROM (iLO Premium feature).

Run-time firmware authentication

- Background scans by iLO
- Event logs, alerting, web UI

Security built into every level

iLO Standard	<ul style="list-style-type: none"> - Silicon Root of Trust - FW supply chain attack detection - FIPS 140-2 Level 1 validation - Secure made BIOS (TAA) - Manual secure recovery - Authenticated updates - Common criteria - Single sign-on 	<ul style="list-style-type: none"> - Secure Start - Measured Boot - UEFI Secure Boot - Agentless management - Remote firmware update - Trusted eXecution Technology - NIST 800-147b BIOS/UEFI protection 	HW options
iLO Advanced	<ul style="list-style-type: none"> - CAC 2-factor authentication - Remote system logs - Remote console - Virtual media 	<ul style="list-style-type: none"> - Directory services - ArcSight unique connector - Kerberos 2-factor authentication 	<ul style="list-style-type: none"> - Chassis intrusion detection - Three-factor rack security - NICs - TPM - Cyber safe TAA SKUs - Smart Array w/ secure encryption
iLO Advanced Premium Security Edition	<ul style="list-style-type: none"> - Automatic Secure Recovery - Runtime FW validation - Secure Erase of User Data - Commercial National Security Algorithms 		

Figure 2-6 iLO licensing

Security is built into every level. [Figure 2-6](#) lists the security features available with each iLO license type. Many of the security features are provided with iLO Standard.

While iLO 5 is responsible for security enforcement, not every company will wish to pay a premium for services that they may not need. Fortunately, most of the security features are included in the iLO Standard license, such as Silicon Root of Trust, Secure Start, Secure Boot and Measured Boot. iLO provides a notification if firmware has been compromised, which will require manual intervention to correct the issue. For most organizations this level of capability is sufficient. The risk of this type of attack is very low and would probably result in triggering a forensic investigation as to how this breach occurred in the first place.

The iLO Advanced license offers advanced security features, with secure remote management. It also includes directory integration, CAC support, and Kerberos authentication to a directory service.

The iLO Advanced Premium Security Edition license provides all the features associated with iLO Standard and Advanced licenses with the addition of the Automatic Secure Recovery, runtime firmware validation, secure erase of user data, and CNSA compliance capability. Secure erase of user data provides erasure of all NAND data on the server, particularly useful when repurposing a server.

Other security features include the HPE Smart Memory check performed during server start-up. This check executes an authentication process and algorithms used to determine if genuine HPE memory is being used. This ensures that the memory is legitimate and does not contain security backdoor code and can be trusted for use in the platform.

Modes for HPE Secure Compute Lifecycle

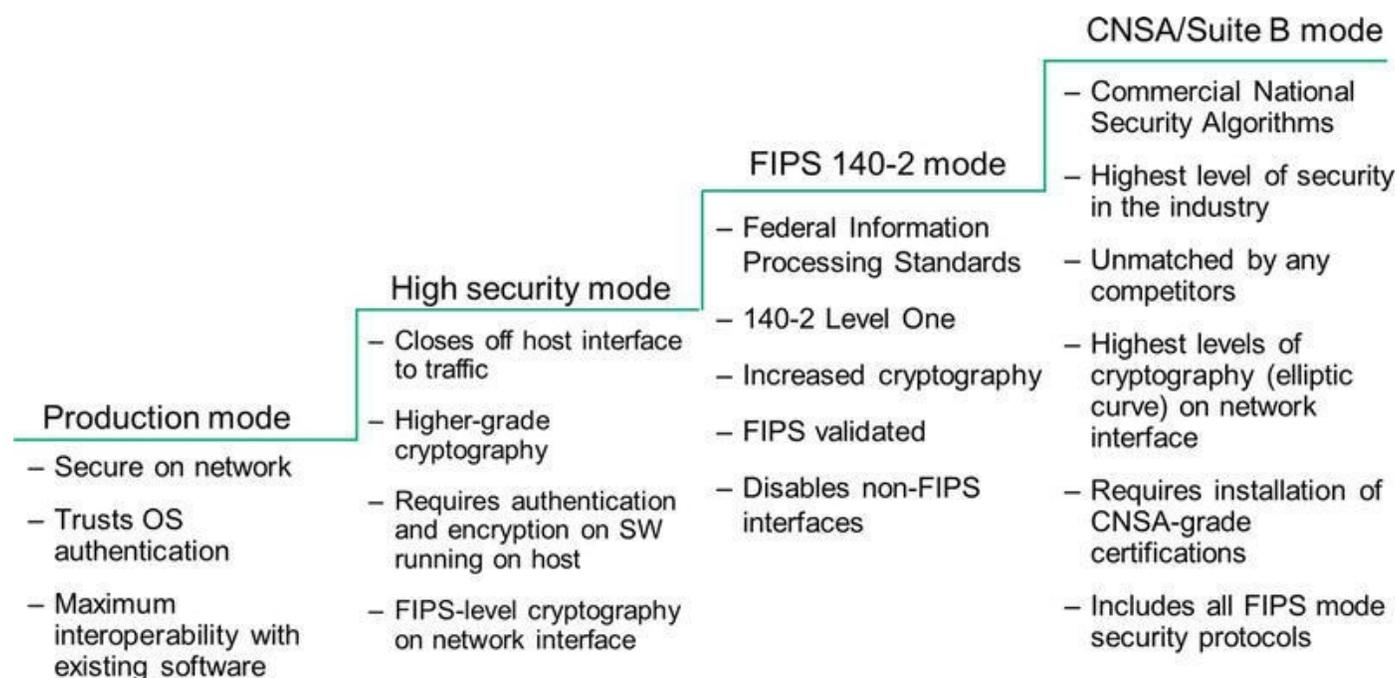


Figure 2-7 Modes for HPE Secure Compute Lifecycle

The capabilities of HPE iLO Standard that comes with every ProLiant Gen10 server gives customers the ability to configure their ProLiant Gen10 servers in one of three security states: production, high security, or FIPS, as shown in [Figure 2-7](#). With the iLO Advanced Premium Security Edition license, customers that need the highest-level encryption capabilities have a fourth security state available to them: CNSA/SuiteB.

As you move up the scale in security, the server enforces stronger encryption rules for web pages, SSH, and network communications. Note that both ends of each network connection must support the encryption rules, or they will not be able to communicate, and some interfaces are shut down to limit potential security threats.

iLO security states

Production (default)

When set to this security state:

iLO uses the factory default encryption settings.

The system maintenance switch setting to bypass iLO security (sometimes called the iLO Security Override switch) disables the password requirement for logging in to iLO.

High security

When iLO is set to this security state:

iLO enforces the use of AES ciphers over the secure channels, including secure HTTP transmissions through the browser, SSH port, iLO RESTful API, and RIBCL. When HighSecurity is enabled, you must use a supported cipher to connect to iLO through these secure channels. This security state does not affect communications and connections over less-secure channels.

Username and password restrictions for iLO RESTful API and RIBCL commands executed from the host system are enforced when iLO is configured to use this security state.

Remote Console data uses AES-128 bidirectional encryption.

The HPQLOCFG utility negotiates an SSL connection to iLO and then uses the strongest available cipher to send RIBCL scripts to iLO over the network.

You cannot connect to the server with network-based tools that do not support TLS 1.2.

The system maintenance switch setting to bypass iLO security (sometimes called the iLO Security Override switch) does not disable the password requirement for logging in to iLO.

FIPS

When iLO is set to this security state:

iLO operates in a mode intended to comply with the requirements of FIPS 140-2 level 1.

iLO enforces the use of AES ciphers over the secure channels, including secure HTTP transmissions through the browser, SSH port, iLO RESTful API, and RIBCL. When FIPS is enabled, you must use a supported cipher to connect to iLO through these secure channels. This security state does not affect communications and connections over less-secure channels.

Username and password restrictions for iLO RESTful API and RIBCL commands executed from the host system are enforced when iLO is configured to use this security state.

Remote Console data uses AES-128 bidirectional encryption.

The HPQLOCFG utility negotiates an SSL connection to iLO and then uses the strongest available cipher to send RIBCL scripts to iLO over the network.

You cannot connect to the server with network-based tools that do not support TLS 1.2.

The system maintenance switch setting to bypass iLO security (sometimes called the iLO Security Override switch) does not disable the password requirement for logging in to iLO.

CNSA/SuiteB

The SuiteB security state (also called CNSA mode) is available only when the FIPS security state is enabled.

When set to this security state:

iLO operates in a mode intended to comply with the SuiteB requirements defined by the NSA.

Security is high enough for systems used to hold United States government top secret classified data.

You cannot connect to the server with network-based tools that do not support TLS 1.2.

The system maintenance switch setting to bypass iLO security (sometimes called the iLO Security Override switch) does not disable the password requirement for logging in to iLO.

New secure server options for Gen10

Table 2-1 Secure server options

Options category	Description	Security feature	Value/Benefit
DRAM memory	32 GB Dual Rank x4 DDR4-2666	Authentication process/ algorithm	Assures genuine HPE memory and no counterfeits
Network Adapter (NIC)	Ethernet 10/25G 2-port adapter	HW RoT, Secure Boot, sanitization, authentication, and device-level firewall	Enables certified trusted and digitally signed firmware for initial configuration and updates; protects the server via the network
Solid State Storage Drive (SSD)	480 GB 6G SATA MU-2 SFF SC SSD	Digitally-signed firmware	Prevents FW attacks
Smart Array Controller	Smart Array SR Controller	HPE Smart Array Secure Encryption License	Data is encrypted at controller level
Data-at-rest encryption	HPE Smart Array SR Secure Encryption E-LTU	License for encryption capability that is integrated into Smart Array controller firmware	Encrypts data for protection against any malicious attack while not in use
Rack and Power	G2 Advanced Enterprise Racks, PDU, and Biometric Locks	Racks with front and rear doors support a variety of electronic and biometric locking solutions for physical security	Three-factor ID

[Table 2-1](#) lists several new secure server options for ProLiant Gen10 servers.

Server network interface card adapter security

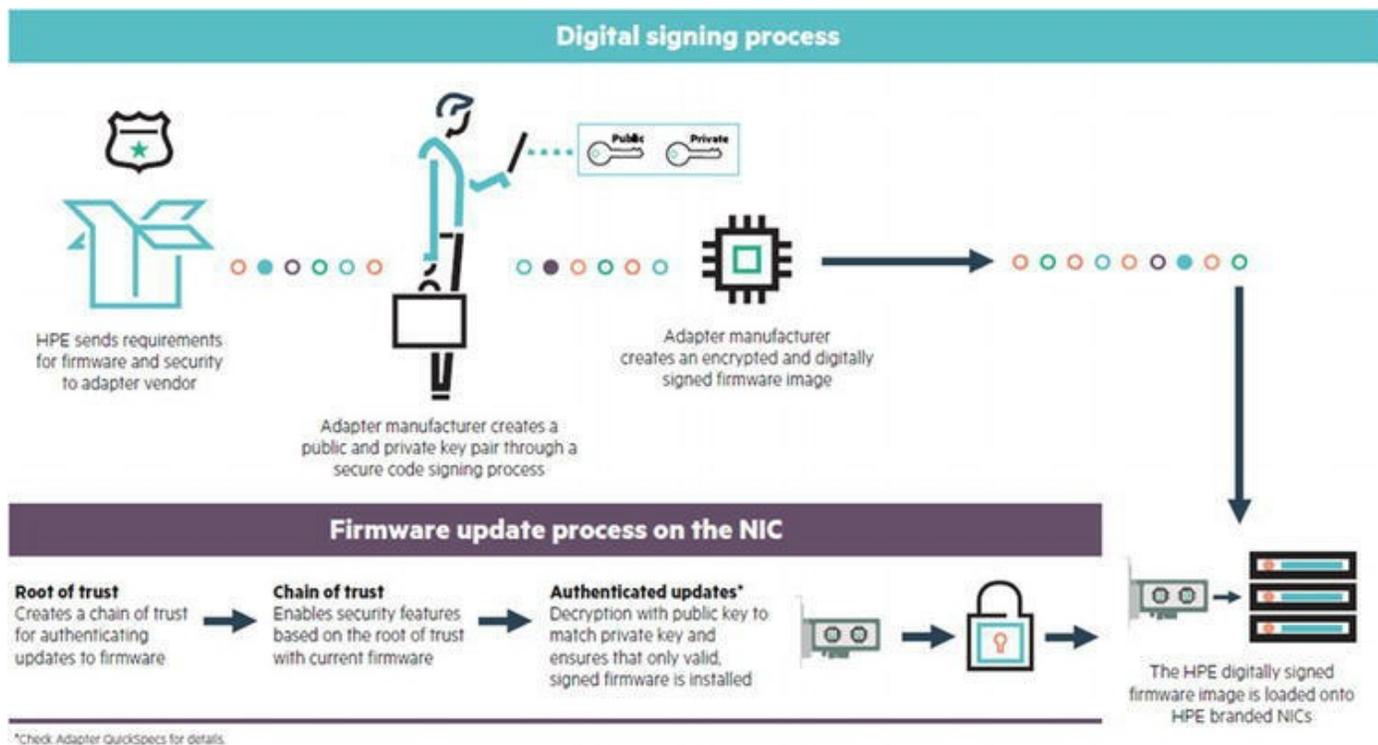


Figure 2-8 How the NIC safeguards firmware

As an extension to the capability provided by iLO 5 in the server, security capability has been extended to other components such as Ethernet network interface cards (NICs). In this case, rogue firmware images are blocked from installation. HPE sends the security requirements to the third-party manufacturer who creates a public and private key pair through a secure code signing process, as shown in [Figure 2-8](#).

Let us take a look at how the NICs help protect, detect, and recover from malicious attacks to the firmware. It all starts with digitally signing the firmware. After the updated firmware from the vendor has been encrypted with a private key, the NIC goes to work. The network adapter's Chain of Trust is created from a "true" hardware Root of Trust, which is essential for validating firmware with public keys embedded in the NIC silicon (for hardware authentication). Authenticated updates for NICs validate that signed firmware is correct and trusted to eliminate any rogue firmware installation.

HPE Smart Array Secure Encryption



- Offered on all HPE Gen8, Gen9, and Gen10 servers
- Data is encrypted at controller level
- Supported on RAID volumes only
- Local and remote encryption key management
- One license per server



Figure 2-9 Enterprise-class controller-based encryption solution for data-at-rest

As shown in [Figure 2-9](#), HPE Secure Encryption is a controller-based, enterprise-class data encryption solution that protects data at rest on bulk storage hard drives and SSDs attached to a compatible HPE Smart Array Controller. The solution is compatible with the HPE Secure Key Manager, and can operate with or without the presence of a key manager in the environment, depending on individual customer needs.

Secure Encryption secures any data deemed sensitive and requiring extra levels of protection through the application of XTS-AES 256-bit data encryption. Secure Encryption only applies to RAID volumes.

HPE SmartCache can be used in conjunction with Secure Encryption. SmartCache enables SSDs to be used as caching devices for hard drive media. Data can be accessed from the solid state drive instead of hard drives. Data stored on the SmartCache drive utilizes the same encryption methods and keys as the originating volume where the data is permanently stored, extending protection to the SmartCache drives.

Learning check

Which iLO license is required to support “Runtime Firmware Validation”?

- A. iLO Advanced license
- B. iLO Standard license
- C. iLO Advanced Premium Security Edition license
- D. iLO Foundation license

HPE Gen10 server physical security features

HPE Gen10 servers offer unique security features embedded in the silicon. The HPE Secure Compute Lifecycle offers innovations in firmware protection, malware detection, and firmware recovery.

Gen10 Trusted Platform Module



Figure 2-10 HPE TPM provides additional flexibility and customer choice

The HPE Trusted Platform Module (TPM), shown in [Figure 2-10](#), works with programs such as Microsoft Windows BitLocker to increase data security by storing the encryption startup key in hardware on the server, which provides a more secure environment by pairing the drive to the server. Pairing the drive to the server helps prevent the encrypted drive from being read if inserted in a different server. The HPE TPM can also store passwords, certificates, and encryption keys that can authenticate server hardware and software through remote attestation while the measured boot capability enhances the effectiveness of anti-malware solutions.

ProLiant Gen10 servers support:

Embedded firmware-based TPM 2.0.

- Optional Discrete TPM Module configurable for TPM 1.2 or TPM 2.0.

Embedded firmware-based TPM 2.0 called Platform Trust Technology (PTT)

- Fully functional as TPM 2.0, but is NOT certified.
- Does NOT support Command Response Buffer (CRB). Supports FIFO.
- Disabled by default.

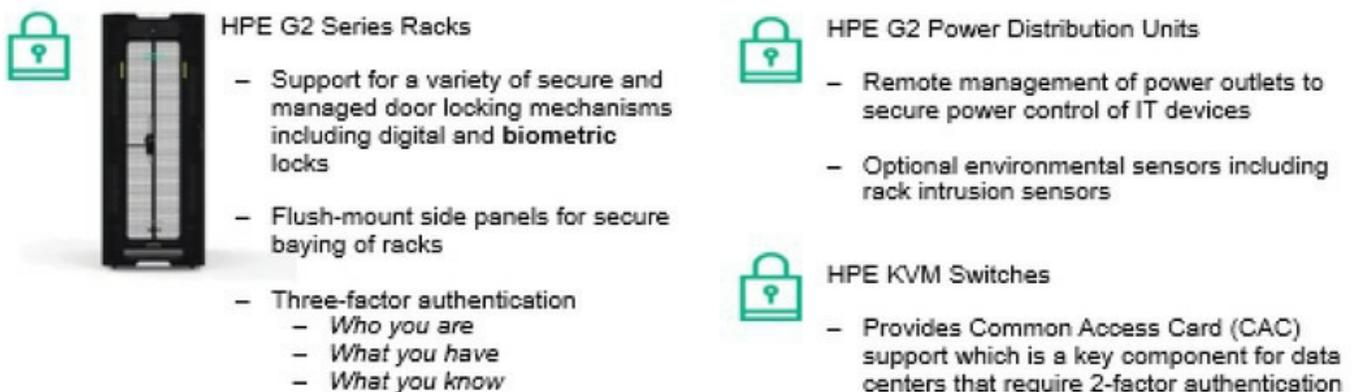
Optional Discrete TPM

- One module will support TPM 1.2 or TPM 2.0 mode.
- Can be configured for FIPS mode.
- TPM is “logically bound” to the platform by the BIOS.

Server intrusion detection

ProLiant Gen10 servers include an option for a chassis Intrusion Detection switch, which detects if the chassis access cover is opened or closed. The iLO management processor monitors the switch and if there is a change (if the access cover is either opened, closed, on AUX power), it creates a log entry noting the intrusion. You can set various alerting mechanisms (Remote SysLog, SNMP, alertmail, and so forth) to be notified of the intrusion. The switch and the iLO reporting occur as long as the server is plugged in, regardless of whether the server is powered on or off.

HPE rack and power infrastructure



The diagram illustrates three components of HPE rack and power infrastructure, each accompanied by a lock icon and a list of security features:

- HPE G2 Series Racks**
 - Support for a variety of secure and managed door locking mechanisms including digital and biometric locks
 - Flush-mount side panels for secure baying of racks
 - Three-factor authentication
 - *Who you are*
 - *What you have*
 - *What you know*
- HPE G2 Power Distribution Units**
 - Remote management of power outlets to secure power control of IT devices
 - Optional environmental sensors including rack intrusion sensors
- HPE KVM Switches**
 - Provides Common Access Card (CAC) support which is a key component for data centers that require 2-factor authentication

Figure 2-11 Advanced security for your critical physical infrastructure

Organizations that are offering co-hosting opportunities are concerned about providing

secure access to owned racks. HPE provides enhanced security features to ensure only authorized personnel have access to owned racks and infrastructure. As shown in [Figure 2-11](#), these features include: CAC, three-factor authentication, and biometric locks.

Learning check

A customer is hosting solutions on behalf of other organizations and is interested in providing rack and server physical access security. For cost reasons, they would rather not have to build secure cages to provide this level of security. How should you respond?

Write a summary of the key concepts presented in this chapter.

Summary

The HPE Security focus is concerned with ensuring the best protection for customers by embedding security technologies. This chapter is largely concerned with the HPE Server Secure Compute Lifecycle with built-in protection, detection, and recovery capability.

HPE ProLiant Gen10 introduces the concept of Silicon Root of Trust. This feature is anchored in the silicon is unique in the industry, giving HPE an impenetrable protection right through the entire supply chain. Therefore, at all stages of the manufacturing process, authenticity and integrity is assured for the server platform.

The iLO secure recovery feature is initiated automatically by iLO, recovering from a known good version of the firmware in the event of the iLO firmware validation failure.

3 HPE Server Technologies and Tools

WHAT IS IN THIS CHAPTER FOR YOU?

After completing this chapter, you should be able to:

Identify the processor technologies embedded in ProLiant Gen10 servers.

- ✓ Intel and AMD[®] processors
- ✓ Memory options
- ✓ Storage

Describe Gen10 Intelligent System Tuning (IST).

- ✓ Workload Matching
- ✓ Jitter Smoothing
- ✓ Core Boosting

Identify ProLiant Gen10 Storage technologies.

Identify ProLiant Gen10 Network technologies.

Describe HPE configuration tools.

Explain on-system management with HPE Integrated Lights-Out 5 (iLO 5).

Prelearning check

You are attending a meeting with a customer to discuss their future data center modernization plans. The conversation turns to Gen10 performance enhancements. You are asked to explain how HPE IST may provide benefit for their virtual desktop infrastructure (VDI) application. How should you respond?

A customer is planning to implement a dense Virtual Machine solution where large memory capacity is key to success. Your customer asks for confirmation that Gen10 servers can provide the capacities they need. How should you respond?

HPE ProLiant Gen10 server innovation

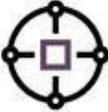
Delivering the World's Most Secure software-defined compute and converged infrastructure to run diverse workloads and applications across traditional and multi-cloud environments					
					
Most Secure Industry Standard Servers <ul style="list-style-type: none"> – Unmatched threat protection through hardware root of trust, extensive standards compliance, and supply chain attack detection – Unparalleled ability to recover firmware and OS after denial of service attempt or detection of compromised code 	Unprecedented high-speed memory capacity with persistence <ul style="list-style-type: none"> – High-capacity data acceleration with flash-backed persistent memory at multi-terabyte capacities for large data-intensive workloads – Second generation of memory-driven compute innovation on the path to the machine 	Intelligent System Tuning <ul style="list-style-type: none"> – Performance tuning to enable more workloads on more cores at a given CPU frequency for greater application licensing efficiency – Predictable latency reduction and balanced workload optimization 	New levels of compute <ul style="list-style-type: none"> – Next generation industry standard CPUs with faster processing, higher speed memory access, enhanced software-defined management, and security – Enhanced GPU levels of performance and choice 	Increased in-server storage density <ul style="list-style-type: none"> – Substantially greater NVMe capacity for large write-intensive workloads needing advanced caching/tiering – Enhanced storage density in servers with more SFF and LFF drives for collaboration and database workloads 	More efficient and easier server management <ul style="list-style-type: none"> – Enables large-scale FW deployment – Improved GUI to simplify management with industry standard APIs – Easy system debug access – Convenient warranty entitlement validation

Figure 3-1 ProLiant Gen10 servers provide relevant innovations

[Figure 3-1](#) shows an overview of the innovations HPE has made with ProLiant Gen10 servers. To meet growing business demands, HPE ProLiant Gen10 rack and tower servers redefine compute economics by delivering more compute and storage capacity, along with less power requirements, and floor space consumption. ProLiant rack and tower servers provide faster compute, memory, and IO performance, coupled with increased storage and networking performance—including lower latency.

Features of ProLiant servers

User-inspired features of ProLiant Gen10 servers include:

HPE SmartMemory—Improved performance, manageability, and efficiency; reduced downtime; and lower energy costs. SmartMemory verifies whether DIMMs have passed the HPE qualification and testing processes and determines if the memory has been optimized to run on ProLiant Gen10 servers. Key technology enhancements offered by double data rate 4 (DDR4) include:

- Increased bandwidth—DDR4 SmartMemory provides up to 2666 MT/s bandwidth for up to a 66% increase in throughput over DDR3 memory.
- Better data rate—The DDR4 specification defines eventual data rates of up to 3200

MT/s, more than 11% faster than the Gen9 2400 MT/s of DDR4 memory speed.

- 1.2-volt operation—All DDR4 memory operates at 1.2 volts, compared to 1.35 or 1.5 volts for DDR3 memory, for significant system power savings, particularly in larger memory configurations.
- 16 banks of memory per rank—Internally, the DRAMs used in DIMMs are organized into arrays of cells defined by banks, rows, and columns. DDR4 memory has 16 banks of memory in a DRAM chip compared to the eight banks in DDR3. This allows an increased number of memory requests that can be queued by the memory controller. It is one of the contributors to the lower latency of DDR4 memory.
- Encoded Rank Selection—DDR4 eliminates the work-around known as rank multiplication that DDR3 employed to enable four ranks of memory on LRDIMMs using the traditional chip select lines. When there are eight or fewer total ranks installed on a memory channel, DDR4 uses the direct chip select mode to address the correct rank. When more than eight ranks are installed, DDR4 uses a four-bit encoded chip select value for rank selection. This encoded value is interpreted by the registers on the DIMMs to determine the correct rank to enable for the memory operation. This new encoded chip select scheme allows DDR4 memory to address up to 24 memory ranks on a memory channel.
- Retry on error—DDR4 memory and new memory controllers will retry a memory request whenever a memory error or address parity error occurs. This reduces the number of system halts that may have occurred due to transient errors in previous generations of memory subsystems.
- Adaptive Double Device Data Correction (ADDDC)—Allows recovery from two sequential DRAM failures on the memory DIMMs, as well as recovery from subsequent single-bit error on the DIMM. ADDDC, also known as **HPE Fast Fault Tolerant Memory**, monitors for DRAM device failures, pinpointing the address affected and reorganizing the way data is stored to create a protected region just large enough to correct the DRAM failure. Each memory region runs independently until an error is detected. Only the combined DDDC mode region runs at reduced performance. Memory runs at max performance until the feature is activated.

Non-volatile DIMM (NVDIMM)—Memory that retains its contents even when the electrical power is removed either from an unexpected power loss, system failure, or normal shutdown. NVDIMM is constructed from DRAM components to provide low-latency data access with NAND Flash storage to provide the backup copy of the data stored in DRAM. The NAND Flash receives data from DRAM after the failure or power loss. This process is powered by the HPE Smart Storage Battery. Data are recovered from flash to DRAM after the system power is restored. Uses include

caching, high-performance storage tier, or database log files.

HPE Standard Memory—is sourced from the highest-quality DRAM. It undergoes the same rigorous testing and authentication process as HPE SmartMemory to ensure it is completely compatible with HPE servers and will perform to industry-defined specifications. Although HPE Standard Memory still offers the performance and reliability you have come to expect from HPE, it is ideal as an affordable solution for SMB and ROBO customers.

HPE Scalable Persistent Memory—You can now do much larger in-memory compute with persistence (TB instead of low 100s of GBs with NVDIMMs). Accelerate applications for fast caching and storage, reduce transaction costs for latency-sensitive workloads, and deploy bigger, more affordable data sets to gain new insights from large memory pools. Data are written to actual server memory marked as Persistent Memory via basic input/output system (BIOS)/CPLD means. A backup power source (800W power supply unit/400W battery backup unit) is used to supply power during a power-loss event and move data from DRAM to Non-Volatile Memory Express Solid State Drives (NVMe SSDs) for persistence.



Note

Requires Flex Slot UPS Power Supply.

HPE Smart Array controllers—Enterprise-class Redundant Array of Independent Disks (RAID) controllers enhance performance, data availability, and storage capacity. The current generation of Smart Array controllers is available in many form factors, including embedded on the system board, as a stand-up card in a PCIe slot, or as a mezzanine card or modular controller for HPE server blades. Smart Array controllers are designed for ProLiant customers who need fault tolerance for their direct-attached storage (DAS) environments or present the disks as-is to the host operating system. They provide reliable RAID fault tolerance for ProLiant storage to maximize system uptime. They also support advanced features such as RAID 10 Advanced Data Mirroring (ADM), which is a feature exclusive to HPE that provides a three-drive mirror. HPE Gen10 Smart Array controllers can operate in mixed, RAID and Host Bus Adapter (HBA) mode. Use both HBA and RAID mode simultaneously on one controller for added flexibility. Mixed mode for Smart Array controllers frees up a PCIe slot for other uses and is now available on P-class and E-class controllers.

HPE SmartDrives—Simplified monitoring, reduced data loss, and better serviceability. The SmartDrive product portfolio consists of both 2.5-inch small form factor (SFF) and 3.5-inch LFF SATA and SAS hard drives and SSDs and features a unique hot-plug

carrier. SmartDrive Authentication technology verifies that users are using genuine HPE drives. Most drives feature the HPE Smart Carrier with intuitive icons to report drive activity at-a-glance and a “do not remove” button that minimizes the potential for data loss caused by human error. There are three categories of drives to choose from: Entry, Midline, and Enterprise.

NVM Express (NVMe)—Non-volatile memory in the form of SSD. NVMe devices exist in the form of standard-sized PCIe expansion cards and as 2.5-inch form-factor devices that provide four-lane PCIe interface. Hosting an entire database on one or more HPE NVMe PCIe SSDs can boost IO, leverage in-memory access, reduce latency, and scale performance in-line with your processing requirements. These features, coupled with HPE Express Bay’s front accessibility and serviceability, create a flexible and dependable solution to proactively address storage needs.

HPE M.2 and M.2 Enablement Kit SSDs—A specification for internally mounted solid-state memory devices. M.2 SSDs are most suitable for boot/swap. This flexible form factor saves hot-pluggable bays for removable SSDs.

HPE Flexible Network Adapters—Easy configuration, fast serviceability. HPE Converged Network Adapters provide high throughput, low processor utilization, and the high small packet performance for Ethernet connectivity. Full hardware offload of Fibre Channel over Ethernet (FCoE) and iSCSI storage protocols for converged networking and storage data connectivity on ProLiant rack-mounted, Apollo, and Synergy servers as well as BladeSystem server blades.

HPE iLO 5 management—Rapid deployment, streamlined management. iLO 5 provides the core foundation and intelligence for all ProLiant servers. iLO can be accessed from any location using a web browser or the iLO 5 mobile application. iLO 5 simplifies server setup, enables health monitoring, provides power and thermal control, and promotes remote administration. It supports the complete lifecycle of all ProLiant servers, from initial deployment to ongoing management and service alerting.

HPE OneView—Converged management software to reduce infrastructure complexity with automation simplicity. HPE OneView is designed to accelerate IT operations for managing servers, storage, and network resources. This innovative platform is converged, software-defined, and automated to reduce operating expenses (OPEX), improve agility, and free up resources for new business initiatives. HPE OneView supports lights-out automation and provides a simple, fast, and efficient path to Infrastructure-as-a-Service and to a hybrid cloud.



Note

See the latest HPE OneView release notes for full product support.

HPE server racks—Efficient and high-performance self-cooling racks designed for maximizing usable cabinet space. HPE Intelligent Series Racks with HPE Location Discovery Services enable ProLiant Gen10 and Gen9 servers to self-identify and inventory to optimize workload placement. Location Discovery Services is a solution that merges the physical rack with IT equipment to provide automated location awareness of installed servers for advanced asset management and workload placement.

HPE power and cooling management—Reduced power requirements, reclaimed resources, and fewer errors. HPE has developed a full line of power-management products that protect and manage computer systems ranging from individual servers to distributed enterprises. Intelligent Power Discovery Services combine an Intelligent Power Distribution Unit (iPDU) and Flexible Slot (Flex Slot) Platinum/Titanium Power Supplies with HPE software to create an automated, energy-aware network between IT systems and facilities. Intelligent Power Discovery Services with iPDUs automatically track power usage and document configurations to increase system uptime and reduce the risk of outages.

HPE ProLiant product lines

HPE ProLiant servers are based on industry-standard x86 architectures and are separated into five main product lines.

DL servers—Density-optimized for stacking in rack-mounted server environments ideal for multi-server deployments

ML servers—Tower- and rack-based servers designed with capacity for maximum internal expansion

BL servers—Blade servers that fit within the HPE BladeSystem

SY servers—Compute modules that fit within the HPE Synergy solution

XL servers—Compute modules that fit within the Apollo 2000 HPC solution

ProLiant servers are also split into several series that denote processor configuration.

The 10, 300, and 400 series comprise single- and dual-socket systems.

The 500 and 600 series comprise quad-socket-capable systems.

Models where the last digit is zero are built with Intel[®] Xeon[®] processors inside (such as DL380), and those with five at the end are built to support AMD Opteron processors (such as DL385).

HPE ProLiant Gen10 benefits

ProLiant Gen10 servers are built using next-generation HPE Server Options. HPE Server Memory, Server Storage, Persistent Memory, Server Networking, and Power Supplies deliver key innovations that enable the performance, security, reliability, and efficiency found in HPE next-generation servers. In turn, HPE servers, storage, and networking are powered and protected by the next-generation Rack and Power Infrastructure.

Table 3-1 Gen10 benefits

	Increase capacity and lower TCO	<ul style="list-style-type: none">• Twice as much compute power per watt using 12 GB/s SAS SSDs and HPE DDR4 SmartMemory• Flexible choices in storage, networking, and power supplies• Improved security across server lifecycle• Persistent Storage at memory speeds• Increased storage capacity, compute density, GPUs
	Move faster with management automation	<ul style="list-style-type: none">• Reliable, secure embedded management• Faster provisioning with software-defined HPE OneView management• Improved GUI and Industry-Standard APIs, faster problem resolution
	Accelerate application performance	<ul style="list-style-type: none">• Support for 1 million IOPS with 12 Gb/s HPE Smart Array fast controllers• Quick data access with up to four times read and write workload acceleration provided by HPE SmartCache• Persistent Storage at memory speeds• IST• Up to 14% better memory performance provided by HPE DDR4 SmartMemory and ProLiant rack servers

[Table 3-1](#) presents some benefits of HPE ProLiant Gen10 servers.

Intel[®] Xeon[®] Processor Scalable Family

Features	Intel® Xeon® E5-2600 v4	Intel® Xeon® Processor Scalable Family (Skylake Server)
Cores Per Socket	Up to 22 cores	Up to 28 cores
Threads Per Socket	Up to 44 threads	Up to 56 threads
L3 Cache	Up to 55 MB	Up to 38.5 MB (non-inclusive)
Socket Interconnect Speed	2x QPI @ 9.6 GT/s	Up to 3x UPI @ 10.4 GT/s
PCIe* Lanes/ Controllers/Speed(GT/s)	40 / 10 / PCIe* 3.0 (8 GT/s)	48 / 12 / PCIe 3.0 (8 GT/s)
Memory Capability	4 channels, 3 DIMMs per channel (RDIMMs, LRDIMMs, or 3DS LRDIMMs)	6 channels, 2 DIMMs per channel (RDIMMs, LRDIMMs, or 3DS LRDIMMs)
Max Memory Speed (MT/s)	Up to 2400	Up to 2666
TDP (W)	55W-145W	70W-205W

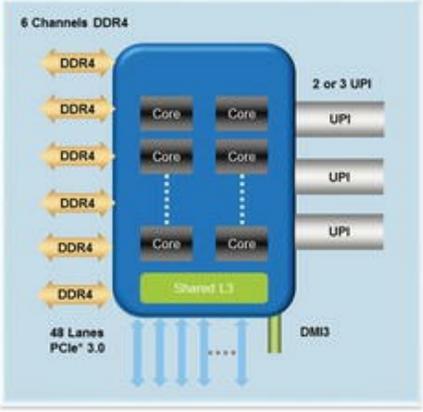


Figure 3-2 Comparison of Intel® Xeon® E5-2600 v4 and Skylake Server

Figure 3-2 details differences between the previous generation of Xeon® processor compared to the Xeon® Processor Scalable Family, specifically Skylake-SP.

On the Skylake-SP, the core count rises from 22 to 28 cores which with Hyperthreading permit up to 56 threads. The L3 cache is reduced to 38.5 MB to create space for the increased per core L2 cache at 1 MB. An architectural change results in the L3 cache operating as non-inclusive.



Note

For more information on cache inclusion policy, enter the following URL into your browser:

https://en.wikipedia.org/wiki/Cache_inclusion_policy#Inclusive_Policy

The inter-processor interconnect speed has increased from 2× QPI (Quick Path Interconnect) at 9.6 GT/s for the E5 processor to up to 3× UPI (Ultra Path Interconnect) at 10.4 GT/s in Skylake-SP.

The most radical change is related to the increase in memory channels, rising to six channels supporting two DIMMs per channel in Skylake-SP instead of four memory channels with three DIMMS per channel in the E5 processor. In both cases, DDR4 RDIMMS, LRDIMMS, or 3DS LRDIMMS are supported. However, the maximum memory speed rising from up to 2400 MT/s for the E5 processor for up to 2666 MT/s in Skylake-SP. The memory DIMM capacity also increases, rising to support 128 GB DIMMS.

The number of IO channels has increased to 48 PCIe 3.0 for Skylake-SP as opposed to 40 PCIe 3.0 for the E5 processor.

The total power consumption has increased though, rising to between 70W and 205W.

Intel® Xeon® Skylake SP (two socket) platform

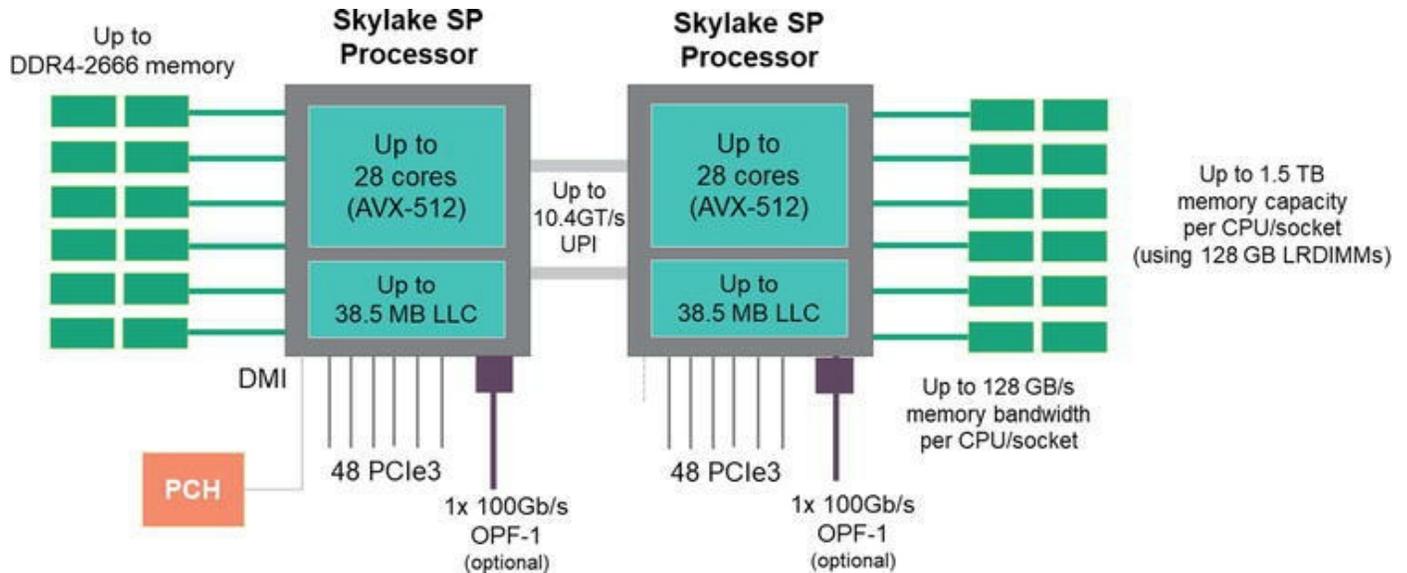


Figure 3-3 Skylake-SP

The Intel® Xeon® Processor Scalable platform provides the foundation for a powerful data center platform that creates significant advances in agility, scalability, and security.

As shown in [Figure 3-3](#), this generation introduces increased memory capacity and performance due to the six memory channel design and the advanced inter-processor communication UPI.

Intel® Xeon® Processor Scalable Family SKUs

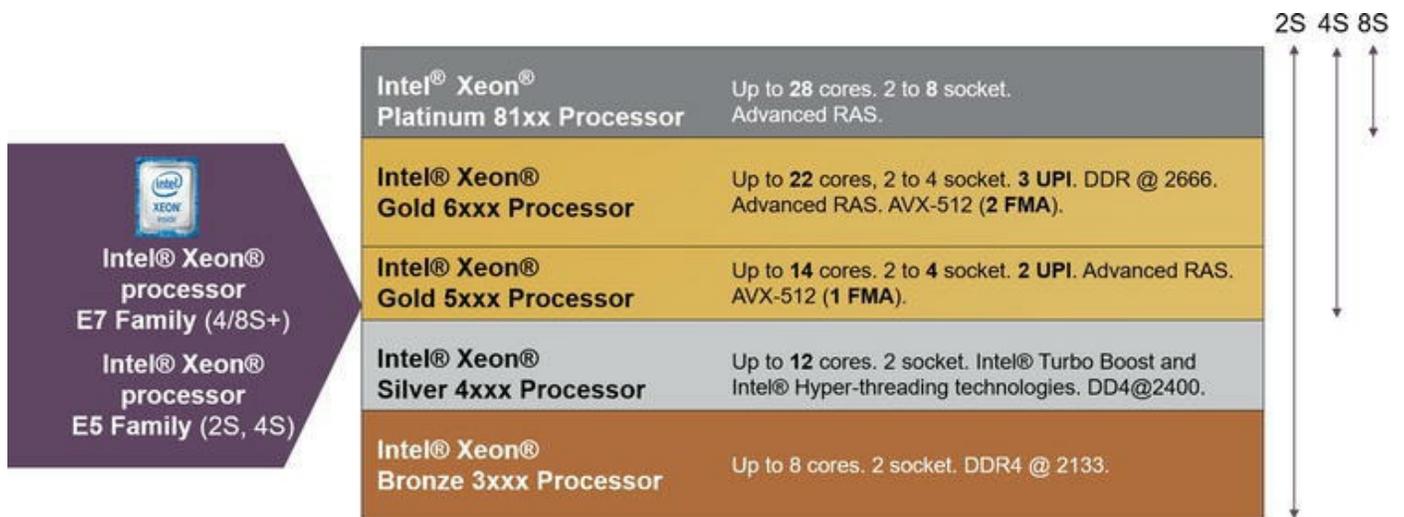


Figure 3-4 Intel® metal-based SKU system: Platinum, Gold, Silver, and Bronze

In previous generations, two- and four-socket processor families were segregated into different product lines. One of the big changes with the Intel® Xeon® Processor Scalable Family is that it includes all the processor models associated with this new generation. The processors from Intel® Xeon® Processor Scalable Family are scalable from a two-socket configuration to an eight-socket configuration. They are Intel’s platform of choice for the most scalable and reliable performance with the greatest variety of features and integrations designed to meet the needs of the widest variety of workloads.

With this transition, Intel has moved to the naming convention Platinum, Gold, Silver, and Bronze as shown in [Figure 3-4](#).

Platinum SKUs offer the best performance across the broadest workload, including:

Mission-critical and real-time analytics applications

Artificial Intelligence (Deep Learning)

Workload-optimized solutions for compute/network and storage in hybrid-cloud deployments and effortless scalability for two-, four-, or eight-socket implementations

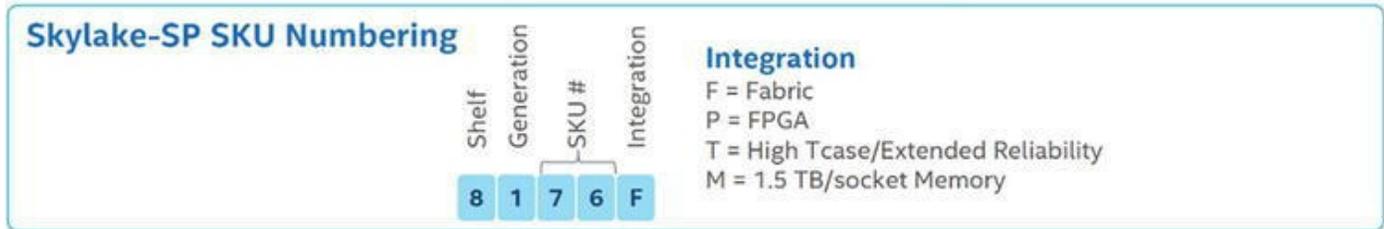
Ultimate in hardware-enhanced security and reliability

Gold SKUs are where many mainstream workloads start seeing breakaway gains in performance (through high-memory speed and capacity support, more interconnect and accelerator engines integrated, and advanced levels of reliability, availability, and serviceability).

Silver SKUs have solid compute capabilities and energy efficiency.

Bronze SKUs enable an entry point to the family (performance increase from the Xeon® E3 family).

Intel® Xeon® Processor Scalable Family—New brand



31xx (Bronze)	41xx (Silver)
• 2S-2UPI	• 2S-2UPI
• 6-ch DDR4 @ 2133	• 6-ch DDR4 @ 2400
• 2 UPI links @ 9.6GT/s	• 2 UPI links @ 9.6GT/s
	• Intel® Turbo Boost Technology
	• Intel® HT Technology
• Intel® AVX-512 (1 512-bit FMA)	• Intel® AVX-512 (1 512-bit FMA)
• 48 lanes PCIe* Gen3	• 48 lanes PCIe* Gen3
• Standard RAS	• Standard RAS

51xx (Gold)	61xx (Gold)	81xx (Platinum)
• 2S-2UPI, 4S-2UPI capability	• 2S-2UPI, 2S-3UPI, 4S-2UPI, 4S-3UPI capability	• 2S-2UPI, 2S-3UPI, 4S-2UPI, 4S-3UPI, 8S-3UPI capability
• 6-ch DDR4 @ 2400	• 6-ch DDR4 @ 2666	• 6-ch DDR4 @ 2666
• 2 UPI links @ 10.4GT/s	• 3 UPI links @ 10.4GT/s	• 3 UPI links @ 10.4GT/s
• Intel® Turbo Boost Technology	• Intel® Turbo Boost Technology	• Intel® Turbo Boost Technology
• Intel® HT Technology	• Intel® HT Technology	• Intel® HT Technology
• Intel® AVX-512 (1 512-bit FMA)	• Intel® AVX-512 (2 512-bit FMA)	• Intel® AVX-512 (2 512-bit FMA)
• 48 lanes PCIe* Gen3	• 48 lanes PCIe* Gen3	• 48 lanes PCIe* Gen3
	• Node Controller Support	• Node Controller Support
• Advanced RAS	• Advanced RAS	• Advanced RAS

Figure 3-5 Intel® Xeon® Processor Scalable Family SKU numbering system

Figure 3-5 shows the Xeon Processor Scalable Family product numbering.

Learning check

A customer is planning to implement a dense Virtual Machine solution where large memory capacity is key to success. Your customer asks for confirmation that Gen10 servers can provide the capacities they need. How should you respond?

HPE IST

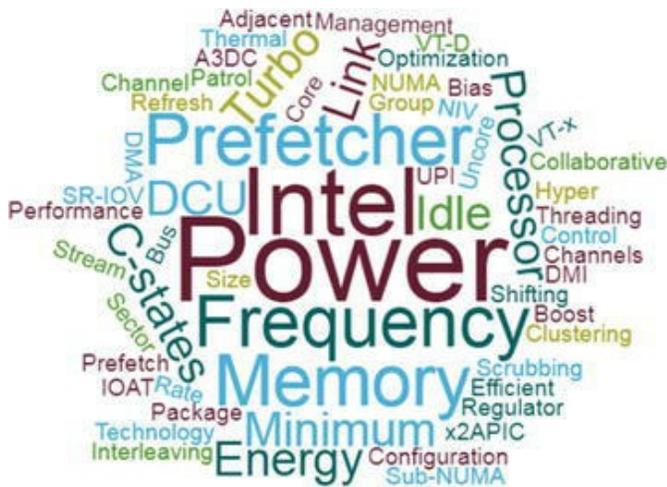


Figure 3-6 HPE IST

Developed through the HPE partnership with Intel, IST is a new set of server tuning technologies that enable you to dynamically configure server resources to match specific workloads for significant performance improvements, real savings, and a more intelligent server environment.

IST consists of three elements, as shown in [Figure 3-6](#). Features can be utilized independently or combined to deliver higher levels of performance and lower costs to your server environment. Let us look at each in greater depth.

Workload matching



- Tuning your server through hundreds of possible BIOS configurations is a challenge.
- Understanding how each setting interacts with one another and impacts your workload can leave users overwhelmed.

Figure 3-7 The challenge with BIOS settings

HPE ProLiant Gen10 servers are capable of supporting a very wide range of workloads. To achieve this degree of flexibility requires accurate selection of BIOS configuration parameters with the aim to provide the best match to workload. As shown in [Figure 3-7](#), this process can be difficult due to the many different permutations and dependencies. There are a multitude of BIOS settings that can be adjusted to change performance, power, latency, memory, and other variables. This could take hundreds of hours to find the right combination.

HPE has provided a simple, quick way to achieve the optimum selection to match the required workload. Selecting a preconfigured workload profile to automatically tune internal server resources to experience up to a 9% performance improvement over server default settings¹.

HPE servers are set in the factory to use the most common performance settings seen across the majority of application workloads, but at the same time, turn on most power-management settings that have minimal impact on overall performance.

You can choose the workload profile to suit low-latency workloads, workloads that heavily use GPUs, workloads that do not require high performance and can, therefore, benefit from power and cooling reducing setting, and many others. Therefore, servers can be tuned to match specific workloads.

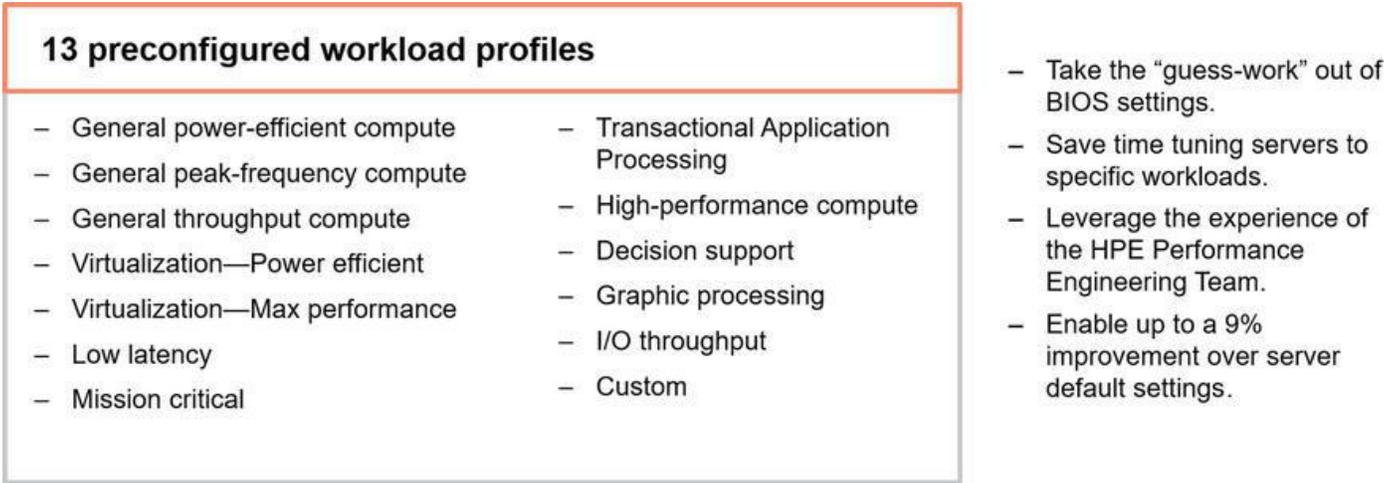


Figure 3-8 Preconfigured workload profiles



Note

¹HPE Internal testing from Performance Engineering Benchmarking team, April 2017.

HPE Workload matching allows you to simply choose a workload profile, and the server will automatically configure the BIOS settings to that workload. You can leverage the experience of HPE’s performance engineering teams and take the guesswork out of BIOS tuning. Of course, you can always set up a profile, and then tweak it to match your specific workload.

[Figure 3-8](#) lists the preconfigured workload profiles. They are described as:

General power-efficient compute (Default)—The default out-of-the-box profile. It uses the most common performance settings seen across the majority of application workloads, but at the same time, turns on most power-management settings that have minimal impact on overall performance.

General peak-frequency compute—This profile setting is intended for general purpose workloads that benefit from achieving the maximum frequency for any individual core at any time. Power-management settings are applied to ensure that any component frequency upside can be achieved readily.

General throughput compute—This profile setting is intended for general purpose workloads where maximum sustained processor throughput is desired. Most power-management settings are disabled.

Virtualization—Power efficient—This profile setting is intended to ensure that all

available virtualization technologies are enabled. Certain virtualization technologies can have possible performance impacts to non-virtualized environments and are often disabled in other profiles.

Virtualization—Max performance—This profile setting is intended to ensure that all available virtualization technologies are enabled and configured for maximum performance.

Low latency—This profile setting is intended for those customers who require the lowest amount of computational latency and a real-time environment. Power management is disabled as well as other management features that introduce any computational jitter.

Mission critical—This profile setting is intended for those customers that value mission-critical reliability, availability and serviceability (RAS) features above computational performance.

Transactional Application Processing—This profile setting is intended for business processing customer environments. These customers are typically utilizing online transaction processing (OLTP) applications that require a database back-end.

High-performance compute (HPC)—This profile setting is specifically aimed at HPC customer environments. These customers typically run in a clustered environment with each node performing at maximum utilization for extended periods of time to run large-scale scientific and engineering workloads.

Decision support—This profile setting is targeted for customers with Enterprise Business Database (Business Intelligence) workloads that are focused on operating and/or accessing data warehouses, such as data mining and online analytical processing (OLAP).

Graphics processing—This profile is intended for configurations that use GPUs and depend on maximum bandwidth between IO and memory.

I/O throughput—This profile is intended for configurations that depend on maximum throughput between IO and memory.

Custom—This profile unlocks any profile dependencies and allows the customer to optimize every setting independently.

What is Jitter?

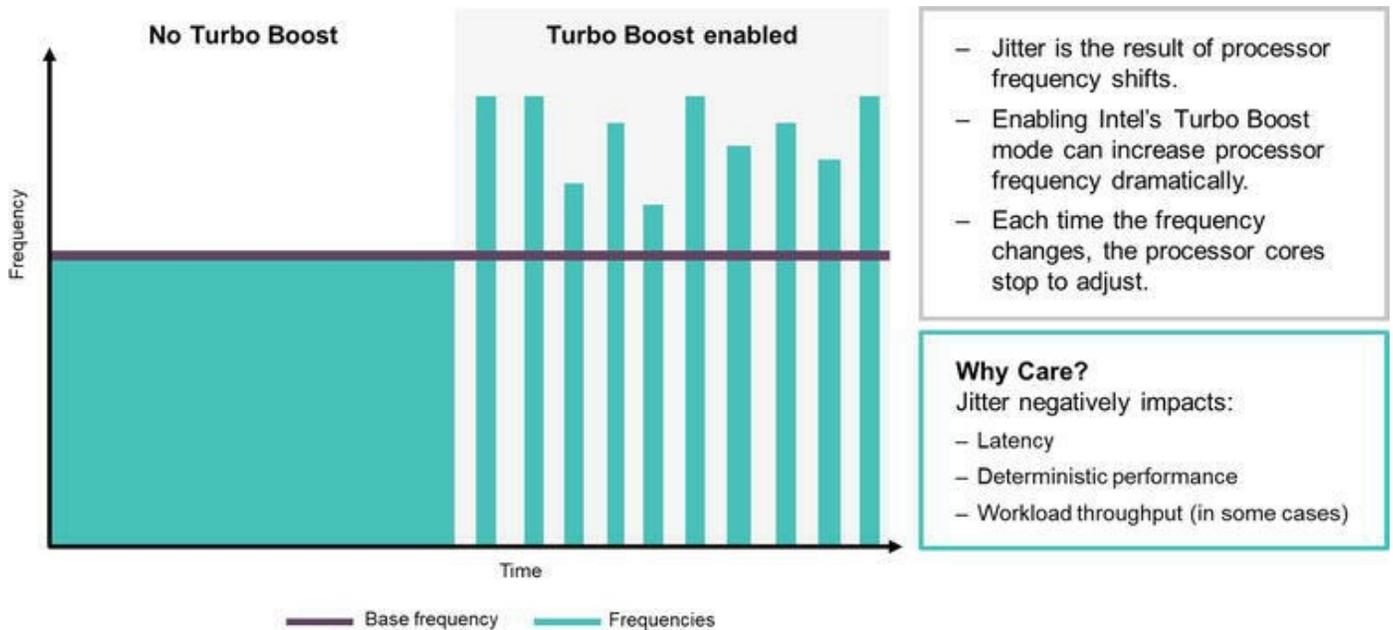


Figure 3-9 Jitter

As shown in [Figure 3-9](#), engaging processor turbo boost can cause frequency fluctuations or “jitter” which results in a constant struggle between maximum output and deterministic performance needs. HPE Jitter Smoothing technology minimizes processor frequency fluctuation to reduce latency and deliver deterministic and reliable performance. In variable workloads where processor frequency changes occur often, Jitter Smoothing can improve overall throughput above Turbo Boost mode alone.

Jitter Smoothing is ideal for high-frequency traders, high-performance computing, and workloads where processor frequency is highly variable. This feature is available on all Gen10 servers with iLO 5 and on iLO Advanced or above license.

What is Jitter and why is it important?

Enabling Intel’s Turbo Boost can increase processor frequency dramatically (for example, from 3.2 GHz to 3.7 GHz).

At high/turbo frequencies, changes in workloads require the processor to change frequency.

Each change in processor frequency requires the processor to utilize thousands of clock cycles to begin to execute at the new frequency.

The time to change processor frequency change can take between 10µs and 20µs (microseconds).

Jitter is:

- Frequency fluctuation
- Processor “downtime” during each frequency change
- Not good

Jitter impacts latency. For example, as the processor thrashes (changes frequency many times), deterministic performance is negatively impacted.

Jitter impacts overall workload throughput (for example, each time the processor changes frequency, it takes 10us to 20us. This adds up if the processor is changing frequencies a lot).

Which customers care about jitter?

High-Frequency Traders (HFT). HFT application algorithms depend on deterministic performance. Latency negatively impacts performance, and therefore, turbo mode is not tolerated.

Jitter Smoothing

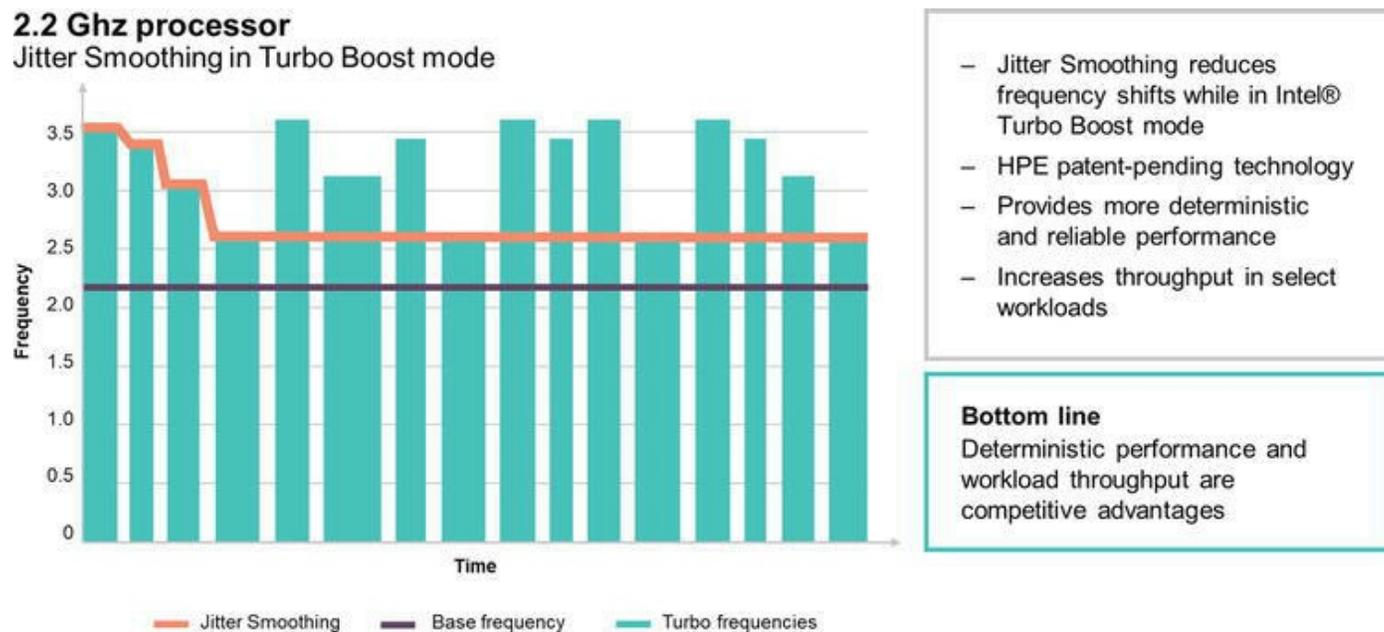


Figure 3-10 Jitter Smoothing in Turbo Boost mode

[Figure 3-10](#) shows the effect of Jitter Smoothing. Jitter Smoothing can be accessed through the RBSU or through the Representational State Transfer (REST) interface for BIOS parameters.

While Jitter Smoothing can be applied with any workload profile, the low latency and

HPC workload profiles are designed to work best with Jitter Smoothing. The difference is that, today, the low-latency workload profile disables Intel's Turbo mode and the HPC workload profile does not.

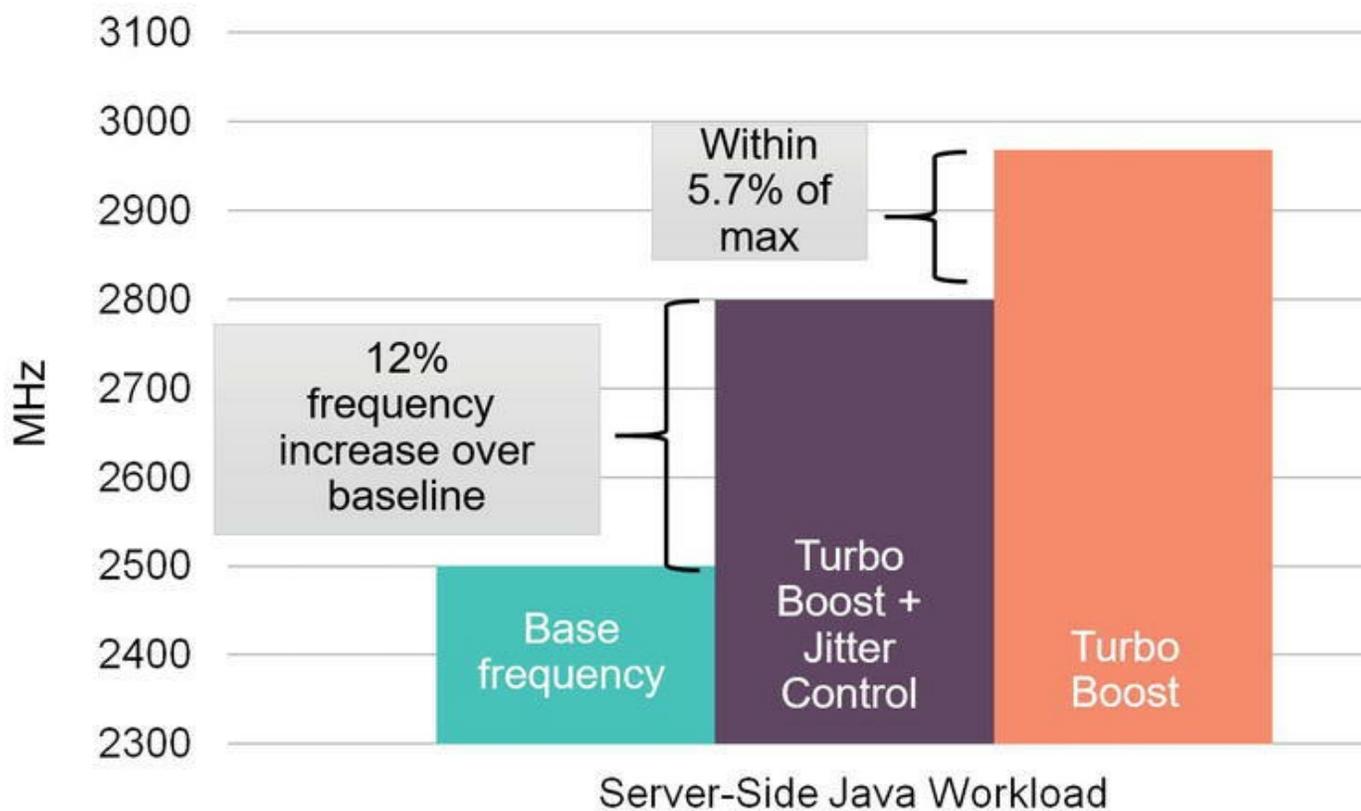


Figure 3-11 Actual results for Server-Side Java Workload

[Figure 3-11](#) shows actual results for a Server-Side Java Workload. The combination of Jitter Smoothing and Turbo Boost mode yielded better performance (higher CPU frequency) than the baseline while retaining deterministic frequencies.

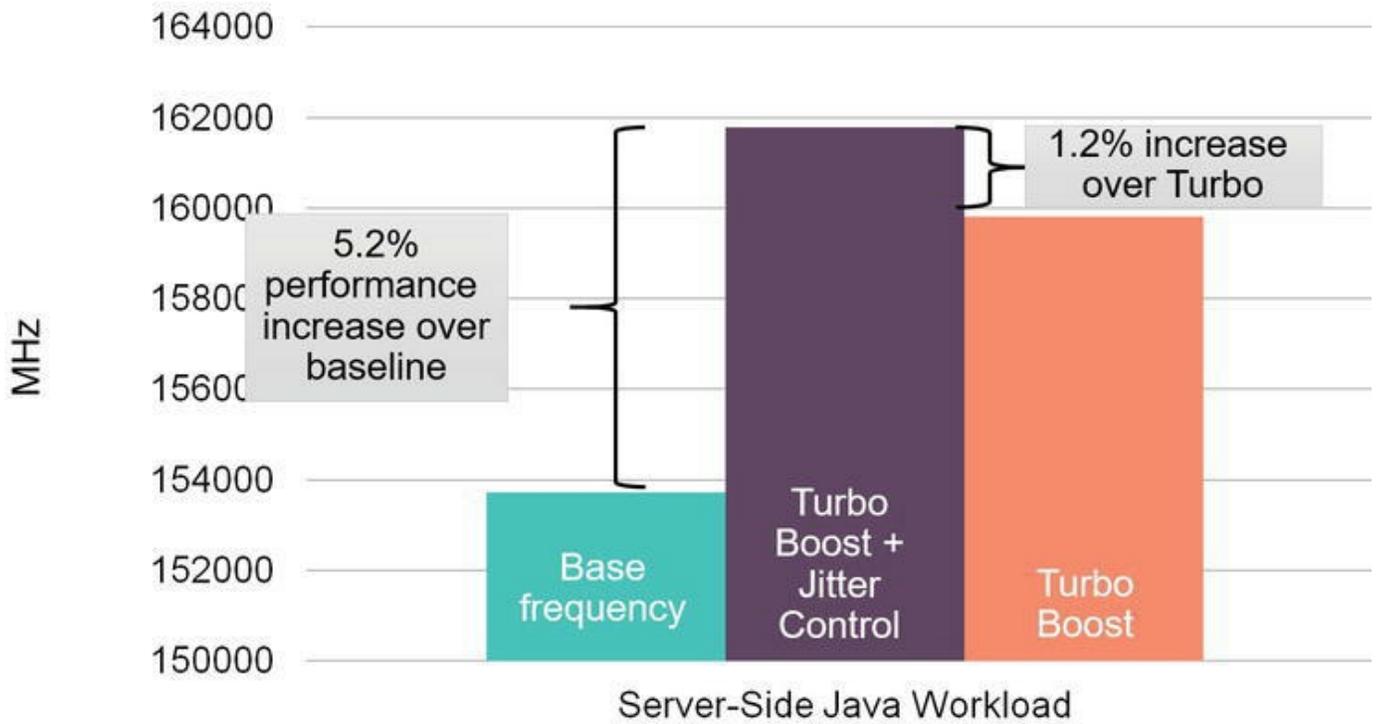


Figure 3-12 More results for Server-Side Java Workload

In some cases, Jitter Smoothing + Turbo Boost mode can improve overall workload throughput above Turbo Boost mode alone, as shown in [Figure 3-12](#).

Turbo Boosting

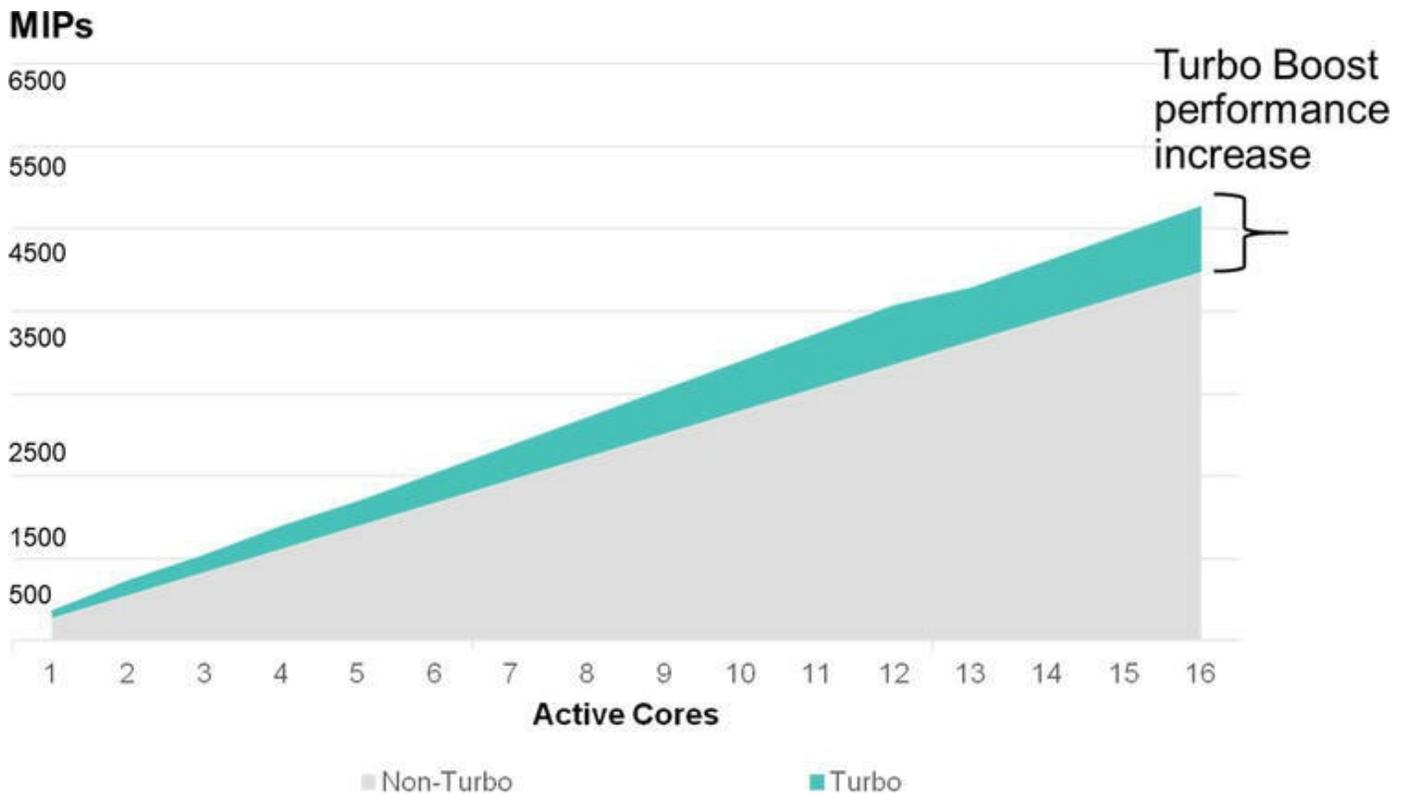


Figure 3-13 Turbo Boosting

[Figure 3-13](#) shows that enabling Intel's® Turbo Boost mode can increase processor frequency and total millions of instructions per second (MIPs). However, as more cores are activated, this reduces the maximum frequency you can achieve, and therefore, MIPs increases are not linear.

Enabling Intel's Turbo Boost mode allows the processor to opportunistically attain higher frequencies (for example, depending on power, cooling, and so forth). However, as more CPU cores are activated, more heat is generated, and therefore, the highest turbo frequencies are not permitted.

For example: if only one core in the CPU is active, and Turbo Boost mode is enabled, the CPU can opportunistically attain the highest frequency specified by Intel. As more cores become active, the thermal design power profile (TDP) of the processor will lower the maximum frequency that can be opportunistically attained.

The key point here is that the standard roadmap Intel® SKU is designed to operate in many environments. Therefore, the profiles are set to reduce maximum turbo frequency to prevent the processor from overheating. While enabling Turbo Boost mode can improve MIPs numbers, the gain in performance is not linear.

Advanced heat sink design and efficient cooling systems are beneficial in ensuring high-frequency selection as more cores become active.

Core Boosting

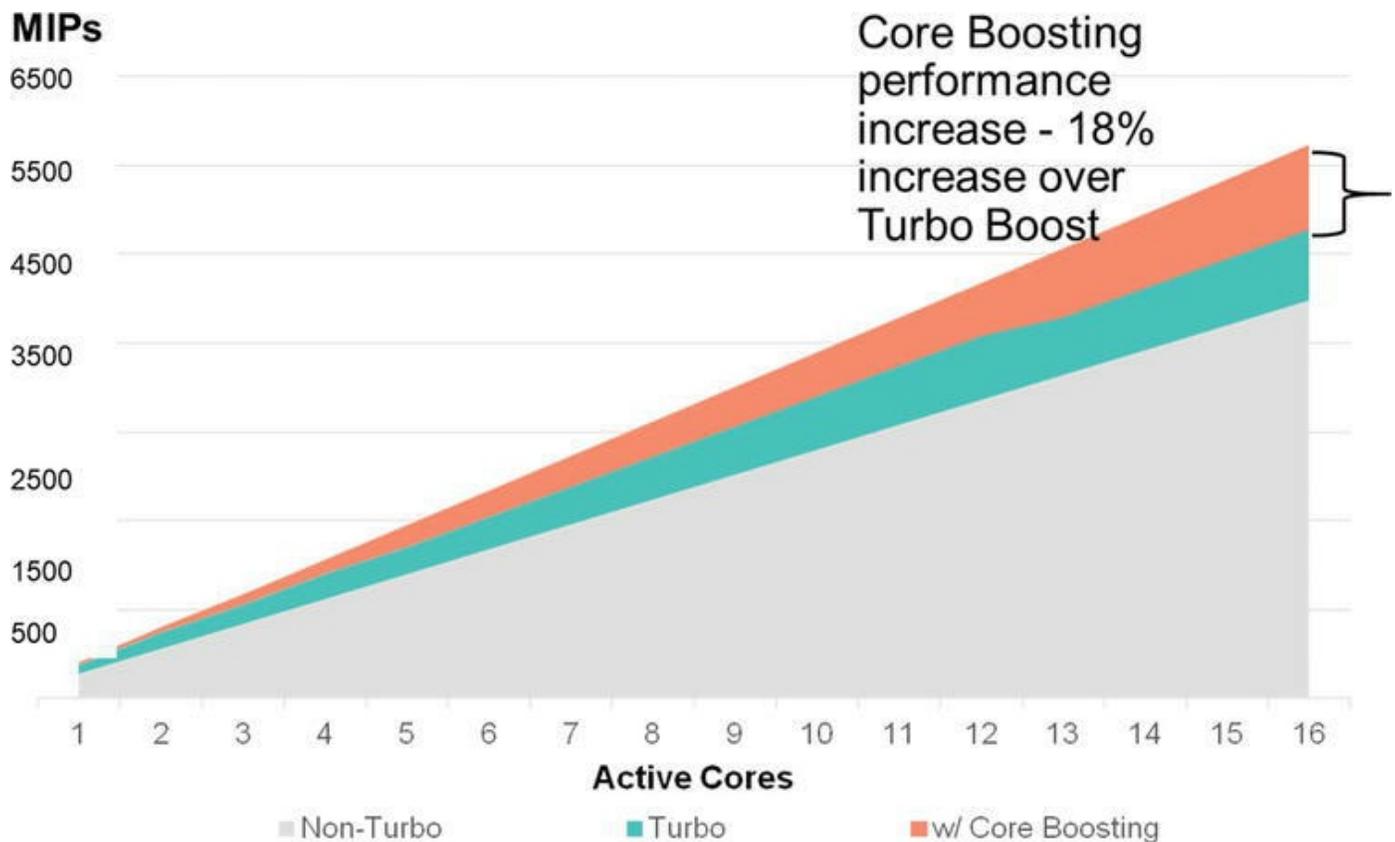


Figure 3-14 Core Boosting

Core Boosting enables higher processor frequencies on more active cores than standard Intel® processor profiles. Higher performance across fewer processors can result in significant savings when it comes to core-based licensing. In some cases, customers with core-based licensing products like Oracle can experience up to \$100K in annual savings or \$500K over five years.

[Figure 3-14](#) shows that Core Boosting enables higher performance across more processor cores than Turbo Boosting alone. This means that performance can be improved by running workloads on more active cores at higher frequencies. Core Boosting works without overclocking, which maintains Intel® warranty and reliability. Licensing costs can be reduced for core- and processor-based licensing workloads.

Core Boosting is ideal for virtualized environments, Big Data workloads, core-based licensing applications, and workloads where performance is a competitive advantage. It is available on Gen10 platforms paired with select Intel® processors, iLO 5, and an iLO Advanced or above license.

Learning check

You are attending a meeting with a customer to discuss their future data center modernization plans. The conversation turns to Gen10 performance enhancements. You are asked to explain how HPE IST may provide benefit for their VDI application. How should you respond?

Convergence of memory and storage

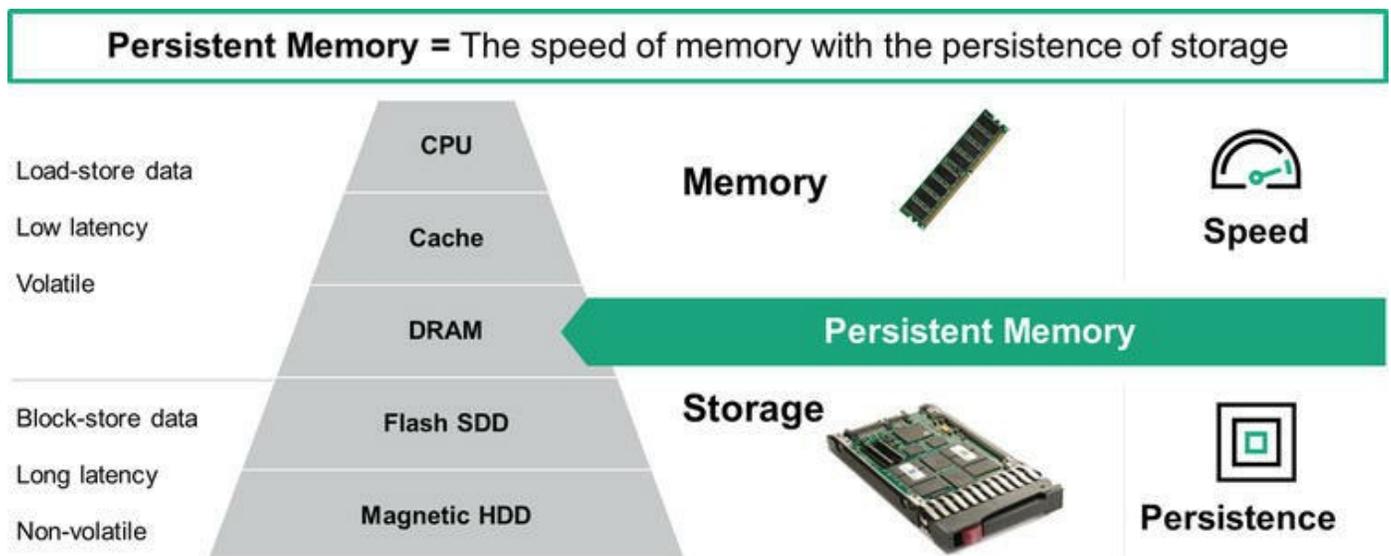


Figure 3-15 Persistent Memory

Recent, huge increases in the power of CPUs has left the rest of the server architecture lagging behind. The goal for future server design is to mitigate these deficiencies or redesign the server so that latencies incurred in the IO subsystem are eliminated. To put this into perspective, a single reference to a directly connected hard drive can be measured in milliseconds, compared to less than tens of nanoseconds when visiting the same data when transferred to DRAM. While SSDs seek to eliminate the delay incurred during the mechanical operation of hard disk drive (HDD), the IO subsystem offers no improvement. One way the HDD/SSD latency can be mitigated is to load more data into a larger memory system. However, the performance gain achieved by doing this is undermined by the volatile nature of current DRAM memory technology, namely, that DRAM loses information in the event of a power failure. The development of PCIe

class of SSD, utilizing NAND Flash memory, has resulted in significant improvement in IO performance. However, until genuine non-volatile DRAM technology becomes widely available, the solution to this problem is the NVDIMM.

The industry-standard innovation NVDIMM-N, previously known as type 1, combines DRAM and NAND Flash onto a single DIMM. As shown in [Figure 3-15](#), while the DRAM provides the conventional large capacity, high endurance, low latency, and volatile memory; the Flash memory is used as persistent store. The relatively low endurance, capacity, and performance of NAND Flash compared to DRAM seem to be at odds, though since NAND is only written upon loss of power or some other trigger, endurance should not be an issue. DRAM contents are restored as part of the BIOS boot process.

HPE innovation includes:

HPE BIOS detects and prevents system errors.

HPE byte-addressable memory provides a standard interface for software vendors.

NVDIMM/Persistency Controller moves data from DRAM to Flash upon power loss or other trigger.

HPE Smart Storage Battery or “UPS Power Supply” provides the backup power to HPE NVDIMMs.

HPE Gen10 DDR4 memory technologies

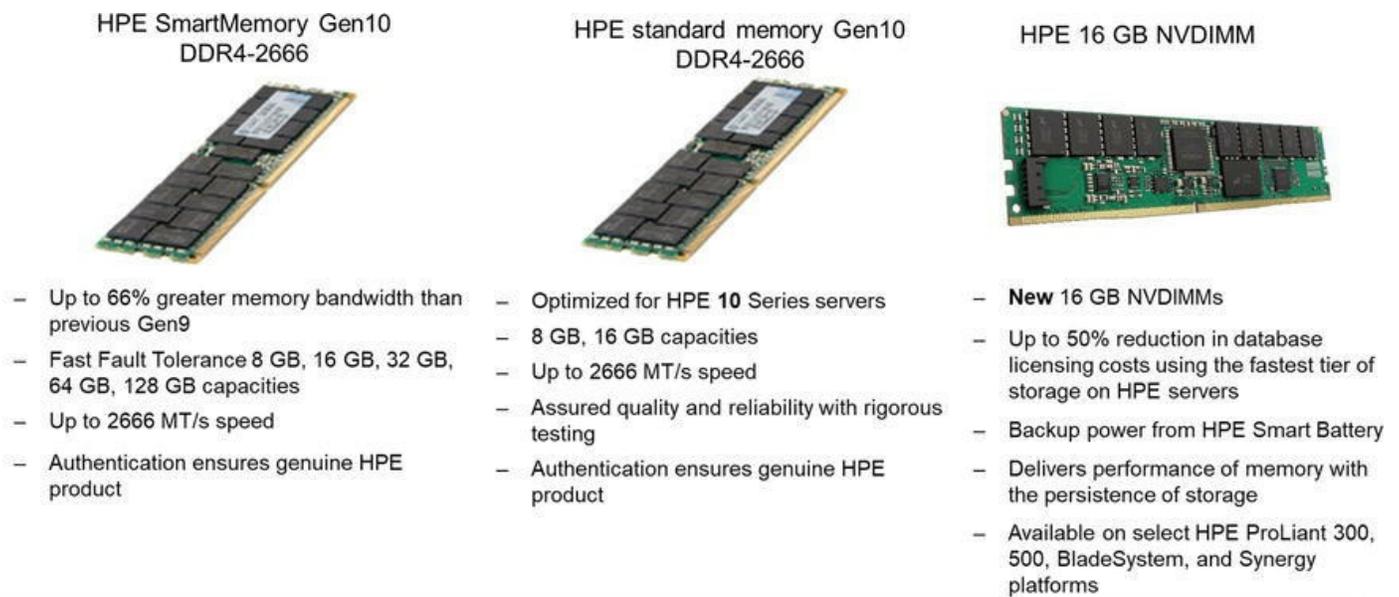


Figure 3-16 Gen10 DDR4 SmartMemory, standard memory, and NVDIMM

[Figure 3-16](#) shows some of the features of HPE SmartMemory, HPE Standard Memory, and HPE NVDIMMs.

HPE SmartMemory

Gen10 DDR4-2666 RDIMM and LRDIMM serve to unlock higher levels of server performance. Features include:

Up to 66% greater memory bandwidth than the previous DDR3 technology

Enables customers to increase total memory capacity by 400%, improving performance for memory-intensive applications

Offering 8 GB, 16 GB, 32 GB, 64 GB, and 128 GB capacities

Up to 2666 MT/s speed

Fast Fault Tolerance reduces unplanned memory outages and proactive DIMM swaps

Proactively identifies issues that may cause uncorrectable errors and unplanned downtime with HPE Advanced Error Detection

HPE SmartMemory is ideal for Enterprise, HPC, and SMB computing environments.

HPE 16 GB NVDIMM

Features of the HPE 16 GB NVDIMM include:

Highest performing Persistent Memory in the market running at DRAM speeds

Up to 50% reduction in database licensing costs, using the fastest tier of storage on HPE servers

Backup power from HPE Smart Battery

Delivers performance of memory with the persistence of storage

Available on select HPE ProLiant 300, 500, BladeSystem, and Synergy platforms

Ideal for databases, write caching, and write-intensive workloads constrained by smaller bottlenecks

HPE Persistent Memory portfolio

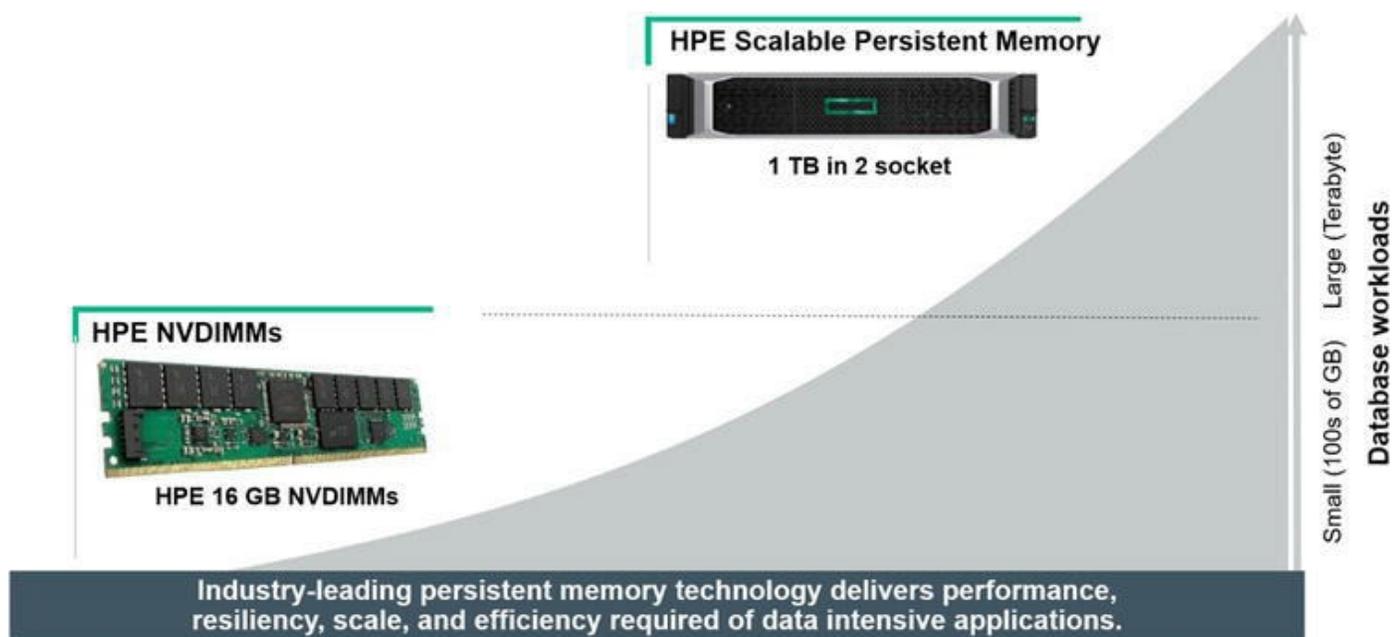


Figure 3-17 HPE Persistent Memory portfolio

As shown in [Figure 3-17](#), the HPE Persistent Memory portfolio includes HPE NVDIMMs and HPE Scalable Persistent Memory.

What is an NVDIMM?

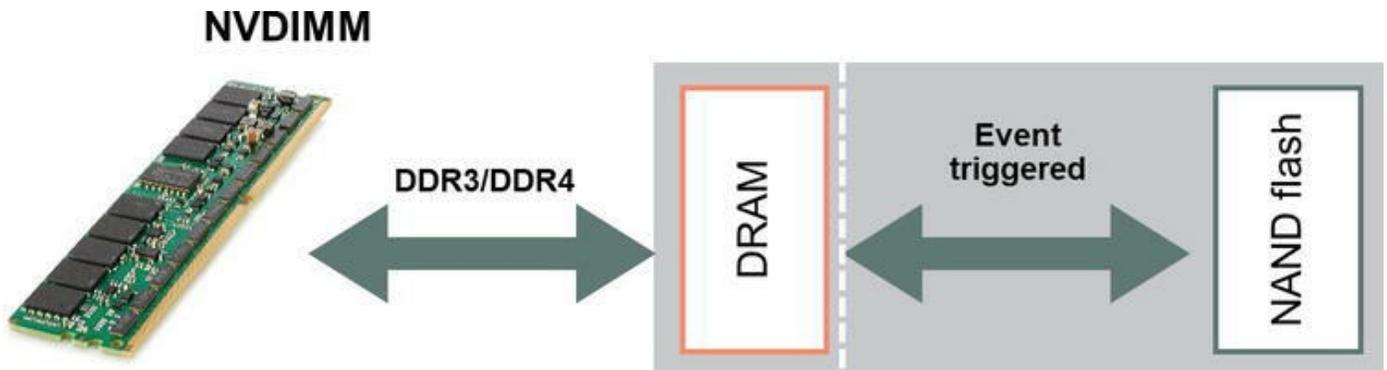


Figure 3-18 NVDIMM architecture

As shown in [Figure 3-18](#), HPE NVDIMM technology combines standard DRAM-based system memory with NAND Flash-based persistency components to achieve both high performance and storage persistency.

In normal operation, the NVDIMM uses high-speed DRAM for servicing all data accesses, and operates like a standard system memory DIMM. A built-in Persistency Controller constantly monitors the system for notification of events that threaten the data in DRAM, such as:

Sudden power loss

Shutdown/restart initiated from the operating system

Catastrophic system errors

Operating system crashes

Notification of any of these events results in the Persistency Controller initiating a backup operation, in which the Persistency Controller transfers all the contents of the volatile DRAM to the onboard NAND Flash device, which stores the data until the event has been resolved.

HPE Scalable Persistent Memory

- Provides OS with Persistent Storage with speed of DRAM
- Uses standard hardware components
 - DDR4 DIMMs
 - NVMe SSD Drives
- Provides greater flexibility than standard NVDIMM-N allowing for larger persistent memory regions. UPS Power Supply provides power for backup on AC power loss

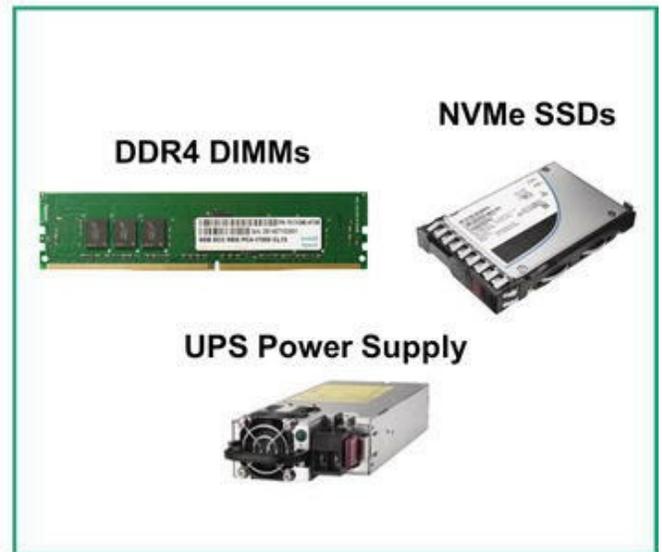


Figure 3-19 HPE Scalable Persistent Memory

An integrated storage solution, HPE Scalable Persistent Memory runs at memory speeds with terabyte capacity using resilient technology. As shown in [Figure 3-19](#), it includes a DRAM layer (DDR4 DIMMs) to accelerate applications, a dedicated flash tier (NVMe SSDs) for persistency, and backup power (UPS Power Supply) to facilitate moving data from DRAM to Flash. HPE Scalable Persistent Memory enables large in-memory compute with persistence.

HPE Scalable Persistent Memory is ideal for:

Large in-memory compute

Checkpoints and restores

Hybrid Transactional/Analytical Processing (HTAP) Real-Time Analytics

Large databases

VSAN and Storage Spaces Direct

Big Data, service providers, performance tier, and virtualizations

HPE Scalable Persistent Memory uses standard components DDR4 DIMMs, NVMe SSD drives, and UPS Power Supply. The Unified Extensible Firmware Interface (UEFI) BIOS settings are configured to identify the Persistent Memory regions.

HPE Scalable Persistent Memory use cases

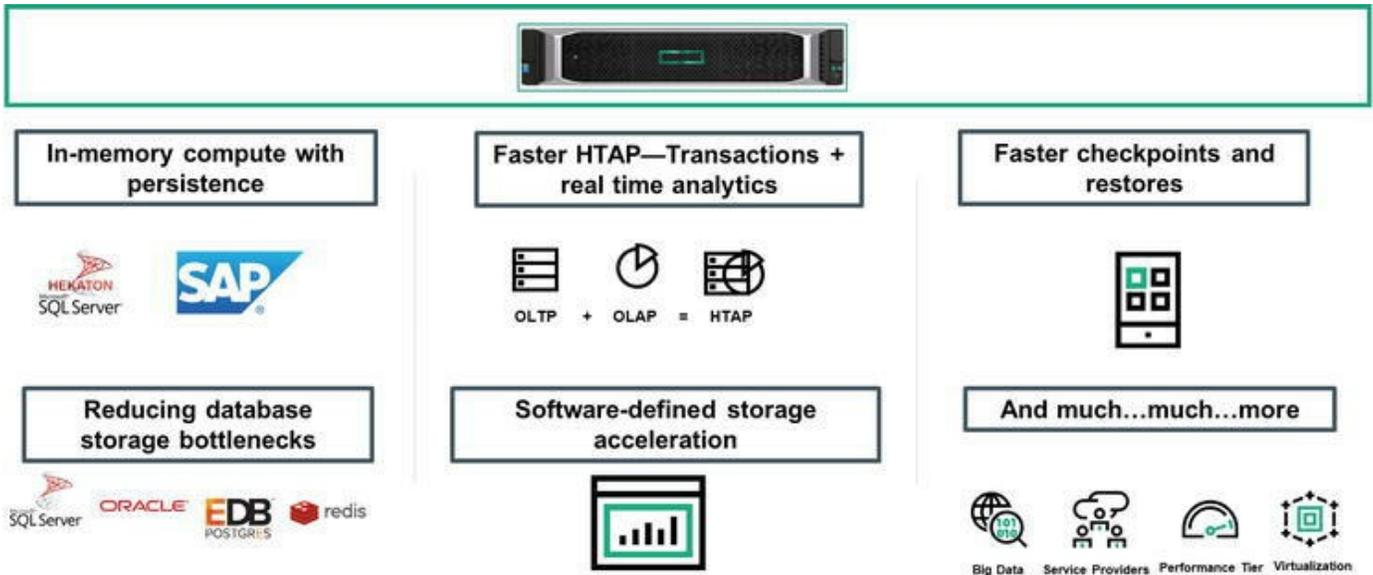


Figure 3-20 Performance of memory—Persistence of storage

Anywhere you have storage bottlenecks today, where you are writing to block storage devices and using a layer of Persistent Memory in place you can see huge benefits. [Figure 3-20](#) shows some of the use cases for HPE Scalable Persistent Memory.

Use cases for HPE Scalable Persistent Memory:

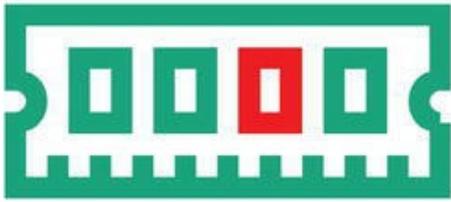
Reducing database storage bottlenecks. Database performance can be vastly improved by using low-latency DRAM to reduce or eliminate storage bottlenecks in lieu of using higher-latency block storage devices.

Many databases have separate operations for the transactional portion of the database and the analytics portions which uses data tables from the transactional portion.

Customers can utilize TB-scale HPE Scalable Persistent Memory to enable real-time analytics—known as HTAP—where they do the transactions and analytics simultaneously.

Software-defined storage (SDS) implementations can also benefit from HPE Scalable Persistent Memory.

HPE Fast Fault Tolerance Memory



Designed to help reduce unplanned outages, system downtime, and even disruptive server crashes

- Monitors for DRAM device failures, pinpointing the address affected and reorganizing the way data is stored to correct the DRAM failure.
- Each memory region runs independently until an error is detected. Only the region with ADDDC or repair activated will run at a reduced performance.
- Memory runs at max performance until feature is activated with no memory bandwidth loss until correction.

Figure 3-21 Seamless repair and recovery

With the introduction of HPE Gen10, new HPE SmartMemory has been introduced that provides increased performance supporting speeds of up to 2666 MT/s, matching the performance capabilities of the new generation of Intel® Xeon® Processor Scalable Family. Also, a RAS feature called Fast Fault Tolerance is available that enables the system to boot with full memory performance while monitoring for DRAM device failures.

Fast Fault Tolerance is a new feature in Gen10 server memory that enables the system to boot with full memory performance while monitoring for DRAM device failures. In the event of a memory failure, the memory subsystem automatically reorganizes the way data is stored in memory to create a protected region just large enough to correct the DRAM failure, while allowing the remaining portions of memory to continue to run at full performance, as shown in [Figure 3-21](#).

Features:

- Reduces unplanned outages and proactive DIMM swaps that prevent unplanned outages
- Algorithm allows each memory region to run independently (high performance) until an error is detected
- Projected to deliver significantly greater memory bandwidth than DDC for random access memory patterns

Improvement over Gen9 solutions— The prior technology (DDDC) is not used by most customers because of a severe performance penalty (all ranks must be configured so they are permanently in Lockstep which incurs a potential memory bandwidth loss of

50%).

How it works:

- DIMMs run fully independent until one DRAM starts to have errors.
- Pair the section of the DIMM that has errors with another section in the same channel.
- Run only that combined section in DDDC mode.
- Only the combined DDDC mode region runs at reduced performance. All other regions still run independently.

Performance:

- All memory runs at max performance until Fast Fault Tolerance is invoked. No memory bandwidth loss until correction is invoked.
- Only the region with ADDDC invoked will run at reduced performance.

HPE Smart Storage Battery



Figure 3-22 HPE Smart Storage Battery

The HPE Smart Storage Battery, examples of which are shown in [Figure 3-22](#), is required to provide backup power to the components that are required to retain volatile information. The HPE Smart Array controllers, HPE NVDIMM, and Scalable Persistent Memory are among the devices that continue to receive power after an interruption in AC power. The HPE Smart Storage Battery capabilities include:

Provide power for HPE Smart Array controllers and HPE 8 GB/16 GB NVDIMMS

Long-life expectancy (seven years) Lithium-ion batteries offering

- 96-watt battery in rack and towers servers

- 12-watt battery in BladeSystem and Synergy servers

Integrated design eliminates thermal and connections issues with alternative backup power sources

- Prewired to memory slots to support NVDIMMs with no extra cabling required

HPE Smart Array controllers and storage

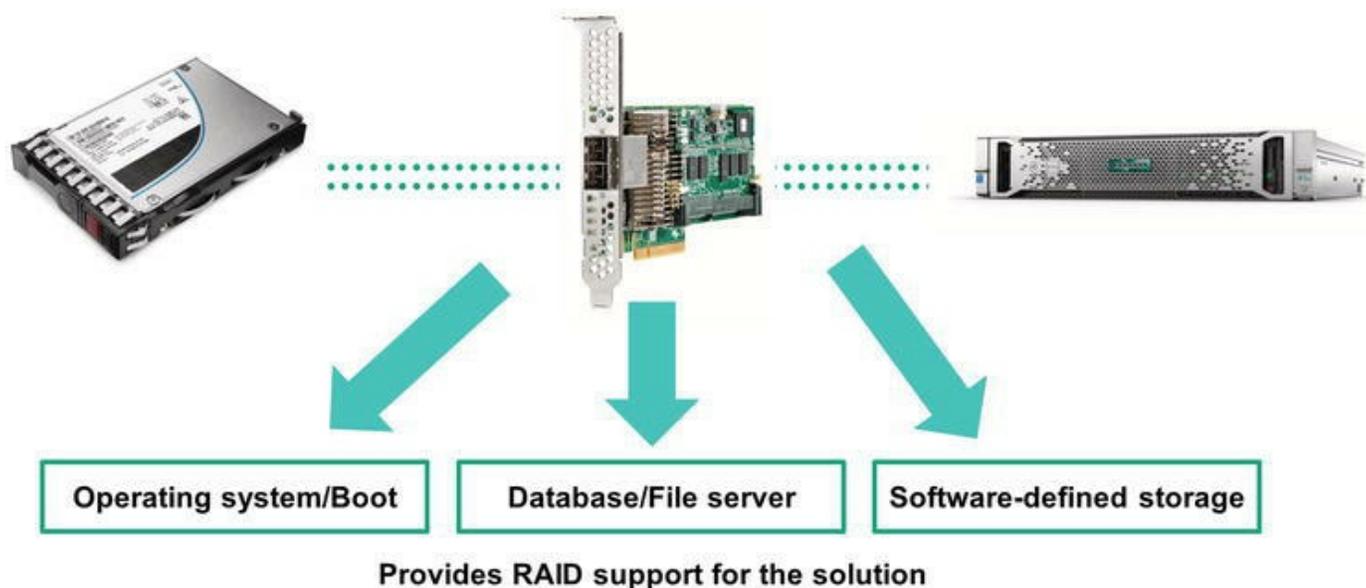


Figure 3-23 HPE Smart Array controllers

The HPE Smart Array controller card is a device that sits between the host system (server) and the storage system (SSD or HDD) and allows them to speak to each other, as shown in [Figure 3-23](#). There are two types: Smart HBA and Smart Array controller.

An HBA is an expansion card that plugs into a slot (such as PCI-e) on the computer system's motherboard and provides fast, reliable non-RAID IO between the host and the storage devices.

HBAs can reliably connect hundreds or even thousands of HDDs, tape, and SSDs to the host.

A RAID controller card is similar to an HBA, but can also add redundancy, optimize performance, reduce latency, or even make smart decisions on whether to store data on an HDD or an SSD cache, depending on user needs. Since these additional tasks consume power and processing speed, RAID controllers are typically more expensive than HBAs and handle fewer devices.

RAID levels

When the controller is operating in RAID mode data is distributed across the drives in one of several ways, referred to as RAID levels, depending on the required level of redundancy and performance. The different schemes, or data distribution layouts, are

named by the word RAID followed by a number, for example, RAID 0 or RAID 1. Each schema, or RAID level, provides a different balance among the key goals: reliability, availability, performance, and capacity. RAID levels greater than RAID 0 provide protection against unrecoverable sector read errors, as well as against failures of whole physical drives.

RAID is a data storage virtualization technology that combines multiple physical disk drive components into a single logical unit for the purposes of data redundancy, performance improvement, or both. This can be achieved under hardware or alternatively software control.

Data are distributed across the drives in one of several ways, referred to as RAID levels, depending on the required level of redundancy and performance.

Different RAID levels are selected to provide a different balance among the key goals of reliability, availability, performance, and capacity. Typically, RAID 0, 1, 10, 5, 6, 50, and 60 are used.

RAID level 0, for example, provides the greatest capacity and read/write performance at the expense of having no redundancy. Therefore, just one disk failure will result in a catastrophic loss of all data.

RAID level 1, data are written identically to two or more drives, thereby producing a “mirrored set” of drives, read operations may reference any drive. The total capacity is equal to a single disk. The array continues to operate as long as at least one drive remains functional.

Table 3-2 RAID levels

RAID level	Function/Application	Limitations
0 Minimum 2 disks	Data are distributed, striped, across separate disk drives.	Highly vulnerable to failure. The entire array will fail if one drive fails.
1 Minimum 2 disks	Mirroring—Identical data written to two or more drives, high-fault tolerance, good performance (higher read performance than RAID 0). Good for use as a boot device.	50% of capacity dedicated to fault protection for a two-drive mirror.
1+0 Minimum 2 disks	Implemented as striped, mirrored disks. Good for database applications requiring high performance and	50% of capacity dedicated to fault protection. Doubles the number of drives required.

fault tolerance.

5 Minimum 3 disks	One set of parity data is striped across all drives. Protects against the failure of any one drive in the array. Good for transaction processing, file and application servers, and virtualization server applications.	Potentially risky for large arrays. Can only withstand the loss of one drive without total array failure. Low-write performance (improved with the presence of controller-based cache memory and flash-based battery backup).
6 Minimum 4 disks	Two sets of parity data are distributed across all drives. Protects against the failure of two drives in an array. Provides higher fault tolerance than RAID 5. Good for 24 × 7 operations that require higher levels of tolerance than RAID 5.	Lower write performance than other RAID levels. Sequential and burst-write performance can be improved with write-back cache flash-based battery backup. Can tolerate up to two drive failures.
50 Minimum 2 groups of 3 disks	Two or more groups of RAID 5 disk arrays with distributed striping across the groups. Combines the straight block-level striping of RAID 0 with the distributed parity of RAID 5. Recommended for applications that require high-fault tolerance, capacity, and random access performance.	A maximum of one drive from each of the RAID 5 groups could fail without loss of data. As the number of drives in a RAID set increases and the capacity of the drives increase, this impacts the fault-recovery time correspondingly as the interval for rebuilding the RAID set increases.
60 Minimum of 2 groups of 4 disks	Two or more groups of RAID 6 disk arrays with distributed striping across the groups. Combines the straight block-level striping of RAID 0 with the distributed double parity of RAID 6, resulting in a RAID 0 array striped across RAID 6 elements. It	A maximum of two drives from each of the RAID 6 groups could fail without loss of data. Like RAID 5, as the number of drives in a RAID set increases and the capacity of the drives

requires at least eight disks.

increase, this impacts the fault-recovery time correspondingly as the interval for rebuilding the RAID set increases.

JBOD

Just a Bunch Of Disks. Presented by a controller in HBA mode. Disks are individually software controlled, or combined, to provide a RAID-like operation by the operating system. Cheaper than hardware RAID, can span multiple clustered servers for better fault tolerance. Sometimes referred to SDS.

Performance is limited by the speed of the processor and the IO subsystem. Redundancy is handled by the software.

[Table 3-2](#) lists the functions, applications, and limitations of different RAID levels.

Gen10 Smart Array portfolio

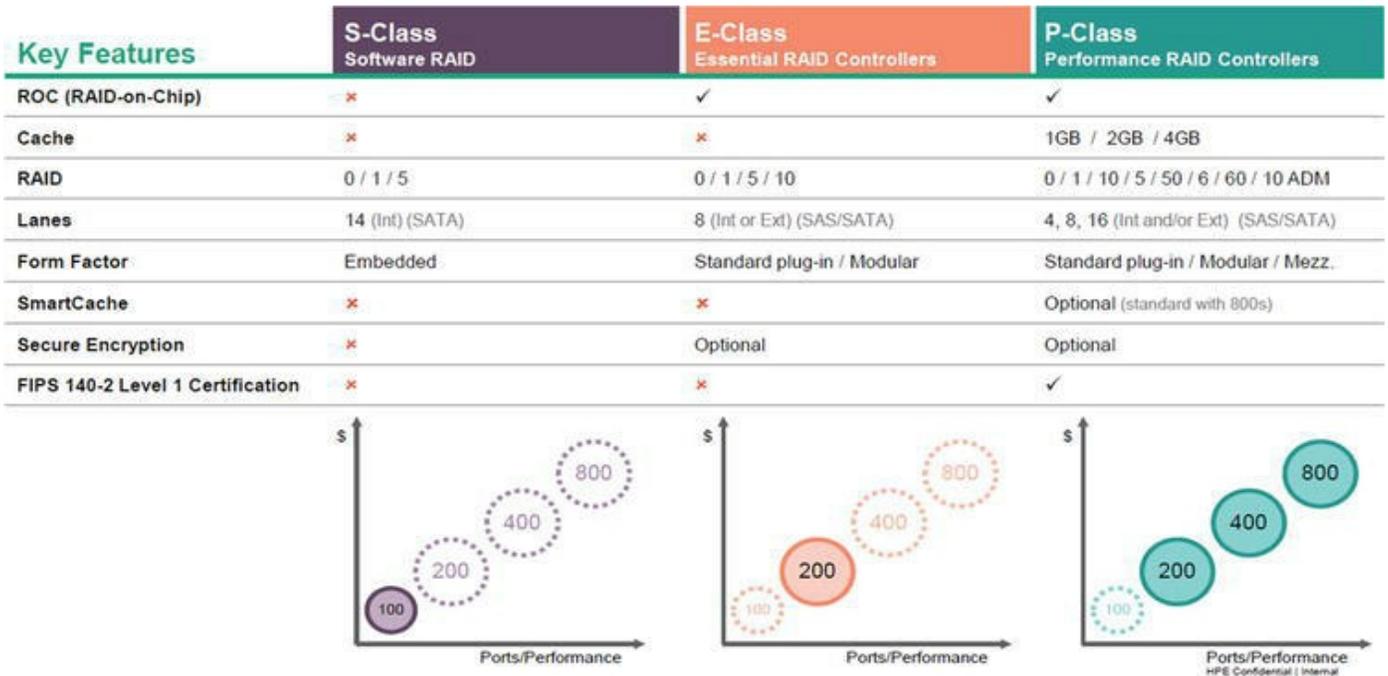


Figure 3-24 Gen10 Smart Array portfolio comparison

HPE Smart Array controllers are ideal for maximizing performance, capacity, and resilience for DAS connected to entry-level and enterprise servers. [Figure 3-24](#) compares the features of the Smart Array portfolio.

Features include:

Mixed mode—Free up a PCIe slot with Mixed mode for Smart Array controllers, offering flexibility to use both HBA and RAID modes simultaneously.

Increased productivity with New UEFI Configuration Tool that reduces time to configure simple RAID volumes.

Save time for newly created RAID 5 or RAID 6 volumes that require parity initialization with Rapid Parity Initialization (RPI).

Enhanced protection for data at rest on all SAS/SATA drives and data security to comply with regulations for sensitive data using New HPE Smart Array SR Secure Encryption.

Better performance—HPE Gen10 controllers deliver up to 1.6M IOPS (4 KB random reads), 65% more performance over the previous generations of controllers.

Less power—HPE Gen10 controller uses up to 46% less power than the previous generation, resulting in power and cooling savings.

Caching solution—HPE Smart Array SR SmartCache accelerates access to data on HDD by up to 4× by caching the hot data on SSDs.

Gen10 Smart Array naming framework

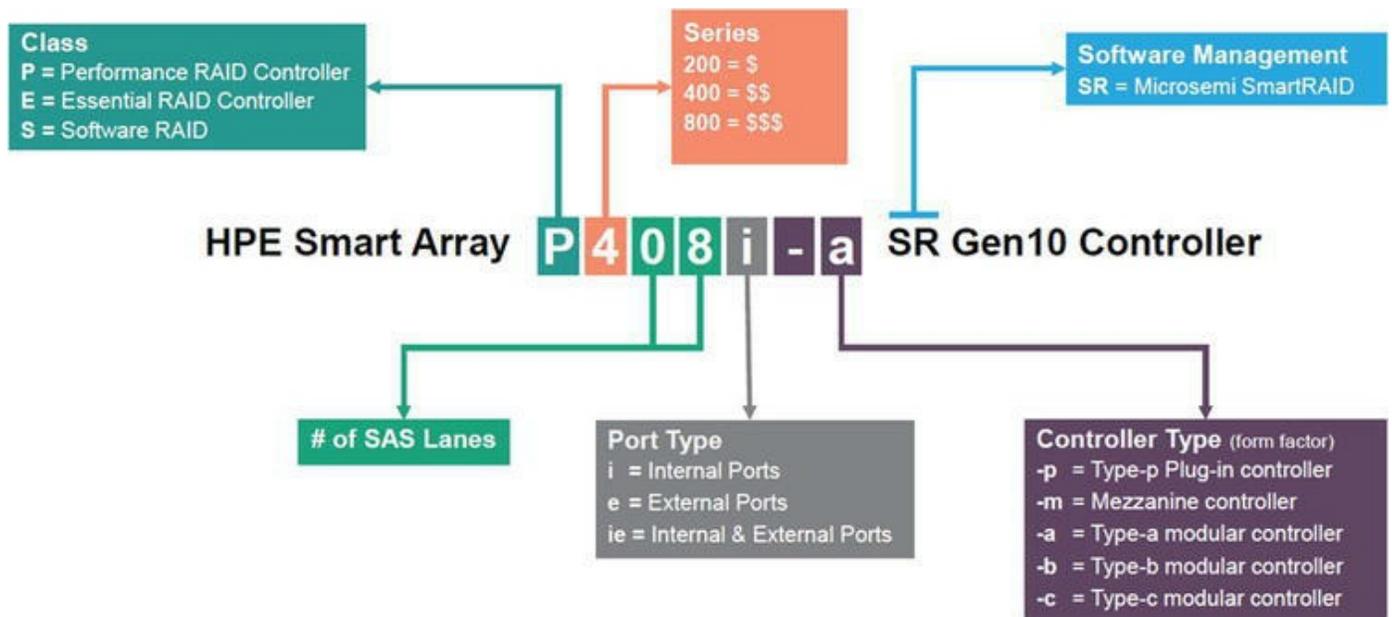


Figure 3-25 Gen10 Smart Array naming framework

[Figure 3-25](#) explains the naming framework of HPE Smart Array products.

ProLiant Gen10 SR Smart Array controllers



Figure 3-26 ProLiant Gen10 SR Smart Array controllers

[Figure 3-26](#) shows several ProLiant Gen10 SR Smart Array controllers. Features

include:

12 Gb/s SAS PCIe 3.0 standard card

Supports internal DAS in ProLiant BL460c Gen10 Server blades

Enhances reliability and storage performance

Optimizes HPE Secure Encryption

Advanced RAID levels with 1 to 4 GB FBWC

Mixed mode operation combines HBA and RAID mode simultaneously



Note

For information on the complete range of HPE Smart Array controllers and Smart HBAs, enter the following URL into your browser:

<https://www.hpe.com/uk/en/product-catalog/servers/smart-array-controllers-and-smart-host-bus-adapters.hits-12.html>

Activity: Features of HPE Smart Array controller

This activity requires you to do an internet search to locate the QuickSpecs for the P816i-a Smart Array controller then answer the following questions:

What is the number of internal SAS lanes and the maximum Flash Back Write Cache memory capacity?

What is the maximum number of logical drives?

What is the form factor of this controller?

What is required to support the Flash Back Write Cache?

Describe the new feature known as “Mixed mode” controller operation?

HPE Server Storage portfolio



Figure 3-27 HPE HDDs

HPE HDDs, shown in [Figure 3-27](#), offer lower \$/GB than flash and provide:

Increased SFF capacity: 15,000 rpm up to 900 GB; 10,000 rpm up to 2.4 TB

Rigorous testing/qualification program, backed by 3.35 million test hours

Features include:

All HDDs include Digitally Signed Firmware and Best in Class (BIC) firmware security features.

Support for Non-Disruptive Updates (NDU) which means that firmware updates do not interrupt access to data or system service.

Prevents unexpected data loss with “Do Not Remove” button on HPE Smart Carriers.

Supported on: HPE ProLiant rack and tower servers, HPE BladeSystem, HPE Apollo, HPE Integrity, and HPE Synergy Compute Modules.

HPE HDDs deliver proven performance at industry-leading capacities backed by high quality and reliability and are the preferred choice for many workloads ranging from small businesses to large enterprise.

HPE HDDs are optimized to deliver maximum performance on HPE server platforms and are proven for applications that demand reliable performance in 24 × 7 operational environments. HPE HDDs offer customers great value per gigabyte without sacrificing performance. In addition, an industry-leading testing and qualification program ensures that our drives are rugged enough for any need.

Table 3-3 HDD classes, workloads, and usage specifications.

Class	Enterprise		Midline	Entry*
Workload	High		Medium	Low
Usage	Mission-critical, high I/O: Email, enterprise resource planning (ERP), Customer relationship management (CRM)		High Capacity, high availability storage: Backup, archive, reference	Low I/O: Boot and backup
Interface	SAS 12 Gb		SAS 12 Gb SATA 6Gb	SATA 6Gb
RPM	15000	10000	7200	7200/5900
Form factor	SFF & LFFc	SFF	SFF & LFF	LFF
Capacities	Up to 900GB	Up to 1.8TB	Up to 10TB	Up to 4TB
Warranty	3 years	3 years	1 year	1 year

As presented in [Table 3-3](#), HPE HDDs are available for three types of workloads: Entry, Midline, and Enterprise. HDDs also provide two interfaces—SAS and SATA—and in a variety of capacities (up to 10 TB and growing).



Note

*Supported on select HPE ProLiant 10 series servers only.

HPE drives are tuned to deliver optimum performance for HPE servers. They undergo an intense qualification process to eliminate firmware and operating system incompatibilities. HPE drive firmware is designed to maximize both functionality and compatibility with HPE ProLiant rack and tower servers, Apollo family, Blade, and Synergy platforms.

Prefailure alert

When used in conjunction with a Smart Array controller, the Self-Monitoring, Analysis, and Reporting Technology (SMART) capable firmware in HPE hard drives enables extensive fault prediction capabilities. If potential problems develop in one of the drives, the Smart Array controller and/or Smart HDD lets you know in advance, so you can have the drive replaced before it fails, under warranty.

Warranty

Enterprise hard drives have a three-year limited warranty regardless of the warranty period for the system in which they are installed.

Midline hard drives have a one-year limited warranty regardless of the warranty period for the system in which they are installed.

HPE solid-state drives

A solid-state drive (SSD) is a non-volatile storage that stores persistent data on solid-state flash memory. SSDs actually are not hard drives in the traditional sense of the term, as there are no moving parts involved.

Similar to a memory stick, there are no moving parts in an SSD. Rather, information is stored in microchips. Conversely, a HDD uses a mechanical arm with a read/write head to move around and read information from the right location on a storage platter. This difference is what makes SSD so much quicker.

Features include:

Flash memory

Low latency 2.5” NVMe SSDs enables fast data access

Reduced latency and higher performance per server with up to 474,000 IOPS for core enterprise workloads

HPE SSD portfolio

Table 3-4 HPE SSD portfolio

	Write intensive	Mixed use	Read intensive
Interface	SAS, SATA, PCIe/ NVMe, M.2	SAS, SATA, PCIe/ NVMe, M.2	SAS, SATA, PCIe/ NVMe, M.2
Endurance	>=10 DWPD	>1 and <10 DWPD	<= DWPD
Typical workload	High read/write applications	Mixed read/write applications	High read/low write applications

HPE SSDs are suited to enterprise environments with highly random data under a variety of write-workload applications. SSDs provide significantly better random read and write IO operations per second (IOPS) compared to HDDs. While sequential read and write throughput is also improved over HDDs, the greatest benefit is recognized in random-data applications. As a result, these high-performance, low-latency, and low-power SSDs provide significant system benefits for applications that previously over-provisioned HDD capacity to achieve better performance.

As presented in [Table 3-4](#), HPE SSDs are available in three broad categories based on their typical target workloads: Read Intensive, Mixed Use, and Write Intensive.

Read-Intensive SSDs are typically the lowest price, with a typical endurance of ≤ 1 drive writes per day (DWPD).

Write-Intensive SSDs typically have the highest write performance, with a typical endurance of ≥ 10 DWPD.

Mixed Use SSDs are for workloads that need a balance of strong read and write performance, with endurance typically >1 and <10 DWPD.

The workloads indicate the number of DWPD that you can expect from the drive. (DWPD is the maximum number of 4K host writes to the entire drive capacity of the SSD per day over a five-year period.)

HPE Write-Intensive SSDs

HPE Write-Intensive 12Gb SAS and 6Gb SATA SSDs provide high write performance and endurance. They are best suited for mission-critical enterprise environments with workloads high in read and writes. Write-Intensive workloads include OLTP, VDI, business intelligence, and Big Data analytics.

HPE Mixed Use SSDs

HPE Mixed Use 12Gb SAS and 6Gb SATA SSDs are best suited for high IO applications with workloads requiring balanced performance between reads and writes.

When paired with HPE ProLiant servers, these SSDs help you meet the challenges of Big Data.

HPE Read-Intensive SSDs

HPE Read-Intensive 12Gb SAS and 6Gb SATA SSDs deliver enterprise features at an affordable price in HPE servers. Entry-level pricing is fueling the rapid adoption of SSDs for read-intensive workloads because the cost per IOPS compares favorably to HDDs. Read-Intensive SSDs deliver great performance for workloads high in reads such as boot/swap, web servers, and read caching, just to name a few.

HPE Read-Intensive M.2 Solid State Enablement Kits

The HPE M.2 Solid State Enablement Kit is a recent addition to the Read-Intensive SSD family and is best suited for operating system boot/swap. The M.2 Solid State Enablement Kit is available in dual and single 64 GB and 120 GB capacities. The kits are compatible with ProLiant Gen9 and Gen10 Blades and currently support a 6 Gb SATA interface. Also, the dual and single 120 GB and 340 GB capacities are available in the M.2 Enablement Kit and currently supports ProLiant ML/DL servers.

Boost performance with HPE NVMe PCIe 2.5” SSDs

With the introduction of NVMe, an industry interface specification for accessing solid-state storage through PCI Express, manufacturers have a set of guidelines that seek to release them from the limitations of previous standards, and also provide a wide range of interoperability benefits.

HPE NVMe PCIe 2.5” SSDs talk directly to your applications via the PCIe bus, boosting IO bandwidth and reducing latency to scale performance in-line with your processing requirements. This means, for example, that you can host your entire database on one or more HPE NVMe PCIe 2.5” SSDs for enhanced in-memory access and performance. NVMe is an industry specification that focuses on efficiency, scalability, and performance.

HPE PCIe Workload Accelerators for ProLiant servers



Figure 3-28 HPE PCIe Workload Accelerators for ProLiant servers

HPE PCIe Workload Accelerators, shown in [Figure 3-28](#), are advanced PCIe card-based storage devices that use solid-state storage technology directly on the PCI bus, ensuring high read and write data rates and accelerated application performance. The associated application performance improvements have a positive impact on business results and the ability to accelerate IO-bound workloads such as databases, virtualization, and data analytics, resulting in significant cost and time savings.

HPE PCIe Workload Accelerators provide consistent, predictable, reliable, sustained high performance for business-critical applications. They talk directly to applications via the PCIe bus, boosting IO and reducing latency to scale performance in-line with processing requirements. This means, for example, that an entire database can be hosted on one or more workload accelerators for enhanced in-memory access and performance.

The accelerators are built on NAND Flash memory technology that uses flash memory cells to store data. A single-level cell (SLC) stores one bit per cell and a multi-level cell (MLC) uses two bits per cell. MLC can store more data within each device but it has lower endurance characteristics than SLC.

Integrated seamlessly with ProLiant ML and DL servers via the PCIe bus, the workload accelerators increase per server performance by eliminating IO bottlenecks. By decreasing the need for external HDD storage arrays together with their software layers, the data center footprint is reduced along with the associated complexity resulting in significant management, maintenance, and power and cooling savings.

HPE PCIe Workload Accelerators are available in the following product categories:

HPE Read-Intensive PCIe Workload Accelerators

- Workloads requiring best latency and read performance including web servers, cloud environments, and databases
- Best PCIe Workload Accelerator IOPS/\$
- Capacities: 1.2 TB, 1.3 TB, 1.6 TB, 3.2 TB, and 6.4 TB

HPE Mixed Use PCIe Workload Accelerators

- Workloads requiring best latency and mixed read and write performance including web servers, cloud environments, OLTP, and databases
- Capacities: 700 GB, 1.0 TB, 1.3 TB, 1.4 TB, 2.6 TB, and 5.2 TB
- Capacities for NVMe: 800 GB, 1.6 TB, and 2.0 TB

HPE Write-Intensive PCIe Workload Accelerators

- Workloads requiring best latency and write performance including database, business intelligence, OLTP, and VDI
- Capacities: 350 GB and 700 GB
- Capacities for NVMe: 800 GB and 1.6 TB



Note

For compatibility and operating system support information, refer to the PCIe Workload Accelerator Datasheet. To access the datasheet, enter the following URL into your browser:

<https://h20195.www2.hpe.com/v2/GetPDF.aspx/c04123190.pdf>

HPE ProLiant server solution technologies

The design technologies embedded in ProLiant servers allow customers to increase application performance and free up resources to focus on business innovation.

Gen10 power strategy

As the number, complexity, and required performance of applications increase, IT executives are finding that the required space, power, and cooling capacity needed to host additional systems are at a premium. HPE has developed intelligent and efficient server and infrastructure technologies that reduce the power and airflow needed to operate the latest ProLiant servers. These technologies reclaim limited space, power, and cooling resources for needed workloads, while also reducing error-prone manual checking and documenting of power and rack configurations.

HPE iPDUs

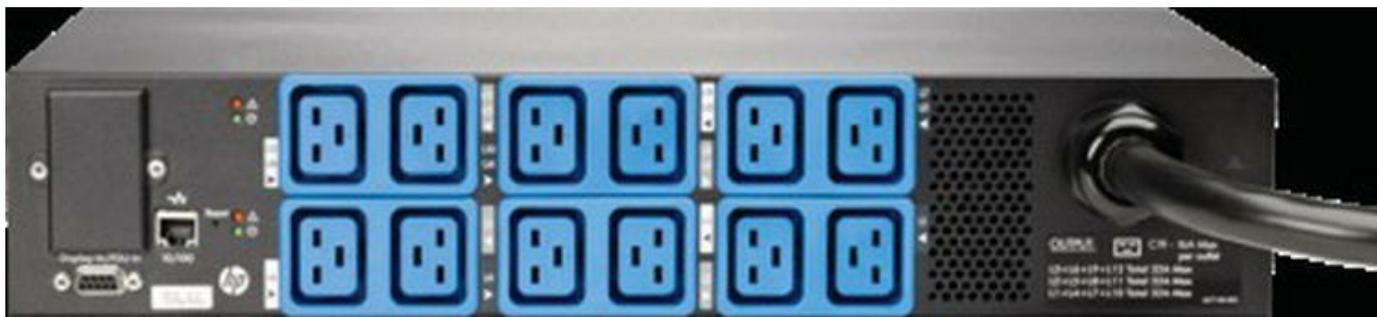


Figure 3-29 HPE iPDU

The HPE iPDU, shown in [Figure 3-29](#), brings state-of-the-art management and control to rack-mounted PDUs to prevent overprovisioning of power from restricting growth in the data center. iPDUs monitor power consumption at the core, load segment, stick, and outlet levels with precision and accuracy in a core and stick architecture. Remote management is built into the PDU core for power cycle ability of individual outlet on the extension bars and high-precision monitoring of current, voltage, wattage, and power that is 99% accurate across the C19 outlets.

The iPDU can help users to track and control power other PDUs cannot monitor. Information is gathered from all monitoring points at half-second intervals to ensure the highest precision. The iPDU can detect a new server even before it is powered on.

An iPDU is ideal for enterprise customers who want to speed power configuration, improve monitoring to reclaim stranded power, and extend data center life. An iPDU enables you to:

Determine the exact power consumption for every component in the rack with less than 1% variation in accuracy.

Monitor and control outlets individually for the ultimate control of power distribution within the rack.

Save valuable space in the rack with dense rack-mountable form factors.

HPE Power Discovery Services with HPE Intelligent Power Discovery combines HPEs, HPE Platinum power supplies, and HPE Insight Control software to control and automate power distribution. When you deploy new ProLiant servers, HPE Intelligent Power Discovery automatically discovers those servers, maps them to the power source, verifies power redundancy, and makes sure everything is hooked up correctly.

HPE iPDU's have a patented modular architecture that improves their flexibility. This building block concept consists of two main parts: the iPDU core unit and extension bars. The iPDU's ship with one core unit, the mountable LED display, and mounting hardware. The six-outlet core unit is zero-U or 1U rack mountable. The 12-outlet core unit is 2U rack-mountable. This architecture:

Communicates with the attached servers to collect asset information for the automatic mapping of the power topology inside a rack

Speeds implementation time

Greatly reduces the risk of human errors that can cause power outages

The iPDU core units and extension bars include the following components:

iPDU core units—Each iPDU core unit incorporates an embedded web engine for remote monitoring and control and includes the LED display unit for local monitoring of current at the rack. Core units can be mounted in HPE Advanced Series Racks or in 1U configurations in any industry-standard rack. Core units are available in a variety of single-phase or three-phase inputs to match any application.

Each core unit has individually monitored International Electrotechnical Commission (IEC) C19 outlets that are individually protected by high-quality commercial grade breakers. Each IEC C19 on the iPDU incorporates out-of-band power line communications (PLC) technology to communicate with either the standard extension bars or the Intelligent extension bars.

PDU extension bars—iPDU extension bars can be plugged into the core unit and mounted directly to the rack frame, spanning the entire length of the rack. All iPDU

extension bars are sold in pairs. There are three types of extension bars:

- PDU Standard Extension Bars—These extension bars have five IEC C13 outlets for connection to external equipment. They can be mounted in the rear of the rack to ease cable management. Individual extension bars do not extend into the maintenance zone between RETMA rails at the rear of the rack. Standard extension bars have power indicators, unit identifier (UID) LEDs, and bright-blue PLC IEC C20 input connectors for use with HPE iPDU core units.
- iPDU Extension Bars G2—The iPDU Extension Bar G2 has an automatic outlet numbering system that maintains consistent outlet identification so that Outlet 1 will always be the top of the Extension Bar, regardless of where it is mounted.

Each Intelligent extension bar is 5U in height and has five individually monitored and switchable IEC C13 outlets. Each IEC C13 outlet:

- ◆ Has a status indicator and UID LED that are activated when scrolling through the local display unit or scrolling remotely using the embedded web engine
- ◆ Incorporates bright-blue PLC connections to communicate with HPE Common Slot Platinum Power Supplies with PLC
- ◆ Is capable of remote power cycling or remote lockout to prevent unauthorized equipment from being installed in the rack
- ◆ Has an automatic 100-millisecond start delay and can be user programmed for restart delays of up to 999 seconds on a per-outlet basis

The inlet cable on the Intelligent extension bar is on the back of the bar to ease installation and improve cable routing to the iPDU core. Intelligent extension bars also have bright-blue PLC IEC C20 input connectors for use with HPE iPDU core units.

Similar to standard extension bars, the Intelligent extension bars can be mounted in the rear of the rack to ease cable management. Standard IEC C14 jumper cables can be used with the Intelligent extension bars for powering non-IPD enabled devices.

Standard PDU Extension Bars (non-intelligent)—Standard PDU Extension Bars can be mounted side-by-side for high-density or redundant applications. They are monitored as a single load segment of five C13 outlets at the IEC C19 outlet on the iPDU core unit. Similar to Intelligent extension bars, the standard extension bars have a power indicator and UID LED that are activated when scrolling through the local display unit or scrolling remotely using the embedded web engine.

HPE Flex Slot power supplies



Figure 3-30 HPE 1800W Flex Slot Platinum Power Supply

HPE Flex Slot power supplies, like the one shown in [Figure 3-30](#), share a common electrical and physical design that allows for hot-plug, tool-less installation into ProLiant Gen10 performance servers. Flex Slot power supplies are rated for Platinum-level certification with efficiency of up to 94% and Titanium-level certification with efficiency of up to 96%. Support for HPE Power Discovery Services, through embedded PLC technology, is available with the 800W Titanium and 1400W/2200W Platinum Plus models. This feature enables each server to communicate identification, location, and power-related data to the iPDU in the rack; this information can then be shared with HPE Insight Control to manage power usage and efficiency in the data center.

Additional benefits of Flex Slot power supplies include:

Titanium-certified power efficiency

- Titanium (96%) power efficiency certification from the 80 Plus program—one of the highest power efficiency certifications in the IT industry
- Decreased data center operating costs by reducing server power requirements and power waste

Flex Slot design

- Tool-less hot-plug design that improves serviceability by allowing quick and easy access to system power supplies
- Common form factor across all ProLiant Gen10 Performance series servers that allows multiple server platforms to share power supply spares, reducing cost and space requirements for spares

Multiple output options

- Certified for high-efficiency operation
- Allows users to right-size a power supply for specific server configurations that helps to reduce power waste, lower overall energy costs, and avoid "trapped" power capacity in the data center
- Support for both low-line and high-line AC input voltages, providing additional flexibility to operate in multiple environments (500W and 800W Platinum only)

Power management

- Supports multiple operating modes to maximize power efficiency when configuring servers with redundant power supplies
- Integrated support for HPE Power Discovery Services that communicates with the iPDU to monitor and manage power usage (Titanium and Platinum only)

Secure data with power supply/battery backup combo

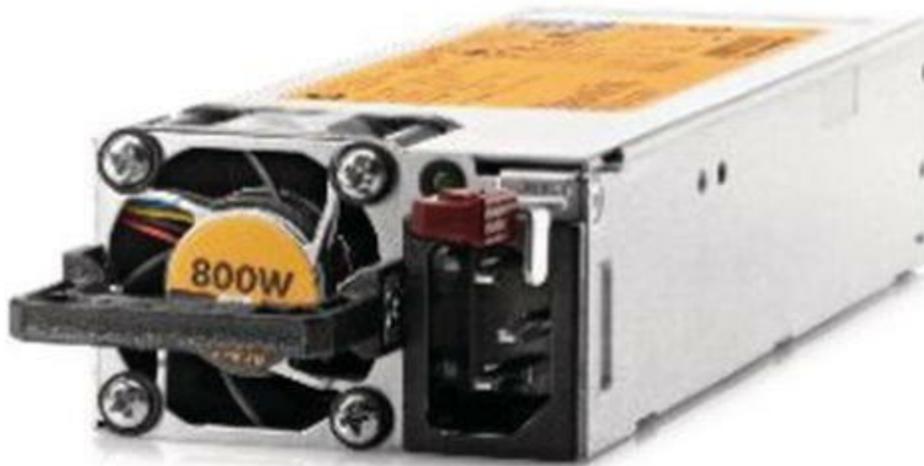


Figure 3-31 HPE 800W Flex Slot Platinum Power Supply

The HPE 800W Flex Slot Platinum Plus Hot Plug Low Halogen Power Supply Kit, shown in [Figure 3-31](#), secures data with 800W Flex Slot power supply with built-in 400W Lithium-ion battery backup for Persistent Memory. Ideal for core enterprise, HPC and Big Data customers looking for higher availability of nodes and data centers seeking new infrastructure designs for greater efficiency.

Learning check

What are functions of the HPE iPDU? Select two.

- A. Determine the exact power consumption for every component in the rack.
- B. Support up to 24 devices through a battery backup unit.
- C. Support for Location Discovery Services.
- D. Use NAND Flash memory technology to consume less memory.
- E. Integrate seamlessly with ProLiant servers via the PCIe bus.

HPE server networking

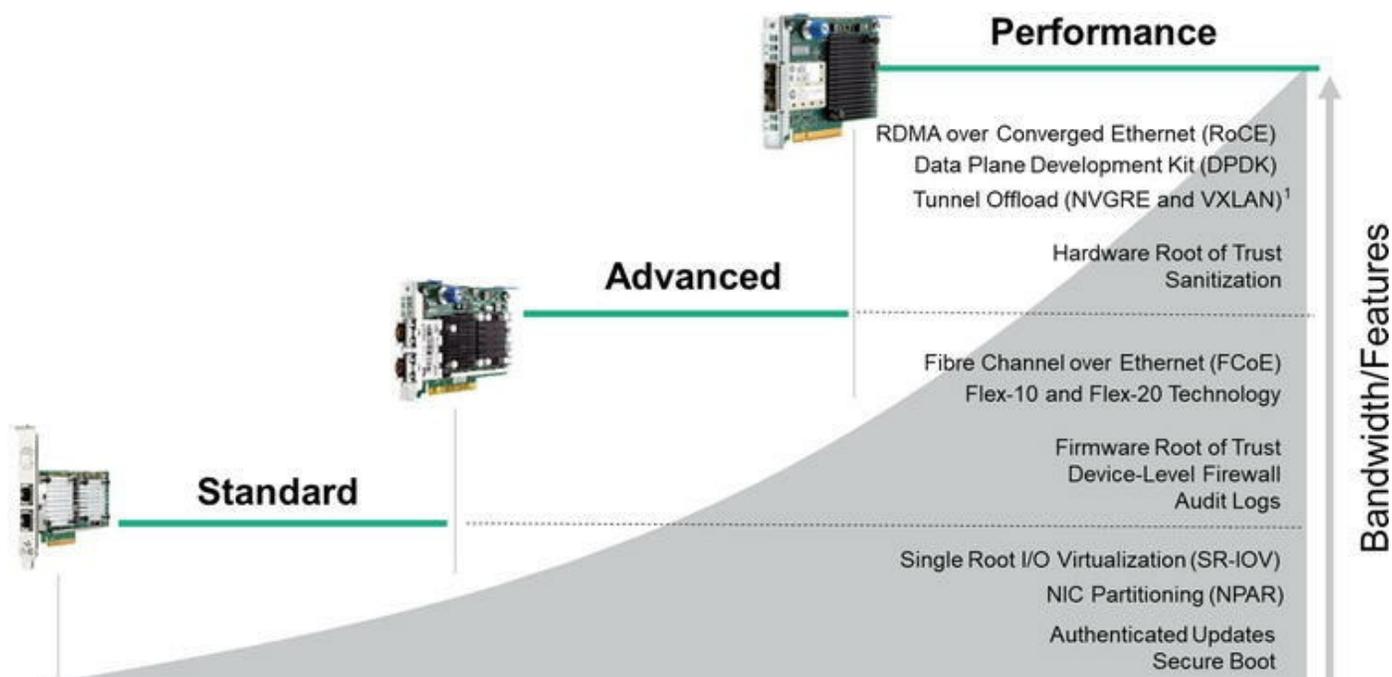


Figure 3-32 Gen10 Network Adapter product portfolio

As shown in [Figure 3-32](#), HPE networking controller models include Standard, Advanced, and Performance series.

The **Standard** series adapters enable a cost-effective Ethernet solution to everyday applications. The economic scalability of these adapters features functionality like Single Root IO virtualization (SR-IOV) for increased performance via direct access to hardware from a virtual environment.

Network Partitioning (NPAR) enables users to “right size” data paths. The Standard series offers 2 and 4 port choices and a broad selection of features at the lowest power and cost per port. The SR-IOV feature enables basic virtualization for expanding network fabric and increasing performance. Other features of the Standard series include:

Increase performance via direct access to hardware from virtual environment with SR-IOV

Prevent malicious attacks with authentication of digitally signed firmware

Ensure system is secure at start-up with Secure Boot

Eliminate host-based software direct memory or storage access with device-level firewall

Forensics recovery with sanitization and audit logs

Broad Support: 1 GbE and 10 GbE on HPE ProLiant DL360, DL380, DL560 Gen10, and select HPE Apollo platforms

The **Advanced** series adapters simplify network and storage topology to build the new hybrid server infrastructure using the CNA FlexFabric technology. The configurable flexibility of these adapters reduces the number of network interface cards required to connect disparate storage and IP networks, reduces the number of cables and switches, and reduces power and cooling costs.

An Advanced series feature is FCoE which reduces the hardware complexity required to connect to disparate storage and IP networks by 50% or more. It also eliminates network sprawl at the server edge and saves up to 47% on upstream ToR (Top of Rack) switch cable connection. Other features of the Advanced series include:

Reduce cabling and maximize bandwidth using NIC Partitioning (NPAR)

Increase performance via direct access to hardware from virtual environment with SR-IOV

Consolidate data traffic using single adapter with FCoE

Prevent malicious attacks with authentication of digitally signed firmware

Eliminate host-based software direct access with device-level firewall

Forensics recovery with sanitization and audit logs

The **Performance** series adapters deliver high bandwidth at low latency. Expressly fast adapters can maximize packet throughput and workload performance with the Data Plane Development Kit support (DPDK).

The 25 GbE adapters are tuned to work with the latest 10/25/50/100 Ethernet standards and with HPE Networking ToR switches. These new 25 GbE adapters maximize performance while auto-negotiating down to 10 Gbps or up to 25 Gbps. In addition, using a 25 GbE enabled server can reduce TCO up to 27% over a 10 GbE server infrastructure while increasing bandwidth by 56%.

Multispeed—10/25/50/100 Gbps: The latest adapter cards are the 25 GbE adapters that enable optimal PCIe slots usage to reduce the total number required to build a 25/50/100 Gbps infrastructure with future support to include speeds beyond 100 Gbps using eight lanes for 400 Gbps.

RDMA over Converged Ethernet (RoCE)—Network-intensive applications like networked storage or cluster computing need a network infrastructure with a high bandwidth and low latency. The advantages of RoCE over other strategies are lower

latency, lower CPU load, and higher bandwidth.

DPDK—The DPDK allows software-based customization and optimization of network performance by using polling instead of traditional interrupt-driven network processing.

Security features—HW Root of trust to create Chain of Trust for DS Firmware Authentication. Sanitization (Secure User Data Erase) renders User and configuration data on the NIC irretrievable so that NICs can be safely repurposed or disposed.

All of HPE Network Adapter Options are based on IEEE networking standard Ethernet protocols to ensure industry-standard compatibility and are easy to configure on the PCIe bus, or as a FlexLOM form factor.



Note

For more information on the HPE networking options portfolio, go to:

<https://www.hpe.com/uk/en/servers/networking.html>

HPE configuration tools

HPE offers a variety of configurator tools to guide you in developing appropriate IT solutions for customers. These tools streamline the ability to select and configure HPE products and to create quotes for you and your customers.

HPE One Config Advanced

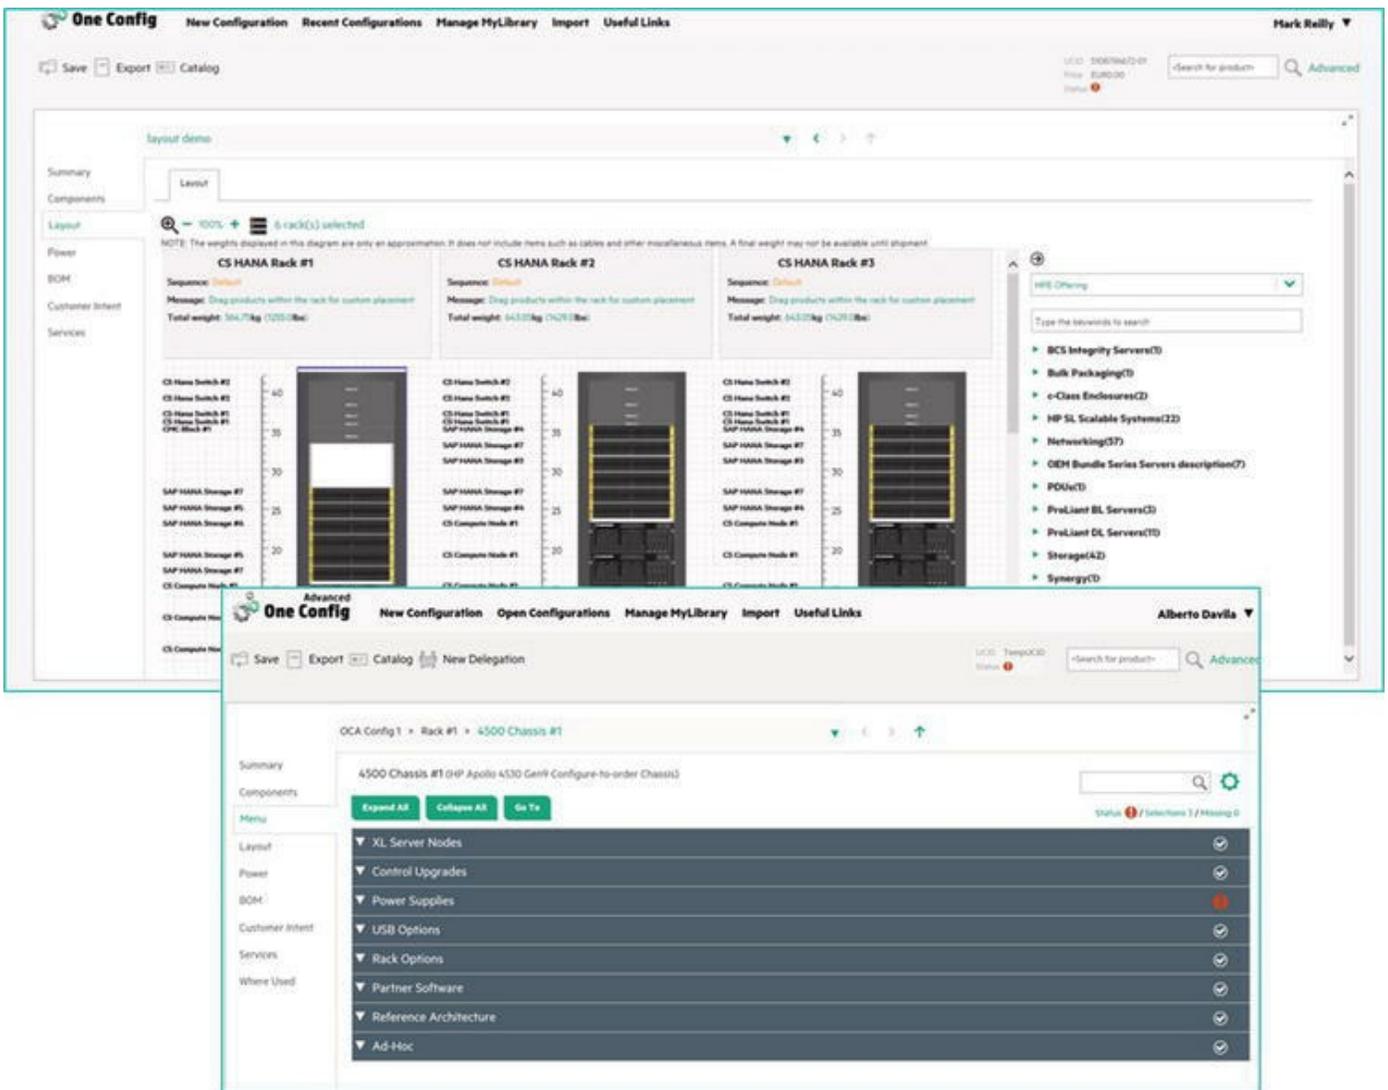


Figure 3-33 HPE One Config Links Advanced

The HPE next-generation configurator which will replace SalesBuilder for Windows (SBW), One Configuration Advanced (OCA), shown in [Figure 3-33](#), allows you to:

Import configurations from other tools, such as SBW

Create new configurations, including complex solutions

Add ad-hoc products

Create custom SKUs

Manage services within the solution

Define customer intent

Set power and network connections

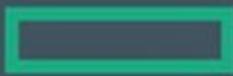
Save and export documents



Note

OCA can be accessed via the HPE Partner Ready Portal at <https://partner.hpe.com>

HPE iQuote Universal



Hewlett Packard
Enterprise

iQuote

Welcome to iQuote Universal

Select country to Login or Register

Please select country

Country not listed?

Request iQuote be made available in your country!

Uganda

Your email address

REQUEST

LOG IN REGISTER

Figure 3-34 HPE iQuote Universal

iQuote Universal, shown in [Figure 3-34](#), is a free service available worldwide for channel partners, customers, and IT consultants. Browse popular HPE server, storage, and networking products and create fast, accurate configurations that you can export to request a quote from your Distributor or Reseller.

iQuote Universal offers all of the same functionality as the subscription service but without host-specific pricing and stock availability information. iQuote Universal is a quoting tool only and does not feature any of the “buy now”/”add-to-basket” features of the hosted version. If it is important for your business to be able to share your own real-time pricing and stock information, then we recommend iQuote Subscription (hosted)

edition.

Features:

The HPE cloud service for channel sales configuration and quoting

Designed to simplify the process of selling HPE products and maximize revenue on every sale

Includes ProLiant servers as well as HPE storage, networking, commercial desktops, laptops, and workstations

Resellers and IT providers can:

- Select a product
- Create a configuration
- Export the information
- Send a validated bill of materials (BOM) to the distributor or supplier



Note

To access iQuote Universal, enter the following URL into your browser:

<https://iquote.hpe.com/asp/signin.aspx>

HPE Power Advisor

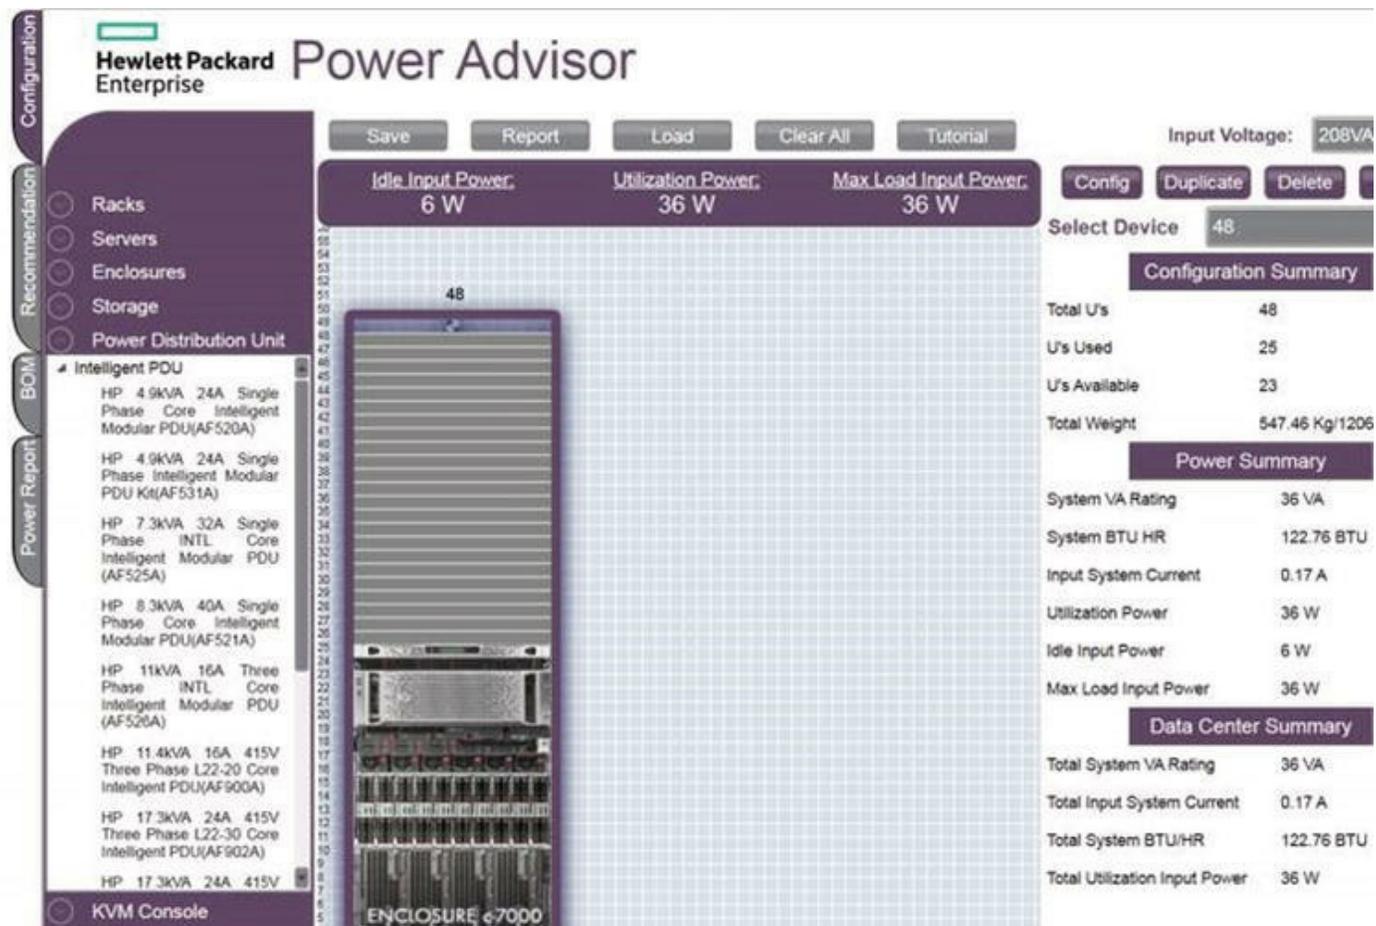


Figure 3-35 HPE Power Advisor

When you are designing and expanding data centers or planning rack deployments, accurate estimates of power requirements are essential to ensure appropriate levels of power and cooling as well as to determine power-related operating costs for the customer's budgetary considerations.

HPE Power Advisor, shown in [Figure 3-35](#), is an easy-to-use tool that estimates data center power requirements for server and storage configurations. Version 9.x includes ProLiant Gen10 servers and options. A downloadable version and online application are available. The Power Advisor online tool supports Google Chrome and Mozilla Firefox. This tool allows you to:

Accurately estimate power consumption of HPE server and storage products.

Select the appropriate power supplies and other system components.

Configure and plan power usage at a system, rack, and multi-rack levels.

Access useful tools including a cost-of-ownership calculator, power report, and BOM.

**Note**

For more information on the Power Advisor enter the following URL into your browser. After the webpage opens, scroll down to HPE Power Advisor.

<https://www.hpe.com/h20195/v2/GetPDF.aspx/4AA6-2925ENW.pdf>

Activity: Features of HPE Power Advisor

The following link can be used to run the Power Advisor online version:

<https://paonline56.itcs.hpe.com/?Page=Index>

This activity requires you to run the HPE Power Advisor tool to build a simple configuration. Please follow the following instructions before answering the activity questions.

Accept the license agreement if necessary, and select the required AC input voltage.

Select the 42U Enterprise G2 Pallet Rack.

From the Server menu select the ProLiant DL380 Gen10, U position 12 and 8 SFF CTO option.

In the GUI, select the DL380 Gen10, and click **Config**.

Work through the configuration menu to add:

- Two CPUs
- Twelve 64 GB memory DIMM
- One storage controller P816i-a
- Flexible LOM 10/25 GB 640FLR expansion
- One 500W Flex Slot Platinum power supplies

Answer the following questions:

Why do you get a warning in RED at the bottom of the screen?

What is required to resolve this warning?

What happens when you click **Save**?

Record the “Total Utilization Power”.

Select the 42U Rack, and record the “Total Weight” from the configuration summary.

If you have time, try adding an ML350 Gen10 and performing similar configuration tasks detailed at the beginning of this activity and answer the same questions.

HPE Solid-State Drive Selector Tool

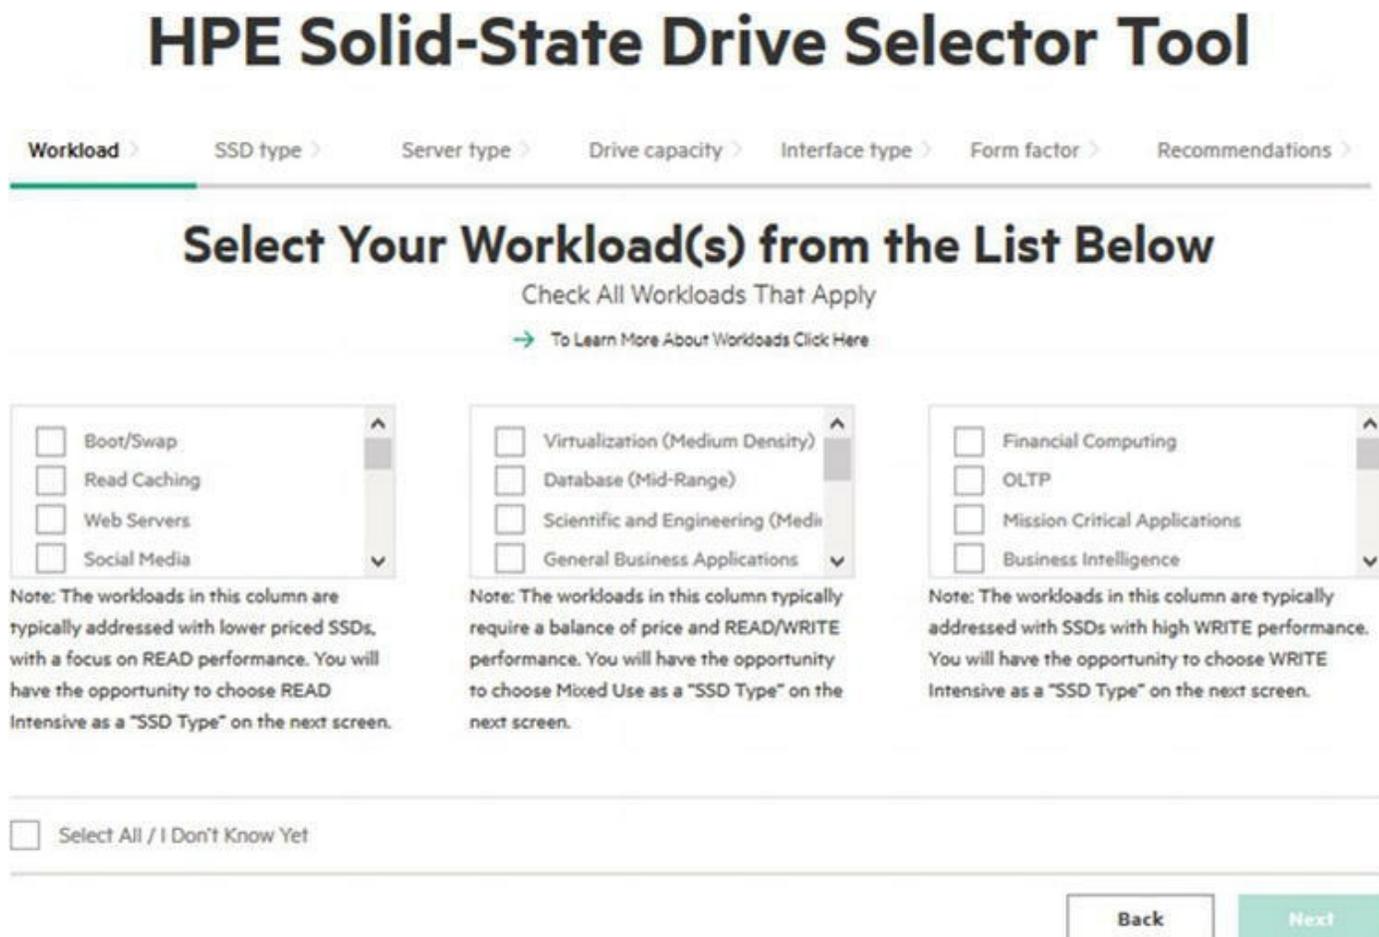


Figure 3-36 HPE Solid-State Drive Selector Tool

The HPE Solid-State Drive Selector Tool, shown in [Figure 3-36](#), is an online tool that helps to guide the process for selecting the right SSD drive for the chosen workload. In addition to the easy-to-follow guided process, the tool includes an alternative, manual driven, expert mode with fine-grained selection process.

Manual selection criteria include:

Server type

Endurance

Max power

Random READ average latency

Random WRITE average latency

SAS, SATA, PCIe/PCIe NVMe



Note

Enter the URL given below into your browser to access the HPE SSD Selector Tool

<https://ssd.hpe.com>

HPE Server Memory Configurator

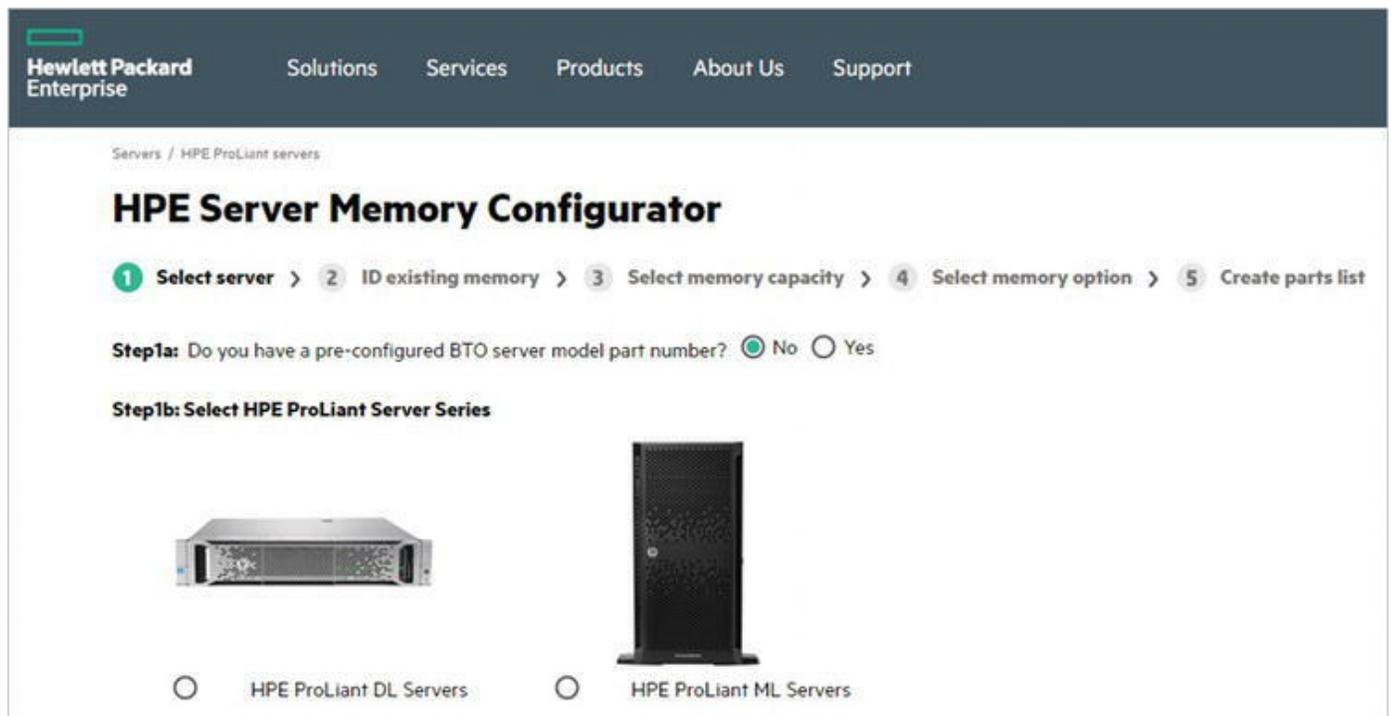


Figure 3-37 HPE Server Memory Configurator

The HPE Server Memory Configurator, shown in [Figure 3-37](#), is a web-based tool used to assist with populating DDR memory in ProLiant servers. This tool provides optimal configurations based on memory population guidelines. Non-optimal configurations are also shown for customers who require a specific memory configuration. Guidance provided is based on customer-provided information and does not guarantee specific performance.

This tool includes a five-step process that provides recommended memory

configurations and RAM module installation locations:

Select the relevant ProLiant server.

Identify existing memory (auto-detect memory with HPE Insight Diagnostics or manually enter current memory).

Select memory capacity.

Select memory option.

Create parts list.



Note

Enter the URL given below into your browser to access the Server Memory Configurator.

<http://h22195.www2.hp.com/DDR4memoryconfig/Home/Legal>

HPE solution sizers

HPE offers several automated tools that assist with recommending a solution environment. The sizing information and algorithms in HPE solution sizers have been developed using testing and performance data on a wide range of HPE servers running solutions from partners such as Citrix, Microsoft, SAP, and VMware. These tools provide a consistent methodology to help determine a “best fit” server for the environment. Sizers are downloaded and run on the user’s personal computer. Updates with the latest information on HPE hardware and solution software are available automatically when the user is connected to the internet and may optionally be installed by the user.

There are several solutions sizers available through the Hewlett Packard Enterprise Information Library, including:

HPE Converged Infrastructure Solution Sizer Suite (CISSS)—Solution sizers from HPE are conveniently available through the CISSS. This suite provides an easy way to install sizers, consolidate the BOM generated by multiple sizings, access reference architectures, and more. You can use the CISS to:

- List the HPE solution sizers and select which ones to install through the Sizer Manager
- Size an application solution using one of the installed solution sizers
- Combine application solutions after two or more solutions have been sized and

saved

- Size an HPE ConvergedSystem solution
- Calculate power requirements for a solution using Power Advisor

HPE Sizer for Server Virtualization—This automated, downloadable tool provides quick and helpful sizing guidance for HPE server and storage configurations running in VMware vSphere 5.0 or Hyper-V R2 environments. The tool allows users to create new solutions; open existing solutions; or use other types of performance data collecting tools, such as the Microsoft Assessment and Planning (MAP) Toolkit, to build virtualized configurations based on HPE server and storage technologies. It enables the user to quickly compare different solution configurations and produces a customizable server and storage solution complete with a detailed BOM that includes part numbers and prices.



Note

To access solution sizers, enter the following URL into your browser:

<https://www.hpe.com/h20195/V2/Getdocument.aspx?docname=4AA6-2732ENW>

HPE Storage Sizer

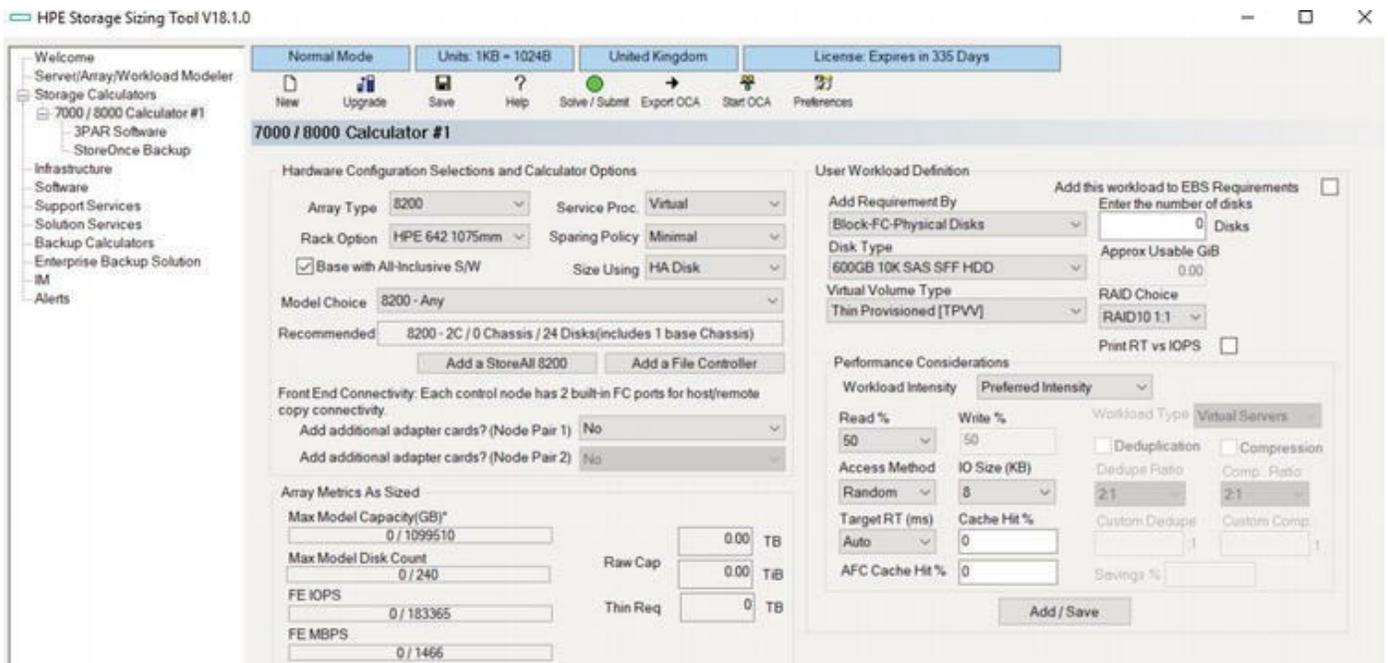


Figure 3-38 HPE Storage Sizer

The HPE Storage Sizer, shown in [Figure 3-38](#), is a sizing tool that helps you design a storage infrastructure to meet the needs of a customer. The Storage Sizer can be downloaded from the HPE website. This is an important feature because it keeps the sizer current and any configuration prepared using this tool will be a valid, fully supported configuration.

The Storage Sizer supports the disk storage subsystem and other storage solutions such as backup systems, network-attached storage (NAS) solutions, and other storage components. The Storage Sizer requires a license.

Storage Sizer provides the following features and benefits:

- Simplifies the process of designing a storage solution

- Applies storage design, licensing, and services rules

- Provides output as a valid, supported configuration that can be imported directly into SBW for quotation

- Provides localized parts and pricing for different geographic regions

- Includes HPE Smart Update technology, which brings new products or functionality to you through an Internet connection

- Encompasses the HPE storage family

- Initiates an update for every product launch as part of the new product introduction process

- Includes new functionality, which was added based on user input, annual surveys, and quarterly focus groups

The Storage Sizer enables you to work with your customers to design a storage infrastructure that will meet their online and offline needs. You can define customer requirements, such as:

- Performance requirements with specific metrics

- Business requirements, such as server consolidation

- Pure capacity requirements

For example, additional requirements might include raw capacity, estimated IOPS, replication and backup criteria, and the number of host ports.

Because the tool applies all the HPE SAN design rules, it provides a valid, supported storage infrastructure to meet the requirements of your customer. Use the Storage Sizer when you are not sure which combination of products will best address customer

requirements. This tool lets you try different solutions.

A helpful wizard interface guides you through the process of sizing a SAN by asking a series of questions about the proposed configuration. This wizard is intended for those who have less experience using the Storage Sizer.



Note

To access the Storage Sizer, enter the following URL into your browser:

<https://sizersllb.itsc.hpe.com/swdsizerweb/>

SAN Design Reference Guide

The SAN Design Reference Guide is the single source for SAN configuration support, interoperability, and best practices. The guide provides access to HPE multi-vendor, end-to-end storage networking architectural information, including:

SAN design rules

SAN topologies and supported configurations

SAN design philosophies, security, and management

HPE best practices

SAN components

- Architecture
 - Configurations
 - Implementation
 - New technologies
-



Note

To access the SAN Design Reference Guide, enter the following URL into your browser:

<https://h20566.www2.hpe.com/hpsc/doc/public/display?docId=c00403562>

Additional HPE configuration tools

Additional HPE configuration tools include:

HPE Switch Selector—This selector tool helps with determining the correct HPE

networking product based on specific requirements.



Note

To access the Switch Selector, enter the following URL into your browser:

http://h17007.www1.hp.com/us/en/networking/products/switches/switch-selector.aspx#.WmXEyKhl_-h

HPE Networking Online Configurator—This configurator enables quick and easy creation of price quotations for HPE networking products. Users can create a complex networking solution with a network map, wiring guide, or inclusion of multi-vendor networks.



Note

To access the configurator, enter the following URL into your browser:

<http://h17007.www1.hp.com/us/en/networking/products/configurator/index>

You can use these tools to access information to support configurations:

HPE PartSurfer—This tool provides fast, easy access to service parts information for a wide range of HPE products.



Note

To access PartSurfer, enter the following URL into your browser:

<http://partsurfer.hp.com/search.aspx>

HPE Customer Self Repair Services Media Library—Users can find media for HPE products using this library.



Note

To access this library, enter the following URL into your browser:

<https://sml-csr.ext.hp.com/>

HPE Single Point of Configuration Knowledge (SPOCK)—SPOCK is the primary portal used to obtain detailed information about supported HPE storage product

configurations. In SPOCK, you will find details of supported configurations (hardware and software) as well as other useful technical documents.



Note

To access SPOCK, enter the following URL into your browser:

<https://h20272.www2.hpe.com/spock/>

Learning check

Which HPE standalone application enables you to determine the power consumption and total weight of a configuration?

- A. Server Memory Configurator
- B. Converged Infrastructure Sizing Suite
- C. Power Advisor
- D. HPE Product Bulletin

HPE server management technologies

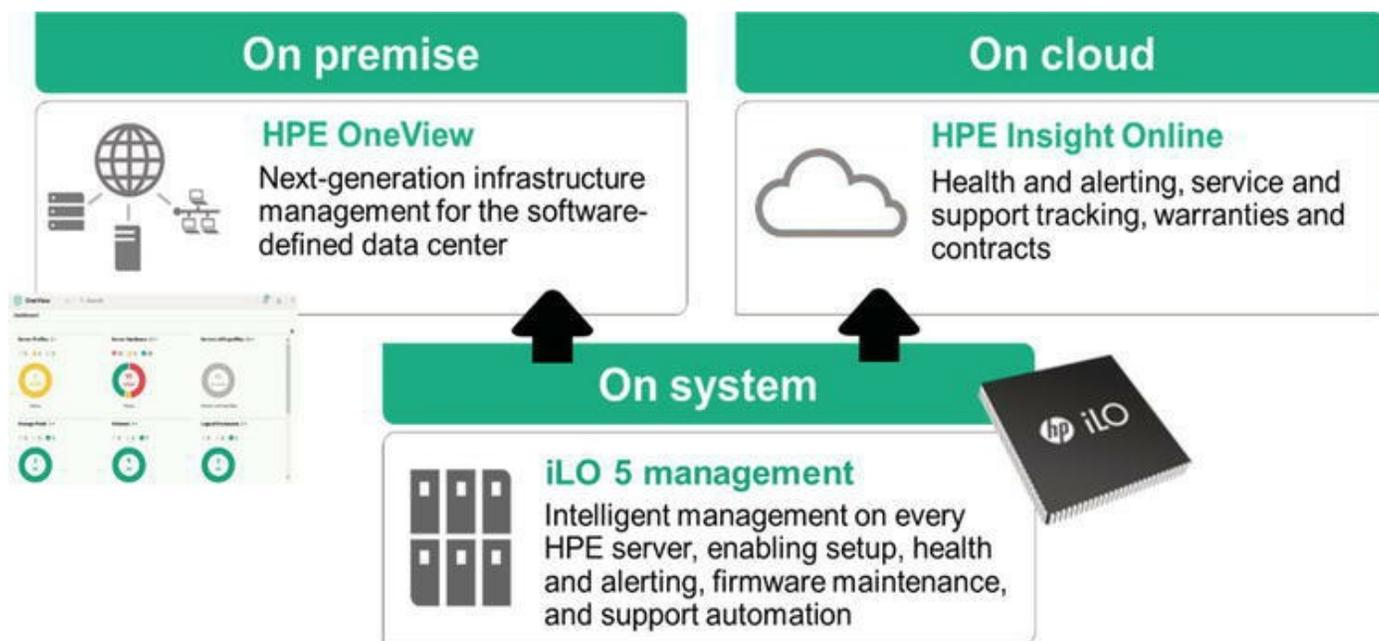


Figure 3-39 On system, on premises, and on cloud offerings

HPE provides a comprehensive set of infrastructure management solutions to help customers operate their increasingly complex data centers. They encompass the lifecycle of critical operations: configuration, provisioning, system health monitoring, firmware updates, and rapid deployment. These offerings belong to one of three categories, as shown in [Figure 3-39](#):

On system—Intelligence on every HPE server enables setup, health and alerting, and firmware maintenance. HPE on-system management also includes the following:

- HPE iLO 5 is embedded into every ProLiant server. It provides essential remote management and control regardless of the servers' state of operation.
- HPE Smart Update Manager (HPE SUM) provides systematic and organized server updates using HPE Service Pack for ProLiant (SPP).
- HPE Intelligent Provisioning (IP) assists in the initial configuration and operating system deployment for ProLiant servers.
- UEFI has replaced BIOS mode for ProLiant Gen10 servers to support server configurations.
- HPE RESTful Interface Tool is a CLI scripting tool that leverages HPE REST application programming interface (API) for secure, remote server configurations at scale.

- HPE Scripting Toolkit for Windows and Linux are utilities that provide scripting for ProLiant server configurations at scale.
- HPE Scripting Tools for Windows PowerShell is a set of utilities that can be used to perform configuration tasks on HPE servers based on PowerShell syntax and scripting models.

On premises—HPE OneView helps you automate and provision an on-demand, hybrid infrastructure. This provides converged management for servers, storage, and networks with automation simplicity.

On cloud—HPE Insight Online enables users to monitor IT devices from anywhere, at any time. This cloud-based management solution is useful for service ticket and warranty tracking, along with health status and alerting for servers, storage, and networking.

On-system management

HPE iLO 5



Figure 3-40 Key functions of iLO 5

HPE iLO 5 provides the automated intelligence to maintain complete server control from anywhere. You can access HPE iLO 5 from any location using a web browser or the iLO 5 mobile application. HPE iLO 5 is embedded into every HPE ProLiant server. It uses dedicated, built-in hardware/firmware for system management and allows access to BIOS settings. iLO 5 enables a system administrator to remotely manage servers even when the server is powered off, regardless of whether the operating system is installed or functional. With iLO 5, system administrators can reinstall the operating system. iLO 5 is a key enabler of embedded remote support, HPE Active Health System (AHS), HPE IP, and HPE Agentless Management.

As shown in [Figure 3-40](#), the key features and functions of HPE iLO 5 are:

Provision—Rapid discovery and remote access features allow customers to inventory and deploy servers using virtual media and iLO Federation remotely with the HPE iLO 5 web interface, remote console, CLI, or mobile application.

Monitor—System health and performance protection with advanced power and thermal control enables customers to achieve maximum power efficiency. In addition, Agentless Management helps to monitor core hardware and related alerts without needing to install agents or providers on the host operating system.

Optimize—An integrated remote console enables users to access, control, and configure system properties.

Support—Core instrumentation that operates whether the operating system is up or

down allows users to view the Integrated Management Log and download AHS logs to send to HPE support for faster problem identification and resolution.

Security—Immutable Silicon Root of Trust for Secure Start with ability to automatically rollback to known good firmware, Common Access Card (CAC) 2-factor authentication support, Run-time Firmware Validation to verify the integrity of iLO and BIOS firmware.

Single interface for server automation

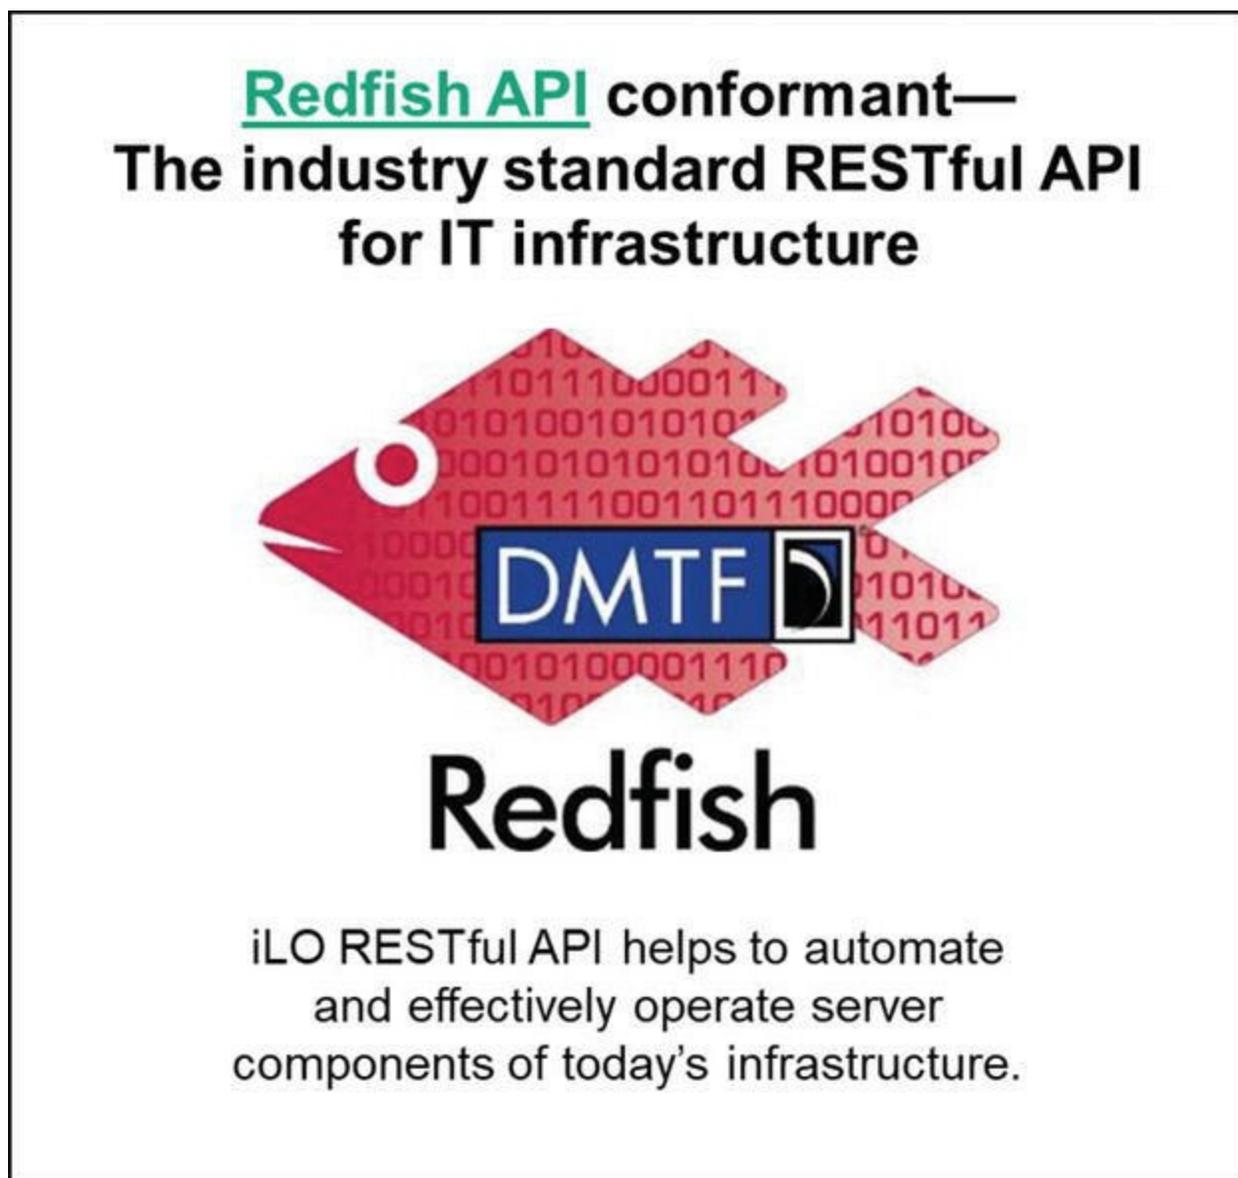


Figure 3-41 Redfish specifies a RESTful interface

When HPE pioneered the original iLO management processor, the industry responded

with commodity baseboard management controllers, typically based on the Intelligent Platform Management Interface (IPMI) protocol. With the iLO 5 Management Engine, HPE innovated again by introducing an interface based on the modern REST approach being favored for systems and software. This so-called RESTful API provides robust control that benefits traditional IT architectures as well as administrators with cloud and web-based data center infrastructures that achieve scalability by deploying large numbers of basic servers. Managing these scale-out architectures poses particular challenges for data center planners and administrators. Traditional interfaces like IPMI are not able to address the scale-out and cloud-based requirements for simplicity and security available in modern programming interfaces.

Redfish takes the work initially implemented as the iLO RESTful API and uses it as a foundation for an open, industry-standard specification and schema for today’s cloud and web-based data center infrastructures sponsored and controlled by the Distributed Management Task Force, Inc. (DMTF), a peer-review standards body recognized throughout the industry, as shown in [Figure 3-41](#). Redfish establishes a new management standard for system control that is scalable, easy to use, and secure. HPE is a founding member of the Redfish specification and most major vendors are now participating in the effort. Redfish defines the industry-standard for the Software Defined Data Center (SDDC) and the effort to modernize heterogeneous data centers. In addition, HPE ProLiant servers expose iLO RESTful API extensions, allowing customers to experience the full range of value-add features available from a programmable interface.

Redfish and Software Defined Compute (SDC) provide the control plane for IT infrastructure enabling customers to program simple configuration and maintenance tasks. At scale, these tasks have been extremely complex, fragile, and time-consuming. Redfish also enables development of new higher-level automation and orchestration features, previously only possible in large-scale, highly customized service provider environments.

Table 3-5 Features and benefits of adopting Redfish

	Features	Benefits
Provision	• Retrieve server information details	• Easy access and configuration for remote management
	• UEFI configuration	• Manage server settings and deploy at scale
	• Manage UEFI Secure Boot	• Software Development Kit (SDK) provides Python sample code available on GitHub to get started with your server deployment
	• Boot Order	

Management

- iLO 5 configuration
- Find iLO NIC MAC/Networking details
- iLO User Account Management
- Virtual Media Mount/Unmount

Monitor	<ul style="list-style-type: none">• Diagnostics• System Memory Details• Smart Array and NIC Inventory• Get/Clear iLO and IML Logs• Get Power Detail	Monitor servers remotely using commonly available tools or custom software
----------------	---	--

Health	<ul style="list-style-type: none">• AHS download	Engage with HPE support using AHS logs
---------------	--	--

Manage	<ul style="list-style-type: none">• Reset Server• Reset iLO• Smart Array configuration with Gen10 Smart Array controllers• Stage and update server components with iLO Repository• Redfish conformance updates for iLO, BIOS, Secure Boot,	Secure data and remove critical information prior to server retirement
---------------	--	--

Redfish is based on current software architecture, methods, tools, and scripting environments used in most IT environments. As presented in [Table 3-5](#), HPE customers can now benefit from adopting the Redfish standard into their HPE ProLiant Gen8, Gen9, and Gen10 installed base, as well as on other non-HPE server products that implement to the standard.

Customers asked for a modern interface in which they expect APIs to use the protocols, structures, and security models that are common in emerging cloud interfaces. Specifically, customers asked for RESTful protocols that express data in JavaScript Object Notation (JSON) formats.

Redfish as “based on the tools and scripting environments most users already have enabling feature-rich remote management while being compatible with the existing tool set. Redfish was built from the ground up to scale to the modern multiple server environments encountered in today’s enterprise, hyper-scale, and cloud infrastructures.”

iLO 5 license types

Every ProLiant server ships with iLO 5 standard features, which simplify server setup, engage health monitoring, monitor power and thermal control, and enable remote administration. Licensing provides additional iLO 4 functionality. Customers can choose from multiple levels of licensing to suit their business needs:

HPE iLO Standard license—The basic iLO 5 license supports all servers. Features include iLO Federation Discovery, remote serial console, pre-boot health summary, iLO 5 reboot, RESTful API, Agentless Management, health monitoring, web-based GUI, and so on. The features included with the Standard license are also included with all other licenses.

HPE iLO Advanced license—These licenses are ideal for enterprise data centers. The HPE iLO Advanced or HPE iLO Advanced for BladeSystem licenses offer premium remote functionality for any ProLiant server. This license enables the full implementation of iLO Federation for features such as Group Firmware Updates, Group Virtual Media, Group Power Control, Group Power Capping, and Group License Activation. Customers can also activate the enhanced integrated remote console; virtual keyboard, video, and mouse (KVM); multi-user console collaboration; advanced security; power and performance management functionality; and many others.

**Note**

HPE iLO Standard and iLO Advanced licenses are also available as iLO Standard for BladeSystem and iLO Advanced for BladeSystem.

HPE iLO Advanced Premium Security Edition—This license includes all that the iLO Advanced license provides with the addition of Automatic secure recovery, Runtime FW validation, Secure Erase of user data, and Commercial National Security Algorithms.

**Note**

To learn more about included features and servers supported with each iLO 5 license, refer to the HPE iLO Licensing Guide. To access this guide, enter the following URL into your browser:

<http://h20564.www2.hpe.com/hpsc/doc/public/display?docId=c04951959>

**Note**

HPE iLO Standard functions out-of-the-box without additional software installation, regardless of the state of operation of your server. To gain access to advanced functionality, you only need to unlock it with an add-on license.

Activity: Exploring the iLO 5 interface

This activity requires you to use Internet Explorer or Chrome to browse to the emulator index page. On the emulator index page select **Exploring the iLO 5 Interface**.

Follow the instructions on the screen. As you progress from page to page, answer each question about the iLO 5 interface. Compare your answers to those presented at the end of the exercise.



Note

Enter the following URL into your browser to visit the iLO 5 emulator page:

<http://hpe.azureedge.net/gen10emulators/index.html>

After using the emulator please answer these questions:

Which types of tasks can be accomplished using the iLO 5 interface?

Which menus or options in the iLO 5 interface will be most useful to your role?

What is the most interesting feature included in the emulator?

Activity: HPE iLO 5 license installation

This activity requires you to browse to the emulator index page. On the emulator index page select **Installing an iLO License**.

Follow the instructions on the screen.



Note

Follow the hyperlink below to visit the iLO 5 emulator page:

<http://hpe.azureedge.net/gen10emulators/index.html>

After using the emulator please answer these questions:

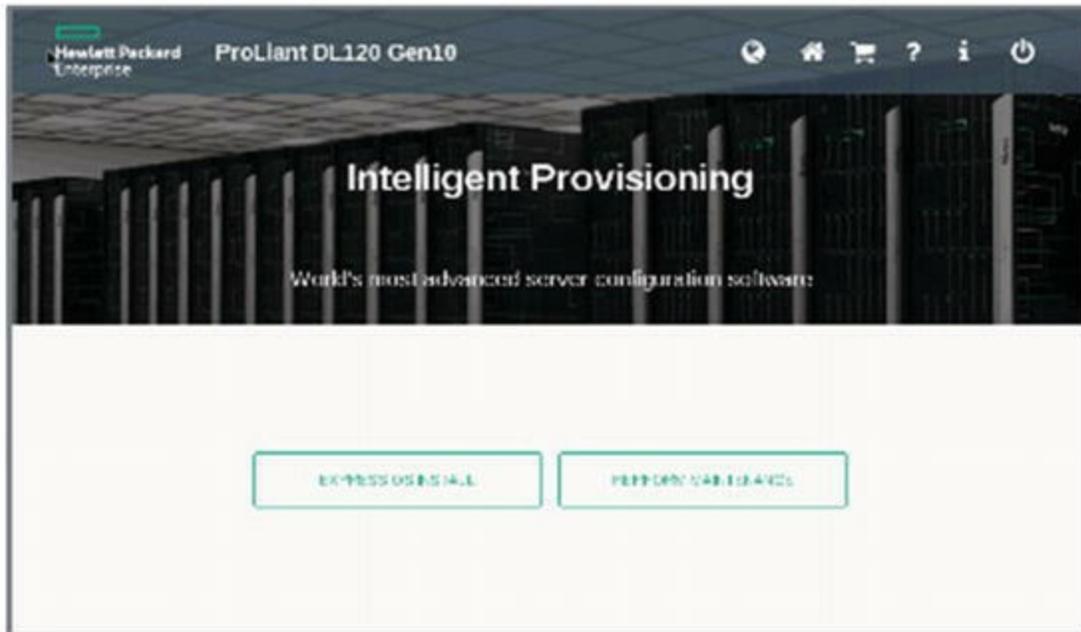
What license did you install?

Which additional features were available following installation of the new license?

What is the most interesting feature included in the emulator?

Intelligent Provisioning

Supported on all Gen10 platforms with iLO



Intelligent Provisioning is a wizard-based embedded server provisioning utility that makes it extremely easy to take your server from out of the box into production.

Figure 3-42 IP is supported on all Gen10 platforms with iLO

As shown in [Figure 3-42](#), IP is a single-server deployment tool embedded in ProLiant Gen10 and Gen9 servers and HPE Synergy Compute Modules that replaces the SmartStart CDs and Smart Update Firmware DVD used with previous generations of ProLiant servers. IP simplifies server setup, providing a reliable and consistent way to deploy servers. IP provides hassle-free provisioning without the need for external media, simple interface to check the availability of updated drivers and firmware, and end-of-life decommissioning is a key security requirement for some customers.

Key new features:

IP is now available via iLO; simply press F10 at boot

No need for any optical discs or external media

Faster updates and improved wizards for complex configurations

- Launch IP up to 3× faster on Gen10 servers
- Up to 22% faster initial set up with IP with Gen10 servers vs. Gen9

Integrated erase utility which can be used to wipe system data and reset hardware settings

Access to all Advanced BIOS (RBSU) settings from IP

Activity: HPE IP

You perform this activity individually. This activity requires you to use Internet Explorer or Chrome to browse to the emulator index page. On the emulator index page select **Deploying an Operating System**.

Follow the instructions on the screen. As you progress from page to page, answer each question about the iLO 5 interface. Compare your answers to those presented at the end of the exercise.



Note

Follow the hyperlink below to visit the emulator index page:

<http://hpe.azureedge.net/gen10emulators/index.html>

After using the emulator please answer these questions:

Which key did you press during power-on self-test (POST) to launch IP?

Which operating system did you install?

What is the most interesting feature included in the emulator?

Firmware management tools

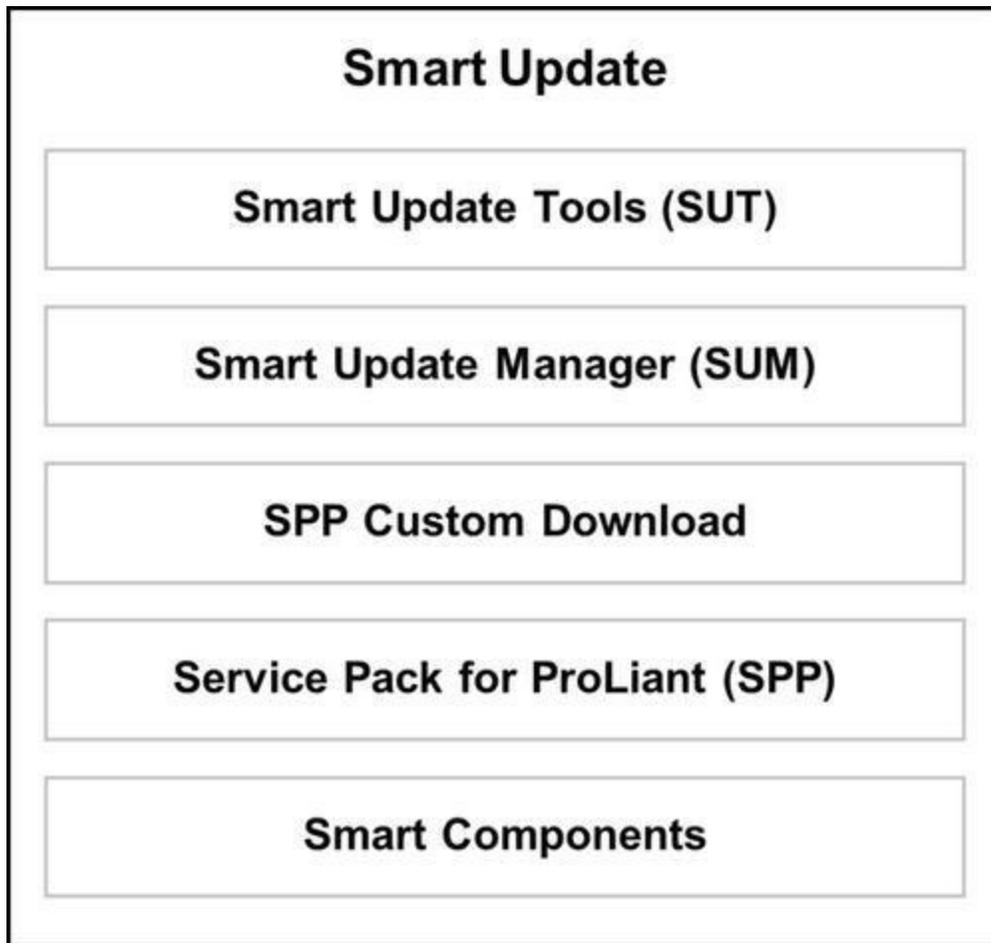


Figure 3-43 Smart Update components

SUM and the SPP combine to create the HPE Smart Update Technology to solve the problem of time-consuming, expensive, and error-prone updates. SUM is an innovative tool for keeping the firmware, drivers, and system software of HPE ProLiant, HPE Synergy, HPE BladeSystem, and HPE Moonshot infrastructures and associated options up to date and secure.

SUM discovers the installed hardware and the current versions of firmware, drivers, and system software; provides an update recommendation; and applies the updates in an efficient order to reduce impact on operations. SUM provides a variety of interfaces for applying updates, so you can pick the interface that meets your needs. Update the servers while they are offline or online, locally or remotely via web browser, or interactively or automatically.

New in HPE Gen10 and HPE Gen9 servers, Smart Update has been enhanced to provide secure, tamper-free updates. HPE firmware, driver, and software updates now include a digital signature that prevents unauthorized or modified updates from being applied. Through integrations with HPE iLO Amplifier Pack and HPE OneView, Smart

Update Technology has been extended to provide simple and secure server updates at scale from a few servers to thousands of servers.

Increases ease of server management

Firmware and driver updates are also smarter than ever with Smart Update. HPE iLO with Integrated Smart Update utilizes the HPE iLO Repository to reduce maintenance windows, overall downtime, and the number of personnel required to execute firmware updates. Enhanced rollback capabilities also ensure peace of mind by enabling IT administrators to revert back to a known good update, so you can be confident and in control of your environment. Staging and scheduling of updates is also a new feature executed through the RESTful API.

SUM allows the maintenance to come from the administrator at a console with a browser-based GUI, CLI, or interactive CLI versus having to physically touch each unit, thus reducing staff travel time.

In combination with the SPP, HPE delivers a complete solution for keeping HPE systems up and running with the latest firmware and drivers.

SUM integrations with HPE OneView and HPE iLO Amplifier Pack allow firmware, driver, and system software updates from within HPE OneView and HPE iLO Amplifier Pack.

Integrated information on dependencies, including those of the HPE Onboard Administrator and HPE Virtual Connect, based on extensive testing of each SPP.

Deployment features including live logs which provide detailed information of target update process.

New features:

Stage and schedule updates ahead of time and activate later

Integrated with iLO to provide secure updates via the management network

Firmware, driver, and system software validated as a complete set (SPP) ensures secure updates are applied efficiently and correctly

Built-in rollback to revert to last known good (bookmarked) update

Tamper-proof updates; firmware updates accessible only through iLO and digitally validated

Convenient update process with minimal manual intervention via Smart Update integrations with iLO and OneView

iLO Amplifier Pack

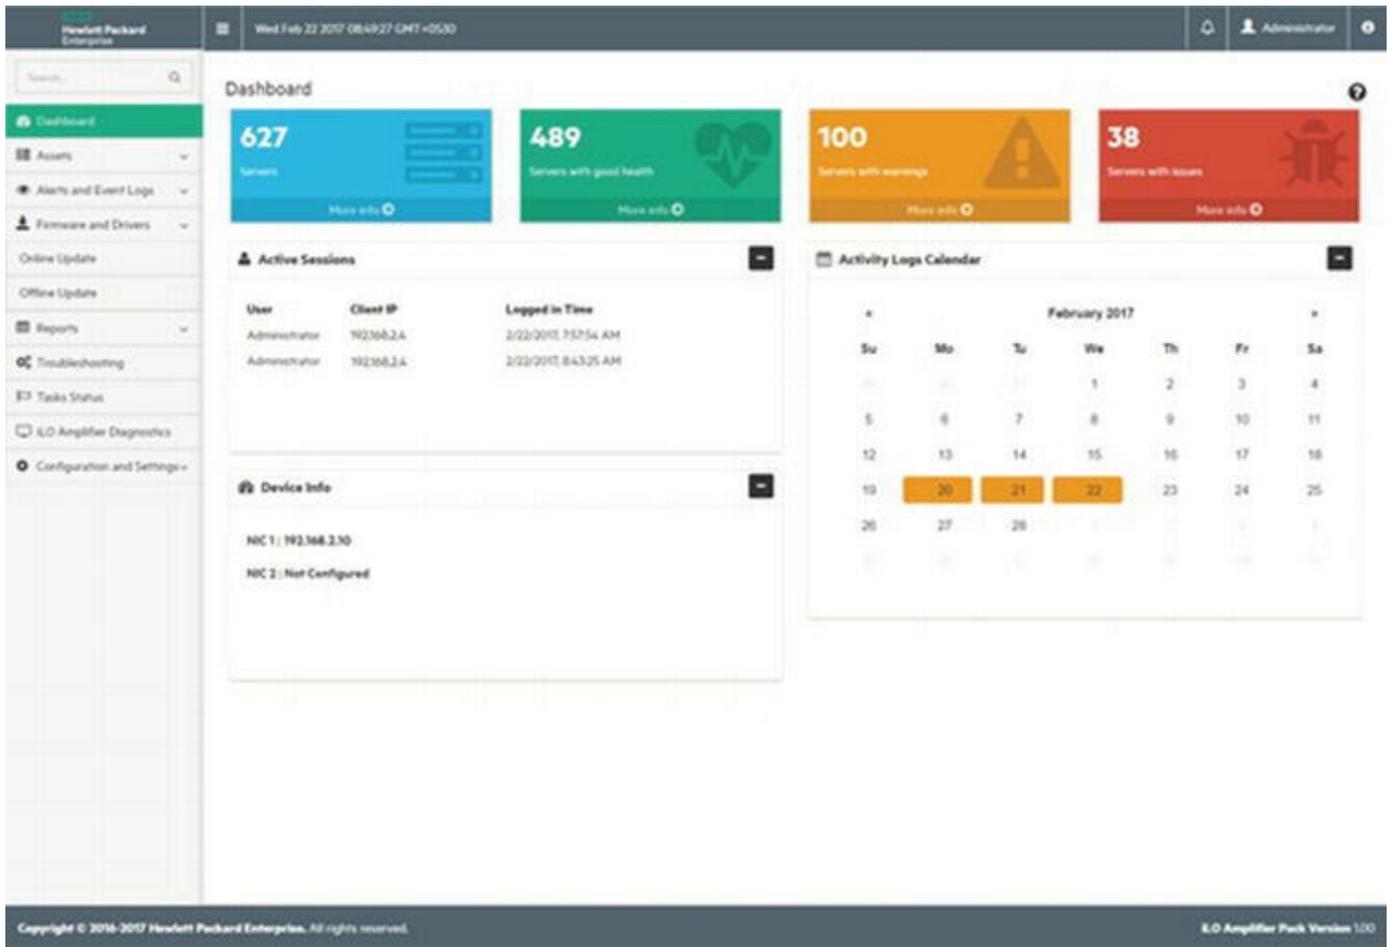


Figure 3-44 iLO Amplifier Pack

The iLO Amplifier Pack, shown in [Figure 3-44](#), is an advanced server inventory and firmware and driver update solution that enables rapid discovery, detailed inventory reporting, and firmware and driver updates by leveraging iLO advanced functionality. The iLO Amplifier Pack performs rapid server discovery and inventory for thousands of ProLiant Gen 10, Gen9, and Gen8 (Rack, Tower, Blade, and Apollo) servers for the purpose of updating firmware and drivers at scale.

iLO Amplifier Pack is a fast, massively scalable inventory and update management tool from HPE for iLO customers. You can use iLO Amplifier Pack to help you manage the following types of common scenarios efficiently and with minimal downtime.

Table 3-6 Features and benefits of iLO Amplifier Pack

Features	Benefits
Discovering • Adding a	iLO Amplifier Pack has a clean, intuitive GUI

- server
- Adding servers in an iLO Federation Group from the Discovery page
- Adding servers in an IPv4 address range
- Adding servers from a CSV file

that is easy to use and can add servers and groups one at a time or thousands at a time. Discovery takes only a few minutes and does not require server downtime.

Monitoring

- Viewing the dashboard
- Viewing the server list
- Viewing inventory details
- Viewing iLO Federation groups
- Viewing server alerts

iLO Amplifier Pack allows you to monitor the overall health of your infrastructure from a single page in your browser. Drill down for detailed information about individual servers or groups.

Reporting

- Viewing the firmware report
- Viewing the iLO license

Use the options from the Reports menu to view and download up-to-date reports.

report

- Viewing the basic device report
- Viewing the hardware inventory report

Managing

- Managing server UID status
- Managing server power status
- Configuring remote syslog
- Mounting virtual media
- Creating and joining iLO Federation groups

Script maintenance is not necessary. You can accomplish tasks on a large scale with no customization upkeep required.

Updating

- Performing an offline firmware update
- Performing an online firmware update
- Running and completing tasks

iLO Amplifier Pack simplifies updating tasks by automating the update process requiring limited user interaction and minimal downtime.

[Table 3-6](#) lists features and benefits of iLO Amplifier Pack.

AHS Viewer

HPE AHS is an essential component of the iLO Management Engine portfolio. It is an industry-first technology that provides continuous, proactive health monitoring of over 1600 system parameters. 100% of configuration changes are logged for more accurate problem resolution. SmartMemory and SmartDrive devices log failure information to accurately document events and avoid inadvertent reuse of failed components. This enables customers to start problem analysis faster and spend less time reproducing or describing errors. Support is easy, just hand off the AHS log file to HPE support and get the issue resolution started. You can also view the error information contained in the AHS Viewer by uploading your AHS file to www.hpe.com/servers/ahsv.

Features:

Read AHS logs

Self-repair some server errors

Firmware analysis, ability to read release notes and download firmware

Raise support cases

Rich Configuration History, Health, and Service Alerts

Raising the bar on quality and customer experience

Faster and accurate problem resolution

Consolidated diagnostic mechanism

Always-on proactive diagnostics rather than reactive

Enhanced customer support with AHS Viewer

The ability to read and troubleshoot AHS Data

HPE recommended actions based on experience and best practice within the tool

On-system management with UEFI

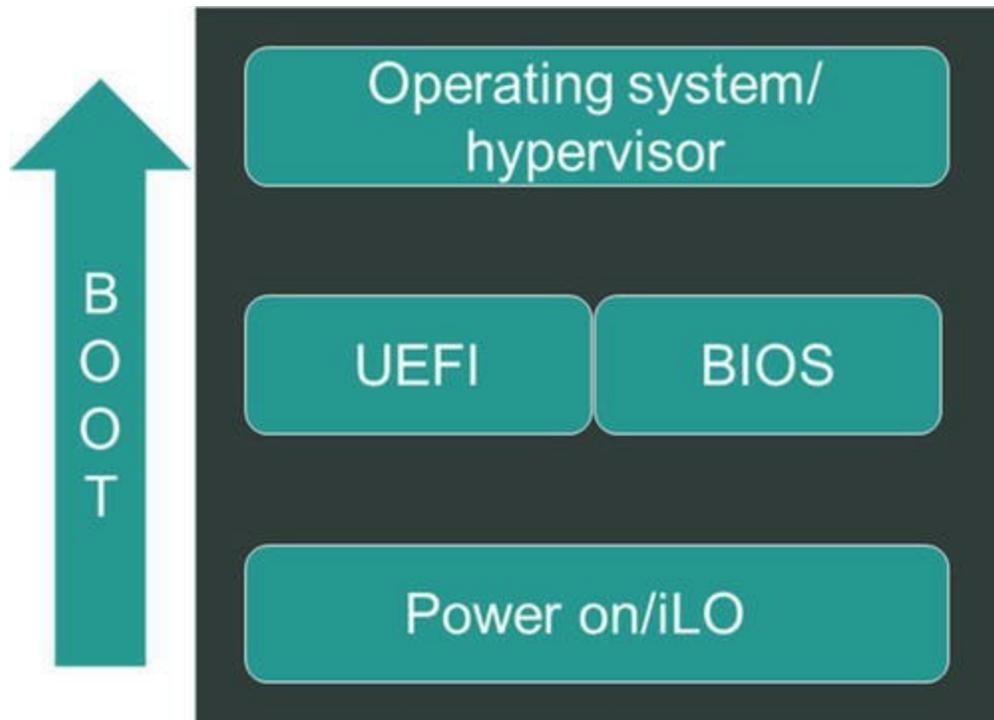


Figure 3-45 Server boot process

The BIOS is software that resides on a chip on the system board within a server. The BIOS starts up the server's components when the server is powered on. It then makes sure the components are functioning properly before passing off that functionality to the operating system as part of the POST.

The UEFI specification is replacing more than 30 years of legacy BIOS evolution and bringing major changes to the pre-boot environment. UEFI is a community effort governed by the UEFI Forum that defines four platform classes:

Class 0—Legacy BIOS-based only

Class 1—Obsolete

Class 2—Legacy BIOS-based and UEFI

Class 3—UEFI only

As shown in [Figure 3-45](#), UEFI and BIOS work together to boot an OS or hypervisor when power is applied to a server.

HPE has adopted UEFI for ProLiant Gen10 and Gen9 servers because the UEFI specification has evolved with an emphasis on stability, security, and compatibility. The goal is to modernize platform firmware and provide an interface that is not architecture specific. UEFI class 2 provides optional support for legacy BIOS mode. UEFI standardizes interfaces:

Between platform firmware and operating system (both boot and runtime)

Within platform initialization firmware

Within the pre-boot UEFI environment/shell

UEFI specifications are platform independent and designed to support multiple platforms and architectures. They promote cross-functionality and support adoption across multiple operating systems.

Gen 10 UEFI enhancements

UEFI is a specification that defines a software interface between a computer's operating system and firmware. UEFI has been developed as a potential replacement for the BIOS firmware interface commonly found on many servers and PCs. In practice, most UEFI firmware images provide legacy support for BIOS services. UEFI can support remote diagnostics and repair of computers, even without another operating system. The UEFI interface includes data tables that contain platform-related information along with boot and runtime service calls that are available to the operating system and running pre-boot applications.

Customer needs:

Secure local or remote boot

Support for new devices like Persistent Memory

Enterprise-class reliability and performance

Key new features:

Secure Start powered by Silicon Root of Trust

Support for Scalable Persistent Memory

Up to 67% reduction in boot time with similar configurations with Gen10 servers vs. Gen9

Enablement of memory RAS features such as Address-Based Memory Mirroring and HPE Fast Fault Tolerant Memory

Smart Array configuration now available on UEFI boot mode

Redfish API conformant

New pre-boot environment user interface optimized for either keyboard or mouse

Workload Profiles for performance optimization



Note

Access the Hewlett Packard Enterprise Information Library to get more information. To access the related documents, enter the following URL into your browser:

<https://www.hpe.com/uk/en/product-catalog/detail/pip.6935826.html>

Activity: UEFI Shell

This activity requires you to browse to the emulator index page. On the emulator index page select **UEFI Shell**.

Follow the instructions on the screen.



Note

Follow the hyperlink below to visit the emulator index page:

<http://hpe.azureedge.net/gen10emulators/index.html>

After using the emulator please answer these questions:

In which menu did you find the Embedded UEFI Shell?

Which option did you use to get a page by page output from the help command?

What did you do to become resident on the fs1 file system?

What did you do to execute the Red Hat boot loader?

What is the most interesting feature included in the emulator?

On-premises management with HPE OneView



Deploy Infrastructure Faster

IT generalists can quickly respond to changing business requirements by rapidly and reliably composing and updating resources using automated templates created by IT specialists.



Simplify Operations

The HPE OneView Global Dashboard simplifies operations by providing a single, unified view of the health of thousands of servers, profiles and enclosures across multiple data centres.



Increase Productivity

Automate resource provisioning, configuration and monitoring with the HPE OneView unified API. Developers and ISVs can use automation for more aligned, responsive service delivery.

Figure 3-46 HPE OneView

As shown in [Figure 3-46](#), HPE OneView lets you take an innovative approach to infrastructure management based on software-defined intelligence. Deploy infrastructure faster, simplify lifecycle operations, and improve productivity with efficient workflow automation, a modern dashboard, and the industry's broadest partner ecosystem.

Simplify hybrid IT environment with a management platform that transforms servers, storage, and networking into software-defined infrastructure. This latest release of HPE OneView incorporates a number of feature enhancements, including extended platform support, simplified firmware management, global dashboard enhancements, common storage management, and several new composable ecosystem integrations.

Through software-defined intelligence, HPE OneView takes a template-driven approach to provisioning, updating, and integrating compute, storage, firmware/device drivers, and networking infrastructure. This approach not only reduces the risk of human error by enabling you to develop the template once and then replicate as needed but also helps to boost the productivity of administrators and software developers. In addition, change operations—such as adding more storage to a service, modifying network connectivity, or updating firmware—can also be implemented via templates so that changes are implemented automatically.

Continuous, automated lifecycle operations reduce cost, save time, and increase time to value for your business. Templates go far beyond just initial deployment. With them, you can also simplify system updates and enforce compliance to ensure infrastructure stability. You can manage deployment plans and create bootable images from capturing, cloning, or customizing golden images. Further, you can enforce compliance using templates to quickly provision, update, or roll back images to minimize maintenance windows.

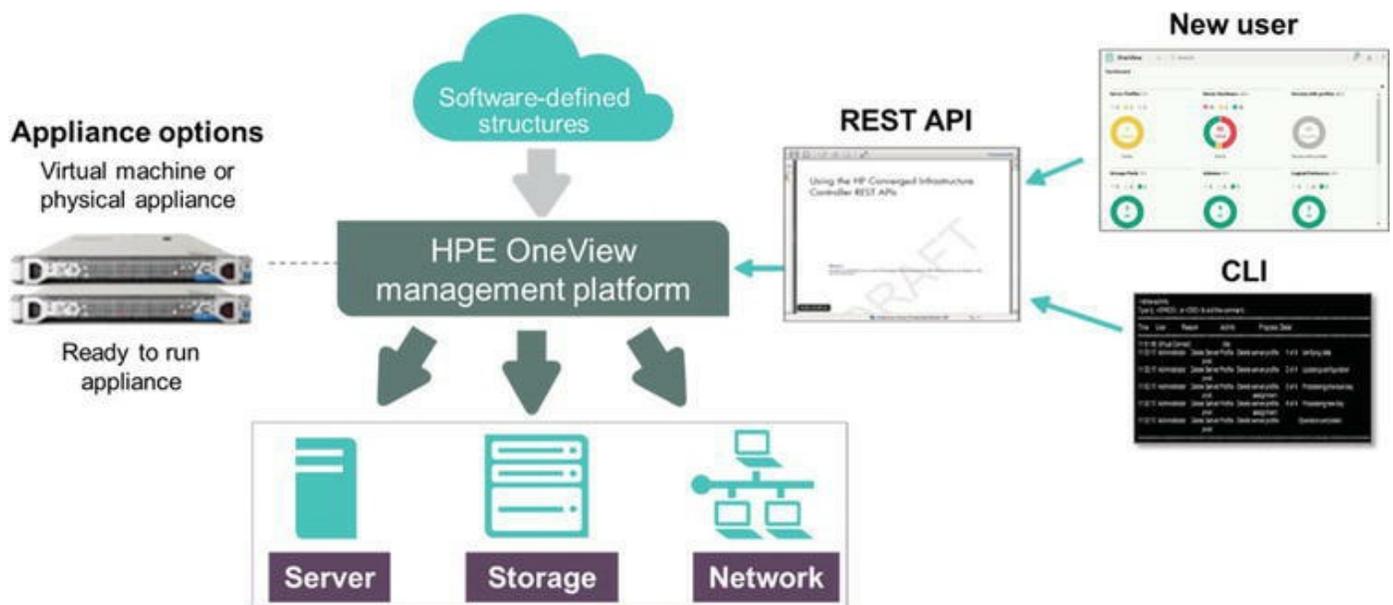


Figure 3-47 HPE OneView has a single, integrated platform

HPE OneView is a software-defined management platform that addresses the challenges of manual operation, human error, and limited extensibility in virtualized HPE BladeSystem and rack server environments. As shown in [Figure 3-47](#), it is a single, integrated platform that provides one view of server, network, and storage resources. HPE OneView is available as a virtual and physical appliance. It captures processes, configurations, and best practices for a simple, integrated user experience. This modern, collaborative approach automates the deployment and management of infrastructure—repeatedly, reliably, and at scale.

HPE OneView accelerates all data center processes. The software-based approach to lifecycle management in HPE OneView automates operations to reduce the cost and time to deliver IT services. The RESTful API allows users to easily create customized workflows and scripts, as well as configuration profiles for push-button builds that instantly deliver resources without mistakes or variation.

HPE OneView provides:

Software-defined process templates for:

- Automated infrastructure configuration and provisioning
- Robust infrastructure health and monitoring

A single, integrated management environment for a converged infrastructure that enables IT teams to work and collaborate in a more natural and automated way

An intuitive interface and powerful search capabilities

Open architecture and an SDK for integration with enterprise management tools and applications



Note

HPE OneView is intended to replace the capabilities found in HPE SIM, HPE Insight Control, and HPE Virtual Connect Enterprise Manager (VCEM).

On-cloud management with HPE Insight Online

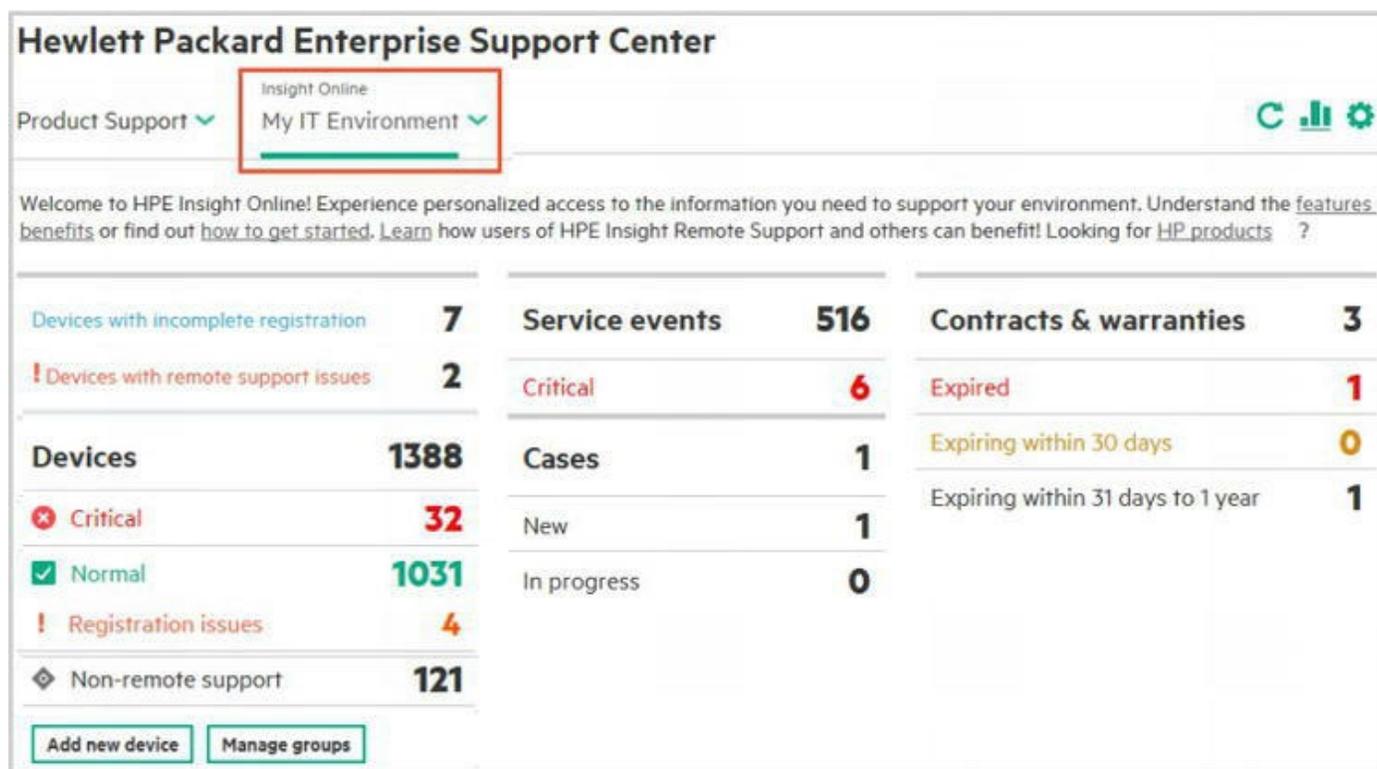


Figure 3-48 HPE Insight Online

HPE Insight Online, shown in [Figure 3-48](#), is a cloud-based infrastructure management and support portal available through the HPE Support Center. Insight Online provides the information you need to monitor devices in the IT environment from anywhere and at any time. Insight Online and related remote support tools are available at no additional cost, as part of the HPE warranty or support agreement.

Insight Online:

Is powered by HPE remote support technology

- Insight Online direct connect for ProLiant Gen10, Gen8, and Gen9 servers

- HPE Insight Remote Support software

Provides online access to:

- Devices
- Service events
- Support cases
- Contracts and warranties

Insight Online provides automated support for a converged infrastructure of servers, storage, and networking devices using a personalized, cloud-based dashboard. The Insight Online dashboard allows you to:

Track service events and support cases by device or user

View device configurations

Monitor HPE contracts and warranties as well as HPE Proactive Care Services credit balances

- Create reports for support cases, contracts, and warranties

Speed problem resolution time by up to 66%

Organize and share information

Access incident trending reports (for Proactive Care Services customers only)

Learning check

Write a summary of the key concepts presented in this chapter.

Summary

HPE ProLiant Gen10 rack and tower servers redefine compute economics by delivering more compute and storage capacity; offering less power and floor space consumption; and providing faster compute, memory, and IO performance.

The design technologies embedded in ProLiant servers allow customers to increase application performance and free up resources to focus on business innovation. Key technologies include HPE PCIe Workload Accelerators, power capping capabilities, HPE iPDUs, HPE Flex Slot power supplies, and HPE Smart Storage Batteries.

HPE provides a comprehensive set of infrastructure management solutions to help customers to operate increasingly complex data centers. These offerings belong to one of three categories:

- On-system
- On-premises
- On-cloud

4 HPE ProLiant Tower Servers

WHAT IS IN THIS CHAPTER FOR YOU?

After completing this chapter, you should be able to describe the high-level features of the following HPE tower servers:

ProLiant ML350 Gen10

ProLiant ML110 Gen10

ProLiant MicroServer Gen10

ProLiant ML10 Gen9

ProLiant ML30 Gen9

ProLiant ML110 Gen9

ProLiant ML150 Gen9

Prelearning check

You are attending a meeting with a customer to discuss their remote office server requirements. Four out of the 10 offices have fewer than 10 employees. They need a low-cost server to provide local print services and a network gateway with simple to use management. How should you respond?

A customer approaches your stand at a trade fair and asks you to suggest the best server solution to provide a low cost, highly available server platform for a stand-alone virtual machine environment. How should you respond?

HPE ProLiant tower servers

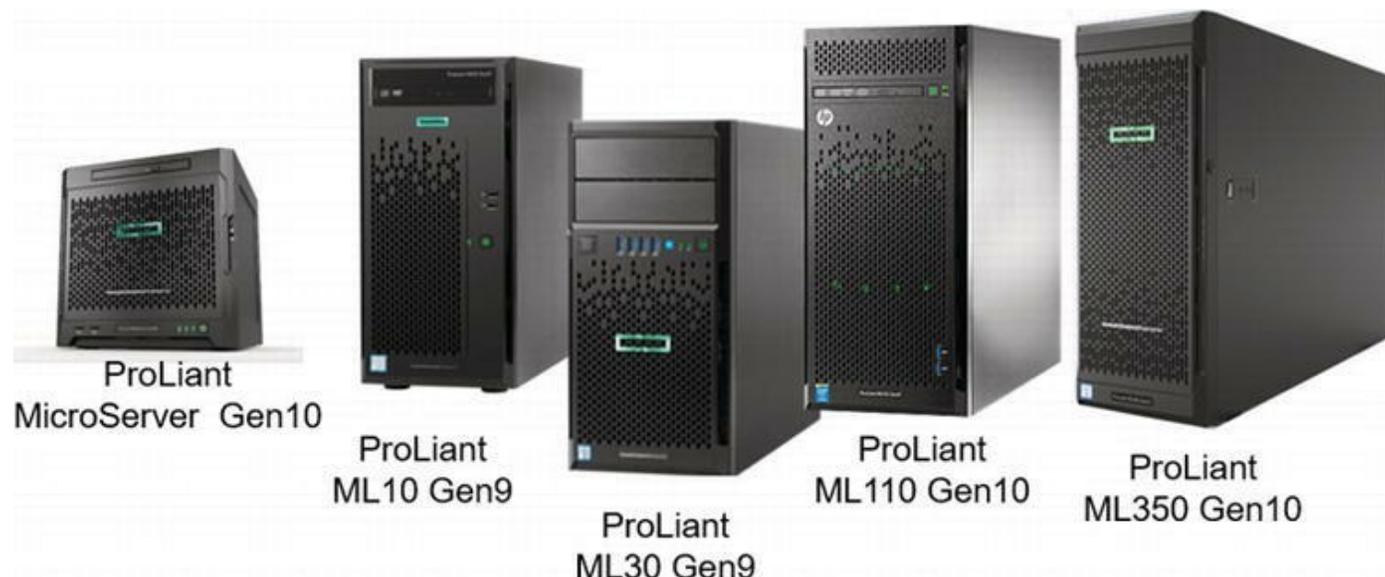


Figure 4-1 HPE ProLiant Gen10 and Gen9 tower servers

ProLiant servers, like the ones shown in [Figure 4-1](#), are based on industry-standard x86 architectures that enjoy industry-wide application support. The HPE commitment to participate in joint development activities with processor manufacturers and HPE solution development, service, and support ensures that ProLiant servers provide a familiar and trusted platform.

HPE ProLiant servers are separated into five main product-lines.

DL servers—Density-optimized for stacking in rack-mounted server environments ideal for multiserver deployments

ML servers—Tower- and rack-based servers designed with capacity for maximum internal expansion

BL servers—Comprise blade servers that fit within the HPE BladeSystem

SY servers—Compute modules that fit within the HPE Synergy solution

XL servers—Compute modules that fit within the HPE Apollo solution

ProLiant servers are also split into several series that denote processor configuration.

The 10, 300, and 400 series comprise single- and dual-socket systems.

The 500 and 600 series comprise quad-socket-capable systems.

Models where the last digit is zero are built with Intel® Xeon® processors inside (such as DL380), and those with five at the end are built to support AMD Opteron processors (such as DL385).

HPE ProLiant ML servers are stand-alone units that contain all of the components required to respond to requests from client computers. These tower servers are relatively compact. They can be used in work areas that are not specifically designed to accommodate servers, meaning that customers do not need a special data room or bays in which to install servers. These simple, robust ProLiant servers provide a complete infrastructure that supports customer business objectives and growth.

Compute innovations in the ProLiant Gen10 tower portfolio include simple management and storage tools, along with proven configurations that provide easy remote access and improved energy efficiencies to lower total cost of ownership (TCO). Integrated with a simplified but comprehensive management suite and industry-leading support, the ProLiant Gen10 tower server portfolio delivers business value, helps to increase IT staff productivity, and expedites service delivery. In addition, the tower portfolio includes financing options, along with a service and channel network to significantly increase the speed of IT operations.

ProLiant ML servers are ideal for maximum internal storage and input/output (IO) flexibility. They are also available in rack deployment options.

HPE ProLiant ML servers

HPE ProLiant Gen10 ML servers are relatively compact, stand-alone tower servers.

The ProLiant Gen10 tower portfolio delivers:

Easy-to-use tools, simple processes, and reliable support to help server administrators keep hardware running

Affordability to increase business agility and help acquire and retain customers

The ProLiant ML Gen10 family includes the following server models:

ProLiant ML350 Gen10

ProLiant ML110 Gen10

ProLiant MicroServer Gen10

The ProLiant ML Gen9 family includes the following server models:

ProLiant ML10 Gen9

ProLiant ML30 Gen9

ProLiant ML110 Gen9

ProLiant ML150 Gen9

HPE ProLiant ML350 Gen10

Introducing

Performance with unmatched capacity and reliability—all in a much smaller form factor



Availability, expandability, and serviceability—a winning combination



Agile infrastructure management for accelerating IT service delivery



The ML350 Gen10 Server is a 2P server providing high performance, expandability, and manageability—ideal for expanding SMBs and enterprise businesses.



Figure 4-2 HPE ProLiant ML350 Gen10 Server

The ProLiant ML350 Gen10, as shown in [Figure 4-2](#), is a dual-processor premium server that delivers a class-leading combination of performance, expandability, and manageability. The ProLiant ML350 Gen10 provides powerful converged management capabilities for the infrastructure lifecycle with embedded server management for provisioning, updating, and diagnostic support through iLO 5.

This server offers the option of eight to 24 small form factor (SFF) drives (8×3 cages) or four to 12 large form factor (LFF) drives (4×3 cages) and increased IO expansion. It features one or two Intel[®] Xeon[®] Skylake Series, four to 28 cores with up to 70% performance gain, and HPE DDR4 SmartMemory offering up to 14% performance increase when compared to the previous generation.

The ML350 Gen10 Server is ideal for enterprise IT infrastructure to mission-critical applications. With the HPE ProLiant ML350 Gen10 Server, you can deploy a single platform to handle a wide variety of enterprise workloads.

General purpose server applications (email, collaboration, and database)—Improve productivity of your organization. Best match for performance and best growth possibilities for growing medium sized businesses

Data warehouse/analytics—Find the information you need, when you need it, to enable better business decisions

Storage-centric applications—Remove bottlenecks and improve performance

Customer relationship management (CRM)—Your customers will gain a 360-degree view of their data to improve customer satisfaction and loyalty

Enterprise resource planning (ERP)—Trust the ML350 Gen10 to help your customers run their business in near real-time

Virtualization—Consolidate server footprint by running multiple workloads on a single ML350

Virtual desktop infrastructure (VDI)—Deploy remote desktop services to provide your customer's workers with the flexibility they need to work anywhere, at any time, using almost any device

SAP—Your customers can streamline their business processes through consistency and real-time transparency into their end-to-end corporate data

ProLiant ML350 Gen10 Server features include:

Up to two Intel® Xeon® Skylake Series, up to 28 cores

Additional support for 12 Gb/s SAS and four embedded 1 GbE NICs with a broad range of graphics and compute options

Up to eight non-volatile memory express (NVMe) PCIe solid state drives (SSDs)

Choices of SFF x24 or LFF x12 hard disk drive (HDD)/SSD

Graphic support of NVIDIA M6000 and AMD FirePro W7100

S100i Storage Controller and a choice of HPE Flexible Smart Array (Adaptive Raid On Chip—AROC) or Smart Host Bus Adapter (HBA) Controllers for performance or additional features

Up to 24 DIMM slots of HPE DDR4 SmartMemory delivering 2666 MT/s (3 TB Max)

Persistent Memory up to 12 NVDIMM options

Large expansion capacity with eight PCIe 3.0 expansion slots, seven USB ports, one Micro SD, 5U rack conversion

Unified Extensible Firmware Interface (UEFI) boot mode, Intelligent Provisioning, and Scripting Toolkit

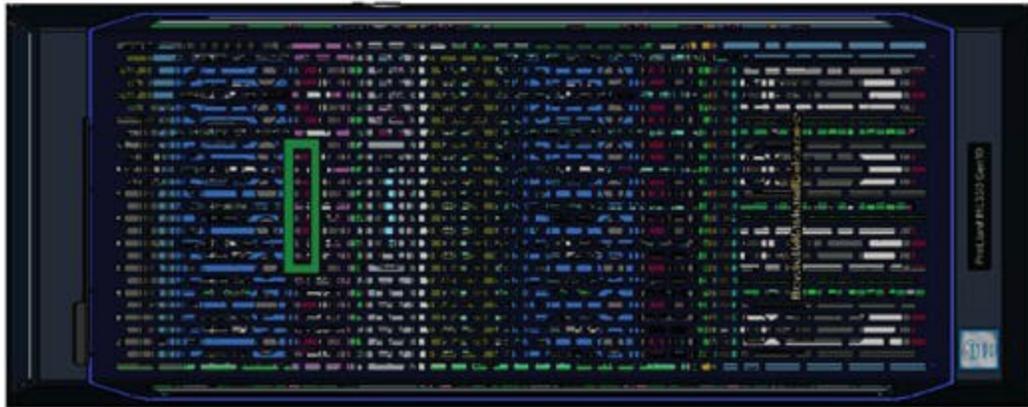
HPE iLO 5, UEFI BIOS and HPE Secure Encryption Capabilities

ASHRAE A3 and A4, lower idle power and ENERGY STAR qualified server configurations

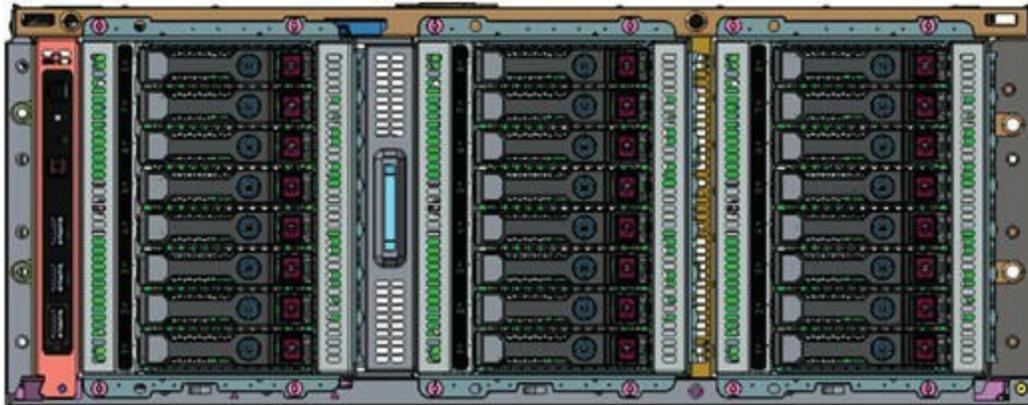
Up to 96% efficiency (Titanium) Flexible Slot redundant power supply (RPS) 500W/800W/1600W new Gen10 RPS and Battery Backup unit

Storage options

Support up to **24SFF / 12LFF**



24 (8*3) SFF Model



16 (8*2) SFF HDD/SSD + 8 PCIe NVMe SSD

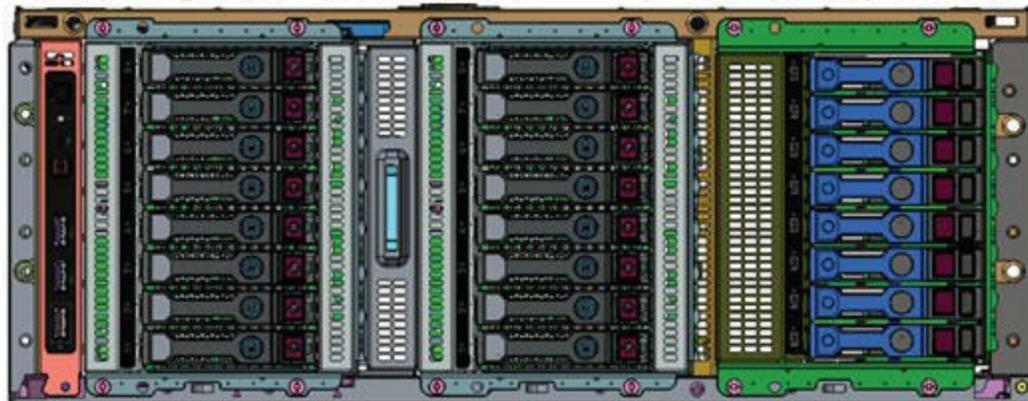


Figure 4-3 HPE ProLiant ML350 Gen10 storage options

[Figure 4-3](#) shows some of the storage options available for the HPE ProLiant ML350 Gen10 Server.

ProLiant ML110 Gen10 Server

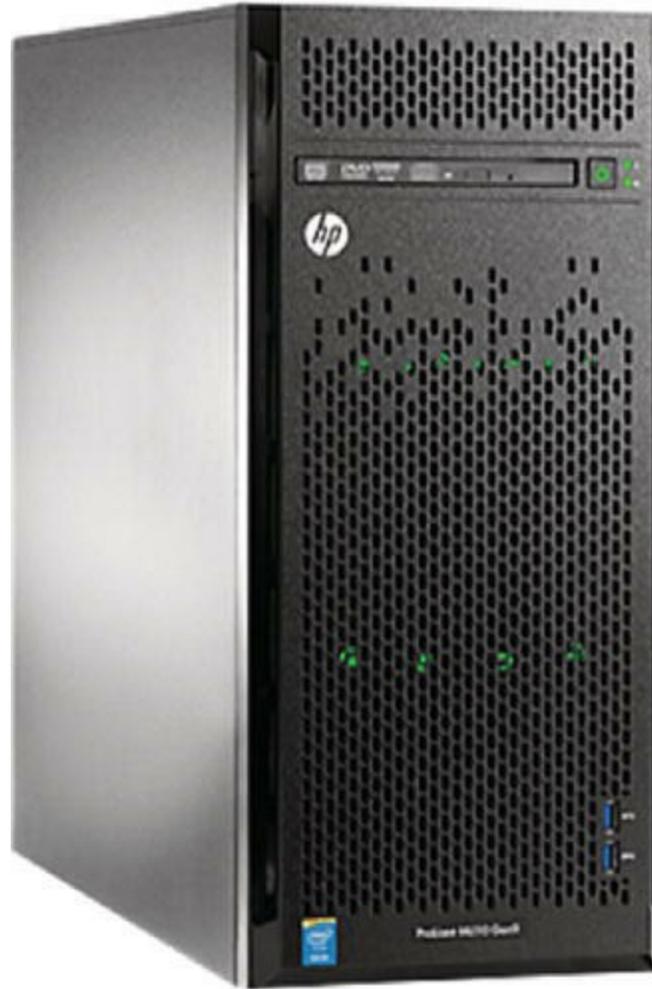


Figure 4-4 ProLiant ML110 Gen10 Server

The ProLiant ML110 Gen10 Server, as shown in [Figure 4-4](#), is a 4.5U tower with a single processor. It features better performance, expansion, and growth than previous one-processor, single-socket tower models. Designed to meet SMB performance compute demands, the ML110 Gen10 delivers exceptional value at an affordable price. It includes up to five PCIe slots, eight LFF, or 16 SFF disk drives, and eight DDR4 DIMM slots with a maximum capacity of 256 GB for growing business needs. This server is ideal for office environments because its chassis depth is less than 19 inches and it weighs less than 25 kilograms.

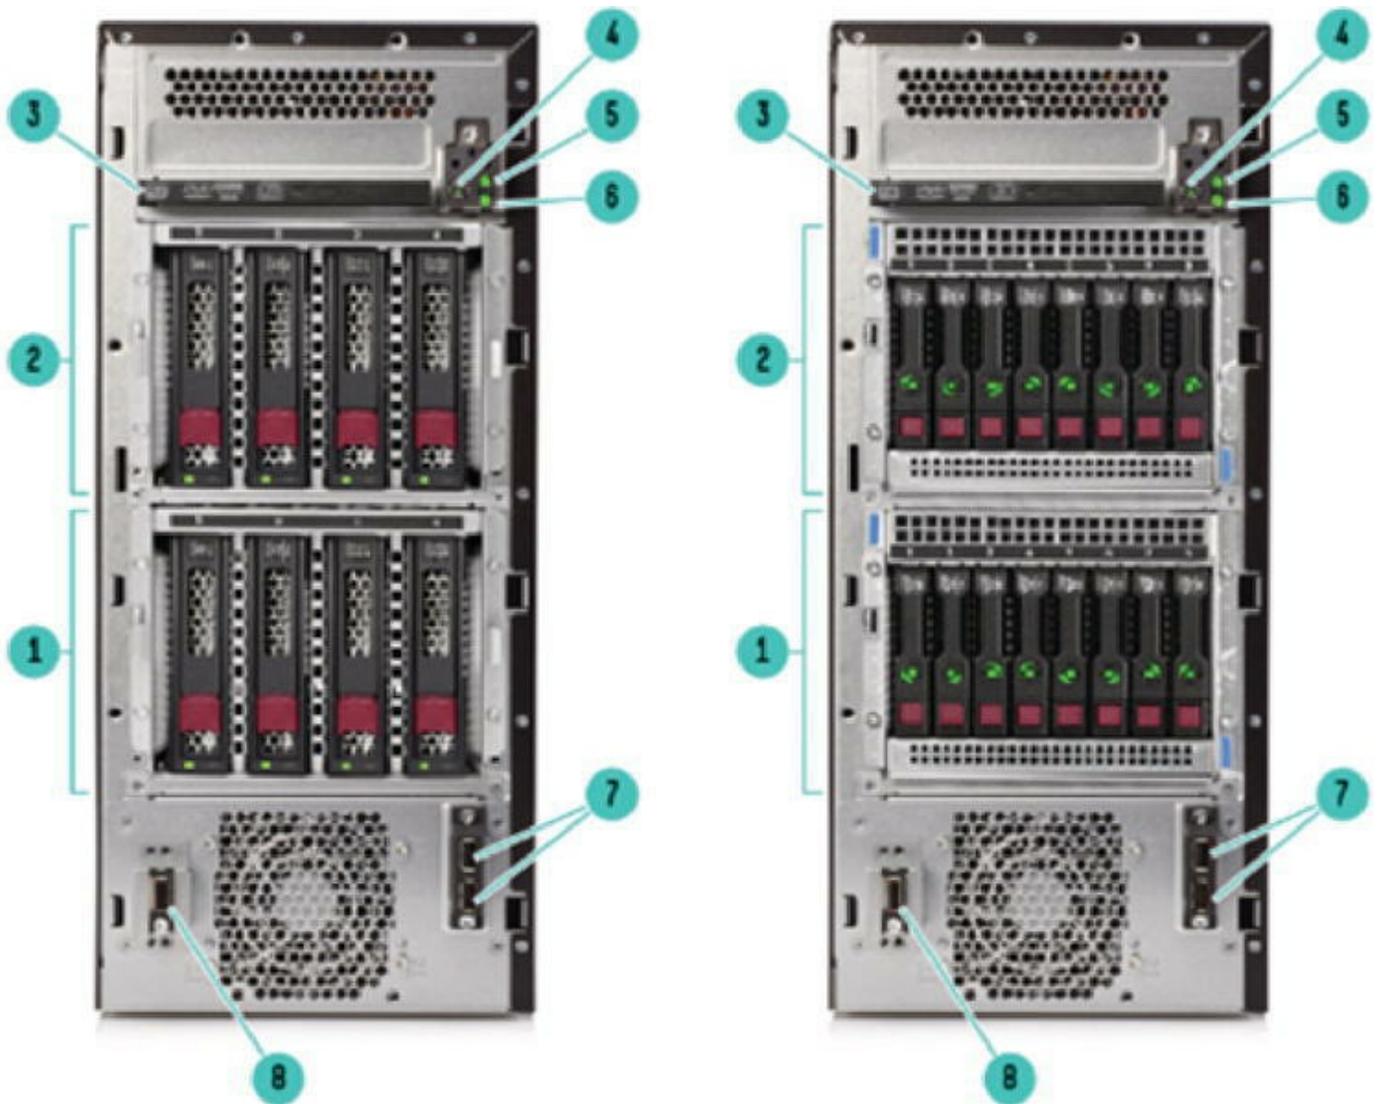


Figure 4-5 HPE ProLiant ML110 front view

[Figure 4-5](#) shows the front view of the ProLiant ML110 with the eight LFF drive option on the left and the 16 SFF option on the right. The numbered components are:

- Drive cage 1
- Drive cage 2 (optional)
- Optical drive (optional)
- Power button/LED
- Health LED
- NIC status LED
- USB 3.0 connectors (x2)

iLO Service Port

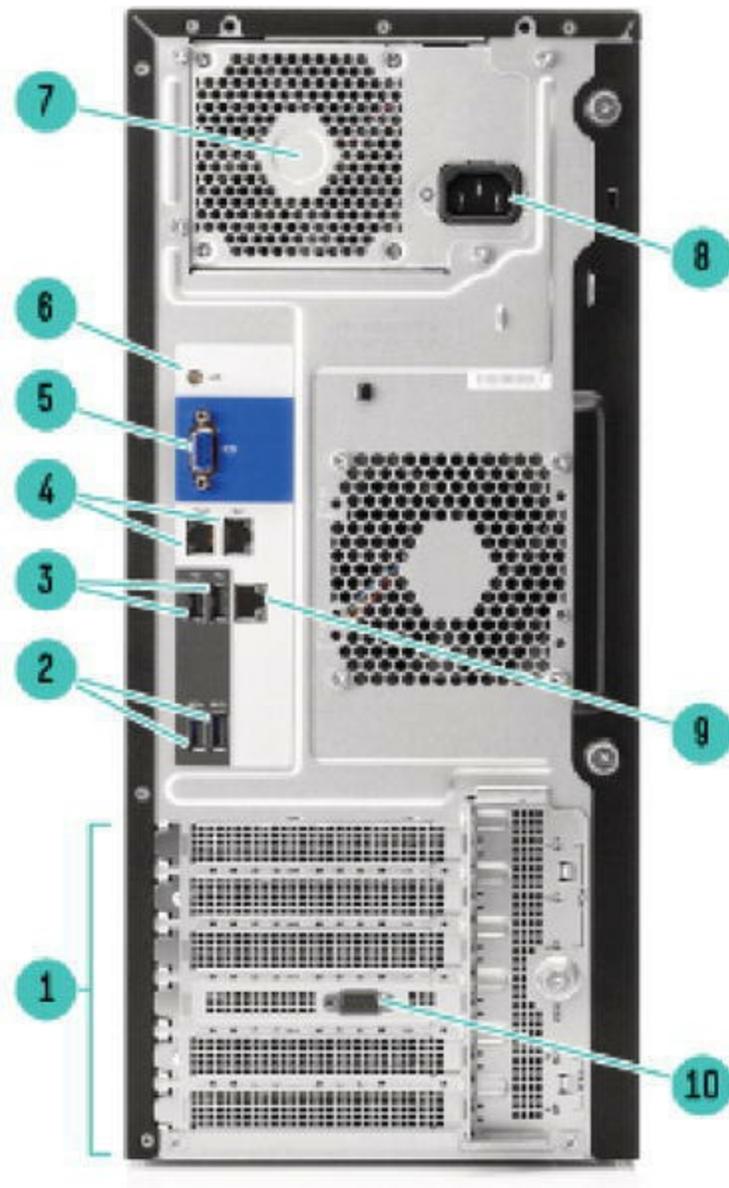


Figure 4-6 HPE ProLiant ML110 rear view

[Figure 4-6](#) shows the rear view of the HPE ProLiant ML110. The numbered components are:

- PCIe 3.0 slots (slots 1-5)
- USB 3.0 connectors (x2)
- USB 2.0 connectors (x2)
- Network RJ-45 ports (x2)
- Video connector

UID button/LED

Power supply bay

Power supply power connection

iLO management port

Serial port (optional)

The ProLiant ML110 Gen10 features include:

One Intel® Xeon® Processor Scalable Gold, Silver, Bronze Series, up to 14 cores, 105W

Up to six DIMM slots of HPE DDR4 SmartMemory delivering 2666 MT/s (192 GB Max)

Five PCIe 3.0 expansion slots

Energy efficiency with ASHRAE A3 compliance. This means that the server can operate at ambient temperatures up to 40 degrees Celsius.

Standard HPE Dynamic Smart Array S100i, HPE Smart HBA Controllers for performance or additional features

16 SFF (8 × 2 cages) hot plug/eight LFF (4 × 2 cages) hot plug/non-hot plug max, HDD/SSD m.2 enablement

2 × 1 GbE embedded + Standup

Rear VGA and Optional Serial, **eight** USB 3.0, one Micro-SD

One standard ATX Multi-Output PSU 350W/550W up to 85% efficiency

Up to two HPE Flex slot RPS 94%–96%, RPS option

A single-wide or double-wide Graphics processing unit (GPU) support

Front iLO Service Port, HPE iLO 5, HPE SUM, HPE RESTful Interface Tool, UEFI

HPE ProLiant ML350 and ML110 Gen10 comparison

Table 4-1 Compute, memory, and storage options

	ProLiant ML350 Gen10 (2P/4U)	ProLiant ML110 Gen10 (1P/4.5U)
Compute	Up to (two) Intel® Xeon® Skylake Series, four to 28 Cores up to 205W (shelf 1/2/3/4), PCIe 3.0, up to (eight) slots	Up to (two) Intel® Xeon® Skylake Series, four to 14 Cores up to 105W (shelf 1/2/3) PCIe 3.0, up to (five) slots
Memory	HPE Smart Memory (24) DDR4 (six channels per CPU), up to 2666 MT/s (3 TB max)	HPE Smart Memory (six) DDR4, up to 2666 MT/s (192 GB max)
Persistent Memory	Up to (12) HPE NVDIMM, HPE Scalable Persistent Memory	None
Storage	Standard HPE Dynamic Smart Array S100i Choice of HPE Flexible Smart Array (AROC) or HPE Smart HBA Controllers for performance or additional features	Standard HPE Dynamic Smart Array S100i HPE Smart HBA Controllers for performance or additional features
FBWC	2 GB DDR3-1866 MHz, 72-bit wide bus at 14.9 GB/s on P450ar	Optional
Battery	HPE DL/ML/SL 96W Smart Storage Battery	HPE DL/ML/SL 96W Smart Storage Battery
HPE SmartDrives	24 SFF (8*3 cages)/12 LFF (4*3 cages) hot plug max, HDD/SSD, non-hot plug 12 LFF—lower cost option, eight NVMe PCIe SSD m.2 enablement	16 SFF (8*2 cages) hot plug/8 LFF (4*2 cages) hot plug/non-hot plug max, HDD/SSD m.2 enablement

Networking	4 × 1 GbE embedded + Standup	2 × 1 GbE embedded + Standup
VGA/Serial/USB Ports/SD	Rear VGA, rear Serial Port, rear Display Port, seven USB 3.0, one Micro-SD	Rear VGA and Optional Serial, eight USB 3.0, one Micro-SD
GPU support	Single/Double-Wide and Active/Passive up to 10.5" (four)	Single/Double-Wide and Active up to 10.5" (two)

Table 4-2 Management, power, and warranty options

	ProLiant ML350 Gen10 (2P/4U)	ProLiant ML110 Gen10 (1P/4.5U)
Management	iLO 5—Enterprise Tier	iLO 5—Mid Tier
Support	HPE Insight Online	HPE Insight Online
Embedded	Front iLO Service Port, HPE iLO 5, HPE SUM, HPE RESTful Interface Tool, UEFI	Front iLO Service Port, HPE iLO 5, HPE SUM, HPE RESTful Interface Tool, UEFI
Power and cooling	(two) 94%–96% Gen9 Flex Slot RPS or new Gen10 RPS 500W/800W/1600W, support 1+1 power redundancy (one) 500W PSU (Beech non-RPS)—lower cost option	(1) Standard ATX Multi-Output PSU 350W/550W up to 85% efficiency (2) HPE Flex slot RPS 94%–96%, RPS option
Industry compliance	ASHRAE A3 and A4, lower idle power, Energy Star	ASHRAE A3, lower idle power, Energy Star (limited configuration)
Power Discovery Services	Yes	N/A
Location Discovery Services	Yes	N/A
Redundant fan	Yes	Yes
Chassis depth	25.5" (Tower/4U Rack, 5U with the T-to-R conversion)	19" (Tower/4.5U Rack)
Serviceability—Easy install rails	Optional—Rack tray	Optional—Rack tray
Warranty	3/3/3	3/1/1

[Tables 4-1](#) and [4-2](#) present a comparison between the ProLiant ML350 Gen10 and ML110 Gen10 servers.

HPE ProLiant ML350 Gen10 to Gen9 comparison

Table 4-3 ProLiant ML350 Gen10 and Gen9 compute, memory, and storage options

	ProLiant ML350 Gen10	ProLiant ML350 Gen9
Compute	Up to (two) Intel® Xeon® Skylake Series, 4/6/8/10/12/14/16/18/20/22/24/26/28 cores (Bronze/Silver/Gold/Platinum) PCIe 3.0, up to (eight) slots	Up to (two) Intel® Xeon® E5-2600 v3/v4 Series 4/6/8/10/12/14/16/18 Cores PCIe 3.0, up to (eight) PCIe 2.0 (one) slots
Memory	HPE Smart Memory (24) DDR4 (six channels per CPU), up to 2666 MT/s (3 TB)	HPE Smart Memory (24) DDR4, up to 2400 MT/s (3 TB) with 128 GB DIMMs
Persistent Memory	Up to (12) HPE NVDIMM, HPE Scalable Persistent Memory	None
Storage	Standard HPE Dynamic Smart Array S100i Choice of HPE Flexible Smart Array (AROC) or Smart HBA Controllers for performance or additional features	Standard HPE Dynamic Smart Array B140i Choice of HPE Flexible Smart Array (AROC) or Smart HBA Controllers for performance or additional features
HPE drives	24 SFF (8 × 3 cages)/12 LFF (4 × 3 cages) HP max, HDD/SSD, NHP 12 LFF—lower cost option, eight NVMe PCIe SSD, and m.2 enablement	48 SFF (8 × 6 cages), 12 LFF (8 × 3 cages) max, HDD/SSD, six NVMe PCIe SSD and m.2 enablement
Networking	4 × 1 GbE embedded + Standup (1/10/25 GbE)	4 × 1 GbE embedded + Standup (1/10/25 GbE)
VGA/Serial/USB Ports/SD	Rear VGA, rear Serial Port, rear Display Port, seven USB 3.0/2.0, one Micro-SD	Rear VGA, rear Serial Port, rear Display Port, seven USB 3.0/2.0, one Micro-SD
GPU support	Single/Double-Wide and Active/Passive up to 10.5” (four)	Single/Double-Wide, Active up to 10.5” (four)

[Table 4-3](#) presents a comparison between the ProLiant ML350 Gen10 and Gen9

compute, memory, and storage options.

HPE ProLiant MicroServer Gen10



Figure 4-7 ProLiant MicroServer Gen10 front view

The HPE ProLiant MicroServer offers a unique form factor that is compact and great for small spaces. The loudest acoustic noise level of a fully loaded HPE MicroServer is 40 decibels, which is about the noise level of a library. Benefits include low acquisition cost, compute optimized for small office, home office, and enabled for graphics-related workloads. [Figure 4-7](#) shows the front view of the ProLiant MicroServer Gen10, with the external view on the left and the internal view on the right. The numbered components for the external view are:

USB 3.0 ports (x2)

Media bay (optional DVD-RW or serial advanced technology attachment [SATA] SSD)

Front bezel lock/unlock indicators

Power on/standby button and system power LED

Health LED

NIC status LED

Drive activity LED

The numbered components for the internal view are:

Hard drive screws

Hard drive bays

ProLiant MicroServer Gen10 features include:

Unique and serviceable design

- Compact and stylish designed for small space
- Easy access to hard drives, memory for simple upgrade
- Simple installation and upgrades with easy access to drives, memory, and PCIe slots

Graphic performance

- Up to 48% better performance with Passmark 2008 measurement than previous generations
- Dual display ports capable of 4K display resolution—4× denser than 1080p FHD—greater clarity, more vibrant, and realistic colors
- Stream media to three devices with optional AMD Radeon Pro WX 2100 PCIe graphics card

Redefine the small business server

- Low acquisition costs
- ClearOS operating system preloaded gives customers an out-of-box ready SMB solution for cloud, gateway, network, and server capability
- Optional read-intensive SATA SSD for boot device
- Optimized compute with AMD Opteron X3000 processor and DDR4 memory

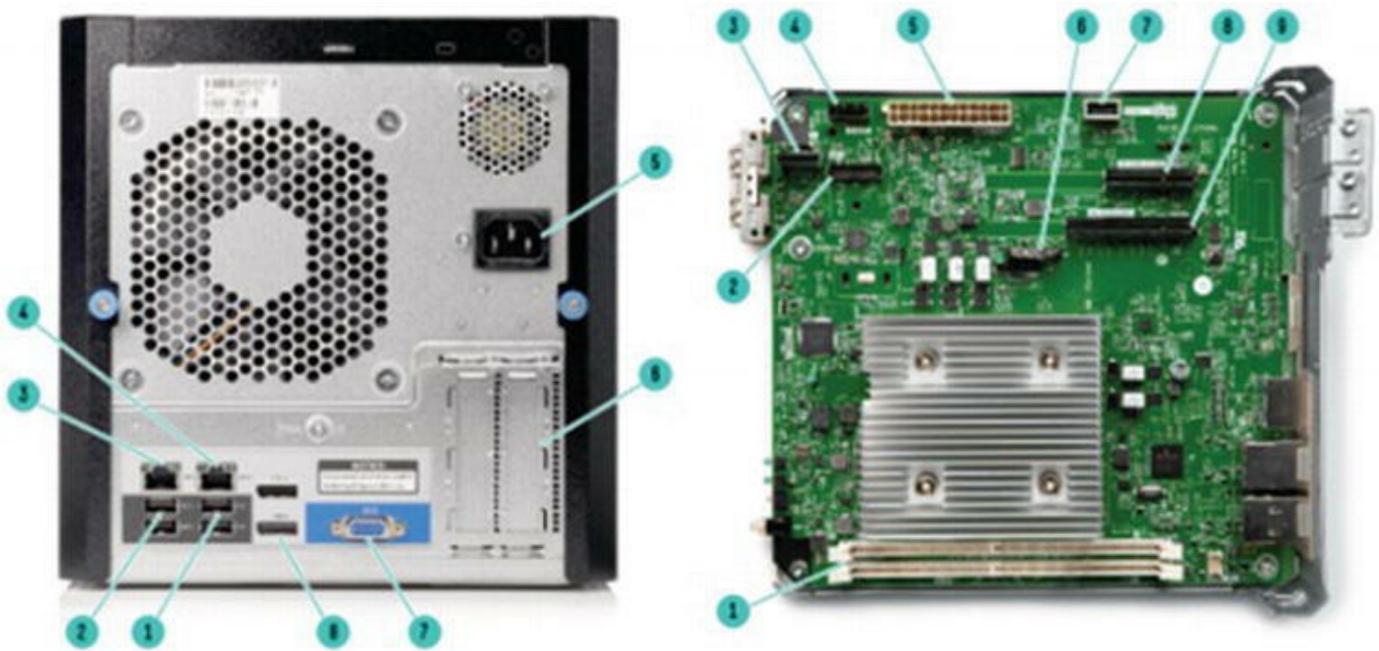


Figure 4-8 ProLiant MicroServer Gen10 rear and internal view

[Figure 4-8](#) shows the ProLiant MicroServer Gen10 rear and internal views, with the rear view on the left and the internal view on the right. The numbered components for the rear view are:

- USB 2.0 ports (x2)
- USB 3.0 ports (x2)
- NIC 1 port
- NIC 2 port
- Power socket
- PCIe Gen3 slots (x2)
- VGA port
- Display ports (x2)

The numbered components for the internal view are:

- DIMM slots (x2)
- Trusted Platform Module (TPM) connector
- LFF/SFF SATA connector
- SATA connector

System board power connector

System battery

Internal USB 2.0 port

Expansion slot 2, PCIe3 x4

Expansion slot 1, PCIe3 x8

ProLiant MicroServer Gen10 options include:

AMD Opteron™ X3216 or X3421 processors

AMD Radeon™ Pro WX 2100 Graphics Card

HPE Smart Array E208i-p SR Gen10 Controller (RAID 0/1/5/10)

Embedded Marvell 88SE9230 PCIe to SATA 6Gb controller (RAID 0/1/10)

HPE Ethernet 10 GB 2-Port 521T Adaptor

Two DIMM slots of HPE Standard Memory DDR4 Unbuffered memory max 32 GB

Five PCIe 2.0 expansion slots

ClearOS on HPE ProLiant

ClearOS is available for preload on HPE ProLiant MicroServer Gen10, ML110, ML30, and DL20 servers making these servers ready to use right out of the box. HPE is also making ClearOS available for download and installation via Intelligent Provisioning on all ProLiant servers 300 series and below. ClearOS is available to download from HPE.com.

HPE ProLiant Gen10 workload-based positioning

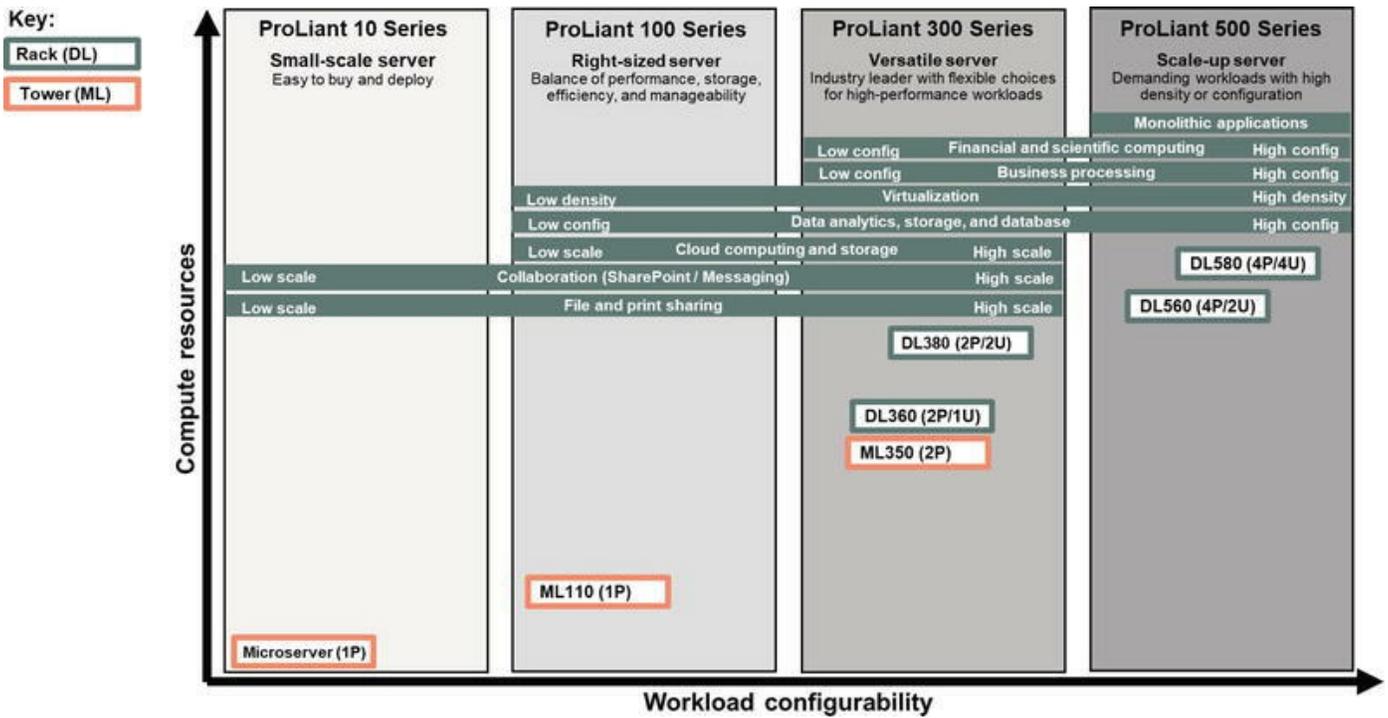


Figure 4-9 HPE ProLiant Gen10 workload-based positioning

Figure 4-9 shows a comparison of the ProLiant ML servers in terms of compute resources, configurability, and types of workloads. The figure also shows the ProLiant DL servers, which are covered in Chapter 5: HPE Rack Server Solutions.

Learning check

Which description is most accurate regarding the ProLiant ML server family?

- A. Ideal for maximum internal storage and IO flexibility
- B. Density-optimized for flexibility and manageability
- C. Ideal for multiserver deployments
- D. Available in rack and cluster models
- E. Ideal for large data center deployments

You are attending a meeting with a customer to discuss their remote office server requirements. Four out of the 10 offices have fewer than 10 employees. They need a low-cost server to provide local print services and a network gateway with simple to use management. What capabilities would be useful to satisfy this request?

ProLiant ML10 Gen9 Server



Figure 4-10 ProLiant ML10 Gen9 Server

The ProLiant ML10 Gen9, shown in [Figure 4-10](#), provides server-level performance in a quiet, compact size that is easy to deploy in home or small business environments. This server provides expandability with internal support for up to six LFF disk drives and increased IO expansion. It offers the lowest acquisition cost for a one-processor tower in the HPE server portfolio.

The workstation-oriented 4U design of the ProLiant ML10 Gen9 provides better performance, powered by just the right amount of processing, storage, memory, and networking to run essential shared applications with the required efficiency. It can be used to secure, collaborate, and provision. The intuitive server management software allows customers to focus on their business instead of managing the server environment.

ProLiant ML10 Gen9 features include:

Intel® Xeon® E3-1200 v5 processor in a 4U chassis offering

Embedded Intel® Active Management Technology (AMT) module with remote console, web browser GUI, and remote power on/off capability

Intel® chipset that supports Intel® Rapid Storage Technology (RST) SATA RAID disk drives

UEFI boot mode, HPE Intelligent Provisioning, and HPE Scripting Toolkit

Optional HPE H241 Smart HBA with the flexibility to run in HBA mode or simple RAID mode

GPU that supports two single-wide GPU card and one double-wide card

ProLiant ML10 Gen9 front view

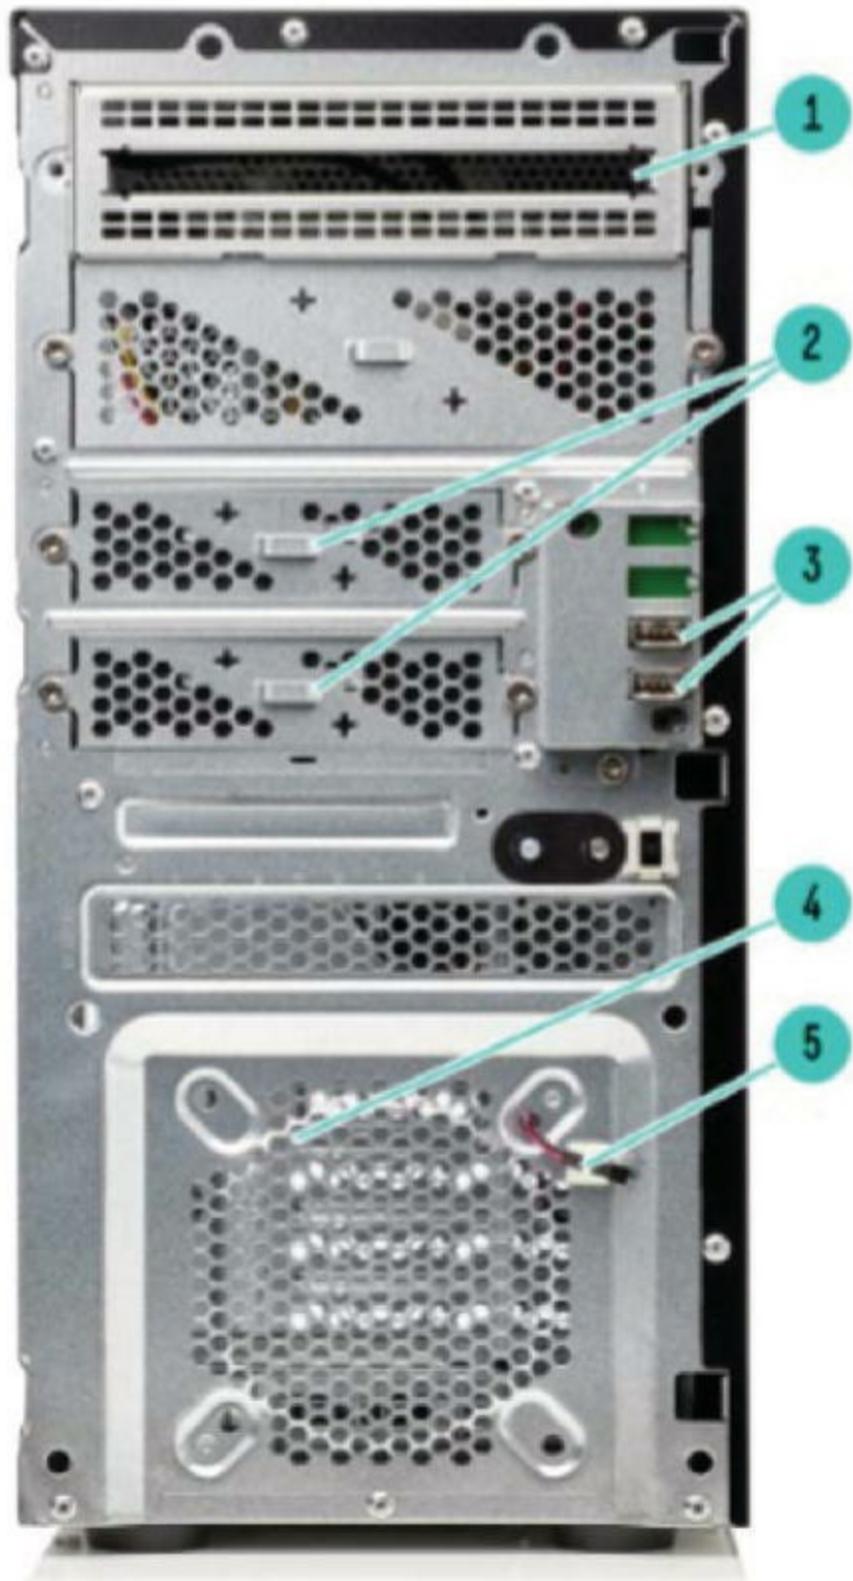


Figure 4-11 ProLiant ML10 Gen9 front view

[Figure 4-11](#) shows the components of a ProLiant ML10 Gen9 (4 LFF configuration)

from the front view. The numbered components are:

Media bay

Hard drive bays

USB 2.0 connectors

Drive cage

Ambient temperature sensor

Enhancements over the previous generation of ML10 server include:

Intel® RST

HPE Smart HBA controllers for external backup

Intel® Advanced Management Technology

UEFI support

ProLiant ML10 Gen9 rear view

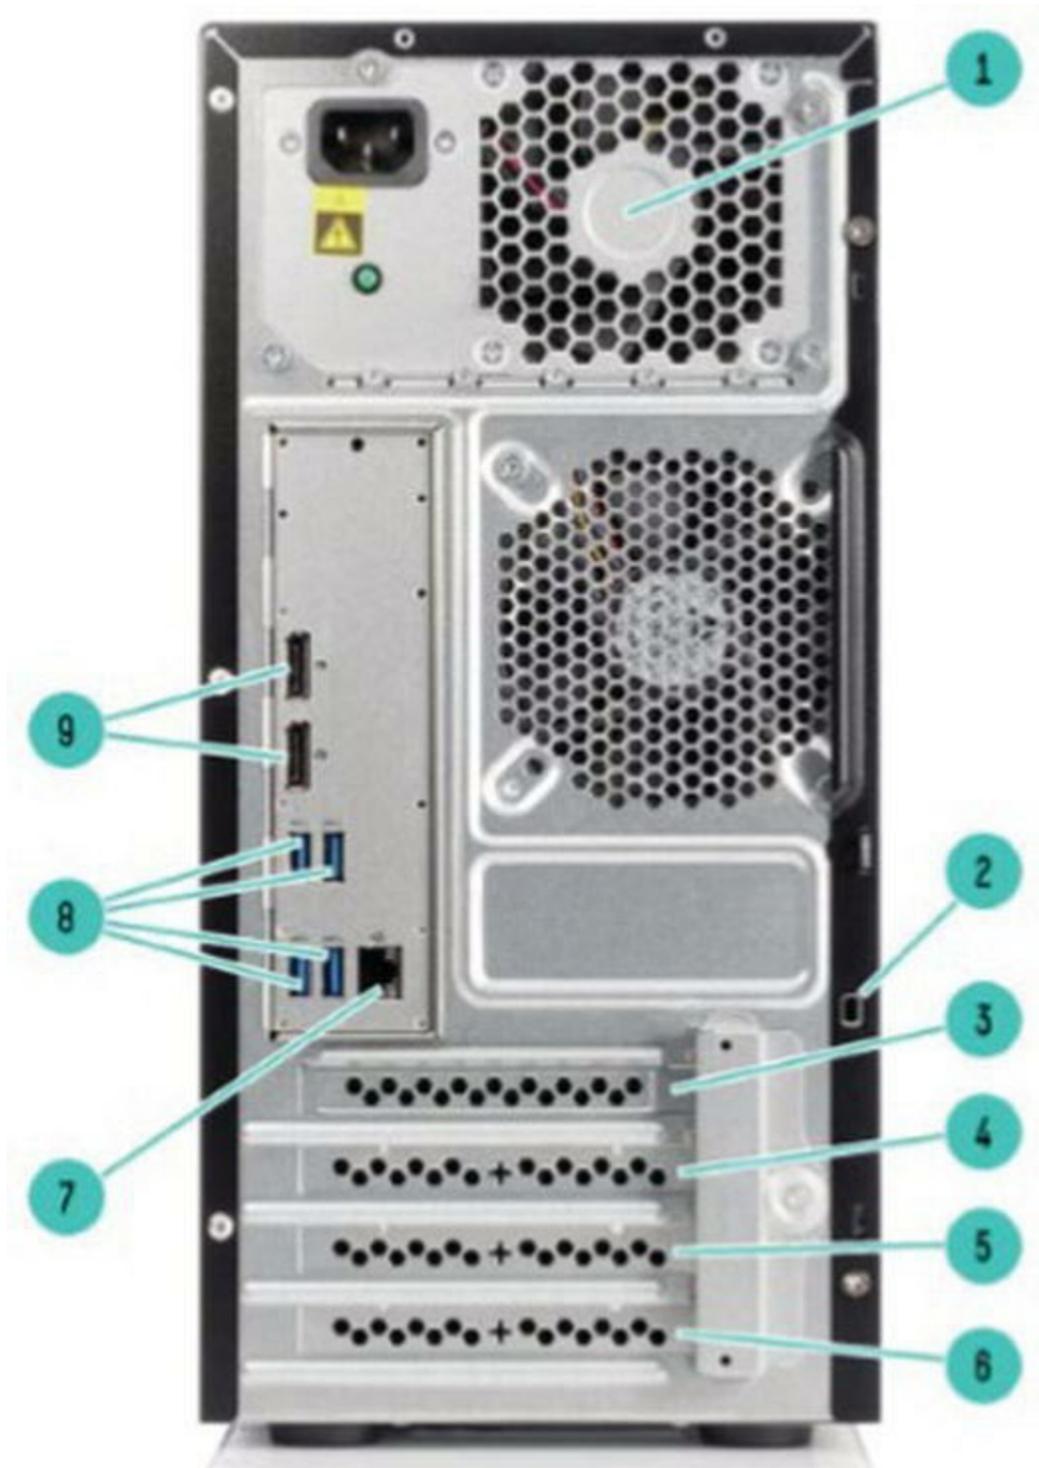


Figure 4-12 ProLiant ML10 Gen9 rear view

[Figure 4-12](#) shows the components of the ProLiant ML10 Gen9 from the rear view. The numbered components are:

Power supply

Kensington security slot

Slot 1, PCIe3 x8

Slot 2, PCIe3 x16

Slot 3, PCIe3 x4

Slot 4, PCIe3 x4

NIC port

USB 3.0 connector

Display ports

Additional enhancements over the previous generation of ML10 server include:

Intel® Ethernet Connection 1219-LM

HPE Standard Memory (4) DDR4, up to 2133 MHz (64 GB max)

ProLiant ML30 Gen9 Server



Figure 4-13 ProLiant ML30 Gen9 Server

The ProLiant ML30 Gen9 Server, shown in [Figure 4-13](#), is a full-featured, affordable, single-socket tower server that provides value performance and expansion for small businesses and remote branch offices. It is the ideal first server for growing businesses.

This server features HPE SmartDrives, with drive options including eight SFF and four LFF drive cages. These options support up to 8 TB per slot, which provides customers with flexibility for increased capacity of local storage.

The ProLiant ML30 Gen9 uses the HPE iLO Management Engine, which delivers a complete set of embedded features (standard on all ProLiant Gen9 Server products). These features include HPE Agentless Management, HPE Active Health System, HPE Intelligent Provisioning, and HPE Embedded Remote Support.

HPE Smart Array options provide RAID mirroring and striping capability to protect

critical data. The flash-backed write cache (FBWC) captures and holds data indefinitely in the event of a power loss, equipment failure, human error, or virus attacks.

The ProLiant ML30 Gen9 also features:

Quad-core Xeon[®] processor or dual-core Pentium Core i3 processor options

Dedicated iLO 4 port and error-correcting code (ECC) memory

2133 MHz DDR4 memory, with up to four DIMMs (64 GB maximum)

Four PCIe slots

Compatibility with 24-port HPE 1810 Series switches

1 GB 332i Ethernet adapter (two ports per controller)

ProLiant ML110 Gen9 Server

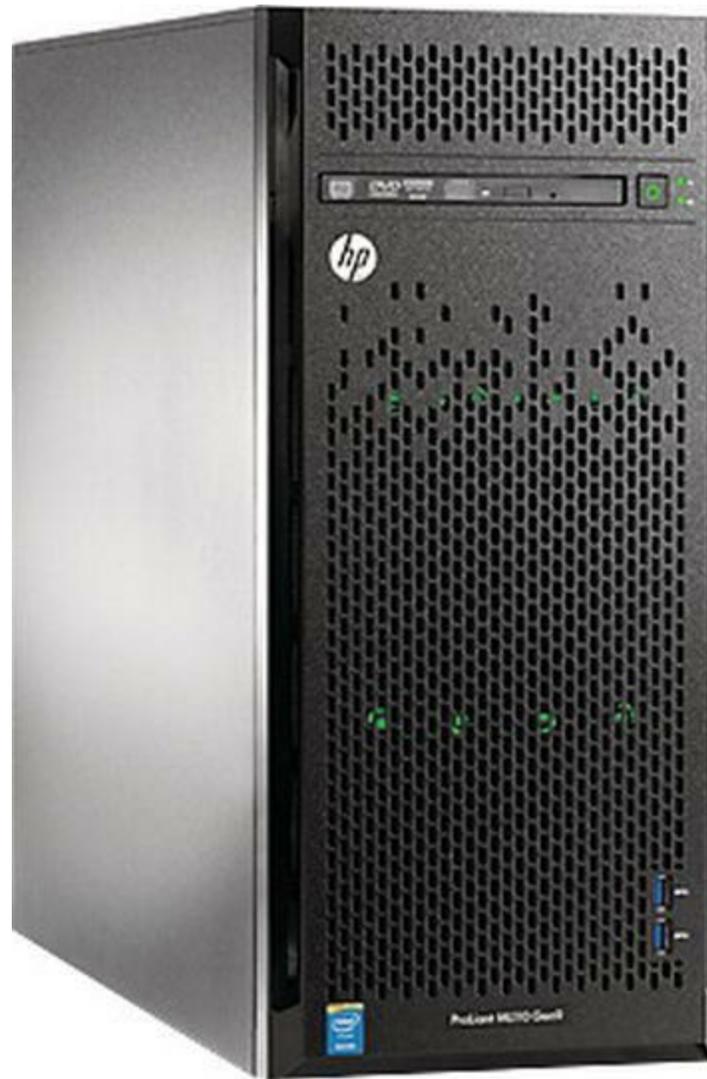


Figure 4-14 ProLiant ML110 Gen9 Server

The HPE ProLiant ML110 Gen9 Server, shown in [Figure 4-14](#), is a small, quiet tower with performance, expansion, and growth at an affordable price to meet most small- and mid-sized businesses' compute demands. The ML110 Gen9 Server is a single-processor tower form factor that can also be placed in a rack server cabinet using 4.5U of rack space. It includes up to five PCIe slots with room to grow with eight LFF or 16 SFF disk drives. Provides eight DDR4 DIMM slots supporting HPE DDR4 SmartMemory with a maximum capacity of 256 GB, making this the ideal server for small- to mid-sized businesses and remote offices/branch offices.

Features include:

Xeon[®] E5-2600 v4 and E5-1600 v3 series processors with up to 14 cores

Energy efficiency with ASHRAE A3 compliance

Embedded SATA HPE Dynamic Smart Array B140i controller for boot and data

Optional Smart Array controllers

Eight USB ports and a redundant power supply option

A single-wide GPGPU

HPE Insight Control

Embedded iLO 4 remote and out-of-band management

ProLiant ML150 Gen9 Server

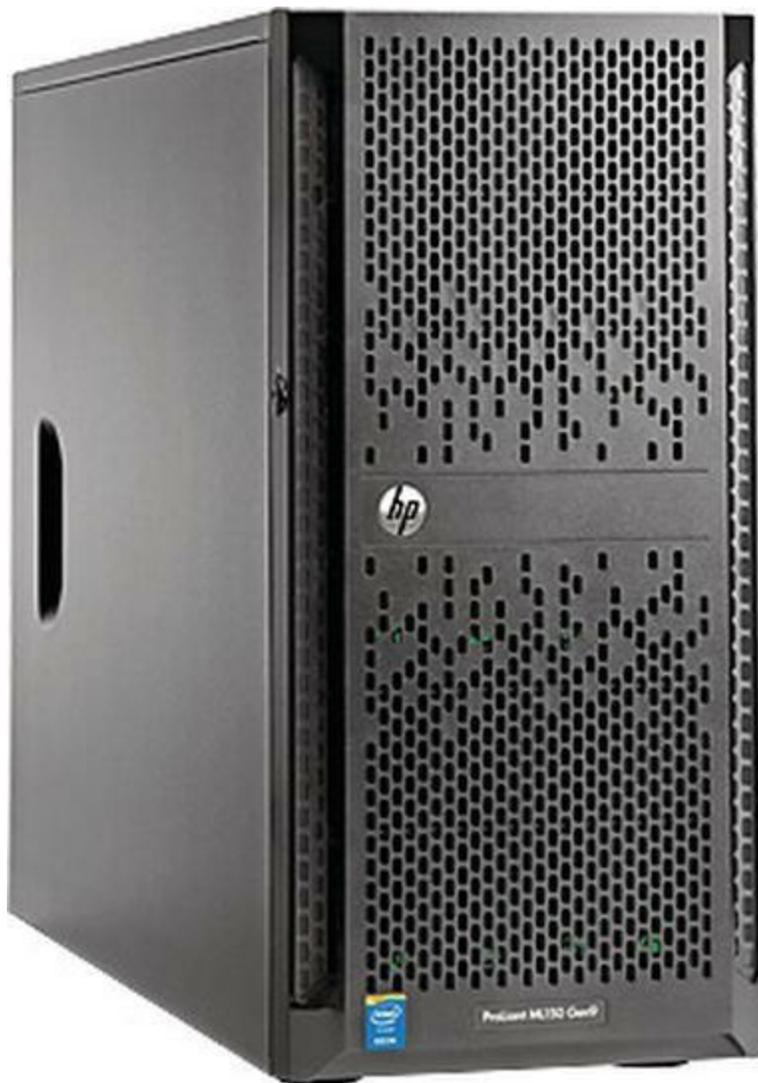


Figure 4-15 ProLiant ML150 Gen9 Server

The ProLiant ML150 Gen9, shown in [Figure 4-15](#), is a dual-socket server that provides configuration flexibility and expansion options to meet a wide range of different capacities, connectivity, and design requirements. This compact server offers up to 10 LFF disk drives or 16 SFF disk drive options and increased IO expansion. It covers a wide range of applications and workloads and addresses the needs of sophisticated SMBs and large enterprise-class server environments, with an essential mix of performance, efficiency, and expandability.

The HPE ProLiant ML150 Gen9 features:

Support for high-capacity hard disk drives (HDDs) (up to 8 TB LFF)

Dual MicroSD, providing a low-cost boot option

Up to two Xeon® processor E5-2600 v4 processors with up to 18 cores

Up to 16 DIMM slots of HPE DDR4 SmartMemory

A choice of:

- Smart Array controllers with 12 Gb/s serial-attached SCSI (SAS) technology
- HPE Smart HBAs with SAS connectivity capable of running HBA mode or simple RAID mode

Six PCIe expansion slots, eight USB ports, and power supply options

Energy efficiency with ASHRAE A3 compliance

Support for AMD FirePro W7100 graphics card

iLO 4 and UEFI boot mode

Table 4-4 ProLiant ML Gen9 workloads

Essential		Performance
• Entry	• Growth	• Traditional
• ML10 Gen9	• ML110 Gen9	• ML350 Gen9
• ML30 Gen9	• ML150 Gen9	•
Workloads		
File, print, and business apps Collaboration (email/workgroup) IT and web Virtualization Graphics—intensive medical imaging and animation rendering	Graphics and VDI apps for financial services, education, scientific research, and medical imaging IT infrastructure support, database, web, and collaboration	Mission-critical apps Graphics and VDI apps

[Table 4-4](#) compares use cases of the ProLiant ML Gen9 models.

ProLiant ML servers are designed to maximize computing power by enabling you to add components such as PCI expansion cards and disk drives inside the chassis. The

flexible, configurable design of ML servers enables their use in enterprise data centers, remote branch offices, and growing businesses. ML servers provide high-availability features to ensure maximum uptime, as well as price and performance flexibility for the customer.

Learning check

A customer approaches your stand at a trade fair and asks you to suggest the best server solution to provide a low cost, highly available server platform for a stand-alone virtual machine environment. How should you respond?

Activity: Matching ProLiant tower servers to customer needs

This activity requires you to apply ProLiant ML Gen10 Servers to customer scenarios. By completing this exercise, you will gain experience determining which servers fit various customer environments, requirements, workloads, and so on.

Using the information provided below, determine which server would best fit your customer needs.

The customer scenarios are:

Customer 1—This business provides cleaning services for homes and businesses. It manages approximately 10 facilities and has about 150 employees. The company is seeking a server solution with the following features:

- Ideal for remote and branch offices
- File, print, and business applications
- Virtualization workloads
- Automation and intelligent management features
- Flexibility to expand in the future

Customer 2—A midsized theatre, acoustics, and digital design consultancy is expanding rapidly. Business leaders are interested in balancing employees' mobility, security, and collaboration needs. The business is interested in a traditional server with the following features:

- Highly available
- Expansion options with large internal storage capacity
- CRM
- Access for remote, on-the-go employees
- Several USB ports
- Improved workload performance and power efficiency

Customer 3—A small accountancy business owns two offices. Each office employs a maximum of 10 employees. The owner is interested in opening a third location and looking for the company's first server solution. The company is seeking an IT solution with the following features:

- Small footprint

- Budget-friendly
- Support for collaboration (email/workgroup)
- Quiet operation

Explain the reasoning behind your server selection for each customer in the space provided.

Activity debrief

Did you discover more than one possible solution? Why or why not?

What additional requirements are important in the server selection process? What else would you ask a potential customer before making a recommendation?

Learning check

Write a summary of the key concepts presented in this chapter.

Activity: Matching ProLiant tower servers to customer needs

While there are no “right” answers when it comes to which server to recommend for a particular set of customer needs, it can be useful to choose a server to start the conversation and build from there. With this in mind, suggested starting points for the customers in the scenario are:

Customer 1—ProLiant ML30 Gen9. This server is a full-featured, affordable, single-socket tower server that provides value performance, expansion, and growth for small businesses and remote branch offices.

Customer 2—ProLiant ML350 Gen10. This is a two-processor (2P) server providing high performance, expandability, and manageability—ideal for expanding SMBs and enterprise businesses.

Customer 3—ProLiant ML10 Gen9. This provides server-level performance in a quiet, compact size that is easy to deploy in small business environments.

If you chose a different server to start the conversation with, that is OK. If you did choose a different server, why did you choose the one you chose?

You should ask your customer some questions to help you refine your recommendation. While not an exhaustive list, questions might include:

What are your future plans and business goals?

What is the projected role of the server?

Do you have a preference for operating system?

What applications are you running?

To which kind of network will the server be attached?

How much are you thinking of investing in the solution?

How much storage space do you need, and is this likely to grow in the near future?

What is the required availability of the server?

Is server price more important than functionality?

Is a rack or tower configuration preferred?

Will backups be performed?

Is power protection needed?

Which kinds of system management tools are needed?

What level of maintenance and support is desired?

Is there a long-term IT strategy in place?

What is the biggest IT problem facing the business today? What does the customer believe are possible solutions? What are the barriers to the solution?

Summary

HPE ProLiant ML servers are relatively compact, simple, and robust. They are integrated with a simplified yet comprehensive management suite and industry-leading support. The ProLiant tower server portfolio:

Delivers business value

Helps increase IT staff productivity

Expedites service delivery

Includes Silicon Root of Trust and Intelligent System Tuning

Offers complete financing options, service offerings, and a channel network

Provides simple management and storage tools, along with proven configurations

Maximizes computing power by enabling you to add components such as PCI expansion cards and disk drives inside the chassis

Includes high-availability features to ensure maximum uptime

Offers price and performance flexibility

The ProLiant ML Gen10 Server family includes:

ProLiant ML350 Gen10

ProLiant ML110 Gen10

ProLiant MicroServer Gen10

The ProLiant ML Gen9 Server family includes:

ProLiant ML150 Gen9

ProLiant ML110 Gen9

ProLiant ML30 Gen9

ProLiant ML10 Gen9

5 HPE Rack Server Solutions

WHAT IS IN THIS CHAPTER FOR YOU?

After completing this chapter, you should be able to:

Describe, at a high level, the features and market position of the HPE ProLiant DL rack server lines.

Describe the HPE SimpliVity 380 Hyperconverged solution.

Describe, at a high level, the features and functions of HPE Intelligent Series Racks.

Prelearning check

You are attending a meeting with a customer to discuss their future data center modernization plans. During the discussion they raise the subject of hyperconverged solutions and that they are considering Microsoft Storage Spaces Direct. When they considered the HPE DL380 Rack Server solution, they were concerned that the platform may not offer sufficient HDD and SSD storage to meet their needs. How should you respond?

You are at a customer meeting and one of the IT team comments that there is a belief that hyperconverged solutions performance is poor due to the CPU bound deduplication and compression functionality. They are considering the HPE SimpliVity 380 solution. How would you address their performance concerns?

HPE ProLiant rack servers



Figure 5-1 ProLiant rack server lines

ProLiant servers are based on industry-standard x86 architectures that enjoy industry-wide application support. The HPE commitment to participate in joint development activities with processor manufacturers and HPE solution development, service, and support ensures that ProLiant servers provide a familiar and trusted platform.

HPE ProLiant Gen10 servers

HPE ProLiant servers are separated into five main product-lines.

DL servers—Density-optimized for stacking in rack-mounted server environments ideal for multi-server deployments

ML servers—Tower- and rack-based servers designed with capacity for maximum internal expansion

BL servers—Comprise blade servers that fit within the HPE BladeSystem

SY servers—Compute modules that fit within the HPE Synergy solution

XL servers—Compute modules that fit within the HPE Apollo solution

ProLiant servers are also split into several series that denote processor configuration.

The 10, 300, and 400 series comprise single- and dual-socket systems.

The 500 and 600 series comprise quad-socket-capable systems.

Models where the last digit is zero are built with Intel® Xeon® processors inside (such as DL380), and those with five at the end are built to support AMD Opteron processors (such as DL385).

The ProLiant DL server families offer rack-mounted servers based on Intel® and AMD® processors. ProLiant DL servers, some of which are shown in [Figure 5-1](#), are:

Density-optimized for flexibility and manageability in rack-mounted server environments

Ideal for multi-server deployments

Available in rack and cluster models

The ProLiant DL Gen10 family of rack servers offers a balance of efficiency, performance, and management. ProLiant Gen10 rack servers incorporate embedded components that provide a rich feature set in a compact chassis size. These servers are ideal in dense, space-constrained, and multi-server environments. They are best for customers who want a small footprint, but with enterprise server features. With ProLiant Gen10 rack servers, you can significantly increase the speed of IT operations and enable IT to respond rapidly to any business need.

Compared to the previous generation ProLiant Gen9 servers, these Gen10 servers feature increased processor core count, memory, and internal storage capacities, as well as the next generation of embedded Smart Array technology. In addition, they are performance optimized for multi-application workloads. ProLiant Gen10 server highlights over previous-generation servers include:

Enhanced manageability and security with the HPE integrated Lights-Out (iLO 5) Management Engine

Up to twice the cache capacity with 4 GB flash back write cache (FBWC) options for indefinite data retention

FlexibleLOM available on ProLiant Gen10 server blades and select ProLiant Gen10 rack-mounted servers

Easier accessibility with enhanced cable routing

Improved serviceability with clear and numbered air baffle designs

ProLiant Gen10 servers offer FlexibleLOMs, an optional flexible network technology that offers customers a choice of 1 Gb, 10 Gb, or 25 Gb base-T Ethernet or converged networking in their embedded adapter.

LAN-on-motherboard (LOM) technology provides essential network connectivity without requiring an optional network card to be installed in an expansion slot. The HPE FlexibleLOM module attaches to a dedicated edge connector on the system board. FlexibleLOM technology allows you to select the connectivity you need and adapt to network changes without using a standard PCIe slot.

ProLiant Gen10 rack server portfolio

		DL10 series	DL300 series	DL500 series
	Security	✓✓✓	✓✓✓	✓✓✓
Improved security across server lifecycle, extensive standards compliance				
	Productivity		✓✓✓	✓✓✓
Persistent storage at memory speeds				
	Velocity	✓	✓✓✓	✓✓✓
Intel® performance, expanded NVMe storage				
	Efficiency	✓	✓✓✓	✓✓✓
Simplified option portfolio, economics		Intelligent System Tuning		
	Flexibility	✓	✓✓✓	✓✓✓
Increased storage capability		Increase storage, compute density, GPUs		
	Simplicity	✓✓✓	✓✓✓	✓✓✓
Improved GUI and Industry Standard APIs, faster problem resolution				

Figure 5-2 ProLiant DL server positioning

ProLiant servers, ranging from one to four socket, offer ideal solutions for front-end, infrastructure, and network-edge workloads, where a single or a few applications make up the main workload for the server. For Windows and Linux environments, the ProLiant server family is the ideal platform for building a converged infrastructure.

With an intelligent and efficient data center based on an HPE Converged Infrastructure strategy, an enterprise can build a strong foundation and protect infrastructure investments. These customer-inspired innovations deliver real business results with a five-month return on investment (ROI).

ProLiant DL rack servers are available in three series—DL10, DL300, and DL500. Although all series are designed to handle multiple workloads—IT infrastructure, web, business applications, collaboration, analytics, Big Data, and more—each family is optimized for specific use cases, as shown in [Figure 5-2](#).



Note

For more information on ProLiant DL rack servers, enter the URL into your browser.

<https://www.hpe.com/uk/en/servers/proliant-dl-servers.html>

ProLiant server product numbers

HPE positions ProLiant servers by two criteria—line and series—to help customers choose the server that best fits their requirements. The numbering of the ProLiant rack server line reflects the following features:

DL500 series Gen10

- Scale from four to 112 cores with up to four Intel[®] Xeon[®] Processor Scalable Family
- 6 TB maximum memory
- Fast memory with (48) HPE Smart Memory DDR4 2666MHz (6.0 TB)
- High-storage density 24 SFF max, HDD/SSD, M.2 enablement kit, and 12 NVMe PCIe SSD option
- Flexible drive bay supporting multiple combinations of NVMe and SAS SFF drives
- Enhanced RAS features for high availability
- Dedicated iLO 5 port
- Support up to two graphical processing units (GPU), eight available IO slots + FlexibleLOM for analytics and HPC applications
- Up to four, 96% efficient Flex Slot power supplies with redundant configurations
- Dynamically tune server performance using with Intelligent System Tuning (IST)

DL300 series Gen10

- Up to two Intel[®] Xeon[®] Processor Scalable family, PCIe 3.0, up to eight available slots + m.2 support on riser
- Four to 56 processor cores
- HPE Smart Memory (24) DDR4 (six channels per CPU), up to 2666 MT/s (3.0 TB max)
- Standard HPE Dynamic Smart Array S100i Choice of HPE Flexible Smart Array or

HPE Smart HBA Controllers for performance or additional features

- 24+6 SFF/12+4+3 LFF + 2SFF max, HDD/SSD, 20 NVMe PCIe SSD, and m.2 enablement
- 4 × 1 GbE embedded + Choice of FlexibleLOM + Standup cards
- Up to two, 96% efficient Flex Slot power supplies with redundant configurations
- Dynamically tune server performance using with IST
- ASHRAE A3 and A4, lower idle power, Energy Star
- Dedicated iLO 5 port

DL20/80 series Gen9

- One or two Xeon® E3-1200 v5 or E5-2600 v4 product family processors (Intel® Core i3 and Intel® Pentium available on the DL20 Gen9)



Note

The ProLiant DL20 server supports E3-1200 processors; ProLiant DL80 servers support E5-2600 processors.

Two to 14 processor cores

64 GB or 256 GB maximum memory

Four or eight DIMM slots

Two large form factor (LFF), four SFF, or 12 LFF/eight SFF hard drives

3/1/1 warranty

HPE ProLiant DL360 Gen10 Server



Figure 5-3 ProLiant DL360 Gen10 Server

The HPE ProLiant DL360 Gen10 Server, shown in [Figure 5-3](#), is a dual-socket 1U rack

server that delivers enhanced price/performance, enterprise-class management, and scalable capabilities to meet the needs of general workloads today and tomorrow. It is ideal for small- to medium-sized businesses (SMBs) and value-focused enterprise customers requiring essential features in an optimized design. The HPE ProLiant DL360 Gen10 Server supports the Intel® Xeon® Processor Scalable Family with up to 28 cores per processor, plus 2666 MT/s HPE DDR4 SmartMemory supporting up to 3.0 TB max, with the added performance that 12 NVDIMMs and 10 NVMe brings. Deploy this dense platform for diverse workloads in space-constrained environments and maintain it with ease by automating the most essential server lifecycle management tasks with HPE OneView and HPE iLO 5.

Each processor socket supports 12 DIMM sockets. Maximum memory capacity is limited by processor selection and DIMM type. Mixing of RDIMM and LRDIMM memory is not supported. The maximum memory capacity is 3 TB per processor assuming 24 × 128 GB 2666 MT/s LRDIMM, 768 GB when populated with 24 × 32 GB @ 2666 MT/s RDIMM memory, and 192 GB 12 × 16 GB NVDIMM @ 2666 MT/s.

Key selling points of the ProLiant DL360 Gen10 include:

Increased performance with IST, Persistent Memory, and large NVMe capability

Increased storage capacities and options to support large storage workloads

PCIe expansion with HPE FlexibleLOM and HPE Flexible Smart Array

The ProLiant DL360 Gen10 Server may be ideal for:

Dynamic workloads in dense virtualized environments

Compute-intensive applications (web caching, data analytics)

Low latency and transactional applications (warehouse and database)

High-Performance Computing focuses on the Financial Services Industry (FSI)

Other key advantages include:

HPE Performance with IST and Persistent memory

Industry-leading security

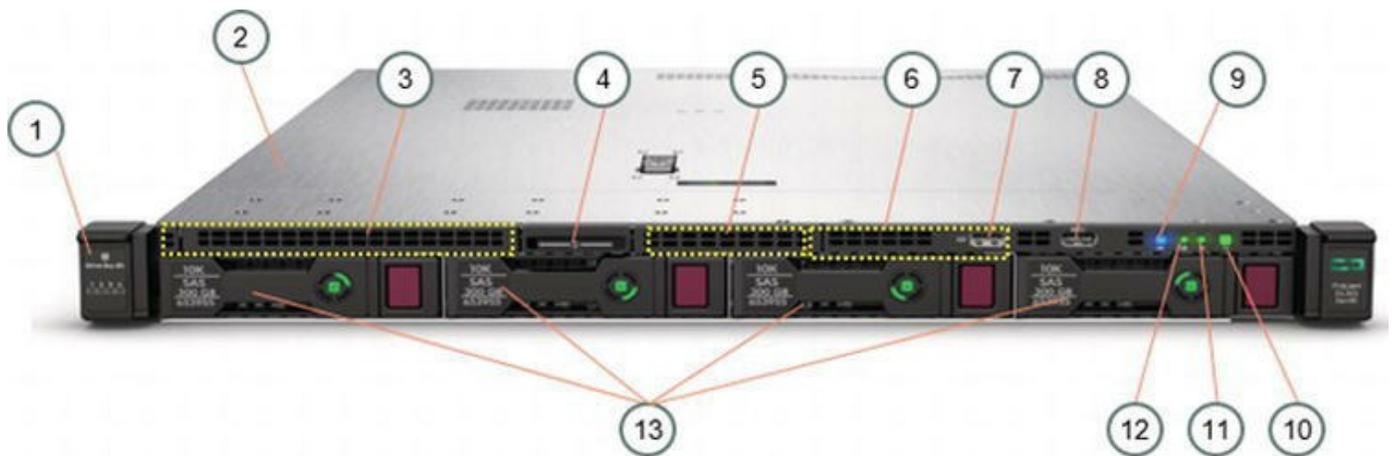
New Premium 10 NVMe chassis

Flexible chassis supporting future upgrades

New rear drive option (single SFF or Dual uFF)

Performance gains with new HPE Smart Array and HPE Smart Memory DDR4 2666 MT/s

HPE ProLiant DL360 Gen10 Server: Four LFF front of system detail



Options:

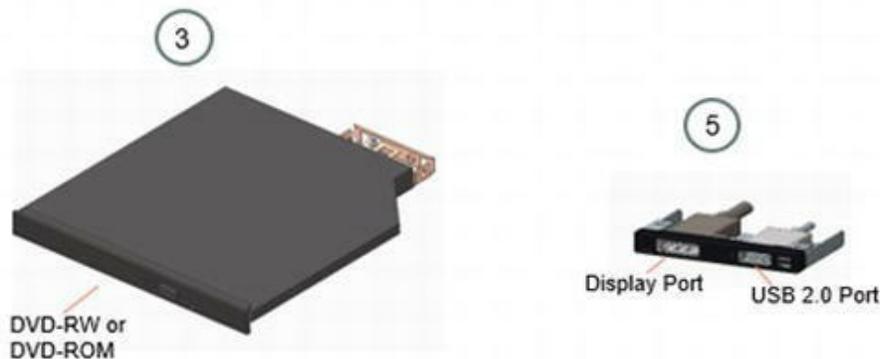


Figure 5-4 ProLiant DL360 Gen10 Server 4 LFF, front view

Key features of the ProLiant DL360 Gen10 as shown in [Figure 5-4](#) are:

Front view:

- | | |
|--|--|
| 1. Drive support label | 7. iLO Service Port |
| 2. Quick removal access panel | 8. USB 3.0 Port |
| 3. Optional: DVD-RW or DVD-ROM (blank shown) | 9. UID button/LED |
| 4. Serial number label pull tab | 10. Power On/Standby button and system power LED |
| 5. Optional: Display port and USB 2.0 port kit (blank shown) | 11. Health LED |
| 6. Optional: System Insight Display (SID)*—std. shown | 12. NIC Status LED |

*This option will lose #7 iLO Service Port

ProLiant DL360 Gen10: Eight SFF front of system detail



Figure 5-5 ProLiant DL360 Gen10 Server eight LFF, front view

Key features of the ProLiant DL360 Gen10 as shown in [Figure 5-5](#) are:

Front view:

- | | |
|---|---|
| Drive support label | 5. UID button/LED |
| Quick removal access panel | 6. Power On/Standby button and system power LED |
| Serial number label pull tab | 7. Health LED |
| Universal media bay options: | 8. NIC status LED |
| a. Option shown: +2 SFF SAS/SATA (total max 10 SFF) | 9. USB 3.0 port |
| b. Optional: +2 SFF NVMe drives | 10. iLO Service Port |
| c. Optional: DVD-RW or DVD-ROM | 11. Eight SAS/SATA/SSD drive bays |
| d. Optional: +2 Dual uFF (4× M.2 cartridges) | |
| e. Optional: Display port + USB 3.0 Port | |

HPE ProLiant DL360 Gen10: New premium 10 SFF NVMe

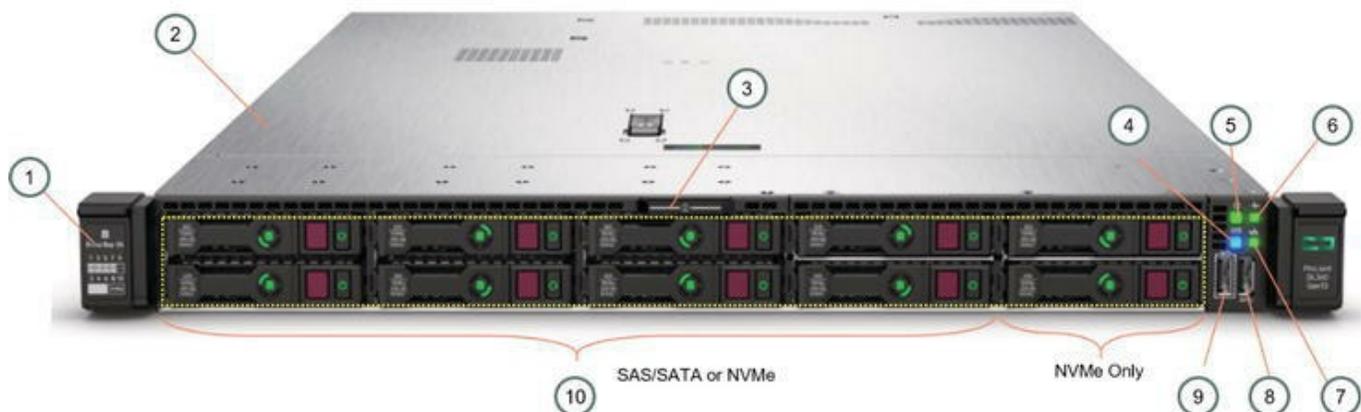


Figure 5-6 ProLiant DL360 Gen10 Server 10 SFF NVMe, front view

Key features of the ProLiant DL360 Gen10 with premium 10 SFF NVMe drives as shown in [Figure 5-6](#) are:

Front view:

- | | |
|---|---|
| 1. Drive support label | 7. NIC status LED |
| 2. Quick removal access panel | 8. USB 3.0 port |
| 3. Serial no. label pull tab | 9. iLO Service Port |
| 4. UID button/LED | 10. Max 10 NVMe drive bays (PCIe direct attached) |
| 5. Power On/Standby button and system power LED | *Mix-n-match SAS/SATA or NVMe: |
| 6. Health LED | -Drive bays 1-8 support SAS/SATA and NVMe |
| | -Drive bays 9 and 10 are NVMe ONLY |

ProLiant DL360 Gen10: Interior system detail—SFF chassis

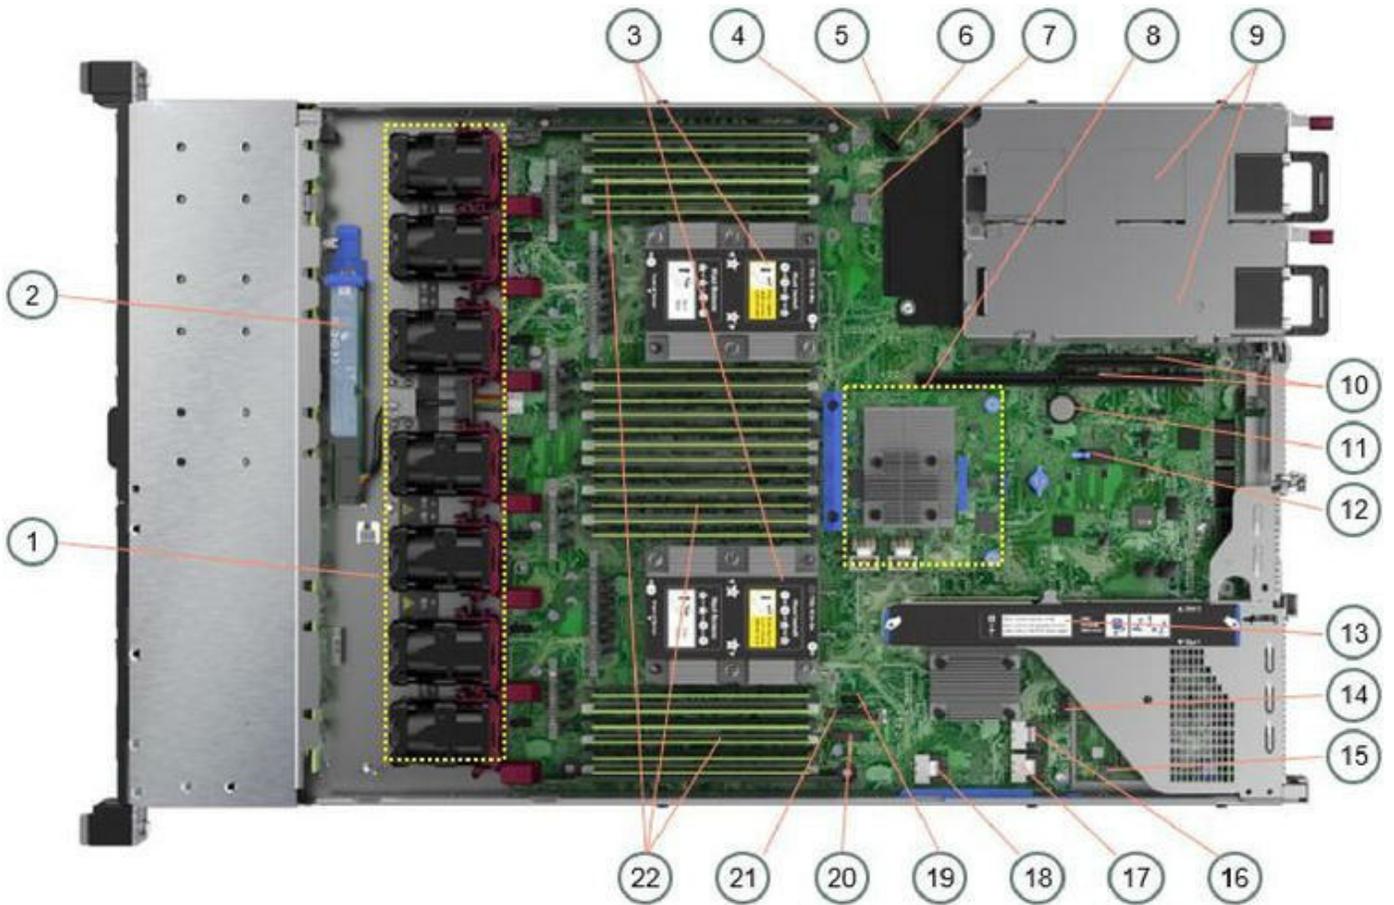


Figure 5-7 ProLiant DL360 Gen10 Server interior view

Key internal features of the ProLiant DL360 Gen10 as shown in [Figure 5-7](#) are:

Internal view:

Standard single rotor hot-plug fans*

- One CPU—five standard fans
- Two CPUs—seven standard fans

Optional: High-performance fans

*For eight SFF and four LFF chassis only, 10 NVMe chassis ships with seven high-performance fans as standard

Optional: Smart Storage Battery

Up to two processors (shown with standard heat sinks)

MicroSD card slot (Dual Micro-SD option available)

Optional: Chassis Intrusion Detection

Hard Drive backplane power connector

Dual internal USB 3.0 connector

Smart Array Controller (Type -a shown)

Up to two Power Supplies for redundant power

Secondary (CPU2) PCIe 3.0 riser*

*Not available on Premium 10 SFF NVMe chassis due to dedicated 10 × 4 NVMe riser for PCIe direct attached

Optional: Low Profile x16 **or**

Optional: Full Height x16

11. System Battery

12. Optional: TPM 2.0

13. Primary (CPU1) PCIe 3.0 riser

Standard: GPU power connector + 1 x16 and 1 x8

Optional: Two SATA M.2 + 2 x16

Optional: 2 x4 NVMe + 1 x16 and 1 x8*

*Only available on eight SFF chassis

14. Optional: Front Display port/USB 2.0

15. FlexibleLOM (supports various NICs)

16. x4 SATA port 1

17. x4 SATA port 2

18. x2 SATA port 3

19. x1 SATA port 4

20. Optical/SATA port 5

21. Front Power USB 3.0 connector

22. DDR4 DIMM slots—Fully populated 24 DIMMs shown

ProLiant DL360 Gen10: Rear view of system detail

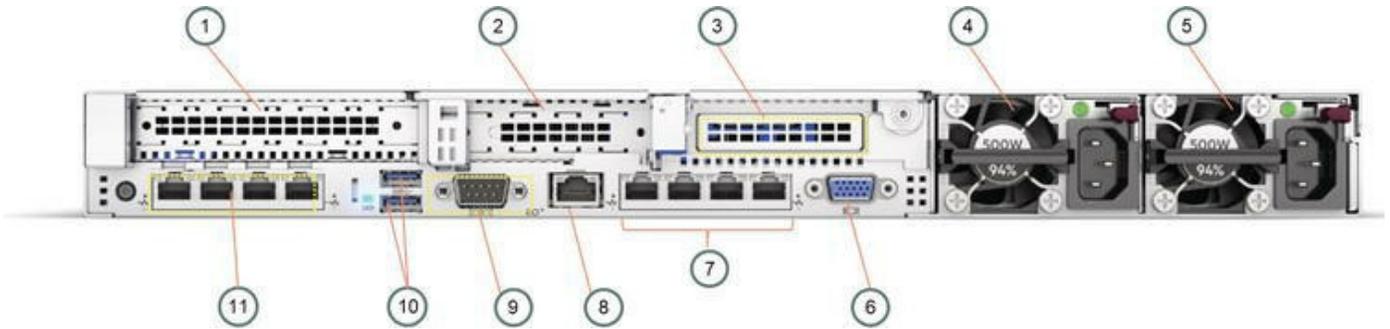


Figure 5-8 ProLiant DL360 Gen10 Server, rear system detail

Key features of the ProLiant DL360 Gen10 as shown in [Figure 5-8](#) are:

Rear view:

Slot 1 PCIe 3.0

Optional: Rear Drive + one SFF or one uFF (2x M.2 cartridges)

Slot 2 PCIe 3.0

Optional: Slot 3 PCIe 3.0 (Requires 2nd processor)

Power Supply 2

Power Supply 1

6. VGA port

7. Embedded 4 x 1 GbE Adapter

8. iLO Management Port

9. Optional: Serial Port

10. USB 3.0 Ports

11. Optional: FlexibleLOM* (shown 4 x 1 GbE)

*Supports various FlexibleLOM NICs up to 25 GbE

Learning check

A customer has expressed a need for a secure two-socket ProLiant rack server to host one of their low-density virtualized applications. Which server should you recommend?

- a. DL380 Gen10
- b. DL180 Gen9
- c. DL560 Gen10
- d. DL360 Gen10

HPE ProLiant DL380 Gen10 Server



Figure 5-9 ProLiant DL380 Gen10

The HPE ProLiant DL380 Gen10 Server, shown in [Figure 5-9](#), is a dual-socket 2U rack server that delivers enhanced performance, enterprise-class management, expandability, and scalable capabilities to meet the needs of general and heavy workloads. Designed for supreme versatility and resiliency and adaptable for diverse workloads making it ideal for multiple environments from Containers to Cloud to Big Data.

The HPE ProLiant DL360 Gen10 Server supports the Intel[®] Xeon[®] Processor Scalable Family with up to 28 cores, plus 2666 MT/s HPE DDR4 SmartMemory supporting up to 3.0 TB maximum, with the added performance that 12 NVDIMMs and 16 NVMe brings.

Key selling points and advantages include:

Increased performance with IST, Persistent Memory, and greater NVMe capability

“Future proof” design keeps up with your business needs

Huge storage footprint for large storage workloads

PCIe expansion with HPE FlexLOM and HPE Flexible Smart Array

Addresses 16 of IDC’s top 18 workloads in a single chassis

HPE Performance with IST and Persistent memory

Industry-leading security

Flexible chassis supporting future upgrades

Performance gains with new HPE Smart Array and HPE Smart Memory DDR4 2666 MHz

Greater NVMe (20) PCIe Storage capacity

Broad GPU support

Industry-leading security

The ProLiant DL380 Gen10 Server may be ideal for:

Enterprise customers with workloads including virtualization, containers, Big Data, storage-centric apps, data warehousing/analytics, CRM, ERP, VDI, SAP, and large storage capacity, such as Microsoft Exchange, and so forth

Virtualization, productivity workhorse with the ability to scale up or down to any environment. Single solution to deploy in multiple environments

HPE ProLiant DL380 Gen10—Front system detail

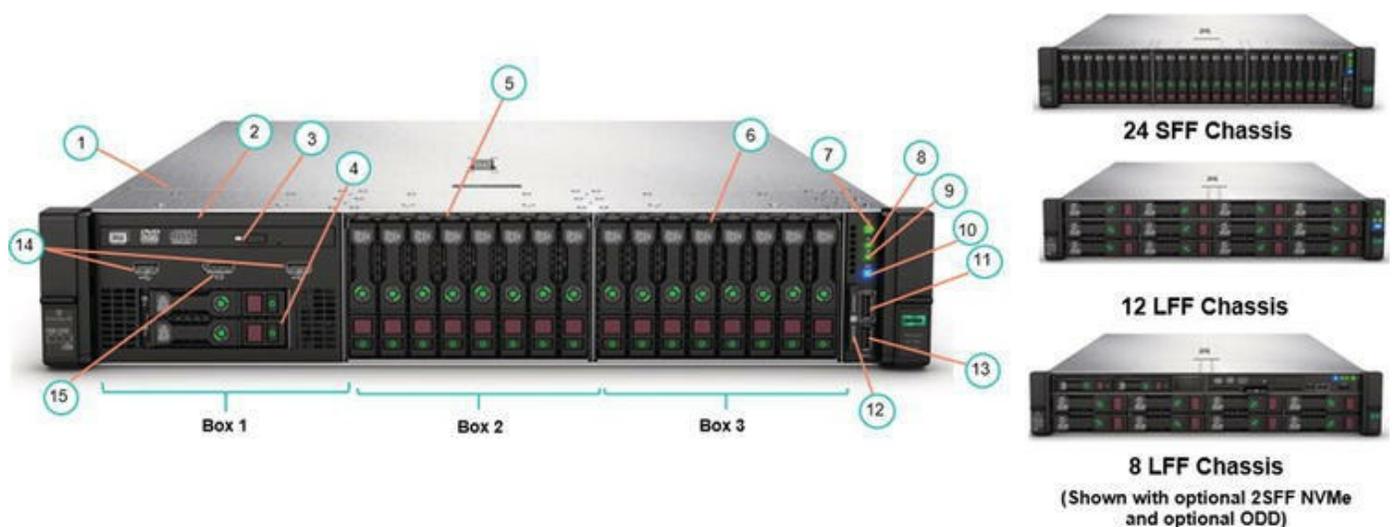


Figure 5-10 ProLiant DL380 Gen10 Server, front view and chassis

Key features of the ProLiant DL380 Gen10 as shown in [Figure 5-10](#) are:

Front view:

- | | |
|---|---|
| Quick removal access panel | 7. Power On/Standby button and system power LED button |
| Optional Universal Media bay. | 8. Health LED |
| Two USB 2.0 and Display port (eight SFF bay or six SFF + two NVMe or eight NVMe optional) | 9. NIC status |
| Optional Optical drive. Requires Universal Media bay | 10. UID button |
| Optional two SFF HDD, requires optional Universal Media bay | 11. iLO Front Service Port |
| Drive Bay 2. NVMe shown (eight SFF, six SFF + two NVMe or eight NVMe PCIe SSD optional) | 12. Serial label pull tag |
| Eight SFF Drive Cage Bay | 13. USB 3.0 |
| | 14. Optional USB 2.0 (via Universal Media Bay) |
| | 15. Optional front display port (Via Universal Media Bay) |

HPE ProLiant DL380 Gen10: Interior of system detail (eight SFF)

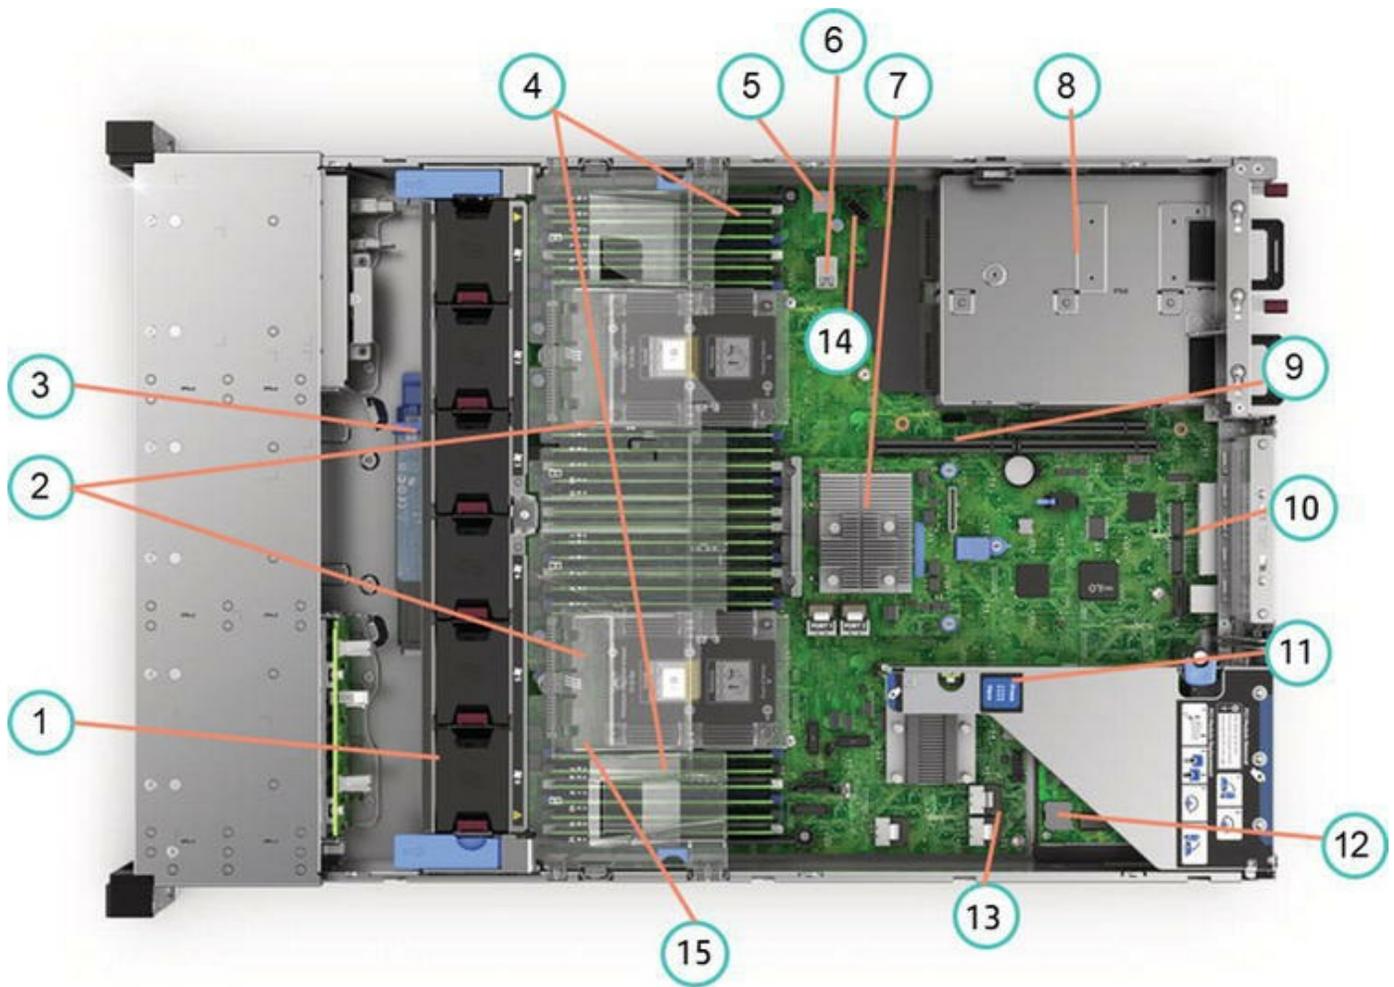


Figure 5-11 ProLiant DL380 Gen10, interior layout

Key features of the ProLiant DL380 Gen10 as shown in [Figure 5-11](#) are:

Internal view:

Fan cage shown with six standard hot-plug fans (High Ambient temperature fans optional)

Two processors, heatsink showing

Optional HPE Smart Storage Battery

DDR4 DIMM slots. Shown fully populated in 24 slots (12 per processor)

MicroSD card slot (Optional Dual Micro-SD option)

Internal USB 3.0 connector

Optional HPE Flexible Smart Array or Smart HBA (P408i-a shown)

8. (Under) Hot Plug redundant HPE Flexible Slot Power supplies

9. Connection for second (optional) riser (requires second CPU)

10. Embedded 4 × 1 GbE NIC

11. Primary PCIe riser, standard (optional double-wide GPU riser)

12. Optional Flexible LOM slot

13. X4 SATA ports (1 and 2).

14. Chassis intrusion detection connector

15. Clear air baffle

HPE ProLiant DL380 Gen10: Rear system detail

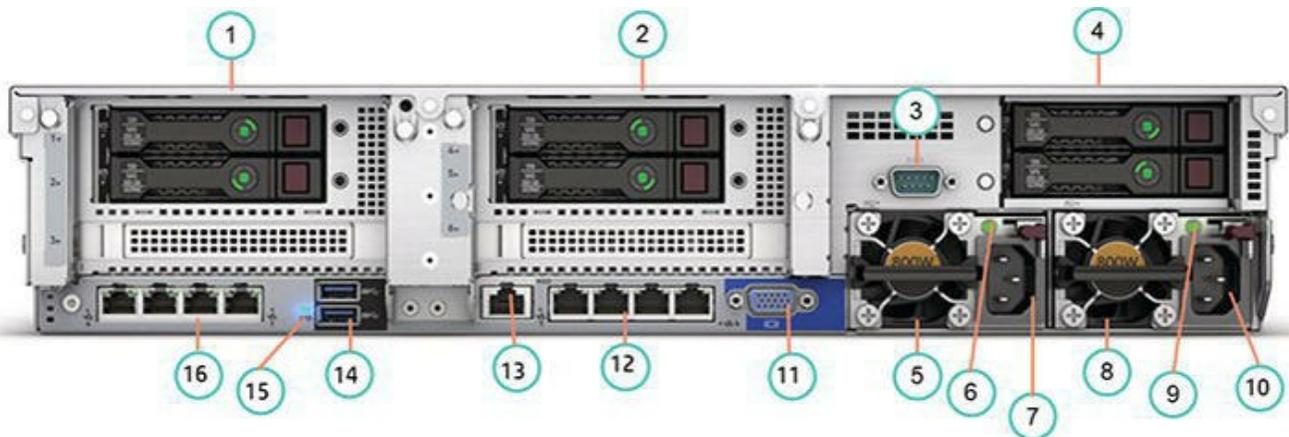


Figure 5-12 ProLiant DL380 Gen10, rear view

Key features of the ProLiant DL380 Gen10 as shown in [Figure 5-12](#) are:

Rear view:

Primary riser. PCI Slots (Slots 1 to 3 top to bottom, riser shipped standard, not shown), optional two SFF rear drives

Secondary riser. PCI Slots (Slots 4 to 6 top to bottom, not shown, requires second riser card, and second processor). Showing optional two SFF rear

Optional serial port

Tertiary Riser (Slots 7 to 8). Optional rear two SFF HDD (supported in 24 SFF or 12 LFF front end)

HPE Flexible Slot Power Supply bay 1 (800W shown)

Power supply Power LED

7. Power supply Power connection
8. HPE Flexible Slot Power Supply bay 2 (800W shown)
9. Power supply Power LED
10. Power supply Power connection
11. VGA connector
12. Embedded 4 × 1 GbE Network Adapter
13. Dedicated iLO management port
14. USB connectors 3.0 (two)
15. Unit ID LED
16. Optional FlexibleLOM ports (shown: 4 × 1 GbE)

HPE ProLiant DL380 Gen10 options

The data center standard for general-purpose compute, the HPE ProLiant DL380 Gen10 Server, delivers the best performance and expandability in the HPE 2P rack portfolio.

Adaptable for diverse workloads and environments, the secure 2P 2U HPE ProLiant DL380 Gen10 delivers world-class performance with the right balance of expandability and scalability. Designed for supreme versatility and resiliency while being backed by a comprehensive warranty make it ideal for multiple environments from containers to cloud to Big Data. Standardized on the industry's most trusted compute platform, making it an ideal server for many environments.

Other key features of the ProLiant DL380 Gen10 in addition to those shown in the preceding graphic are:

Greater chassis flexibility with up to 20 NVMe drives supported

Four LFF mid-tray bringing total LFF storage capacity to over 190 TB

HPE Persistent Memory at over 1 TB scale

Expanded GPU support to 3 × double-wide (DW) or 5 × single-wide (SW) cards

Additional boot/drive/rear options: SATA M.2; dual uFF SSD (2 × M.2 cartridges)

Intel® Xeon® Processor Scalable Family from four to 28 cores; 85 to 205W; 1.8 to 3.6 GHz

HPE DDR4 SmartMemory up to 2666 MT/s

Security features: iLO 5 (Silicon Root of Trust); Chassis Intrusion Detection; TPM 2.0; digitally signed FW

There are several drive bay options:

One to three SFF drive bay supports eight SFF or NVMe drives for a total of 24 with additional six SFF rear drive bay option to total 30 SFF drives

One to two SFF drive bays with the optional Universal Media bay including space for two additional SFF devices for a total of 18 with additional six SFF rear drive bay option to total 24 SFF drives

Or the 12 LFF with optional four LFF mid-plane and optional three LFF + two SFF rear drive bay to total 19 LFF drives + two SFF drives

All models come with the S100i Smart Array Controller with embedded software RAID support for 12 drives. The S100i uses 14 embedded SATA ports, but only 12 are accessible as two are reserved to support the two M.2 options on the primary riser.

Single processor models typically ship with four standard fans. The second processor option kit contains two additional fans. The eight LFF chassis ships with six standard fans. Six high-performance fans are shipped for all other chassis combinations and are also required for the rear drives, graphics (GPU) card, or NVMe configurations.

Each processor socket supports 12 DIMM sockets. Maximum memory capacity is limited by processor selection and DIMM type. Mixing of RDIMM and LRDIMM memory is not supported. Maximum memory capacity with 24 × 128 GB LRDIMMs is 3.0 TB, and 768 GB when populated with 24 × 32 GB RDIMM memory.

The storage controller on the DL380 Gen10 Server is one of the following depending on the model:

Entry model—HPE Dynamic Smart Array S100i SR SW RAID Controller, operating in UEFI mode only.

Base model—HPE Smart Array P816i-a 16-port and P408i-a 8-port SR Gen10 Controllers are available with Smart Storage battery included. This controller is in the Adaptive RAID on chip slot, which enables SAS drives and supports FBWC for data retention.

Performance and high-performance models—HPE Smart Array P408i-a SR Gen10 Controller including Smart Storage battery. This controller enables SAS drives and supports FBWC for enhanced data retention.

ProLiant DL380 Gen10 Servers offer customers a choice of HPE 1 Gb Ethernet 4-Port 331i Adapter plus optional HPE FlexibleLOM or stand up card, available in Entry and Base models. In the Performance models, the HPE 1Gb Ethernet 4-Port 331i Adapter plus HPE Ethernet 10/25 Gb 2-port 640FLR-SFP28 Adapter or the 2-port 631FLR-SFP28 adapters are available. A range of NIC cards is also available to enhance networking capabilities.

Flex slot 500W, 800W, and 1600W power supplies are available in Platinum Low Halogen models. In addition, there are 800W Universal, Titanium, and -48VDC versions.

Computational and Graphic Accelerators include:

HPE NVIDIA Quadro P2000, P4000, and P6000 GPU Module versions

NVIDIA Tesla M10 Quad GPU Module

HPE NVIDIA Tesla P4 8GB Module

HPE NVIDIA Tesla P40 24GB Module

HPE NVIDIA Tesla P100 PCIE 16GB Module

HPE Storage options include:

Emulex Fibre Channel HBAs

- HPE StoreFabric SN1200E 16 Gb single and dual Port Fibre Channel Host Bus Adapter
- HPE StoreFabric SN1600E 32 Gb single and dual Port Fibre Channel Host Bus Adapter

QLogic Fibre Channel HBAs

- HPE StoreFabric SN1100Q 16 Gb Single and dual Port Fibre Channel Host Bus Adapter
- HPE StoreFabric SN1600Q 32 Gb Single and dual Port Fibre Channel Host Bus Adapter

Converged network Adapters

- HPE StoreFabric CN1100R Dual Port Converged Network Adapter
- HPE StoreFabric CN1100R 10GBASE-T Dual Port Converged Network Adapter

- HPE StoreFabric CN1200E 10 Gb Converged Network Adapter
- HPE StoreFabric CN1200E 10GBASE-T Dual Port Converged Network Adapter

HPE Gen10 Chassis Intrusion Detection kit is available for physical protection. This provides a physical connection from the chassis board and hood and detects any physical intrusion into the chassis, providing security during the entire supply chain process of shipping, receiving, distribution, and operation.

HPE ProLiant DL380 Gen10: System/block diagram

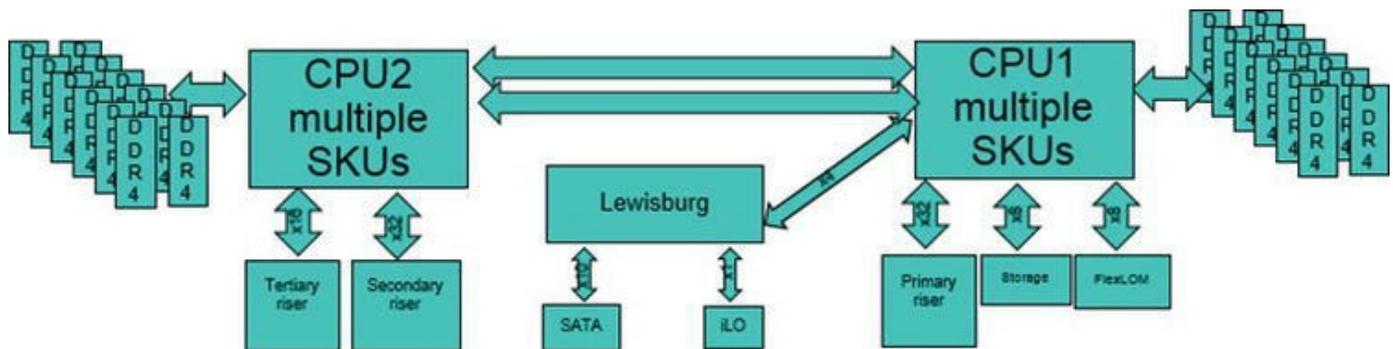


Figure 5-13 ProLiant DL380 Gen10, system/block diagram

The system block diagram in [Figure 5-13](#) shows the two-socket architecture for the HPE ProLiant DL380 Gen10. It emphasizes the distribution of IO, IO risers and memory and shows that half of the DDR4 system memory, primary riser, storage, and network FlexLOM is provided via CPU1. CPU2 is required to provide the remainder of system memory and the optional secondary and tertiary IO risers.

Activity: Using HPE QuickSpecs

This activity allows you to match the ProLiant DL360 and DL380 Gen10 rack server to customer requirements. Please do an internet search to locate the DL360 and DL380 Gen10 QuickSpecs to answer the following questions:

What is the maximum number of NVMe drives supported in the HPE ProLiant DL360 Gen10 solution?

Your customer is considering upgrading their network edge to support 10 Gb with the option to support higher bandwidths in the future. Which option should you include in the ProLiant DL380 Gen10 build to meet these requirements?

HPE ProLiant DL560 Gen10 Server



Figure 5-14 ProLiant DL560 Gen10 Server

The HPE ProLiant DL560 Gen10 Server, shown in [Figure 5-14](#) is a high-density, four-socket (4S) server with high performance, scalability, and reliability, all in a 2U

chassis. Supporting the latest Intel® Xeon® Scalable processors, the HPE ProLiant DL560 Gen10 Server offers impressive processing power, up to 6 TB of fast memory, IO of up to eight PCIe 3.0 slots, plus the intelligence and simplicity of automated management with HPE OneView and HPE iLO.

The HPE ProLiant DL560 Gen10 Server is designed for business-critical workloads and virtualized environments that require compact and powerful computing capacity. The ideal high-density four-socket server for multi-workload compute, the ProLiant DL560 Gen10 delivers performance, scalability, and expandability in a dense 2U footprint.

Features of the ProLiant DL560 Gen10 include:

Scale from four to 112 cores with up to four Intel® Xeon® Processor Scalable Family

Fast memory with (48) HPE Smart Memory DDR4 2666 MHz (6.0 TB)

Up to 24 NVDIMMs for performance improvement to accelerate data management

Dynamically tune server performance using with IST

2× NVMe drives and up to 64 TB delivering direct connect performance

High-storage density 24 SFF max, HDD/SSD, M.2 enablement kit, and 12 NVMe PCIe SSD option

Support up to two GPU, eight available IO slots + FlexibleLOM for analytics and HPC applications

Up to four 96% efficient Flex Slot1 power supplies with redundant configurations

Modular 2 + 2 processor design—scale as needed

Flexible drive bay supporting multiple combinations of NVMe and SAS SFF drives

Secure system management with iLO 5 and UEFI to protect assets and data from attack

Enhanced RAS features for high availability

Ideal use cases include:

Server consolidation, virtualization environments with a focus on data-intensive applications

Structured data management (SAP HANA, Oracle, SQL)

Business-critical workloads (ERP, CRM)

HPE ProLiant DL560 Gen10—Front system detail

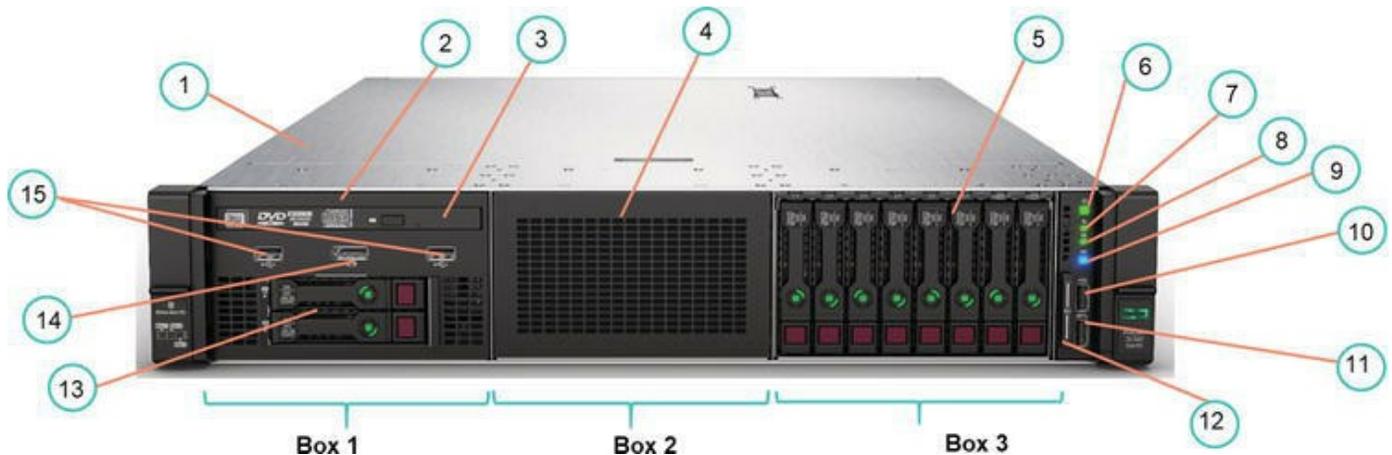


Figure 5-15 ProLiant DL560 Gen10 Server, front view

The key features of the ProLiant DL560 Gen10 as shown in [Figure 5-15](#) are:

Front view:

- | | |
|----------------------------------|--|
| Quick removal access panel | 7. Health LED |
| Optional Universal Media bay. | 8. NIC status |
| Two USB 2.0 and VGA standard | 9. UID button |
| (eight SFF bay or six SFF + two | 10. iLO front service port |
| NVMe optional) | 11. USB 3.0 |
| Optional Optical drive. Requires | 12. Serial label pull tag |
| Universal Media bay | 13. Optional two SFF HDD, requires optional |
| Drive Box 2. (eight SFF, six SFF | Universal Media bay |
| + two NVMe or eight NVMe PCIe | 14. Optional front display port (via Universal |
| SSD optional) | Media Bay) |
| Drive Box 3. (eight SFF, six SFF | 15. Optional USB 2.0 (via Universal Media |
| + two NVMe optional) | Bay) |
| Power On/Standby button and | |
| system power LED button | |

HPE ProLiant DL560 Gen10—Interior of system detail

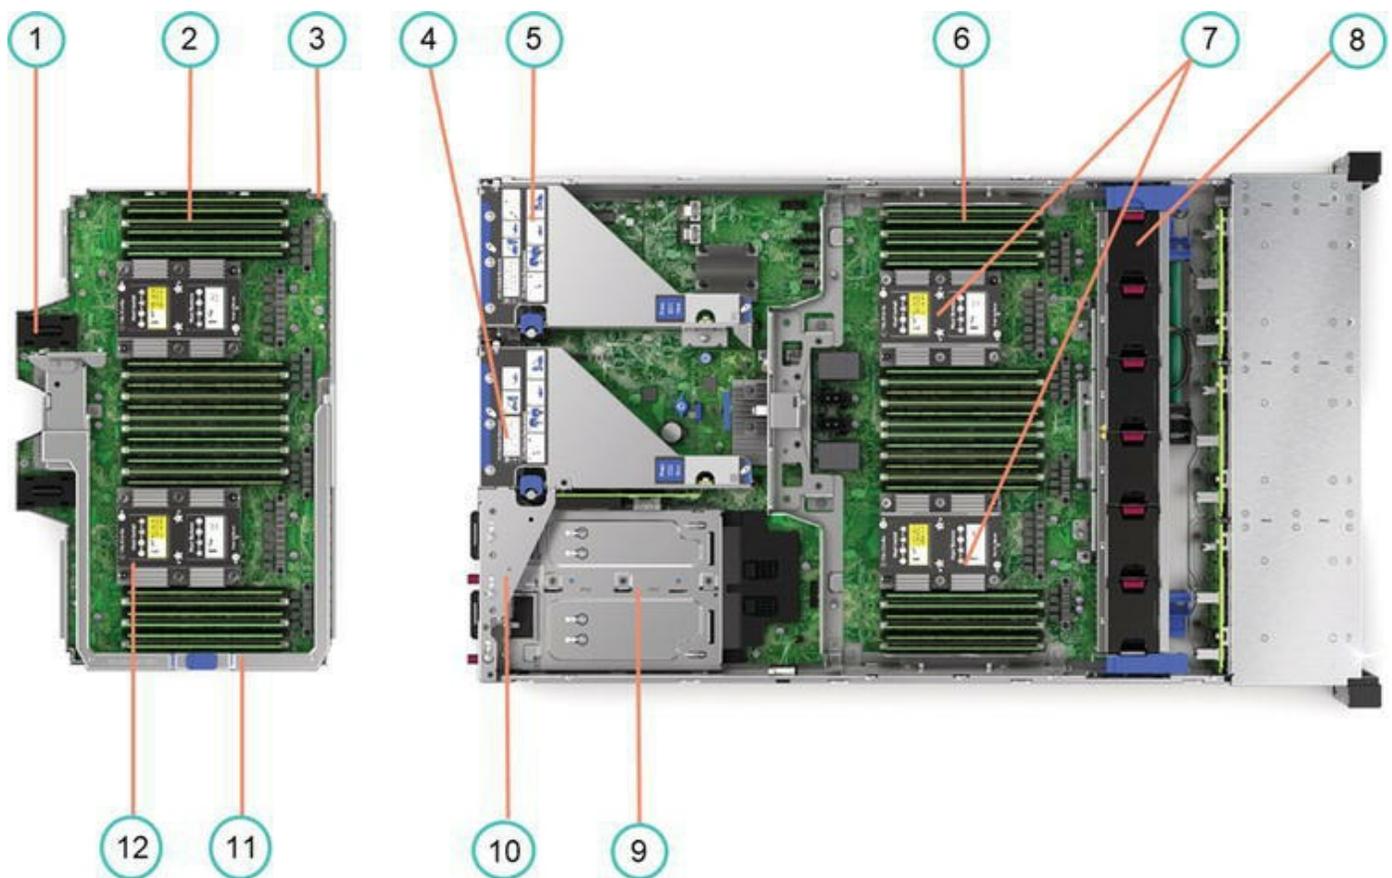


Figure 5-16 ProLiant DL560 Gen10 Server interior layout

The HPE ProLiant DL560 Gen10 Server shown in [Figure 5-16](#) has a flexible processor tray allowing you to scale up from one to four processors when you need to, saving on upfront costs and the flexible drive cage design supports up to 24 SFF SAS/SATA with a maximum of 12 NVMe drives.

The ProLiant DL560 supports up to 48 DIMM slots—24 on the system board shown on the right of [Figure 5-16](#) and 24 on the processor mezzanine board shown on the left of [Figure 5-16](#)—which support up to a total of 6 TB 2666 MT/s DDR4 HPE Smart memory. HPE DDR4 SmartMemory improves workload performance and power efficiency while reducing data loss and downtime with enhanced error handling. Choice of HPE FlexibleLOM Adapters offers a range of networking bandwidth (1 GbE–25 GbE) and fabric so you can adapt and grow to changing business needs.

DL560 supports up to eight PCIe 3.0 expansion slots for GPUs and networking cards offering increased IO bandwidth and expandability.

Up to four, 96% efficient 800W or 1600W HPE Flex slot power supplies are supported which enable higher power redundant configurations and flexible voltage ranges. The slots provide the capability to trade-off between 2 + 2 power supplies or use as extra

PCIe slots.

The key features of the ProLiant DL560 Gen10 as shown in [Figure 5-16](#) are:

Internal view:

- | | |
|--|---|
| Left connector used for DL560 four-port NVMe Mezzanine card (Daughtercard) | 7. Two Processors, heatsink showing |
| DDR4 DIMM slots. Shown fully populated in 24 slots (12 per processor) | 8. Fan cage shown with six standard hot-plug fans |
| Upper CPU Mezzanine Board Kit | 9. (Under) Hot Plug redundant HPE Flexible Slot Power supplies |
| Optional secondary PCIe riser | 10. Optional Tertiary riser |
| Default primary PCIe riser | 11. Handle for removing upper CPU Mezzanine Board Kit |
| DDR4 DIMM slots on upper CPU mezzanine board kit. Shown fully populated in 24 slots (12 per processor) | 12. Two Processors, heatsink showing on upper CPU mezzanine board kit |



Note

The optional four-port NVMe Mezzanine card 874633-B21 supports a maximum of eight NVMe drives and does not consume a PCIe slot. It goes on top of the HPE DL5x0 Gen10 CPU Mezzanine Board Kit (872222-B21) and requires a four-processor configuration. See DL560 QuickSpecs for further detail.

HPE ProLiant DL560 Gen10: Rear system detail

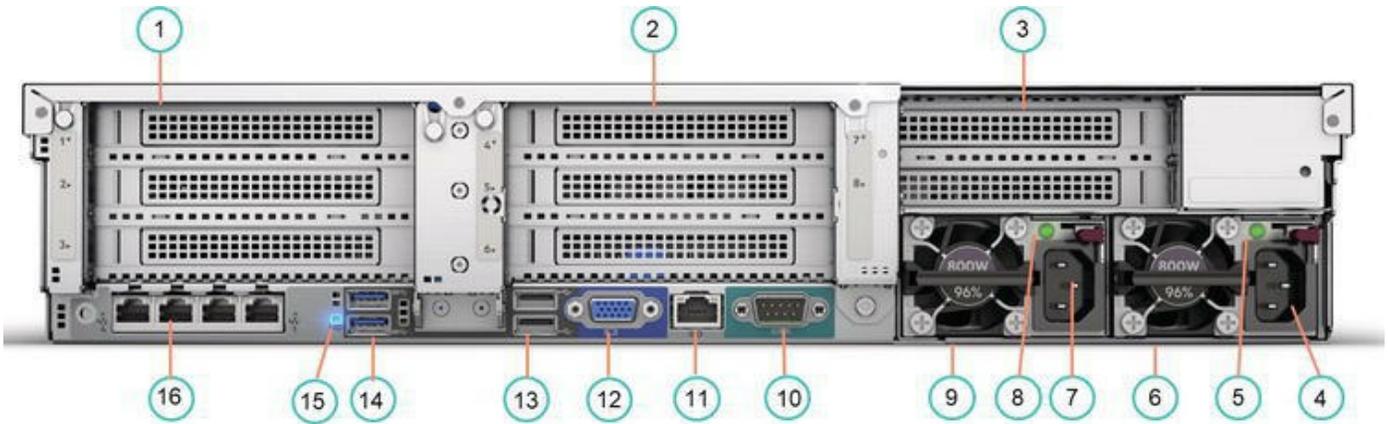


Figure 5-17 ProLiant DL560 Gen10 Server, rear view

Key features of the ProLiant DL560 Gen10 as shown in [Figure 5-17](#) are:

Rear view:

- | | |
|---|---|
| PCIe Slots (Slots 1 to 3 top to bottom, riser shipped standard) | 16. FlexibleLOM ports (Port 1 on right side) |
| PCIe Slots (Slots 4 to 6 top to bottom, requires second riser card and second processor) | 15. Unit ID LED |
| PCIe Slots (Slots 7-8 top to bottom), requires tertiary riser card and second processor, Not available with 4× Flex Slot power supplies | 14. USB connectors 3.0 (two) |
| Power supply Power connection | 13. USB connectors 2.0 (two) |
| Power supply Power LED | 12. VGA (video) connector |
| HPE Flexible Slot Power Supply bay 1 (800W PS shown) | 11. Dedicated iLO connector |
| Power supply Power connection | 10. Serial Port |
| | 9. HPE Flexible Slot Power Supply bay 2 (800W power supply shown) |
| | 8. Power supply Power LED |

HPE ProLiant DL560 Gen10: System/block diagram

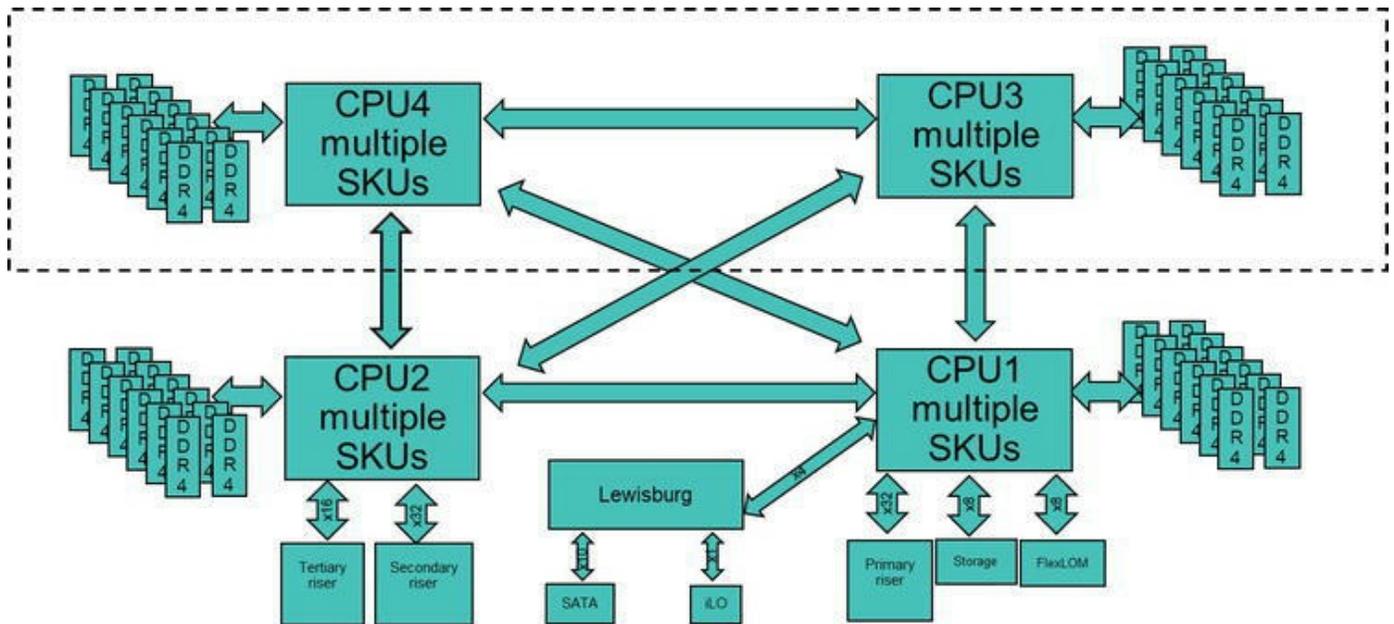


Figure 5-18 ProLiant DL560 Gen10 Server, system/block diagram

The system block diagram in [Figure 5-18](#) shows the four-socket architecture for the HPE ProLiant DL560 Gen10. It shows the distribution of IO, IO risers and memory, and that one-quarter of the system memory, primary riser, storage, and network FlexLom is provided by CPU1. CPU2 is required to provide the optional secondary and tertiary IO risers. Each of the processors contributes one-quarter of the system memory capacity.

HPE ProLiant DL560 Gen10 options

No hard drives ship standard on this server. Box 3 supporting eight SFF SAS/SATA is shipped as default. The eight hot-plug SFF SAS/SATA hard drive bays are available in entry and base models, and 16 in performance models, upgradable to a maximum of 24. If one of the two optional Universal Media Bay Kits is chosen, a maximum of 16 SFF drives is available instead of 24. The standard Universal Media Bay Kit occupies an eight SFF drive bay cage and includes an optional optical disk drive; another version has an optional two SFF drive bay as well. The eight NVMe drive option can only be placed in Bay 2 and a further two located in Bay 3. All preconfigures models come with embedded software RAID support for 10 SATA drives. Optional HPE Smart Array Controllers can be added.

The storage controller on the DL560 Gen10 Server is one of the following depending on the model:

Entry model—HPE Dynamic Smart Array S100i SR SW RAID Controller, operating in UEFI mode only.

Base model—HPE Smart Array E208i-a, E208i-p, and E208e-p SR Gen10 Controller. This controller is in the Adaptive RAID on Chip slot, which enables SAS drives and has 2 GB FBWC for data retention.

Performance model—HPE Smart Array P408i-a, P408e-p, and p816i-a SR Gen10 Controller. This controller enables SAS drives and has 4 GB FBWC for enhanced data retention. Performance RAID Controllers require the HPE Smart Storage Battery (sold separately).

ProLiant DL560 Gen10 Servers offer network technology, which offers customers a choice of 1 GbE four-port 331 FLR, FlexFabric 10 GbE two-port 533FLR-T, base-T Ethernet, or converged networking in their embedded adapter, available in Entry and Base models. In the Performance models, the 10/25 GbE two-port 640FLR-SFP28 and 10/25 GbE two-port 631FLR-SFP28 adapters are available. A range of NIC cards is also available to enhance networking capabilities.

Entry models have a single HPE 1600W Flex Slot Platinum Hot Plug Low Halogen Power Supply kit. Base and Performance models have redundant HPE 1600W Flex Slot Platinum Low Halogen Hot Plug Power Supply kit. Also available 4x800W, -48V DC, and 227VAC/380VDC power input 800W Flex Slot Platinum Low Halogen power supply kits.

HPE Flex Slot power supplies allow for commonality of power supplies across ProLiant servers and are designed to provide the highest power supply efficiency without degrading system performance. HPE Flexible Slot Power Supplies provide up to 96% power efficiency with 80Plus Platinum-certified power supplies that lower server power requirements and reduce power waste in your data center. The power supply form factor is 25% smaller than previous generation HPE Common Slot Power Supplies, providing more space to add additional server options and improve IT performance. All Flex Slot power sources are UL, CE Mark Compliant, hot-plug, and support-redundant configurations.

HPE Flex Slot Platinum power supplies are required when enabling the HPE Intelligent Power Discovery (IPD) solution. IPD is the first technology to create an automated, energy-aware network between IT systems and facilities. This allows companies to reclaim millions of dollars in wasted power capacity and downtime costs across data centers.



Note

Mixing power supplies in the same server is not supported. All power supplies must be of the same output and efficiency rating. If you install non-matched power supplies, you will get errors and operations will fail. For additional information, refer to the HPE Flexible Power Supplies site by entering the following URL into your browser:

<https://www.hpe.com/us/en/product-catalog/servers/power-supplies/pip.hpe-flexible-slot-power-supplies.7268787.html>

HPE ProLiant DL580 Gen10 Server



Figure 5-19 ProLiant DL580 Gen10 Server

The HPE ProLiant DL580 Gen10 Server, shown in [Figure 5-19](#), is a high-density, four-socket (4S) server with high performance, scalability, and reliability, all in a 4U chassis. Supporting the latest Intel® Xeon® Scalable processors, the HPE ProLiant DL580 Gen10 Server offers greater processing power, up to 6 TB of fast memory, IO of up to 16 PCIe 3.0 slots, plus the intelligence and simplicity of automated management with HPE OneView and HPE iLO 5.

The HPE ProLiant DL580 Gen10 Server is the ideal server for business-critical workloads and general 4P data-intensive applications where the right performance is paramount.

Features of the ProLiant DL580 Gen10 include:

Scale from four to 112 cores with up to 4 Intel® Xeon® Processor Scalable Family

Fast memory with (48) HPE Smart Memory DDR4 2666MHz (maximum 6.0 TB)

Up to 24 NVDIMMs for performance improvement to accelerate data management

Dynamically tune server performance using with IST

2× more NVMe drives and up to 64 TB delivering direct connect performance

High-storage density 24 SFF max, HDD/SSD, M.2 enablement kit, and 12 NVMe PCIe SSD option

Support up to two GPU, eight available IO slots + FlexibleLOM for analytics and HPC applications

Up to four 96% efficient Flex Slot1 power supplies with redundant configurations

Modular 2 + 2 processor design—scale as needed

Flexible drive bay supporting multiple combinations of NVMe and SAS SFF drives

Secure system management with iLO 5 and UEFI to protect assets and data from attack

Enhanced RAS features for high availability

Customers looking to deploy SAP HANA in small instances for dev/test/QA or in production environments have the option to leverage ProLiant DL580 Gen10 servers supporting two to four socket configurations with up to 6 TB of memory (using 128 GB DIMMs). This configuration is ideally suited for smaller HANA workloads and can support multiple HANA use case scenarios.

Ideal use cases include:

Server consolidation, virtualization environments

Structured data management (Oracle, SQL)

Business-critical workloads (ERP, CRM)

Graphics-intensive workloads (VDI)

HPE ProLiant DL580 Gen10—Front system detail

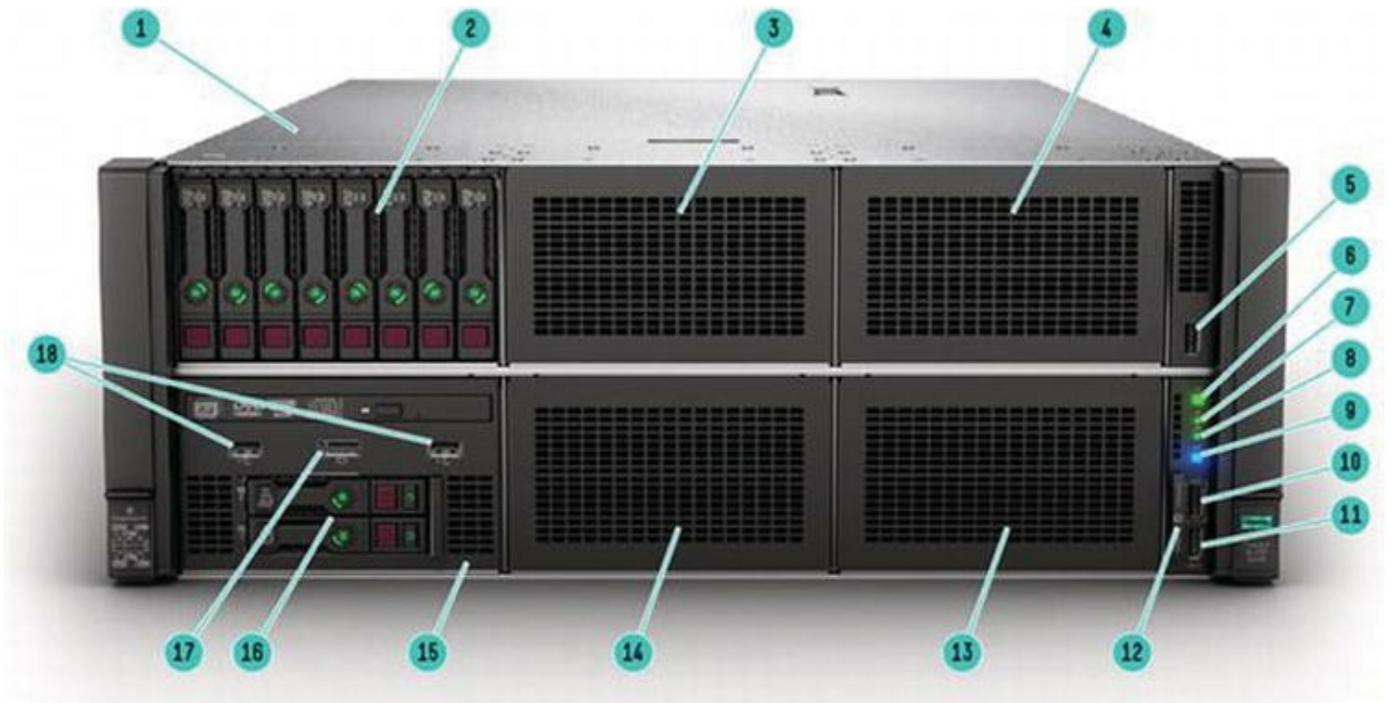


Figure 5-20 HPE ProLiant DL580 front view

The DL580 Gen10 is the ideal server for business-critical workloads, virtualization, server consolidation, database, business processing, graphics intensive, and general 4P data-intensive applications where the right performance is paramount.

Key features of the ProLiant DL580 Gen10 as shown in [Figure 5-20](#) are:

- | | |
|---|---|
| Quick removal access panel | 9. UID button |
| Box 1 (eight SFF, six SFF + two NVMe or eight NVMe (supports only four NVMe drives) SSD optional) | 10. iLO Front Service Port |
| Box 2 (eight SFF, six SFF + two NVMe or eight NVMe PCIe SSD optional) | 11. Front USB 3.0 port |
| Box 3 (eight SFF, six SFF + two NVMe or eight NVMe PCIe SSD optional) | 12. Serial label pull tag |
| Front USB 3.0 port | 13. Box 6 (eight SFF) |
| Power On/Standby button and system power LED button | 14. Box 5 (eight SFF) |
| Health LED | 15. Box 4 (eight SFF or Universal Media bay) |
| NIC status | 16. Optional two SFF HDD, requires optional Universal Media bay |
| | 17. Optional front display port (via Universal Media Bay) |
| | 18. Optional USB 2.0 (via Universal Media Bay) |

HPE ProLiant DL580 Gen10—Rear system detail

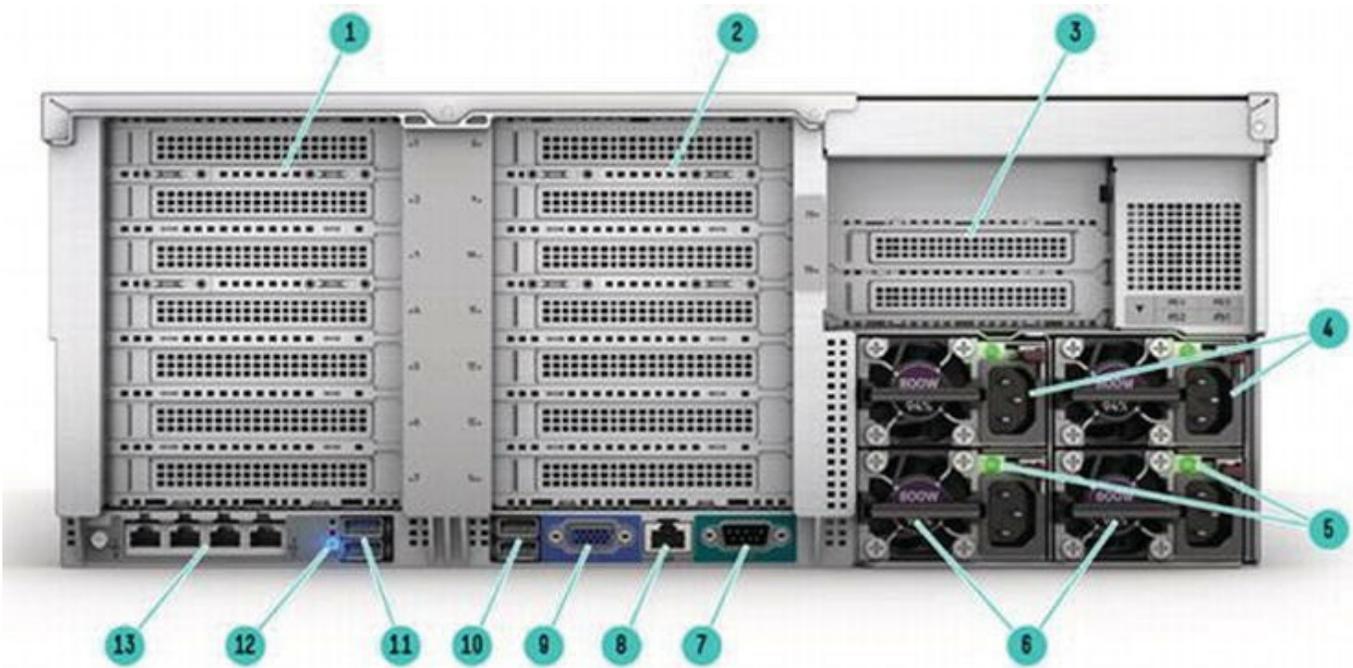


Figure 5-21 HPE ProLiant DL580 rear view

Key features of the ProLiant DL580 Gen10 as shown in [Figure 5-21](#) are:

- | | |
|---|--|
| PCIe Slots (Slots 1-7 top to bottom), requires primary riser | 6. HPE Flexible Slot Power Supply, 800W PS shown (max. four) |
| PCIe Slots (Slots 8-14 top to bottom), requires secondary riser (includes tertiary riser) | 7. Serial connector |
| PCIe Slots (Slots 15-16 top to bottom), requires tertiary riser (included with secondary riser) | 8. Dedicated iLO connector |
| Power supply Power connection (max. four) | 9. VGA (video) connector |
| Power supply Power LED (max. four) | 10. USB connectors 2.0 (two) |
| | 11. USB connectors 3.0 (two) |
| | 12. Unit ID LED |
| | 13. FlexibleLOM ports (Port 1 on right side) |

HPE ProLiant DL580 Gen10—Interior of system detail

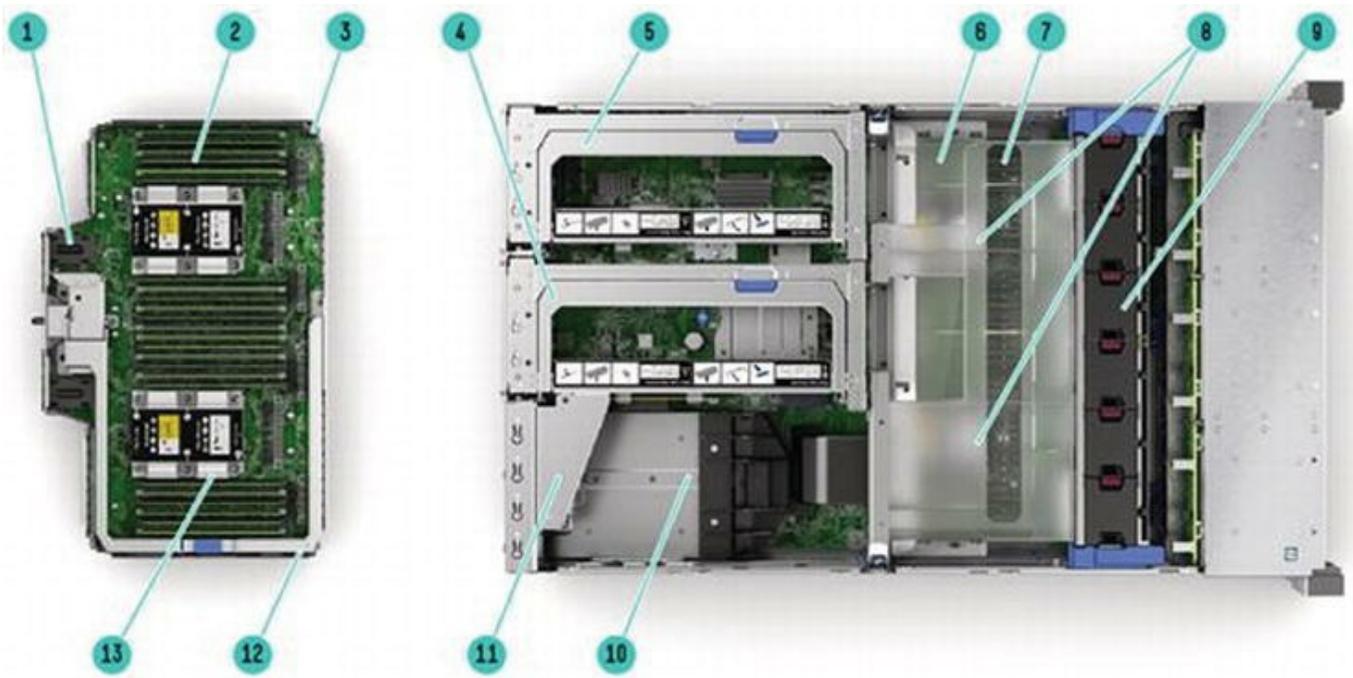


Figure 5-22 HPE ProLiant DL580 interior view with upper CPU mezzanine tray

Key features of the ProLiant DL580 Gen10 as shown in [Figure 5-22](#) are:

- | | |
|---|---|
| Left connector used for DL580 four-port NVMe Mezzanine card (Daughtercard) | 8. Two Processors (under the air baffle) |
| DDR4 DIMM slots. Shown fully populated in 24 slots (12 per processor) | 9. Fan cage shown with 12 standard hot-plug fans |
| Upper CPU Mezzanine Board Kit | 10. (Under) Max. four Hot Plug redundant HPE Flexible Slot Power supplies |
| Optional secondary PCIe riser (includes tertiary riser) | 11. Optional Tertiary riser (included with secondary riser) |
| Optional primary PCIe riser | |
| Air baffle | |
| DDR4 DIMM slots on CPU board kit. Shown fully populated in 24 slots (12 per processor) under the air baffle | 12. Handle for removing upper CPU Mezzanine Board Kit |
| | 13. Two Processors, heatsink showing on upper CPU mezzanine board kit |

Additional features of the ProLiant DL580 Gen10 are:

Support for the Intel® Xeon® Processor Scalable Family supporting up to four processors with core count ranging from four to 28 per processor

Delivers up to 4.8X more disk drives (max. 48) and up to 4X more NVMe drives (max. 20) compared to the ProLiant DL580 Gen9 Server. Up to 6 TB of 2666 MT/s HPE DDR4 Smart Memory modules for increased performance for data-intensive application workloads

Massive expandability with up to 16 PCIe slots compared to nine in previous generations with support for up to four full length/full height GPU cards

DDR4 Registered (RDIMM) or Rank Load Reduced (LRDIMM). With 12 DIMM sockets per processor, six channels per processor and two DIMMs per channel for a total of 48 DIMM slots



Note

LRDIMM and RDIMM technologies are distinct memory types and cannot be mixed within a server.



Note

HPE memory from previous generation servers are not qualified or warranted with this server. HPE SmartMemory is required to realize the memory performance improvements and enhanced functionality for ProLiant Gen10 servers. For additional information, refer to the HPE SmartMemory QuickSpecs by entering the following URL into your browser:

<https://h20195.www2.hpe.com/v2/GetHTML.aspx?docname=c04355083>

The HPE ProLiant DL580 Gen10 servers offer a flexible network technology—FlexibleLOMs, which offers customers a choice of 1 Gb, 10 Gb, 25 Gb, or 10 Gb base-T Ethernet or converged networking in their embedded adapter. A range of NIC cards is also available to enhance networking capabilities:

Entry model—HPE Ethernet 1Gb four-port 331FLR Adapter

Base model—HPE FlexFabric 10Gb two-port 535FLR-T Adapter

Performance model—HPE FlexFabric 10/25 Gb two-port 640FLR-SFP28 Adapter

The storage controller on the DL580 Gen9 server is one of the following depending on

the model:

Essential RAID

- HPE Smart Array E208i-p SR Gen10 Controller
- HPE Smart Array E208e-p SR Gen10 Controller

Performance RAID

- HPE Smart Array P408i-p SR Gen10 Controller
- HPE Smart Array P408e-p SR Gen10 Controller

HPE ProLiant Gen10 workload-based positioning

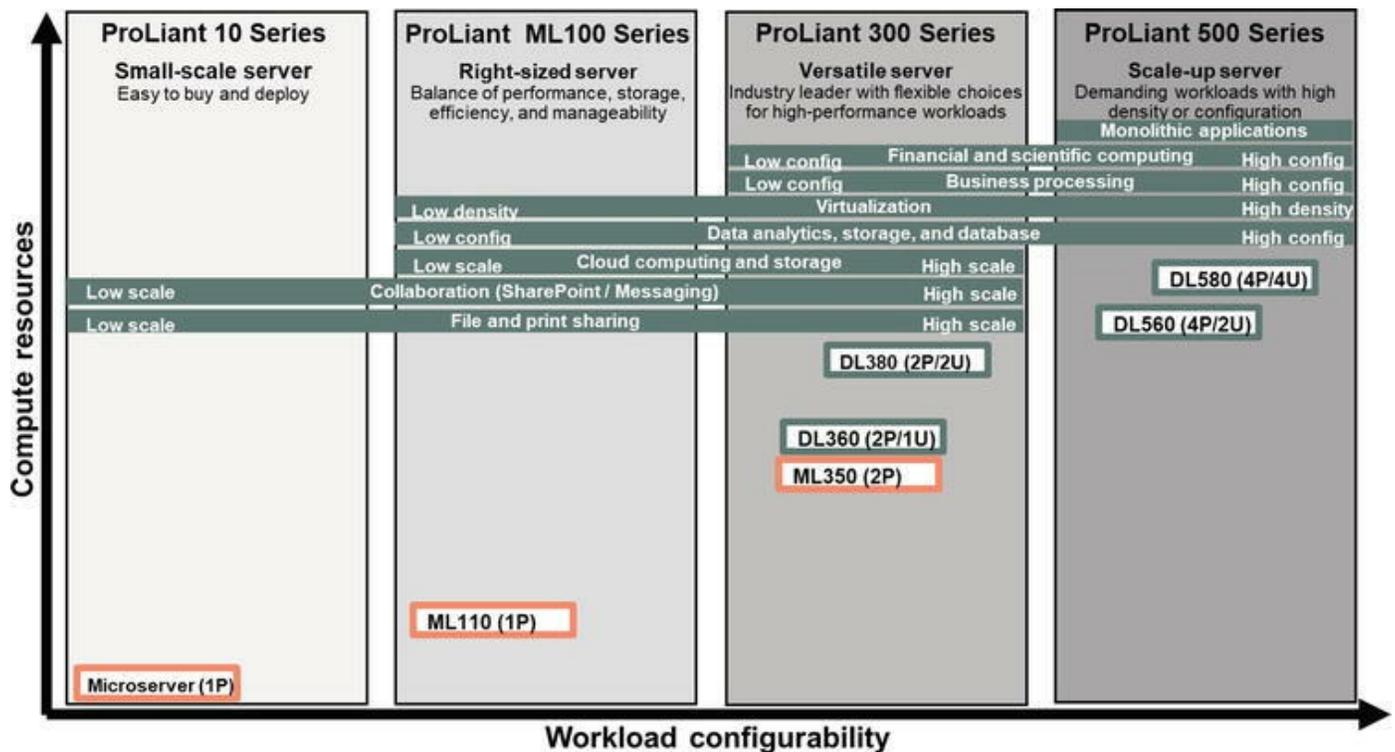


Figure 5-23 HPE ProLiant Gen10 workload-based positioning by series

To choose the right system for business needs, you should start by answering the right questions: What is the consumption model? What is the adoption strategy? What workloads are these systems accommodating? [Figure 5-23](#) outlines where each ProLiant Gen10 model fits into these considerations.

Looking first at the consumption model, decide if the customer needs preconfigured systems or if they have an IT staff able to build-out the data center. Next, think about

density. If they have limited physical space, they will need to ensure that compute density is maximized. Examine what form factor is typically purchased such as rack, tower, blade platforms, or Synergy. Alternatively, does the workload require hardware that is purpose-built to solve their compute needs?

Next, look at the adoption strategy. Are they growing at a consistent pace over time or is this a new project that is being deployed all at once? For example, if the project is a 2400-seat VDI migration, are they migrating 200 seats a month target completion in one year or are they migrating all 2400 seats at one time.

Finally, consider the workloads and applications that the systems will be running. The broad HPE compute portfolio can accommodate every workload, and with service and finance, offerings can deliver the right support and economics as well.

Learning check

You are attending a meeting with a customer to discuss their future data center modernization plans. During the discussion they raise the subject of hyperconverged solutions and that they are considering Microsoft Storage Spaces Direct. When they considered the HPE DL380 Rack Server solution, they were concerned that the platform may not offer sufficient HDD and SSD storage to meet their needs. How should you respond?

A customer has expressed a need for a four-socket ProLiant rack server to host one of their business-critical applications? Which server should you recommend?

- A. DL380 Gen10
- B. DL180 Gen9
- C. DL560 Gen10
- D. DL360 Gen9

HPE ProLiant DL385 Gen10



Figure 5-24 HPE ProLiant DL385 Gen10

The HPE ProLiant DL385 Gen10 server, shown in [Figure 5-24](#), is an enterprise-grade, dual-socket 2U rack server that delivers enhanced performance, enterprise-class management, expandability, and scalable capabilities to meet the needs of general and

heavy workloads. A secure and flexible server that delivers advanced performance for virtualized and memory-centric workloads.

The HPE ProLiant DL385 Gen10 Server supports the AMD® EPYC™ 7000 series server processor for data center with up to 32 cores, plus eight memory channels for HPE DDR4 2666 MT/s SmartMemory supporting up to 4.0 TB max.

Ideal for virtualization, general-purpose compute, memory-centric workloads, high core count workloads, SharePoint Storage, and Enterprise applications.

Table 5-1 HPE ProLiant DL385 Gen10 features

CPU	EPYC up to 64 cores (up to 128 threads)
Memory	32 DIMMs (up to 4 TB)
Storage Controller	Integrated S100i and Flexible Smart Array
Storage Capacity	30 SFF or 19 LFF + 2 SFF or 24 NVMe
Networking	4 × 1 GbE, FlexibleLOM, Standup
Management	iLO 5
Power	Redundant Flex Slot
Additional	Embedded M.2, USB, and MicroSD

HPE ProLiant DL385 Gen10 options

Table 5-2 HPE ProLiant DL385 Gen10 options

DL385 Gen10	
Compute	Up to (two) AMD® EPYC® 7000 Series Processor family, (32C), PCIe 3.0, up to eight available slots + m.2 support embedded
Memory	HPE Smart Memory (32) DDR4 (eight channels per CPU), up to 2667 MT/s (4 TB max)
Persistent Memory	Not currently supported
Storage	Standard HPE Smart Array S100i Software RAID for M.2 connectors Choice of HPE Smart Array Essential and Performance RAID Controllers for performance or additional features
HPE	24+6 SFF/12+4+3 LFF + 2 SFF max, HDD/SSD, 24

SmartDrives	NVMe PCIe SSD, and m.2 connector embedded on system board
Networking	4 × 1 GbE embedded + Choice of FlexibleLOM + Standup
VGA/Serial/USB Ports/SD	Front Display Port Opt, Rear VGA & Optional Serial, 5 USB 3.0, +2 USB 2.0 optional, Dual Micro-SD
GPU Support	Single/Double-Wide (five) & Active/Passive up to 10.5" (three)
Management	
Converged	HPE OneView and HPE iLO Advanced
Support	HPE Insight Online with enhanced mobile app
Embedded	HPE iLO 5, SUM, RESTful Interface Tool, UEFI
Power & Cooling	Up to 96% efficiency (Titanium) with Flexible Slot PS
Industry Compliance	ASHRAE A3 and A4, lower idle power, Energy Star
Power/Location Discovery Services	Supported/Unsupported
Chassis Depth	26.75" (SFF), 28.75" (LFF)
Serviceability— Easy Install Rails	Standard
Warranty	3/3/3

[Tables 5-1](#) and [5-2](#) list options and features available for the ProLiant DL385 Gen10.

HPE ProLiant DL385 Gen10—Front system detail

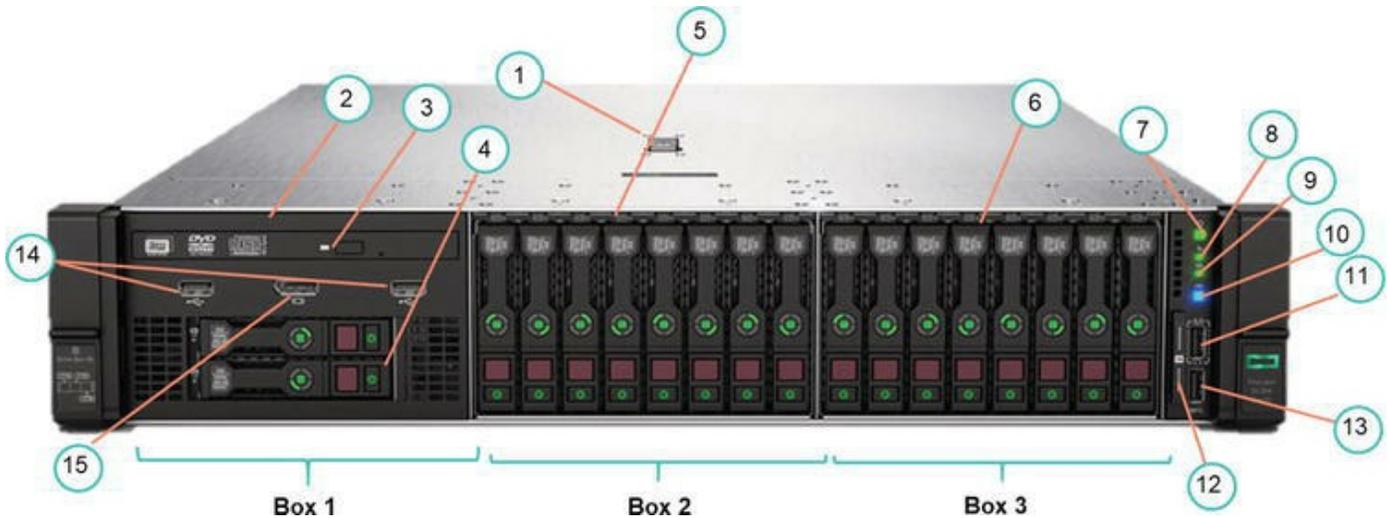


Figure 5-25 HPE ProLiant DL385 front view

Key features of the ProLiant DL385 Gen10 as shown in [Figure 5-25](#) are:

Front view:

Quick removal access panel

Drive Bay 1. Optional Universal Media bay. Two USB 2.0 and Display port standard shown (eight SFF bay or six SFF + two NVMe or eight NVMe optional)

Optional Optical drive. Requires Universal Media bay

Optional two SFF HDD. Requires optional Universal Media bay.

Drive Bay 2. NVMe shown (eight SFF, six SFF + two NVMe or eight NVMe PCIe SSD optional)

Drive Bay 3. NVMe shown (eight SFF, six SFF + 2NVMe or eight NVMe PCIe SSD optional)

Power On/Standby button and system power LED button

8. Health LED

9. NIC status

10. Unit ID button

11. iLO front service port

12. Serial label pull tag

13. USB 3.0

14. Optional USB 2.0 (via Universal Media Bay)

15. Optional front display port (via Universal Media Bay)

HPE ProLiant DL385 Gen10: Interior of system detail (eight SFF)

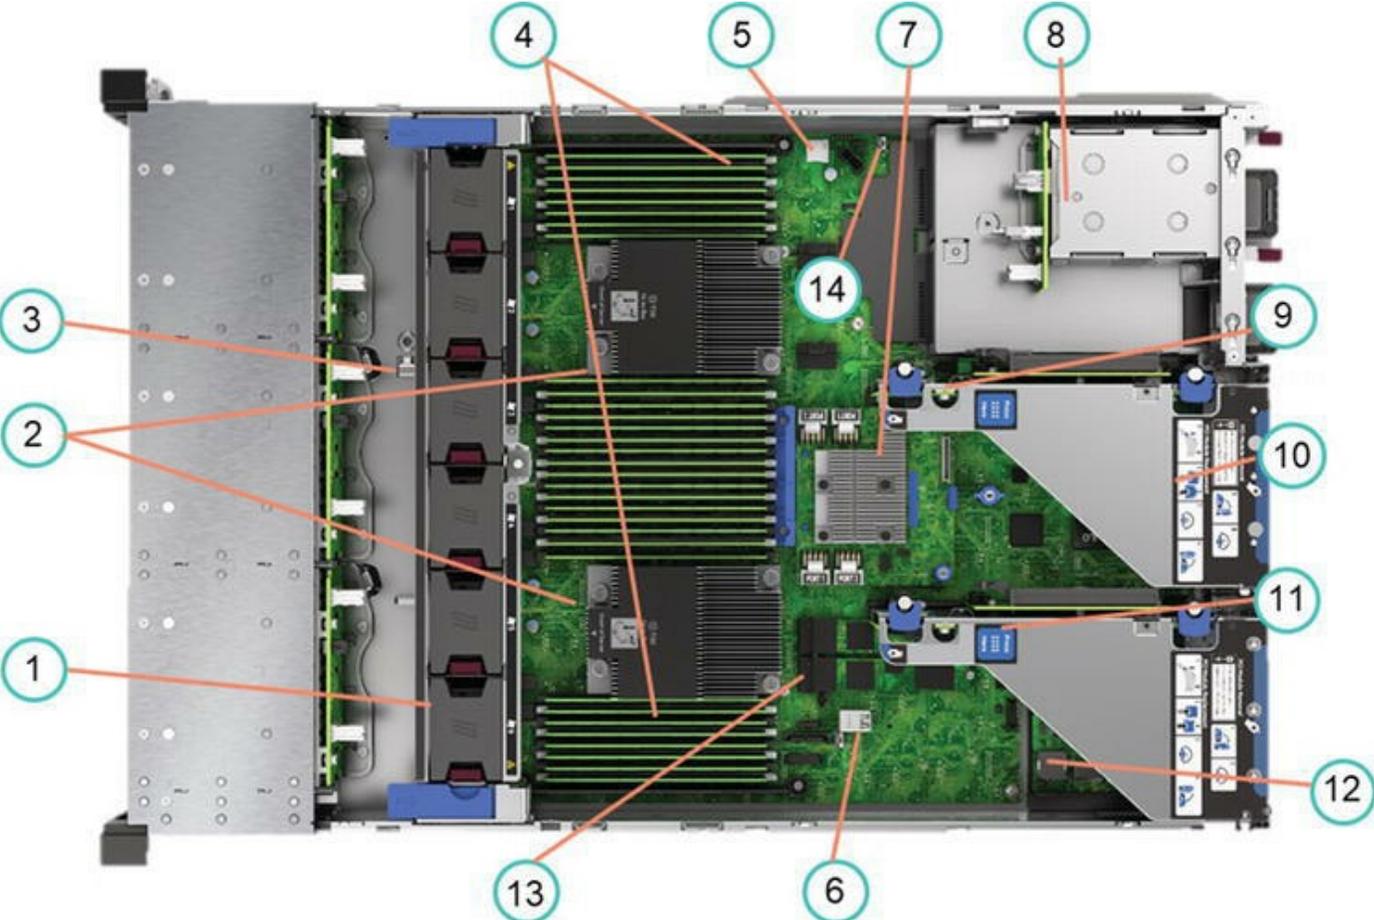


Figure 5-26 HPE ProLiant DL385 interior view

Key features of the ProLiant DL385 Gen10 as shown in [Figure 5-26](#) are:

Internal view:

Fan cage shown with six standard hot-plug fans (high-performance fans optional)

Two Processors (heatsinks shown)

Optional: Smart Storage Battery location (Battery not shown)

DDR4 DIMM slots. Shown fully populated in 32 slots (16 per processor)

MicroSD card slot (Optional dual MicroSD option)

Internal USB 3.0 connector

Optional HPE Flexible Smart Array Controller (P408i-a shown)

8. (Under) Hot Plug redundant HPE Flexible Slot Power supplies

9. Connection for second (optional) riser (requires second CPU)

10. Embedded 4 × 1 GbE NIC

11. Primary PCIe riser, standard (optional double-wide GPU riser)

12. Optional Flexible LOM slot

13. Embedded M.2 connectors

14. Chassis intrusion detection connector

HPE ProLiant DL385 Gen10—Rear view system detail

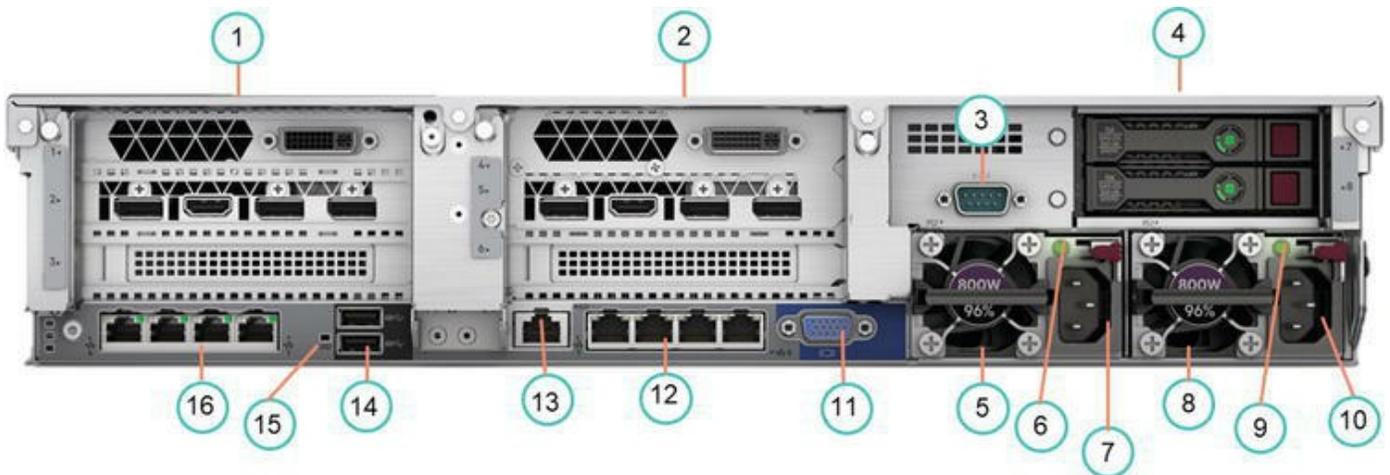


Figure 5-27 HPE ProLiant DL385 rear view

Key features of the ProLiant DL385 Gen10 as shown in [Figure 5-27](#) are:

Rear view:

Primary Riser. PCI Slots (Slots 1-3 top to bottom, riser shipped standard, not shown) Optional two SFF rear drives	8. HPE Flexible Slot Power Supply bay 1
	9. Power supply power LED
	10. Power supply power connection
Secondary Riser. PCI Slots (Slots 4-6 top to bottom, not shown, requires second riser card, and second processor). Optional two SFF rear drives	11. VGA connector
	12. Embedded 4 × 1 GbE Network Adapter
	13. Dedicated iLO management port
	14. USB connectors 3.0 (two)
Optional serial port	15. Unit ID LED
Tertiary Riser (Slots 7-8 top to bottom, not shown). Optional rear two SFF HDD	16. Optional FlexibleLOM ports (Shown: 4 × 1 GbE)
HPE Flexible Slot power supply bay 2	
Power supply power LED	
Power supply power connection	

Main memory encryption and encrypted virtualization

AMD 7000 series processors include a dedicated security processor within the server System on Chip (SoC) product. This Secure Processor is an integrated ARM[®] Cortex[®] A5 that sits alongside the main CPU. It provides a dedicated secure space to run multiple security-related functions that require full isolation and can be offloaded from the main CPU. One of these functions is secure OS, which is the starting point for the AMD Secure Root of Trust technology, one of three basic categories of security offered by the EPYC SoC. The other two are AMD Secure Run and AMD Secure Move technologies. These are summarized in [Table 5-3](#).

Table 5-3 AMD Security on Chip features

Secure Root-of-Trust (AMD Secure Processor, Secure Boot)	Prevents the use of rootkits/bootkits that may inject malicious code prior to an OS loading creating a hardware root of trust, enabling only known and trusted software to be loaded and run—from initial boot load through BIOS load.
Secure Run (Secure Memory Encryption, Secure Encrypted Virtualization)	Memory scraping and cold boot attacks are both modes of attack that can be used when data running in the main system memory is not encrypted. Secure Memory Encryption (SME) encrypts system memory. Secure Encrypted Virtualization (SEV) isolates the hypervisor and guest VMs to prevent access to data in shared guest data areas.
Secure Move (SEV-enabled servers, APIs, third-party key management)	Migrating VMs within the data center, to off-premises private cloud, or to the public cloud can be insecure. SEV-enabled servers can establish a secure channel between them and send memory encryption keys to the remote platform.

Secure Root-of-Trust is a well-known security requirement for bare-metal system boots. In addition to securely booting the native system (including authenticating AMD Secure Processor code and platform BIOS code), AMD also supports secure boot for SEV virtual machines (VMs) and can prove that the boot image for those machines has not been tampered with by the hypervisor or another third-party. The system authenticates and loads code for the AMD Secure Processor to perform key management.

Secure Run is a technology that provides for encryption of data in the main system memory. This is an expansion beyond other encryption technologies that secure “data at rest” (on a disk or solid-state drive) or “data in motion” (on a network). The focus of Secure Run is the two central technology additions to AMD EPYC: SME and SEV.

AMD integrates technology in the CPU, requiring no application software modifications. Some OS and hypervisor-enablement is required for both SME and SEV.

SME uses dedicated hardware in the on-die memory controllers with an Advanced Encryption Standard Engine to encrypt data written to DRAM and decrypt the data when read. Encryption keys, managed by the AMD Secure Processor, are randomly generated on each system reset and not visible to software running on the main CPU cores.

SEV directly addresses the problem of the unrestricted access that the hypervisor has to data and code. When using SEV, neither hypervisor nor guest operating system has access to the resource of the other, meaning that lower-privileged code is protected and no longer dependent on higher-privileged code for startup and execution. The hypervisor and guest layers are still able to communicate, but only through much more tightly controlled communication paths.

Secure Move relies on the establishment of a secure channel between two SEV-enabled platforms so that the hypervisor can implement migration and snapshot functions securely. VMs may be migrated to other SEV-enabled platforms within the data center itself, to private off-premises cloud, to a disaster recovery solution, or to the public cloud. With the secure connection established, SEV firmware sends the guest’s memory encryption keys over to the remote platform, which can then run the guest.

HPE ProLiant Gen9 rack server portfolio

Table 5-4 HPE ProLiant Gen9 rack server portfolio workload matching

Essential		Performance	
Entry	Growth	Traditional	Scale up
DL80 Gen9	DL180 Gen9	DL380 Gen9	DL580 Gen9
DL60 Gen9	DL160 Gen9	DL360 Gen9	DL560 Gen9
DL20 Gen9	DL120 Gen9		
<ul style="list-style-type: none"> - File and print - Messaging - Academic and research 	<ul style="list-style-type: none"> - No SQL, Hadoop - Low VM density 	<ul style="list-style-type: none"> - Mid to high VM density - Mission-critical apps 	<ul style="list-style-type: none"> - Large databases - UNIX alternative - Monolithic applications

[Table 5-4](#) lists the ProLiant Gen9 rack server portfolio and the types of workloads that they are ideal for.

HPE ProLiant DL20 Gen9 Server

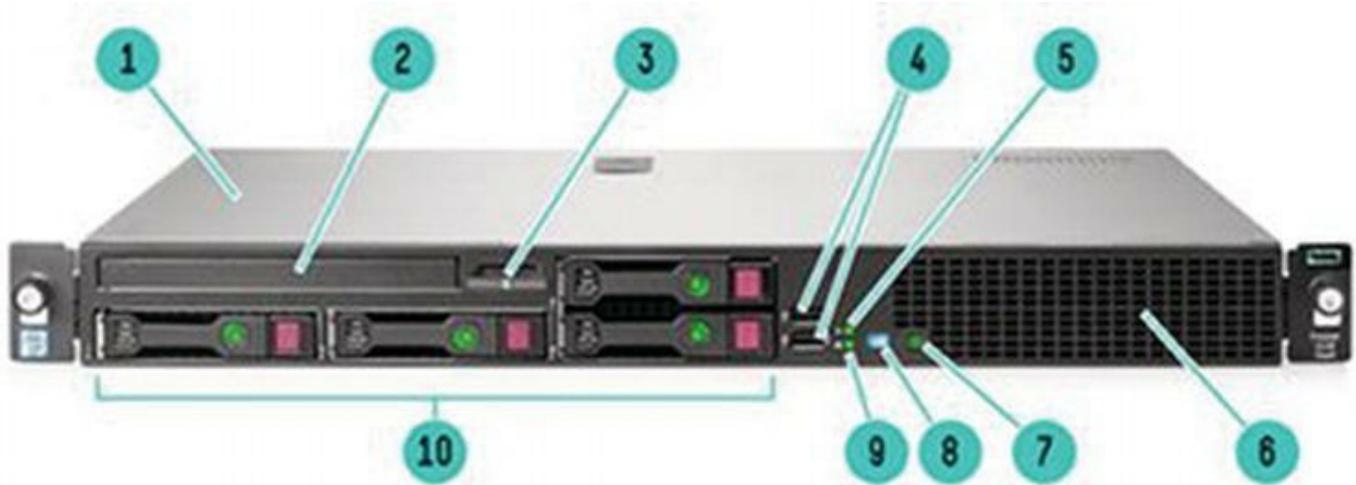


Figure 5-28 HPE ProLiant DL20 Gen9 Front view

The dense and compact HPE ProLiant DL20 Gen9 Server powered by Intel[®] Pentium, Core i3, or Xeon[®] E3-1200 v5 processors provides a unique blend of enterprise-class capabilities and value. It offers outstanding configuration flexibility to meet a variety of business requirements at an affordable price.

Key features of the HPE ProLiant DL20 Gen9 Server shown in [Figure 5-28](#) are:

- | | |
|---|--|
| 1. Access panel | 7. Power On/Standby button and system power LED button |
| 2. Optical drive | 8. UID LED |
| 3. Serial number/iLO 4 information pull tab | 9. NIC status LED |
| 4. USB 2.0 connectors | 10. Hard drive bays |
| 5. Health LED | |
| 6. Front two fan assembly | |

The storage controller on the DL20 Gen9 Server is one of the following depending on the model:

Entry and base models—Dynamic Smart Array B140i Controller

Performance models—H240 FIO Smart Host Bus Adapter

HPE ProLiant DL60 Gen9 Server: Internal view

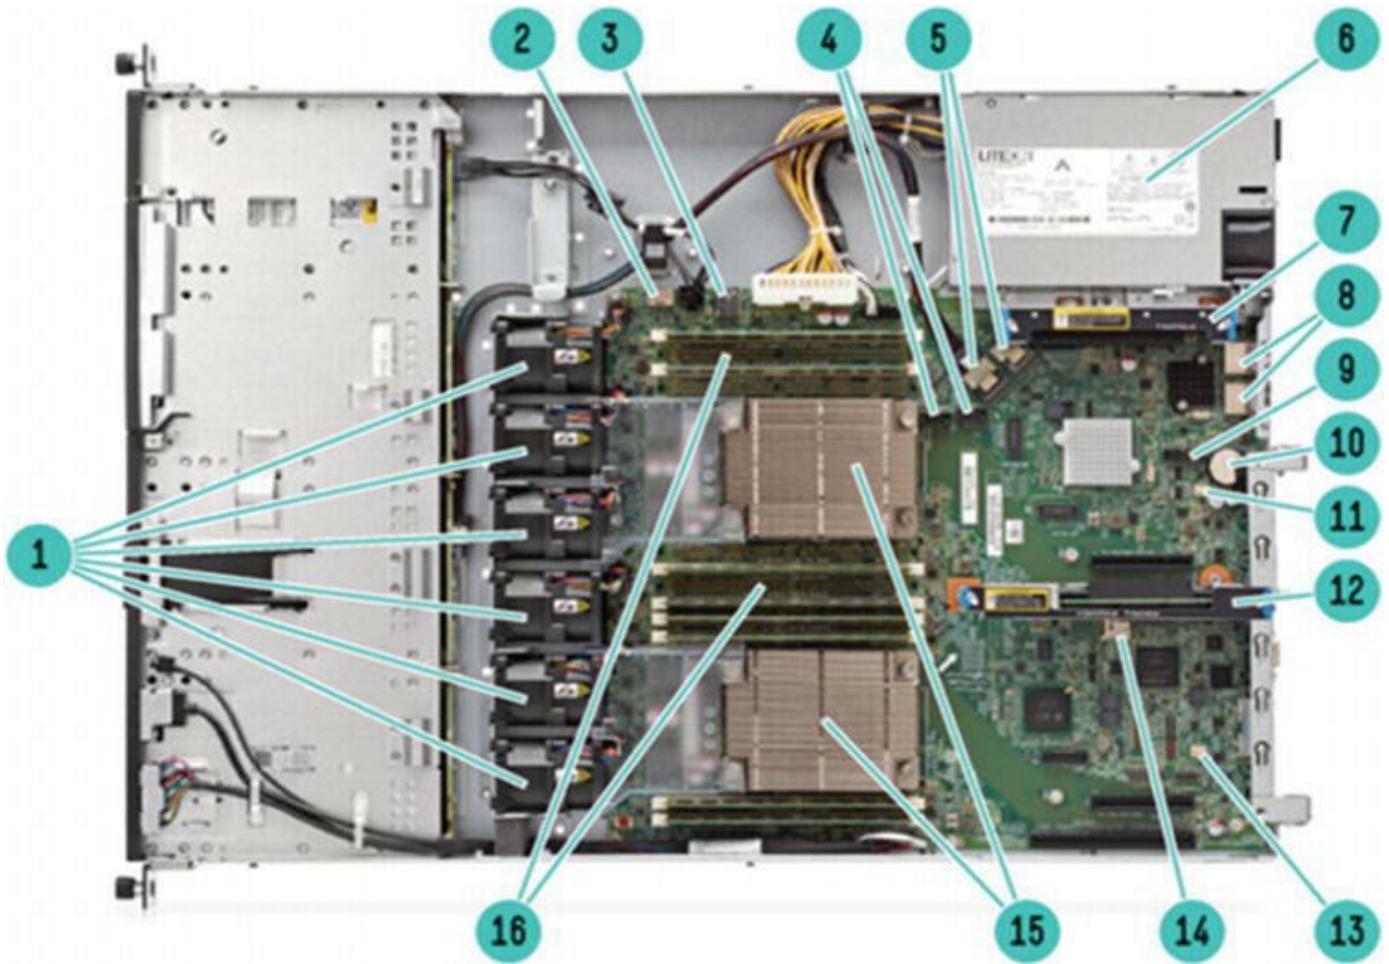


Figure 5-29 ProLiant DL60 Gen9 Server, internal view

Key features of the HPE ProLiant DL60 Gen9 Server shown in [Figure 5-29](#) are:

- Hot-swappable fans (up to six)
- HPE Smart Storage Battery connector
- Internal USB 3.0 connector
- SATA connector
- Mini-SAS connectors
- Power supply bay
- Primary PCIe riser board connector
- Embedded 2 × 1 GbE network adapter
- TPM connector
- 10. System battery
- 11. Storage backup power connector for PCIe Slots 1 and 2
- 12. Secondary PCIe riser slot
- 13. Storage backup power connector for PCIe Slot 3
- 14. MicroSD slot
- 15. Processor heat sinks (processors with HPE Smart Socket placed below heat sink)
- 16. DDR4 DIMM slots (up to eight)

HPE ProLiant DL80 Gen9 Server

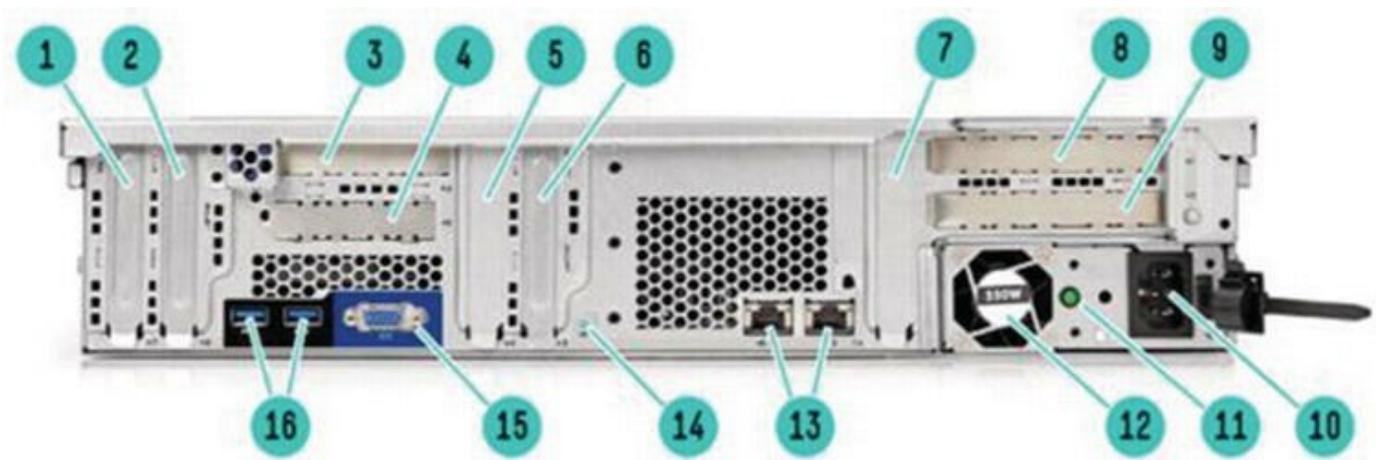


Figure 5-30 HPE ProLiant DL80 Gen9 rear view

The HPE ProLiant DL80 Gen9 Server is designed with affordable storage, networking and processor scalability, and manageability in a 2U server for cost-conscious service provider and SMBs. Apart from basic IT infrastructure workloads, this server is ideal for running storage capacity-intensive applications such as cold storage and backups.

Key features of the HPE ProLiant DL80 Gen9 Server shown in [Figure 5-30](#) are:

- Slot 7 PCIe3 x16 low-profile, standup
- Slot 6 PCIe3 x8 low-profile, standup
- Slot 4 PCIe3 x8 low-profile, riser (optional)
- Slot 5 PCIe3 x8 FlexibleLOM riser (optional)
- Slot 4 PCIe3 x16 low-profile, standup
- Slot 3 PCIe3 x8 low-profile, standup
- Slot 1 PCIe3 x16 low-profile, standup
- Slot 1 PCIe3 x16 full-height, full-length GPU riser (optional) or Slot 1 PCIe3 x8 (8, 4, 1) full-height, half-length riser (optional)
- 9. Slot 2 PCIe3 x8 full-height half-length riser (optional)
- 10. Power supply power connection
- 11. Power supply status LED
- 12. Non-hot-plug power supply
- 13. Embedded 2 × 1 GbE network adapter
- 14. UID LED
- 15. Video connector
- 16. USB 3.0 connectors

The storage controller on the DL80 Gen9 Server is one of the following depending on the model:

Entry models—Dynamic Smart Array B140i Controller

Base models—H240 FIO Smart Host Bus Adapter

HPE SimpliVity 380 Server



Figure 5-31 HPE SimpliVity 380 Server

For many years, the traditional IT model has been to build out separate towers of technology, covering servers/virtualization, storage, and networking. Hyperconverged infrastructure (HCI) has changed the landscape and the thinking around IT deployment. Hyperconverged systems, like the HPE SimpliVity 380 shown in [Figure 5-31](#), deliver all parts of a typical IT infrastructure out of a single set of hardware and software that performs the previously segregated tasks of storage, networking, and compute. At an operational level, hyper-converged infrastructure simplifies the deployment, ongoing management, and refresh of hardware, while moving to a more cloud-like consumption model. Additional capacity can usually be implemented by deploying extra hardware nodes and adding them to the existing configuration. The protection of data across the hardware is managed by software, providing graceful recovery from component or server failure.

Hyperconverged infrastructure provides significant benefits:

Data efficiency—Hyperconverged infrastructure reduces storage, bandwidth, and IOPS requirements.

Elasticity—Hyperconvergence makes it easy to scale out/in resources as required by the business demands.

VM-centricity—A focus on the VM or workload as the cornerstone of enterprise IT with compute, memory, storage, and networking revolving around individual VMs.

Data protection—Ensuring that data can be restored in the event of loss or corruption is a key IT requirement, made far easier by hyper-converged infrastructure.

VM mobility—Hyperconvergence enables greater application/workload mobility.

High availability—Hyperconvergence enables higher levels of availability than possible in legacy systems.

Cost efficiency—Hyperconverged infrastructure brings to IT as sustainable step-based economic model that eliminates waste.

The true power of hyperconvergence comes from full consolidation of software and hardware devices. Hyperconvergence is the ultimate in an overall trend of convergence brought to market in recent years. The aim is to bring simplicity to increasingly complex data centers.

Hyperconvergence solves common business problems

The built-in capabilities in hyperconverged infrastructure technology help organizations address a range of critical challenges for managing infrastructure.

Data center consolidation—Minimize data center footprint for lower upfront costs, greater operational efficiency, and improved performance.

Data protection and disaster recovery—Protect data in a fraction of the time required by legacy solutions with end-to-end data protection and automated disaster recovery.

VDI scalability—Scale up VDI with unmatched performance. Get more virtual desktops on less hardware and globally backup, restore, or clone a VM in three clicks or less.

VM vending—Deploy VMs in a few clicks to add capacity with customizable, repeated templates for the provisioning of those machines.

Eliminate storage complexity—Employing Direct-Attached Storage (DAS) in place of complex shared storage solutions. Distributing storage IO across two or more cluster nodes for high-speed availability and recovery.

Hyperconvergence offers significant lower CapEx as a result of lower upfront prices for infrastructure, lower OpEx through reductions in operational expenses and personnel, and faster time-to-value for new business needs. In the context of IT management, newly emerging infrastructure generalists—people with broad knowledge of infrastructure and business needs—can easily support hyperconverged infrastructure. Organizations no longer need to maintain separate islands of resource engineers to manage each aspect of the data center.

HPE SimpliVity combines the entire infrastructure into a simple, flexible building block to reduce the cost and complexity of the IT environment and deliver the technology organizations need.

HPE SimpliVity architecture

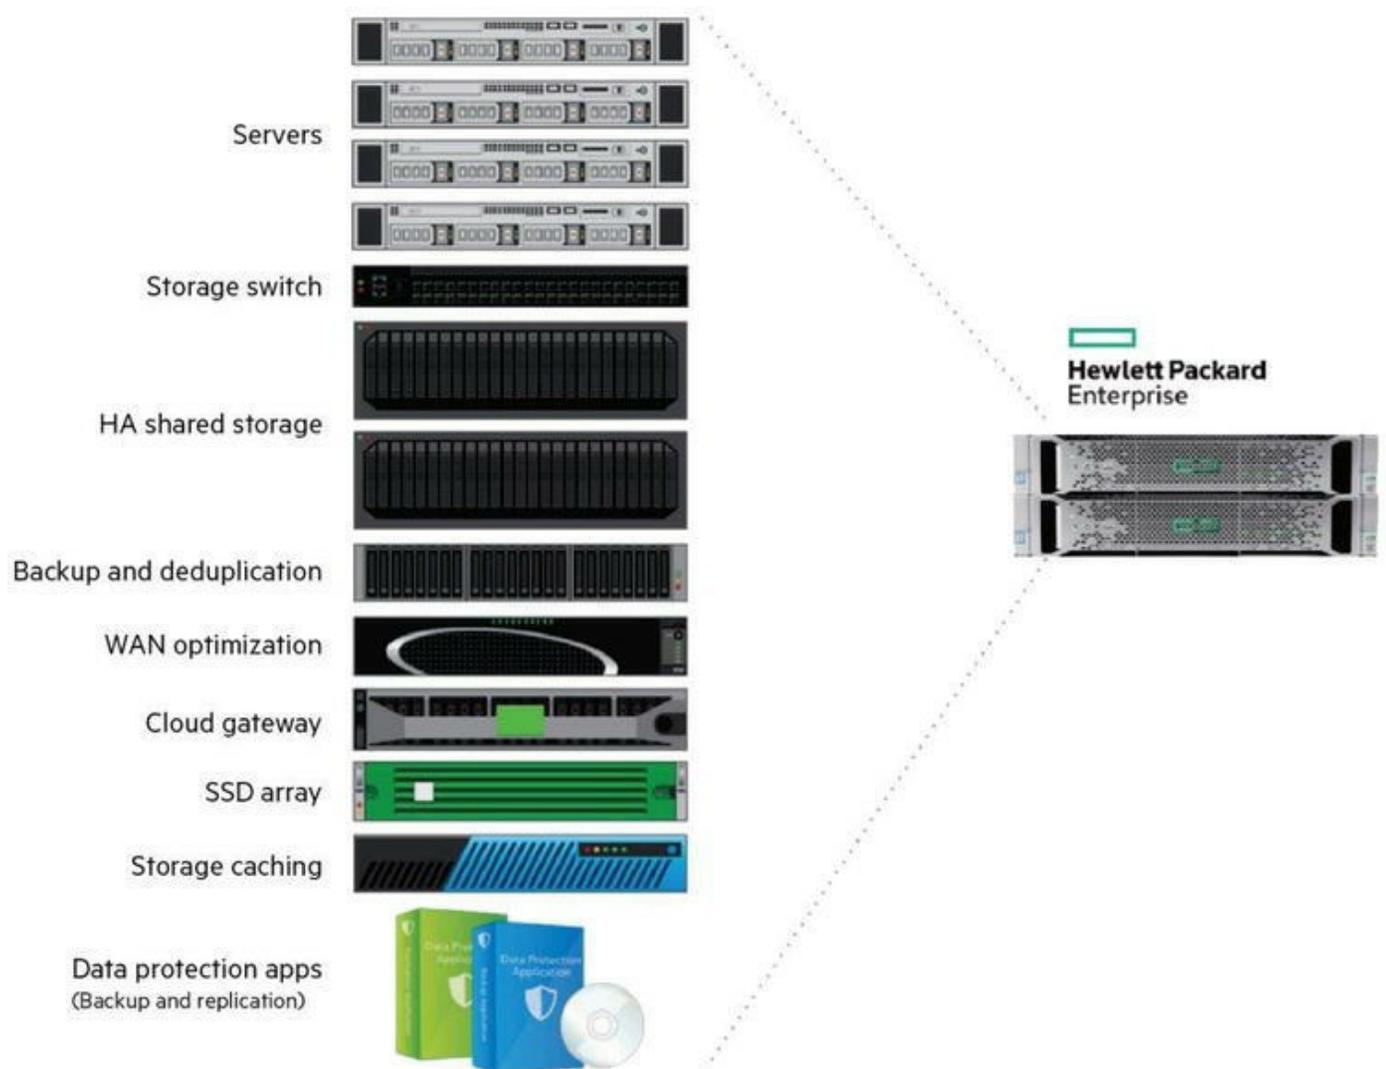


Figure 5-32 HPE SimpliVity building blocks

HPE SimpliVity is a software-defined hyperconverged infrastructure solution that runs on x86 industry-standard servers to deliver turnkey hyper-converged infrastructure for the software-defined data center. Clustering multiple HPE SimpliVity nodes form a shared resource pool that delivers high availability, mobility, and efficient scaling of performance and capacity.

As shown in [Figure 5-32](#), HPE SimpliVity technology assimilates all IT infrastructure and services below the hypervisor into a single, scalable 2U building block. Two or more deployed clusters of HPE SimpliVity building blocks form a federation—a massively scalable pool of shared resources.

HPE SimpliVity building blocks

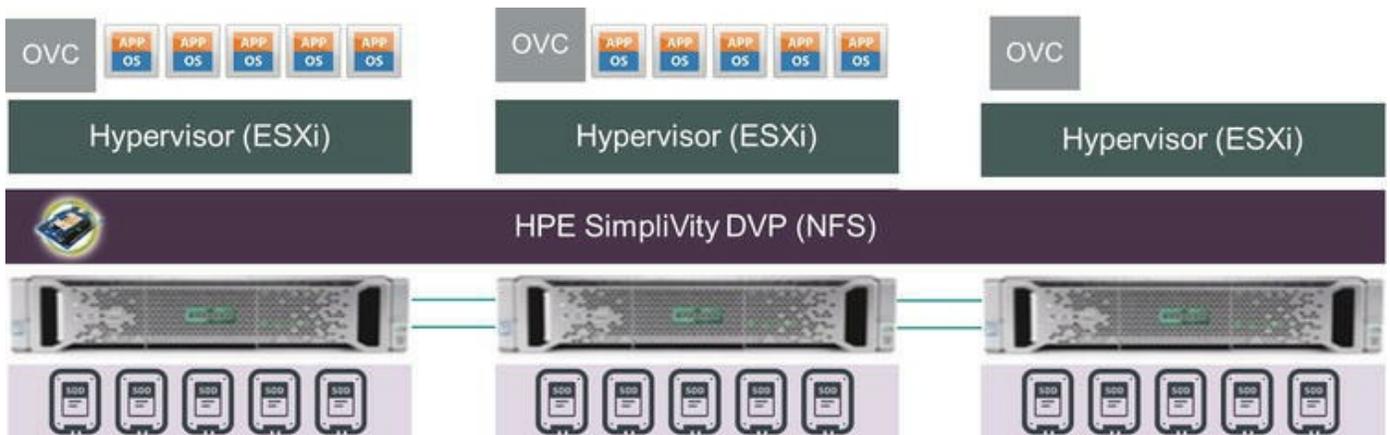


Figure 5-33 HPE SimpliVity building blocks

HPE differentiation from other hyperconverged infrastructure is the HPE SimpliVity Data Virtualization Platform (DVP), which enables a single shared resource pool across multiple sites and provides highly efficient data storage and mobility. [Figure 5-33](#) shows the basic building blocks of a hyperconverged solution. HPE SimpliVity DVP consists of the HPE OmniStack Virtual Controller and the HPE OmniStack Accelerator Card, which enables accelerated data efficiency, global unified management, and built-in data protection.

HPE SimpliVity DVP features

Guaranteed data efficiency	Built-in resiliency, backup, and disaster recovery	Global VM-centric management and mobility
<ul style="list-style-type: none"> • Always-on compression and deduplication • All data at inception, globally • Offloaded to HPE OmniStack Accelerator • Guaranteed 90% capacity savings across primary storage and backup* 	<ul style="list-style-type: none"> • Full logical backups with near zero overhead • Guaranteed 60-second restore of 1TB VM* • Granular RTOs and RPOs from hours to seconds • Simple, affordable offsite DR 	<ul style="list-style-type: none"> • Policy-based, VM-centric management • No LUNS, shares, or volumes • Right-click operations • Native tool integration • Single view of all data centers and ROBOs

Figure 5-34 Features of the HPE SimpliVity DVP

[Figure 5-34](#) shows key features of the HPE SimpliVity DVP.

HPE SimpliVity data architecture

The HPE SimpliVity data architecture is a globally aware file system and object store with data optimization techniques and data services. It abstracts data from its underlying hardware, enabling hypervisor, and x86 hardware independence. It shifts policies and management from an infrastructure component orientation, usually LUNs or datastores, to the application workload and VM level.

Data efficiency

To foster mobility, data efficiency is key. Data deduplication and compression facilitate lighter-weight, mobile data structures. Optimization further helps to manage data intelligently based on the anticipated usage of data by the applications.

Deduplication

Deduplication is the process of reducing disk workload, replication bandwidth, and storage capacity by eliminating redundant data and retaining only one unique instance of the data on disk. Deduplication helps to deliver a lightweight, mobile data structure and, therefore, is key to solving the complexity crisis by addressing the root cause.

Compression

Compression is the process of reducing the size of a given data set, providing a more condensed format. It provides complementary value to deduplication by further reducing the disk workload, replication bandwidth, and storage capacity.

Optimization

Optimization is the intelligent treatment of data on disk based on its anticipated use by an application. Systems that can identify file types and make real-time decisions about where to store that data can achieve overall improved storage efficiency, performance, and bandwidth usage.

Deduplication, compression, and optimization have several key benefits that address the core requirements of today's data center.

Dramatic bandwidth reduction—Deduplicating, compressing, and optimizing of all data

and keeping it in this state when transferring to a remote site makes the transfer itself more efficient and saves significant bandwidth resources.

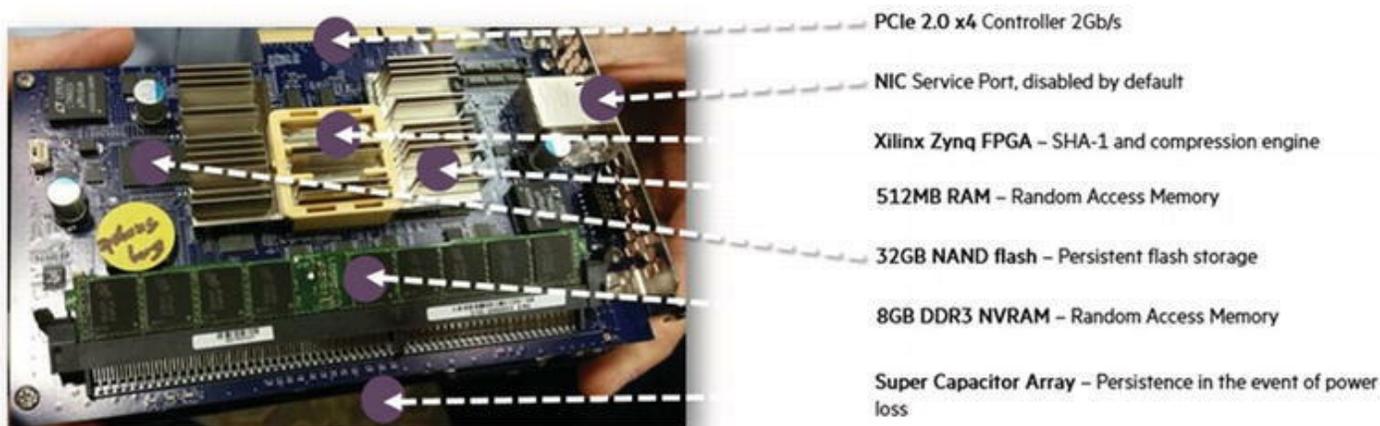
Enhanced data mobility—A fundamental principle of server virtualization is the mobility of VMs, but the unoptimized data structures can significantly block mobility in a traditional infrastructure environment due to their size. When the data is deduplicated and compressed, it is easier to create point-in-time copies, move VMs from one system to another, and move data to different data centers across the globe.

Efficient storage utilization—Required capacity can be reduced significantly in standard primary use cases based on the effective use of deduplication, compression, and optimization.

Enhanced performance—Since data is deduplicated, compressed, and optimized in real time before data is written to disk, it helps eliminate unnecessary writes to disk that saves IOPS and improves overall performance.

SimpliVity deduplication is done at inception, once and forever. SimpliVity’s DVP delivers inline deduplication, but without a performance penalty. That is because the OmniStack Accelerator allows all of the “heavy lifting” to be offloaded from the host CPUs. This leaves as much CPU as possible available to run the business applications. The OmniStack Accelerator also delivers extremely predictable performance, which allows SimpliVity’s hyperconverged infrastructure to deliver predictable performance for business applications.

OmniStack Accelerator Card (Tiered Index Accelerator)



- Power loss triggers a save of TIA memory to on-board flash

Figure 5-35 OmniStack Accelerator Card (Tiered Index Accelerator)

HPE OmniStack Accelerator Card (Tiered Index Accelerator)

Rather than taking an existing data architecture and trying to build in deduplication, compression, and optimization, HPE took the inverse approach. As a first step, it designed the core technology that performs deduplication and compression on primary data, in real-time, without impact to performance or latency, and built an entire globally federated data architecture around that foundation. This foundation manages the resulting fine-grained data elements across a global federation of systems.

In doing so, it addressed most of the core requirements for truly effective deduplication, compression, and optimization for the primary production infrastructure system and beyond, including:

Executing in real-time

Processing once and maintaining this efficient state through the entire lifecycle of the data (no need for a hydration or dehydration inefficiencies)

Deduplicating across all tiers of data within a system

Deduplicating across all data sets

Maintaining the initial deduplicated state across all locations, including the public cloud

Executing without an impact on the performance

The HPE SimpliVity DVP performs deduplication, compression, and optimization in real-time, as the data is first written to the datastore. This is in contrast to post-process deduplication, which allows data to be written first without deduplication and at some later stage, performs the deduplication process. The big problem with post-process deduplication is that it uses more IOPS since data is written, read, deduplicated/compressed, and written again.

The HPE OmniStack Accelerator Card, as shown in [Figure 5-35](#), handles the heavy lifting, delivering the required processing power without the high costs. It is a uniquely architected PCIe module that processes all writes and manages the compute-intensive tasks of deduplication and compression and allows the x86 CPUs to run customers' business applications.



Note

It is important to note that HPE OmniStack Accelerator Card never contains the only copy of a piece of data within a data center. HPE SimpliVity technology can handle the failure of an entire hyperconverged building block and can continue to serve data at a level of performance that allows mission-critical enterprise applications to continue to run.

HPE Intelligent Series Racks



Figure 5-36 HPE Intelligent Series Racks

The HPE Intelligent Series Rack, shown in [Figure 5-36](#), is designed to meet the current and future requirements of demanding data centers. It offers innovative intelligence capabilities for asset management, unparalleled structural integrity, cooling and cable management advances, and a wide choice of power and switching options.

The HPE Intelligent Series is the first rack family to offer optional location discovery intelligence for advanced asset management. This capability integrates with iLO 4 and iLO 5 management and HPE Intelligent Power Distribution Units (iPDUs) to communicate the U space, row, and rack location of each server. It also helps you manage and place workloads by carefully mapping current demands and available resources.

All Intelligent Series Racks are industry-standard EIA-310 19 inch-RETMA. They support all HPE rack-mountable equipment as well as third-party equipment designed to mount in 19-inch wide racks.

HPE Intelligent Network Racks are designed for dense network equipment. These racks have the front vertical rails moved back 75 mm to facilitate front to rear cabling and have additional bristle-covered cable pass-through slots in the front and on top of the racks to prevent mixing of hot and cold air.

These racks are available in standard 600 mm and 800 mm widths. The 800 mm wide Intelligent Network Racks include slotted vertical cable channels on both sides in the rear of the rack. Intelligent Air Duct Racks work with the HPE Rack Air Duct Kit to direct hot exhaust air from the rear of the rack to air returns in the ceiling. These racks are perfect for racks up to 14 kW where the mixing of hot and cold air in the data center needs to be eliminated to improve efficiency without the cost of additional fans.

Learning check

You are at a customer meeting and one of the IT team comments that there is a belief that hyperconverged solutions performance is poor due to the CPU bound deduplication and compression functionality. They are considering the HPE SimpliVity 380 solution. How would you address their performance concerns?

Write a summary of the key concepts presented in this chapter.

Summary

ProLiant servers are based on industry-standard x86 architectures.

ProLiant DL Gen10 rack servers incorporate embedded components that provide a rich feature set in a compact chassis size. They are ideal in dense, space-constrained, and multi-server environments for customers who want enterprise server features occupying a smaller footprint, enhanced built-in security, manageability, reliability, and serviceability.

HPE SimpliVity technology assimilates all IT infrastructure and services below the hypervisor into a single, scalable 2U building block. The HPE SimpliVity 380 solution offers unparalleled scale, simplicity and single shared resource pool across multiple sites and provides highly efficient data storage and mobility. The massive storage scale offered by the SimpliVity 380 solution coupled with the HPE OmniStack Accelerator Card, enables accelerated data efficiency, global unified management, and built-in data protection.

6 High Performance Computing Solutions

WHAT IS IN THIS CHAPTER FOR YOU?

After completing this chapter, you should be able to:

Explain the importance of High-Performance Computing (HPC).

Describe the portfolio of HPE HPC.

Explain the Apollo 2000 HPC solution.

Describe the Apollo 10 series HPC solutions.

Prelearning check

At a trade fair, a customer approaches you to explain that they are a small startup company doing research in developing organic battery technologies for the motor car industry. They have a limited budget to invest in the next generation of HPC. They are interested in technologies supplied by HPE. How should you respond?

You are attending a meeting with the IT procurement manager of a large research university. Attending the meeting is the head of a small department performing medical research for the next generation of antibiotics. They already have an HPE Apollo 2000 solution, and they are looking to upgrade from Gen9 servers and scale to a larger solution. How should you respond?

Supercomputing and HPC

Let us begin the discussion of HPC with some definitions.

What is supercomputing?

- Supercomputing efficiently solves extremely complex or data-intensive problems by concentrating the processing power of multiple, parallel computers.

Supercomputing definition

- Supercomputing involves a system working at the maximum potential performance of any computer, typically measured in Petaflops (where one Petaflop is one quadrillion [thousand trillion] or 10^{15} floating point operations per second).

Supercomputer use cases

- Sample use cases include deep learning (DL), genomics, astronomical calculations, simulation, life science research, and pharmaceuticals.

How does HPE approach supercomputing?

- HPE approaches supercomputing through an HPC architecture.
- HPE aims to overcome traditional cost barriers to drive a wider adoption of supercomputing.
- HPC solutions empower innovation at any scale, building on purpose-built HPC systems and technologies, applications, and support services.

If you are new to the concept of HPC, you are probably wondering what all the fuss is about. Just what is HPC?

The term HPC most generally refers to the practices of aggregating computing power in a way that delivers much higher performance than one could get out of a typical desktop computer, workstation, or server to solve large problems in science, engineering, or business.

Traditional proprietary Vector Supercomputers were hugely expensive, esoterically designed, physically large machines from companies such as Control Data Corporation and Cray, designed to work on mankind's biggest problems in science and engineering. Applications include origins of the universe, new cancer drugs, genomics, and advanced simulation. Most of these esoteric supercomputer manufacturers have either failed or have been acquired by other companies.

Organizations that rely on these solutions for research or commerce now consider Commercial-Off-The-Shelf (COTS) solutions as a possible replacement for the traditional supercomputer. The benefits include: lower cost, flexibility, short

development cycle, scale, and performance. Some research organizations such as Lawrence Livermore Labs successfully built their own COTS solution as a proof of concept. This kind of solution is not without its challenges, including:

Coordinating booting tens or hundreds of servers

Managing a cluster containing tens or hundreds of servers

Providing high-speed data exchange between cluster nodes

Coordinating distributed data processing

Providing the space, power, and cooling necessary to accommodate the hundreds or even thousands of servers

Today the situation is changing. Single-chip, multi-core CPUs are many times more powerful compared to supercomputers of the recent past. Many more companies are experiencing huge benefits from recent advances in high-speed networks and improved microprocessor performance. These advances mean that clusters or arrays of servers are becoming an appealing and cost-effective vehicle for parallel computing. Clusters, built using commodity hardware and software components, are now playing a major role in redefining the concept of supercomputing.

HPE approaches supercomputing through HPC architecture. HPC makes it possible to overcome traditional cost barriers to supercomputing. You can choose how much compute power you want to concentrate in HPC clusters. HPE HPC solutions empower innovation at any scale, building on purpose-built HPC systems and technologies solutions, applications, and support services.



Note

Learn more about how HPE approaches supercomputing and how these solutions are benefiting key applications by entering the following URL into your browser:

<https://www.hpe.com/uk/en/solutions/hpc-high-performance-computing.html>

HPE purpose-built portfolio for HPC

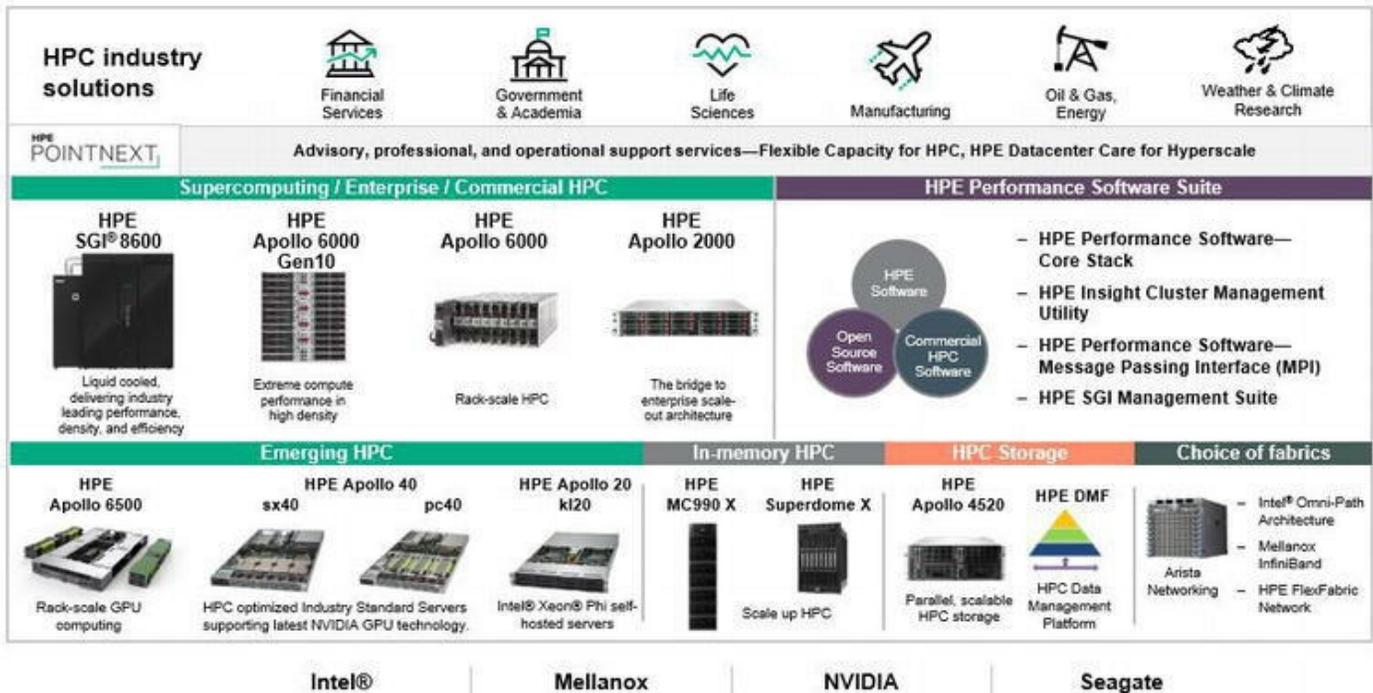


Figure 6-1 HPE purpose-built portfolio for HPC

As shown in [Figure 6-1](#), the HPE portfolio for HPC includes HPE SGI (from the recent acquisition of Silicon Graphics), HPE Apollo 20, 40, 2000, 4000, and 6000, HPE MC990 X, HPE Superdome X, and HPE Superdome Flex (not shown). In this chapter, we focus on the HPE Apollo 10 series (20 and 40) and the HPE Apollo 2000.

The HPE Apollo family is a set of purpose-built compute offerings designed to address the needs of HPC workloads as well as Big Data Analytics, Object Storage, and Enterprise Service Provider needs, as shown in [Figure 6-1](#). The emerging scale-out type of architectures driven by the cloud for Big Data Analytics and Object Storage require simple storage at massive scale, substantial configuration flexibility, and a high degree of density optimization. The HPC type of workloads requires significant parallel processing performance in combination with increased data center infrastructure efficiency to maximize the performance per unit of data center resource consumption (power/cooling/space/administrators). As the need for HPC moves beyond the traditional domain of academia and research institutions to mainstream enterprise applications to drive business innovation, customers need the capability to shift from traditional platforms to scale-out solutions and HPC solutions in a nondisruptive manner to better serve their business needs.

The HPE Apollo family supports a broad range of HPC deployments from small, entry

configurations in the enterprise with Apollo 2000 to the most complex supercomputing applications with the HPE SGI 8600.

HPE Apollo 2000 is the ideal enterprise bridge for customers who are embarking on the scale-out journey for the enterprise and need a solution for small deployments of HPC as well as general purpose workloads in a nondisruptive manner. The Apollo 6000 is a more powerful solution that delivers rack-scale efficiencies and is designed for midsize to large HPC deployments with a higher level of infrastructure efficiency.

HPE Apollo 4000 systems are purpose-built for large-scale deployments of software-defined object and clustered storage, analytics, or active archives.

HPE Apollo 6000 delivers HPC capabilities for the next level of processing requirements in midsize and large enterprise with rack-scale efficiencies in terms of power, cooling, and space together with targeted workload optimization to enhance performance.

HPE SGI 8600 is a liquid-cooled supercomputer that can provide more than 600 teraflops of pure Intel Xeon Processor Scalable Family compute performance per rack, and grow to tens of thousands of nodes with minimal increase in system overhead. It is ideal for very complex numerical simulations, DL, and mathematical modelling applications that require massively parallel computing capabilities.

HPE Apollo 2000—Enterprise bridge to HPC



Figure 6-2 Flexible scale-out architecture in a traditional 2U form factor

As organizations move to the Idea Economy, even the most conservative enterprise customers are looking for ways to save space and become more efficient. Think about insurance companies. Many of them now provide their customers with apps, and to deal with the huge compute and storage demands required to support them, these companies are looking for nondisruptive ways to introduce more efficient shared infrastructure into their existing data centers. The Apollo 2000 delivers twice the density of traditional rack mount systems and the efficiency of a shared infrastructure but maintains a familiar form factor, the same racks, cabling, serviceability access, operations, and system management. No retraining of personnel or cost of change for introducing efficient, space saving, and scale-out architecture.

As shown in [Figure 6-2](#), the Apollo 2000 System is a density-optimized, 2U shared infrastructure chassis for up to four ProLiant Gen9 or Gen10 independent, hot-plug servers with all the traditional data center attributes—standard racks and cabling and rear-aisle serviceability access. A 42U rack can hold up to 20 Apollo r2000 series chassis accommodating up to 80 servers per rack.

The Apollo 2000 System offers the configuration flexibility to support a variety of workloads, from remote site systems to large HPC clusters and everything in between. It can be deployed cost-effectively starting with a single 2U, shared-infrastructure chassis to meet the configuration needs of a wide variety of scale-out workloads.

With Apollo 2000 System servers, there is flexibility to tailor the system to the precise needs of each workload with computer and flexible IO and storage options. Apollo 2000 System servers can be “mixed and matched” within a single chassis to support different applications, and it can even be deployed with a single server, leaving room to scale as the customer’s needs grow.

In addition, the Apollo 2000 chassis comes with four new generation single rotor fans and an additional four fans can be added for redundancy. The power can be managed by the Advanced Power Manager (APM) and optional rack-level manager. HPE APM automatically discovers hardware components and enables bay level power on and off, server metering, and aggregate dynamic power capping. HPE APM consolidates Ethernet access to all resident iLOs, and provides a secure single sign-on. HPE APM also provides asset management capabilities for shared infrastructure.

HPE Apollo 2000—r2800 Chassis—24 SFF



Figure 6-3 HPE Apollo r2800 Chassis

The Apollo 2000 System brings HPE ProLiant XL170r and XL190r Gen10 server technology, including iLO 5, into this 2U, multi-server chassis. It is the follow-on product to the SL2500.

The HPE ProLiant XL170r Gen10 Server and the HPE ProLiant XL190r Gen10 Server offer more configurations choices that cover a much wider range of scale-out workloads than previous generations. Storage and IO flexibility enable customers to optimize for performance or economy—the right compute for the right workload.

The r2800 Chassis shown in [Figure 6-3](#), with 24 SFF drives, allows customer to choose how they allocate the hard drives across the server nodes. It also provides the following:

Up to four expansion slots in the XL190r support accelerators or other full-size cards

Mix and match trays to build a unique solution or partially populate leaving room for growth in the future

Configuration flexibility for variety of workloads

- Mix and match servers for workload optimization
- HPC performance with accelerators, top bin CPUs, fast HPC clustering

- Storage flexibility and a broad range of IO options for workload optimization

Density-optimized for traditional data centers

- Up to four powerful servers in 2U chassis—2X the density of 1U servers
- Traditional racks and cabling for existing data centers

Cost-effective in any configuration

- Simple at scale—it is ProLiant Gen10
- Same ProLiant enterprise-class management and operational tools
- HPE iLO 5 management saves administration time and cost
- APM enables more efficient capacity per rack
- HPE Insight CMU to monitor, manage, and optimize compute clusters of any size

Apollo 2000 system offerings

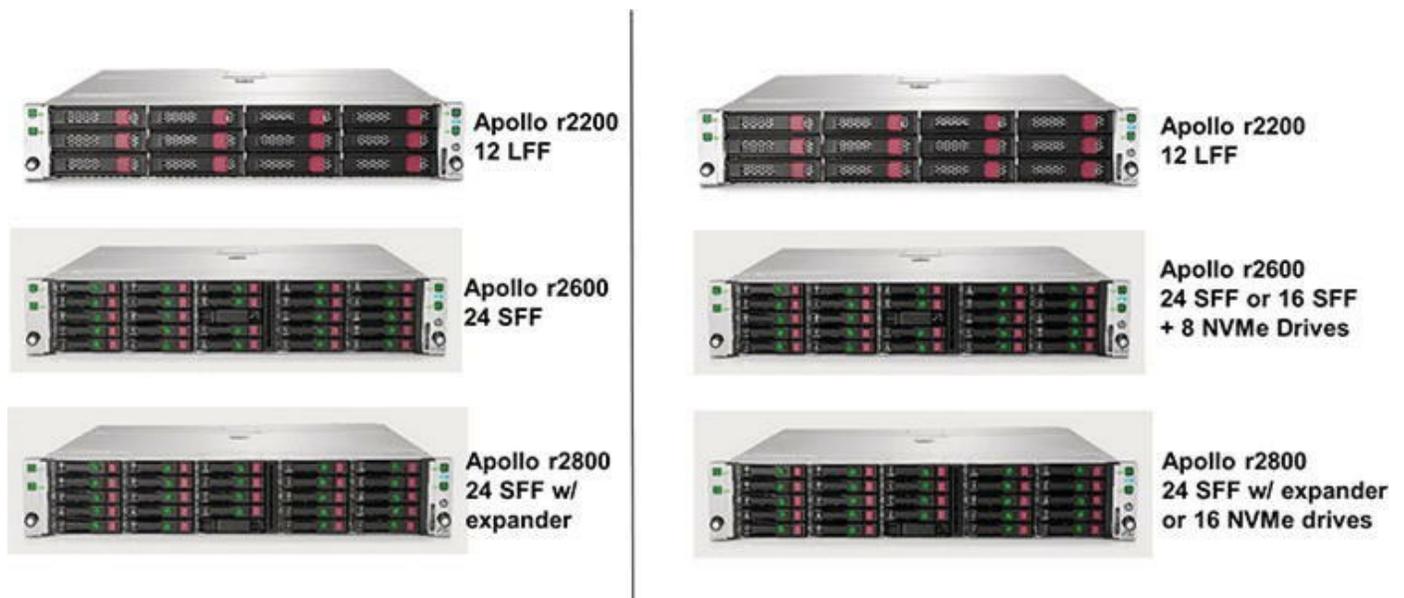


Figure 6-4 Apollo 2000 system offerings

Apollo 2000 supports up to four servers in 2U and offers three different enclosures to address different needs, as shown in [Figure 6-4](#).

Apollo r2200 with 12 LFF—If the focus is low-cost, high capacity, and storage is the key, go with Apollo r2200 large form factor (12) 3.5” hot-plug SAS, SATA HDD, or SSDs allocated equally across server nodes. Current capacity of LFF drives is 10 TB.

Apollo r2600 with 24 SFF—If storage and workload performance is the focus, go with Apollo 2600 small form factor, 24 SFF 2.5” hot-plug SAS, SATA HDD, or SSDs allocated equally across servers. Current capacity in SFF is up to 3.84 TB. Alternatively, if ultimate performance is required, the 16 SFF + 8 NVMe drive option may be the best choice.

Apollo r2800 chassis—If the customer likes to control the number of hard drives that are allocated to each server, SFF with SAS expander supports flexible drive mapping, enabling custom drive allocations to match workloads. An individual server can have from zero to 24 drives allocated. For ultimate performance, the 16 NVMe drive option may be the best choice.

The Rack Consolidation Management (RCM) module allows you to aggregate the iLO at the chassis level. You can daisy chain them and connect just one cable to the Top of Rack (ToR) management switch, eliminating the need of multiple iLO cables.

HPE Apollo 2000—r2800 Chassis with four 1U nodes

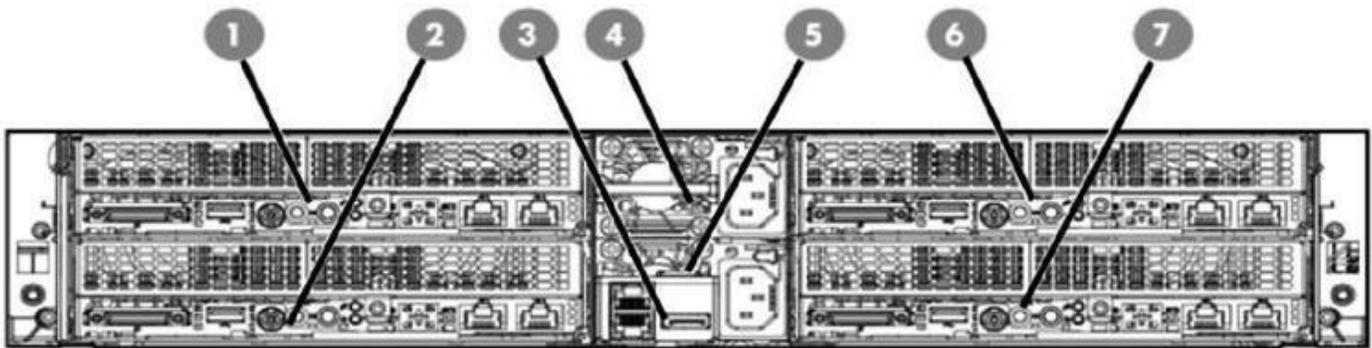


Figure 6-5 Apollo 2000—r2800 Chassis with four 1U nodes

[Figure 6-5](#) shows the major components of the r2800 chassis with four 1U nodes. The numbered components are:

Node 4

Node 3

RCM module

Power supply 2

Power supply 1

Node 2

Node 1

HPE Apollo 2000—r2800 Chassis with two 2U nodes

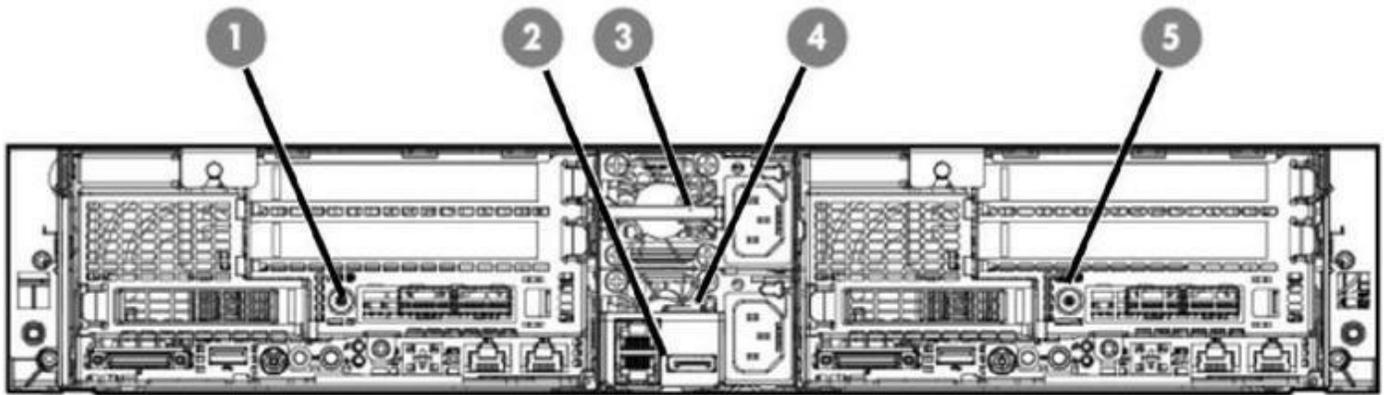


Figure 6-6 Apollo 2000—r2800 Chassis with two 2U nodes

[Figure 6-6](#) shows the major components of the r2800 chassis with two 2U nodes. The numbered components are:

Node 3

RCM module

Power supply 2

Power supply 1

Node 1

HPE Apollo 2000—r2800 chassis fans

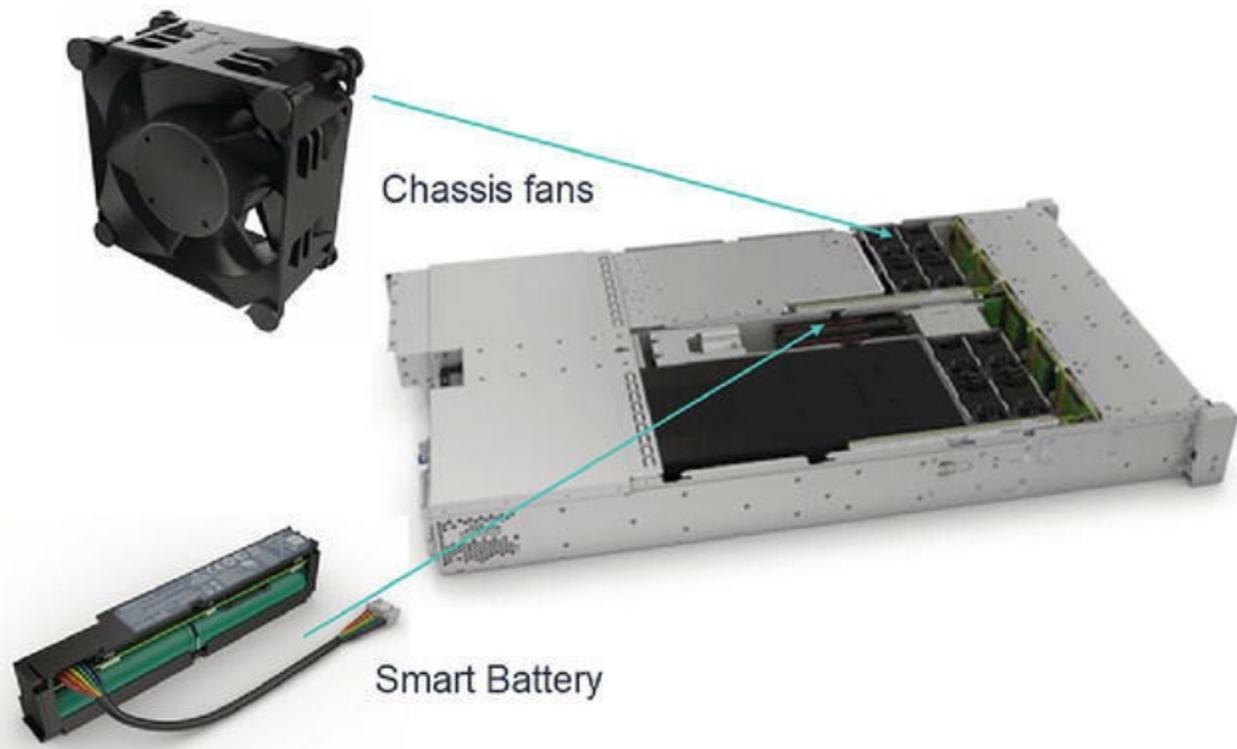


Figure 6-7 r2800 chassis fans and smart battery

[Figure 6-7](#) shows the location of the r2800 chassis fans and HPE Smart Battery.

The enclosure comes standard with 10 hot-plug cooling fans to provide the best overall cooling for the entire enclosure under many conditions, options, and workloads.

The HPE Smart Battery provides central battery backup for up to 24 devices in case of power outage, including the Flash Backed Write Cache module of HPE Smart Array Controllers.

HPE Apollo 2000—r2800 RCM Module



Figure 6-8 r2800 RCM Module

The RCM Module kit, shown in [Figure 6-8](#), allows iLO aggregation at the chassis level, and also has a port for the HPE APM. The RCM Module consists of 2 iLO ports for redundancy. It can also be daisy chained to connect to a ToR management switch.

HPE ProLiant XL170r and XL190r Server—Gen 10

Chassis	r2200, r2600 and r2800 chassis—New NVMe Backplane for r2800 chassis (4 per 1U node)
Processor	Xeon® Scalable Family series processor
Memory	16 DDR4 DIMMs, 2666 MHz HPE SmartMemory
IO slots	2 I/O slots for a choice of networking and clustering options including : 1 GB, 10 GB, 100 GB Ethernet, 100 Gb/s EDR InfiniBand, 100Gb/s Omni-Path, Fibre Channel and options for either 1 PCIe (x16) slot plus a FlexibleLOM or 2 PCIe (x16) slots
Storage	Dual SATA host based M.2 2280 SSDs-internal
Storage controllers	HPE Smart Array E208i-p, E208e-p, P408e-p, or P408i-p SR Gen10 Controller
Management	HPE iLO 5 (Each server will have its own dedicated iLO) HPE Apollo Platform Manager (rack level mgmt)



Figure 6-9 ProLiant XL170r and XL190r features

[Figure 6-9](#) highlights some of the features of HPE ProLiant XL170r and XL190r servers. Features include:

HPE Apollo r2000 Series Chassis accommodates four independently serviceable HPE ProLiant XL170r Gen10 Servers with up to 80 servers in a 42U rack with one to four servers per chassis.

HPE Apollo r2000 Series Chassis accommodates two independently serviceable HPE ProLiant XL190r Gen10 Servers with up to 40 servers and 80 integrated accelerators in a 42U rack with either one or two servers per chassis.

ProLiant XL170r and XL190r Gen10 Servers have the same configuration options for CPU and memory. The main difference between the two models is that the XL190r has more PCIe slots providing support for additional expansion cards.

HPE Apollo 2000—XL170r Gen10



Density-optimized general purpose cloud server with all-in-one high-performance computing/storage

Figure 6-10 ProLiant XL170r Gen10 Server

The Apollo r2000 series chassis accommodates four independently serviceable ProLiant XL170r Gen10 servers with up to 80 servers in a 42U rack with one to four servers per chassis. The ProLiant XL170r Gen10 Server shown in [Figure 6-10](#) is a 1U half-width, two-processor server. Rack scale power efficiency is achieved with HPE Apollo Platform Manager for rack, chassis, server, and component-level management.

The ProLiant XL170r Gen10 Server configuration options include:

Performance and efficient CPUs—Intel[®] Xeon[®] Scalable Family series processor

options ranging from Bronze to Platinum with choices from four to 26 cores, 1.7 GHz to 3.6 GHz CPU speed, and power ratings between 85 and 150 watts

16 memory DIMM slots, DDR4 2666 MT/s HPE SmartMemory

Two IO slots for a choice of networking and clustering options including: 1 Gb, 10 Gb, 100 Gb Ethernet, 100 Gb/s EDR InfiniBand, 100Gb/s Omni-Path, Fibre Channel, and options for either one PCIe (x16) slot plus a FlexibleLOM or two PCIe (x16) slots

HPE Apollo 2000—XL170r Gen10: Top view

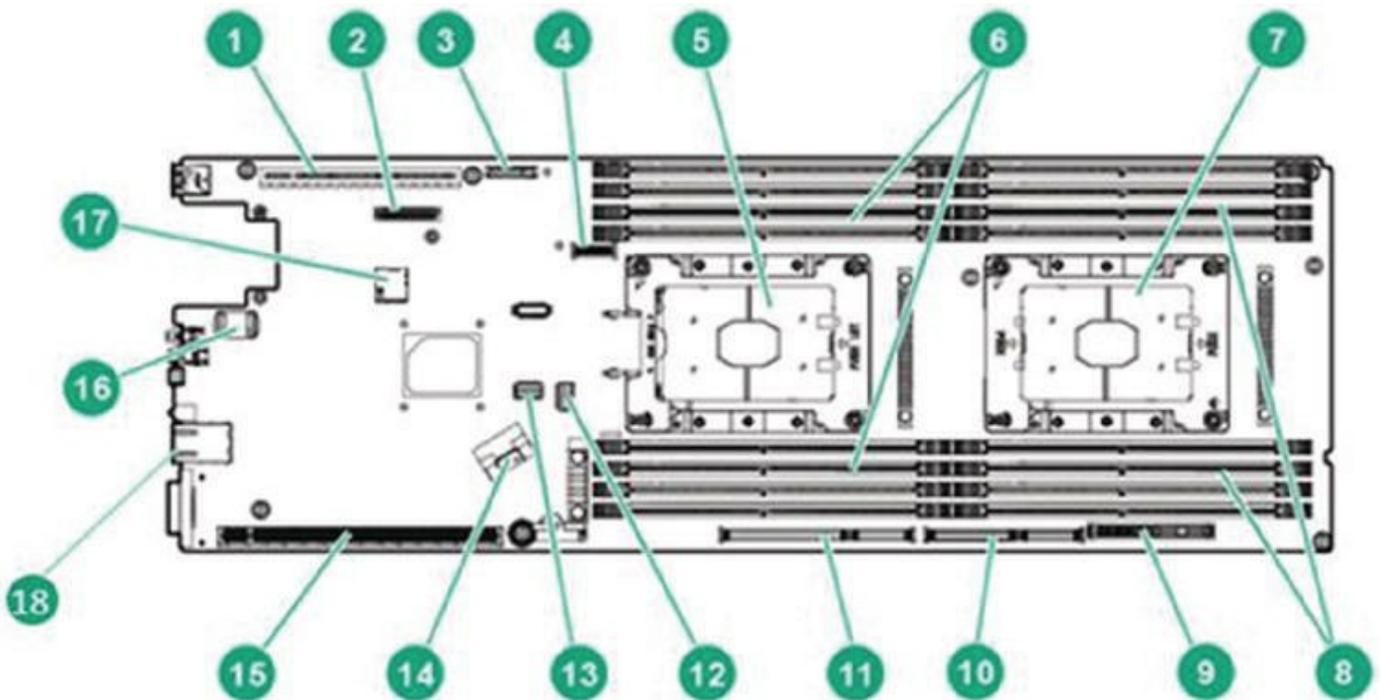


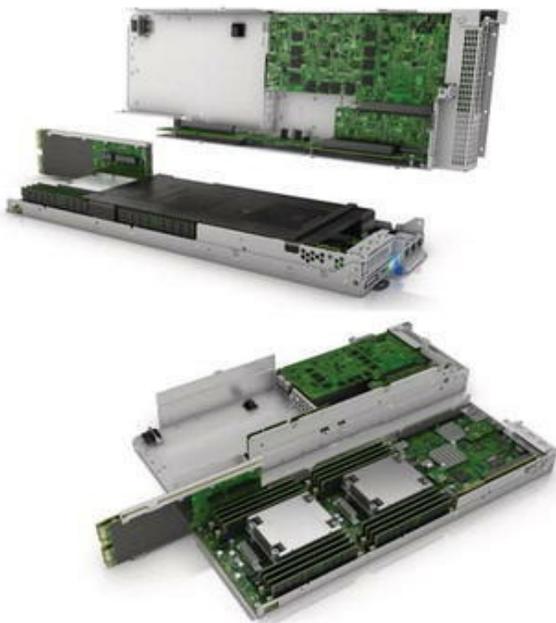
Figure 6-11 HPE Apollo 2000—XL170r Gen10 top view

[Figure 6-11](#) shows the major components of the XL170r Gen10 server. The numbered components are:

Top view:

1. Primary riser slot 1
2. mLOM connector
3. System battery
4. M.2 SSD riser connector
5. Processor 1
6. DIMMs for processor 1
7. Processor 2
8. DIMMs for processor 2
9. Bayonet board slot
10. Secondary riser slot 4
11. Secondary riser slot 3
12. Slimline connector (SATA x4)
13. System maintenance switch
14. Slimline connector (SATA x8)
15. Secondary riser slot 2
16. TPM connector
17. microSD slot
18. iLO dedicated port

HPE Apollo 2000—XL190r Gen10



Density-optimized all-in-one high-performance computing/storage server equipped with high-performance computing (GPUs or coprocessors) for enterprises, ISP, and HPC

Figure 6-12 ProLiant XL190r Gen10 Server

The ProLiant XL190r Gen10 Server shown in [Figure 6-12](#) is a 2U half-width, two-processor server with the same configuration options as the XL170r for CPU and memory, but has additional PCIe slots in multiple configurations providing support for additional expansion cards and two integrated accelerators per server (GPUs or coprocessors).

This server leverages Intel's® latest Xeon Scalable Family series processors options ranging from Bronze to Platinum, increasing performance up to 40% over previous generations, and DDR4 HPE SmartMemory, at speeds of up to 2666 MHz and 1.5 TB maximum which boosts bandwidth and efficiency up to 66% over previous generation servers.

The dense and flexible HPE Apollo 2000 Chassis can accommodate up to two 2U hot-swap server trays to meet various workload needs while lowering TCO.

Rack scale power efficiency is fuelled by the optional HPE Apollo Platform Manager for rack, chassis, server, and component-level management.

HPE Apollo 2000—XL190r Gen10: Top view

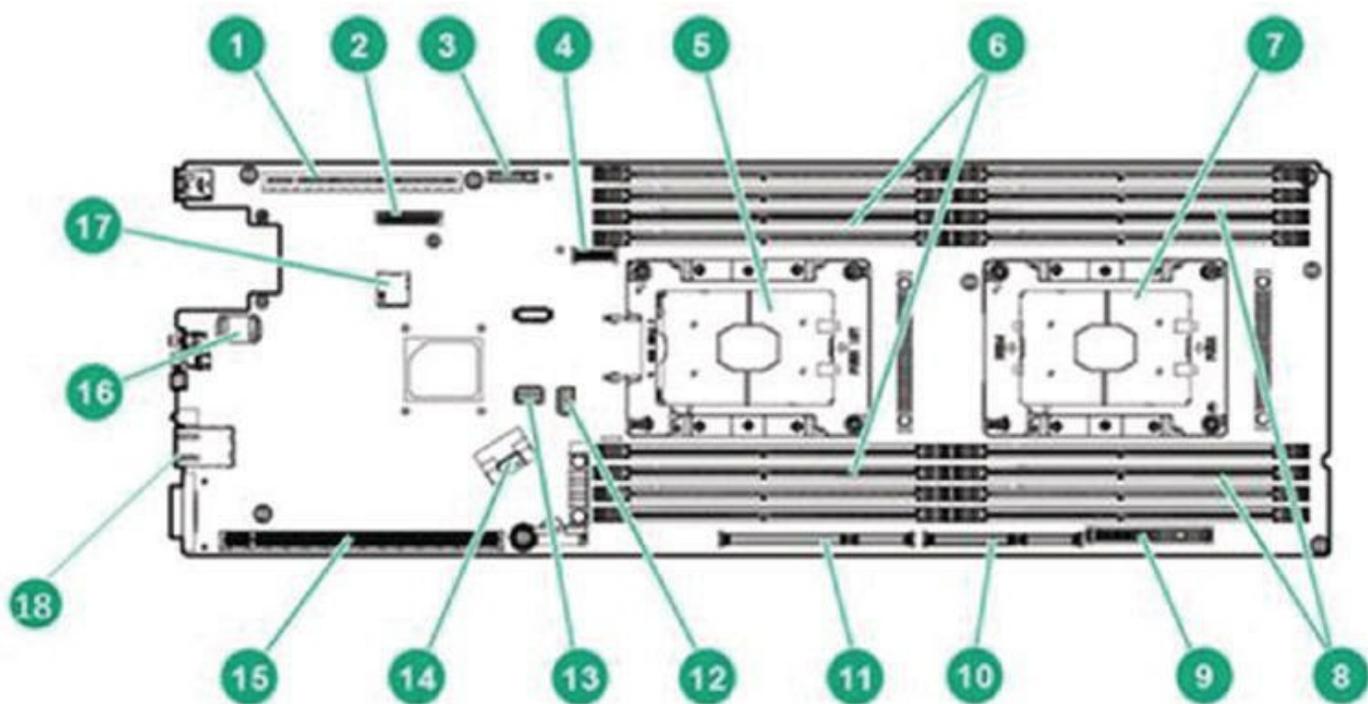


Figure 6-13 HPE Apollo 2000—XL190r Gen10 top view

[Figure 6-13](#) shows the major components of the XL190r Gen10 server. The numbered components are:

Top view:

1. Primary riser slot 1
2. mLOM connector
3. System battery
4. M.2 SSD riser connector
5. Processor 1
6. DIMMs for processor 1
7. Processor 2
8. DIMMs for processor 2
9. Bayonet board slot
10. Secondary riser slot 4
11. Secondary riser slot 3
12. Slimline connector (SATA x4)
13. System maintenance switch
14. Slimline connector (SATA x8)
15. Secondary riser slot 2
16. TPM connector
17. microSD slot
18. iLO dedicated port

The Apollo 2000 chassis can accommodate two of the XL190r trays; however, you can mix and match XL170r and XL190r, if necessary.

Learning check

In the HPE Apollo r2600 24 SFF chassis, how are the 24 disks allocated to each of the XL170r and the XL190r Servers?

You are attending a meeting with the IT procurement manager of a large research university. Attending the meeting is the head of a small department performing medical research for the next generation of antibiotics. They already have an HPE Apollo 2000 solution, and they are looking to upgrade from Gen9 servers and scale to a larger solution. How should you respond?

An HPE ProLiant XL170r provides two IO slots for a choice of networking and clustering options which includes Ethernet 1 Gb, 10 Gb, and 100 Gb. What are the alternatives? (Select two.)

- A. 100 Gb/s EDR InfiniBand
- B. Myrinet cluster interconnect
- C. Quadrics QsNet network
- D. Scalable Coherent Interconnect
- E. 100Gb/s Omni-Path

HPE Apollo 10 Series

Modern scientific discovery requires very high computing power and the capability to deal with huge volumes of data. Businesses are also increasingly relying on the power of supercomputers to invent innovative solutions, reduce cost, and decrease time to market for products and services. However, the nature of computing is changing with an increasing number of data-intensive critical applications: by 2020, it is estimated that 25 billion devices will be connected and will generate over two zettabytes of traffic every year.

Data is the most valuable currency in the Idea Economy, where success favors the businesses that get their ideas to market the fastest. Data-driven organizations achieve significant gains by implementing continuous improvement and predictive planning to improve customer satisfaction, quickly inform big decisions, eliminate waste, and reduce risk.

This flexible, density-optimized system is ideal for HPC, Big Data Analytics, service providers, and enterprises. It also supports a wide range of configurations for just about any workload.

As HPC extends its reach across many types of industries, it also extends to solutions of many sizes. Today, about half of the HPC market still involves systems that cost more than \$500,000, but the other half is devoted to midsized and small HPC deployments. While enterprises recognize the value of HPC in sharpening their competitive edge, every dollar spent on the solution must justify itself.

Compute power density is the key to delivering HPC with an attractive TCO. A dense solution reduces power and cooling bills, as well as the physical footprint. This last factor plays a crucial role in saving the customer money. A data center has a finite size, but demands for more compute power are always increasing. Expanding to a new data center can be prohibitively expensive, so any solution that can pack more power into a smaller space can save companies millions.

HPE Apollo 10 series makes HPC and DL affordable. Low entry costs combine with a flexible industry-standard platform, custom solutions, and a choice of components. Apollo 10 series servers readily integrate into standard cluster environments.

HPE Apollo 20 Family—kl20

Server	kl20—2U 4 servers
Power	Two 2000W platinum level
Processor	Four Intel® Xeon® Phi™ Processor (1 per server) (230W)
Memory	24 DDR4 DIMMs, 2400 MHz (6 per server)
Network	Dual-port GigE controller (Intel® I350)
IO slots	Eight PCIe 3.0 x16 low-profile (2 per server)
Storage	12 x LFF drives (3 per server)
Management	Integrated IPMI 2.0



Figure 6-14 HPE Apollo kl20 Server with Intel® Xeon® Phi processor

Compact 2U industry-standard form factor houses four independent self-hosted Xeon® Phi nodes. Each Intel® Xeon® Phi has up to 72 cores and can deliver greater than five teraflops enabling greater than 20 teraflops of performance in 2U.

As shown in [Figure 6-14](#), the HPE Apollo kl20 delivers a dense and flexible computing platform optimized for HPC applications. Powered by the Intel® Xeon® Phi processor and with support for both Intel® Omni-Path and Mellanox InfiniBand interconnect fabric options, the Apollo kl20 provides the ultimate performance for compute-intensive, highly parallelized workloads that use large data sets. Apollo kl20 compute nodes are built for dense, HPC environments providing up to 272 cores, 1.5 TB of memory, and 96 TB of storage capacity in a compact 2U form factor to create best solutions for accelerating data-intensive computations, such as DL, life sciences, or material sciences.

HPE Apollo 40 Family—sx40 Gen10

Server	sx40—1U 4 NVIDIA P100 SXM2
Power	Two 2000W titanium level
Processor	Intel® Xeon® Processor Scalable Family (165W); 3 UPI
Memory	12 DDR4 DIMMs, 2666 MHz
Network	Dual-port 1 GigE controller (Intel® I350)
IO slots	Two PCIe 3.0 x16 from PLX Two PCIe 3.0 x16 from CPU
Storage	2 x SFF SATA drives
Management	Integrated IPMI 2.0



Figure 6-15 HPE Apollo sx40 Server with Intel® Xeon® processor

The HPE Apollo sx40 system, shown in [Figure 6-15](#), is a 1U dual socket server featuring up to four NVIDIA® Tesla® GPUs in SXM2 form factor and based on the Intel® Xeon® Processor Scalable Family. The HPE Apollo sx40 Server is an optimized industry-standard server supporting DL and HPC workloads, using the SXM2 form factor to provide increased available GPU memory bandwidth and performance. Enjoy quicker and easier deployments with systems configured, integrated, and tested by HPE. Using industry-standard IPMI 2.0 system management and the HPE Performance Software Suite, the server is easy to use and manage. The HPE Apollo sx40 Server delivers cost-effective DL and HPC performance with high-compute density (four GPUs per 1U rack space), helping to turn data into insights and insights into actions.

HPE Apollo 40 Family—pc40 Gen10

Server	Pc40—1U 4 NVIDIA Tesla P100 SXM2 GPUs
Power	Two 2000W titanium level
Processor	Intel® Xeon® Processor Scalable Family (165W); 3 UPI
Memory	12 DDR4 DIMMs, 2666 MHz
Network	Dual-Port 1GigE controller (Intel® I350)
IO slots	Four PCIe 3.0 x16 for GPUs Two PCIe 3.0 x16 expansion slots
Storage	2 x SFF SATA drives
Management	Integrated IPMI 2.0



Figure 6-16 HPE Apollo pc40 Server with Intel® Xeon® processor

The HPE Apollo pc40 Server shown in [Figure 6-16](#) is an optimized, industry-standard server for DL and HPC workloads that uses the PCIe form factor for greater choice in GPU selection. Enjoy quicker and easier deployments with systems configured, integrated, and tested by HPE. Using industry-standard IPMI 2.0 system management and the HPE Performance Software Suite, the server is easy to use and manage. The HPE Apollo pc40 Server delivers cost-effective DL and HPC performance with high-compute density (up to four GPUs per 1U rack space), helping to turn data into insights and insights into actions.

The HPE Apollo pc40 Server features up to four NVIDIA Tesla PCIe graphic processing units (GPUs) to accelerate mixed-application HPC computing as well as DL workloads. The PCIe form factor can allow for a greater selection of accelerators and GPUs than alternate form factors.

Optimize the server for required workloads by choosing from the available CPUs in the Intel® Xeon® Processor Scalable Family, up to twelve 2666 MHz DDR4 DIMMs, optional network adapters, and up to two SFF hard drives or solid-state drives.

Fully integrated and factory-tested clusters are delivered with a complete cluster management software stack to simplify deployment and management and can be custom configured for perfect right-sizing.

Multiple network topologies are supported in cluster builds: Fat Tree, Hypercube, and Enhanced Hypercube for increased flexibility.

Learning check

At a trade fair, a customer approaches you to explain that they are a small startup company doing research in developing organic battery technologies for the motor car industry. They have a limited budget to invest in the next generation of HPC. They are interested in technologies supplied by HPE. How should you respond?

HPC use cases



Figure 6-17 Industries where HPC is critical

The description of HPC can sound abstract, but the problems that HPC solves are quite real. [Figure 6-17](#) shows some of the industries where HPC is critical. A weather model predicts rain, enabling you to choose to take your umbrella to avoid a soaking. HPC no longer belongs to large research facilities; enterprises across many verticals have recognized that to out-compute is to out-compete. By embracing HPC, they can bring better products to market more quickly.

Electronics continue to make vast strides in smartphones and mobile computing become smarter, cars become more efficient and connected, and hardware manufacturers pack more and more power into smaller packages. Engineers could not keep up this pace without the aid of computers themselves. Manufacturers use computer-aided engineering and electronic design automation to simulate and design better chips.

In healthcare, HPC helps researchers' model systems at the level of the ecosystem or the molecule. Pharmaceutical researchers use HPC to design drugs that are safer and more effective. Scientists use HPC to develop greener ways to run the world—from more efficient solar chips to new types of batteries.

You see the benefit of HPC in almost every movie, where CGI effects convince us to believe in the unbelievable. HPC also has a home in the music industry, in which HPC is helping to produce better quality.

Financial institutions require HPC to inform decisions such as where to invest money or which types of loans to make. This type of HPC is called Monte Carlo simulation—a simulation that informs decisions that are influenced by many variables, some of them random. Monte Carlo simulations are not just about finance. A retail company might need to choose the best location to open a new branch or select the right goods in which

to stock the shelves. A software company might need to decide how much to devote to developing a particular project. An aircraft manufacturer might need to simulate the effect that a lightning strike will have on the carbon fiber composite building material.

The race to develop the next-generation battery technology is a prime example of out-compute to out-compete. Successive approximation (also known as trial and error) resulted in Sony taking 10 years to develop the first commercial lithium-ion battery. Depletion of fossil fuels is driving massive research into new battery technology. The race is on to produce fast charging, lightweight, high power, and safe low-cost batteries. Start-up companies can join the search for the right combination of materials and manufacturing technique to produce future batteries. Cost efficient, high-performance HPE HPC solutions are the ideal partner.

Why GPU and coprocessor acceleration?

Designed for accelerating the embarrassingly parallel tasks common to many HPC applications

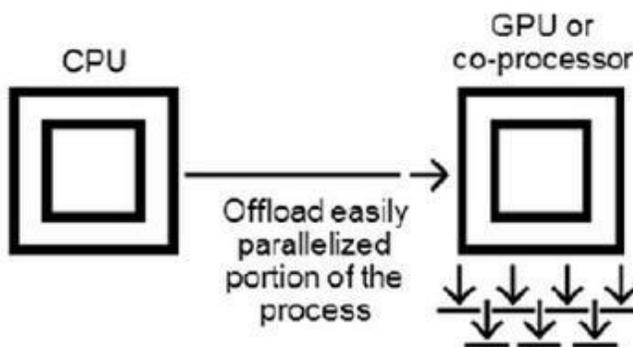


Figure 6-18 CPUs and GPUs were designed for different purposes

As shown in [Figure 6-18](#), CPUs were designed to meet the needs of many different types of workloads, including single-threaded processes and multi-threaded ones. GPUs, on the other hand, were originally designed for just one purpose: rendering graphics. Rendering each pixel constituted one task, separate from other tasks. So GPUs were optimized for multi-threading, rendering as many pixels as possible in parallel. Many HPC applications also feature workloads that can be parallelized and divided into many threads. These applications benefit highly from running on a CPU that is enhanced with a GPU.

The NVIDIA GPUs can boost performance up to 10 times depending on the application. The K40 provides up to 1.4 TFLOPS, 12 GB memory, and 288 Gb/s memory

bandwidth, and the K80, up to 8.73 single-precision TFLOPS and 2.91 double-precision TFLOPs with NVIDIA GPU boost, 24 GB memory, and 480 Gb/s. The AMD S9150 provides 5.07 single-precision TFLOPS, 2.53 double-precision TFLOPS, 16 GB memory, and up to 320 Gb/s memory bandwidth.

Instead of GPU accelerators, you can install Intel® coprocessors. The coprocessor consists of a dense group of cores (60 for the Intel® Phi 5110P and 61 for the Intel® Phi 7120P) and generous memory and memory bandwidth (8 GB and 320 Gb/s for the 5110P, 16 GB and 352 Gb/s for the 7120P). These coprocessors, like GPUs, are also optimized for highly parallelized tasks and can speed those tasks with an additional 2.0 to 2.4 single-precision teraflops (TFLOPs) or 1.0 to 1.2 double-precision TFLOPs.

Before you choose accelerators, it is crucial that you verify that the customer’s application is architected to take advantage of that accelerator.

DL and GPU market—Projections

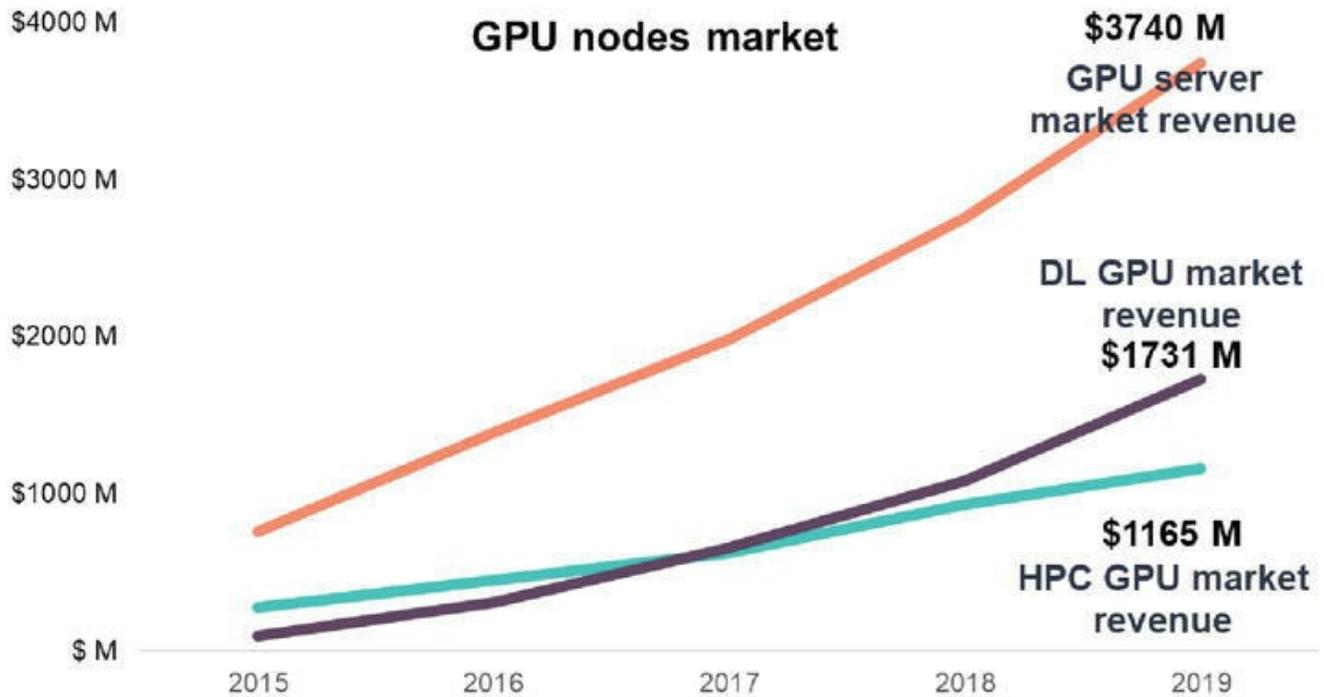


Figure 6-19 DL and GPU market projections

[Figure 6-19](#) shows projections for the DL and GPU market up to 2019.

It shows that:

GPU server market expanding rapidly to ~\$4B by 2019

- HPC market growth relatively modest
- DL passes HPC ~2017

DL providing deep insight into data and processes for all industry verticals

- Government: Autonomous systems, intelligence, and analysis
- Oil and gas: Automated seismic
- Manufacturing: Autonomous systems

GPU accelerators—HPC and DL

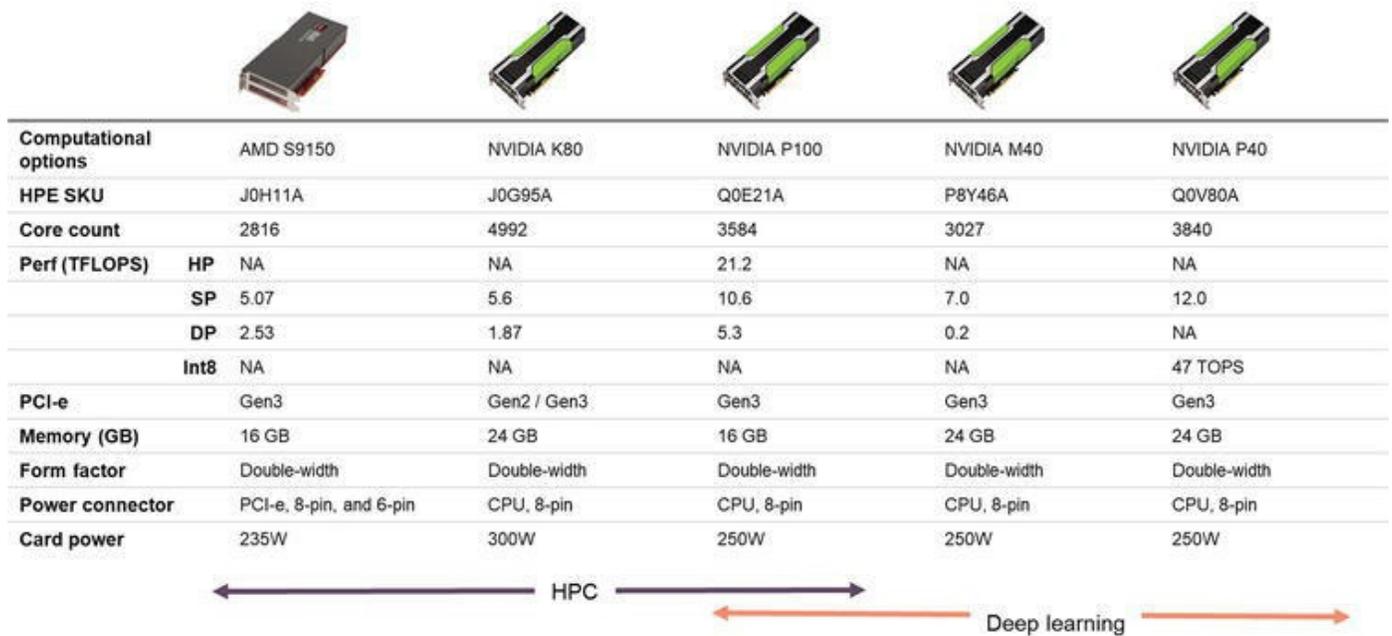


Figure 6-20 GPU accelerators

[Figure 6-20](#) shows the range of GPU accelerators available for HPC and DL.

Using HPE Apollo 2000 for storage

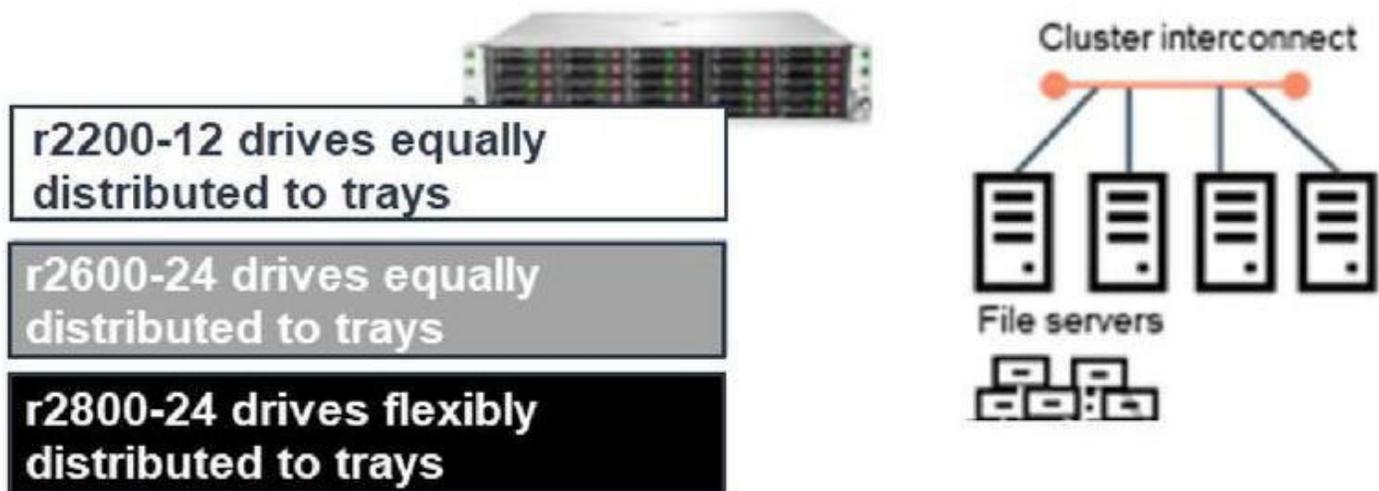


Figure 6-21 Apollo 2000 chassis provides a higher density of storage

An Apollo 2000 solution might be chosen to use only the local storage rather than shared storage. These HPC clusters tend to be relatively small, and as shown in [Figure 6-21](#), the Apollo 2000 chassis provides a relatively high density of storage, sometimes enabling it to meet the customer's needs.

Apollo 2000 storage sharing options:

Share the local storage in a smaller HPE Apollo 2000 cluster

Provide shared storage for HPE Apollo 6000 cluster (alternative to the Apollo 4000)

HPE Apollo 6000 compute trays have their own storage. The HPE 2000 chassis, on the other hand, provides the storage to installed trays.

HPE Apollo r2200 Chassis—Provides 12 LFF drives, equally distributed (three SATA or SAS HDDs or SDDs per XL170r or six per XL190r)

HPE Apollo r2600 Chassis—Provides six SATA or SAS HDDs or SDDs per XL170r or 12 per XL190r

HPE Apollo r2800 Chassis—Provides 24 drives like the r2600; however, you can choose how many you want to allocate to each server

- For example, you can select one or two servers to act as the file servers and assign all or most of the drives to them.

HPE Apollo 2000 networking choices

For highly parallelized, multi-threaded HPC applications, where jobs are shared across

many nodes in a cluster, the interconnections between compute nodes can act as a bottleneck, slowing down the computation. Properly provisioning the interconnections, on the other hand, will give the cluster a performance boost.

HPE Apollo 2000 XL170r and the XL190r IO options include Ethernet and InfiniBand. Which type is chosen will depend on whether the customer's application implements Message Passing Interface (MPI), a standardized and portable message-passing system designed to function on a wide variety of parallel computing architectures, or Remote Direct Memory Access over Converged Ethernet (RoCE) rather than Remote Direct Memory Access (RDMA) over InfiniBand.

Whenever HPC applications are highly parallelized such as with MPI, the interconnect must deliver high throughput and low latency. The InfiniBand options for the Apollo 2000 compute trays can provide 10 or 56 Gbps. InfiniBand also delivers extremely low latency and avoids the traditional IO stack and instead uses RDMA to connect nodes at the memory level, essentially extending internal fabric between nodes. However, if Ethernet is preferred, certain HPE FlexFabric adapters support RoCE, making Ethernet a good option for the interconnect. These HPE adapters also feature offloading of traffic processing, which prevents precious compute power being consumed by processing traffic. However, the chosen network switches must be able to support RoCE.

Questions to consider when choosing:

Does the customer have a preference for Ethernet or InfiniBand?

How parallelized is the HPC application?

- InfiniBand or RoCE-capable Ethernet (FlexFabric) for the best throughput and the lowest latency?

What availability requirements does the customer have? Are two NICs required for redundancy?

What type of shared storage solution is used?

- Will the interconnect also carry storage (use 10 GB rather than 1 GB)?
- Is Fibre Channel over Ethernet (FCoE) required?

Learning check

Write a summary of the key concepts presented in this chapter.

Summary

As organizations move to the New Style of IT, even the most conservative enterprise customers are looking to save space and become more efficient. The Apollo 2000 delivers twice the density of a traditional rack server.

The HPE ProLiant Gen10 XL170r and XL190r for HPE Apollo chassis offer more configuration choices that cover a wide range of scale-out workloads.

Modern scientific discovery requires very high computing power and the capability to deal with huge volumes of data. Businesses are also increasingly relying on the power of supercomputers to invent innovative solutions, reduce cost, and decrease time to market for products and services.

The HPE Apollo 20 Family servers deliver a dense and flexible computing platform optimized for HPC applications.

7 HPE BladeSystem Solutions

WHAT IS IN THIS CHAPTER FOR YOU?

After completing this chapter, you should be able to:

Explain how HPE BladeSystem solutions meet data center challenges.

Provide an overview of the currently shipping HPE ProLiant server blades.

Describe, at a high level, BladeSystem enclosures and their components, including:

- ✓ Onboard Administrator (OA)
- ✓ Insight Display
- ✓ Server blades

Summarize BladeSystem interconnect options.

Explain HPE OneView and HPE BladeSystem integration.

Prelearning check

A conversation with a customer reveals that they are considering an upgrade to their BladeSystem solution they had installed eight years ago. The data center edge has been upgraded to 20 Gb Ethernet and they are investing in a 16 Gb FC SAN. They ask you if they can upgrade the BladeSystem to support these new technologies. How should you respond?

A customer is rolling out a production environment that will consist of up to 30 c7000 enclosures. They are looking for a central management solution that will allow them to perform template-driven server deployment and enclosure lifecycle management. How should you respond?

HPE BladeSystem



Figure 7-1 HPE BladeSystem

As shown in [Figure 7-1](#), HPE BladeSystem is a modular infrastructure platform that converges server, storage, and network fabric to accelerate operations and speed delivery of applications and services running in physical, virtual, and cloud computing environments. The BladeSystem portfolio comprises server blades, workstation blades, enclosures, interconnects, and multiple storage options such as tape and storage blades.

A server blade is a self-contained server that fits into an enclosure with other blades. Sometimes known as a chassis, this enclosure provides the power, cooling, connectivity, and management to each server blade. The server blades contain only the core processing elements, making them hot-swappable. HPE refers to the entire package as a BladeSystem.

The core infrastructure is shared, meaning capital costs can be significantly lower. Blades share power, cooling, network, and storage infrastructure at the BladeSystem

enclosure level, so additional equipment is not needed for each server. This results in a dramatic reduction in the number of power distribution units (PDUs), power cables, LAN and SAN switches, connectors, adapters, and cables. The newest-generation technologies can be incorporated by simply changing the components that need to be changed.

HPE ProLiant Gen10 and Gen9 server blades offer a flexible, embedded storage controller strategy and controller-based encryption, as well as optional redundant microSD cards. ProLiant Gen10 and Gen9 server blades also provide improved storage options with USB 3.0 (internal) and increased performance capacity, compared with the previous generation of server blades.

Any of the server blades can be enhanced with a variety of mezzanine cards, including Ethernet, serial attached SCSI (SAS), Fibre Channel, and InfiniBand (IB) options. For each type of connection, HPE offers interconnect modules, such as HPE Virtual Connect modules. You can manage the entire infrastructure from a central location with HPE OneView.

BladeSystem meets data center challenges

BladeSystem solutions continue to meet the changing needs of data center customers by:

Allowing easier, faster, and more economical changes to server and storage setups without disrupting LAN and SAN domains

Lowering application deployment and infrastructure operations costs by reducing the number of IT architecture variants

Reducing connectivity complexity and costs

Lowering purchase and operations costs when adding or replacing compute and storage capacity

Supporting third-party component integration with well-defined interfaces, such as Ethernet NICs and switches, Fibre Channel host bus adapter (HBA) and switches, and IB host channel adapters (HCAs)

BladeSystem enables IT to:

Consolidate—A single modular infrastructure integrates servers, storage, networking, and management software that can be managed with a common, consistent user experience.

Virtualize—Pervasive virtualization enables customers to run any workload, meet high-availability requirements, and support scale-out and scale-up efforts. It also enables

businesses to create logical, abstracted connections to LANs and SANs.

Automate—Freeing up IT resources for more important tasks enables customers to simplify routine tasks and processes, saving time while maintaining control.

HPE BladeSystem—The ultimate converged infrastructure

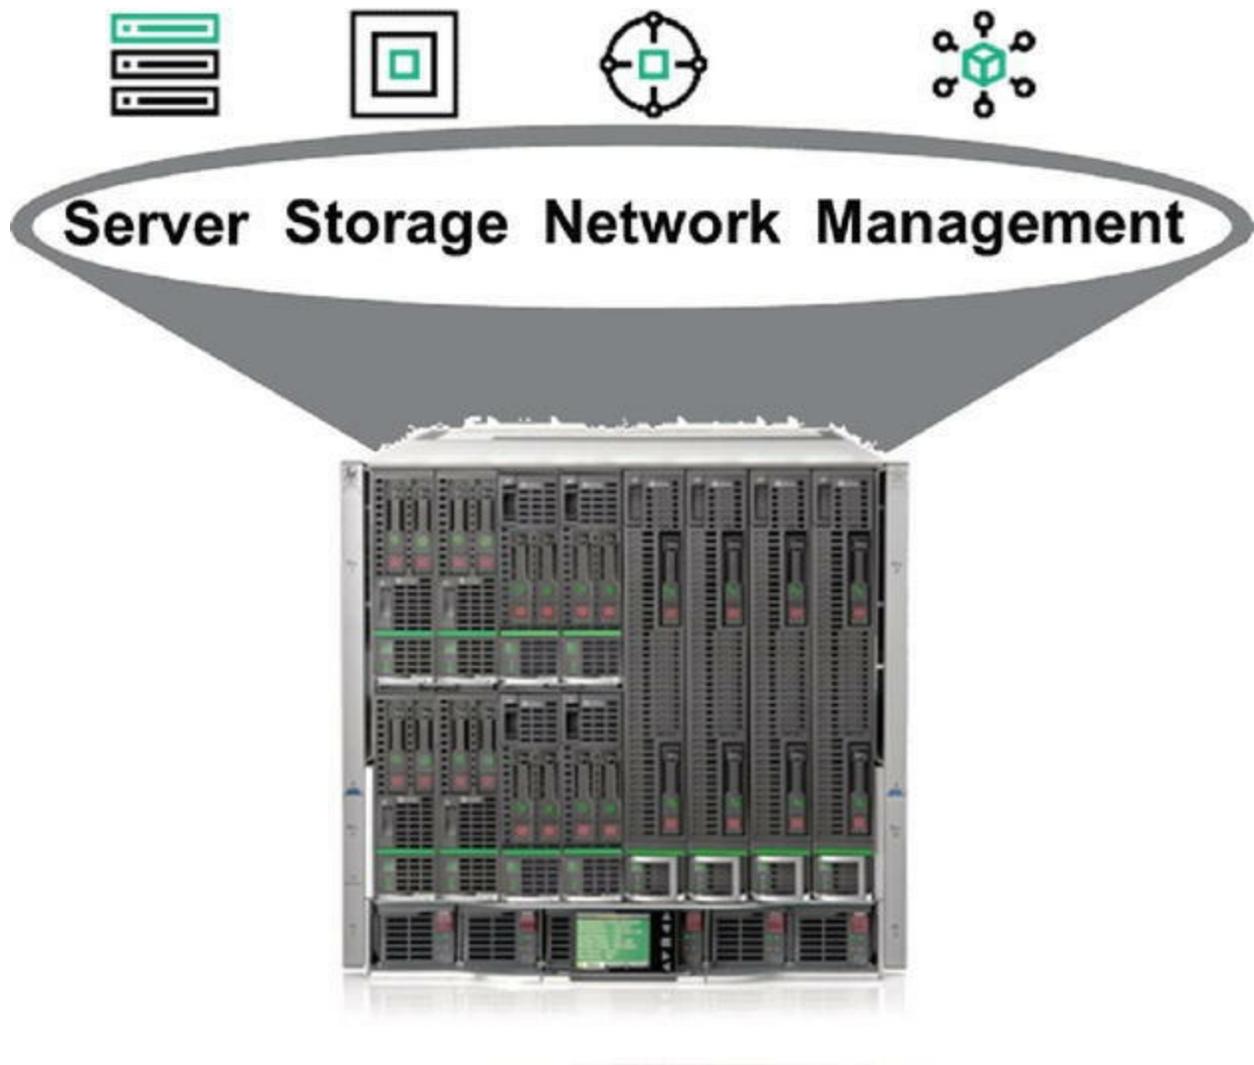


Figure 7-2 BladeSystems share servers, storage, network, and management

As shown in [Figure 7-2](#), BladeSystems share power, network, and storage infrastructure at the enclosure level. Because equipment is not needed for each server blade, there is a drastic reduction in PDUs, power cables, LAN and SAN switches, connectors, adapters, and cables.

You can introduce the newest-generation technologies by simply changing the components that need to be changed instead of the “rip and replace” approach used for traditional infrastructures. In addition, compared to a traditional environment, the time needed for routine server administrative tasks with BladeSystem can be reduced significantly. With HPE OneView, you can provision servers quickly, manage health proactively, control servers from anywhere, and manage power confidently.

Target industries and businesses best suited for BladeSystem include:

Financial

Telecom

Web

HPC

Hosting

Enterprise

SMB

BladeSystem Gen10 and Gen9 solutions

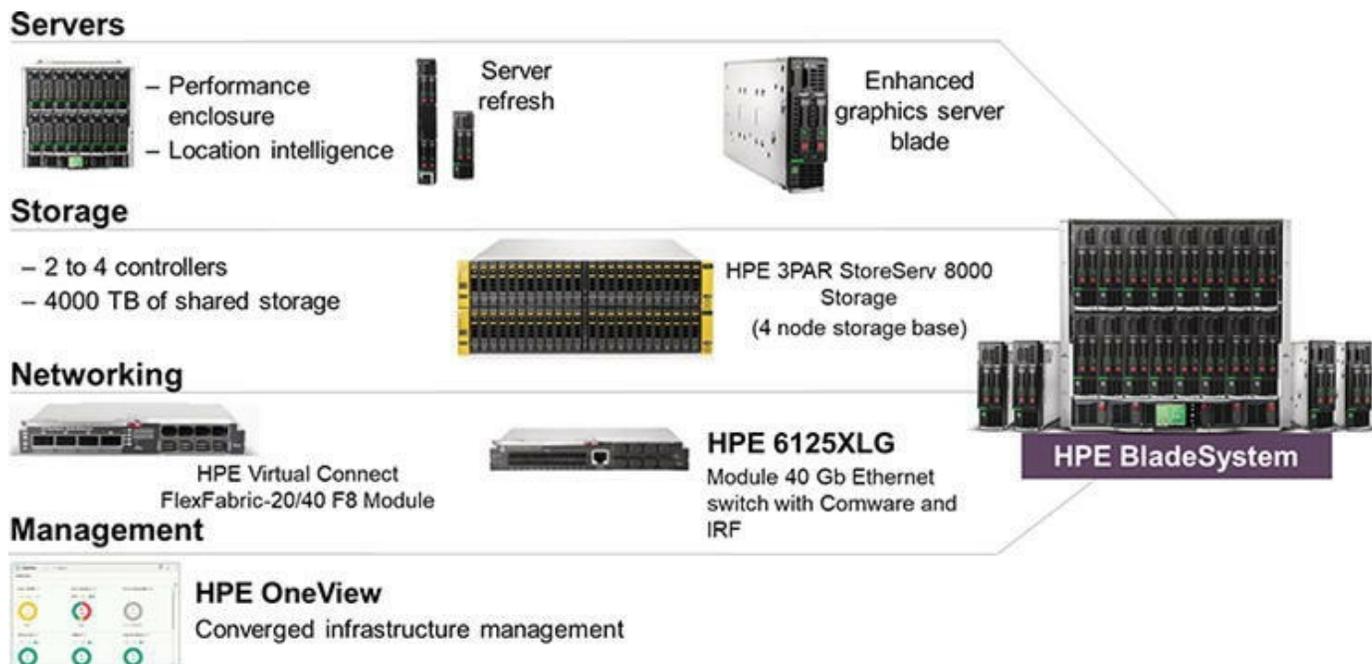


Figure 7-3 The highly adaptive and scalable BladeSystem c7000 enclosure

The BladeSystem c7000 enclosure is adaptive and scalable, as shown in [Figure 7-3](#). It transitions with the IT environment and includes modular server, interconnect, and storage components. The enclosure is 10U high and holds full-height and/or half-height server blades that may be mixed with storage blades, plus redundant network and storage interconnect modules. The enclosure includes a shared high-speed passive midplane with aggregate bandwidth of up to 7 Tbps for wire-once connectivity of server blades to network and shared storage.



Note

7 Tbps is calculated as follows: 16 half-height device bays \times 16 lanes per device bay \times 14 Gbps per device bay to interconnect bandwidth per lane \times 2 for duplex = 7168 Gbps = 7 Tbps. The midplane also provides additional bandwidth through device bay to bay crosslinks, interconnect crosslinks, and management signals.

Power is delivered through a passive pooled-power backplane that enables the full capacity of the power supplies to be available to the server blades for improved flexibility and redundancy. Power input is provided with a very wide selection of AC and DC power subsystems for flexibility in connecting to data center power.

You can populate a BladeSystem c7000 enclosure with these components:

Up to eight full-height or 16 half-height server, storage, or other option blades

Up to eight interconnect modules (four redundant fabrics) featuring a variety of industry standards, including:

- Ethernet
- Fibre Channel
- Fibre Channel over Ethernet (FCoE)
- IB
- iSCSI
- SAS

Up to six hot-plug power supplies supporting N + 1 and N + N redundancy

Choice of a single-phase AC, three-phase AC, universal AC–DC high voltage, or a –48V DC power subsystem

Up to 10 hot-plug HPE Active Cool 200 fans

Redundant BladeSystem OA management module (optional active-standby design)

BladeSystem with HPE OneView



Figure 7-4 HPE BladeSystem and HPE OneView provide a single infrastructure and management platform

BladeSystem and HPE OneView provide a single infrastructure and single management platform with federated intelligence that is automated for rapid service delivery, as shown in [Figure 7-4](#). BladeSystem and HPE OneView deliver leading infrastructure convergence, availability with federation, and agility through data center automation.

HPE OneView provides additional functionality and integrates with:

Provisioning

- Representational State Transfer (REST) application programming interface (API)
- Unified Extensible Firmware Interface (UEFI)
- HPE Intelligent Provisioning
- HPE integrated Lights-Out (iLO) 5
- HPE Scripting Toolkit

Monitoring

- HPE Active Health System
- HPE Agentless Management
- Direct connect or HPE Insight Remote Support
- Warranty and service event tracking

Optimizing

- Smart Updates with HPE Smart Update Manager (HPE SUM)

A federated system of blades

BladeSystem and HPE OneView form a federated system. This means that each BladeSystem chassis becomes part of a single management environment. Any workload can be moved to any chassis as long as the receiving blade has the right configuration. It becomes a single, bigger system of blades. In this environment, each blade participates with the other blades, but is a failure domain unto itself. This keeps failure domains small and limited to a chassis, compared to competitive environments where failures in one place can affect everything else.

BladeSystem with HPE OneView is:

Converged—The combination of these solutions encompasses compute, network, storage, virtualization, and management components in one infrastructure, reducing costs. BladeSystem helps businesses:

- Reduce data center costs
- Shift investment from routine maintenance to innovation
- Maximize availability by reducing downtime
- Accelerate enterprise workload deployment such as virtualization and cloud computing

Federated—The chassis is a single canvas or system. Although federated, HPE OneView creates one infrastructure so that you can move server profiles across multiple chassis, and any compute blade is available for any appropriate workload in the whole system. It is not hierarchal, which creates a safer environment because failure domains are smaller.

Automated—Daily lifecycle management is easy with autonomic provisioning, proactive health monitoring, and virtual machine (VM) failover. A valuable benefit is the ability to automate the environment.

ProLiant BL server blades

A server blade is an independently functioning server with all the necessary components integrated into a single board, including processors, memory, network adapters, and optional hard drives. The ProLiant BL Gen10 line of server blades supports the latest Intel processors along with a wide variety of key features:

HPE SmartMemory @ 2666 MT/s DDR4 LRDIMM, RDIMM, and support for non-volatile DIMMs (NVDIMMs)

HPE Smart Storage—Workload optimized, flexible storage

- HPE Smart HBA SATA, SAS, or solid-state drives (SSDs) small form factor (SFF) drives, depending on the model
- Expanded storage options with hot-plug uFF drives and larger internal M.2 options
- HPE Smart Array Controllers
 - ◆ S100i SR SW RAID (Embedded SATA Software RAID/No Cache) Microsoft Windows Server only
 - ◆ HPE Smart Array P204b-i (12 G SAS, 1 GB FBWC)
 - ◆ Smart Array P408e-m controller, depending on model



Note

The HPE Gen10 NVMe FIO Setting (873373-B21) is required to support SFF NVMe SSDs within the system. This option is not compatible with the HPE Smart Array P204i-b. HPE recommends the use of a dual M.2 SSD kit for boot when using this option.

The HPE Smart Array S100i SR Controller (chipset SATA) comes standard with the HPE BL460c Gen10 10 Gb/20 Gb FLB configure to order (CTO) Blade. If the HPE Smart Array P204i-b controller is not chosen, a SATA cable will be provided to support SATA devices for the two internal drive bays. If RAID is required when using the S100i, please choose the HPE FIO Enable Smart Array S100i SR Setting (product number 784308-B21).

- Improved storage options with support for USB 3.0
- Smart Storage Battery with 12W shared backup power

Boot devices

- Slot for two M.2 storage devices

- Optional redundant microSD card

Optional HPE OneView system management

Scalable blade form factors and device bays

BladeSystem server blades are delivered in two form factors: half-height and full-height. Server blades can be installed (and mixed with other server blades) in BladeSystem c7000 enclosures.

The half-height and full-height server blade form factors, which scale vertically in the c7000 enclosure, provide several benefits, including reduced cost, increased reliability, improved ease-of-use, and configuration flexibility. Placing full-height form factors, half-height form factors, or both in the same enclosure lets you exploit this flexibility. You can use two half-height bays together (one above and one below) to accommodate a full-height, single-wide server blade.

Common server blade configurations include 16 two-processor half-height server blades, eight four-processor full-height, single-wide server blades, or a mixture of the two form factors.

The ProLiant BL server portfolio



Figure 7-5 ProLiant server blades deliver value across workloads

ProLiant server blades deliver value across workloads to support a wide range of applications, from business infrastructure to mission-critical, on a single platform. [Figure 7-5](#) shows the ProLiant BL server portfolio.

ProLiant Gen10 and Gen9 server blades build intelligence and automation in every

server and across the data center. They continuously analyze thousands of system parameters to optimize application performance, proactively improve uptime, and offer insight into every aspect of an IT infrastructure.

Server virtualization continues to provide significant return for businesses and as VM workloads increase, it is critical that the infrastructure can scale to meet these demands. BladeSystem and ProLiant server blades provide the ideal platform for virtualization. ProLiant server blades are designed for virtualization from the ground up, delivering more memory and superior networking capability for virtualized environments and compute intensive application workloads.

Key features of ProLiant Gen10 server blades compared with the previous generation are shown in [Figure 7-6](#).



Figure 7-6 ProLiant Gen10 server blades key features

ProLiant BL460 Gen10

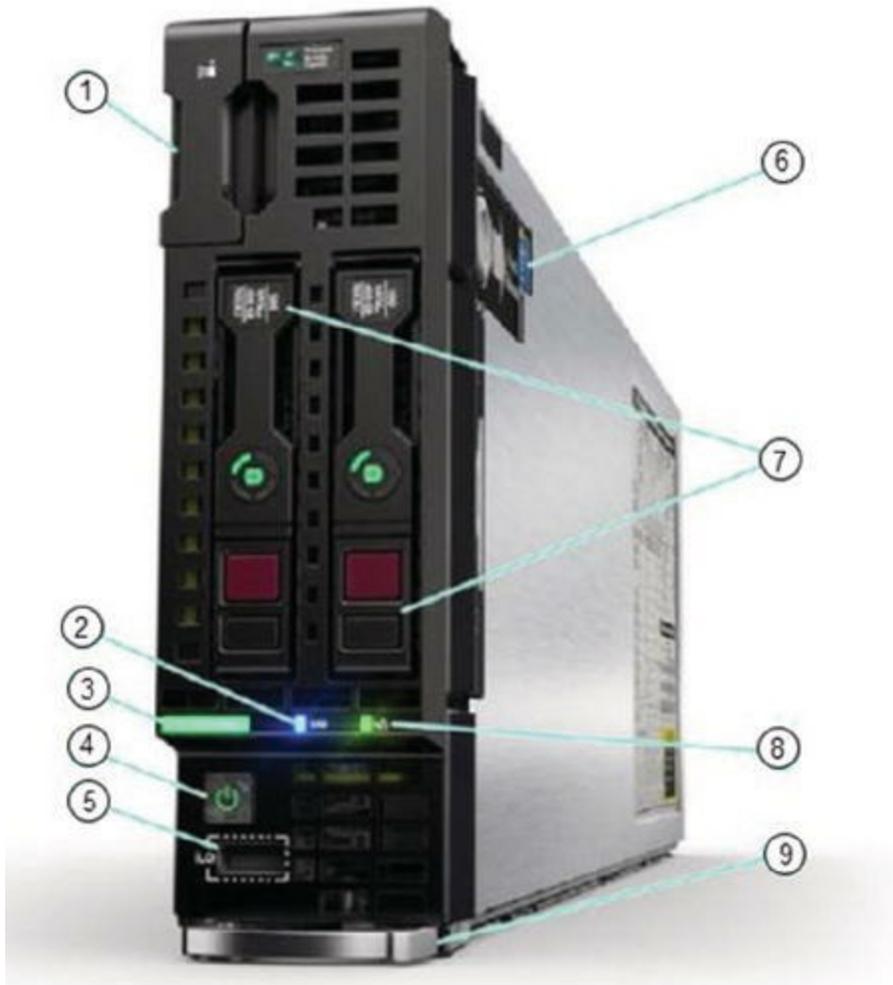


Figure 7-7 ProLiant BL460 Gen10 front view

[Figure 7-7](#) shows the front view and features of the ProLiant BL460 Gen10. Key features are:

Front view:

Serial label pull tab (SUV connector located behind)

Unique identification (UID) LED

Health status bar LED

Power button

iLO USB connection

Quick access panel

Front hot-plug drive bays

NIC activity LED

Server blade release lever

ProLiant BL460c Gen10

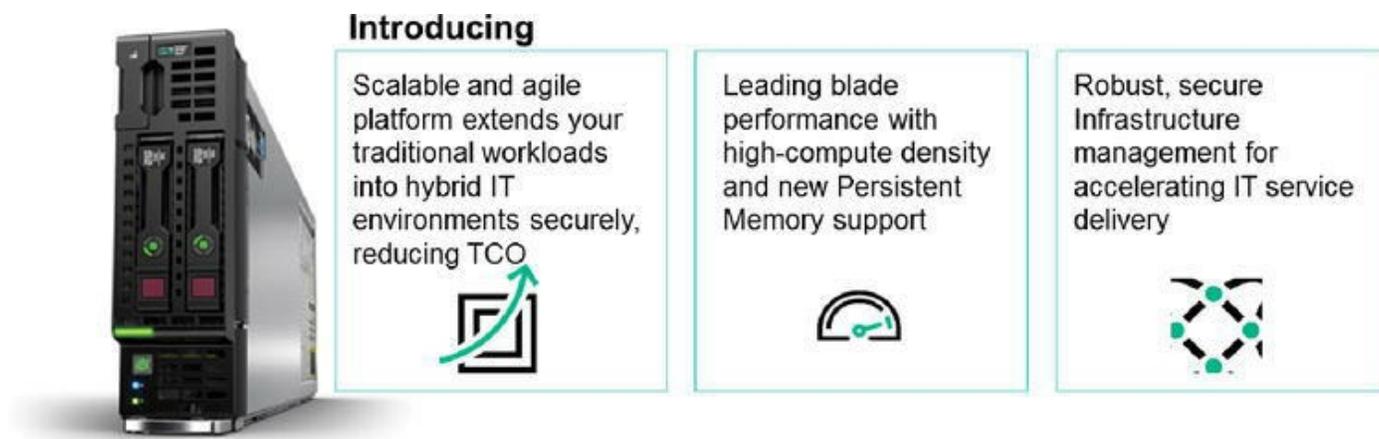


Figure 7-8 The standard to securely accelerate services delivery for converged workloads

Deliver increased agility and performance to your customer's business with the HPE ProLiant BL460c Gen10 Blade Server, which streamlines converged operations, boosts business performance, and increases ROI through centralized deployment and management. From virtualization through unified communications & collaboration (UC&C) and cloud environments, boost workload performance, and resource utilization for optimal services delivery service-level agreements (SLAs), with the standard for converged data center architectures.

HPE ProLiant BL460c Gen10 Server supports:

Up to two Intel® SkyLake Bronze, Silver, Gold, and Platinum (150W and below) processors

HPE Smart Memory (16) DDR4, up to 2666 MT/s (1 TB max)

Support for two 16 GB NVDIMMs (Type 1)

Standard HPE Dynamic Smart Array S100i

Optional HPE Smart Array P204i-b or NVMe pass through

uFF HPE SmartDrives

Two NVMe PCIe SSD option

Optional M.2 Support

HPE iLO 5

As shown in [Figure 7-8](#), benefits of the ProLiant BL460c Gen10 include:

Leading-edge performance and density

- Up to 25% performance increase bin-to-bin with 150W processors and DDR4-2666 memory over Gen9 blade servers
- Up to 65% more performance with 1.6 million IOPS from Gen10 storage controllers
- Dynamically tune compute performance for specific workloads with Intelligent System Tuning

Scale across hybrid IT

- Modular pay-as-you-go approach to add compute and fabrics that grows with your business
- Drive workloads and applications with support for Unified API for the application ecosystem, including Docker and others

Enables Path to Composable

- Delivers key attributes of the Composable experience
- Simplify IT operations fully managed by HPE OneView 3.x to quickly deploy converged infrastructures

HPE ProLiant BL460c Gen10 to BL460c Gen9 comparison

Table 7-1 HPE ProLiant BL460c Gen10 to BL460c Gen9 comparison

Specifications	Gen9	Gen10
Processor	Intel® Xeon® E5-2600 v3 product family Intel® Xeon® E5-2600 v4 product family	Intel® Xeon® Processor Scalable family up to 150W per socket) Platinum/Gold/Silver/Bronze product family
Processors/cores/speed	One or two processors; 22, 20, 18, 16, 14, 12,10, 8, 6, 4 cores; 3.5 GHz	One or two processors; 26, 24, 22, 20, 18, 16, 14, 12,10, 8, 6, 4 cores; 3.6 GHz

Memory (type, max, slots)	Support up to 2400 MT/s DDR4 SmartMemory; 2 TB Max with 128 GB DIMMs, 16 DIMM slots	Support up to 2666 MT/s DDR4 SmartMemory ; 2 TB Max with 128 GB DIMMs, 16 DIMM slots; Support up to 2× Type-1 16 GB NVDIMM
Drive bays	(Two) SFF SAS/SATA/SSD or (Two) SFF NVMe (optional) SSD	(Two) SFF SAS/SATA/SSD or (Two) SFF NVMe (optional) SSD or (Four) uFF SSD
Network controller*	10 Gb 534M, 536FLB, 560FLB FlexFabric Adapter 2 Ports per controller; 20 Gb 630/650FLB, 630/650M FlexFabric Adapter 2 Ports per controller	10 Gb 534M, 536FLB, 560FLB FlexFabric Adapter 2 Ports per controller; 20 Gb 630/650FLB, 630/650M FlexFabric Adapter 2 Ports per controller
Graphics	WS460c Gen9	Future option on BL460c
Infrastructure management	iLO Management (standard), Intelligent Provisioning (standard), iLO Advanced for BladeSystem (optional), HPE OneView Advanced (optional), HPE Insight Control (optional)	Front USB iLO port , iLO Management (standard), Intelligent Provisioning (standard), iLO Advanced for BladeSystem (optional), HPE OneView Advanced (optional), HPE Insight Control (optional)

Power supply	Enclosure based: Up to (six) each: 2650W* Platinum, 2400W Platinum, 2400W Gold, 2250W	Enclosure based: Up to six each: 2650W* Platinum, 2400W Platinum, 2400W Gold, 2250W
Storage controller*	One HPE Smart Array P244br Controller with 1 GB Flash Backed Write Cache (FBWC) or HPE H244br Smart HBA.	Standard HPE Dynamic Smart Array S100i optional HPE Smart Array P204i-b with 1 GB Flash Backed Write Cache (FBWC) or NVMe pass through
Warranty (parts, labor, onsite support)	3/3/3	3/3/3

[Table 7-1](#) provides a comparison between ProLiant BL460c Gen9 and Gen10 servers.

Learning check

What is the maximum number of server blades that can populate a single c7000 enclosure?

- A. Up to four full-height and 16 half-height blades
- B. Up to eight full-height or 16 half-height blades
- C. Up to eight full-height and 16 half-height blades
- D. Up to 16 full-height and eight half-height blades

BladeSystem enclosure



Figure 7-9 BladeSystem c7000 enclosure

A BladeSystem solution starts with an enclosure to serve as the foundation for a converged infrastructure. Enclosures offer power, cooling, network, and storage capabilities in one chassis.

The BladeSystem c7000, shown in [Figure 7-9](#), is 10U high and holds up to 16 server or storage blades along with redundant network and storage switches. It includes a shared, multi-terabit, high-speed midplane for wire-once connectivity of server blades to network and shared storage.

Power requirements include rack-level PDUs or data center UPSs. Power is delivered through a pooled-power backplane, which ensures that the full capacity of the redundant hot-plug power supplies is available to all blades. Universal high-voltage power allows customers to reduce data center capital and operating expenditures.

The c7000 is ideal for:

Large, dynamic data centers and applications

Demanding workloads

An environment with multiple servers and network designs

Businesses that:

- Need between eight and 16 server or storage blades per enclosure
- Experience rapid growth in the IT environment and purchase servers frequently
- Require rack-level PDUs or data center UPS
- Require the highest levels of availability and redundancy
- Need multiple rack-based shared storage arrays for server blades

BladeSystem c7000 enclosure

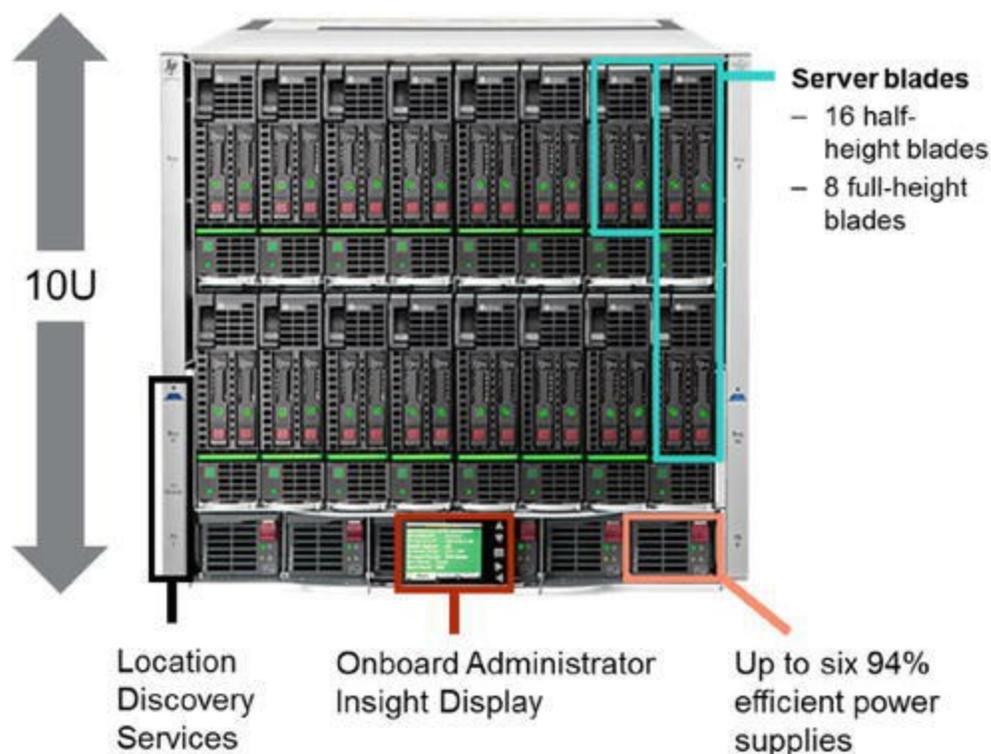


Figure 7-10 Front view of c7000

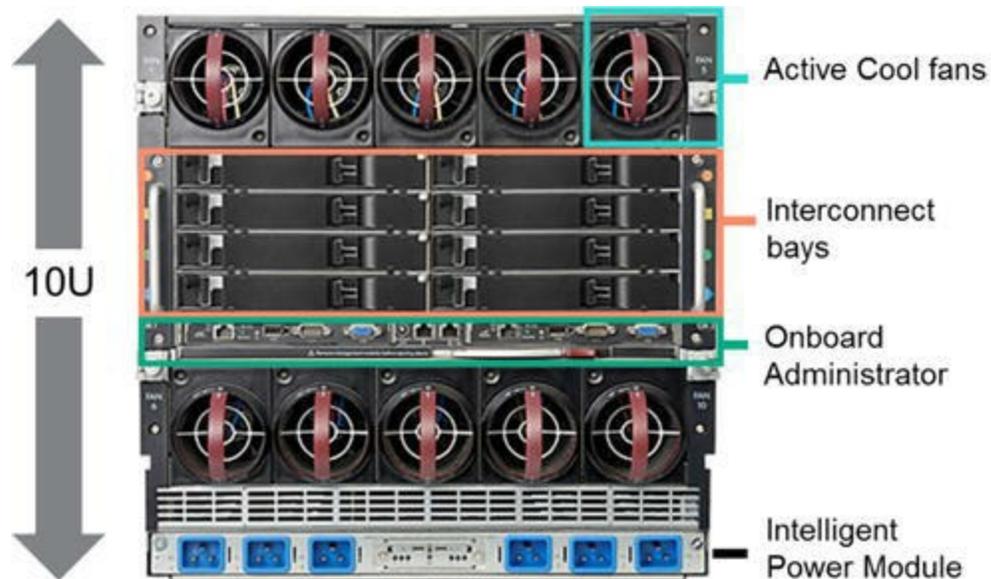


Figure 7-11 Rear view of c7000

[Figures 7-10](#) and [7-11](#) show the location of components on a BladeSystem c7000 enclosure. From the front view, you can see the server blades, which is where the iLO management processor resides. You can view management and troubleshooting data from the OA on the Insight Display screen. From the back, you can see the interconnect bays, where storage can be connected. A BladeSystem enclosure supports connecting an HPE 3PAR StoreServ storage system directly. However, all other storage systems must be connected using a fabric. The OA module resides under the interconnect bays, and power management is at the bottom.

The BladeSystem c7000 enclosure delivers power through the power subsystem. It includes:

- A wide choice of power input options including single-phase, high-line AC; three-phase, high-line AC; universal AC/DC high voltage; or -48V DC

- Up to six hot-plug power supplies with up to 2650W each and support for N + 1 and N + N redundancy

- A passive-pooled-power backplane ensuring that the full capacity of the power supplies is available to many components for improved flexibility and redundancy

The 2650W power supply offers a wide-ranging input supporting 277V AC and 380V DC, in addition to a power input module with Anderson Power Products (APP) Saf-D-Grid connectors. The increased power supply wattage and midplane bandwidth aligned with intelligent infrastructure technologies such as Platinum Power Supplies, Intelligent Power Module, and HPE Location Discovery Services have enhanced the foundation for converged infrastructure.

The BladeSystem c7000 enclosure has redundant signal paths between servers and interconnect modules. The enclosure signal midplane and separate power backplane are passive (have no active components). Separating the power delivery in the backplane from the high-speed interconnect signals in the midplane results in reduced thermal stress to the signal midplane.

The c7000 enclosure and the components within it optimize the cooling capacity through unique mechanical designs. Airflow through the enclosure is managed to ensure that every device gets cool air and does not sit in the hot exhaust air of another device. Air only goes where it is needed for cooling. Fresh air is pulled into the interconnect bays through a slot in the front of the enclosure. Ducts move the air from the front to the rear of the enclosure, where it is then pulled into the interconnects and the central plenum, and then exhausted out the rear of the system.

OA module



Figure 7-12 c7000 OA module

Unique to BladeSystems, the HPE OA shown in [Figure 7-12](#) is the enclosure management processor, subsystem, and firmware base used to support the BladeSystem enclosure and all the managed devices contained within the enclosure. Each enclosure holds one or two OA management modules. If present, the second module acts as a redundant controller in an active-standby mode.

A BladeSystem OA management module is built into the enclosure with the following functions:

Control of multiple enclosures

A wizard-based initial setup for easy configuration

Integrated access to:

- All server blade iLOs from a single cable
- Interconnect bay device management ports from the single BladeSystem OA cable

Single sign-on capability for all devices in the enclosure

Role-based security locally and with Lightweight Directory Access Protocol (LDAP) directory services

The OA:

Provides a secure, single point of contact for users performing basic management tasks on server blades or switches within the enclosure

Monitors and manages elements of the enclosure such as shared power, shared cooling, IO fabric, and iLO

Reports:

- Asset and inventory information for the devices in the enclosure
- Thermal and power status, including real-time actual power usage per server and per enclosure

Is fully integrated into all HPE system management applications

Performs initial configuration steps for the enclosure

Enables runtime management and configuration of the enclosure components

Informs users of problems within the enclosure through email, Simple Network Management Protocol (SNMP), or the Insight Display

Offers web-based and command line interface (CLI) manageability

- Drives all management features through the two Inter-Integrated Circuit (I2C) and Intelligent Chassis Management Bus (ICMB) interfaces
- Aggregates up to 16 iLO ports in a c7000 enclosure
- Simplifies cable management
- Provides a GUI to launch individual server iLO management interfaces

iLO 5 monitors each server blade to control enclosure fans and provide optimal cooling for each server. It passes sensor data to the OA, which the OA uses to manage the

system's power use. The rear of each module has a blue LED that can be enabled (locally and remotely) and used to identify the enclosure from the back of the rack.

The OA can be managed locally, remotely, and through HPE Systems Insight Manager (SIM) tools. It also provides local and remote management capability through Insight Display and browser access. The OA module for the c7000 enclosure is available in two versions: with keyboard, video, and mouse (KVM) support and without KVM. The firmware for both versions is the same. Each OA module has three ports: a network, USB, and serial port. Double data rate 2 (DDR2) OA models also have a VGA connector.

Insight Display

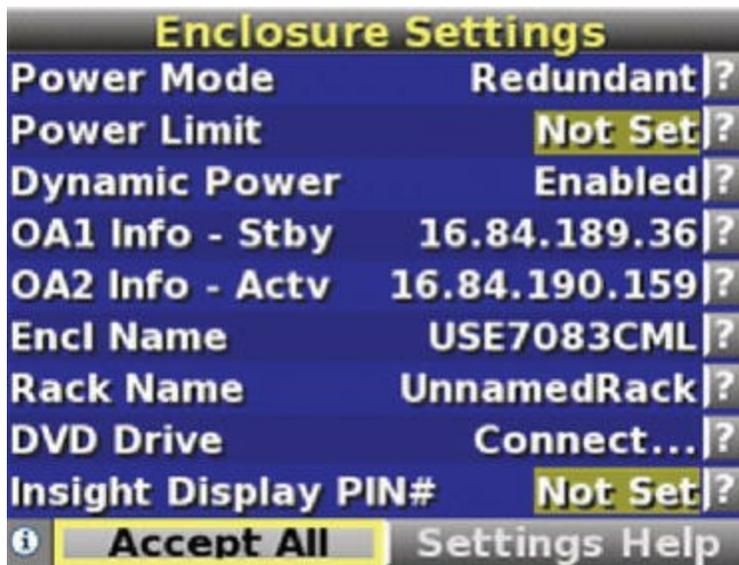


Figure 7-13 Insight Display view

BladeSystem Insight Display, shown in [Figure 7-13](#), is an LCD panel on the front of the enclosure that provides easy setup and management functions. It is designed for configuring and troubleshooting while you are standing next to the enclosure in a rack. It displays a quick glance of enclosure settings, the health status of all bays in the enclosure, and diagnostic information if the OA detects a problem in the enclosure. A green background indicates that everything in the enclosure is properly configured and running within specifications.



Note

The Insight Display configures important settings in the OA, including the IP address of the OA module. You can access the Insight Display directly through the buttons on the panel or remotely through the OA GUI.

The integrated Insight Display is linked to the OA with KVM for local system setup and management. The LCD panel slides left to right to allow access to power supply bays 3 and 4, which are behind the panel in the enclosure.

Main Menu

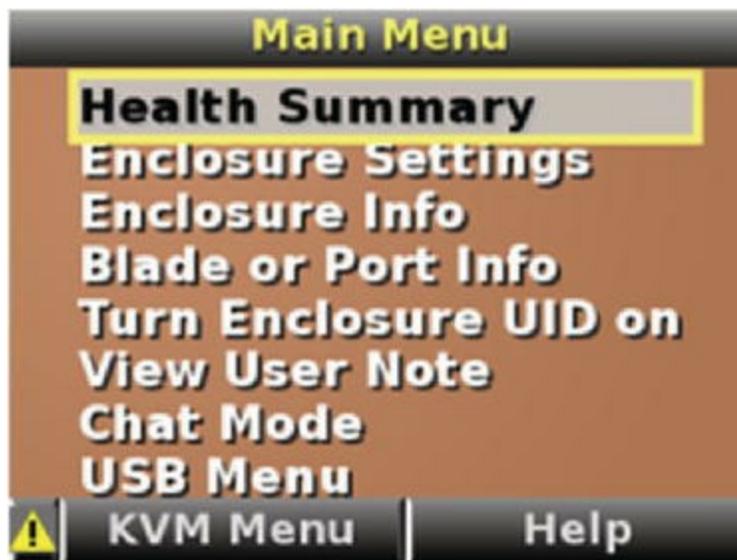


Figure 7-14 Insight Display Main Menu

From the Insight Display Main Menu, you can navigate to submenus by pressing the **Left**, **Right**, **Down** and **Up** arrow buttons.

As shown in [Figure 7-14](#), the Main Menu items include:

Health Summary—Displays the current condition of the enclosure. When an error or alert condition is detected, the Health Summary screen displays the total number of error conditions and locations.

Enclosure Settings—Enables you to configure the enclosure and update its settings. You can also make changes directly from the rack. From the Enclosure Settings screen, you can configure:

- Power settings such as the power mode and the power limit
- Standby and active OA IP addresses
- Enclosure name
- Rack name
- Insight Display lockout PIN number

Enclosure Info—Displays information about the enclosure, such as active OA IP address, current health status of the enclosure, and the ambient temperature.

Blade or Port Information—Displays information about a specific blade. From this screen, you can view ports used by a specific blade, embedded NICs, and interconnect modules.

Turn Enclosure UID on—Displays **Turn Enclosure UID Off** when the enclosure UID is active, and displays **Turn Enclosure UID on** when the enclosure UID is off.

View User Note—Provides helpful information such as contact phone numbers. You can change this screen using the remote OA user web interface. Both the background bitmap and the text can be changed.

Chat Mode—Allows the remote administrator to send messages to an enclosure Insight Display to communicate with another user. The technician can use the Insight Display buttons to select from a set of prepared responses or dial in a custom response message.

USB Menu—Offers links to update OA firmware or to restore or save an OA configuration.

HPE Thermal Logic technologies

The BladeSystem c7000 enclosure uses several HPE Thermal Logic technologies, including mechanical design, built-in power and thermal monitoring, and control capabilities. Thermal Logic technologies yield significant power and cooling savings compared to traditional rack- and tower-based servers. Thermal Logic technologies also provide an instant view of power use and temperature at the server, enclosure, or rack level. They automatically adjust power and thermal controls to reduce power usage while maintaining adequate cooling for the devices to enable high availability.

Thermal Logic c7000 enclosure technologies include the following elements and capabilities:

HPE Active Cool 200 fans

HPE Parallel Redundant Scalable Enclosure Cooling (PARSEC) architecture

Power and thermal monitoring

3D Sea of Sensors

HPE Active Cool 200 fans

Quite often, dense, full-featured, SFF servers use very small fans for localized cooling in the specific areas. Because the fans generate low airflow (in cubic feet per minute or CFM) at medium backpressure, a single server often requires multiple fans to provide adequate cooling. If each server blade contains several fans, installing many server blades together in an enclosure can result in a significant cost and space overhead.

A second solution for cooling is to use larger, blower-style fans for an entire enclosure. The fans generate high airflow, but they typically require higher power input and more space. They are loud and designed for the largest load in an enclosure. As a result, designers may have to sacrifice server features to allow large, high-power fans to fit in the enclosure. Even then, allowing for adequate airflow to the servers without leakage, over provisioning, or bypass is a challenge.

To overcome these issues in the BladeSystem c7000 enclosure, HPE designed a new type of fan that delivers high airflow and high pressure in a SFF that can scale to meet future cooling needs. HPE has 20 patents for the Active Cool fan technology and implementation.

Active Cool 200 fans can cool 16 server blades using as little as 150W of power. Active Cool 200 fans use ducted fan technology with a high-performance motor and impeller to deliver high CFM at high pressure. The fan includes a bell-mouth inlet with a specially designed impeller and a stator section that also provides cooling fans for the motor and acoustic treatments at the rear of the fan. The unique shape of the fan generates high-volume, high-pressure airflow at even the slowest fan speeds with low noise levels and reduced power consumption.

Learning check

Name components of HPE BladeSystem. (Select three.)

- A. Uninterruptable Power Supply
- B. OA
- C. Insight Control server provisioning
- D. Server blades
- E. Interconnect modules
- F. Location Discovery Engine

BladeSystem interconnect options

The HPE BladeSystem provides eight interconnect bays at the rear of the chassis that can accommodate up to four pairs of identical interconnect modules which when configured correctly form four redundant fabrics. A minimum of one pair of identical interconnect modules is required to create a single redundant fabric.

Scalable interconnect bays

Whether you need basic connectivity for a remote office or a high-bandwidth, low-latency link for a high-performance computing cluster, BladeSystems offer an enclosure connectivity option to meet your needs. The c7000 enclosure supports the industry's broadest portfolio of networking solutions spanning Ethernet, Fibre Channel, IB, SAS, iSCSI, and FCoE. It includes up to four redundant interconnect fabrics supported simultaneously within the enclosure.

BladeSystem interconnect technology features HPE Virtual Connect providing wire-once connectivity to simplify and converge your server-edge connections. The interconnect bays in the enclosure offer network redundancy in the form of two single-wide or double-wide form factor slots for efficient use of space and improved performance. The eight interconnect bays are designed to accommodate two single-wide redundant interconnect modules in a scale-out configuration or a larger, double-wide interconnect module for higher performance, greater bandwidth interconnect modules.

HPE Interconnect modules connect servers and VMs to data and storage networks over Ethernet, Fibre Channel, and iSCSI protocols.

Ethernet interconnects

- HPE Ethernet modules enable up to 2×20 Gb bandwidth to each server and up to 40 Gb speeds to the data center infrastructure.
- HPE Virtual Connect FlexFabric Ethernet modules provide Ethernet connection plus connection to SAN storage with native 8 Gb Fibre Channel and FCoE support.
- A pair of Virtual Connect (VC) modules consolidates multiple physical ports into virtual ports over each of the 20 Gb server ports.

IB modules

- HPE supports 56 Gbps Fourteen Data Rate (FDR) and 40 Gbps 4X Quad Data Rate (QDR) IB switch products.
- The HPE InfiniBand FDR 2-port 545M mezzanine adapter is based on the Mellanox Connect-IB technology. This HCA utilizes a PCIe 3.0 x16 interface and two ports of

FDR IB to provide more than 100 Gb/s of throughput.

- HPE Fibre Channel Pass through support 8 Gb Fibre Channel connectivity.

HPE 6 Gb SAS switch module

- The HPE 6 Gb SAS BL Switch for BladeSystem enclosures is an integral part of HPE direct connect SAS storage, enabling a shared (zoned) SAS storage solution.
- The SAS architecture combines P408e-m SR (with cache module) HPE Smart Array Controller in each server.
- The 6 Gb SAS BL switches connect to supported HPE storage enclosures for SAS storage.
- These switches contain an embedded Virtual SAS Manager (VSM) GUI and CLI interface used to zone, monitor, and update SAS fabric devices.

Fibre Channel interconnects

- HPE Virtual Connect Fibre Channel modules support native FC connectivity. With NPIV¹ enabled FC switches, it is possible to share a single physical Fibre Channel HBA port among multiple virtual ports with high availability and automatic failover.
- HPE Fibre Channel switch modules support native 16 Gb and 8 Gb Fibre Channel connectivity.
- HPE Fibre Channel Pass through support 8 Gb Fibre Channel connectivity.



Note

¹N_Port ID virtualization (NPIV) is a technology that defines how multiple servers can share a single physical Fibre Channel port identification.

Fabric Extender

- The Cisco FEX for HPE BladeSystem provides an extension of the Cisco Nexus switch fabric to the HPE server edge. This module forwards traffic to the parent Cisco Nexus switch over eight 10 Gb uplinks.

HPE BladeSystem signal midplane

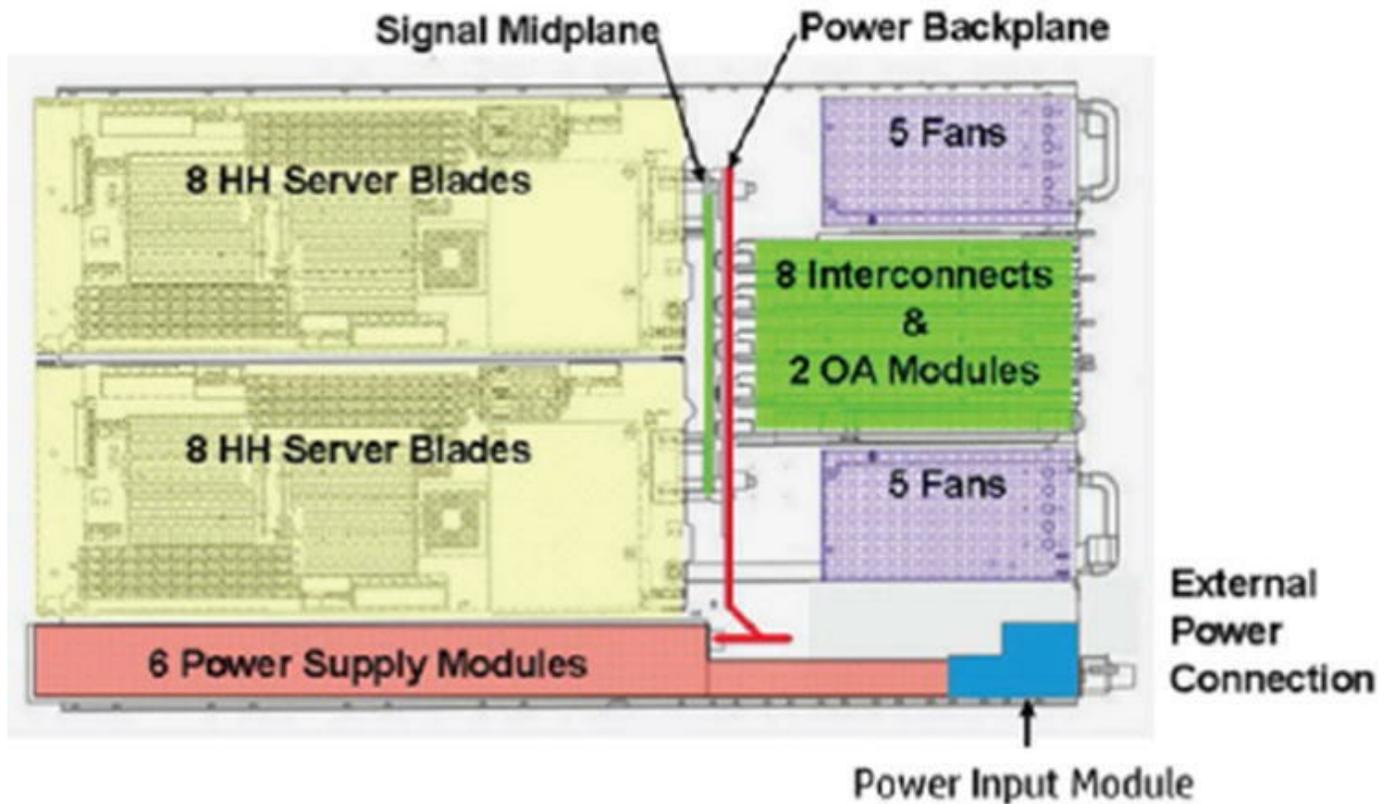


Figure 7-15 HPE BladeSystem signal midplane

The HPE BladeSystem signal midplane provides redundant signal paths between servers and interconnect modules, as shown in [Figure 7-15](#). The enclosure midplane and separate power backplane are passive (no active components).

The rear of the enclosure includes eight interconnect bays that can accommodate eight single or four redundant interconnect module pairs. The interconnect modules plug directly into these interconnect bays. Each c-Class Enclosure requires at least two interconnect switches, VC, or pass-thru modules, side-by-side, for a fully redundant configuration.

The interconnect bays offer network redundancy in the form of two single-wide form factors for efficient use of space and improved performance. The four pairs of interconnect bays are designed to accommodate two single-wide interconnect modules or a single double-wide module for higher performance, greater bandwidth interconnect modules. The signal midplane, in addition to connecting the interconnect modules to the server blades, provides redundant cross-links to adjacent interconnect modules.

Each of the (16) 10 Gb or 20 Gb high-density connectors carry eight data channel connections from each server to the interconnect module bays. A half-height server blade has eight ports provided by a single connector and the full-height server blade has 16 ports provided by two connectors. Each pair of interconnect module bays supports 2

× 16 downlink connections.

NonStop signal midplane scalability and reliability

A key component of the BladeSystem c7000 enclosure is the IO infrastructure NonStop signal midplane that provides the internal wiring between the server or storage blades and the interconnect modules.

The NonStop signal midplane is a passive board that uses serializer/deserializer (SerDes) technology to incorporate multiple protocols and provide point-to-point connectivity between device bays and interconnect bays. To provide high reliability, HPE designed the NonStop signal midplane as a completely passive board. It has no active components along the high-speed signal paths. The midplane consists primarily of traces and connectors. While there are a few components on the midplane, they are limited to passive devices that are very unlikely to fail. The only active device is an Electrically Erasable Programmable Read-Only Memory (EEPROM), which the OA uses to get information such as the midplane serial number. If this device fails, it does not affect the NonStop signal midplane.

The NonStop signal midplane has eight 200-pin connectors to include eight individual switches, four double bay switches, or a combination of the two. It provides the flexibility of 1×, 2×, or 4× connections from the server blade mezzanine cards that connect to the interconnect bays.



Note

If the midplane is replaced for any reason, the EEPROM needs to be reprogrammed with the enclosure serial number. The OA CLI command is used to achieve this.

In the BladeSystem c7000 Platinum Enclosure, the aggregate bandwidth between device bays and interconnect bays is 7 Tbps across the NonStop signal midplane. The NonStop signal midplane in c7000 Platinum Enclosures features signal rates of up to 14 Gbps per lane (with each lane consisting of four SerDes transmit/receive traces). Each half-height server bay has 16 lanes to the interconnect bays delivering a cross-sectional bandwidth of up to 224 Gbps per direction or 448 Gbps full-duplex. This provides the 7 Tbps total aggregate bandwidth for the 16 half-height server bays (448 Gbps per bay × 16 bays = 7168 Gbps = 7 Tbps).

The aggregate bandwidth is measured between the device bays and interconnect bays. It does not include additional traffic capacity between interconnect modules (interconnect

crosslinks), device bays (device bay crosslinks), server blades, and OA management modules.

For the c7000 Platinum Enclosure, one of the areas of improved design focus was on high-speed signal integrity. Obtaining the 14 Gbps level of bandwidth per lane required special attention to high-speed signal integrity.

Benefits of the c7000 Platinum enclosure include:

Using general proven practices for signal integrity to reduce end-to-end signal losses across the signal midplane

Moving the power into an entirely separate backplane to independently enhance the NonStop signal midplane

Providing a method to set optimal signal waveform shapes in the transmitters, depending on the topology of the end-to-end signal channel

By taking advantage of the similar four-trace, differential SerDes transmit and receive signals, the NonStop signal midplane is compatible with either network semantic protocols (such as Ethernet, Fibre Channel, and IB) and memory-semantic protocols (PCI Express), using the same signal traces. Compared to the previous enclosure midplane, the midplane in the c7000 Platinum Enclosure includes the following server blade to interconnect bandwidth improvements:

Hundred percent networking bandwidth increase with 20 Gb Ethernet (10 Gb to 20 Gb Ethernet support)

Hundred percent storage bandwidth with 16 Gb Fibre Channel (8 Gb to 16 Gb Fibre Channel support)

Forty percent increase in IO bandwidth with 4× FDR IB (40 Gb 4× QDR to 56 Gb 4× FDR IB support)

Mezzanine and FlexLOM cards



HPE FlexFabric 20 Gb
2-port 650M Adapter



HPE FlexFabric 20 Gb
2-port 630M Adapter



HPE FlexFabric 20 Gb 2-
port 630FLB Adapter

Figure 7-16 HPE FlexFabric 650M, 630M, and 630FLB Adapters

ProLiant server blades use two types of mezzanine cards to connect to interconnect

fabrics such as Fibre Channel, Ethernet, SAS, or IB.

Type A and Type B mezzanine cards differences include:

Type A mezzanine card is accepted in both expansion slot 1 and 2 (and 3)

Type B mezzanine cards are only accepted in expansion slot 2 (and 3)

The HPE FlexFabric 20Gb 2-port 650M adapter for ProLiant server blades, shown in [Figure 7-16](#), provides full-featured, high-performance, converged 20 Gb Ethernet that accelerates IT services and increases data center efficiency across a broad range of IO intensive enterprise, cloud, and Telco workloads.

When connected to HPE Virtual Connect FlexFabric 20/40 F8 Modules, the HPE FlexFabric 650M adapter provides 20 GbE performance. It offers a rich set of offload technologies including overlay network tunneling and storage, increasing virtualization performance and host system efficiency. The HPE 650M is a dual-speed adapter that can also operate at 10 GbE when connected to HPE Virtual Connect Flex-10/10D Module or HPE 6125XLG Ethernet Switch.

The HPE FlexFabric 20Gb 2-port 650M adapter supports Converged Network Utility (CNU), a manageability application to configure converged network adapters (CNAs) and Ethernet adapters on HPE servers. This host-based utility supports for both GUI and CLI (scriptable) and can be used to configure Ethernet, FCoE, iSCSI, and NPAR-related features/functionality on multiple operating system platforms, including Windows and Linux. CNU is able to configure multiple HPE adapters from various network controllers at the same time. Users can benefit easier setup steps, shorter reboot time, and one-stop solution for multiple adapters via CNU.

The HPE FlexFabric 630M is a two-port 20 GbE mezzanine adapter, featuring the first generation of 20 Gb Ethernet offering in a single-chip solution on a mezzanine form factor that further reduces power requirements for two ports of 20 Gb Ethernet. It is designed for use with HPE BladeSystem c-Class Gen8 and Gen9 blade server platforms' type A and Type B mezzanine slots.

It provides full-duplex, high-performance Ethernet connectivity with support for HPE Virtual Connect Flex-20 blade interconnect technology, allowing each 20 GbE port to be divided into four physical NICs and optimize bandwidth management for virtualized servers.

The HPE FlexFabric 630M adapter, in conjunction with HPE Virtual Connect FlexFabric technology, helps to extend the benefits of virtualization beyond the server and into the rest of the infrastructure.

The HPE 630M supports enterprise-class features such as VLAN tagging, adaptive interrupt coalescing, MSI-X, NIC teaming (bonding), Tunnel Offload (NVGRE,

VxLAN), Receive Side Scaling (RSS), jumbo frames, and PXE boot. It also supports virtualization features such as SR-IOV, Network Partitioning (NPAR), VMware NetQueue, and Microsoft VMQ.

The HPE FlexFabric 20Gb 2-port 630FLB adapter features the first generation of 20 Gb Ethernet offering in a single-chip solution on a FlexibleLOM form factor, further reducing power requirements for two ports of 20 Gb Ethernet. It is designed for use with HPE BladeSystem c-Class servers. It provides full-duplex, high-performance Ethernet connectivity with support for HPE Virtual Connect FlexFabric blade interconnect technology, allowing each 20 GbE port to be divided into four physical NICs and optimize bandwidth management for virtualized servers. The HPE 630FLB FlexFabric network adapter, in conjunction with HPE Virtual Connect FlexFabric technology, helps to extend the benefits of virtualization beyond the server and into the rest of the infrastructure.

Combining both the FlexibleLOM and Mezzanine adapters in a single BladeSystem ProLiant server can provide up to 80 Gb/s of converged bidirectional Ethernet bandwidth, helping to alleviate network bottlenecks.

HPE Dynamic Smart Array P204i SR Gen10 Controller



Figure 7-17 HPE Dynamic Smart Array P204i-c SR Gen10 Controller

The HPE Dynamic Smart Array P204i-c SR Gen10 Controller shown in [Figure 7-17](#), supporting 12 Gb/s SAS and PCIe 3.0, is ideal for maximizing performance while supporting advanced RAID levels with 1 GB flash backed write cache (FBWC). This controller operates in Mixed Mode which combines RAID and HBA operations simultaneously. It has four internal SAS lanes, allowing connection to SAS or SATA drives, and provides enterprise-class storage performance, reliability, security, and efficiency needed to address your evolving data storage needs. This controller offers enhanced security by encrypting data-at-rest on any drive with HPE Smart Array SR Secure Encryption.

HPE Dynamic Smart Array P408e-m SR Gen10

Controller



Figure 7-18 Dynamic Smart Array P204i SR Gen10 Controller

The HPE Smart Array P408e-m SR Gen10 Controller shown in [Figure 7-18](#), supporting 12 Gb/s SAS and PCIe 3.0, is ideal for maximizing performance on externally attached drives while supporting advanced RAID levels with 2 GB FBWC. This controller operates in Mixed Mode which combines RAID and HBA operations simultaneously. It

has eight external SAS lanes, allowing connection to externally attached SAS or SATA drives, and it provides enterprise-class storage performance, reliability, security, and efficiency needed to address evolving data storage needs. This controller offers enhanced security by encrypting data-at-rest on any drive with HPE Smart Array SR Secure Encryption.

The Gen10 controllers are supported by the HPE Smart Storage Battery, which supports multiple devices and is sold separately.

Standard features include:

Storage interface (SAS/SATA)

Eight SAS lanes across an external SAS port

12 Gb/s SAS, 6 Gb/s SATA technology

Mix-and-match SAS and SATA drives to the same controller

Support for SAS tape drives, SAS tape autoloaders, and SAS tape libraries

RAID 0, 1, 5, 6, 10, 50, 60, 1 ADM, 10 ADM (Advanced Data Mirroring)

Mixed Mode (RAID logic drives and HBA physical drives simultaneously)

Legacy and UEFI boot operation

UEFI System Utilities (storage configuration)

Up to 238 physical drives and up to 64 logical drives

HPE Smart Array SR SmartCache (optional license)

HPE Smart Array SR Secure Encryption (optional license)

HPE SSD Smart Path, Rapid Parity Initialization (RPI), Rapid rebuild, Drive Sanitize
Performance Optimization

VC FlexNIC capabilities

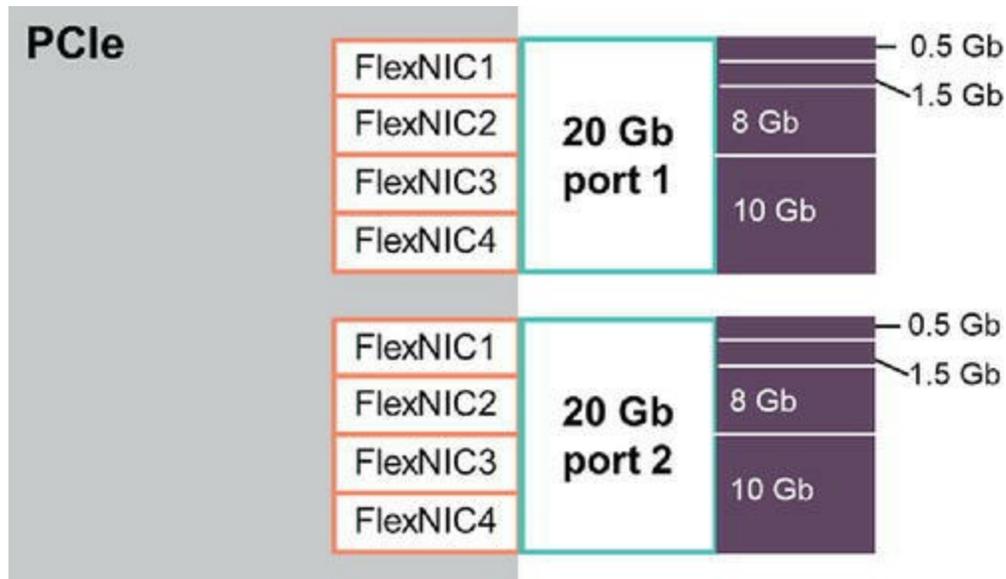


Figure 7-19 Partitioning a link into several smaller-bandwidth FlexNICs

Flex-10/Flex-20 and FlexFabric adapters allow you to partition a 10 Gb/20 Gb link into several smaller-bandwidth FlexNICs, as shown in [Figure 7-19](#). VM applications often require increased network connections per server, which increases network complexity and reduces the number of server resources. VC addresses this issue by enabling you to divide a 10 Gb/20 Gb network port into four independent FlexNIC server connections.

A FlexNIC is a physical PCIe function that presents itself to the system ROM, operating system, and hypervisor as a discrete physical NIC with its own driver instance. A FlexNIC is not a virtual NIC contained in a software layer.

With FlexNICs, you can:

Configure bandwidth on each FlexNIC from 100 MB up to 10 Gb/20 Gb.

Dynamically adjust the bandwidth in 100 Mb increments without requiring a server reboot.

Provide the right amount of bandwidth based on application needs.

Correctly provision bandwidth. You no longer need to overprovision or underprovision. By optimizing bandwidth (setting the minimum and maximum values for individual FlexNICs), VC allocates unused bandwidth from FlexNICs to those FlexNICs with bandwidth demands that exceed minimum values. The minimum ensures guaranteed bandwidth all the time, and maximum is the best effort, depending upon available bandwidth in other FlexNICs. Flex-20-supported adapters can support full-speed protocols such as 10 GbE and 8 Gb Fibre Channel simultaneously.

VC tells the network adapter how to configure each of the four physical functions. Then

the adapter defines each of those physical functions, provisions them into the operating system as individual NICs, and allocates the appropriate bandwidth.

Learning check

A conversation with a customer reveals that they are considering an upgrade to their BladeSystem solution they had installed eight years ago. The data center edge has been upgraded to 20 Gb Ethernet and they are investing in a 16 Gb FC SAN. They ask you if they can upgrade the BladeSystem to support these new technologies. How should you respond?

HPE BladeSystem Fabric connectivity and port mapping

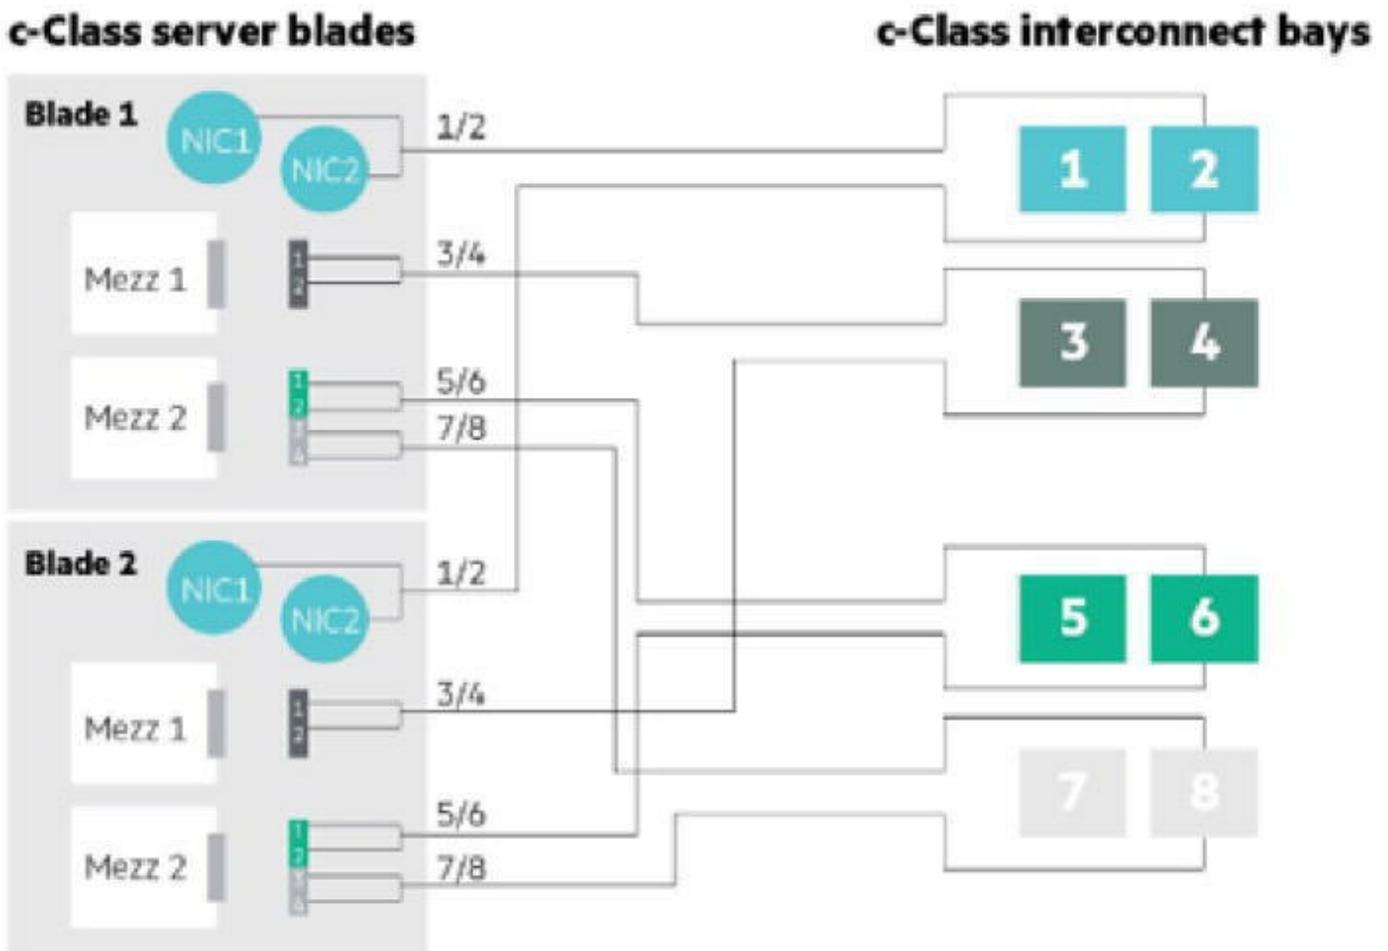


Figure 7-20 How interconnect bays link to server blades

c7000 enclosure device interconnect bay links

The IO signals from the onboard network adapters and mezzanine cards in the server blades in the front of the enclosure are sent to the interconnect modules in the rear of the enclosure via the NonStop passive midplane. Since the NonStop midplane is passive with no active components, these connections between the device bays and the interconnect bays are hard-wired. As a result, the server blade onboard network adapter ports and mezzanine cards are matched to the appropriate interconnect module bays, as shown in [Figure 7-20](#).

For example, an Ethernet switch or Ethernet pass-thru module is used in an interconnect bay that connects to the onboard Ethernet network adapter ports. Likewise, a Fibre Channel switch would be installed in an enclosure interconnect bay that corresponds to a Fibre Channel mezzanine card in a server blade mezzanine slot.

Interconnect bays 1 and 2 are reserved for Ethernet switches or Ethernet pass-thru modules that connect to the server blade onboard network adapter (NIC) ports. The onboard NIC ports may be either embedded Ethernet ports (ProLiant G7 servers and earlier) or ports supplied on a FlexibleLOM adapter (starting with ProLiant Gen8 servers). Mezzanine cards are added for additional IO connectivity and connect to interconnect bays 3 through 8. Redundant interconnect modules are placed adjacent to one another in interconnect bays 3/4, 5/6, and 7/8. For port mapping purposes, it does not matter which device bay houses a server blade. The built-in NIC ports and mezzanine card ports connect to the same interconnect bays.

Port mapping differs slightly between full-height and half-height server blades due to the additional mezzanine cards on the full-height version. The OA and HPE OneView software include tools that simplify the process of mapping mezzanine ports to interconnect module ports.

VC technology

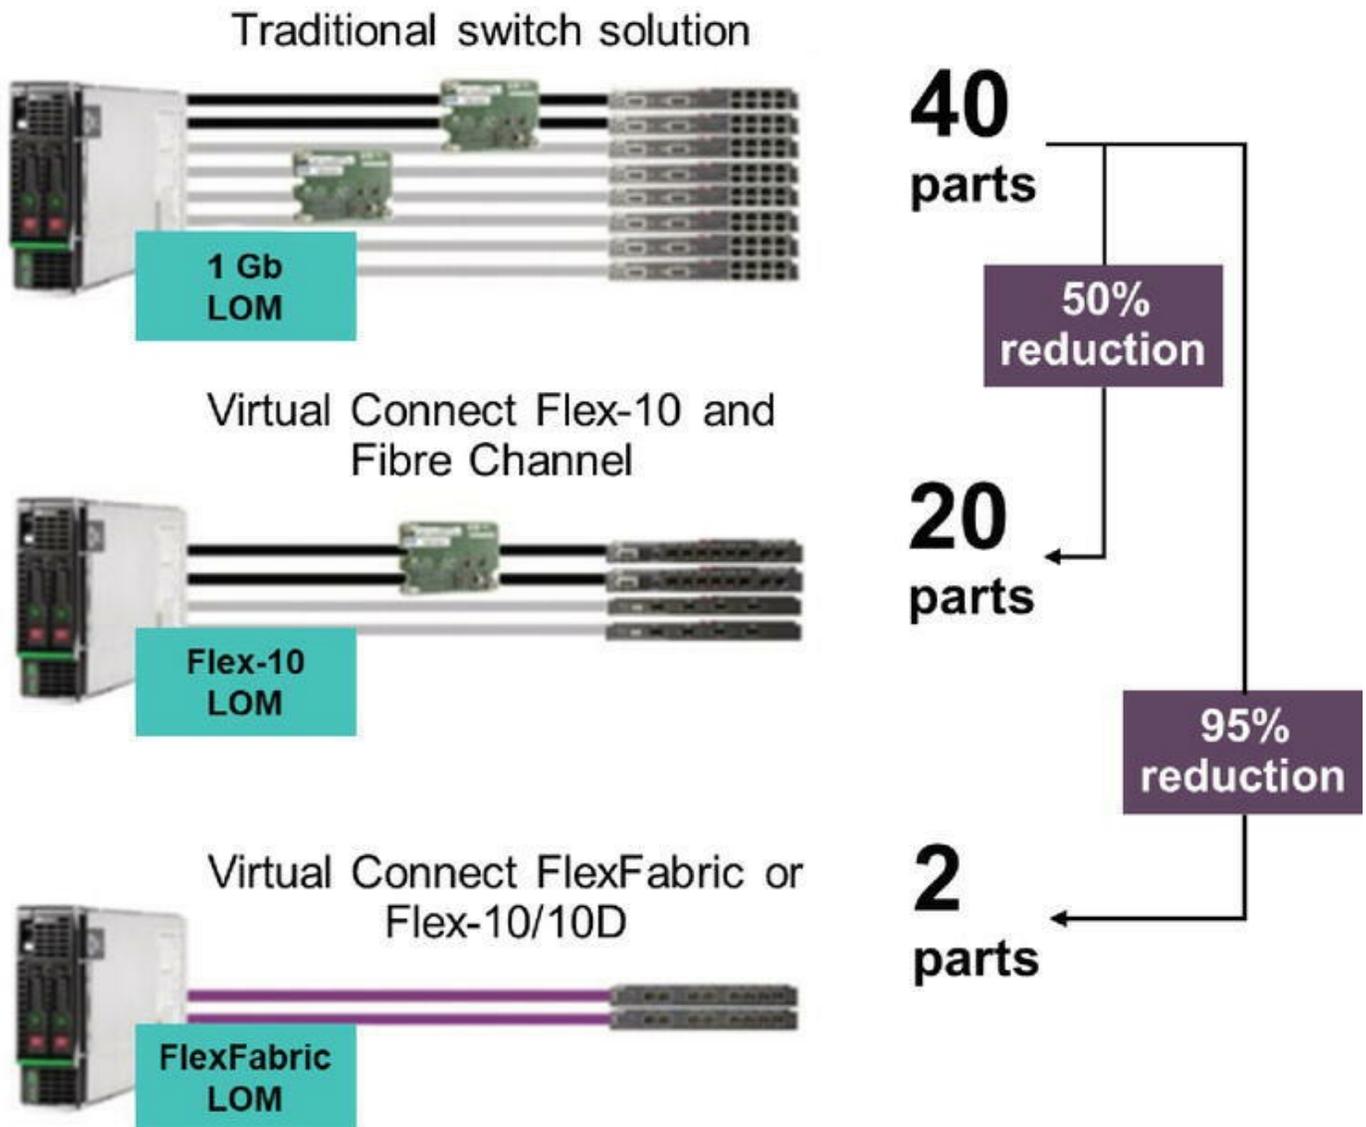


Figure 7-21 FlexFabric solutions reduce cost and components, compared to traditional switch solutions

HPE developed VC technology to simplify networking configuration for server administrators working in a BladeSystem environment. As shown in [Figure 7-21](#), this greatly reduces costs and components compared to traditional switch solutions. VC technology puts an abstraction layer between servers and external networks so that the LAN and SAN see a pool of servers rather than individual servers. Server administrators can physically wire the uplinks from the enclosure to its network connections once, and then manage the network addresses and uplink paths using software.

With VC, network and storage administrators can establish all LAN and SAN connections one time during deployment and do not need to make connection changes later if the servers are changed. After the LAN and SAN connections are made to the

pool of servers, you can use VC Manager (VCM) to create an IO connection profile for each server. When servers are deployed, added, or changed, VC keeps the IO profile for that LAN and SAN connection constant.

Instead of using the default media access control (MAC) addresses for all NICs and default World Wide Names (WWNs) for all HBAs, VCM creates bay-specific IO profiles, assigns unique MAC addresses and WWNs to these profiles, and administers them locally. VC uses pools of unique MAC addresses, WWNs, virtual serial numbers, and server profiles to establish constant server-to-network connections. VC controls the MAC addresses and VLAN connections as well as SAN connections (World Wide Identifiers [WWIDs], zones, and boot parameters) to allow the complete LAN and SAN connection information to be tied to a server or server bay. The server connection profiles contain MAC, WWN, virtual serial number, and boot-from-SAN definitions that are assigned to the enclosure bays and not the physical servers. The physical servers use network assignments and the server profiles instead of default burned-in NIC or HBA addresses. Even if a server is replaced, the MAC and WWN assignments for the enclosure will remain the same, and changes are invisible to the network.

Converging and consolidating server, storage, and network connectivity onto a common fabric with a flatter topology and fewer switches lowers costs. Direct-attach capabilities for HPE 3PAR StoreServ storage systems enabled by Virtual Connect FlexFabric allow customers to flatten and simplify the architecture in the data center. You can move the storage network from an edge-core implementation to an edge implementation directly connected to storage. As shown in the graphic, FlexFabric technology reduces hardware requirements by 95% for a full enclosure of 16 virtualized servers—from 40 components needed in traditional solutions to two FlexFabric modules.

BladeSystem Virtual Connect modules



Virtual Connect FlexFabric-20/40 F8 Module for c-Class BladeSystem



Virtual Connect 16 Gb 24-port Fibre Channel Module for c-Class BladeSystem



Virtual Connect Flex-10/10D module for c-Class BladeSystem



Virtual Connect 16 Gb 20-port Fibre Channel Module for c-Class BladeSystem



Virtual Connect Flex-Fabric 10 Gb/24 port module for c-Class BladeSystem

Figure 7-22 HPE Virtual Connect Ethernet modules and Virtual Connect Fibre Channel modules

BladeSystem supports HPE Virtual Connect Ethernet modules and Virtual Connect Fibre Channel modules as shown in [Figure 7-22](#). Other supported interconnects are:

Ethernet switches

Fibre Channel switches

Pass-thru modules

IB products

Blade switches

Switches offer a traditional approach to administering a network. The primary value in switches is cable consolidation through high-speed uplinks. Each interconnect module also provides external connectors that vary based on its design. Interconnect modules in the c7000 enclosure are available in two widths:

Single-wide modules provide 16 internal ports with each connecting to a separate device bay in the front of the enclosure.

Double-wide modules provide 16 double-wide internal ports enabling connectivity to IB and other four-lane, high-speed interconnects.

In the BladeSystem c7000 enclosure, pairs of single-wide interconnect modules installed in adjacent horizontal bays provide redundant connectivity for dual-port interfaces in each device bay. Adjacent interconnect modules also have high-speed

cross-connect capability through the enclosure's signal midplane. For double-wide interconnects such as IB, two modules are installed in bays 5/6 and 7/8 to provide redundant high-bandwidth connectivity.

Ethernet and Fibre Channel pass-thru modules are available for those who need direct one-to-one connections between servers and either a LAN or a SAN. Ethernet and Fibre Channel pass-thru modules provide 16-port, transparent, one-to-one port connectivity between the server and an external switch.

Activity: Assessing HPE BladeSystem Innovations: 20 Gb adapter

Watch a video about the benefits of BladeSystem. After watching the video, answer the questions below:



Note

View the video by entering the following URL into your browser:

<https://www.youtube.com/watch?v=euDYfNmYI-I>

What do businesses need to satisfy the increased networking bottleneck occurring as a result of large VM deployment on HPE BladeSystem?

How is the cost reduction achieved for LAN and SAN implementation?

VC network configuration

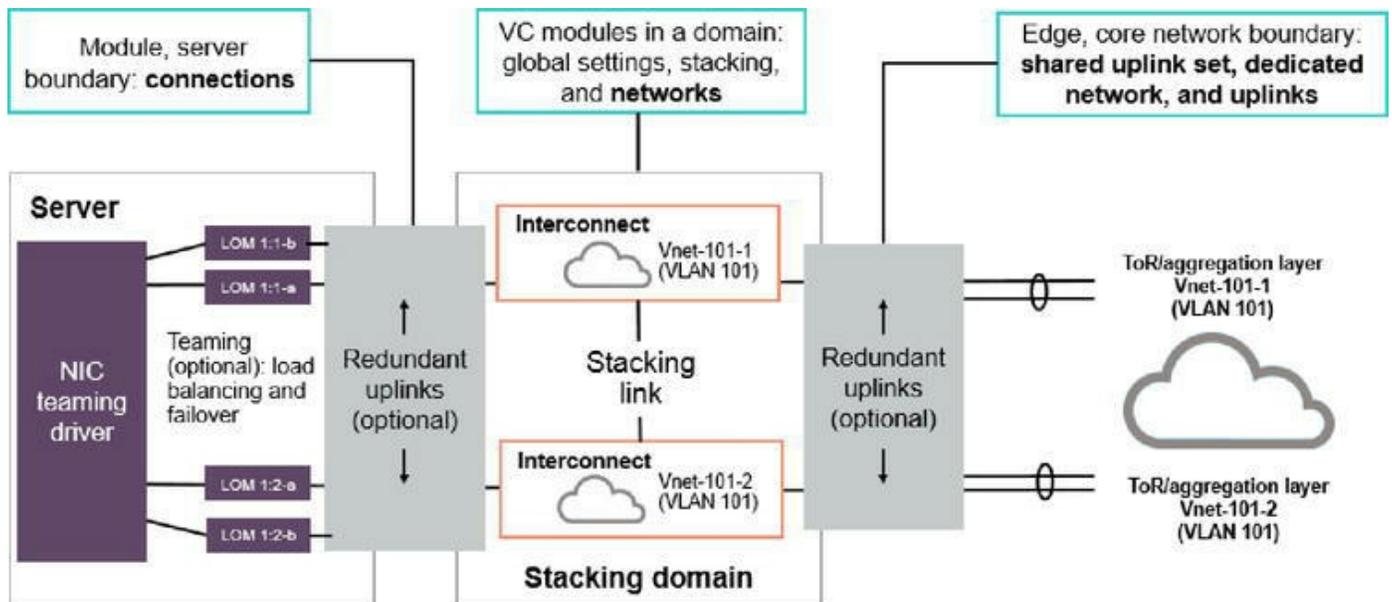


Figure 7-23 VC virtualizes the connections between the server and the network infrastructure

VC is a portfolio of interconnect modules, adapters, embedded software, and optional management applications. Its architecture boosts the efficiency and productivity of data center server, storage, and network administrators. As shown in [Figure 7-23](#), VC virtualizes the connections between the server and the network infrastructure (server-edge virtualization) so that networks can communicate with pools of BladeSystem servers. This virtualization allows you to move or replace servers rapidly without requiring changes or intervention by the LAN and SAN administrators.

VC is standards based and complies with existing and emerging standards for Ethernet, Fibre Channel, and converged networks. The VC portfolio includes:

VC interconnect modules—These modules plug directly into the interconnect bays located in the rear of the BladeSystem enclosure. The FlexFabric-20/40 F8, Flex-10/10D, and Fibre Channel modules connect to server blades through the enclosure midplane. The Ethernet-based modules support 1 Gb, 10 Gb, or 40 Gb on uplinks and 1 Gb, 10 Gb, or 20 Gb on downlinks. This enables you to purchase 1 Gb small form-factor pluggable (SFP) modules and upgrade to 10 Gb SFP+ transceivers when the rest of the infrastructure is ready to support it. In addition, Virtual Connect FlexFabric-20/40 F8 modules offer 2 Gb, 4 Gb, or 8 Gb Fibre Channel on uplinks with Flexports. These universal ports can be configured as Ethernet or Fibre Channel ports.

Flex-10, Flex-20, and FlexFabric adapters—These adapters are available as either LAN-on-motherboard (LOM) devices or mezzanine cards. VC technology also works with 1 GbE adapters and FlexibleLOM devices for ProLiant BL Gen8 servers. A

FlexibleLOM uses a special slot/connector on the motherboard and lets you choose the type of NIC that is “embedded” on the ProLiant Gen8 server.

VC modules enable you to add, remove, or change servers without making corresponding changes to the LAN or SAN. As a result, you can add or change servers in minutes instead of days or weeks.

With VC, you can:

Significantly reduce cabling and server connection complexity.

Maintain constant end-to-end connections to preferred networks and fabrics.

Separate server administration from LAN and SAN administration.

Enable system administrators to be more self-sufficient, so they can add, replace, or modify servers in minutes without impacting production LANs and SANs.

Relieve LAN and SAN administrators from server-centric maintenance.

VC interconnect modules:

Reduce the number of cables required for an enclosure (compared to using pass-thru modules).

Reduce the number of edge switches that LAN and SAN administrators must manage.

Allow pre-provisioning of the network so server administrators can add, replace, or upgrade servers without requiring immediate involvement from the LAN or SAN administrators.

Enable a flatter, less hierarchical network, which decreases equipment and administration costs, reduces latency, and improves performance.

Deliver direct server-to-server connectivity within the BladeSystem enclosure, optimizing traffic flow.

Provide direct-attach SAN and dual-hop FCoE capabilities to extend cost benefits further into the storage network.

By leveraging FCoE, a Layer 2 Ethernet technology, for connectivity to existing Fibre Channel SANs, you can reduce the number of switch modules and HBAs required within the server blade. This further reduces cost, complexity, power, and administrative overhead.

VC management

VC management tools provide the framework that allows administrators to easily set up

and monitor network connections, server profiles, and how the networks map into VMs. These tools are designed to manage the entire converged infrastructure:

VCM is the embedded manager within Flex-10/10D and FlexFabric interconnect modules. It is the core application for configuring and managing single and multi-enclosure VC domain environments. Access to VCM occurs over the same out-of-band Ethernet connection used to access the OA in BladeSystem c-Class enclosures and server blade iLO management engine connections on a separate management network.



Note

VCM manages a single domain of up to four enclosures. HPE OneView replaces VCM and VCEM when the enclosure is imported into HPE OneView. After the enclosure is managed by HPE OneView, VCM is no longer available and cannot be accessed.

VC Enterprise Manager (VCEM) is an optional software application that lets you manage a maximum of 250 VC domains (up to 1000 BladeSystem enclosures and 16,000 server blades). VCEM centralizes connection management and workload mobility for server blades that use VC to access LANs, SANs, and converged network infrastructures. It increases productivity and reduces operating costs by allowing data centers to respond faster to workload and configuration changes. Built on the VC architecture, it provides a central console to administer network address assignments, perform group-based configuration management, and rapidly deploy and move server connections.

HPE OneView is a central integrated management platform that enables IT teams to work and collaborate in a natural and automated way. HPE OneView also provides a software-based approach to lifecycle management, which automates operations to reduce the cost and the time required to deliver IT services. HPE OneView supports an open development platform designed to rapidly adapt to changing business needs. This programmable platform, built on the REST API, allows you to scale beyond data center walls, all the way to the cloud.

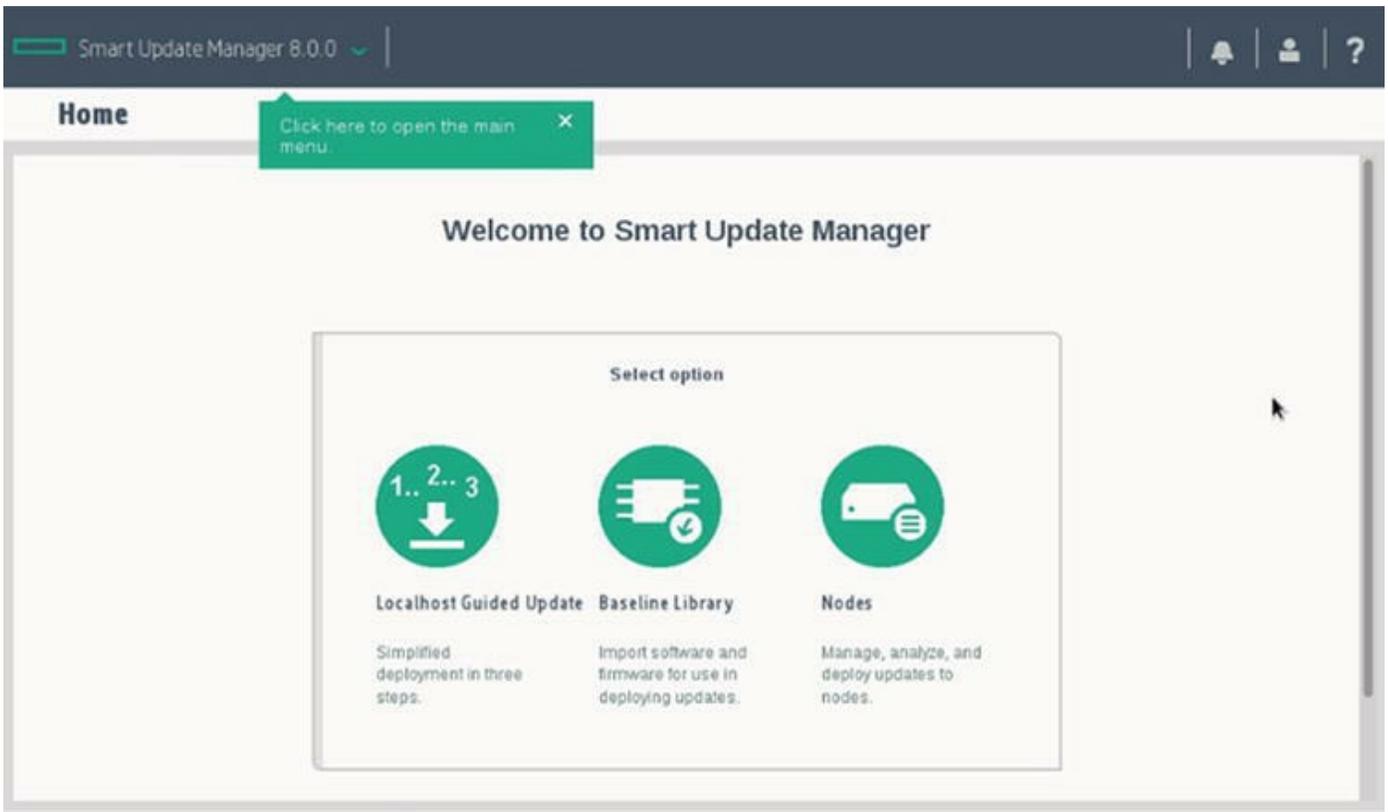
Administrators can also use the VC Support Utility (VCSU) to remotely upgrade VC Ethernet and Fibre Channel module firmware in BladeSystem c7000 enclosures. VCSU enables administrators to perform other maintenance tasks remotely using a stand-alone command line utility. When the utility initiates a firmware upgrade process, VCSU performs an automatic health check, and then all modules are updated at the same time. The utility displays a message indicating that an update is in progress. After the module firmware updates are complete, VCSU activates all of the modules.

**Note**

To download VCSU, enter the URL into your browser and search for “Virtual Connect Support Utility”.

<https://support.hpe.com/hpesc/public/home>

VC firmware



```
-----  
HP BladeSystem c-Class Virtual Connect Support Utility  
Version 1.7.1 (Build 21)  
Build Date: Nov 23 2012 05:51:02  
Copyright (C) 2006-2012 Hewlett-Packard Development Company, L.P.  
All Rights Reserved  
-----  
Please enter action ("help" for list):
```

Figure 7-24 HPE SUM home screen and VCSU CLI

The HPE BladeSystem c-Class VCSU enables administrators to upgrade VC-Enet and VC-FC firmware and to perform other maintenance tasks remotely on both BladeSystem c-Class c7000 enclosures using a standalone command line utility, as shown in [Figure 7-24](#). When the utility initiates a firmware upgrade process, VCSU performs an automatic health check, and then all modules are updated at the same time. The utility displays a message indicating that an update is in progress and the percentage completed. After the module firmware updates are complete, the utility activates all of the modules.

Service Pack for ProLiant

SUM supports VC firmware update. All of the VC interconnect modules need to be in sync for the update to be successful. Failure to achieve synchronization will cause the firmware update to be aborted.

VCSU can be used as a manual alternative to using Service Pack for ProLiant (SPP) for firmware update.

Download and install the VCSU utility.

VCSU utility supports a command line and an interactive mode.

VCSU modes

VCSU supports CLI and interactive mode. VCSU interactive mode is invoked without any parameters. Users are prompted to enter missing command line parameters. For all of the commands provided by VCSU, the username, password, and IP address of the enclosure OA must be provided. In some cases, the VC domain username and password must also be specified.

Display all of the modules in the enclosure and ensure that connectivity exists between the utility and the remote enclosure. Enter `version`.

Display all modules in the target enclosure and which modules can be updated. Enter `report`.

Before performing firmware update take a configuration backup using the `configbackup` command.

VCSU update process automatically performs a health check on the VC domain and will abort if discrepancies are discovered. Run the `healthcheck` command to determine the health status and correct any issues before running the update.

Update the firmware on supported VC-Enet and VCupdate.



Note

VCSU only updates supported VC modules that are able to be updated and require an update. The update process takes approximately 20 minutes for VC-Enet modules and five minutes for VC-FC modules. These updates are performed simultaneously.

VC with HPE OneView

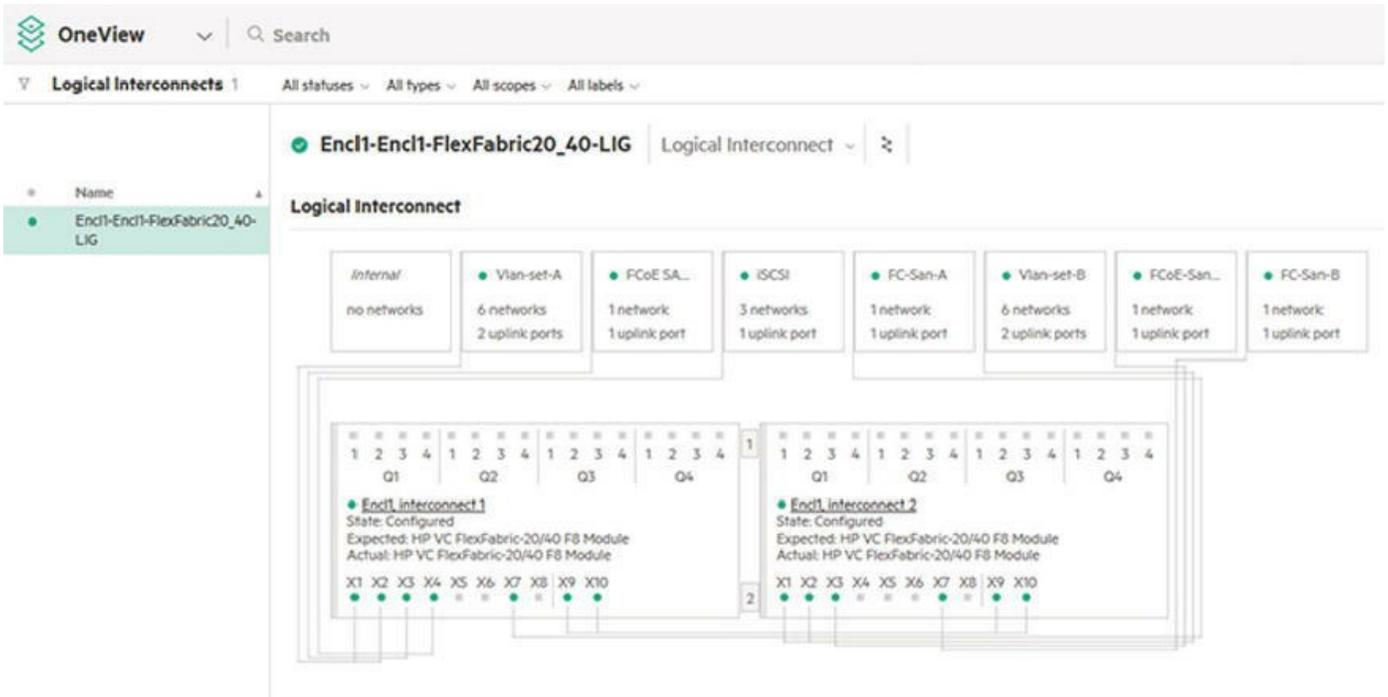


Figure 7-25 HPE OneView manages VC configuration

The HPE OneView management platform, shown in [Figure 7-25](#), becomes the VC Domain manager when the enclosure containing VC interconnect modules is added. OneView responsibilities include:

Enables and disables downlink and manages VC configuration

Deploys previously created configuration to newly added enclosures

Manages VC firmware

Associates networks and SAN fabric connections to the VC uplink connections

When OneView becomes the domain manager all direct access to the VCM GUI or command line is disabled

Learning check

What is a FlexNIC?

- A. A flexible LOM that provides RoCE blades
- B. A physical PCIe function that presents itself to the operating system as a discrete physical NIC
- C. A way of dividing a 10 Gb or 20 Gb port into up to eight partitions
- D. A virtual NIC contained in a software layer

A customer is rolling out a production environment that will consist of up to 30 c7000 enclosures. They are looking for a central management solution that will allow them to perform template-driven server deployment and enclosure lifecycle management. How should you respond?

Summary

BladeSystem solutions enable businesses to consolidate, virtualize, and automate their IT environments. This portfolio meets the changing needs of data center customers by:

- Allowing easier, faster, and more economical changes to server and storage setups without disrupting LAN and SAN domains
- Lowering application deployment and infrastructure operations costs by reducing the number of IT architecture variants
- Reducing connectivity complexity and costs
- Lowering purchase and operations costs when adding or replacing compute and storage capacity
- Supporting grid computing and SOA
- Supporting third-party component integration with well-defined interfaces

ProLiant BL server blades are independently functioning servers with all the necessary components integrated into a single board, including:

- Processors
- Memory
- Network adapters
- Optional hard drives

A BladeSystem enclosure serves as the foundation for a converged infrastructure. The BladeSystem c7000 enclosure offers power, cooling, network, and storage capabilities in one chassis.

Whether customers need basic connectivity for a remote office or a high-bandwidth, low-latency link for a high-performance computing cluster, BladeSystems offer an enclosure connectivity option to meet customer needs.

8 HPE Synergy

WHAT IS IN THIS CHAPTER FOR YOU?

After completing this chapter, you should be able to:

Provide a high-level overview of HPE Synergy.

List and describe the components of HPE Synergy:

- ✓ Composer
- ✓ Compute modules
- ✓ Frames
- ✓ Fabric
- ✓ Storage modules
- ✓ Image Streamer

Describe the Synergy management subsystem and HPE OneView support.

Prelearning check

A client is evaluating the HPE Synergy solution as a replacement for their multi-stacked c7000 enclosure solution. Among many questions about the Synergy product, they specifically are interested in the management solution redundancy capability. How should you respond?

A customer is evaluating a competitor's solution and has raised concerns that Synergy does not appear to have a solution for the rapid deployment of the host operating system to the Synergy compute nodes. How should you respond?

HPE Synergy

HPE Synergy is a single infrastructure of physical and virtual pools of compute, storage, and network fabric resources and a single management interface that allows IT to instantly assemble and reassemble resources in any configuration. Synergy eliminates hardware and operational complexity so IT can deliver infrastructure to applications faster and with greater precision and flexibility.

Emergence of bi-modal IT

Many CIOs and data center system administrators feel challenged to deliver traditional business applications while at the same time standing up new applications such as mobile and cloud-native apps that drive revenue. Ops-driven and cost-focused, traditional IT environments make it difficult to deliver faster value to the business—IT cannot move fast enough for today's application delivery goals.

Traditional business applications are built to run the business. They include applications such as enterprise resource planning (ERP) and other large databases that have been prepackaged and pretested and typically go through one or two release cycles a year. IT has been built around these for the last 20–30 years.

Apps-driven and agility-focused, new IT environments deliver a new type of apps and services that drive revenue and new customer experiences via mobility, Big Data, and cloud-native technologies. These new apps challenge IT to maintain a digital enterprise in a digital economy alongside traditional applications. However, maintaining two different sets of infrastructure—one designed for traditional apps and another designed for cloud-native apps—increases costs and complexity. This approach is not sustainable.

Gartner calls the strategy of maintaining an existing infrastructure for traditional applications while transitioning to infrastructure and tools for the cloud-native applications **bi-modal computing**. The HPE vision is to pull both together with one infrastructure that can deliver the applications while enabling the agility of infrastructure that customers are seeing today in the cloud but on-premises.

A Composable Infrastructure optimizes intelligence and automation via infrastructure-as-code to seamlessly bridge traditional and new IT environments for huge gains in application speed and operational efficiency. It offers a single infrastructure ready for any physical, virtual, or containerized workload including enterprise applications and cloud-native apps. The HPE Composable Infrastructure:

Reduces operational complexity

Accelerates application deployment

Is designed for today and architected for the future

The industry's first platform architected for composability

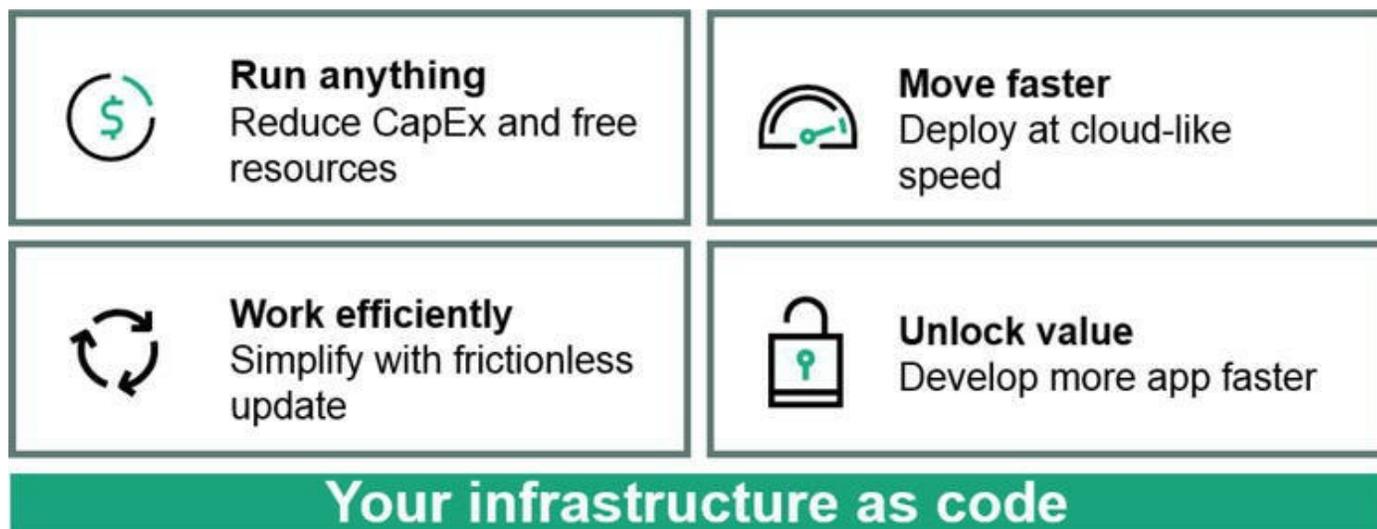


Figure 8-1 A new category of infrastructure to bridge traditional and cloud-native apps

HPE Synergy, the first platform built specifically for a Composable Infrastructure, offers an experience that empowers IT to create and deliver new value instantly and continuously. It reduces operational complexity for traditional workloads and increases operational velocity for the new breed of applications and services. Through a single interface, Synergy composes physical and virtual compute, storage, and network fabric pools into any configuration for any application. As an extensible platform, it enables a broad range of applications and operational models such as virtualization, hybrid cloud, and DevOps.



Note

For more information on DevOps see:

<https://en.wikipedia.org/wiki/DevOps>

With Synergy, IT can become not just the internal service provider but the business partner to rapidly launch new applications that become the business. As shown in [Figure 8-1](#), with Synergy, IT can continuously:

Run anything—Optimize any application and store all data on a single infrastructure with fluid pools of physical and virtual compute, storage, and network fabric.

Move faster—Accelerate application and service delivery through a single interface that precisely composes logical infrastructures at near-instant speeds.

Work efficiently—Reduce operational effort and cost through internal software-defined intelligence with template-driven, frictionless operations.

Unlock value—Increase productivity and control across the data center by integrating and automating infrastructure operations and applications through a unified application program interface (API).

Technically describe infrastructure-as-code—Provision bare metal infrastructure with one line of code—in the same way as virtual machines (VMs) and cloud.

The Synergy platform is unique because it enables customers to create a completely stateless infrastructure. Only Synergy allows configuration and provisioning of compute, fabric, storage, and hypervisor and operating system images as part of a single server profile template. This enables all of the pieces to be configured in one place without changing between multiple tools. It also allows the infrastructure to be provisioned and reprovisioned as needed based on the requirements of workloads that it is hosting.

Composable Infrastructure

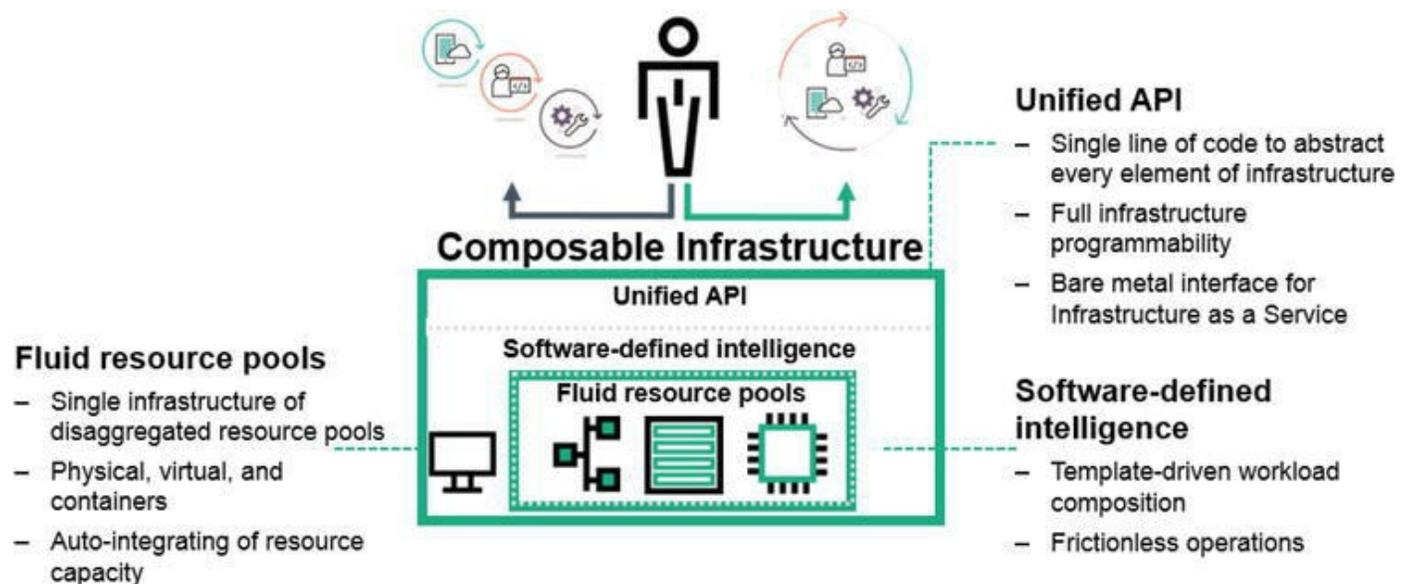


Figure 8-2 The three elements of a Composable Infrastructure

Synergy Composable Infrastructure delivers high performance and composability for the delivery of applications and services. It simplifies network connectivity using disaggregation in a cost-effective, highly available, and scalable architecture. Composable fabric creates a pool of flexible network fabric capacity that can be configured almost instantly to rapidly provision infrastructure for a broad range of applications.

A composable fabric architecture has three key elements as shown in [Figure 8-2](#).

A Composable Infrastructure starts with fluid resource pools, which is essentially a single structure that boots up, ready for any workload with fluid pools of compute, storage, and network fabric that can be instantly turned on and flexed.

The second element is software-defined intelligence, which means embedding intelligence into the infrastructure and using workload templates to tell it how to compose, recompose, and update quickly, in a very repeatable, frictionless manner.

The third element uses all of these capabilities and exposes them through a unified API, which allows infrastructure to be programmed like code so it can become infrastructure as a service.

By delivering a highly flexible, high-performance pool of composable resources centered around a unified API to simplify and speed up deployment and management, customers are able to derive value much faster than was possible with legacy solutions. This is useful for businesses who just want to update existing apps or code new apps and be able to derive infrastructure directly as code.

Synergy enables you to:

Compose on-demand with compute and storage to meet business needs

- Graceful addition of frames in single or multiple steps
- Software-defined intelligence using HPE OneView logical resources
- Logically viewed as a single flat Virtual Connect (VC) module

Implement seamless network connectivity using an easy-to-manage, cost-efficient, resilient fabric

- Extend networking to satellite enclosures without adding hops
- Ultra-low latency for east-west traffic with 2.56 Tb/s throughput in any configuration



Note

For more information on HPE Synergy, watch this video, which explains Synergy in two minutes. To access the video, enter the following URL into your browser:

<https://www.youtube.com/watch?v=8tbNdGa2iS0>

You can also watch an eight-minute video that provides a guided walkthrough of Synergy. To access the video, enter the following URL into your browser:

https://www.youtube.com/watch?v=_rCI4yFtLQ0

Increase productivity and control across the data center

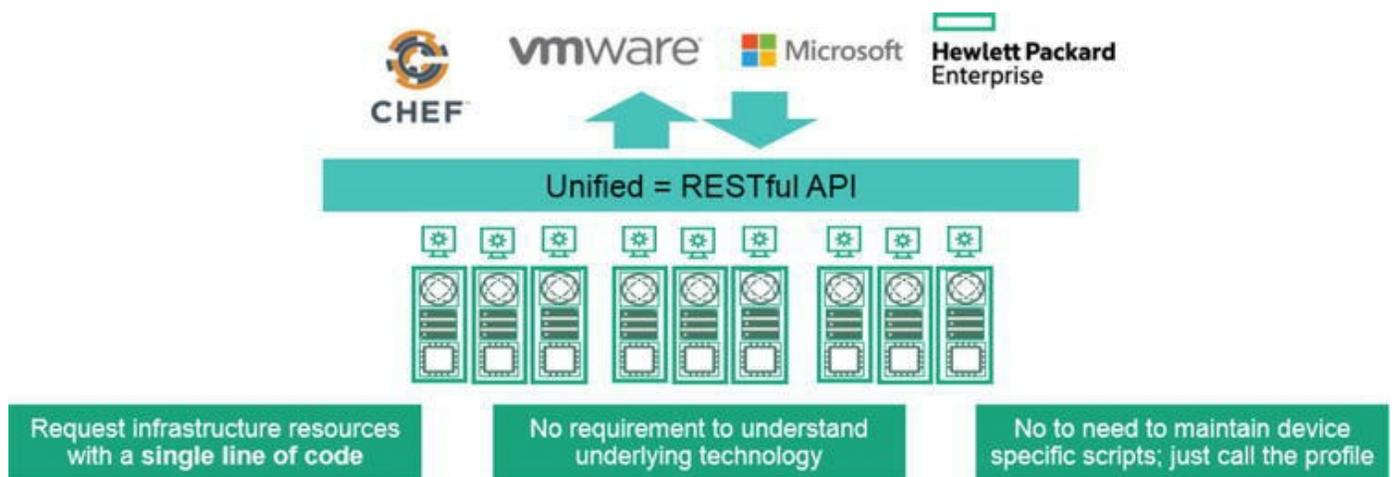


Figure 8-3 Unified API

As shown in [Figure 8-3](#), the Unified API enables the Synergy platform to be configured and controlled programmatically. All of the Synergy management functions delivered through the Synergy Composer are available through the RESTful API, which is compatible with the APIs that HPE OneView customers are using today. This means that application owners or IT operations teams using tools such as HPE Operations Orchestration, Chef, Ansible, VMware vCenter, or Microsoft System Center do not need to configure workflows that make low-level calls to many different device-specific APIs.

Whereas a single traditional server can require as many as 500 or more individual calls to low-level tools to get the infrastructure configured properly, you can now provision a complete infrastructure including the infrastructure personality through a single REST call, in a single line of code. The ability to programmatically invoke a server profile

template means that administrators do not need to be experts in the underlying infrastructure, and they do not have to maintain a large library of complex, device-specific scripts.

With a traditional infrastructure, every device with every manager has its own API. Each API, in turn, has its own API format, data format, and error code format. As a result, configuring this long list of different interfaces can be complex. However, automating the process involves automating across different interfaces with different data formats and different error codes, configuring individual bits and bytes.

Synergy eliminates that complexity with a high-level unified API that brings together all the resources, compute, storage, and network fabric under a single API with a single data format and a single error code format. Rather than having to configure bit by bit, byte by byte, you need only one line of code to have the server configured and running. This not only boosts the speed to provision, but also boosts the productivity of the operations team that is coding automation scripts, and developers who want to automate application infrastructure operations as well.

HPE Software, VMware, Microsoft, and others have a collection of capabilities in both the traditional IT world and in the Idea Economy world where applications are being developed quickly. Synergy integrates with tools such as Chef, Docker, and OpenStack that work the way that people want to develop infrastructure and applications. The single-unified API can plug into the traditional Ops-driven world with tools from VMware and Microsoft, or plug into the new apps-driven world to provision applications and infrastructure together.

Learning check

What are the three key elements of a composable architecture?

Synergy components

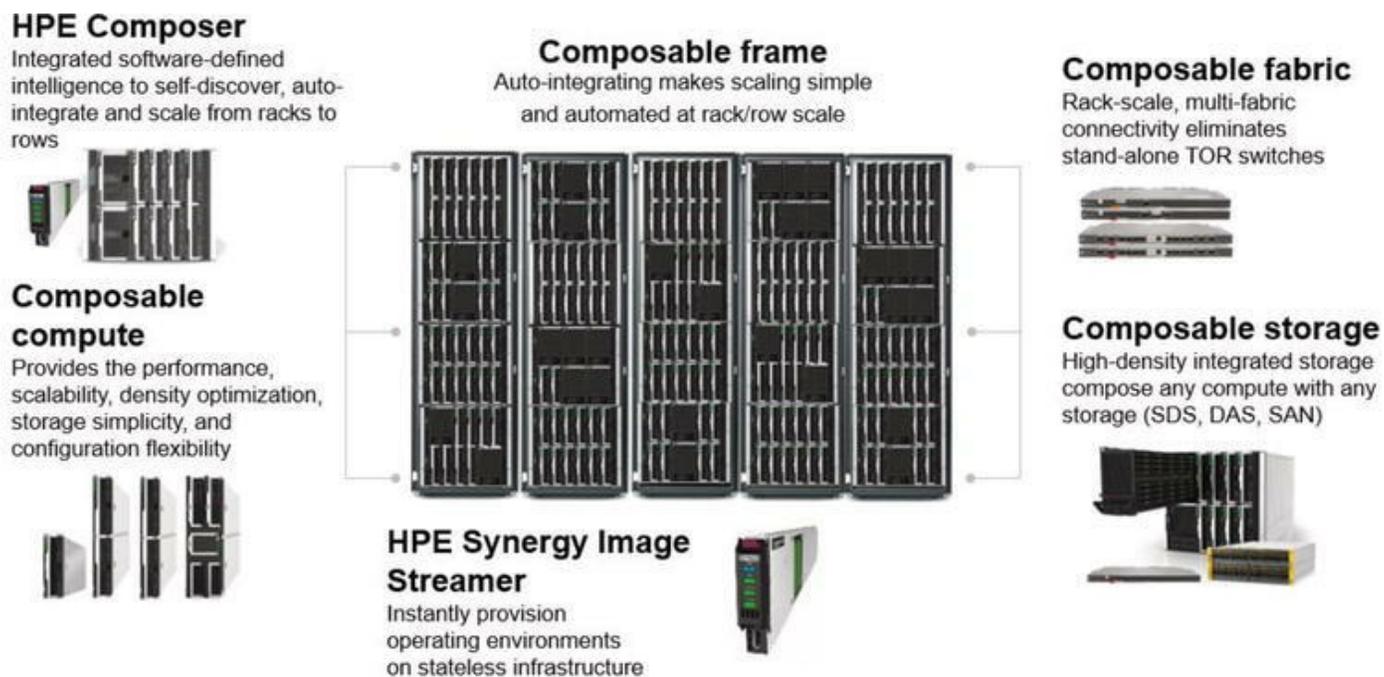


Figure 8-4 HPE Synergy components

With HPE Synergy, compute, storage, and fabric are always available as single pools of resources that can be instantly configured according to the specific needs of each application. As shown in [Figure 8-4](#), the key components of Synergy are:

Composer—Composer is an appliance that runs an embedded version of HPE OneView specifically for management of Synergy resources. It leverages integrated software-defined intelligence to accelerate operations through a single interface.

Image Streamer—Image Streamer enables stateless computing and very rapid provisioning, updating, and reprovisioning of server operating systems and applications. In effect, physical servers can be treated like VMs.

Compute—Compute capacity can be configured for physical, virtual, or container-based workloads.

Frame—Synergy frames may be linked into larger groups or domains of frames to form a dedicated management network, increasing resources available to the business and improving IT efficiency as the size of the Composable Infrastructure grows—achieving both CapEx and OpEx economies of scale.

Storage—Consolidate high-density integrated storage for direct-attached storage (DAS), storage area network (SAN), software-defined storage (SDS), and file, block, and object storage.

Fabric—Rack-scale multi-fabric connectivity eliminates stand-alone top-of-rack (ToR) switches.

Synergy Composer

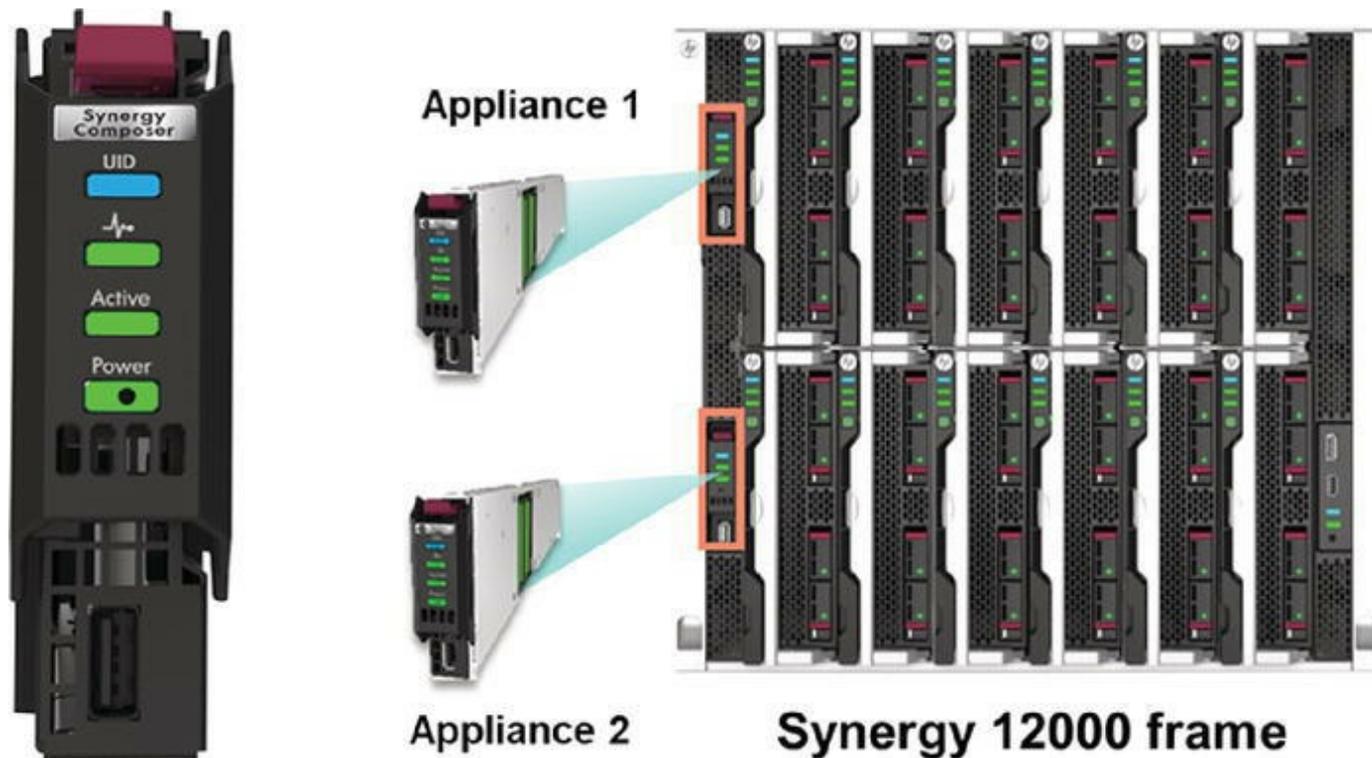


Figure 8-5 Synergy Composer plugs into an appliance bay in the Synergy frame

Composer is a physical appliance integrated within the Synergy frame. It plugs into an appliance bay in the side of the Synergy frame, as shown in [Figure 8-5](#).



Note

Composer does not use a compute module slot.

Composer provides a single interface for assembling and reassembling flexible compute, storage, and fabric resources in any configuration. Its infrastructure-as-code capability accelerates transformation to a hybrid infrastructure and provides on-demand delivery and support of applications and services with consistent governance, compliance, and integration.

Composer uses an embedded version of the HPE OneView management solution to

manage compute modules, fabrics, and storage, which is the essence of software-defined intelligence in Synergy.

Composer deploys, monitors, and updates the infrastructure from one interface and one unified API. It allows IT departments to deploy infrastructure for traditional, virtualized, and cloud environments in a single step, in just a few minutes. Resources can be updated, flexed, and redeployed without service interruptions. This allows infrastructure to be deployed and consistently updated with the right configuration parameters and firmware versions, streamlining the delivery of IT services and the transition to a hybrid cloud. Its reduced complexity and faster service delivery times ultimately enable IT to better respond to changing business needs.

Server templates are a powerful way to quickly and reliably update and maintain existing infrastructure. Composer uses templates to simplify one-to-many updates and manage compute module profiles. This feature adds inheritance to the process, meaning updates can be made once, in the template, and then propagated out to the profiles created from that template. Elements that can be updated via a template include firmware, BIOS settings, local RAID settings, boot order, network configuration, shared storage configuration, operating systems, and many others.

Composer templates also provide **monitor and flag** capabilities with remediation. Profiles created from the template are monitored for configuration compliance. When inconsistencies are detected, an alert is generated indicating that the profile is out of compliance with its template. When a new update is made at the template level, all profiles derived from that template will be flagged as inconsistent. From there, the user has complete control over the remediation process for bringing individual modules or multiple systems back into compliance.

HPE Synergy Image Streamer



Figure 8-6 HPE Synergy Image Streamer

The Synergy Image Streamer, shown in [Figure 8-6](#), represents a new approach to deployment and updates for Composable Infrastructure. This management appliance works with HPE Synergy Composer to provide fast, software-defined control over physical compute modules with operating system and application provisioning. Image Streamer enables truly stateless computing combined with instant-on capability for quick deployment and updates.

Image Streamer is an optional appliance that plugs in to the appliance bays on the side of the Synergy frame. Image Streamer is always implemented as redundant pairs of appliances to provide a highly available boot environment. Image Streamer can be accessed from the unified API for programmatic control by qualified users.

Image Streamer and Composer offer a combined set of features that can be used to rapidly deploy and update infrastructure:

Redundant Image Streamer appliances with failover and high availability

Auto-discovery of Image Streamer by Composer

Software-defined integration with server profiles

Stateless operation with IP addresses assigned to bootable images

Highly available boot and run environment

Private cloud readiness with VMware ESXi and Docker-enabled Linux

Highly available image archive and compliance to the latest verified images

Tools for personalization and customization of images

Golden image editing capabilities

User privacy with rights and privileges from Composer

Programmatic access through the unified API

These capabilities support IT areas that need highly available image storage or programmatic access and control over the infrastructure.

Synergy compute module portfolio

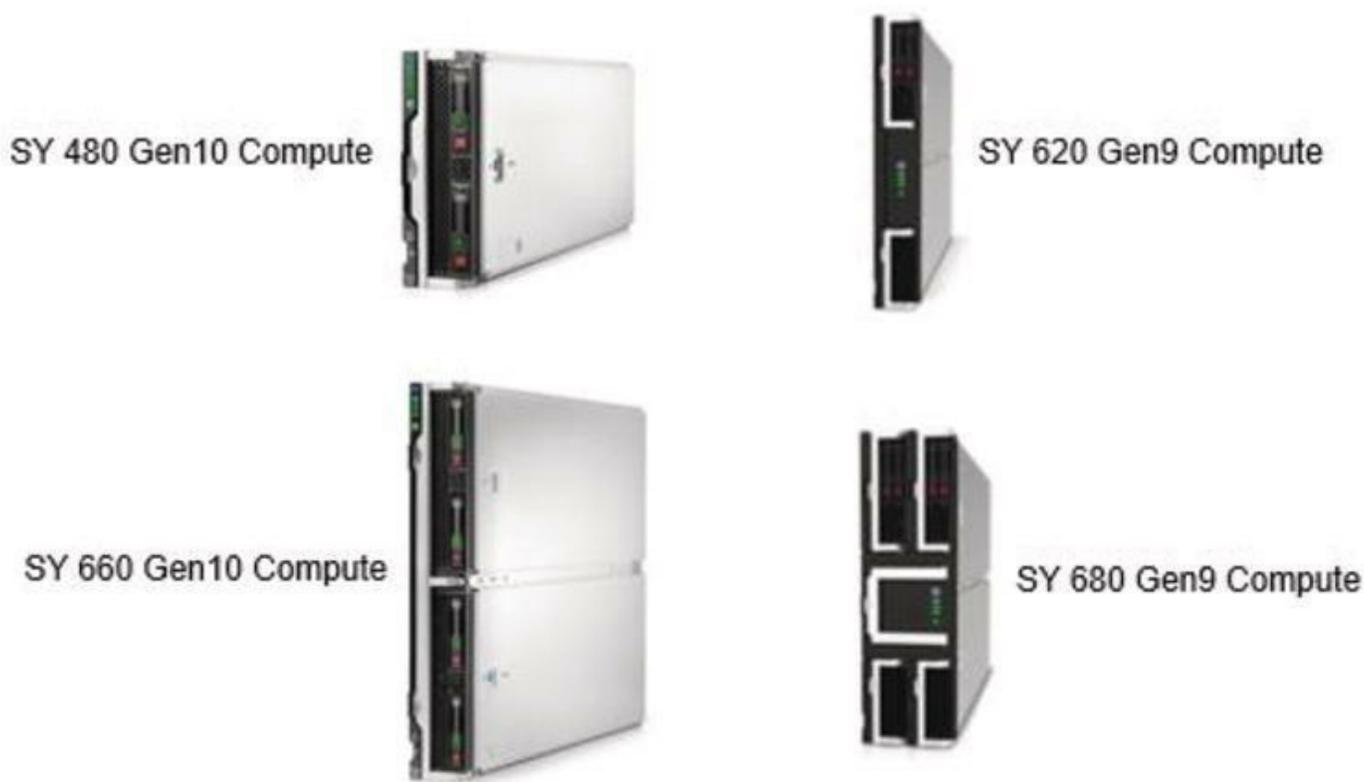


Figure 8-7 Synergy compute Gen10 and Gen9 modules

The flexible design of Synergy composable compute modules optimizes internal storage choices to match workload requirements. [Figure 8-7](#) shows the Synergy compute module portfolio. Synergy supports both two-socket and four-socket compute modules, which provide the performance, scalability, density optimization, storage simplicity, and configuration flexibility to power a variety of workloads, including business processing, IT infrastructure, web infrastructure, collaborative, and high-performance computing.

The SY 480 Compute Module has a flexible front bezel design with two HPE

SmartDrive bays, each of which supports up to two drives per bay or four micro form factor drives with built-in RAID with a choice of serial-attached SCSI (SAS), serial AT attachment (SATA), hard disk drive (HDD), and flash storage. The SY 660 includes four HPE SmartDrive bays for double the amount of internal storage when compared to the SY 480. If internal storage is not required, such as in environments using the Synergy Image Streamer, there is an option to order compute modules with no drives for a stateless configuration.

For workload acceleration, PCIe NVMe SSDs from HPE combine the high throughput and low latency of a storage workload accelerator with the flexibility and convenience of a standard form factor drive, ideal for web scaling, cloud, OLTP, Big Data, or business intelligence. Synergy also supports software-defined storage services and externally attached storage options.

The compute modules converge traffic over high-speed 10/20 Gb connections using HPE Converged Network Adapters (CNAs). When connected to an HPE Virtual Connect SE 40Gb F8 Module, each module provides up to eight adjustable connections (seven data and one storage, or all data) to each 20 Gb compute module port.

HPE Synergy 480 Gen10 Compute Module

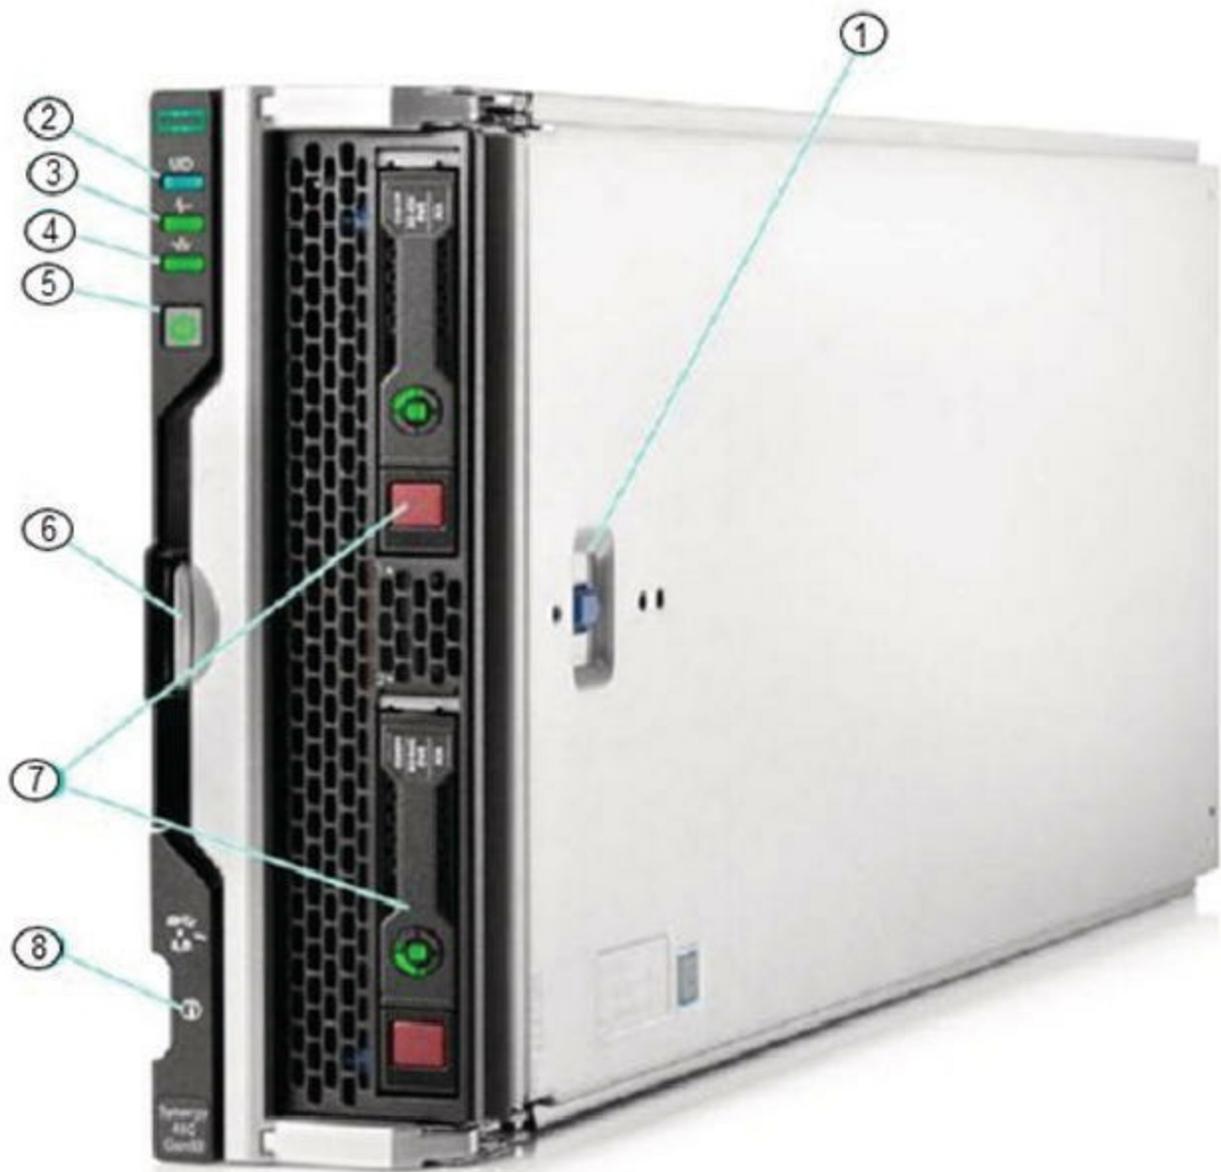


Figure 8-8 HPE Synergy 480 Gen10 front view

[Figure 8-8](#) shows the front view and features of the HPE Synergy 480 Gen10 Compute Module. Numbered components are:

Quick Access Panel

UID LED

Health Status LED

Mezzanine NIC status LED

Power On/Standby button and system power LED

Compute module handle release latch

Removable drive cage with two hot-plug drive bays

External USB 3.0 connector & iLO USB connection (behind serial label pull tab)

The HPE Synergy 480 Gen10 Compute Module delivers superior capacity, efficiency, and flexibility in a two-socket, half-height (HH) form factor to support demanding workloads. It delivers superior capacity, efficiency, and flexibility to power more demanding workloads and increases VM density by providing a full range of processor choices, right-sized storage options, and a simplified IO architecture. It is designed to optimize general-purpose enterprise workload performance including business processing, IT infrastructure, web infrastructure, collaborative, and high-performance computing in physical and virtualized environments while lowering costs within a Composable Infrastructure.

The HPE Synergy 480 Gen10 Compute Module has 24 DIMM slots and memory capacity up to 3.0 TB with 128 GB LRDIMMs. It supports the entire Intel new Intel Xeon Scalable Family of processors—3100, 4100, 5100, 6100, and 8100 series.

Key features and advantages of the HPE Synergy 480 Gen10 Compute Module, when compared to the Gen9 equivalent, include:

Increased performance and scalability with more processor choices and more cores and higher internal transfer rates

Improved designs for more memory capacity per processor, along with higher speeds and Persistent Memory

Expanded flexible storage and capacity with flexible drive solutions and more RAID options

Greater IO bandwidths across PCIe mezzanine slots

HPE Performance with IST and Persistent Memory

Industry-leading security

Performance gains with new HPE Smart Array and HPE Smart Memory DDR4 2666 MHz

High-performance, flexible networking with up to 25 Gb or 50 Gb IO to each compute module

Greater NVMe PCIe Storage capacity

Broad graphics processing unit (GPU) support for density and performance

The HPE Synergy 480 Gen10 Compute Module is ideal for enterprise customers with

Workloads including virtualization, containers, Big Data, storage-centric Apps, data warehousing/analytics, CRM, ERP, virtual desktop infrastructure (VDI), SAP, and large storage capacity, such as Microsoft Exchange, and so forth.

Gen10 vs. Gen9 quick comparison: HPE Synergy 480 Compute Module

Table 8-1 Comparison of HPE Synergy 480 Gen9 and Gen10

Specifications	Gen9	Gen10
Processor	Intel® Xeon® E5-2600 v3 product family Intel® Xeon® E5-2600 v4 product family	Intel® Xeon® Processor Scalable family up to 205W) Platinum/Gold/Silver/Bronze product family
Processors/cores/speed	One or two processors; 22, 20, 18, 16, 14, 12, 10, 8, 6, 4 cores; 3.5 GHz	One or two processors; 28, 26, 24, 22, 20, 18, 16, 14, 12, 10, 8, 6, 4 cores; 3.6 GHz
Memory (type, max slots)	Support up to 2400 MT/s DDR4 SmartMemory; 1.5 TB Max with 64 GB DIMMs, 24 DIMM slots	Support up to 2666 MT/s DDR4 Smart Memory; 3.0 TB Max with 128 GB DIMMs, 24 DIMM slots
Persistent Memory	None	16 GB NVDIMM, (12 total DIMMs are NVDIMM enabled)
Drive Bays	(Two) SFF SAS/SATA/SSD, (Two) SFF NVMe SSD or (Four) uFF	(Two) SFF SAS/SATA/SSD, (Two) SFF NVMe SSD, (Four) uFF or (2) M.2 SATA
DAS	213 TB max: Up to 71 drives per controller	1.5 PB max: Up to 200 drives per controller, in frame

Network Controllers*	HPE Synergy 3820C 10/20 GbE CNA 2 Ports; 2820C 10 GbE CNA 2 ports; HBAs: 3530C 16G FC 2 ports; 3830C 16G FC 2 ports	HPE Synergy 3820C 10/20 GbE CNA 2 Ports; 2820C 10 GbE CNA 2 ports; HBAs: 3530C 16G FC 2 ports; 3830C 16G FC 2 ports; 10 Gb PTM 6810c 25/50 Ethernet
Storage Controllers*	HPE SATA B140i; Smart Array P240nr or HBA H240nr, HPE Smart Array P542D	Software RAID—S100i 6G SATA; E208i-c HBA, 12G SAS; SA RAID—E208i-c, 12G SAS and P416ie-m Hybrid HBA/RAID Mezzanine 12G SAS
Graphics	Mezzanine MXMs: NVIDIA M6, M3000SE and AMD S7100X; Multi-MXM Expansion Module with NVIDIA or AMD MXM options; or PCIe Expansion Module with NVIDIA or Quadro M5000 PCIe GPUs	Mezzanine MXM: NVIDIA M6, M3000SE and AMD S7100X MXM; Multi-MXM Expansion Module with NVIDIA or AMD MXM options; or PCIe Expansion Module with NVIDIA M60, M10, or Quadro P6000 GPUs
Management	OneView 3.0, HPE iLO 4, Insight Control	OneView 3.x, HPE iLO 5, Insight Control
Power and Cooling	Frame-based (96% Platinum)	
Warranty	3/3/3—3-year parts, 3-year labor, and 3-year onsite	

[Table 8-1](#) presents a comparison between the Synergy 480 Gen9 and Gen10 Compute Modules.

HPE Synergy 660 Gen10 Compute Module

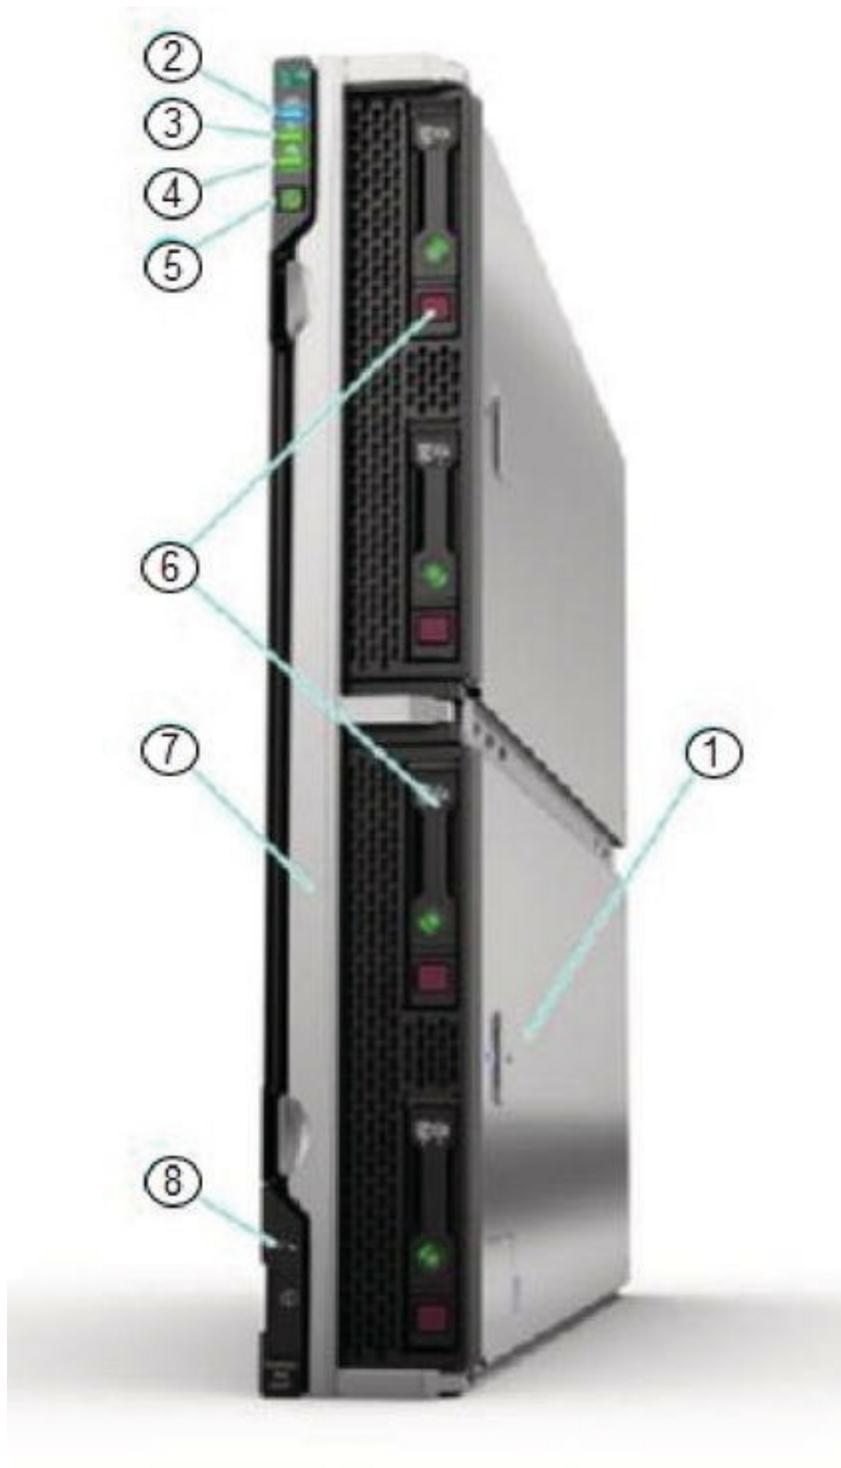


Figure 8-9 HPE Synergy 660 Gen10 front view

[Figure 8-9](#) shows the front view and features of the HPE Synergy 660 Gen10 Compute Module. Numbered components are:

Quick Access Panel

UID LED

Health Status LED

Mezzanine NIC Status LED

Power On/Standby Button and System Power LED

Zero (0) to Four (4) Hot-Plug Drive Bays (or up to eight uFF with Smart Drive bay adapter)

Compute Module Handle Release Latch

External USB Connector and iLO USB (behind serial label pull tab)

The HPE Synergy 660 Gen10 Compute Module delivers agility, control, and security in a four-socket, full-height form factor to support demanding workloads and virtualization density. Handles data-intensive workloads with uncompromised performance and exceptional value. It supports, high-performance, high-density Xeon Scalable Family of processors, 48 DIMM slots providing up to 6 TB of available memory, flexible IO fabric connectivity, and right-sized storage options. It offers a large memory footprint and powerful processors that are required to support more demanding workloads like in-memory and structured databases. Its HPE DDR4 SmartMemory 2666 MHz offers up to a 25% performance increase over the previous generation.

Key features and advantages of the HPE Synergy 660 Gen10 Compute Module include:

Extreme compute performance in a pool of composable resources to run any enterprise workload

Increased performance with 25/50G networking, Persistent Memory, and NVMe capability

Next-generation reliability and availability for multi-workloads

Fast, highly automated deployment and update of compute, storage, and fabric workload resources using HPE Synergy Composer powered by OneView

Four-socket composable form factor and maximum DDR4 memory footprint for scale-up and data-intensive workloads

Persistent Memory for faster database and analytics workloads

Fast, simplest provisioning and updating via highly automated, template-based, workload-centric profiles

High-performance, flexible networking with up to 25 Gb or 50 Gb IO to each compute module

The HPE Synergy 660 Gen10 Compute Module is ideal for:

Large-scale server consolidation and virtualization where price/performance and large memory footprints are key drivers

Large-scale database workloads with core-based licensing

Efficient performance for HPC workloads requiring high-power CPU capabilities

Larger memory footprint for virtualization, data management, front-end business, and data analytics

Gen10 vs. Gen9 quick comparison: HPE Synergy 660 Compute Module

Table 8-2 Comparison of Synergy 660 Gen9 and Gen10

Specifications	Gen9	Gen10
Compute	Two or four Intel® Xeon® Broadwell processors	Two or four Intel® Xeon® Processor Scalable family
Memory	6 TB max: HPE Smart Memory (48) DDR4, up to 2400 MT/s	6 TB max: HPE Smart Memory (48) DDR4, up to 2666 MT/s
Persistent Memory	None	16 GB NVDIMM, (24 total DIMMs are NVDIMM enabled)
Local Storage	<ul style="list-style-type: none"> – Zero to four SFF SAS/SATA/HDDs or NVMe SSDs (12 TB max) – Up to eight uFF drives; – One USB 3.0, and one internal USB 3.0 and uSD 	<ul style="list-style-type: none"> – Diskless – Zero to four SFF SAS/SATA/HDDs or NVMe SSDs (12 TB max) – Up to eight uFF drives – Two USB 3.0, and one internal USB 3.0 and uSD – up to four internal M.2 drives for more diskless or dual boot options (up to 7.68 TB)
DAS	213 TB max: Up to 71 drives per	1.28 PB max: Up to 200 drives per controller, but 160 fit in frame

	controller	
Storage Controllers	Price: HPE SATA B140i Performance, price: HPE Smart Array P240nr or HPE HBA H240nr External, price, performance: HPE Smart Array P542D	– Software RAID—S100i 6G SATA – Essential RAID—E208i-c HBA, 12G SAS – Performance RAID—P408i-c, 12G SAS 8in Lanes-2G Cache, and P416ie-m Mezzanine 12G SAS 8in/8ex Lanes-2G Cache
GPU Support	None	None
Networking	Six x16 PCIe 3.0 mezzanines: 40 Gb Master and 20 Gb Satellite ICMs; 16 Gb FC	Six x16 PCIe 3.0 mezzanines: 25/50/100 Gb Master and 25/50 Gb Satellite ICMs; 25/50/100 Gb Ethernet; 10 Gb PTM; 16/32 Gb FC
Management	OneView 3.0, HPE iLO 4, Insight Control	OneView 3.x, HPE iLO 5, Insight Control
Power and Cooling	Frame-based (96% Platinum), ASHRAE A3, Energy Star	
Warranty	3/3/3—3-year parts, 3-year labor, and 3-year onsite	

[Table 8-2](#) presents a comparison between the Synergy 660 Gen9 and Gen10 Compute Modules.

HPE Synergy 620 and HPE Synergy 680 Gen9 Compute Modules

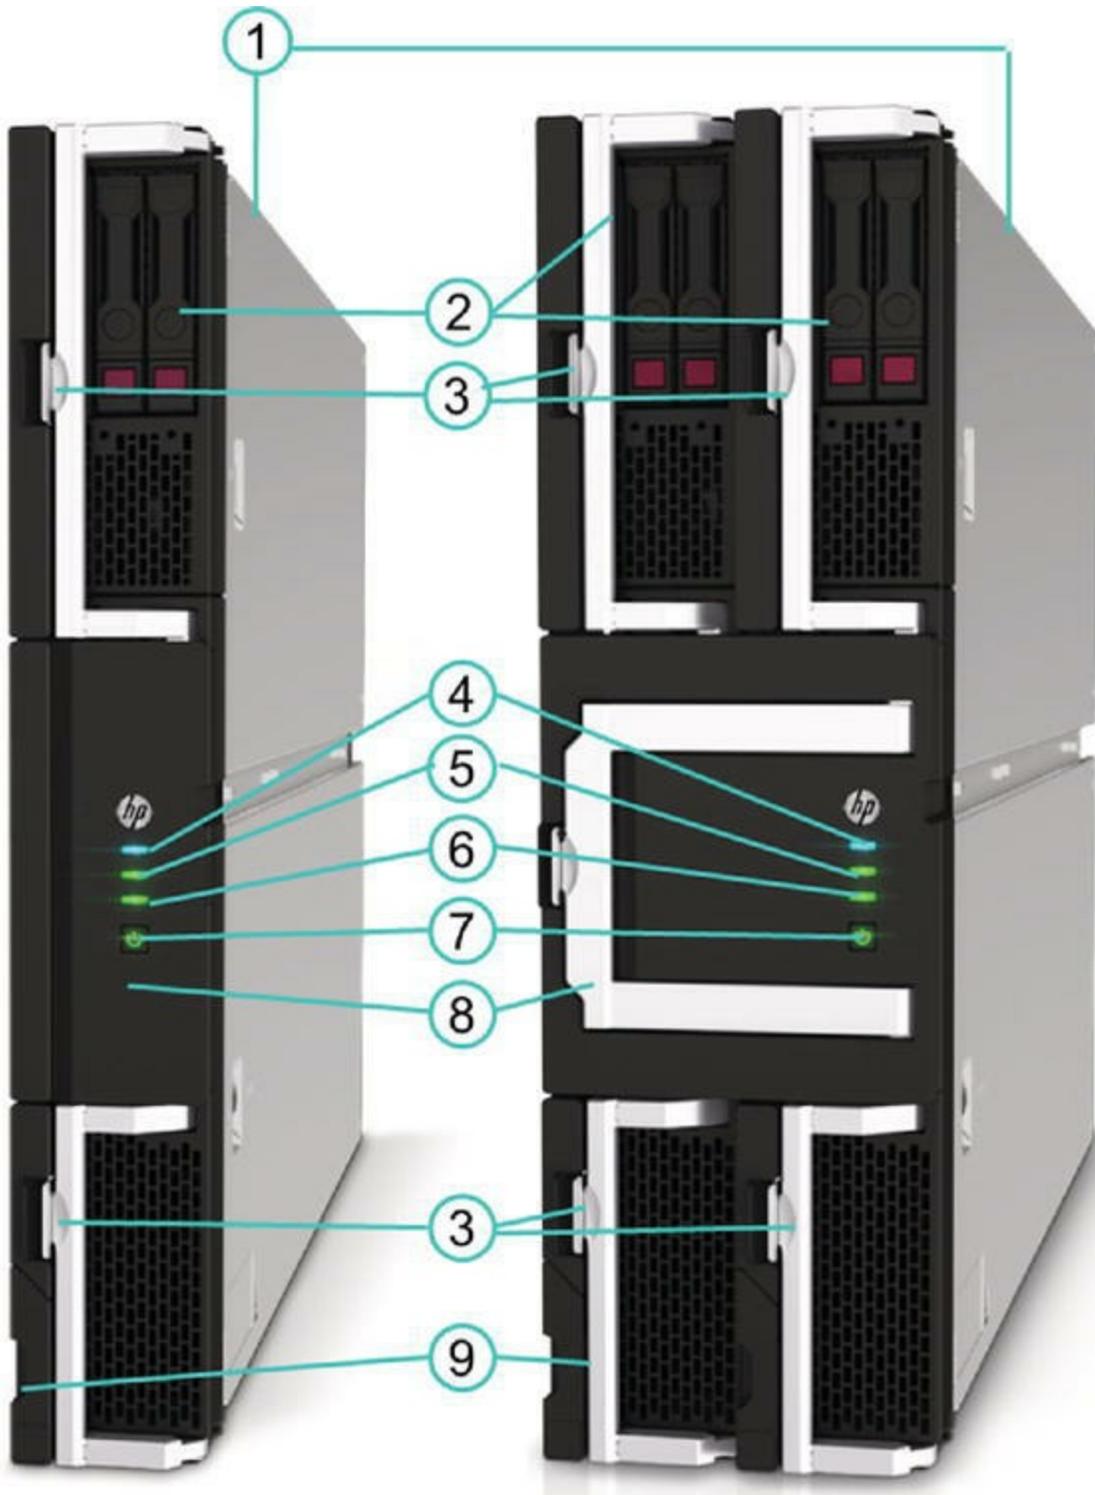


Figure 8-10 HPE Synergy 620/680 Gen9 Compute Modules front view

[Figure 8-10](#) shows the front view and features of the HPE Synergy 620 and 680 Gen9 Compute Modules. Numbered components are:

Quick Access Panel

Hot-plug drive bays for SFF, uFF, or NVMe drives (diskless also available)

Compute module release latches

UID LED

Health status LED

Mezzanine NIC status LED

Power On/Standby button and system power LED

Compute module Link* (release latch only on four-socket system)

External USB 2.0 connector (behind serial label pull tab)

*The compute module link enables scaling from a two-socket to a four-socket system.

These, full-height, single-width and double-width compute modules support the full range of Intel Xeon E7-4800 and E7-8800 v4 processors. They are designed to meet the needs of almost any enterprise IT tier and high-density VM workload. These two-socket and four-socket x86 compute modules are ideal for financial, insurance, healthcare, manufacturing, and retail enterprises that require mission-critical levels of availability, extended versatility, and real-time performance.

Run any workload better with HPE Synergy Gen10

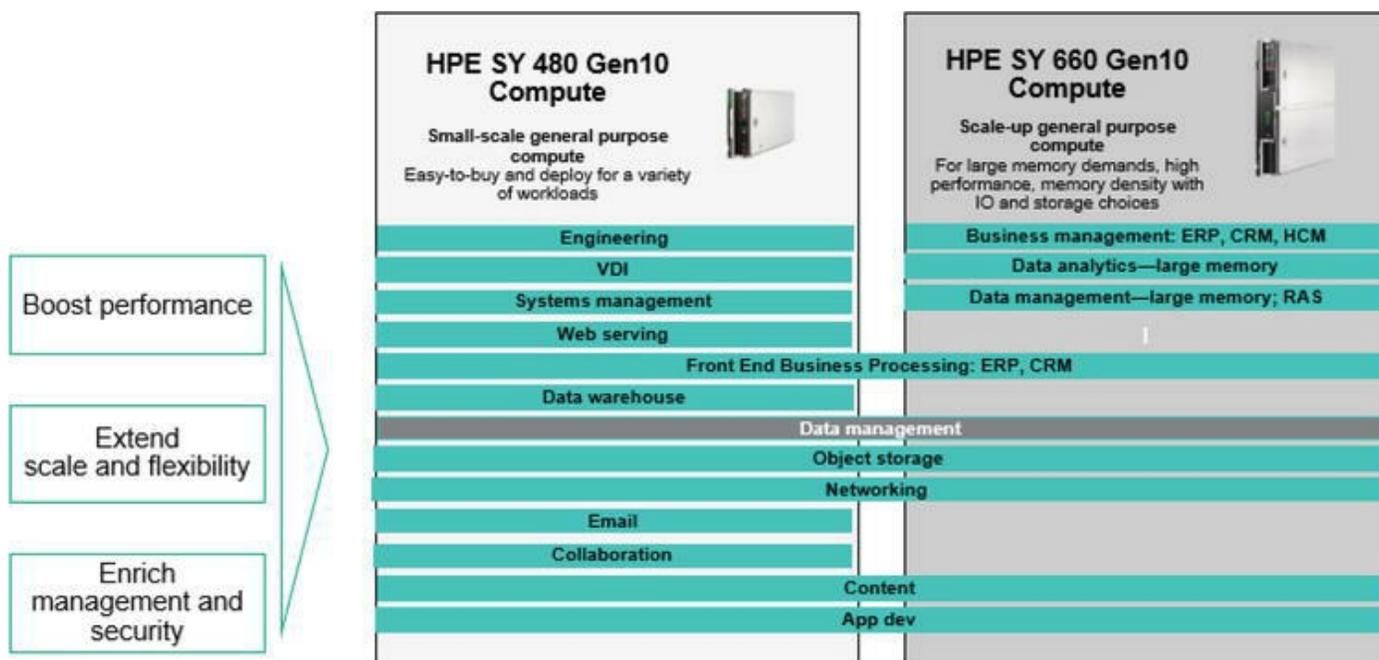


Figure 8-11 HPE Synergy SY480 Gen10 and SY660 Gen10 comparison

The HPE Synergy SY480 Gen10 Compute Module delivers superior capacity, efficiency, and flexibility in a two-socket, HH form factor to support demanding workloads. Ideal for virtualization, containers, Big Data, and unmatched performance for hybrid-IT solutions. In addition, support for GPU expansion modules and more GPU options to further enable VDI workloads.

The HPE Synergy SY660 Gen10 Compute Module delivers a four-socket, full-height form factor for higher performance and scalability for demanding, enterprise data-intensive, and cloud-application workloads. In addition, it has twice the processor, memory, IO capacity for the most data-intensive apps.

[Figure 8-11](#) lists workload examples for SY 480 Gen10 and SY 660 Gen 10 Compute Modules.

SY 480 Gen10 and SY660 Gen 10 compute modules support these additional features:

Composable solution optimized for any workload

Single, software-defined management

Built-in intelligence for auto-discovery and frictionless scaling

Rapidly provision infrastructure and applications in minutes

Flexible configuration with local storage or truly stateless compute

Persistent Memory and fast bandwidth for data-intensive applications

Learning check

A customer is evaluating a competitor's solution and has raised concerns that Synergy does not appear to have a solution for the rapid deployment of the host operating system to the Synergy compute nodes. How should you respond?

HPE Synergy 12000

The Synergy 12000 frame reduces complexity through an intelligent auto-discovery infrastructure and delivers performance to accelerate workload deployment. As the building block for a Synergy infrastructure, a Synergy frame offers substantial expansion and scalability. Each frame supports up to three fully redundant fabric bays for many fabric interconnect module (ICM) combinations.

Synergy 12000 frame front view

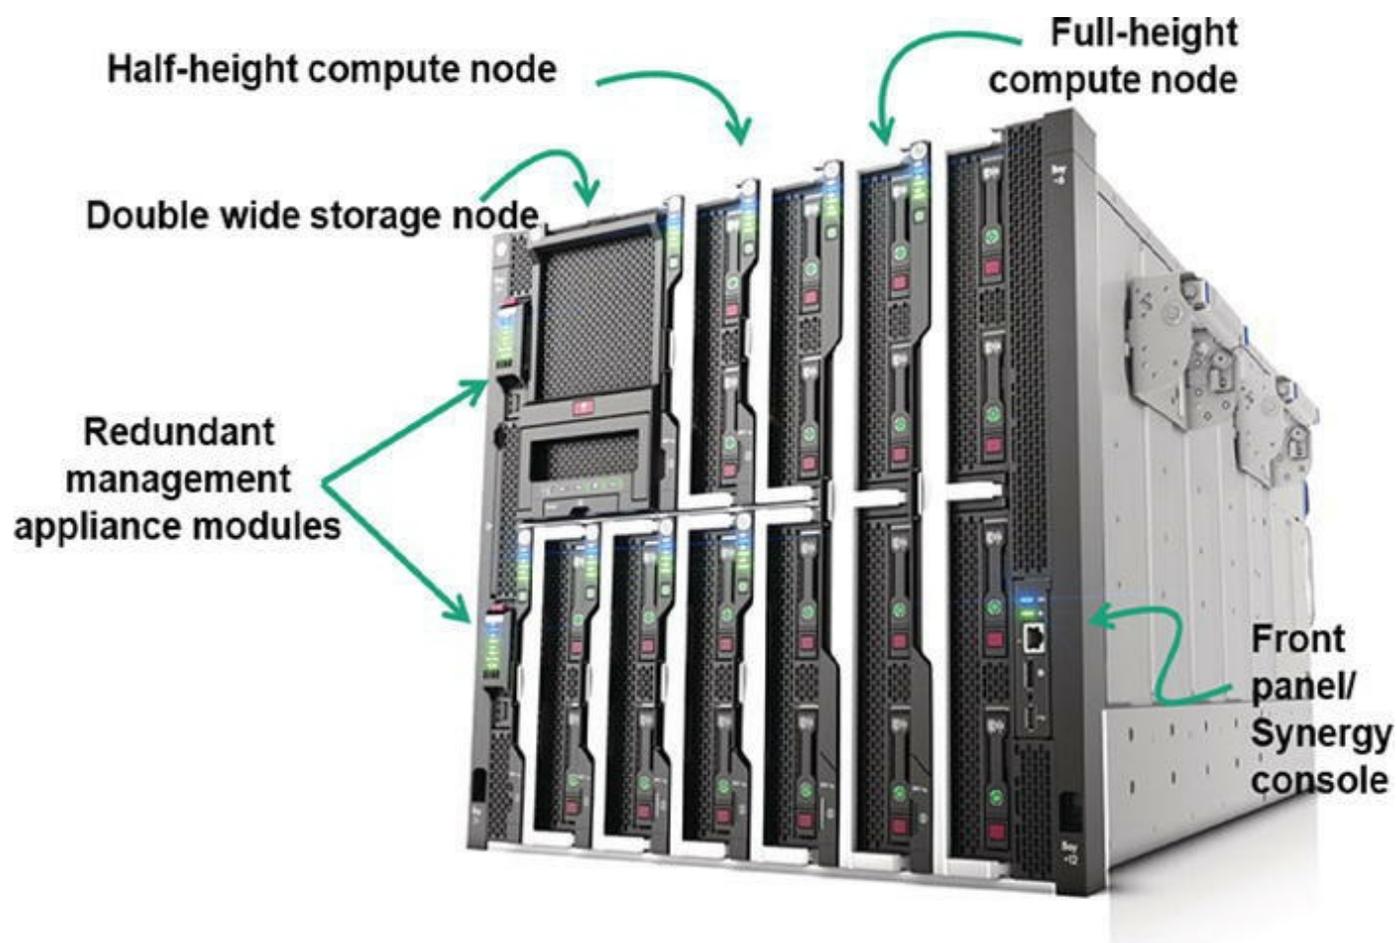


Figure 8-12 Synergy 12000 frame front view

The HPE Synergy 12000 frame shown in [Figure 8-12](#) is a key element of Synergy. The frame accommodates compute, storage, network fabric, and management in a single infrastructure to significantly reduce cost and complexity while delivering significant performance gains to accelerate workload deployment. It provides the base for an intelligent infrastructure with embedded management and scalable links for expansion

as business demand requires.

The HPE Synergy 12000 frame is a uniquely designed system that physically embeds management as code into an intelligent infrastructure to offer management and composability of integrated compute, storage, and fabric resources. Whether resources are in a single frame or multiple linked frames, the system offers composability of all resources.

A Synergy frame's unique design physically embeds Synergy Composer with HPE OneView management software to compose compute, storage, and fabric resources in any configuration. Synergy frames may be linked into larger groups or domains of frames to form a dedicated management network, increasing resources available to the business and IT efficiency as the size of the infrastructure grows—achieving both CapEx and OpEx economies of scale.

The Synergy 12000 frame helps make installation easy by using the standard power feeds of blade systems today. It supports up to six 96% Titanium Efficient, 2650W power supplies that offer redundant N + N, N + 1 power setup. The Synergy frame features an efficient cooling system with 10 built-in fans included in every frame.

Synergy 12000 frame rear view

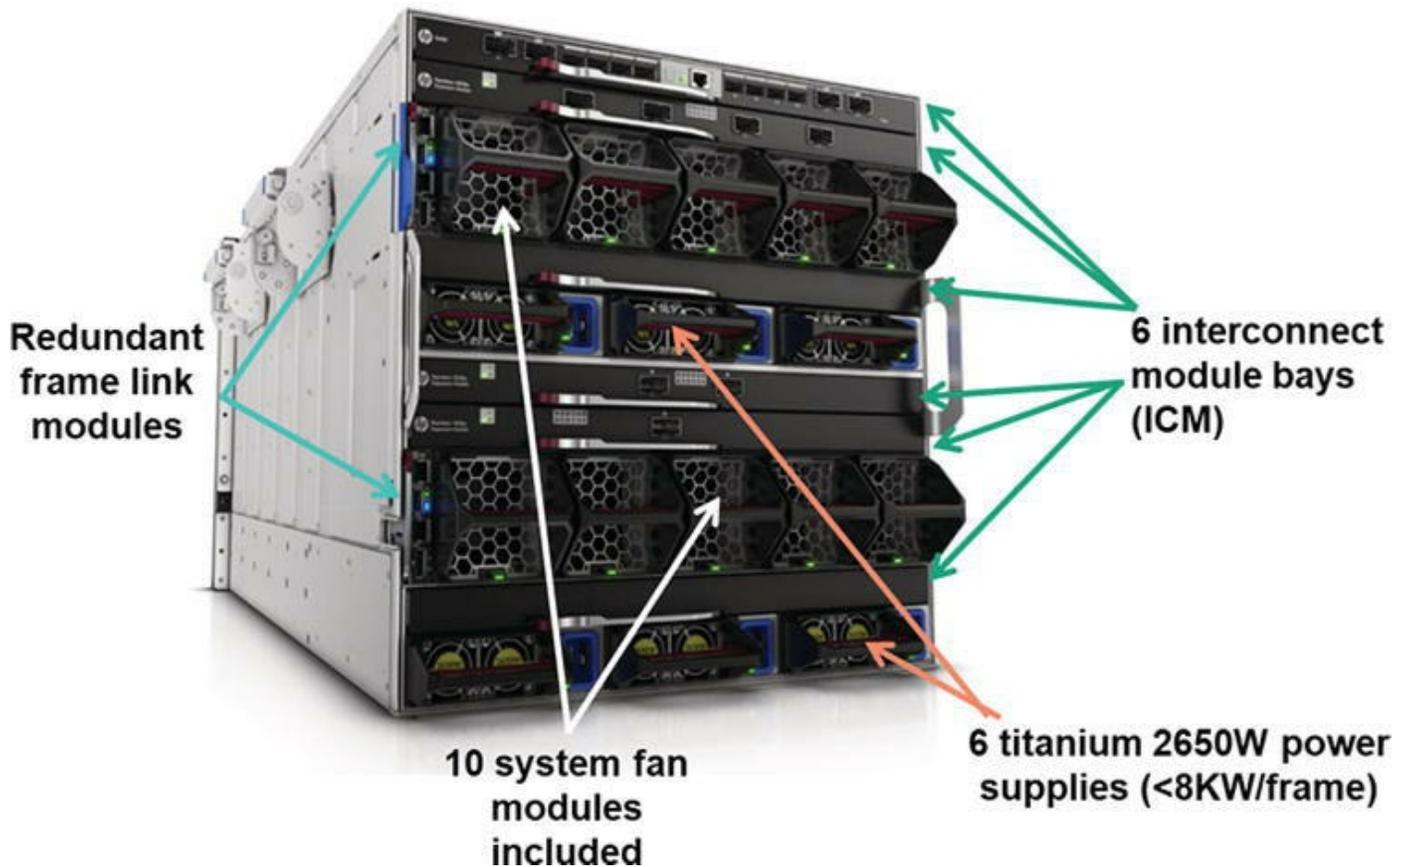


Figure 8-13 Synergy 12000 frame rear view

As shown in [Figure 8-13](#), a Synergy frame includes:

Redundant frame link modules (FLMs) (one is supplied, the second is optional)—provides links to multiple frames through a private air-gapped management network

10 hot-plug system fan modules—provides cooling for the frame

Up to six titanium power supplies—provides 7950W of redundant power line support

Six ICM bays—for full redundancy of 3 fabrics

HPE Synergy fabric configuration

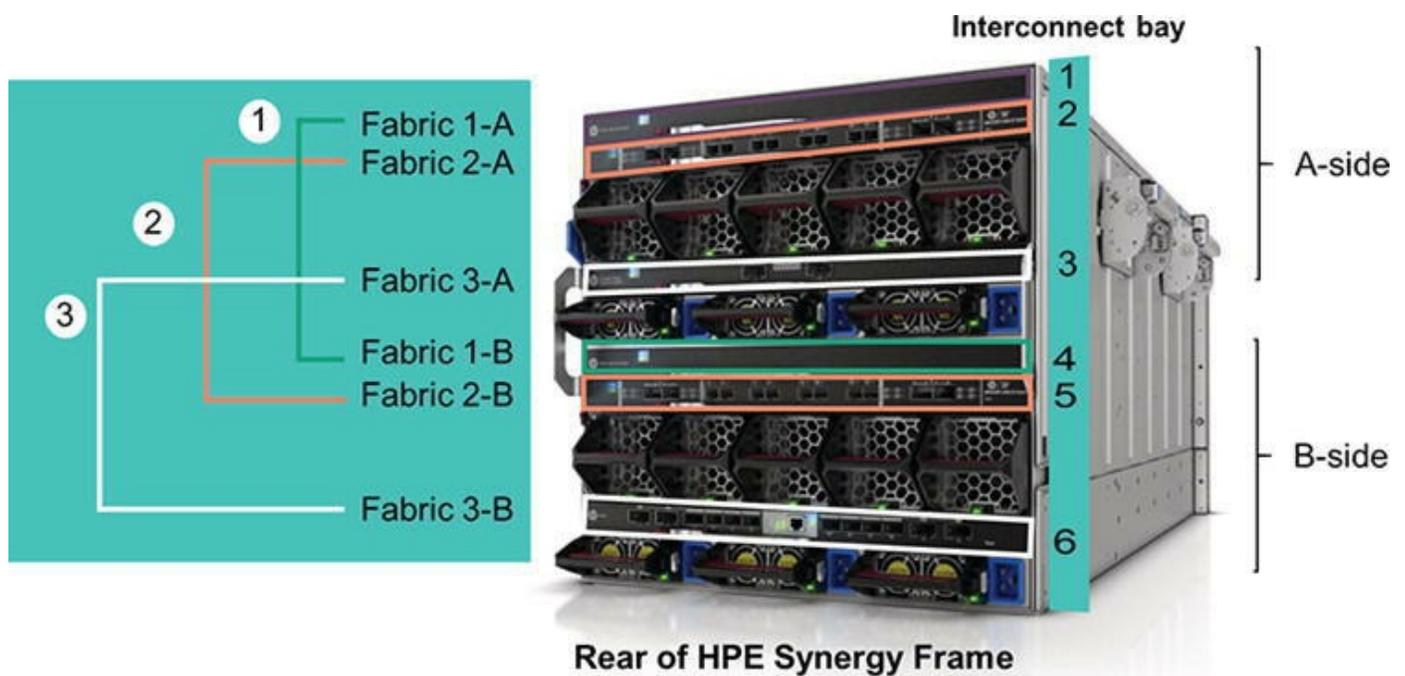


Figure 8-14 HPE Synergy supports three redundant fabrics

[Figure 8-14](#) shows that Synergy supports three redundant fabrics:

Fabric 1 with ICMs in interconnect bay 1 (fabric 1A) and 4 (fabric 1B)

Fabric 2 with ICMs in interconnect bay 2 (fabric 2A) and 5 (fabric 2B)

Fabric 3 with ICMs in interconnect bay 3 (fabric 3A) and 6 (fabric 3B)

You can deploy three redundant fabrics supported by SAS modules, Ethernet modules (including both composable and traditional options), or fiber channel modules (including both composable and traditional options), or a combination of all three.

The SAS modules provide compute modules with connectivity to in-frame storage modules (HPE Synergy D3940). When these modules are required, you must place them in Fabric 1.

Ethernet modules provide Ethernet and perhaps also converged storage connectivity. Start by using Fabric 3 for these modules. Then move to Fabric 2 and Fabric 1 if you need to add more compute module adapter cards and supporting fabrics. The Ethernet modules also permit a multiframe architecture using master and satellite modules distributed across multiple frames.

Managing HPE Synergy

Software-defined architecture auto-discovers and self-assimilates all Synergy resources for immediate use with template-driven operations. Template-driven Synergy management reduces manual interaction and human error by implementing change operations automatically and capturing best practices into templates.

The unique design of a Synergy Frame physically embeds Synergy Composer powered by HPE OneView to compose compute, storage, and fabric resources in any configuration. Synergy frames may be linked into larger groups or domains of frames to form a dedicated management network, increasing resources available to the business and IT efficiency as the size of the infrastructure grows—achieving both CapEx and OpEx economies of scale.

Synergy management subsystem

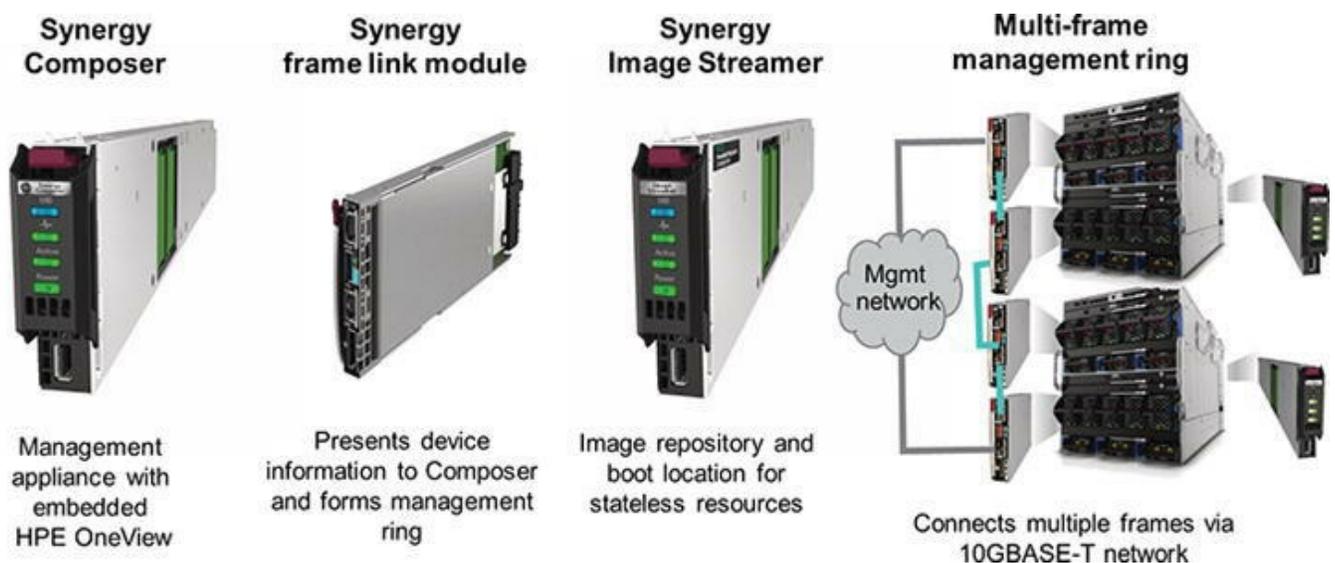


Figure 8-15 Synergy management subsystem

As shown in [Figure 8-15](#), the management subsystem comprises Synergy Composer, Synergy Image Streamer, and FLMs. A single Synergy Composer manages one frame or multiple racks of frames linked through the FLMs. The number of Synergy Composers selected determines the number of frames that can be linked and managed. The use of two Synergy Composer modules is recommended for redundancy and high availability.

Synergy Composer is a management appliance that directly integrates into the frame of the system and provides a single interface for assembling and re-assembling flexible

compute, storage, and fabric resources in any configuration. Its infrastructure-as-code capability accelerates transformation to a hybrid infrastructure and provides on-demand delivery and support of applications and services with consistent governance, compliance, and integration.

Synergy Image Streamer is a new approach to deployment and updates for Composable Infrastructure. This management appliance works with Synergy Composer for fast software-defined control over physical compute modules with operating system provisioning. Image Streamer enables true stateless computing combined with instant-on capability for deployment and updates. This management appliance deploys and updates infrastructure at extreme speed.

Synergy Image Streamer



Figure 8-16 Synergy Image Streamer

To accelerate workload deployment, Synergy can use the optional Image Streamer, as

shown in [Figure 8-16](#), a physical appliance repository of bootable (golden) images that can be streamed across multiple compute modules in a matter of seconds. This unique capability enables Synergy to set up and update infrastructure quickly and consistently. This is significantly faster than the traditional, sequential process of building compute modules—physical provisioning followed by operating system or hypervisor installation. It is ideal for instances such as web-scale deployments where IT needs to provision an operating environment across a large number infrastructure blocks.

Traditional server deployment is a sequential process of provisioning the physical hardware, followed by provisioning an operating system or hypervisor. Traditional memory-based server deployments use general deployment/provisioning tools for service operating system deployment, which uses a RAM-based operating system and is also known as a **pre-boot (pre-install) environment**.

Image Streamer enables a true stateless operation by integrating server profiles with golden images (operating environment and IO drivers) and personalities (operating system and application) for rapid implementation onto available hardware. The fast deployment and compliance management capabilities leverage software-defined intelligence and are accessible via the Unified API. These capabilities set HPE Image Streamer apart from traditional approaches.

Updates to highly replicated physical compute nodes with their operating environments at extreme speeds enables Image Streamer to deliver fast virtualized image changeovers (for use in Test and Dev, DevOps, multiple PaaS) or secure boot and image compliance (for use in defense, government, or financial services institutions). These capabilities are ideal for web-scale deployments where IT needs to provision an operating environment across a large number infrastructure blocks.

Image Streamer ensures high availability by providing redundant repositories of bootable images, which are used in a secure manner. These golden images can be rapidly cloned to create unique bootable images for compute nodes. It enables Synergy to quickly deploy a new compute module or update an existing one. This is far faster than the traditional, sequential process of building servers—physical provisioning followed by operating system or hypervisor installation, IO drivers, and application stacks.

Administrators using Image Streamer can design bootable images for compute nodes, with the operating system and application stacks included, for ready-to-run environments.



Important

Image Streamer requires a minimum of three Synergy frames with redundant Composers for operation and must be implemented as redundant pairs. This minimal system requires four cables, two transceivers, and one ICM for complete operation.

Synergy FLM



Figure 8-17 Synergy FLM

The Synergy FLM, shown in [Figure 8-17](#), is the integrated frame resource information control point and link to multiple frames. FLMs are critical to auto-discovery and growth.

Synergy FLMs provide a dedicated 10 Gb air-gapped management network for resource composability and information, health, and management reporting. The FLMs provide control points for providing resource and health information of the frames to the management appliances. The FLM reports asset and inventory information for the devices in the frame as well as thermal and power information, including real-time actual power usage per module and per frame.

A management port on each FLM provides access to the management network. Link

ports on each FLM link modules together for multi-frame setups. This frame link topology for linked systems offers debug and diagnostics benefits to the systems.

Synergy scales up to 20 frames for efficient growth. You can add multiple frames linked together through a simple yet secure management network to establish a management ring. An infrastructure of multiple frames offers substantial resource capability that can be managed as easily as a single frame. There are two FLM slots in each frame for full redundancy. One FLM is included in every Synergy Frame, and a second can be added for redundancy and to enable multiframe connectivity.

Management architecture

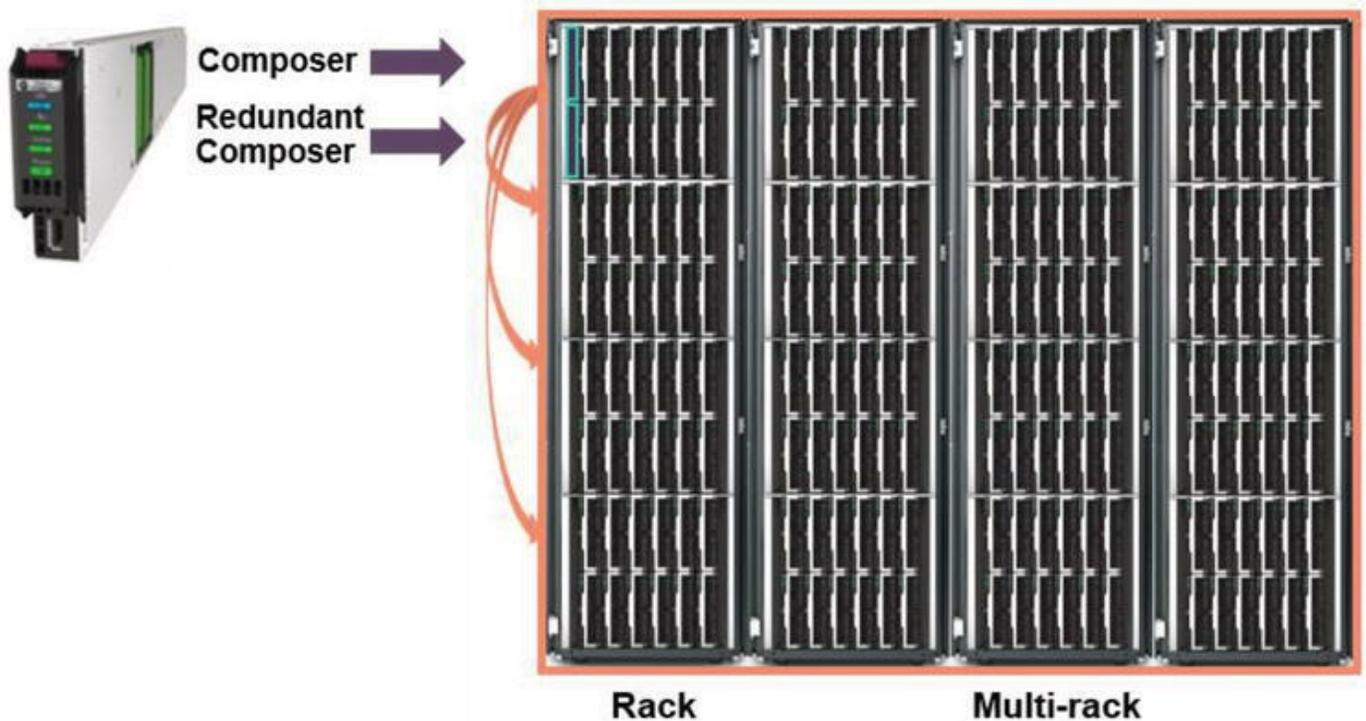


Figure 8-18 Management architecture

Synergy Composers are embedded management controllers that use HPE OneView to provide separated management that is built into the platform. This separated enclosure management does not impact compute modules during updates. Because there is no separate network with separate ports, the result is lower cost over scale, isolated from the data plane. This architecture enables you to update the frames without touching the compute modules. Additional benefits include:

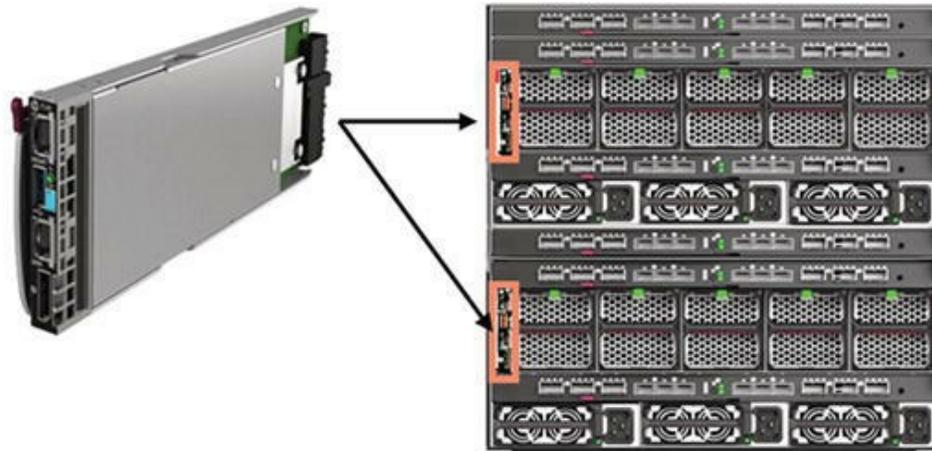
Redundant/clustered Composers as shown in [Figure 8-18](#)

Ability to federate up to 21 enclosures in a single instance for scale

Integrated out-of-band management for safety, fewer ports, and lower cost

You can manage up to 21 frames by installing management appliances. You also can scale over time by configuring any unused appliance slots as image managers or future orchestrators running on one of these appliances.

Reliable, efficient multi-frame management



Rear of frame

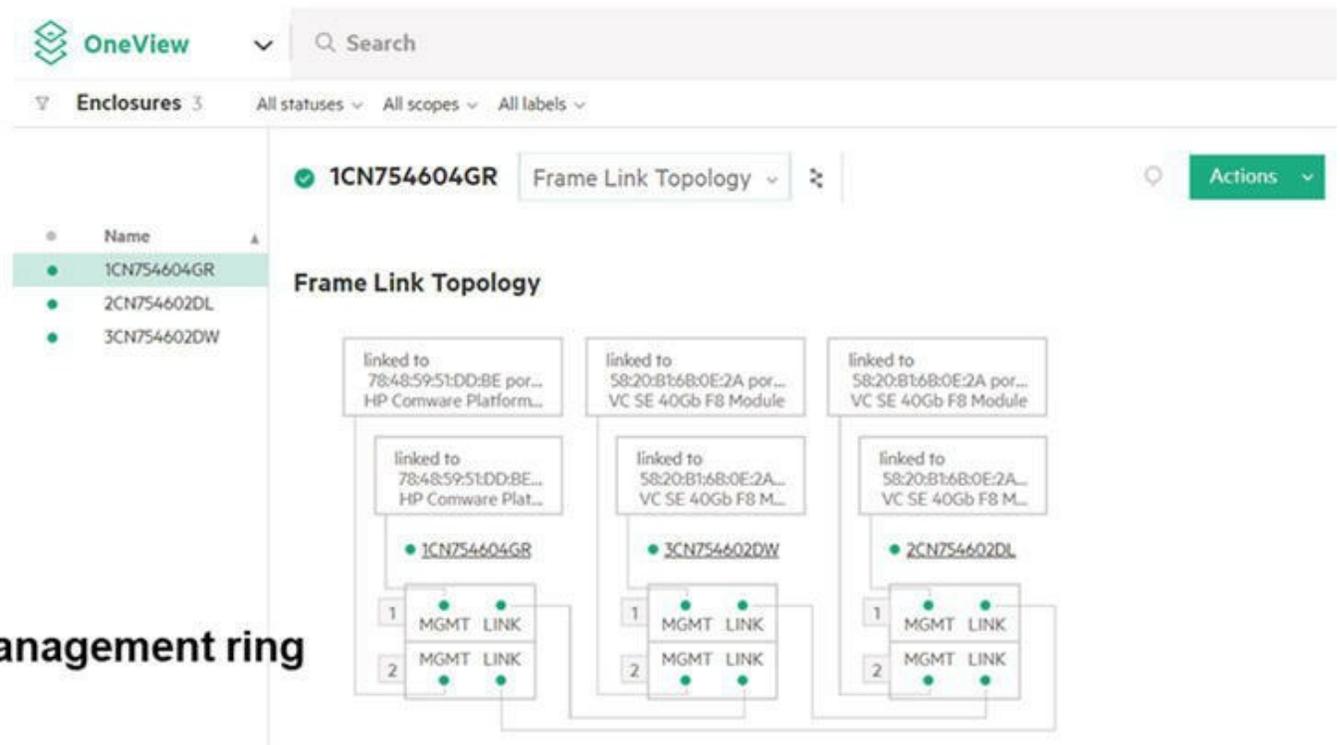


Figure 8-19 HPE Synergy FLM multi-frame management

The top of [Figure 8-19](#) shows where the FLMs are installed in the rear of the Synergy frame. The bottom part of [Figure 8-19](#) shows the frame link topology as viewed in HPE OneView. The FLM link ports (LINK) are connected in a daisy chain, and the management ports (MGMT) are connected to VC Modules and switches connected to the data center management network.

**Note**

Synergy Image Streamer requires the use of the MGMT port from both of the FLMs in the frame where the Image Streamer appliance is installed to be connected to a VC Module, as shown in [Figure 8-19](#). When Image Streamer is not present, the MGMT ports connect only to upstream management network switches (labeled HP Comware Platform in [Figure 8-19](#)).

Learning check

A client is evaluating the HPE Synergy solution as a replacement for their multi-stacked c7000 enclosure solution. Among many questions about the Synergy product, they specifically are interested in the management solution redundancy capability. How should you respond?

Match the Synergy management subsystem component with its function.

Composer	The resource controller and management link that forms the multi-module management ring
Image Streamer	Provides a single interface for assembling and reassembling flexible compute, storage, and fabric resources in any configuration
Frame link module	A repository of bootable images

HPE Synergy fabric and ICMs

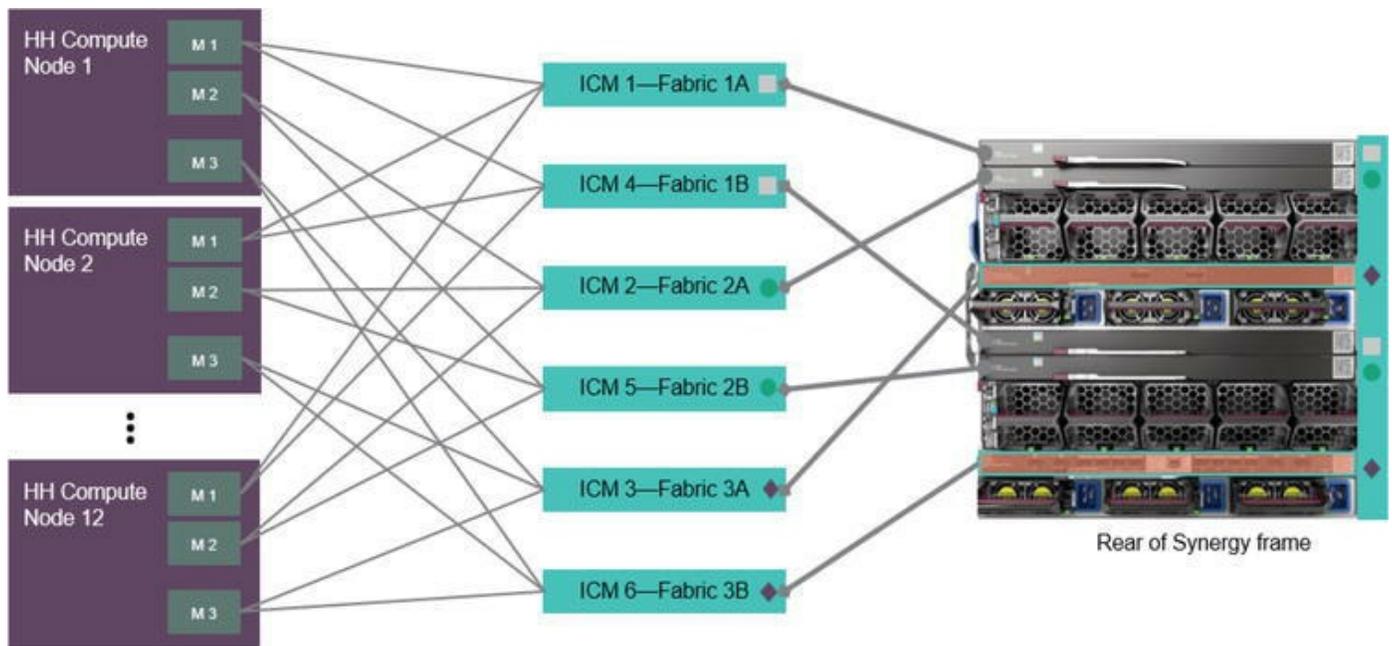


Figure 8-20 HPE Synergy frame ICM mapping

[Figure 8-20](#) gives an overview of HPE Synergy frame ICM mapping.

Each HH compute module can be configured with one, or up to three, mezzanine card interface modules, occupying slots M1-M3. The Synergy backplane supports two high throughput interface ports for each of the mezzanine card slots and are connected to a pair of redundant ICM bays.

The interconnect bays are numbered one to six, top to bottom. Fabric one is constructed from a redundant pair of ICMs that are located in bays one and four, fabric two is formed from interconnect bays two and five, and fabric three consists of interconnect bays three and six. Fabric three forms the default Ethernet fabric for one or up to five Synergy frames, creating a disaggregated, rack-scale fabric.

The disaggregated, rack-scale Synergy design uses a master/satellite architecture to consolidate data center network connections, reduce hardware and management complexity, and scale network bandwidth across multiple frames. This architecture enables Synergy frames to establish a single logical fabric, interconnecting compute modules at high speeds and low latency.

The master modules provide the uplinks for the complete logical fabric, consolidating the connections to the data center network. By reducing hardware requirements, eliminating the need for ToR switches, and consuming fewer ports at the aggregation layer, this architecture reduces costs and management complexity. The master module contains the intelligent networking capabilities that extend to satellite frames through the satellite modules.

Synergy composable fabric provides high performance and composability for the delivery of applications and services. It simplifies network connectivity using fabric disaggregation in a cost-effective, highly available, and scalable architecture. Synergy composable fabric creates a pool of flexible fabric capacity that is configurable almost instantly to provision infrastructure for a broad range of applications.

Synergy frictionless scaling

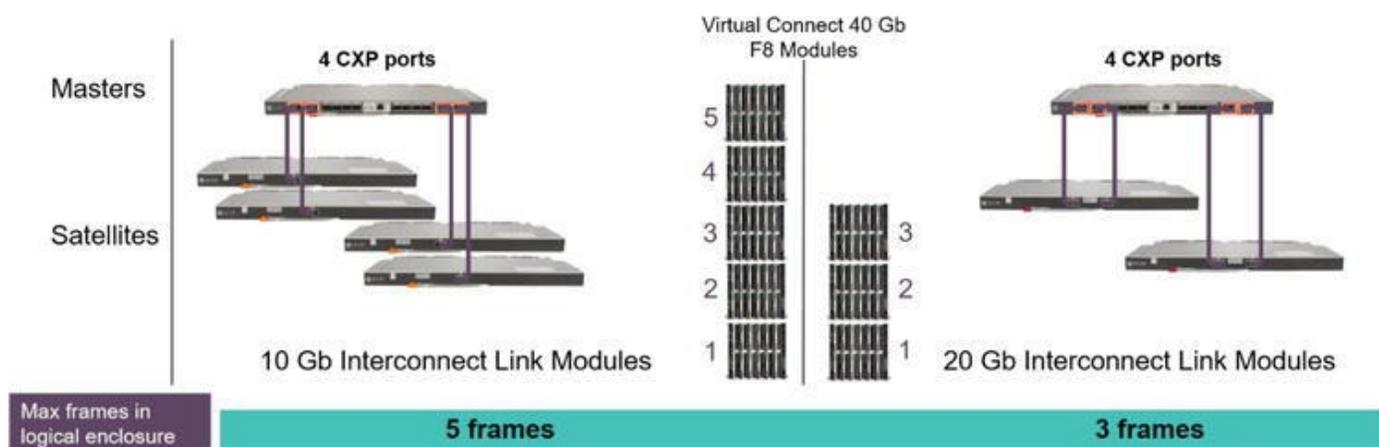


Figure 8-21 Composable fabric for high-density workloads

With Synergy composable fabric, scaling is fast, simple, and requires no downtime. When you add a new frame using a satellite interconnect link module (ILM), the new frame is automatically discovered as an extension of the existing fabric, and the east/west design (no-hop with satellite modules) scales so the performance of the existing workload is not negatively impacted.

The satellite module extends HPE composable fabric to additional frames (known as satellite frames) and replaces fixed ratio of interconnects in each frame by extending the fluid pool of network resources from the master module.

The built-in rack-scale fabric, shown in [Figure 8-21](#), allows a single fabric to span the entire rack all at once, eliminating the need for plugging in ToR switches. As satellite frames are added to a stack, they consume ports on master frames instead of ToR switches. When 10 Gb ILMs are used, a maximum of five frames can be connected. When 20 Gb ILMs are used, a maximum of three frames can be pooled.

Advantages of a master/satellite architecture include:

- 40% or more lower fabric hardware costs

10 Gb and 20 Gb bandwidth with a future path to 40 Gb/100 Gb to compute modules

Flexible bandwidth allocation in 100 MB increments

Ethernet, Fibre Channel (FC)/Fibre Channel over Ethernet (FCoE), internet Small Computer System Interface (iSCSI)

Zero-touch change management

Upgrade with minimum downtime

Synergy composable fabric's enterprise-grade availability allows IT to implement frictionless infrastructure changes with one-tool, one-step template-based operations from the Composer. Change operations such as modifying network connectivity or updating firmware are automatically implemented from the Composer GUI or unified API while the infrastructure is online, which significantly reduces manual interactions and errors, ensures real-time compliance, and minimizes downtime.

Synergy multi-fabric connectivity enables you to:

Reduce complexity, simplify hardware, and lower costs

Increase performance with flat, rack-scale fabric

Simplify operations with frictionless scaling

Improve availability with resilient edge fabrics

HPE Synergy 25/50 Gb Converged Network Adapter



Figure 8-22 HPE Synergy 6810C 25/50 Gb Converged Network Adapter

Synergy provides high-performance converged networking connections in mezzanine form factor. Synergy compute modules support two types of connector supporting mezzanine cards—type “C” or type “D.” Type C slots support only type C mezzanine cards, while type D mezzanine slots support both type C or type D cards.

The HPE Synergy 6810C 25/50 Gb Converged Network Adapter (CNA), shown in [Figure 8-22](#), supports high performance 25 or 50 Gbps connectivity for HPE Synergy compute modules, converged traffic for FCoE, iSCSI, or Ethernet.

Other advanced features of this adapter include:

RDMA over Converged Ethernet (RoCE v2, RoCE v1) used for exchange of information in high-performance cluster solutions (HPC)

Tunnel Offload technologies (NVGRE and VXLAN)

Single-root input/output virtualization (SR-IOV)

IPv6 Acceleration

Precision Time Protocol

Active Health Systems support

Authentication of digitally signed firmware

Additional security features, including:

- UEFI Secure Boot
- Sanitization
- Audit Logs
- Device-level Firewall

HPE Synergy 10/20 Gb Converged Network Adapters

Several CNAs are available in mezzanine form factor for customers who require 10/20 Gb networking connections. These include the HPE Synergy 3520C (10/20 Gb), the HPE Synergy 2820C (10 Gb), and the 3820C (10/20Gb).

HPE Synergy 3520C

The HPE Synergy 3520C 10/20 Gb Converged Network Adapter is designed for traditional IT, private/public cloud, and Telco customers and provides easy-to-manage, flexible, with one of the highest performing converged 10 or 20 Gb Ethernet connectivity within the HPE Synergy Composable Infrastructure. It is a key element in Hewlett Packard Enterprise composable fabric connecting pools of compute resources to networks. HPE Synergy 3520C also offers a rich set of offload technologies such as overlay network tunneling and storage offload to improve virtualization performance and compute host efficiency. With Flex-20 Technology, it converges Ethernet and FCoE onto a single connection, which simplifies IO hardware requirements and reduces costs.

Other key features include:

Converged traffic (FCoE or iSCSI and Ethernet)

DPDK support (PMD)

Tunnel Offload with VXLAN/NVGRE

7× FlexNICs + 1× FlexHBA/FlexNIC per port

Orchestration of reliable adapter firmware updates for an entire HPE Synergy infrastructure from a single tool, HPE Synergy Composer

HPE Synergy 2820C

The HPE Synergy 2820C 10Gb Converged Network Adapter is a key element in HPE composable fabric connecting pools of compute resources to networks with reliable, high-performing converged 10 Gb Ethernet connectivity. It has storage IO functionality along with single root IO virtualization (SR-IOV) capabilities. With Flex-10 Technology, it converges Ethernet and FCoE onto a single connection simplifying hardware and reducing costs.

Other key features include:

Converged traffic (FCoE or iSCSI and Ethernet)

Flex-10 Technology to simplify IO hardware and reduce costs

NPAR, SR-IOV, Receive Side Scaling (RSS), jumbo frames, and PXE boot

Orchestration of reliable adapter firmware updates for an entire HPE Synergy infrastructure from a single tool, HPE Synergy Composer

HPE Synergy 3820C

The HPE Synergy 3820C 10/20Gb Converged Network Adapter is a key element in HPE composable fabric connecting pools of compute resources to networks with reliable, high-performing converged 10 or 20 Gb Ethernet connectivity. With Flex-20 Technology, the Synergy 3820C converges Ethernet and FCoE onto a single connection simplifying hardware and reducing costs. Each adapter replaces a 10 Gb NIC and 8 Gb FC HBA simplifying IO hardware and reducing costs.

Other key features include:

Converged traffic (FCoE or iSCSI and Ethernet)

Flex-20 Technology to simplify IO hardware and reduce costs

NPAR, SR-IOV, RSS, jumbo frames, and PXE boot

Orchestration of reliable adapter firmware updates for an entire HPE Synergy infrastructure from a single tool, HPE Synergy Composer

Synergy fabric portfolio choices

	Composable (OneView managed)	Traditional (CLI managed)	CNAs/HBAs
Converged	 <ul style="list-style-type: none"> - HPE VC SE 40 Gb F8 Module - HPE Synergy 10/20 Gb Interconnect Link Module 	 <ul style="list-style-type: none"> - HPE 40 Gb F8 Switch Module - HPE Synergy 10/20 Gb Interconnect Link Module 	 <ul style="list-style-type: none"> - Requires CNA <ul style="list-style-type: none"> - Ethernet, FCoE, 8 Gb FC - License for FC
Native fibre channel	 <ul style="list-style-type: none"> - HPE Virtual Connect SE 16 Gb FC Module 	 <ul style="list-style-type: none"> - Brocade 16 Gb/24 FC Switch Module - Brocade 16 Gb/24 FC Switch Module Pwr Pk - Brocade 16 Gb/12 FC SAN Switch Module 	 <ul style="list-style-type: none"> - Requires HBA <ul style="list-style-type: none"> - 16 Gb FC - Optional Brocade licensing
Benefits	<ul style="list-style-type: none"> - Fully-composable management by OneView - Disaggregated fabric topology 	<ul style="list-style-type: none"> - Monitored by OneView - No disaggregated fabric 	

Figure 8-23 Synergy offers composable/non-composable and converged/native choices

Synergy provides a composable fabric for high-density workloads. Disaggregated rack-scale architecture reduces costs and simplifies networking, enabling you to:

Precisely match workload performance needs with low latency, multi-fabric, multi-speed architecture

Provide continuous service delivery with high-availability resilient fabric

There are two sets of decision points when choosing a fabric module for a Synergy configuration:

Composable (managed by HPE OneView) or traditional (non-composable—managed through a command line interface [CLI])

Converged or native

Synergy supports redundant fabrics, which can be composable or traditionally managed with the interconnects shown in [Figure 8-23](#). The options for the fabric modules are:

HPE Virtual Connect SE 40Gb F8 module for HPE Synergy

HPE Virtual Connect SE 16Gb FC Module for HPE Synergy

HPE Synergy 10Gb Interconnect Link Module

HPE Synergy 20Gb Interconnect Link Module

HPE Synergy 40Gb F8 Switch Module

Brocade 16Gb FC SAN Switch Module for HPE Synergy

HPE Synergy 10Gb/40Gb Pass-Thru Modules (not shown in [Figure 8-23](#))

Mellanox SH2200 Switch Module for HPE Synergy (not shown in [Figure 8-23](#))

Synergy fabric advantages and benefits

Table 8-3 Positioning Synergy interconnects

	Virtual Connect	Ethernet/SAN switches	Pass-thru modules
When to use	Server admin supports a Virtual Connect solution Enterprise infrastructure management desired BladeSystem customer Cisco networking both LAN and SAN	Networking team influences networking choice at the edge Rich networking features for instance L2/L3 for LAN and fabric servers for SAN at edge Customer has HPN or Brocade networking gear	Server admin wants 10 Gb connectivity for a single frame Clear demarcation between server and network admin Networking team controls all networking aspects
Advantages	Software-defined infrastructure template-based management Wire-once: Easiest server moves, adds, changes Simplest, most flexible infrastructure to manage Flat SAN with 3PAR storage Minimal impact on existing SAN/LAN infrastructure	CLI based interface can be used with home-grown scripting tools A high degree of flexibility	Simple to install and manage Cost-effective with single frame

[Table 8-3](#) presents features and advantages of VC, Ethernet/SAN switches, and pass-thru modules.

Mellanox SH2200 Switch Module for HPE Synergy

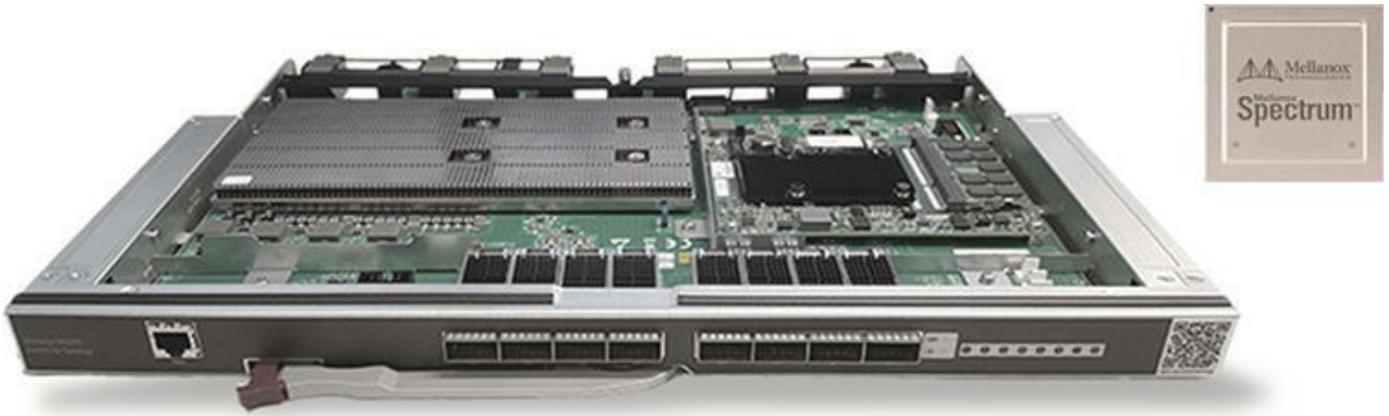


Figure 8-24 Mellanox SH2200 Switch Module for HPE Synergy

The Mellanox SH2200 high-speed Ethernet switch for HPE Synergy, shown in [Figure 8-24](#), delivers high performance, deterministic low latency network fabric targeting HPC, FSI, cloud, and Telco/NFV environments. Composable Infrastructure is not supported for this Ethernet switch.

The Mellanox SH2200 Switch Module for HPE Synergy delivers high-performance, high-speed, low latency 25/50 GbE connectivity to each of the Synergy compute nodes, and 40/100 GbE to upstream network switches. The module delivers up to 2.8 Tbps switching capacity and 8.4 Billion Packet per Second (BPPS) throughput for the most demanding data center applications. It offers advanced and comprehensive Layer 2 and Layer 3 feature sets, as well as hardware-based network virtualization support. Network reliability is improved with resilient fabric connectivity to upstream switches using Multi-module LAG (MLAG) technology to reduce single points of failure.

Synergy storage

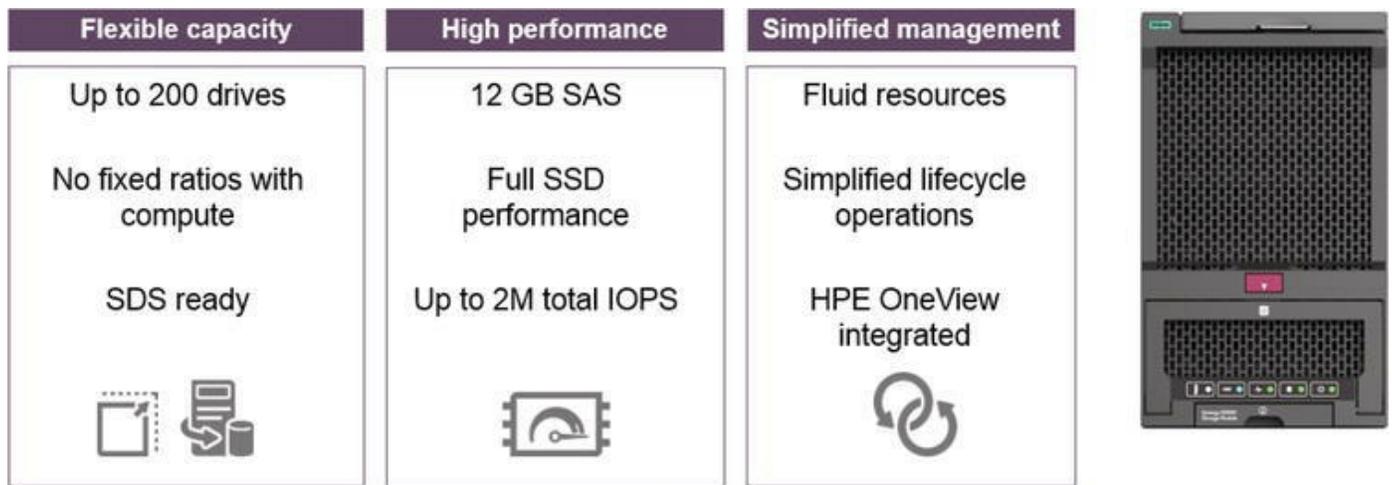


Figure 8-25 Fully integrated, performance capable, composable storage with Synergy

Synergy D3940 composable storage features complete integration into the frame, driving simplicity and reliability. Each storage module has a maximum 40 drives, and Synergy Composer can zone up to 200 storage drives to any compute module, as shown in [Figure 8-25](#). For DAS-centric workloads, the non-blocking SAS fabric allows full utilization of flash storage performance for local applications.

A Synergy frame combines compute, storage, and fabric in a single physical intelligent infrastructure. A combination of storage options—including fully integrated internal storage modules, directly connected and fully orchestrated SAN storage, and software-defined storage—two-socket and four-socket compute modules, and a variety of redundant fabric modules allows it to meet a wide range of workload requirements, so customers can easily deploy the entire infrastructure needed to run an application and store its data.

There are no fixed ratios—storage resources are pooled for composition with compute resources. Synergy storage can be configured with HPE StoreVirtual virtual storage appliance (VSA) to offer flexible capacity for a wide range of uses and workloads:

Sized for performance

- 4× 10 TB VSA license
- 4× two-socket compute nodes with P416e-m SR Gen10 controllers, no local drive
- 1× D3940 with (40) 2 TB drives
- Dual IO adapters, Dual 12 G SAS connection modules

Scaled with growth

- 8× 10 TB VSA license
- 8× two-socket compute nodes with P416e-m SR Gen10 controllers, no local drive
- 2× D3940, 40× 2 TB drives in each, dual IO adapters, and dual 12 G SAS connection modules

Scaled for capacity

- 4× 50 TB VSA license
- 4× two-socket compute node with P416e-m SR Gen10 controller, two local drives
- 4× D3940, 40× 2 TB drives in each
- Dual IO adapters in each D3940, and dual 12 G SAS connection modules

HPE Synergy D3940 Storage Module

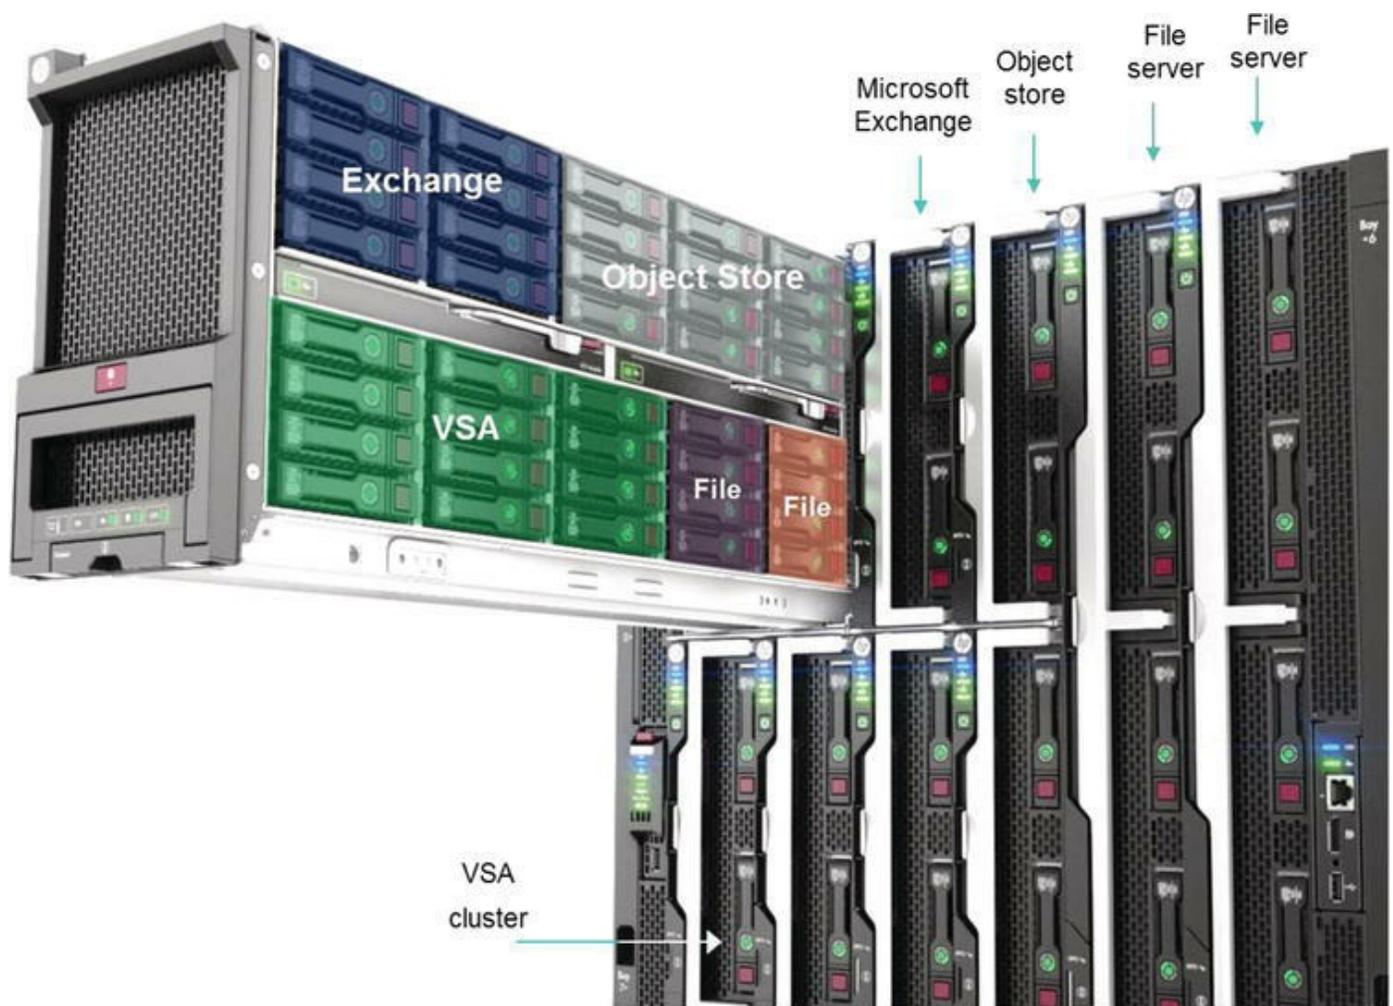


Figure 8-26 HPE Synergy D3940 Storage Module

The HPE Synergy D3940 Storage Module shown in [Figure 8-26](#) is the foundational element of Synergy storage. Within a frame, an administrator can use the Synergy D3940 Storage Module to create logical drives for any compute module in the enclosure, remap connections, or change personalities—all done with stored profiles. Storage module resources can be shared to multiple compute modules across Synergy frames using software such as HPE StoreVirtual VSA.

The Synergy D3940 Storage Module with one or two Synergy 12 Gb SAS connection modules allows you to establish an efficient pool of storage within a Composable Infrastructure which can be easily allocated to meet the individual needs of server workloads and business demands.

Up to five Synergy D3940 Storage Modules can be configured per Synergy 12000 frame, delivering 200 drives in total. Connection to compute modules is delivered through a non-blocking SAS fabric of single or dual connection modules in the frame and single or dual IO modules in each storage module.

The storage module provides an industry-leading density of 24 drives per rack U of space with 40 small form factors drives per module. Each module supports redundant IO adapters for failover and as a composed resource under HPE OneView 3.0 and later features non-disruptive updates and an “any-to-any” composability methodology, breaking historical drive to compute ratios and offering a truly flexible pool of resources. Performance features include:

4× controller bandwidth (96 Gb) compared with HPE BladeSystem c-Class (24 Gb)

2× drive enclosure bandwidth (192 Gb) compared with the HPE D3700 enclosure (96 Gb)

HPE Dynamic Bandwidth Optimizer for 6G SATA

Synergy D3940 Storage Module as DAS = 2M IOPs

Innovation from end to end



Figure 8-27 HPE Synergy storage represents the best of HPE innovation from end to end

Synergy D3940 storage is an end-to-end non-blocking 12 Gb SAS fabric that delivers on the promise of full solid-state drive (SSD) performance with up to 2 million total input/output operations per second (IOPS). As shown in [Figure 8-27](#), the fabric includes:

D3940 Storage Module

12 Gb SAS connection modules

Smart Array controller(s)

SSD or NVMe drives

The 12 Gb SAS connection module, residing in ICM bays 1 and 4, creates a non-blocking fabric for storage traffic routed from storage controllers inside the compute modules. Traffic is passed through 12 SAS ports; each has four 12 Gb/s channels for an aggregated total of 48 Gb/s per port.

HPE Smart Array P416ie-m SR Gen10 controller



Figure 8-28 HPE Smart Array P416ie-m SR Gen10 controller

The HPE Smart Array P416ie-m SR Gen10 controller, shown in [Figure 8-28](#), supports 12 Gb/s SAS and PCIe 3.0 and provides enterprise-class storage performance, reliability, security, and efficiency needed to address your evolving data storage needs while supporting advanced RAID levels with 2 GB flash backed write cache (FBWC). This controller operates in Mixed Mode which combines RAID and HBA operations simultaneously and offers enhanced security by encrypting data-at-rest on any drive with HPE Smart Array SR Secure Encryption. It has eight external SAS lanes, allowing

connection to SAS or SATA drives in the HPE Synergy D3940, and eight internal SAS lanes, allowing connection to drives inside the HPE Synergy compute modules.

Remote support in Synergy via embedded HPE OneView



Figure 8-29 HPE OneView provides integrated remote support

HPE OneView provides integrated remote support that is part of the Composer appliance using the agentless remote support that is part of iLO, as shown in [Figure 8-29](#).

Key features of remote support in Synergy are:

Appliance registration for remote support—Customer opt-in, contact information, and support partner details

Service incident forwarding and tracking in HPE OneView—Case status updates

Insight Online Access (optional)—With customer opt-in, all devices are visible in the Insight Online portal

Solution aware support ticket routing—With Converged Systems solution deployments, support tickets are routed to Converged Systems support

Collections, including Active Health System data—The Active Health System log is included in the automated event transmittal

Learning check

Match the Synergy component with its function.

Composer	Eliminates the need for stand-alone ToR switches
Frame	Embeds the HPE OneView management solution to manage compute modules, fabrics, and storage
Compute module	Provides the performance, scalability, density optimization, storage simplicity, and configuration flexibility
Fabric	Reduces complexity through an intelligent autodiscovery infrastructure

Write a summary of the key concepts presented in this chapter.

Summary

Synergy is a single infrastructure of physical and virtual pools of compute, storage, and fabric resources and a single management interface that allows IT to instantly assemble and reassemble resources in any configuration.

Synergy eliminates hardware and operational complexity so IT can deliver infrastructure to applications faster and with greater precision and flexibility.

The key components of Synergy are:

- Composer
- Compute
- Frame
- Storage
- Fabric

The Synergy management subsystem comprises Synergy Composer, Synergy Image Streamer, and the FLMs.

Synergy Composer uses software-defined intelligence with embedded HPE OneView to aggregate compute, storage, and fabric resources.

9 HPE On-system Management

WHAT IS IN THIS CHAPTER FOR YOU?

After completing this chapter, you should be able to:

Explain how to prepare an HPE ProLiant server system for deployment.

Name the tools used during the ProLiant boot process and server configuration.

- ✓ System Utilities
- ✓ BIOS/Platform Configuration

Name the HPE on-system management tools used to prepare a ProLiant system for deployment.

- ✓ iLO 5
 - ✓ Intelligent Provisioning
 - ✓ HPE Smart Update Manager (HPE SUM)/Service Pack for ProLiant (SPP)
 - ✓ RESTful Interface Tool
 - ✓ Agentless Management
 - ✓ HPE Smart Storage Administrator (SSA)
-

PRELEARNING CHECK

You are designing a solution for a client that will involve installing servers in a rack. The client needs local access to iLO but due to power availability problems in the data center the support staff cannot use KVM.

Which Gen10 feature would be useful to satisfy this need?

A customer approaches your stand at a trade fair and informs you that they are looking to upgrade their G7 servers to Gen10 servers. They ask if they still need to use the Smart Start CD to deploy Windows?

How should you respond?

Preparing a ProLiant system for deployment

To prepare a ProLiant system for a deployment, you need to:

Install hardware.

Configure iLO.

Launch Unified Extensible Firmware Interface (UEFI) and perform server configuration.

Boot Intelligent Provisioning.

Perform firmware update.

Validate environmental standards.

Configure Direct Access Storage.

Consider storage networking.

Test the installation (using UEFI embedded diagnostics).

Installing hardware

The decision to select a new ProLiant server often can depend on the ability to configure or upgrade that server based on workload or performance requirements.

Hardware components that can be configured or upgraded

Configuring a system by adding or upgrading components can deliver optimal performance for a given unique set of requirements. ProLiant servers contain several components that can be configured or upgrade to meet business and technical requirements.

IO cards

IO cards include host bus adapters (HBAs), which are interface cards that connect host or storage devices to a storage area network (SAN). They are frequently used for Fibre Channel and internet Small Computer System Interface (iSCSI)-based traffic. Network interface cards (NICs) are cards that connect a computer to a network. They are used for IP-based LAN traffic.

A FlexHBA is a PCIe physical function on the FlexFabric adapter that you can configure to handle storage traffic. The server ROM, operating system (OS), and hypervisor

recognize the PCIe function as an HBA device. You can assign storage traffic (Fibre Channel or SCSI) as a FlexHBA only to the second physical function of each FlexFabric adapter port because this is the port used for storage access.

Installing DIMMs

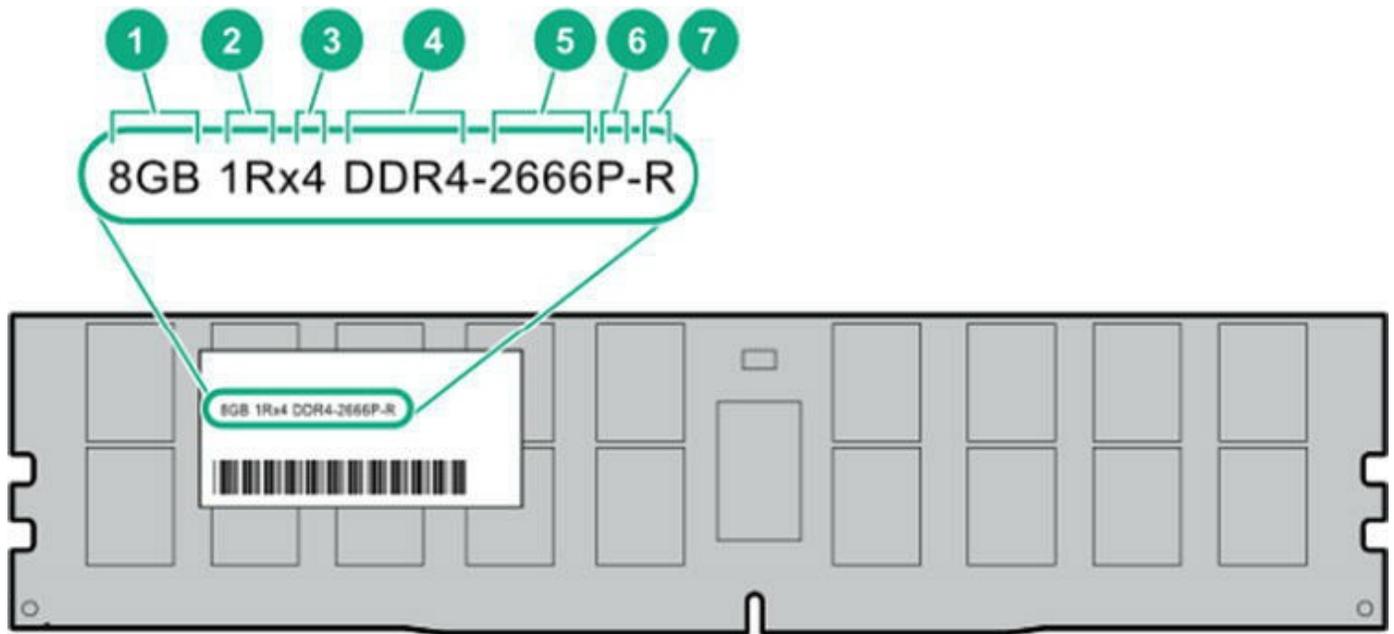


Figure 9-1 Decoding dual in-line memory module configuration

Item	Description	Definition
1	Capacity	8 GB 16 GB 32 GB 64 GB 128 GB
2	Rank	1R = Single rank 2R = Dual rank 4R = Quad rank 8R = Octal rank
3	Data width on DRAM	x4 = 4-bit x8 = 8-bit x16 = 16-bit
4	Memory generation	PC4 = DDR4
5	Maximum memory speed	2133 MT/s 2400 MT/s 2666 MT/s
6	CAS latency	P = CAS 15-15-15 T = CAS 17-17-17 U = CAS 20-18-18 V = CAS 19-19-19 (for RDIMM, LRDIMM) V = CAS 22-19-19 (for 3DS TSVLRDIMM)
7	DIMM type	R = RDIMM (registered) L = LRDIMM (load reduced) E = Unbuffered ECC (UDIMM)

If you are adding extra memory to the server, the placement is important. For example, if the server has four slots and you have two memory modules (dual in-line memory modules or DIMMs), the modules must be inserted into specific slots in a specific order or the server will not boot properly. Refer to the diagram on the hardware or the server's user guide for specific memory configurations. Be sure not to mix DIMM types (UDIMMs, RDIMMs, and LRDIMMs) if the server does not support this.

To configure memory properly, an understanding of single-, dual-, and quad-rank DIMMs is helpful. Some DIMM configuration requirements are based on these classifications. [Figure 9-1](#) shows how to identify the specifications of a DIMM.

A single-rank DIMM has one set of memory chips that is accessed while writing to or

reading from the memory. A dual-rank DIMM is similar to having two single-rank DIMMs on the same module, with only one rank accessible at a time. A quad-rank DIMM is, effectively, two dual-rank DIMMs on the same module. Only one rank is accessible at a time. The server memory control subsystem selects the proper rank within the DIMM when writing to or reading from the DIMM.

Dual- and quad-rank DIMMs provide the greatest capacity with the existing memory technology. For example, if current dynamic random-access memory (DRAM) technology supports 8 GB single-rank DIMMs, a dual-rank DIMM would be 16 GB, and a quad-rank DIMM would be 32 GB.

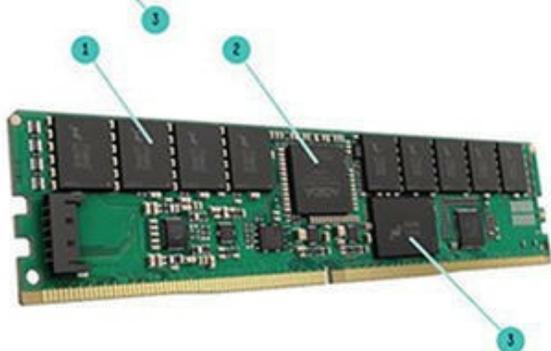
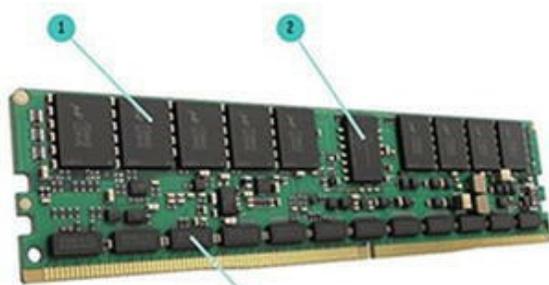
LRDIMMs are labeled as quad-rank DIMMs. There are four ranks of DRAM on the DIMM, but the LRDIMM buffer creates an abstraction that allows the DIMM to appear as a dual-rank DIMM to the system. The LRDIMM buffer isolates the electrical loading of the DRAM from the system to allow for faster operation. This allows faster memory operating speed compared to quad-rank RDIMMs.



Note

UDIMM is unregistered memory; RDIMM is registered DIMM; and LRDIMM is load-reduced.

NVDIMM



Item	Front view	Rear view
1	DRAM	DRAM
2	Register: Command/Address control for operations	FPGA: Controller for NVDIMM
3	MUX: Bus isolation for backup and restore	NAND Flash: Persistent store for the NVDIMM

HPE Smart Storage Battery provides a backup power source in the event of a surprise power loss allowing any data in flight on the DRAM to be moved to the non-volatile NAND Flash.

Population rules for NVDIMM may differ from model to model.

Figure 9-2 Non-volatile dual in-line memory module front and back

The non-volatile dual in-line memory module (NVDIMM) shown in [Figure 9-2](#) combines DRAM and NAND flash memory. The NAND flash memory retains a backup of the volatile data from the DRAM in the event that electrical power is removed, for example, from an unexpected power loss, system crash, or normal shutdown. The Smart Storage battery powers the NVDIMM until the content of the volatile memory is copied to the flash memory.

Gen10 DIMM installation guidelines

HPE Gen10 DL360 / DL380 / DL560* Servers

2 Slots per Channel

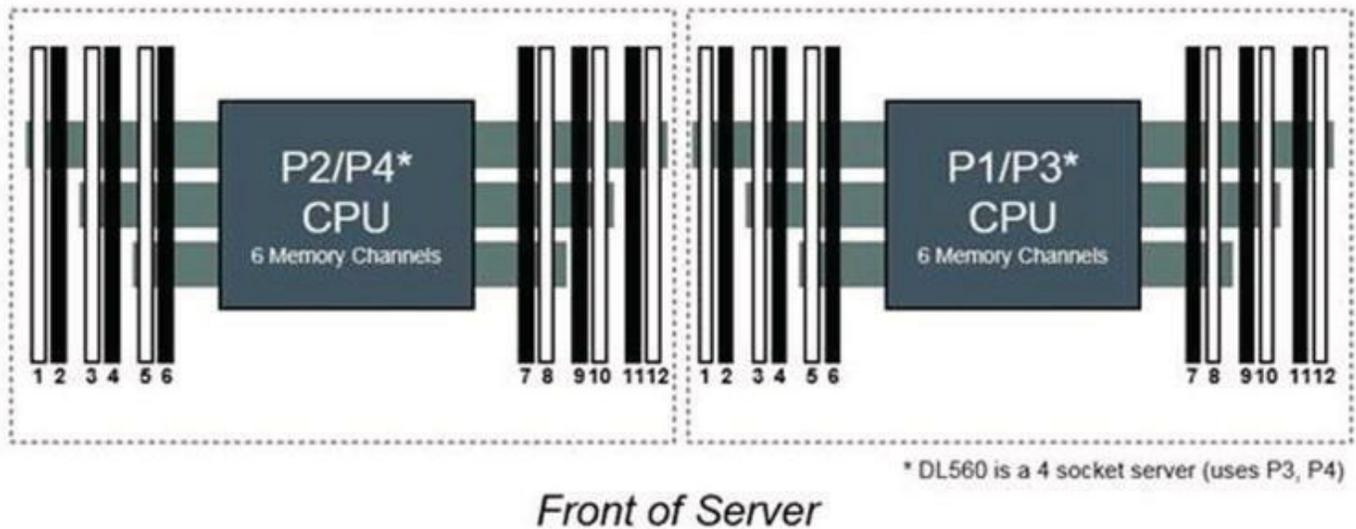


Figure 9-3 ProLiant Gen10 DIMM channels and slots

[Figure 9-3](#) shows the memory slot configuration for HPE ProLiant DL360/DL380/DL560 Gen10 servers. In this figure, the first memory slots for each channel on each processor are the white ones. Note that the HPE ProLiant DL560 is a four-socket server and uses processors 1, 2, 3, and 4.

When installing memory in a server, you should always follow the DIMM slot population guidelines for the server you are working on. For a complete list of guidelines, refer to the user guide that shipped with the server or that is available online.

Sample key guidelines are as follows:

Install DIMMs only if the corresponding processor is installed.

Insert DIMMs into specific slots in a specific order for the server to boot properly.

NVDIMMs occupy slots closest to the CPU.

When two processors are installed, balance the DIMMs across the two processors.

White DIMM slots denote the first slot of a channel (Ch-0 DIMM-0, Ch-1 DIMM-0, Ch-2 DIMM-0) etc.

Do not mix RDIMMs and LRDIMMs.



Note

The guidelines shown are for the ProLiant DL380 Gen10 Server.



Warning

Care should be taken when handling SmartMemory DIMMs and NVDIMMs. Do not bend or twist the DIMM. Insert the DIMM perpendicular to the socket, aligned to the location notch, and apply pressure to the top edge until an audible click is detected.

Gen10 DIMM installation guidelines for a single processor

HPE ProLiant Gen10 12 slot per CPU DIMM Population Order												
DIMMS	Left six DIMM slots						Right six DIMM slots					
1 DIMM								8				
2 DIMMs								8		10		
3 DIMMs								8		10		12
4 DIMMs			3		5			8		10		
5 DIMMs			3		5			8		10		12
6 DIMMs	1		3		5			8		10		12
7 DIMMs	1		3		5		7	8		10		12
8 DIMMs			3	4	5	6	7	8	9	10		
9 DIMMs	1		3		5		7	8	9	10	11	12
10 DIMMs	1		3	4	5	6	7	8	9	10		12
11 DIMMs	1		3	4	5	6	7	8	9	10	11	12
12 DIMMs	1	2	3	4	5	6	7	8	9	10	11	12

Figure 9-4 Gen10 DIMM installation guidelines for a single processor

Figure 9-4 shows the Gen10 DIMM installation guidelines for a single processor.

The Intel Xeon Skylake Scalable Processor supports six memory channels, each channel consists of two DIMM sockets. The numbers in the white and black boxes represent the DIMM position relative to DIMMs one to twelve, left to right. For example: for a single DIMM, populate Channel 0, DIMM 0, identified by the number 8 (DIMM slot number 8) in the top row of the table above.



Note

DIMM counts of 5, 7, 9, 10, and 11 results in an unbalanced configuration and is not recommended.

Configuring or upgrading other components

Other hardware components that can be configured or upgraded include:

Hard drives—These data storage devices are installed internally within the server. Some are hot-pluggable, and others are non-removable. Configure the Smart Array Controller to create the boot disk and data disks.

Fans—ProLiant ML and DL servers typically have fans pre-installed. If you have only one processor installed in a two-socket server, fan blanks must be installed. You must install fans into a BladeSystem enclosure. Always be sure to install the right fan in the right place. Refer to the installation guide for more details.

Cabling—Refer to the server’s documentation for cabling guidelines and considerations so that you can make informed decisions about cabling the server and hardware options to optimize performance.

Configuring the Smart Array Controller

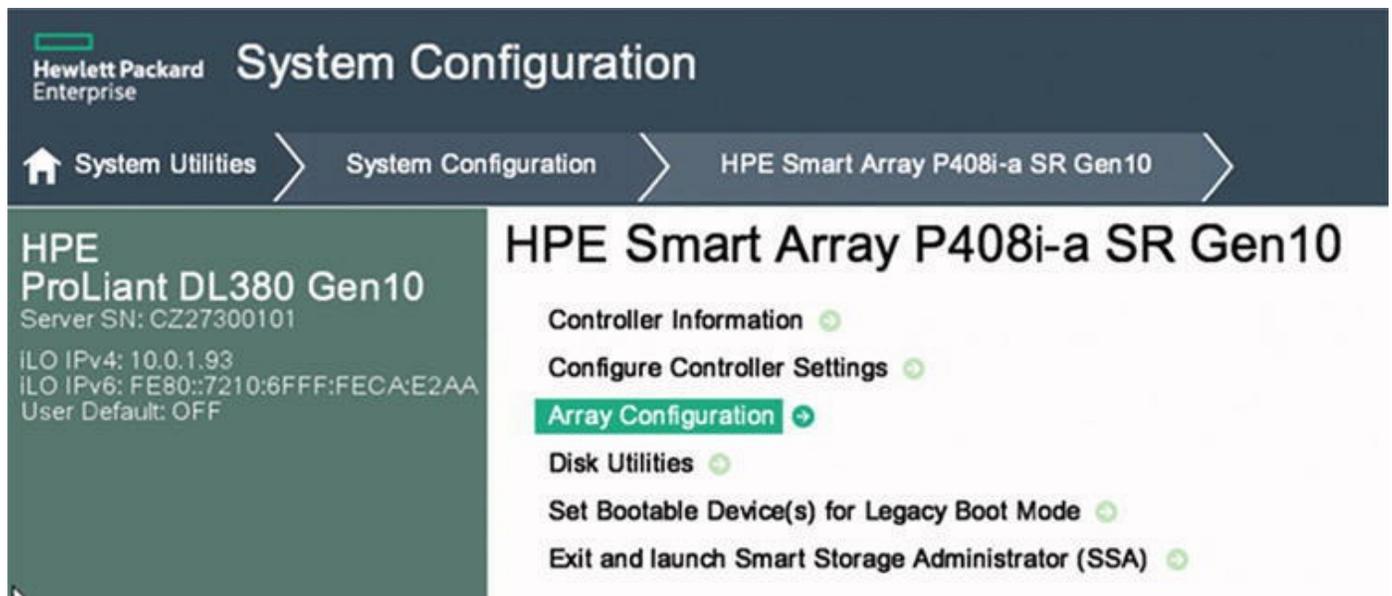


Figure 9-5 UEFI based Smart Array Controller configuration options

In the ProLiant Gen10 solution, there are several different options for configuring the Smart Array Controller. The most comprehensive tool for array configuration and general lifecycle management is the SSA. The SSA can be accessed from Intelligent Provisioning by selecting the maintenance option or can be launched from the System Utilities menu in the UEFI console.

As shown in [Figure 9-5](#), UEFI provides quick Smart Array Configuration by selecting **System Utilities** → **System Configuration** → **Embedded RAID controller** → **Array Configuration**.

Other options allow the configuration to be cleared, determine the state of the backup power source, modify controller settings, and select the default boot device. The Controller Information menu item provides a detailed view of the controller status.

Environmental standards

When installing server equipment, select a location that meets the environmental standards described in the user guide for that server. These standards cover the following basic categories:

Space—To allow for servicing and adequate airflow, observe the following space and airflow requirements when deciding where to install a rack:

- Leave a minimum clearance of 63.5 cm (25 in) in front of the rack.
- Leave a minimum clearance of 76.2 cm (30 in) behind the rack.
- Leave a minimum clearance of 121.9 cm (48 in) from the back of the rack to the back of another rack or row of racks.

Airflow—Most HPE servers draw in cool air through the front door and expel warm air through the rear door. Therefore, the front and rear rack doors must be adequately ventilated to allow ambient room air to enter the cabinet, and the rear door must be adequately ventilated to allow the warm air to escape from the cabinet. When vertical space in the rack is not filled by a server or rack component, the gaps between the components cause changes in airflow through the rack and across the servers. Cover all gaps with blanking panels to maintain proper airflow.

Temperature—The maximum recommended ambient operating temperature for many servers is 35°C (95°F). The temperature in the room where the rack is located must not exceed 35°C (95°F). Certain configurations of some servers are compliant with ASHRAE A3 or A4, which means that they can operate in ambient temperatures of 40°C (104°F) or 45°C (113°F), respectively. See the QuickSpecs or user guide of the server

for more information.

Power—Installation of the equipment must comply with local and regional electrical regulations governing the installation of information technology equipment by licensed electricians. When installing more than one server, you might need to use additional power distribution devices to provide power to all devices safely. Observe the following guidelines:

- Balance the server power load between available AC supply branch circuits.
- Do not allow the system AC current load to exceed 80% of the branch circuit AC current rating.
- Do not use common power outlet strips for this equipment.
- Provide a separate electrical circuit for the server.

Electrical grounding—The server must be grounded properly for proper operation and safety. Requirements vary according to regional codes. Because of the high ground-leakage currents associated with multiple servers connected to the same power source, HPE recommends the use of a PDU that is either permanently wired to the building's branch circuit or includes a non-detachable cord that is wired to an industrial-style plug. National Electrical Manufacturers Association (NEMA) locking-style plugs or those complying with IEC 60309 are considered suitable for this purpose. Using common power outlet strips for the server is not recommended.

Rack requirements—To reduce the risk of personal injury or damage to the equipment, be sure that:

- The leveling jacks are extended to the floor.
- The full weight of the rack rests on the leveling jacks.
- The stabilizing feet are attached to the rack if it is a single-rack installation.
- The racks are coupled together in multiple-rack installations.
- Only one component is extended at a time. A rack might become unstable if more than one component is extended for any reason.



Caution

Always plan the rack installation so that the heaviest item is on the bottom of the rack. Install the heaviest item first and populate the rack from the bottom to the top.

Learning check

What must you do if you have only one processor installed in a two-socket server?

- A. Install fan blanks
- B. Install a blanking panel in the second socket
- C. Balance the DIMMs across both processor sockets
- D. Designate one socket as active and the blank as passive

At server management

Ease of management is a prominent feature of HPE ProLiant servers and Gen10 enhancements include a new iLO Service Port.

iLO Service Port

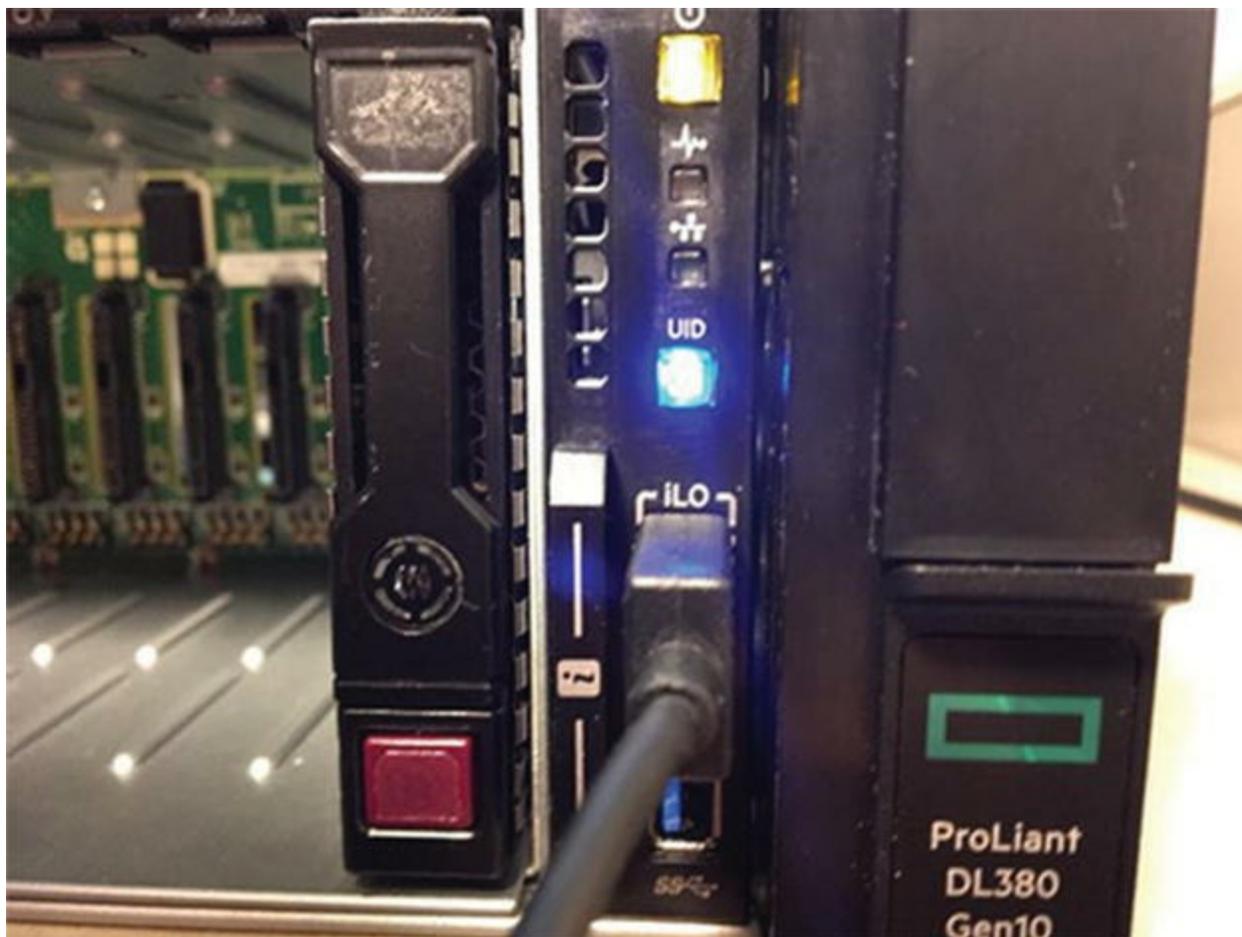


Figure 9-6 iLO Service Port

As shown in [Figure 9-6](#), the iLO Service Port is a USB port with the label iLO on the front of ProLiant Gen10 servers and Synergy Gen10 compute modules.

When you have physical access to a server, you can use the Service Port to do the following:

Download the Active Health System Log to a supported USB flash drive. When you use this feature, the connected USB flash drive is not accessible by the host OS.

Connect a client (such as a laptop) with a supported USB to Ethernet adapter to access the iLO web interface, remote console, CLI, iLO RESTful API, or scripts.

When you use the iLO Service Port:

Actions are logged in the iLO Event Log.

The server UID blinks to indicate the Service Port status. You can also retrieve the status by using a REST client and the iLO RESTful API.

You cannot use the Service Port to boot any device within the server, or the server itself.

You cannot access the server by connecting to the Service Port.

You cannot access the connected device from the server.

Use the following procedure to connect to iLO through the iLO Service Port:

Use a supported USB to Ethernet adapter to connect a client to the Service Port (the USB port labeled iLO, on the front of the server). The client NIC is assigned a link-local address. This process might take several seconds.

Connect to iLO through a browser, the CLI, or a scripting utility by using the IPv4 address 169.254.1.2.

Securing the iLO Service Port

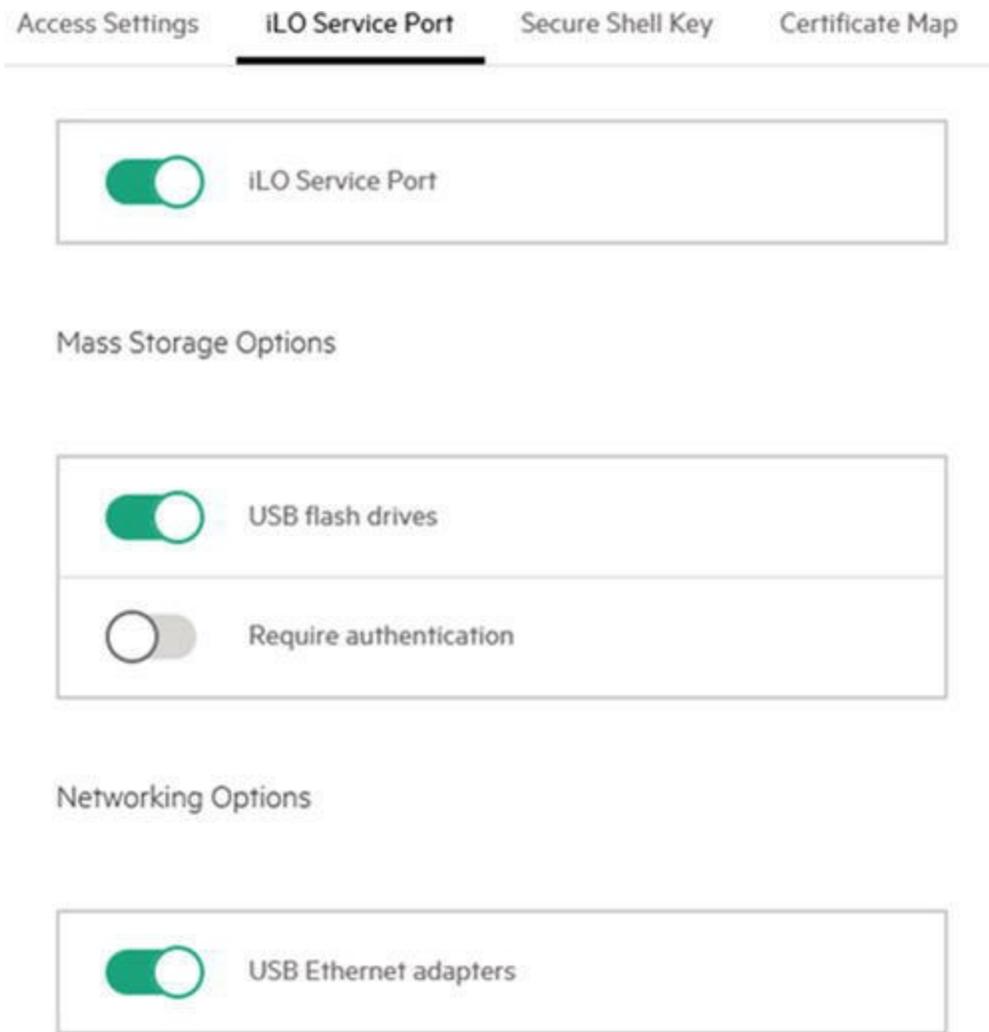


Figure 9-7 iLO Service Port options

iLO 5 allows you to enable or disable the iLO Service Port, as shown in [Figure 9-7](#). The default setting is enabled. When this feature is disabled, you cannot configure the features in the Mass Storage Options or Networking Options sections on this page. Do not disable the iLO Service Port when it is in use. If you disable the port when data is being copied, the data might be corrupted.

Learning check

You are designing a solution for a client that will involve installing servers in a rack. The client needs local access to iLO but due to power availability problems in the data center the support staff cannot use KVM. Which Gen10 feature would be useful to satisfy this need?

HPE iLO 5 management technologies

HPE iLO 5 provides the core foundation and intelligence for all ProLiant Gen10 servers. iLO 5 is ready to run and does not require additional software installation. HPE iLO can be accessed from any location using a web browser. iLO 5 management technologies simplify server setup, enable health monitoring, provide power and thermal control, and promote remote administration. iLO 5 management technologies support the complete lifecycle of ProLiant servers, from initial deployment to ongoing management and service alerting.

Setting up and connecting to iLO

Before setting up an iLO management processor, you must decide how to handle networking and security. The following questions can help you configure iLO:

How will iLO connect to the network?

- Typically, iLO is connected to the network through a dedicated management network or a shared connection on the production network.

Will NIC Teaming be used with the Shared Network Port configuration?

- NIC Teaming is a feature you can use to improve server NIC performance and reliability.

How will iLO acquire an IP address?

- To enable iLO access after it is connected to the network, the iLO management processor must acquire an IP address and subnet mask. You can use a dynamic address or a static address.

What access security is required, and what user accounts and privileges are needed?

- Methods available for managing access to iLO include local accounts, directory services, and CAC smartcard authentication.

What tools will you use to configure iLO?

- iLO supports various interfaces for configuration and operation. These include the iLO web interface, ROM-based setup, Intelligent Provisioning, iLO RESTful API, HPE OneView, HPE Scripting Toolkit, and several other scripting and CLI tools.



Note

For information about additional configuration options, see the “Setting up iLO” section of the HPE iLO 5 User Guide, which is available by entering the following URL into your browser:

https://support.hpe.com/hpsc/doc/public/display?docId=a00026409en_us

Using the iLO web interface

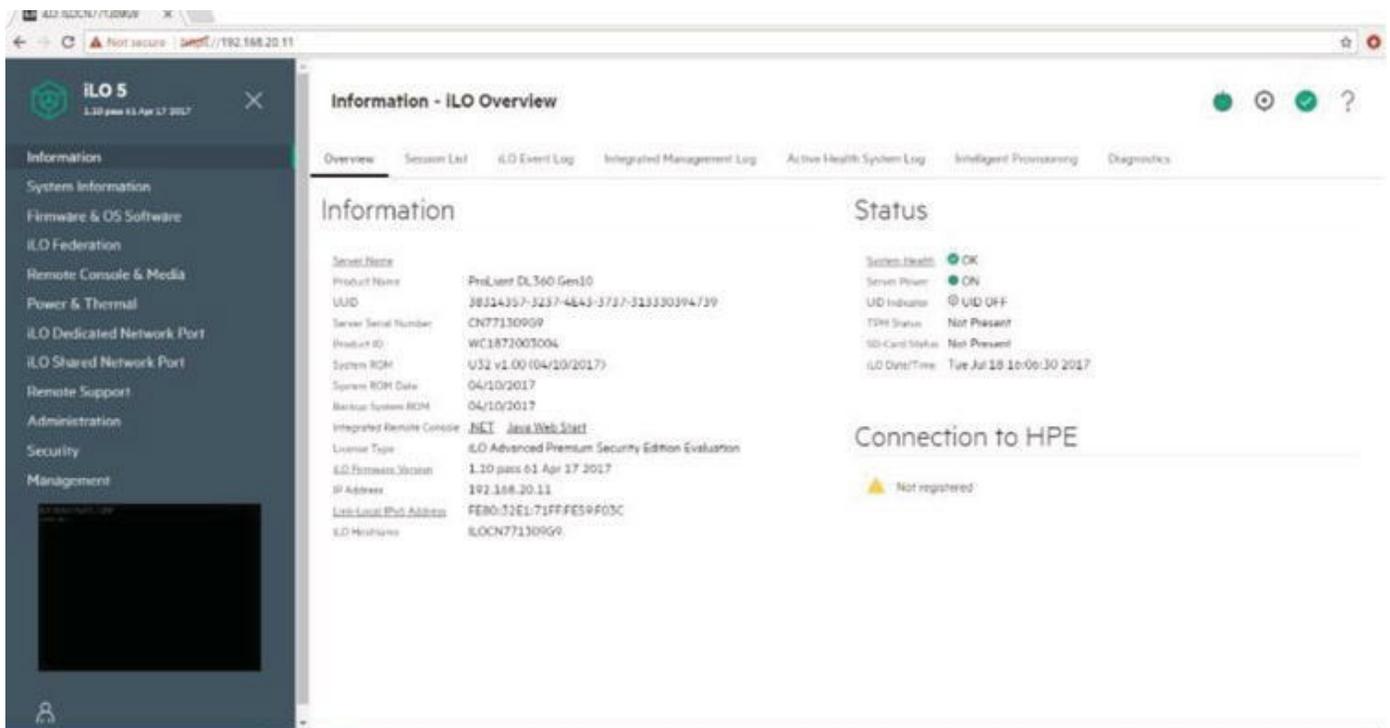


Figure 9-8 iLO 5 Overview screen

The iLO web interface allows you to configure and manage iLO. You can access iLO from a remote network client by using a supported browser and providing the default

DNS name or IP address.

As shown in [Figure 9-8](#), navigational links include:

Information

System Information

Firmware and Software

iLO Federation

Remote Console & Media

Virtual Media

Power & Thermal

iLO Dedicated Network Port

iLO Shared Network Port

Remote Support

Administration

Security

Management

The iLO web interface groups similar tasks for easy navigation and workflow. The interface is organized with a navigation tree in the left pane. To use the web interface, click an item in the navigation tree, and then click the name of the tab you want to view.

iLO Federation

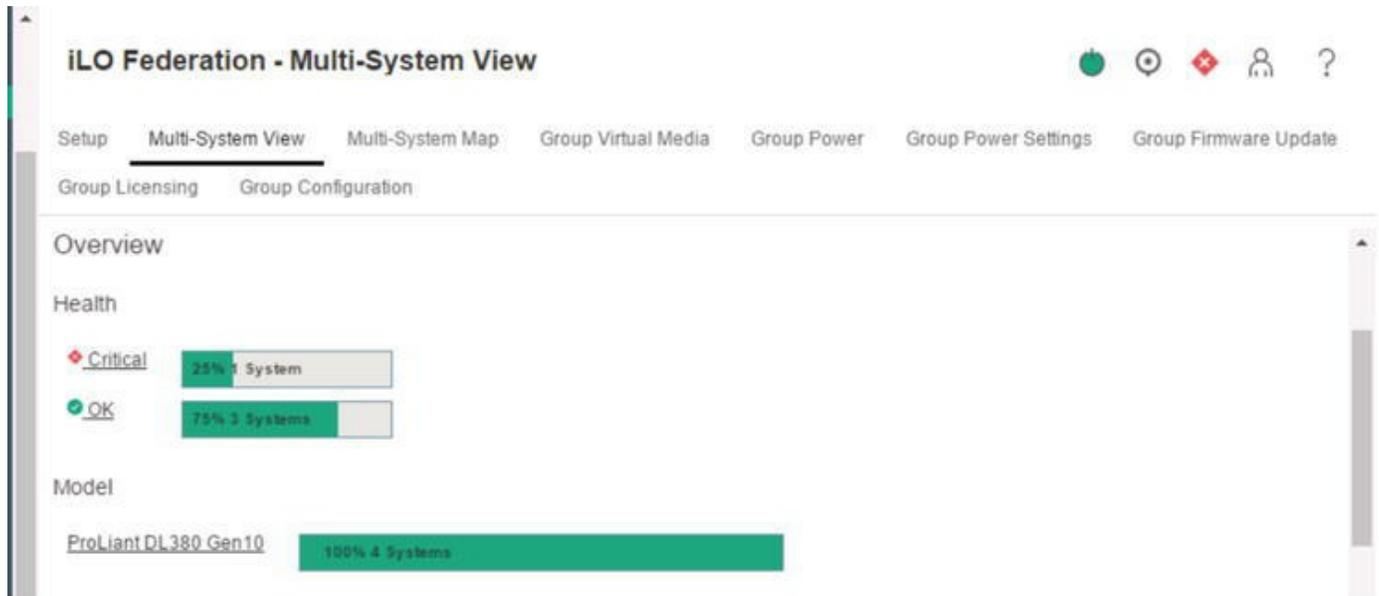


Figure 9-9 Viewing group iLO Federation Multi-System View

iLO Federation Multi-System View is shown in [Figure 9-9](#).

iLO Federation:

Enables you to manage multiple servers at one time

Offers built-in rapid discovery of all iLOs, group configurations, and group health status

Enables you to determine which servers have iLO license

Allows you to unify the system management of thousands of servers, regardless of location, from one system running the iLO web UI

Supports all ProLiant Gen10, Gen9, and Gen8 servers

To remotely manage groups of servers at scale, iLO Federation offers built-in rapid discovery of all iLOs, group configurations, group health status, and ability to determine which servers have iLO licenses. With an HPE iLO Advanced license, you can enable the full implementation of iLO Federation management for features such as Group Firmware Updates, Group Virtual Media, Group Power Control, Group Power Capping, and Group License Activation. iLO Federation lets you unify the system management of thousands of servers, regardless of location, from one system running the iLO web UI (iLO 4 1.40 and later) for any ProLiant server in an iLO Federation management group.

ProLiant boot process and server configuration

The boot process is the initial set of operations a server performs after electrical power to the server is switched on. Configuration begins immediately following the boot process, setting up relationships between functional units of the system.

Server boot process

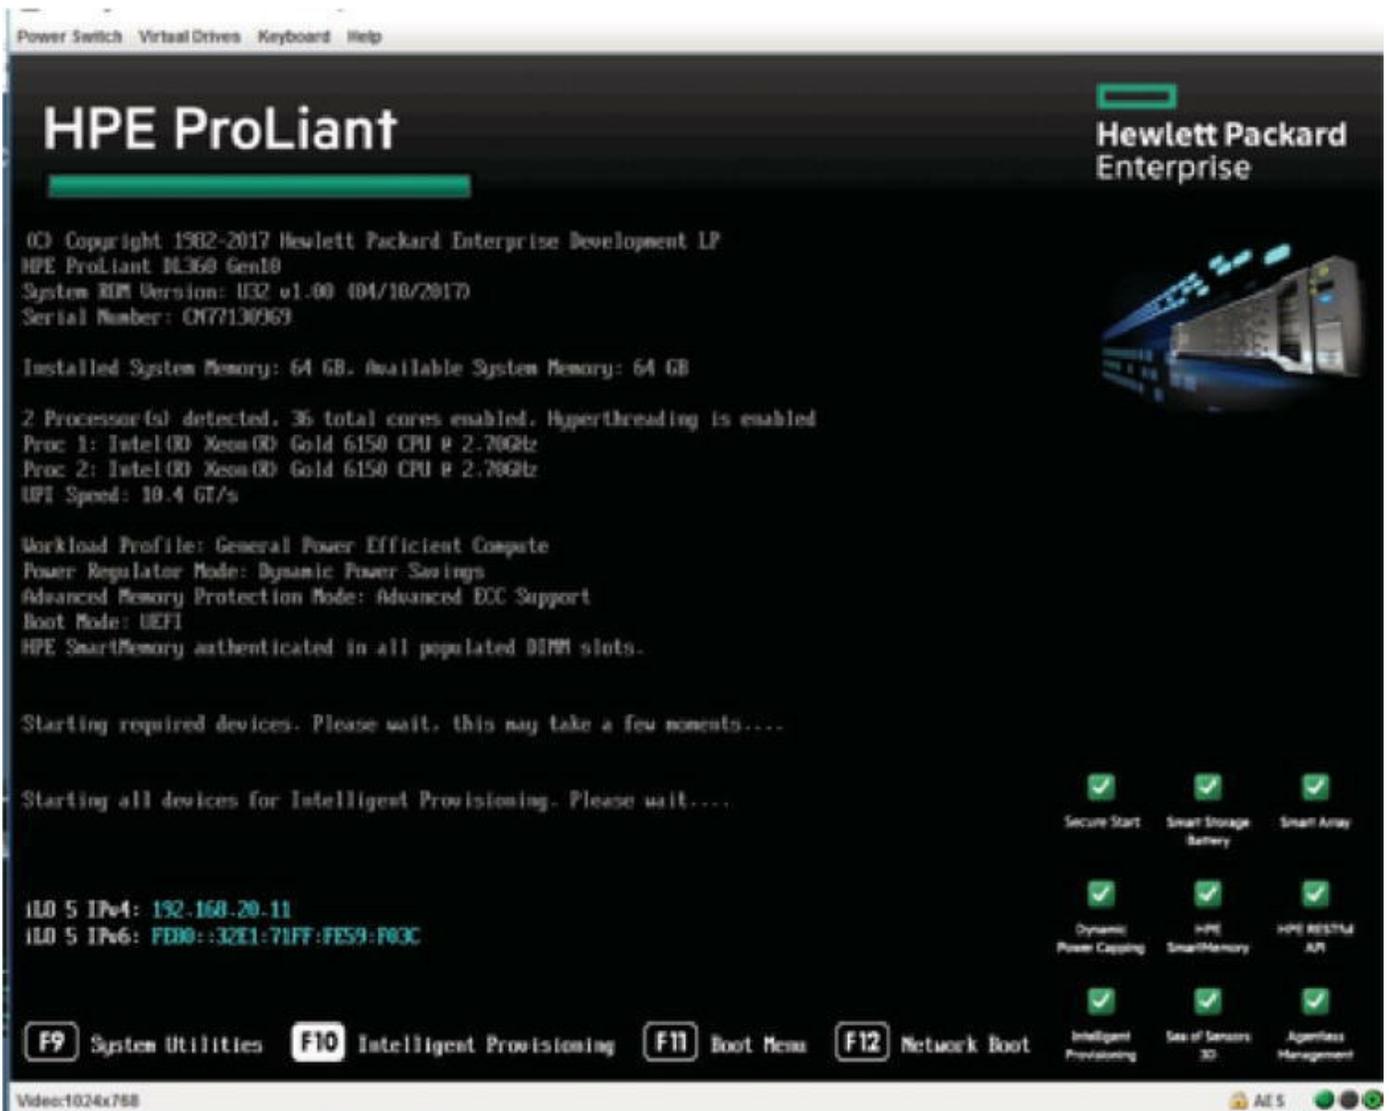


Figure 9-10 ProLiant POST screen

The following steps are general instructions for booting up ProLiant Gen10 servers:
Install the server in the rack or the server blade in the enclosure.

- Connect all peripheral devices.
- Connect the power cord to the power supply.
- Connect the Ethernet cable.

Press the **Power On/Standby** button.

Near the end of the boot process, the Power-On Self-Test (POST) sequence runs and the POST screen displays, as shown in [Figure 9-10](#). This screen is visible for several seconds before the system attempts to boot from a supported boot device. During this time, you can do the following:

- To modify the server configuration ROM default settings, press **F9** to display the System Utilities screen.
- If you do not need to modify the server configuration and are ready to install the OS software, press **F10** to access Intelligent Provisioning.
- Access the Boot menu by pressing the **F11** key.
- Force a network boot by pressing the **F12** key.

To exit the System Utilities screen and reboot the server, select one of the following options:

- **Press F7** or click the **F7: Load Defaults** button at the bottom of the screen.
- **Press F10** or click the **F10 Save** button, this saves any changes pending
- **Press F12** or click the **F12: Save and Exit** button, in the pop-up select cancel or click OK to **Exit and resume system boot**—To exit the system and continues the normal boot process. The system continues through the boot order list and launches the first bootable option in the system.
- **Click the Exit** button, in the pop-up shown in [Figure 9-11](#) select **Cancel** or **OK** to reboot the system.

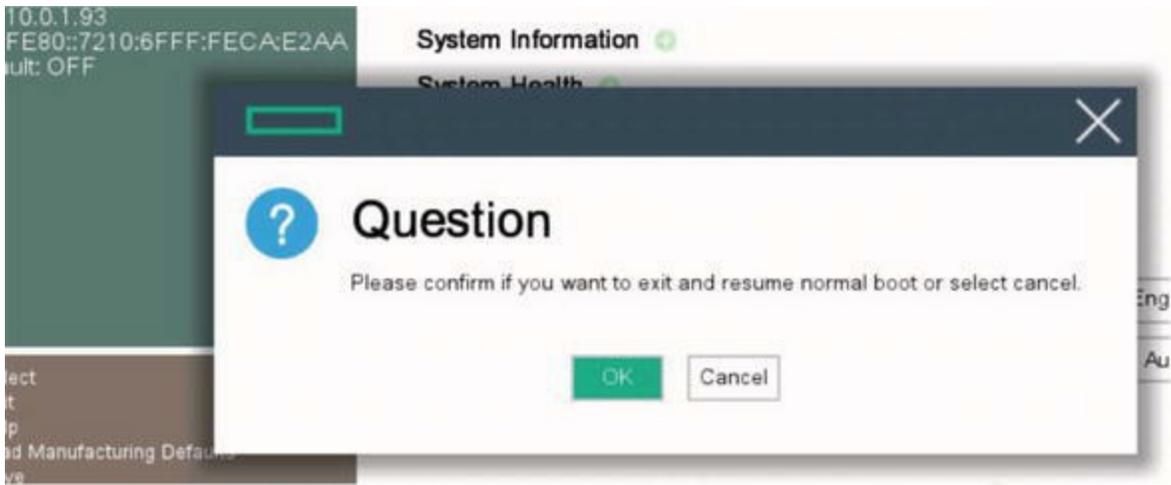


Figure 9-11 Click OK to confirm reboot

ProLiant Gen10 preboot environment

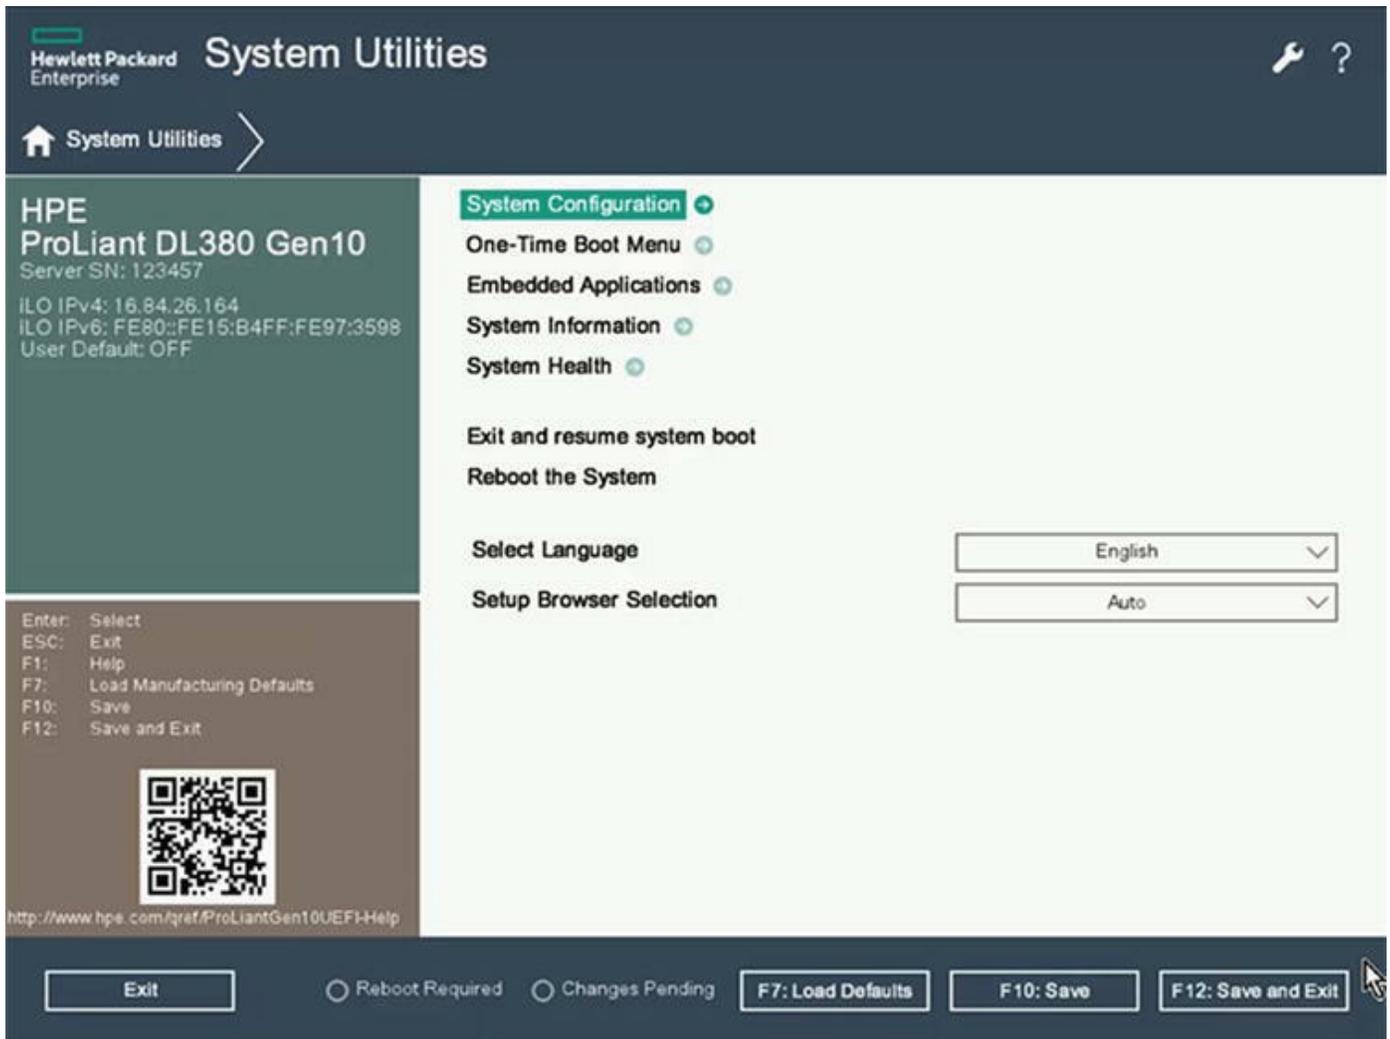


Figure 9-12 ProLiant System Utilities

The System Utilities screen, shown in [Figure 9-12](#), is the main screen in the UEFI menu-driven interface. The System Utilities screen displays menu options for the following configuration tasks:

System Configuration—Displays options for viewing and configuring the BIOS/Platform Configuration (RBSU) menu and the iLO 5 Configuration Utility.

One-Time Boot Menu—Displays options for selecting a boot override option and running a UEFI application from a file system.

Embedded Applications—Displays options for viewing and configuring embedded applications, including Intelligent Provisioning and firmware updates.

System Information—Displays options for viewing the server name and generation, serial number, product ID, BIOS version and date, power management controller, backup BIOS version and date, system memory, and processors.

System Health—Displays options for viewing the current health status of all devices in the system.

Exit and resume system boot—Exits the system and continues the normal booting process.

Reboot the System—Exits the system and reboots it by going through the UEFI Boot Order list and launching the first bootable option in the system. For example, you can launch the UEFI Shell, if enabled and listed as the first bootable option in the list.

Select Language—Enables you to select a language to use in the user interface. English is the default language.

BIOS/Platform Configuration

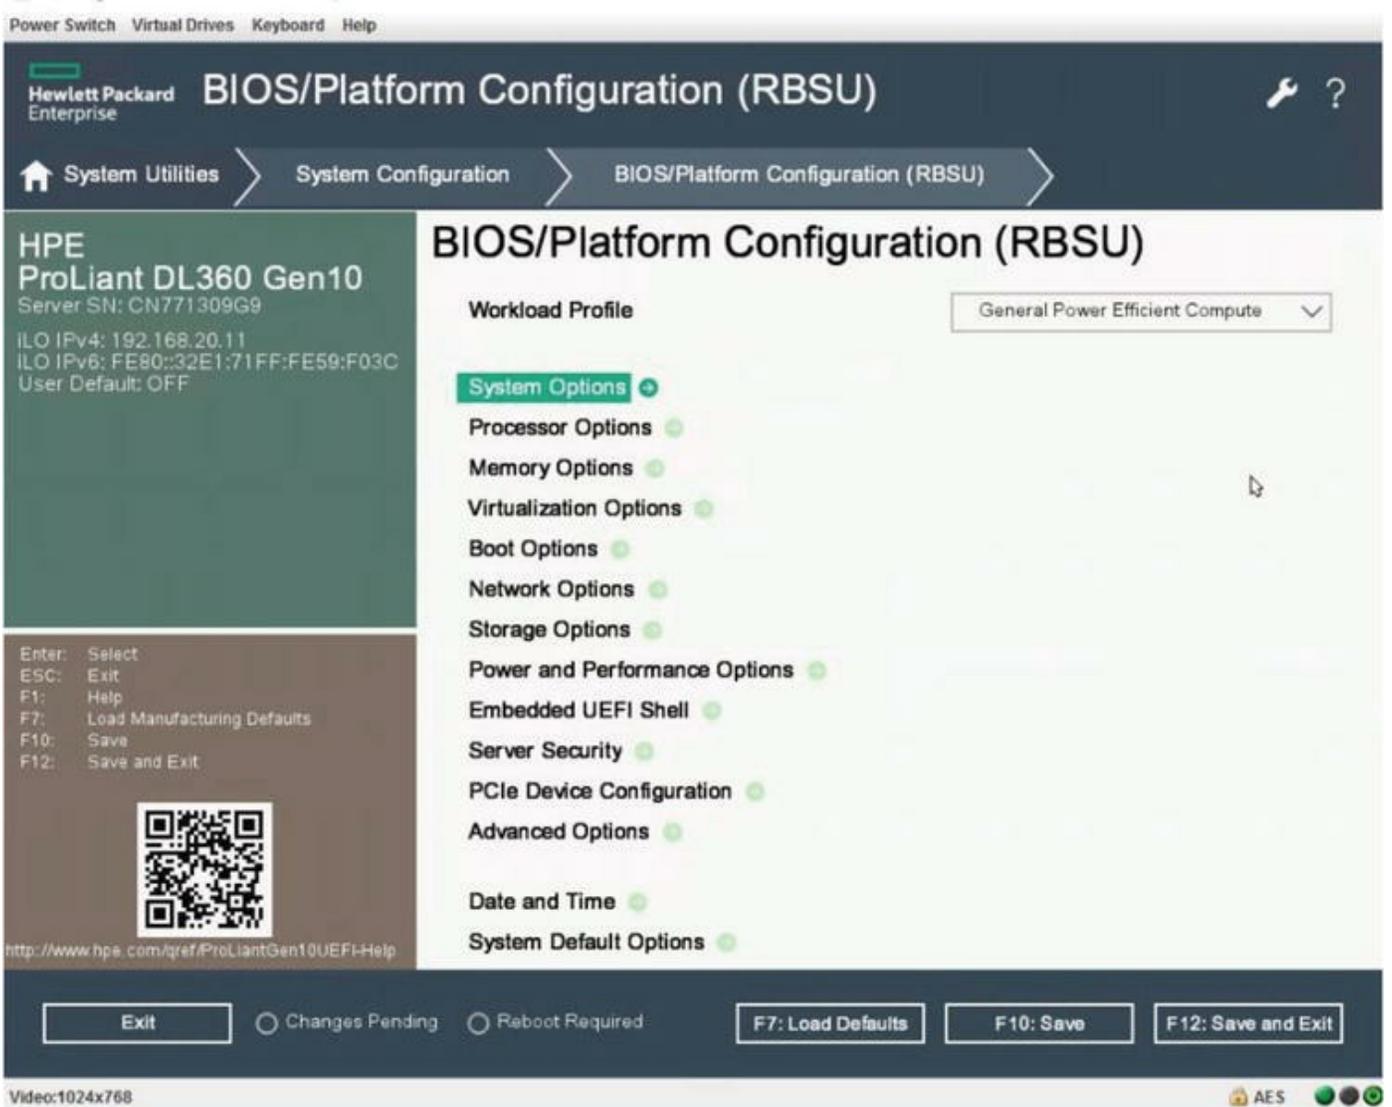


Figure 9-13 BIOS/Platform Configuration

The BIOS/Platform Configuration menu shown in [Figure 9-13](#) replaces the ROM-Based Setup Utility (RBSU) on previous generation ProLiant servers. Use this menu to access and use both UEFI and legacy BIOS options, including:

Workload Profile—Select this option to choose a workload profile for power and performance optimization

System Options—This menu displays several options, including:

- **Boot Time Optimization**—Choose to enable/disable Extended Memory Test, Memory Fast Training, UEFI POST Discovery Mode, and Memory Clear on Warm Rest.
- **USB Options**—Chose to determine USB control, USB Boot Support, enable Virtual Install Disk, and enable Internal SD Card Slot.

- Server Availability—Allows various server availability options to be selected or enabled. Automatic Server Restart (ASR) status and timeout delay. POST F1 Prompt control, Automatic Power-on, and Power-On Delay options can be set.
- Server Asset Information—The Server Information, Administrator Information, Service Contact information, and Custom POST Message options can be set here.

Processor Options—Allows Intel Hyper-Threading control, Cores enabled, and distributed interrupt handling control.

Memory Options—Allows Advanced Memory Protection control and other advanced memory options to be set.

Virtualization Options—Allows control over the Intel Virtualization Technology (Intel VT) to be controlled.

Boot Options—This menu displays several options, including:

- Boot Mode—Use this option to set the boot mode for the system. ProLiant Gen10 servers provide two boot mode configurations: UEFI Mode and Legacy BIOS Mode. Certain boot options require that you select a specific boot mode. By default, the boot mode is set to UEFI Mode.



Important

The boot mode you use must match the OS installation. If not, changing the boot mode can impact the ability of the server to boot to the installed OS. You must reboot the server after making changes to the boot mode.

- UEFI Optimized Boot—Use this option to enable or disable UEFI Optimized Boot, which controls the video settings that the system BIOS uses. If you are running Microsoft Windows 2008 or Windows 2008 R2 OSs and the system is configured for UEFI Mode, this option must be set to **Disabled**. Legacy BIOS Mode components are needed for video operations in Windows. Boot Mode must be set to **UEFI Mode** when this option is enabled. UEFI Optimized Boot must be enabled in order to enable and use Secure Boot, and to operate VMware ESXi.

Network Options—From this menu, you can configure network boot options, the preboot network settings, and iSCSI and VLAN settings.

Storage Options—This option is used to configure the embedded serial AT attachment (SATA) Controller options, the embedded Storage boot policy, Fibre Channel/FCoE scan policy, and NVM Express Options.

Power and Performance Options—This option allows the Dynamic Power savings mode, Intel Turbo Boost, and other power settings to be configured.

Embedded UEFI Shell options—The Embedded UEFI Shell is a preboot command-line environment for scripting and running UEFI applications, including UEFI boot loaders. The Embedded UEFI Shell also provides CLI-based commands you can use to obtain system information, and to configure and update the system BIOS. Enabling this option adds the Embedded UEFI Shell to the UEFI boot options.

Server Security—Use this menu to set the power-on and admin passwords, to enable or disable the Intelligent Provisioning F10 prompt, the One-Time Boot Menu (F11 Prompt), and Backup ROM Image Authentication.

PCI Device Enable/Disable—Use this option to enable or disable embedded and add-in PCI devices. Disabling devices re-allocates the resources (memory, IO, and ROM space and power) that are normally allocated to the device. By default, all devices are enabled.

Advanced options—This menu option controls the ROM Selection—Option to Use Current ROM or switch to backup ROM. Other option to control POST debug message reporting and advanced service options are found here.



Note

UEFI system configuration options vary by Gen10 platform. You might not see some options if they are not available on your particular server.

Decreasing deployment time by significantly reducing boot time

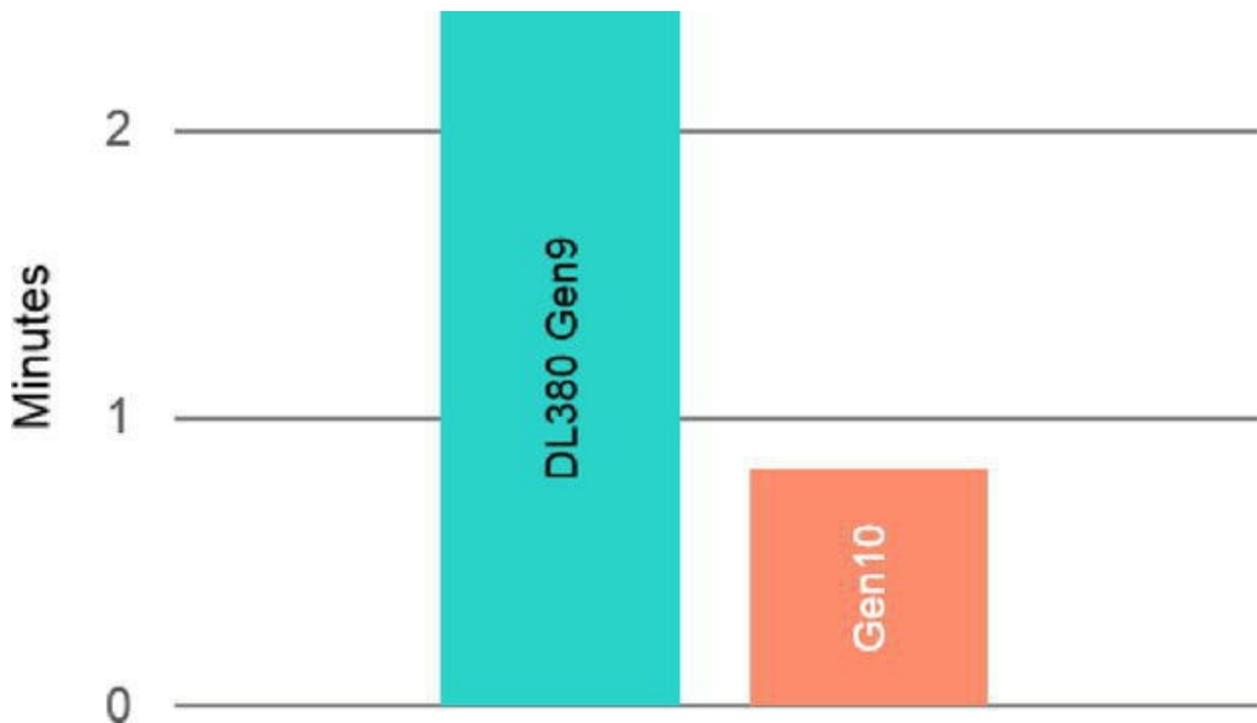


Figure 9-14 Gen10 decreases deployment time by significantly reducing boot time

The boot time optimizations control various tasks that may be performed during the boot process. The aim is to reduce the time taken to boot. However, this may be at odds with a need to ensure that the server is fully tested for use before loading and running the OS.

As shown in [Figure 9-14](#), boot time is significantly reduced by as much as 67% by modifying the boot time optimizations.

Extended Memory test

Memory Fast Training

UEFI POST Discovery Mode

Memory Clear on Warm Reset

Performing shell operations in UEFI

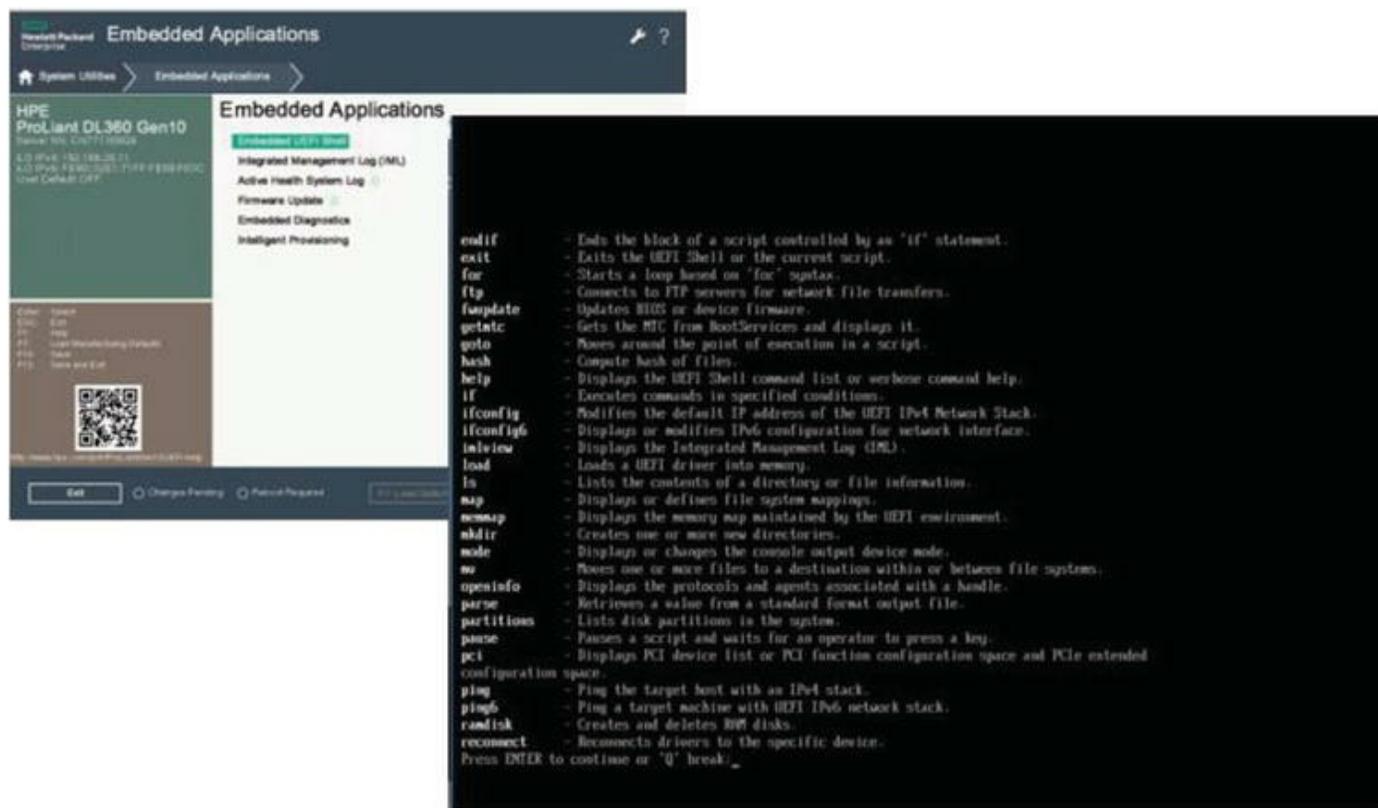


Figure 9-15 UEFI Shell interface

The system BIOS in all ProLiant Gen10 and Gen9 servers and Synergy compute modules includes an Embedded UEFI Shell in the ROM, as shown in [figure 9-15](#). Based on the UEFI Shell specification, the Shell environment provides an API and a CLI that enables scripting, file manipulation, and obtaining system information.

The Shell also runs other UEFI applications. These features enhance the capabilities of the UEFI System Utilities. Access to the Embedded UEFI Shell is enabled by default.

A small selection of UEFI Shell commands and their purpose are listed below:

sysconfig BIOS configuration

fwupdate ROM, storage, and NIC adapter firmware updates

boot Boots or displays boot options

help Lists all of the UEFI Shell commands

help ls -b Displays help for the ls command, the -b option formats the screen output in pages

BIOS RESTful API moves to industry-standard Redfish



```
"iprest command added to your path"
"Type iprest --help for usage"

C:\Windows\System32\iprest
iprest : HP RESTful Interface Tool version 1.3.0
Copyright (c) 2015 Hewlett Packard Development Company L.P.

iprest > types
Enumerating data from iLO... Done
WARNING: Cache is activated. Session keys are stored in plaintext.
type options:
AccountService 1.0.0
BaseNetworkAdapter 1.1.0
Bios 1.0.0
Collection 1.0.0
ComputerSystem 1.0.0
EthernetNetworkInterface 0.92.0
FirmwareService 1.0.0
MemoryInventory 1.0.0
PowerMeter 1.0.0
SystemConfiguration 0.10.0
System 1.2.0
```

Figure 9-16 Industry-standard Redfish helps standardize how configuration settings are scripted

The iLO GUI can be used to perform the initial configuration of the server for deployment in the data center. However, when there are many servers to configure and prepare a scripted deployment method can be more efficient. iLO 5 supports REST APIs that can be accessed using various scripting tools, as shown in [Figure 9-16](#). Features of the industry-standard Redfish include:

BIOS RESTful API introduced in Gen9 ProLiant.

- Provides ability to configure all BIOS configuration settings.

Standardization of BIOS configuration via Redfish lead by HPE.

- BIOS schema definitions moved to Redfish standard.

ProLiant Gen10 will be deprecating the HPE version of these objects.

- Gen10 only supports industry standard Redfish.
- Causes change to BIOS schemas and attribute registry definitions.
- Minor changes to customer scripts using the RESTful API directly.

HPE REST Tool abstracts most changes.

Most scripts using the HPE REST Tool will not require modification.

Learning check

Which tool should you use to configure the boot options on a ProLiant Gen10 server?

- A. HPE OneView
- B. RBSU
- C. Secure Boot mode
- D. BIOS/Platform Configuration menu

Viewing and updating firmware from the iLO 5 web interface

iLO 5 fully integrates and secures the firmware update process. Firmware images are uploaded into the iLO flash memory for immediate or scheduled deployment. These images can also be used to recover firmware in the unlikely event that the option ROM firmware is corrupt or has been compromised.

The iLO web interface supports the option to perform firmware update and supports the following firmware and software management features:

Viewing installed firmware

Using the Flash Firmware controls to install firmware on the local managed server

Using the Group Firmware Update feature to install firmware on multiple servers in an iLO federation group

Accessing the iLO with integrated Smart Update features. This version of iLO supports the following actions:

- Manage the iLO repository and add saved components to the installation queue
- View and remove install sets and add them to the installation queue
- View and remove components from the installation queue

The best practice is to use SUM to manage the installation queue. You can use the iLO web interface to update the queue by adding or removing an individual component.

Supported firmware types:

iLO firmware or Language Packs—Requires iLO reset, triggered automatically

System ROM/BIOS—Requires a server boot

Chassis firmware—Takes effect immediately

Power Management Controller—Does not require server reboot or a system reset

Programmable Logic (complex programmable logic device [CPLD])—Requires server reboot

Backplane

Firmware updates enhance server and iLO functionality with new features, improvements, and security updates.

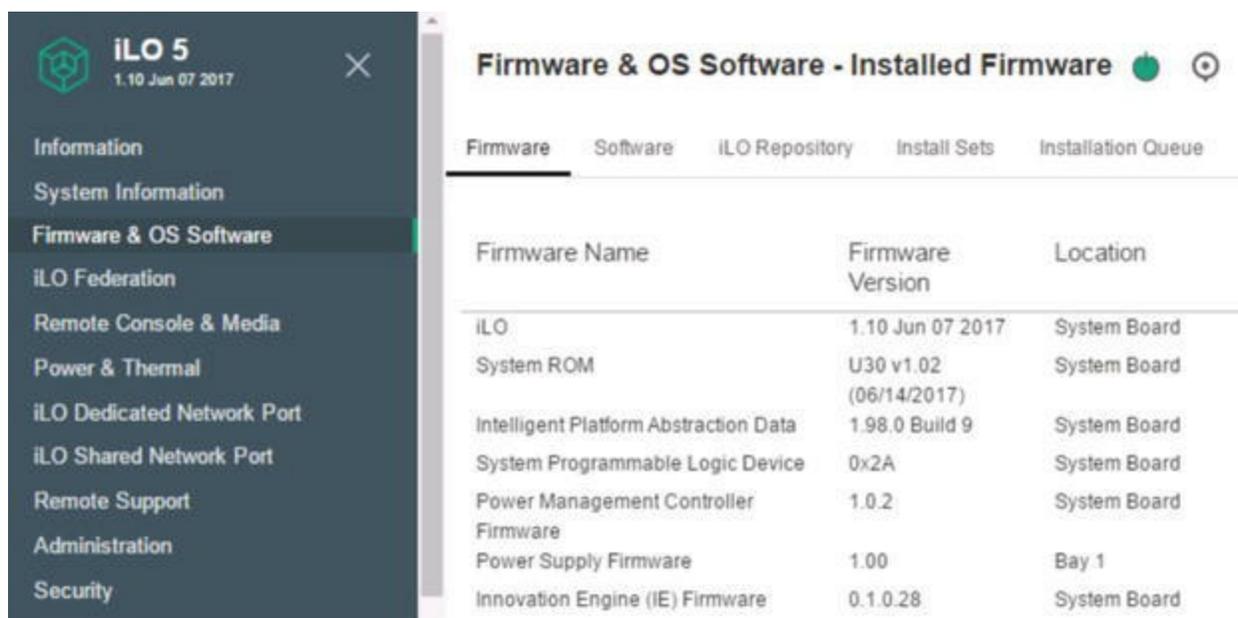


Figure 9-17 iLO 5 Firmware and OS Software view

[Figure 9-17](#) shows the iLO 5 Firmware screen. You can update firmware using the following methods:

Online firmware update—When you use an online method to update firmware, you can perform the update without shutting down the server OS. Online firmware updates can be performed in-band or out-of-band. HPE recommends using online mode when possible.

- In-band—Firmware is sent to iLO from the server host OS. The HPE ProLiant Channel Interface Driver is required for in-band firmware updates. During a host-based firmware update, iLO does not verify login credentials or user privileges because the host-based utilities require a root login (Linux and VMware) or Administrator login (Windows).
- Out-of-band—Firmware is sent to iLO over a network connection. Users with the

Configure iLO Settings privilege can update firmware by using an out-of-band method. If the system maintenance switch is set to disable iLO security, any user can update firmware with an out-of-band method.



Note

Out-of-band management increases the stability and security of system management for firmware updates and agentless management.

Offline firmware update—When you use an offline method to update the firmware, you must reboot the server by using an offline utility. You can use the following offline firmware update methods:

- SPP—Use the SPP to install firmware.
- Scripting Toolkit—Use the Scripting Toolkit to configure several settings within the server and update firmware. This method is useful for deploying to multiple servers.

iLO manual reboot

An iLO manual reboot enables you to reboot iLO without disrupting the host server OS. To perform an iLO manual reboot, use the unit identifier (UID) button:

Five to nine seconds requests a graceful iLO reboot.

10 seconds or longer demand a hardware iLO reboot.

All ProLiant server blades (through the Onboard Administrator GUI) and ProLiant 300 series servers support an iLO manual reboot. This may be required when troubleshooting iLO. You should use this method as a last resort, instead of pulling the power cord. Other, gentler methods include:

XML script (reset_rib.xml)

IPMI Cold Reset

SSH CLI (cd map1, reset)

iLO web GUI (diagnostics page)

Agents

HPE on-system management tools

HPE server management offerings that help you prepare a system for deployment are considered on-system management tools. These offerings include a set of server management capabilities that are embedded on all HPE servers. They are designed to meet the needs of any organization, from enterprise to smaller IT environments.

On-system management includes products and tools available to all ProLiant server customers, such as:

HPE iLO 5—Functions out-of-the-box and provides essential remote management and control regardless of the servers' state of operation

HPE SUM/SPP solution—Provides systematic and organized server updates through SPP or SUM

HPE Intelligent Provisioning—Assists in the initial configuration and OS deployment for ProLiant servers

UEFI—Is the BIOS Mode for ProLiant Gen9 and Gen10 servers that aids with server configurations

HPE RESTful Interface Tool—Leverages HPE RESTful API for secure, remote server configurations at scale

HPE Scripting Toolkit—Is a set of utilities that provide scripting for ProLiant server configurations at scale

HPE Intelligent Provisioning

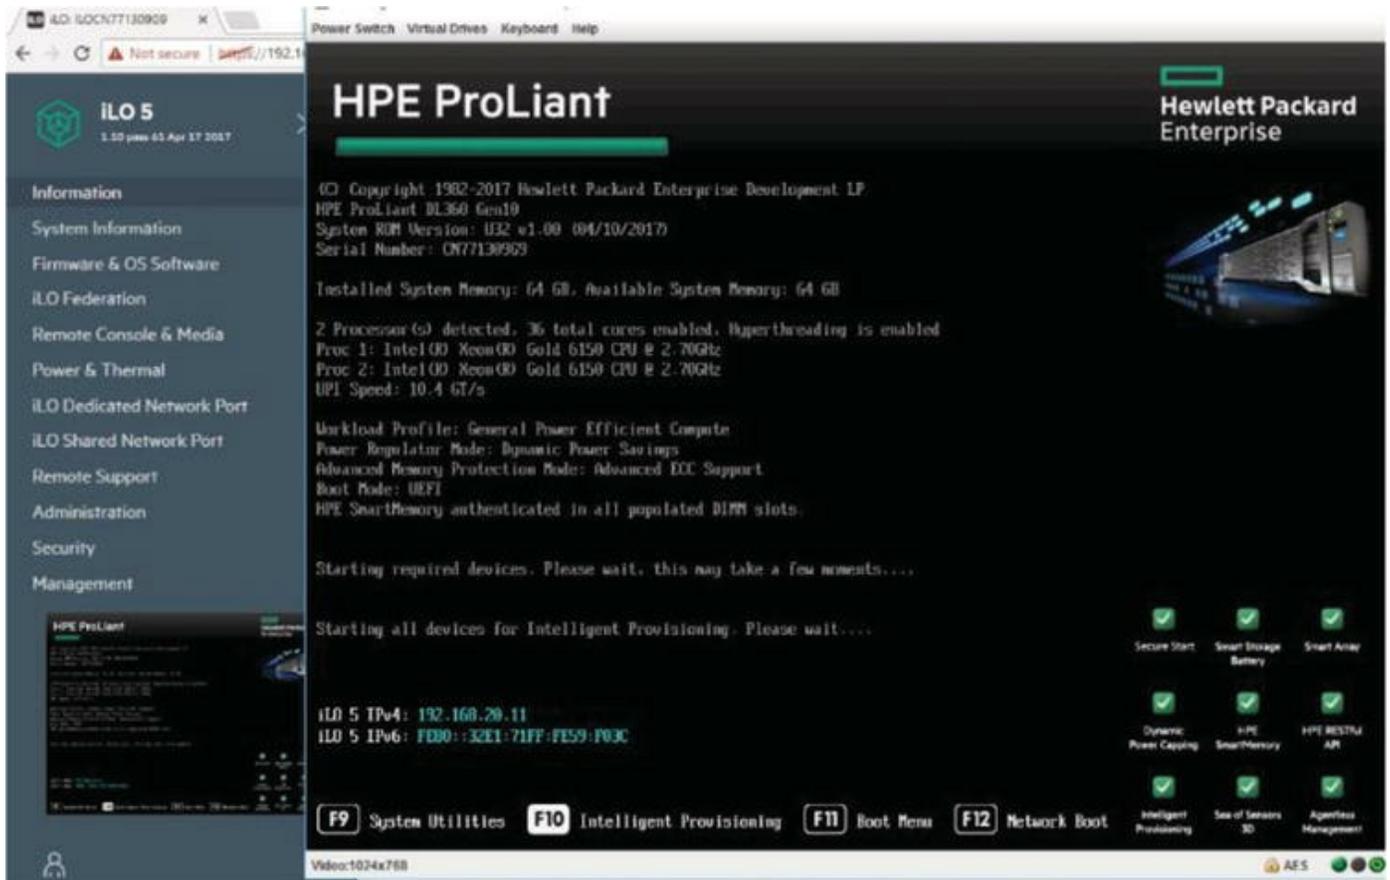


Figure 9-18 Accessing HPE Intelligent Provisioning

Intelligent Provisioning is a single-server deployment tool embedded in ProLiant Gen10 servers that simplifies ProLiant server setup, providing a reliable and consistent way to deploy ProLiant server configurations. It is a preboot startup and diagnostics tool embedded in all currently shipping ProLiant Gen9 and Gen10 servers. It accelerates system configuration so you can quickly and easily get your system up and running. In addition, the built-in smart update capability recognizes if the system software is out of date and automatically downloads the latest update for you.

Residing on a NAND chip, Intelligent Provisioning enables you to quickly deploy servers without having to rely on physical configuration media. All the firmware, drivers, and tools you need are preloaded and ready to deploy. Intelligent Provisioning eliminates many server setups, configuration, and deployment steps, so you can deploy servers much faster than before.

Use Intelligent Provisioning to:

- Deploy an OS to Smart Array-based storage

- Manage and update necessary drivers not included in base OS media

Perform maintenance-related tasks such as updating firmware and configuring iLO and HPE Smart Array controllers

Create and edit a collection of configuration settings, save them in a portable package, and deploy them to many servers via iLO 5 or a USB key

Deploy an OS to two SD cards supported by HPE

To launch Intelligent Provisioning, power up the server and press the **F10** key during system boot, as shown in [Figure 9-18](#). You can also access Intelligent Provisioning from the System Utilities screen. Select **Embedded Applications** and press **Enter**. Then select **Intelligent Provisioning** and press **Enter**. The Intelligent Provisioning home screen is displayed. From here, select one of the following menus to use Intelligent Provisioning:

Configure and Install menu—Enables you to configure hardware and install an OS

Perform Maintenance menu—Enables you to access the configuration utilities you need to maintain a ProLiant server

When logging on to Intelligent Provisioning for the first time, the Activation screen is viewable after the Setting Preferences screen. To activate Intelligent Provisioning and make it available during POST, follow the steps to set preferences and register the server with Insight Remote Support. To exit Intelligent Provisioning, reboot the server by clicking the power icon.



Note

For more information, refer to the HPE Intelligent Provisioning User Guide. Access the user guide by entering the following URL into your web browser;

https://h20564.www2.hp.com/hpsc/doc/public/display?docId=emr_na-a00017037ja_jp



Note

A Smart Start CD is no longer required since Gen8 servers. It has been replaced by Intelligent Provisioning which is embedded in the flash ROM. To run this deployment environment, you simply press the F10 key during startup and wait for Intelligent Provisioning to load.

Intelligent Provisioning startup menu

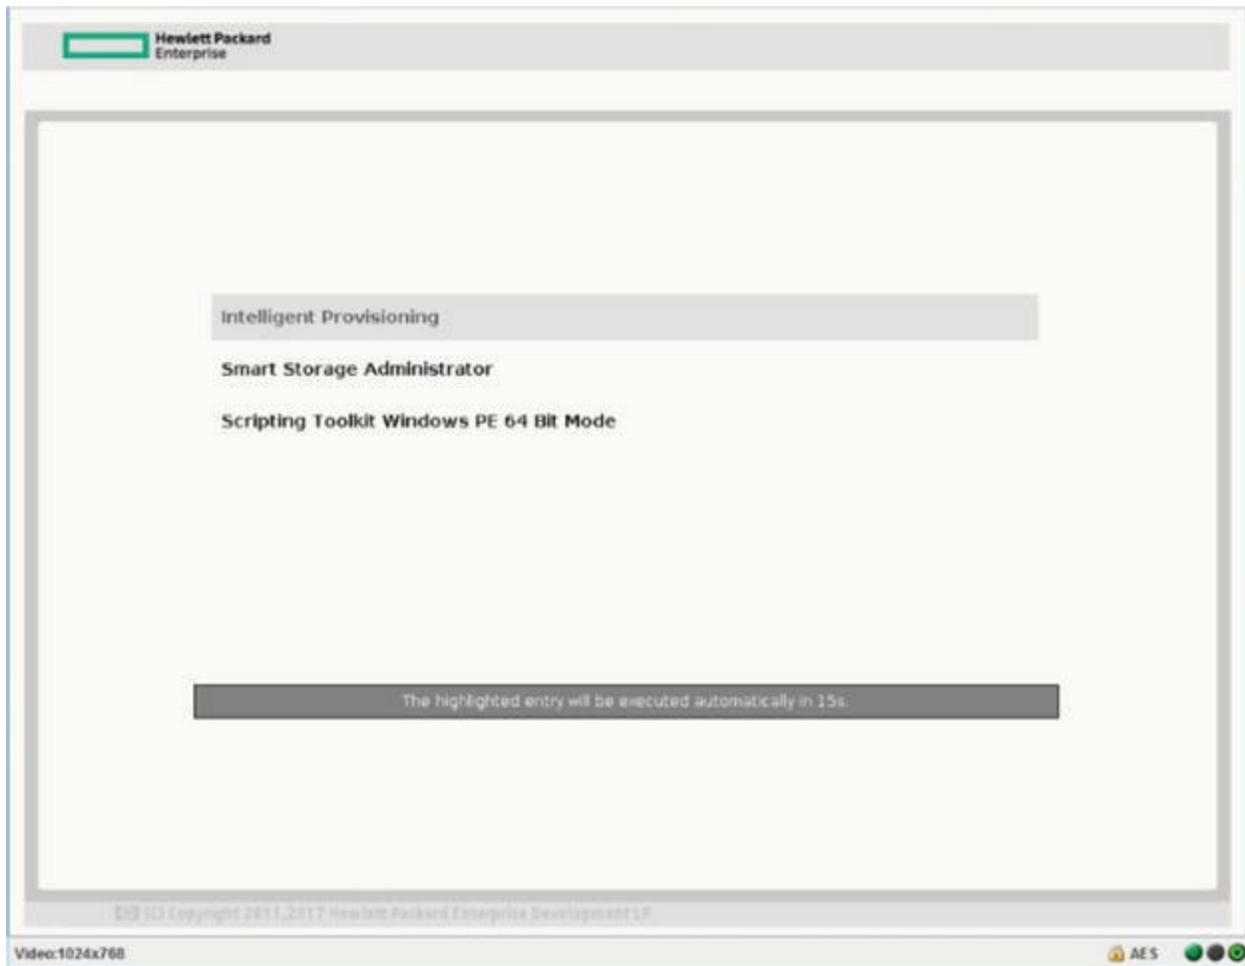


Figure 9-19 Intelligent Provisioning startup menu

As shown in [Figure 9-19](#), Intelligent Provisioning startup menu options are:

Intelligent Provisioning

SSA

Scripting Toolkit Windows PE 64-bit mode

The highlighted option starts automatically in 15 seconds.

Home screen

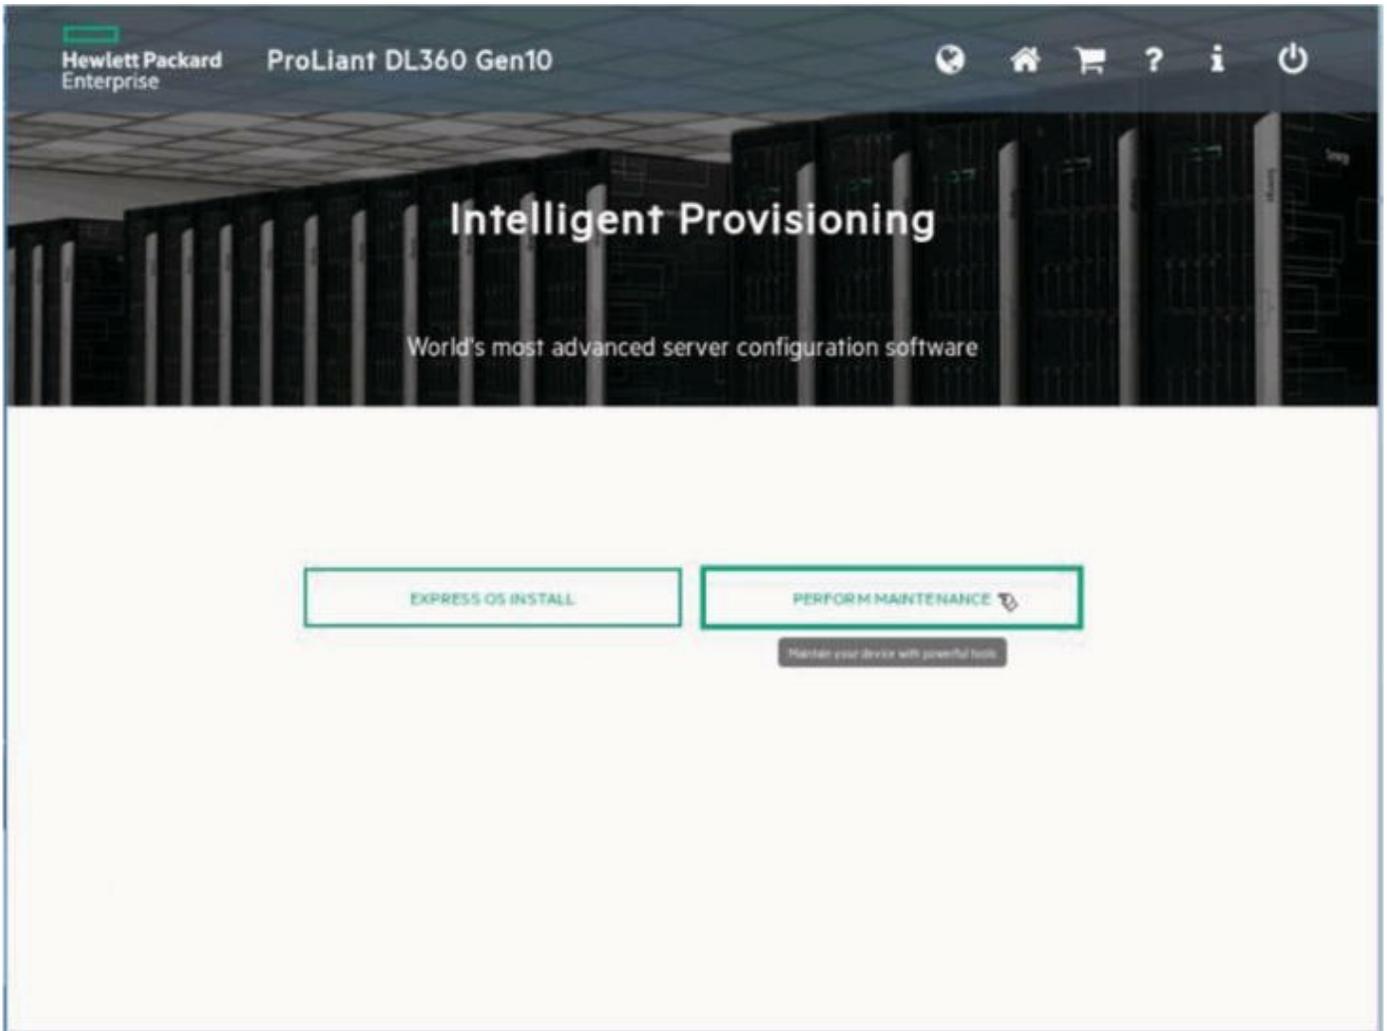


Figure 9-20 Intelligent Provisioning home screen

[Figure 9-20](#) identifies two options to be used when beginning the server deployment process, after the Intelligent Provisioning boot process has completed and the first time configuration has been performed. The first option **Express OS Install** provides assisted OS installation. The second option **Perform Maintenance** provides a selection of tools that can be used to complement the OS install and perform maintenance tasks. These include:

Attempt Firmware Update

BIOS Configuration (RBSU)

Intelligent Provisioning Preferences

iLO Configuration

Active Health System Log

System Erase and Reset

Deployment Settings

RAID Configuration (launch SSA)

Intelligent Provisioning Perform Maintenance options

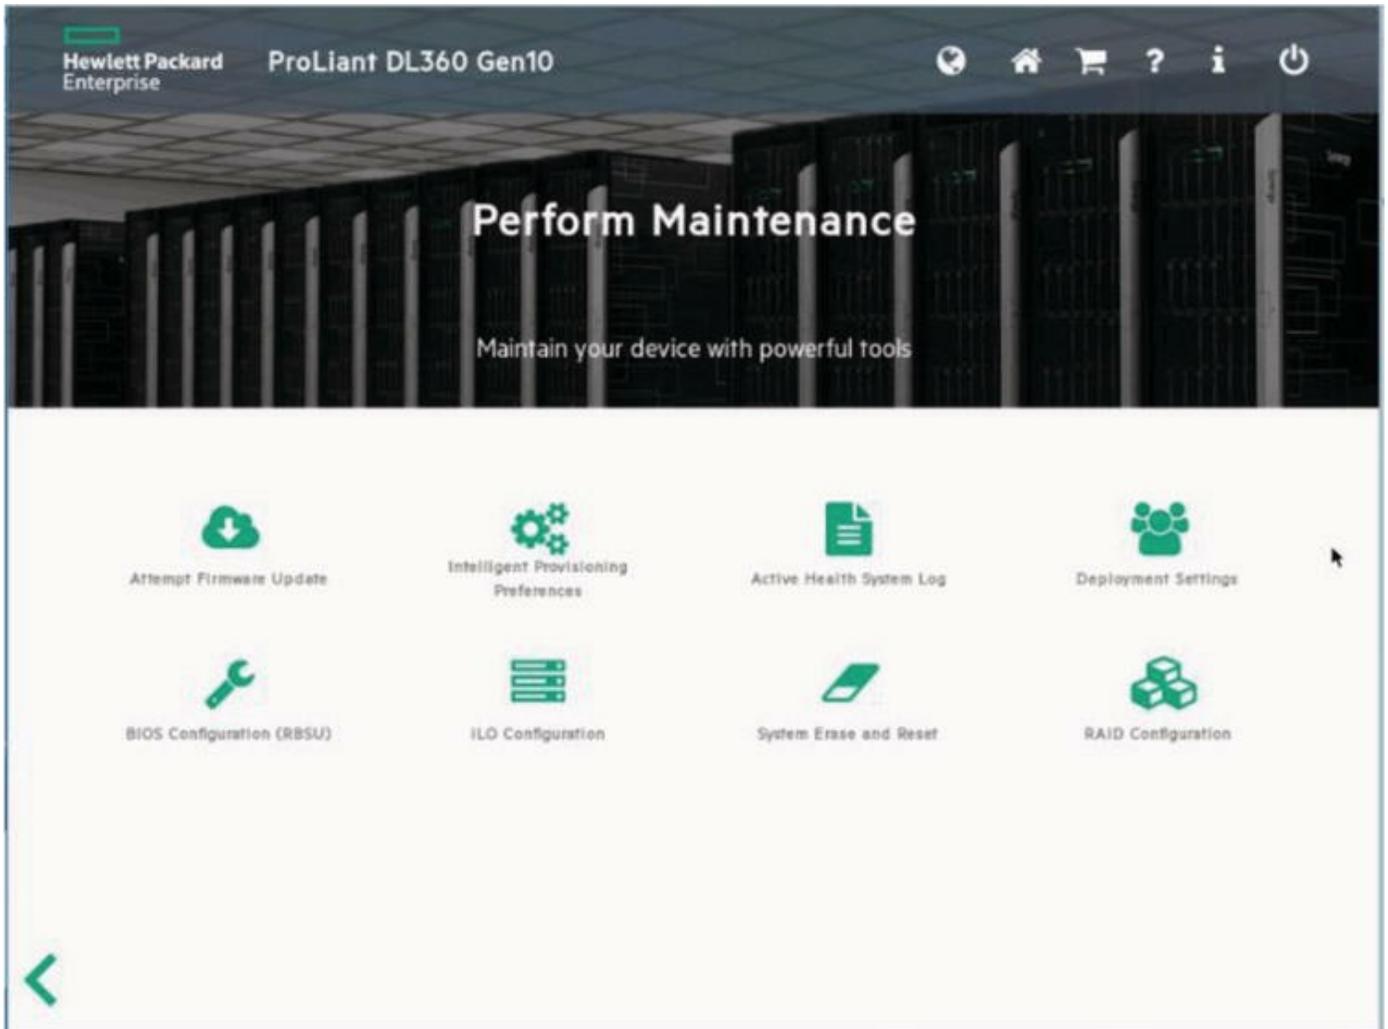


Figure 9-21 Intelligent Provisioning Maintenance screen

Clicking the **Perform Maintenance** link on the Intelligent Provisioning home screen displays the screen shown in [Figure 9-21](#), with several options for configuration utilities needed to maintain ProLiant servers, such as:

Attempt Firmware Update—Intelligent Provisioning manages the firmware update process. The SPP firmware repository URL is configured in the Provisioning

preferences and requires network interface configuration.

Setting Intelligent Provisioning Preferences—When Intelligent Provisioning is run for the first time on a server, this is the first screen that Intelligent Provisioning displays. Use this screen to change basic preferences, including the interface and keyboard languages, network and share setting, system date and time, and software update settings. The EULA is also accessible from this screen.

Downloading Active Health System—HPE Active Health System monitors and records changes in the server hardware and system configuration. Active Health System assists in diagnosing problems and delivering rapid resolution when server failures occur. The Active Health System Log, together with the system monitoring provided by Agentless Management or Simple Network Management Protocol (SNMP) Pass-thru, provides continuous monitoring of hardware and configuration changes, system status, and service alerts for various server components. The Agentless Management Service (AMS) is available in the SPP, which is a disk image (.iso) that you can download from the HPE website. The Active Health System Log can be downloaded manually from iLO or Intelligent Provisioning and sent to HPE.

Using Deployment Settings—Deployment Settings enable you to create a server configuration package that you can deploy to one or more ProLiant Gen9 servers and server blades using a USB key and iLO Scripting. Using Deployment Settings is an alternative to using the HPE Scripting Toolkit.

BIOS Configuration (RBSU)—This feature provides the capability to set and change the server's BIOS settings.

Using iLO Configuration Utility—Use the iLO Configuration Utility to change iLO settings through Intelligent Provisioning instead of through the iLO web interface. You must configure iLO network settings to register the server with HPE Insight Remote Support. If the display is not as expected or an undefined error message is displayed, it is possible that the change was being made at the same time through the iLO web interface. Press **F5** to refresh the screen display.

Using the Erase Utility—Use the Erase Utility to clear hard drives and the Active Health System logs, and to reset the RBSU settings in the UEFI System Utilities. The Secure Erase option provides secure erase functionality for the internal system storage and hard disks following the guidelines outlined in DoD 5220.22-M. Secure erase overwrites all block devices attached to the system through applying random patterns in a three-pass process. These block devices include hard disks, storage systems attached to the server, as well as the internal storage used by iLO. Depending on the amount of storage installed on a system, the secure erase process can take many hours or even days to complete.



Notice

Perform a backup before running the Erase Utility. The utility sets the system to its original factory state, deletes the current hardware configuration information, including array setup and disk partitioning, and erases all connected hard drives completely. Based on your selections in this utility, the Erase Utility erases data from the specified areas of the system. Existing data and system configuration settings are lost.

Secure Erase should be used with extreme caution and only when a system is being decommissioned. The secure erase process will reset iLO and lose all licenses stored there, reset BIOS settings in many cases, and delete all AHS and warranty data stored on the system. iLO will reboot multiple times after the process is complete. Disconnect any FCoE, iSCSI, external Serial-attached SCSI (SAS), and Fibre Channel storage before using secure erase, unless they should also be erased.

RAID Configuration—Use this to launch SSA.

Intelligent Provisioning OS installation

The steps to install the Intelligent Provisioning OS are:

Select Install OS

Mount the media in iLO

Confirm selection

Enter customization

Confirm settings

Launch installation

Initiate reboot

To use Intelligent Provisioning to configure the hardware and install an OS on a ProLiant server, follow the prompts in the Express OS Install menu to complete the steps on these screens:

Hardware Settings—At the first configuration screen, enter global settings to control power use, software updates, and RAID array configuration. Then select whether to perform a software update before the OS is installed. Additional options allow you to

enter array configuration specifications for the server's storage subsystem, select a target controller, and select the install method.

OS Information—Depending on the OS you are installing, you might be prompted to enter or confirm information regarding the partition size, keyboard language, or product key. During installation, you can deploy HPE StoreVirtual VSA, which enables you to create fully featured shared storage on a virtualized server. VSA software is a virtual machine that supports hypervisor environments. It provides shared storage for VMware ESX/ESXi and Microsoft Hyper-V hypervisor environments. When deploying VSA, you can retrieve a free 1 TB license key and activate it.



Note

The deployment of VSA is supported only during recommended and custom installations for Windows, and only during custom installations for VMware. It is not supported during manual installation. The server must be connected to a network to deploy VSA.

Review—The Review screen displays your hardware and OS settings. Continuing past this screen installs the OS and configures the server.



Notice

Continuing past this screen resets the drives to a newly installed state and installs the selected OS. Any existing information on the server is erased. This does not affect the first-time setup because there is no data present on the server.

Learning check

A customer approaches your stand at a trade fair and informs you that they are looking to upgrade their G7 servers to Gen10 servers. They ask if they still need to use the Smart Start CD to deploy Windows? How should you respond?

HPE Smart Update solution stack

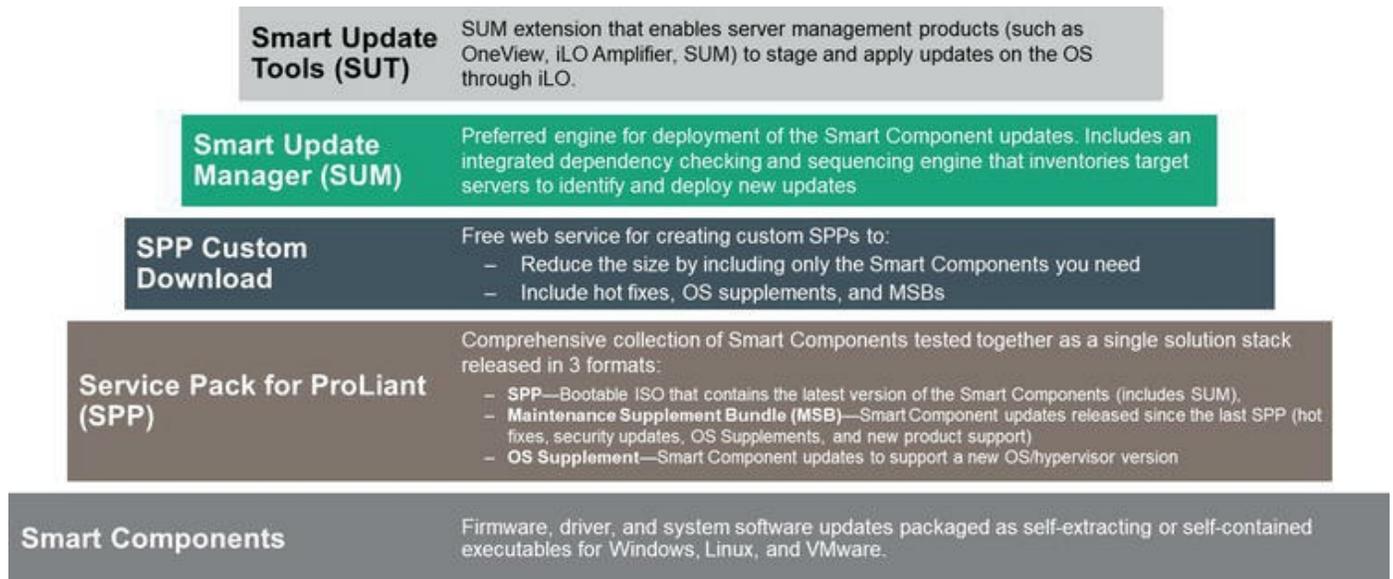


Figure 9-22 HPE Smart Update solution stack

The HPE Smart Update solution stack, shown in [Figure 9-22](#), is designed to help with keeping the firmware, drivers, and system software of HPE server infrastructure and associated options up-to-date and secure.

HPE Smart Update Tools

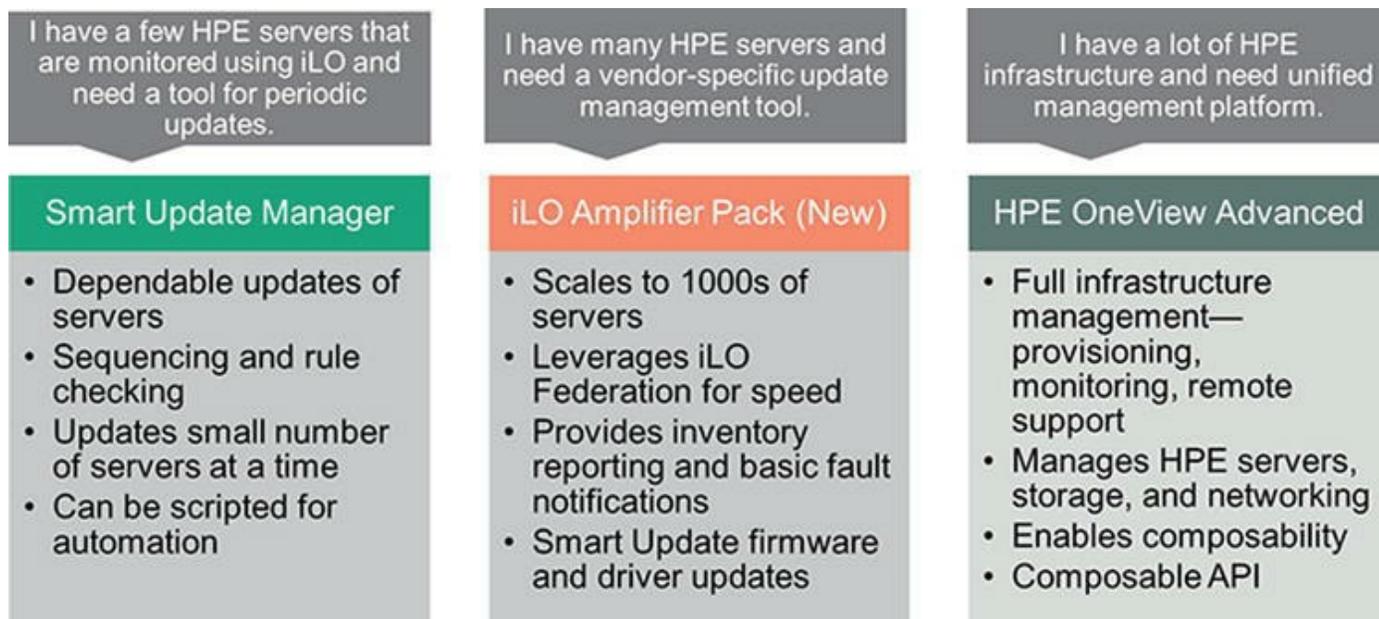


Figure 9-23 Smart Update Tools and how they help customers

HPE Smart Update Tools (SUT) are a SUM extension that enables HPE OneView and HPE iLO Amplifier Pack to stage, schedule and apply updates automatically to reduce IT operation overhead. SUT is an OS utility that provides the ability to perform online firmware and/or driver updates via the HPE iLO management network without the need for OS credentials. SUT supports Windows, Linux (Red Hat and SUSE), and VMware ESXi. [Figure 9-23](#) shows the benefits and use cases of SUT.

SUT enables HPE OneView and iLO Amplifier Pack to stage and apply updates automatically to reduce IT operations and supports dependency and sequencing management between drivers and firmware.

Additional supported features are:

Support for HPE OneView external repository

Support for HPE OneView staging, scheduling, and deployment of updates

Support for HPE iLO Amplifier Pack

Updates using SUT are performed via the HPE iLO management network without the needs for OS credentials. SUT limits the number of reboots required for activation.

HPE SUM

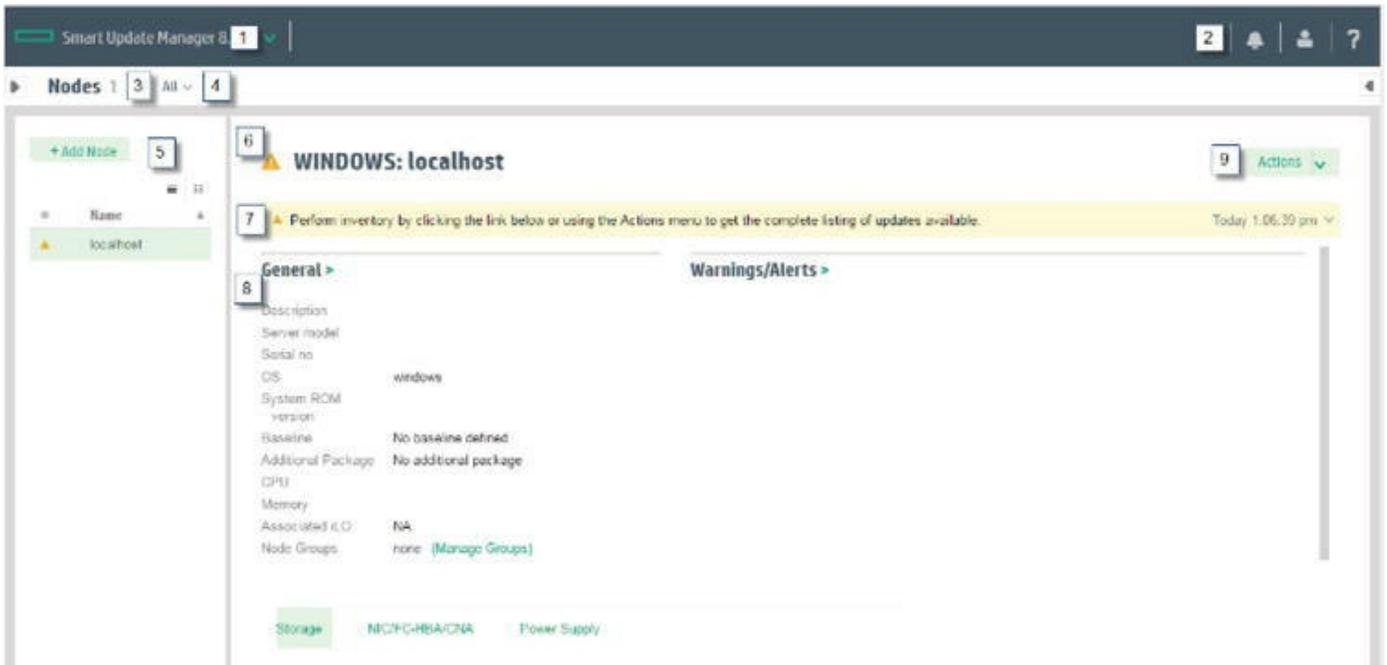


Figure 9-24 HPE SUM GUI

HPE SUM is a program you can download and use to deploy firmware, software, and drivers for ProLiant servers. It has an integrated discovery engine that finds installed hardware and current versions of firmware and software on target servers and identifies associated targets that should be updated in conjunction with each other, to avoid interdependency issues. It is preloaded with all relevant interdependency information from our extensive lab testing before each SPP release.

SUM provides a web-based GUI and a command-line scriptable interface for flexibility and adaptability.

When you launch the GUI, the system uses the default web browser. [Figure 9-24](#) shows the major areas of the GUI. Screens options vary based on the node type, and the legend below describes them.

Main menu—provides links to all HPE SUM screens

Activity list—Displays the status of recent activities

Screen name—Displays the name of the screen

Screen filters—Filters screen objects

Screen list—Displays a list of items on the screen

Baseline or node name—Lists the baseline or node that is selected

Informational highlight—Information you might need to perform or know about a

selected item

Item details—Provides full details about a selected item

Actions—Lists the available actions

Key features of SUM include:

Discovery engine that finds installed versions of hardware, firmware, and software on nodes.

SUM deploys updates in the correct order and ensures that all dependencies are met before deploying an update.

Interdependency checking.

Automatic and step-by-step localhost Guided Update process.

Web browser-based mode.

Ability to create custom baselines and ISOs.

Support for iLO Repository (Gen10 iLO 5 nodes only).

Simultaneous firmware and software deployment for multiple remote nodes.

Local offline firmware deployments with SPP deliverables.

Extensive logging in all modes.

HPE SUM installs updates in the correct order and ensures that all dependencies are met before it deploys each update. It prevents version-based dependencies from destroying an installation, and it ensures that firmware updates are handled in a manner that reduces any downtime required for the firmware update process.

Deploying firmware and software updates together is recommended; however, if you must deploy them separately:

Read/review all release notes.

Update drivers.

Update additional software, such as agents and utilities.

Update firmware.

If you are installing a new server, update the firmware before installing the OS.

SUM does not require an agent for remote installations because it copies a small, secure Simple Object Access Protocol (SOAP) server to the target server for the duration of the installation. After the installation is complete, the SOAP server and all remote files associated with the installation (except installation log files) are removed. HPE SUM

copies the log files from the remote targets back to the system where SUM is executed.

HPE SPP

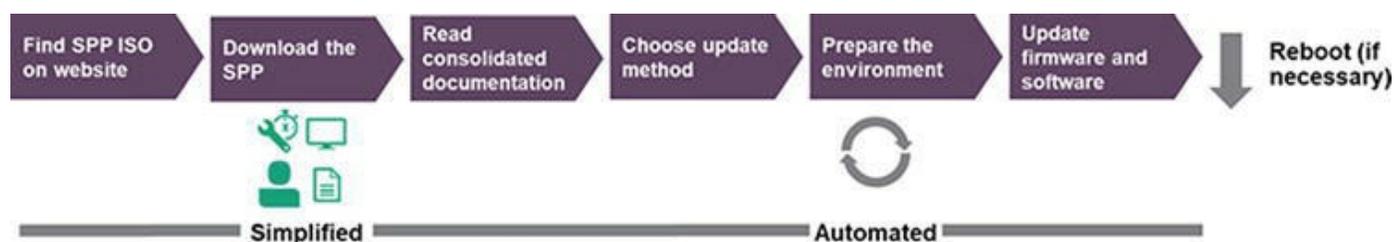


Figure 9-25 HPE SPP simplifies firmware and driver updates

The SPP includes components for HPE ProLiant server and enclosures systems, software, and firmware.

SUM deploys components to supported Windows, Linux, and VMware OSs. [Figure 9-25](#) shows the simplified process using SPP. The single SPP image contains a comprehensive collection of firmware and system software components including drivers, agents, tools, and utilities. These components are tested, managed, and deployed together as a single solution.



Note

Each SPP release includes a version of SUM you can use to deploy the SPP components. Access the SUM website by entering the following URL into your browser;

<http://www.hpe.com/servers/sum-download>

Consisting of SUM and SPP, which reduces time-consuming, expensive, and error-prone updates. The HPE Smart Update website serves as a portal for SUM and SPP.

By enabling firmware and system software to be updated online and integrating these updates into one operation, SPP with SUM offers faster updates of individual ProLiant servers and dramatically faster updates of entire BladeSystem enclosures. SPP with SUM as the deployment engine provides a comprehensive approach to firmware and system software maintenance. Together they provide better operating stability and ensure maximum uptime.

When updating ProLiant servers, you might need to update the firmware on the

following components:

System ROM

iLO

Onboard Administrator and Virtual Connect (VC) modules

Broadcom, Intel, Mellanox, and QLogic NICs

SAS and SATA hard drives

Smart Array controllers

Emulex, QLogic, and Brocade Fibre Channel HBAs and converged network adapters (CNAs)

HPE 6 GB and 12 GB SAS BL switches

Although SPP is a delivery mechanism for firmware and software components for ProLiant servers, SUM is the deployment tool used to plan, schedule, and deploy the components delivered with SPP.



Note

For more information on HPE Smart Update, enter the following URL into your browser;

<http://h17007.www1.hpe.com/us/en/enterprise/servers/solutions/info-library/index.aspx?cat=smartupdate&subcat=sum#.WcDT1tFrztM>

SPP is pretested for component dependencies and is customizable for the environment. You can use it as a baseline and customize it by selecting specific components for deployment to establish a custom baseline. SPP is delivered as a full ISO and is supported for one year.

SPP runs in several modes:

Online mode—Runs on a Windows or Linux-hosted OS.

Offline mode—Server boots to the SPP ISO (bootable ISO only).

- Automatic mode—Firmware components update without interaction.
- Interactive mode—Onscreen instructions guide the user to update firmware components.

SPP supports Microsoft Windows, Red Hat and SUSE Linux, and VMware OSs.

HPE Firmware update technology

All Firmware is securely updated via iLO and orchestrated by SUM.

Firmware update sent to system through iLO network

- Uses iLO authentication and authorization
- iLO checks digital signature and integrity of all firmware images

SUM orchestrates delivery

- Dependency checking
- Sequencing of updates

Images are staged to the 4 GB non-volatile flash partition

RESTful API support for firmware updates

HPE SUM and iLO Federation

When you log in to HPE SUM, it automatically searches for iLO Federation groups on connected networks. SUM searches each group and displays the nodes that respond to the search. You use the Edit scalable update group screen to enter the IP address and user credentials for one node in the iLO Federation group that SUM uses as the interface for inventory and deployment to the nodes in the group.

When you select a group, SUM displays information about the group, including a description, server types, number of servers, and installed firmware versions. SUM only deploys system ROM and iLO firmware to iLO Federation group nodes.

The SUM iLO Federation feature relies on proper configuration of iLO Federation groups before you launch SUM. Multiple iLO Federation groups with the same name or fragmented iLO Federation groups results in SUM only working with a portion of the expected systems.

SUM supports scalability with iLO Federation, including the ability to:

Automatically discover iLO Federation groups on the management network.

Update the iLO and ROM firmware on ProLiant servers in the iLO Federation group through the iLO.

Update all appropriate firmware on ProLiant servers in the iLO Federation group using the iLO using offline firmware deployment.

Deploy VMware vSphere Installation Bundles (VIBs) and Linux firmware RPM

Package Managers (RPMs).

HPE RESTful Interface Tool

The HPE RESTful API is a simple programmatic web interface that allows you to create, read, update, and delete your server infrastructure. This is included in HPE iLO 4 firmware and later to perform tasks such as configuration, inventory, and monitoring of ProLiant servers through iLO.

The RESTful Interface is the next generation infrastructure management interface for iLO. It is based on industry best practice HTTP and web-security standards.

The RESTful Interface is a command-line script tool to harness the API.



Note

For more information on the HPE RESTful API, refer to the document titled “Managing Hewlett Packard Enterprise Servers Using the HP RESTful API.” To download this guide, enter the following URL into your browser;

<http://h20564.www2.hp.com/hpsc/doc/public/display?docId=c04423967>

Activity: Updating firmware with SUM

This activity requires you to use Internet Explorer or Chrome to browse to the emulator index page. On the emulator index page select **Updating Firmware with Smart Update Manager**.

Follow the instructions as they appear on the screen.



Note

Access the emulator by entering the hyperlink into your Chrome or Internet Explorer browser.

<http://hpe.azureedge.net/gen10emulators/index.html>

Activity debrief

Why is it necessary to change the firmware option from force to forced?

Could the reboot after update have been delayed for longer than one minute?

Agentless Management

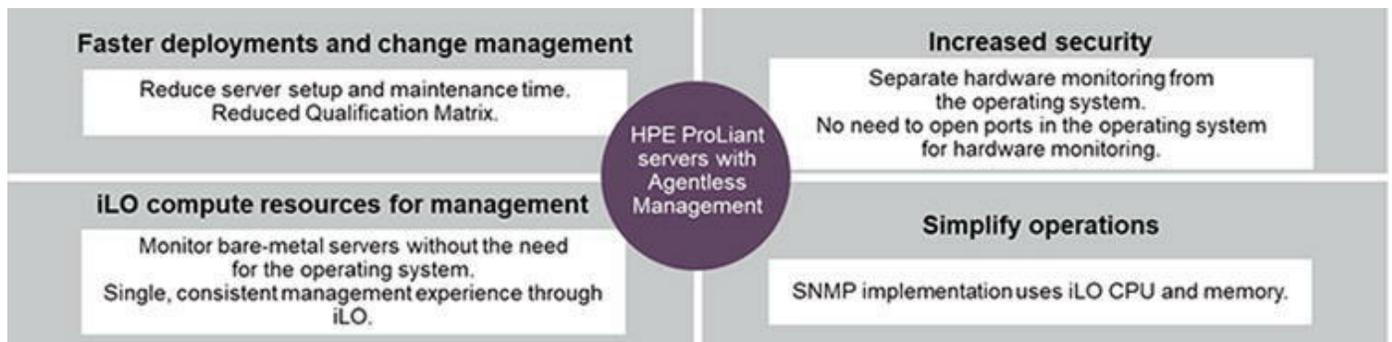


Figure 9-26 iLO 5 Agentless Management

With the Agentless Management solution, the management software (SNMP and Redfish) operate within the HPE iLO firmware instead of the host's OS freeing up memory and processor resources on the host's OS for use by server applications. No agents in the OS means no opening up networks ports reducing security threats, especially for those platforms that might interface directly to the internet. [Figure 9-26](#) shows the benefits of Agentless Management.

With Agentless Management, in addition to monitoring all key internal subsystems (such as thermal, power, and memory), iLO sends SNMP management alerts directly to any data center information management system such as HPE OneView, even with no host OS installed or running. The optional AMS can collect and send OS specific information as well as information of devices that cannot communicate with the iLO, to the iLO through the PCIe bus.

Agentless Management not only simplifies agent management regardless of the host OS, but it also provides an iLO-dedicated management network, isolated from the regular data network. Isolating agent network traffic is not possible in a typical agent-based solution.

When you install a Windows or Linux OS, SPP and Intelligent Provisioning automatically install the AMS. HPE delivers AMS for VMware through:

HPE Support Center (Windows, Red Hat Enterprise Linux, SUSE Linux Enterprise Server, and VMware) at: <https://support.hpe.com/hpesc/public/home>

Customized HPE VMware ISO images that are released on the HPE Software Depot: <https://h20392.www2.hpe.com/portal/swdepot/index.do>

Agentless Management benefits

Agentless Management, health monitoring, and alerting begin working the moment you supply power to the server. Agentless Management runs on the iLO hardware,

independent of the host OS and processor.

The benefits of Agentless Management include:

An intelligent asynchronous device, iLO, on the physical server providing a single network target to monitor and manage the device.

SNMP implementation uses the iLO CPU and memory, thereby freeing up the resources that can be used by the application stack.

Uses out-of-band communication for increased security and stability.

A consistent GUI for configuration regardless of the host OS.

A well-documented and robust Scripting Toolkit based on industry-standard XML to deploy and configure many servers from a single configuration point such as HPE Systems Insight Manager (SIM) and HPE Insight Control. The Scripting Toolkit is designed to integrate into most custom deployment and configuration environments.

The ability to segregate the management network from the data network using the dedicated iLO network port for out-of-band communication, increased security and stability.

The option for the management network and the data network to be physically the same if you use the ProLiant servers' "shared" system networking port capability, giving you a choice for configuration.

No requirement for you to install any software to manage communication across host LAN ports or for you to enable management network protocol ports on the host OS. This significantly reduces potential security threats, especially for those platforms that might interface directly to the Internet.

Monitors Smart Array and disks attached to it from the iLO directly.

Provides logical and physical view of disk information for disks connected to Smart Array.

Customer value points include:

Faster deployments and change management

Increased security

Improved performance

Simplified operations

Preboot System Health Summary

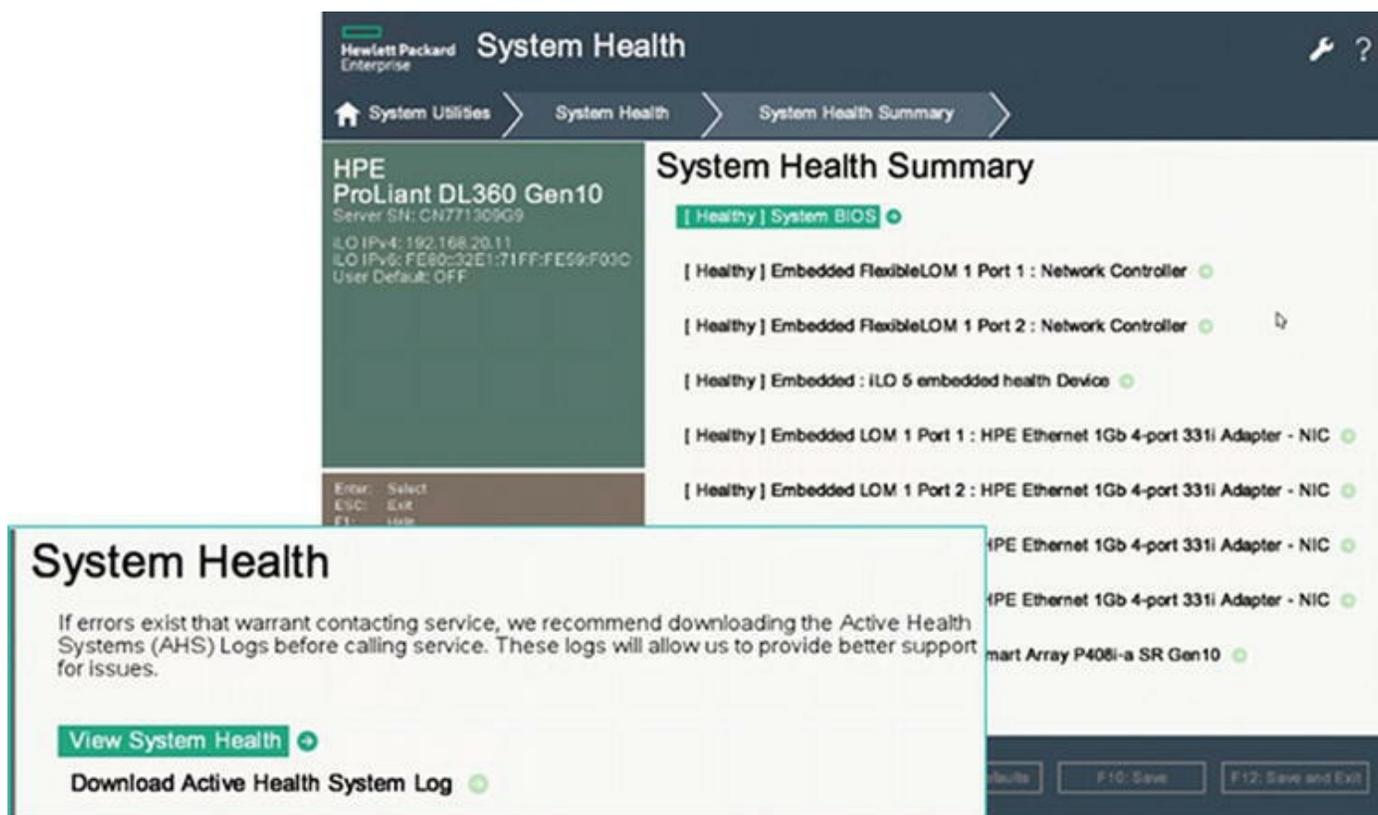


Figure 9-27 HPE ProLiant Gen10 System Health Summary

As shown in [figure 9-27](#), an overview of the Gen10 server health can be viewed from the **System Utilities** → **System Health Summary** screen. The System Health screen can be accessed by pressing **F9** during the server preboot phase. If there are errors displayed in the System Health Summary, the Active Health System Log will provide a more comprehensive report.

Active Health System

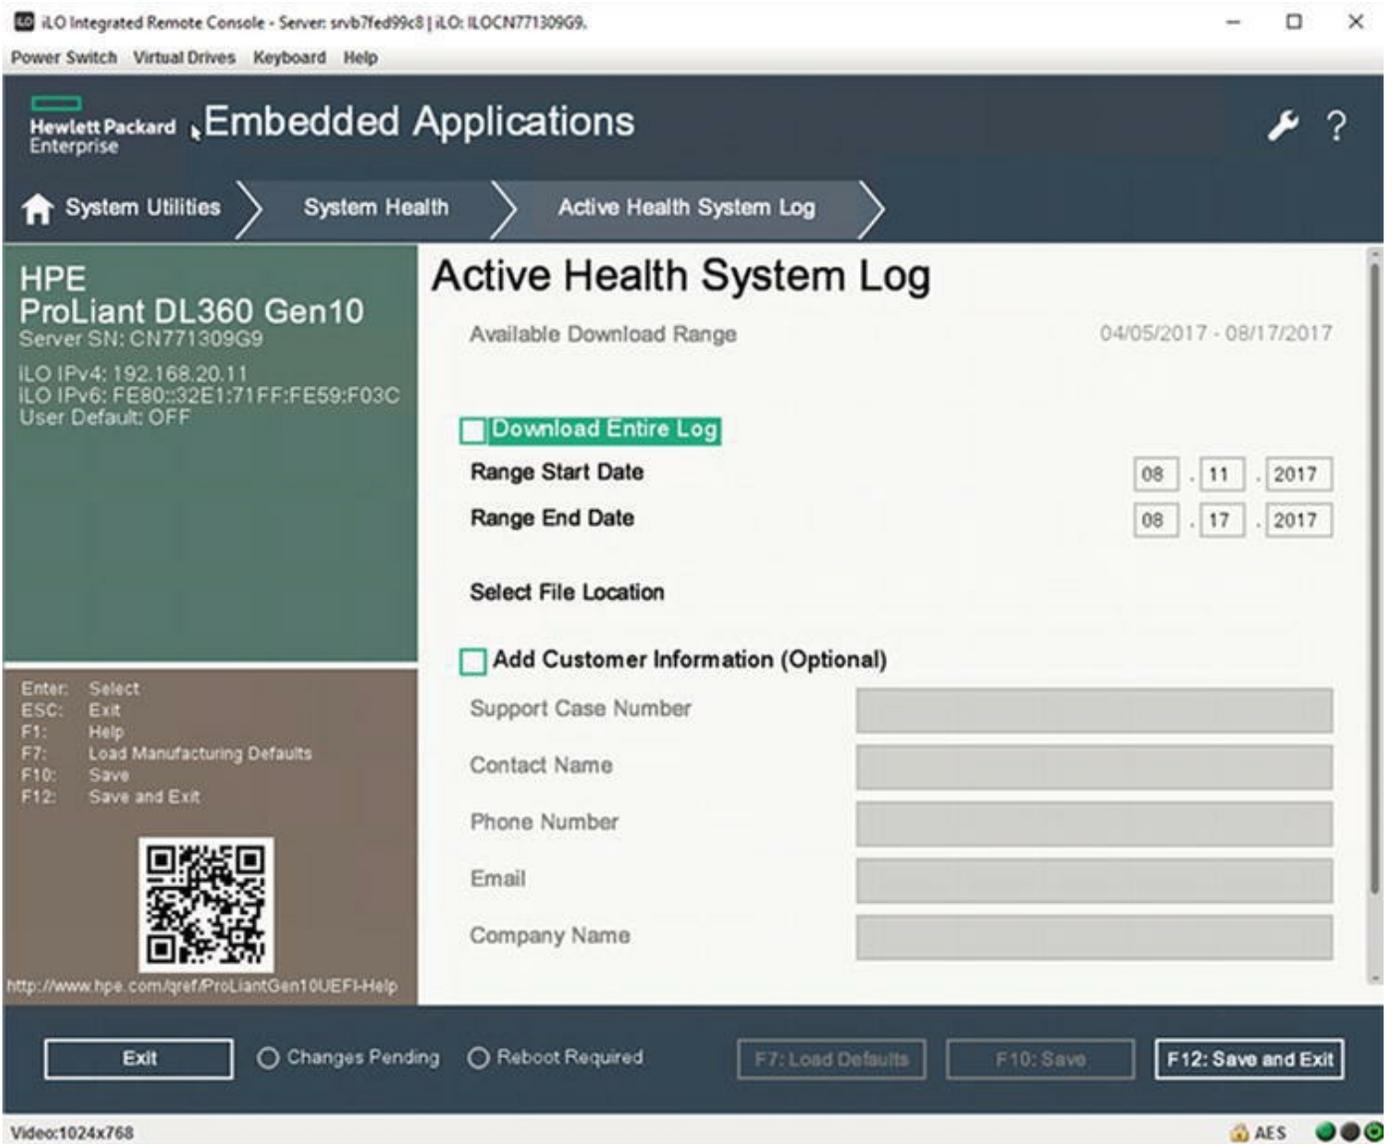


Figure 9-28 Procedure to obtain the Active Health System Log

Active Health System is the 24/7 control center for your server. With fast diagnostic data collection and the richest, most relevant data, it is the fastest way to get your system back online and keep it running optimally. The Active Health System on ProLiant Gen10 servers monitors and records changes in the server hardware and system configuration.

The Active Health System provides:

Continuous health monitoring of over 1600 system parameters

Logging of all configuration changes

Consolidated health and service alerts with precise timestamps

Agentless monitoring that does not affect application performance

It takes less than five minutes to export the Active Health System Log, as shown in [Figure 9-28](#), and send it to an HPE Support professional to help you resolve your issue.

Previously, a system issue without an obvious root cause required you to run diagnostic tools to try to isolate the cause. Although these tools often do a good job of providing the necessary information, they can only be used after the fact and often look only at subsystems individually. Circumstances occur where these tools cannot provide the information needed to isolate the root cause.

UEFI Server Hardware Diagnostics

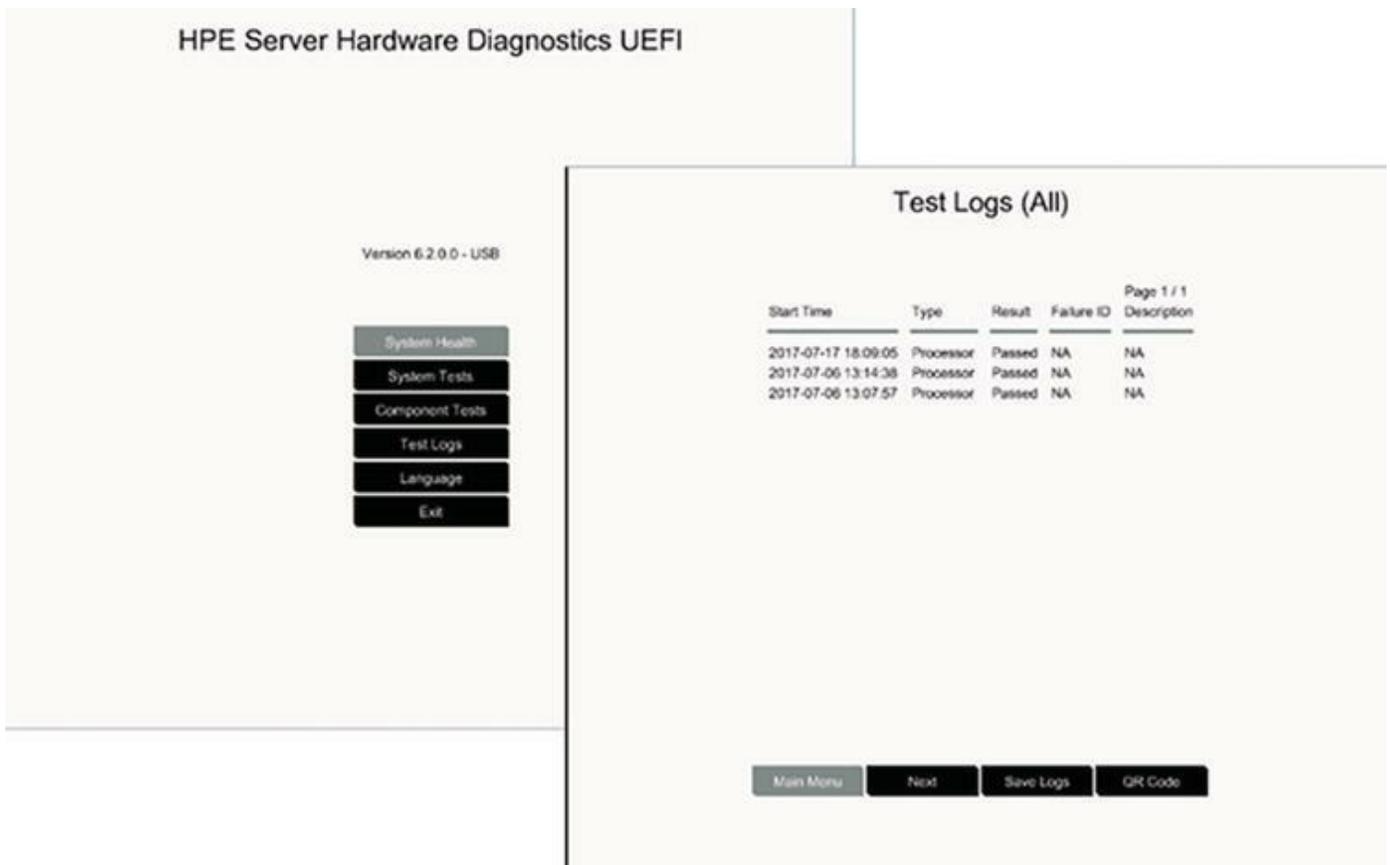


Figure 9-29 UEFI Server Hardware Diagnostics

As shown in [figure 9-29](#), UEFI Server Hardware Diagnostics can be used to test the integrity of server components such as processor, memory, hard drives, and system board. These component tests are called offline diagnostics because the server OS is not running during testing.

It is possible to view test logs for all previous hardware tests.

iLO System Information: Health Summary

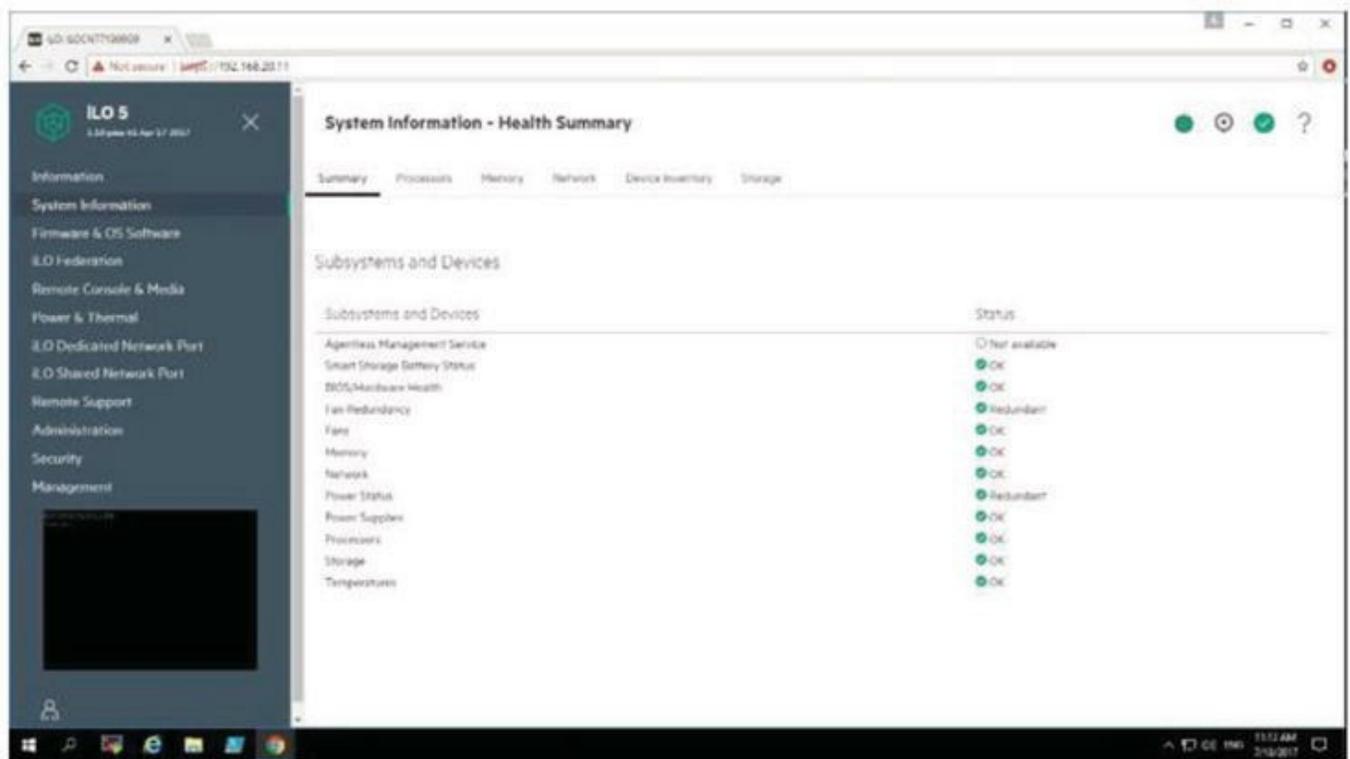


Figure 9-30 System Information Health Summary screen

The Health Summary screen displays the status of monitored subsystems and devices, as shown in [Figure 9-30](#). Depending on the server type, the information on this page varies. To view health summary information, navigate to the **Information** → **System Information** page, and click the **Summary** tab.

Summarized status information is available for:

Systems and Devices

Processors

Memory

Network

Device Inventory

Storage

Each of the categories listed above has its own dedicated tab, and selecting each tab, in turn, provides detailed information on processors, memory, and so on.

Active Health System Log

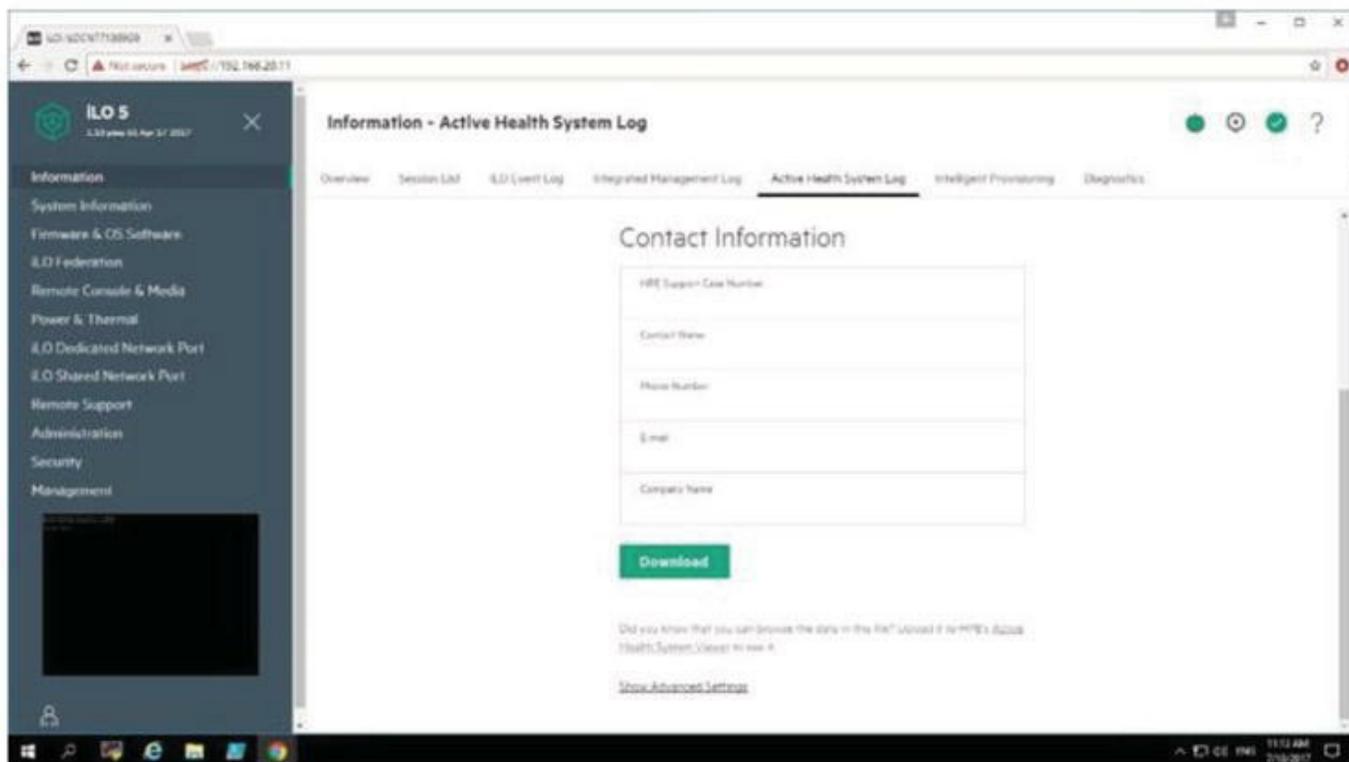


Figure 9-31 Creating Active Health Log

The Active Health System monitors and records changes in the server hardware and system configuration.

The Active Health System provides:

Continuous health monitoring of over 1600 system parameters

Logging of all configuration changes

Consolidated health and service alerts with precise timestamps

Agentless monitoring that does not affect application performance

The data collected by the Active Health System is stored in the Active Health System Log. The data is logged securely, isolated from the OS, and separate from customer data. When the Active Health System Log is full, new data overwrites the oldest data in the log. It takes less than five minutes to download the Active Health System Log and send it to a support professional to help you resolve an issue.

[Figure 9-31](#) shows the iLO 5 Active Health System Log page, from which the entire log, or the log for a specific date range, can be downloaded.

Server storage configuration utilities

Starting with Intelligent Provisioning 1.50, the SSA replaces the HPE Array Configuration Utility (ACU) with an updated design and new features and functionality for Smart Storage initiatives. You should only need to make minimal changes to existing ACU scripts such as calling the appropriate binary or executable to maintain compatibility.

SSA

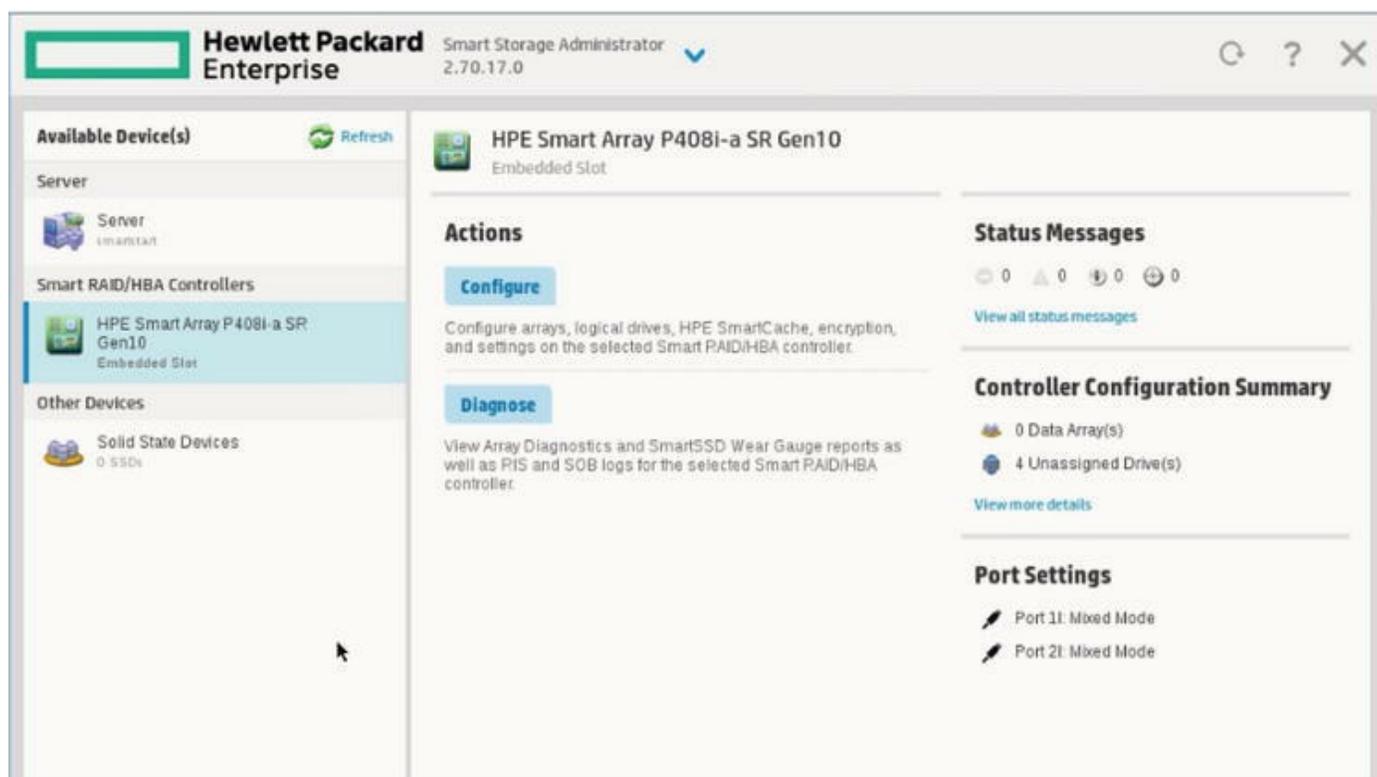


Figure 9-32 SSA

SSA is an application that is accessed via Intelligent Provisioning. As shown in [Figure 9-32](#), it helps you configure, manage, diagnose, and monitor ProLiant Smart Array controllers and HBAs. SSA is the main tool for configuring arrays on Smart Array controllers.

SSA exists in three interface formats:

GUI

CLI

Scripting

All formats provide support for standard configuration tasks. SSA also provides support for advanced configuration tasks, but some of its advanced tasks are available in only one format. The diagnostic features in SSA are also available in the stand-alone software HPE SSA Diagnostics Utility CLI.

SSA polls any drives that are present and builds an appropriate array for those drives. For example, if two drives are connected to the Smart Array card, the setup defaults to RAID 1. HPE recommends selecting this option when initially provisioning a server.

With SSA, you can also configure the cryptographic features of HPE Secure Encryption, enable HPE SSD Smart Path, and create different optimization and priorities associated with the drives and controllers.

SSA is accessible both offline and online:

Accessing SSA in the offline environment—In offline mode, you can configure or maintain detected and supported ProLiant devices, such as optional Smart Array controllers, integrated Smart Array controllers, and RAID array controllers. Some SSA CLI features are only available in the offline environment, such as setting the boot controller and boot volume.

You can launch SSA:

- With HPE Intelligent Provisioning
- During POST
- From an ISO image

Accessing SSA in the online environment—This method requires an administrator to download the SSA executables and install them. You can run SSA online after launching the host OS.



Note

If an HPE Smart Array Controller has been added or is embedded in the system, the controller defaults to a RAID configuration based on the size and number of hard drives installed.

SSA diagnostics

- Create diagnostic report
- Display diagnostic report and save for analysis by HPE
- Replaces the ACU
- SSD Wear Gauge
 - Total SSD drives with Wearout status
 - Total Solid State NVMe drives
 - Total Smart Array SSDs
 - Total Solid State Drives

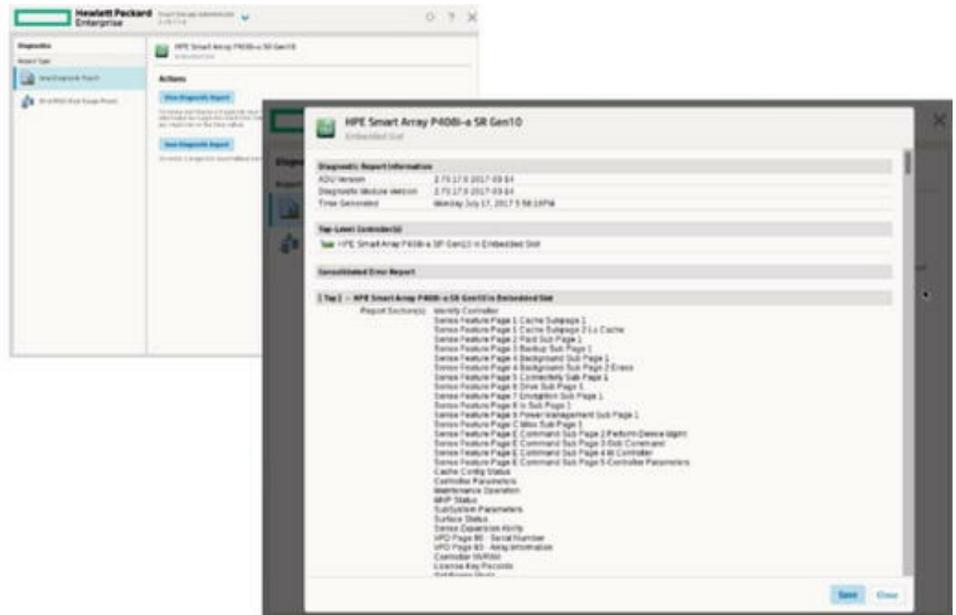


Figure 9-33 SSA Diagnostic

The HPE SSA Diagnostics feature has replaced the Array Diagnostic Utility supported by SmartStart v8.20 and earlier.

As shown in [Figure 9-33](#), HPE SSA generates the following reports and logs:

Array Diagnostic Report—This report contains information about all devices, such as array controllers, storage enclosures, drive cages, as well as logical, physical, and tape drives. For supported solid state drives, this report also contains SmartSSD Wear Gauge information.

SmartSSD Wear Gauge report—This report contains information about the current usage level and remaining expected lifetime of solid state drives attached to the system.

Serial output logs—This log details the serial output for the selected controller.

For each controller, or for all of them, you can select the following tasks:

View Diagnostic Report

Save Diagnostic Report

View SmartSSD Wear Gauge Report

Save SmartSSD Wear Gauge Report

For the view tasks, HPE SSA generates and displays the report or log. For the save tasks, HPE SSA generates a report without the graphical display.

For either task, you can save the report. In online and offline environments, HPE SSA

saves the diagnostic report to a compressed folder, which contains an XML report, a plain text report, and a viewer file so you can display and navigate the report through a web browser.

Performing a diagnostic task

Open HPE SSA.

Open the Diagnostics panel by doing one of the following:

- Select a device and click Diagnose in the quick navigation menu.
- Select an available device from the Home screen, and then click **Diagnose** under the available options.

Select a report type. For this example, use the Array Diagnostic Report selection.

Select Array Diagnostic Report.

The Actions panel for Array Diagnostic Report appears.

Each HPE SSA Diagnostics report contains a consolidated view of any error or warning conditions encountered. It also provides detailed information for every storage device, including the following:

Device status

Configuration flags

Firmware version numbers

Physical drive error logs

HPE SSA Diagnostics never collects information about the data content of logical drives. The diagnostic report does not collect or include the following:

File system types, contents, or status

Partition types, sizes, or layout

Software RAID information

OS device names or mount points

Activity: Creating logical volumes using the SSA GUI

This activity requires you to use Internet Explorer or Chrome to browse to the emulator index page. On the emulator index page select **Creating logical volumes using the GUI**.

Follow the instructions as they appear on the screen.



Note

Access the emulator by entering the following URL into your browser;

<http://hpe.azureedge.net/gen10emulators/index.html>

Activity debrief

What RAID level was selected for the new logical drive?

Which Smart Array Controller did you select?

Activity: Creating logical volumes using the RESTful Interface Tool

This activity requires you to use Internet Explorer or Chrome to browse to the emulator index page. On the emulator index page select **Creating a RAID1 logical volume using the RESTful Interface Tool**.

Follow the instructions as they appear on the screen.



Note

Access the emulator by entering the following URL into your browser:

<http://hpe.azureedge.net/gen10emulators/index.html>

Activity debrief

What drive numbers did you select for your new logical drive?

Why did your new logical drive not appear the first time you looked for it with the `smartarray --controller=1 --logicaldrives` command?

Learning check

Write a summary of the key concepts presented in this chapter.

Summary

To prepare a ProLiant system for a deployment, you need to:

- Install hardware.
- Configure iLO.
- Validate environmental standards.
- Consider storage networking.

The boot process is the initial set of operations a server performs after electrical power is switched on. Configuration begins immediately following the boot process, setting up relationships between functional units of the system.

HPE server management offerings that help you prepare a system for deployment are considered on-system management tools.

HPE uses the HPE SSA Configuration Utility to configure an array with a Smart Array Controller.

10 HPE On-premises and On-cloud Management

WHAT IS IN THIS CHAPTER FOR YOU?

After completing this chapter, you should be able to:

Describe HPE server management.

Explain the HPE OneView approach to infrastructure lifecycle management.

Describe the architecture of HPE OneView.

Explain how HPE OneView integrates with partner products such as:

- ✓ HPE OneView for VMware vCenter
- ✓ HPE OneView for Microsoft System Center
- ✓ HPE OneView for Red Hat Enterprise Virtualization (RHEV)

Explain how to perform remote management with HPE on-premises and on-cloud management tools.

- ✓ HPE iLO Advanced
 - ✓ HPE Insight Online
 - ✓ HPE Insight Remote Support
-

Prelearning check

You are visiting a customer who tells you that they have recruited a new assistant system administrator who is tasked with streamlining their infrastructure management. They currently rely on Insight Manager, but they want to use scripting to perform server deployment. What should you advise?

HPE server management

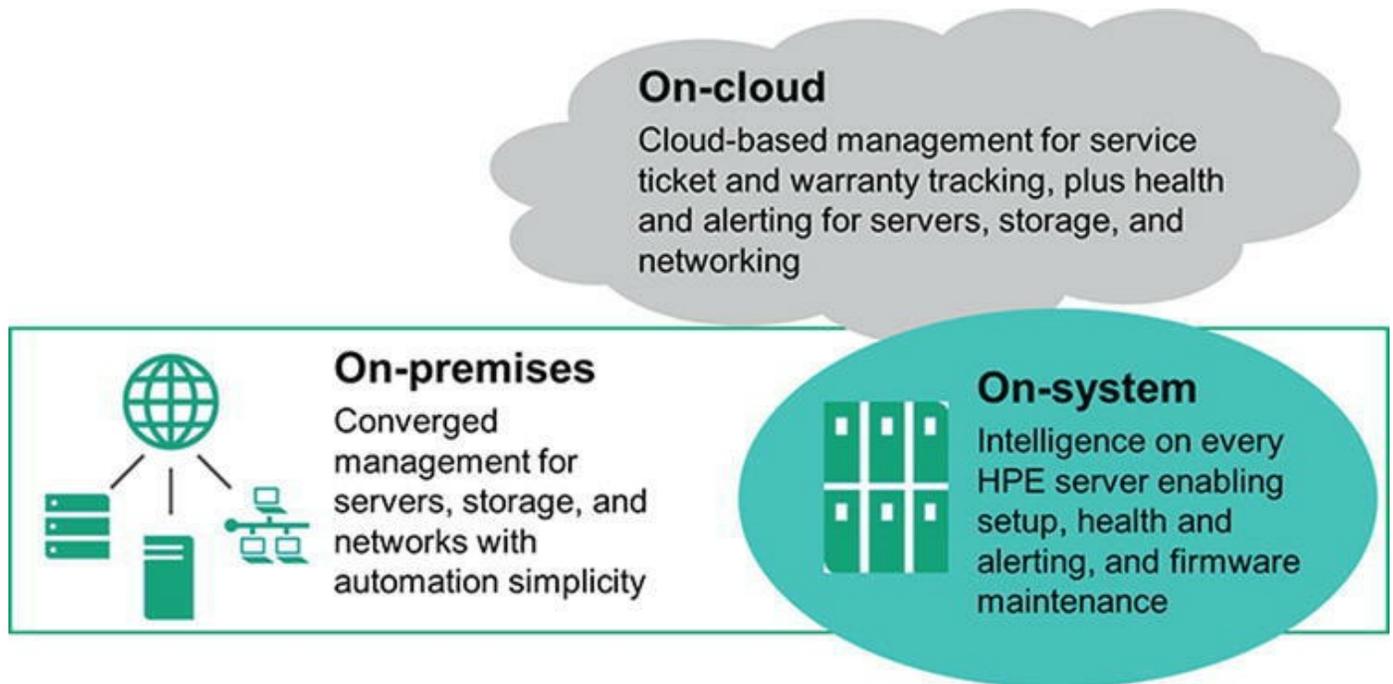


Figure 10-1 HPE management offerings

HPE provides a comprehensive set of management offerings to help companies meet their management requirements at every stage of the server lifecycle. As shown in [Figure 10-1](#), these offerings belong to one of three categories of infrastructure lifecycle management solutions that allow IT administrators to design, monitor, and manage data systems.

On-system—Intelligence on every HPE server enabling setup, health and alerting, and firmware maintenance, HPE on-system management also includes:

- HPE Smart Update Manager (SUM) provides systematic and organized server updates via HPE Service Pack for ProLiant (SPP).
- HPE Intelligent Provisioning assists in the initial configuration and operating system deployment for ProLiant servers.
- Unified Extensible Firmware Interface (UEFI) is the BIOS mode for ProLiant Gen10 servers that aids with server configurations.
- HPE RESTful Interface Tool is a command line interface (CLI) scripting tool that leverages the HPE RESTful application programming interface (API) for secure, remote server configurations at scale.

On-premises—Converged management for servers, storage, and networks with

automation simplicity:

- HPE OneView
- HPE iLO Advanced

On-cloud—Cloud-based management for service ticket and warranty tracking, plus health and alerting for servers, storage, and networking:

- HPE Insight Online
- HPE Insight Remote Support

On-premises management with HPE OneView

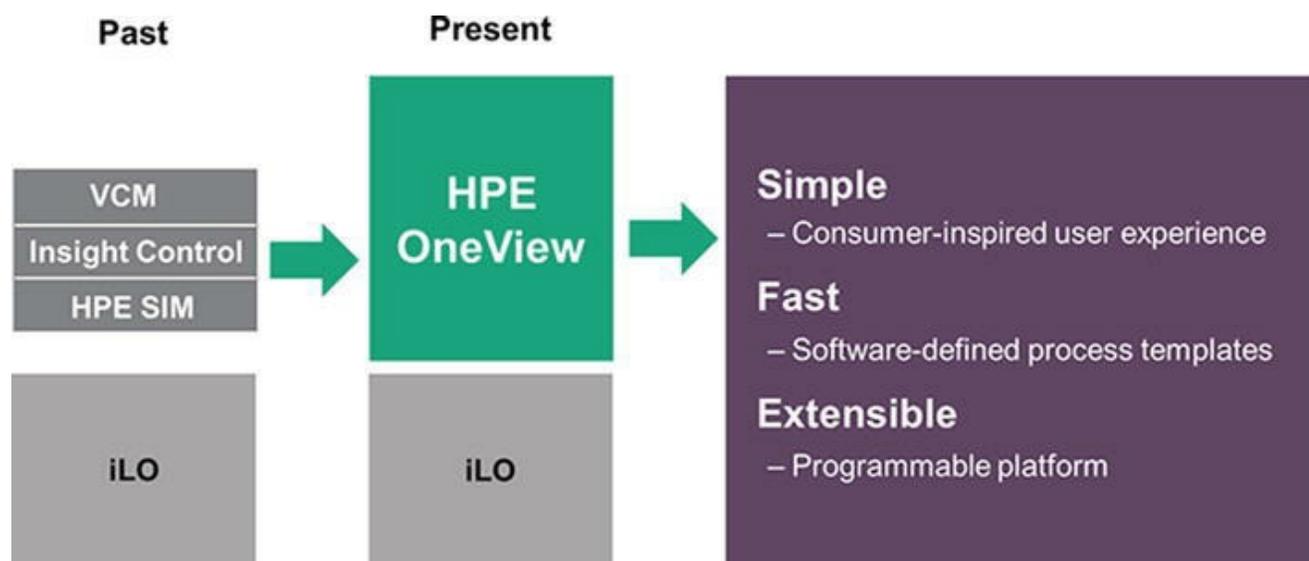


Figure 10-2 On-premises management with HPE OneView

HPE OneView provides a resource-oriented solution that focuses on the entire hardware lifecycle from initial configuration to ongoing monitoring and maintenance. HPE OneView is purpose-built to support scenarios such as:

Deploying bare-metal servers

Deploying hypervisor clusters from bare metal

Performing ongoing hardware maintenance

Responding to alerts and outages

As shown in [Figure 10-2](#), HPE OneView supersedes Virtual Connect Manager (VCM), Insight Control, and HPE Systems Insight Manager (SIM), and integrates with iLO to provide a simple, fast, and extensible management platform.

HPE OneView uses templates that automate the configuration and propagation of server, storage, and networking profiles. Its underlying infrastructure components—including networking, power management, and servers—support virtualization, cloud computing, Big Data, and mixed computing environments. This converged, software-defined, and automated platform reduces operating expenses (OPEX), improves agility, and frees up resources for new business initiatives.

This approach provides the following benefits:

Software-defined approaches to systems management create reliable repeatability to help prevent unplanned outages caused by human error or device failure.

- Profiles and groups capture best practices and policies to help increase productivity and enable compliance and consistency.
- Manage the infrastructure programmatically by using APIs built on industry standards such as Representational State Transfer (REST). These APIs are easily accessible from any programming language, and software development kits are provided for interfaces, PowerShell, and Python scripts.

Automation can streamline the delivery of IT services and speed transition to Infrastructure-as-a-Service (IaaS) and hybrid cloud delivery.

- As an intelligent hub, HPE OneView provides closed-loop automation with consistent APIs, data modeling, and state change message bus (SCMB).
- HPE OneView also supports lights-out automation.

Convergence reduces the number of tools required to learn, manage, deploy, and integrate the infrastructure.

- A single, open platform supports multiple generations of HPE ProLiant BL servers, HPE 3PAR StoreServ storage, HPE ConvergedSystem, and HPE BladeSystem solutions, as well as HPE and Cisco Top of Rack (ToR) switches.
- Microsoft System Center integration adds a single integrated System Center Virtual Machine Manager Fabric Management/Storage Add-in with support for partial domains, Fibre Channel over Ethernet (FCoE) dual hop, server profile templates, and logical enclosures.
- VMware vCenter Server integration adds server profile-based deployment and automated HPE StoreOnce deployment for secure backup and recovery. You can automate control of compute, storage, and networking resources using VMware vCenter or Microsoft System Center without detailed knowledge of every device. Tasks, processes, and projects are accomplished faster and with more consistency than the older patchwork approaches to management.
- RHEV gives administrators insight into and control over their HPE infrastructure while supporting their Red Hat virtualized environment from a single screen.

Lifecycle management automates day-to-day responsibilities by simplifying time-consuming tasks, which leads to increased productivity and reduced operational costs. This capability includes:

- Server profile templates
- Firmware and driver updates
- Logical tracking and linking of resources

Designed for simplicity

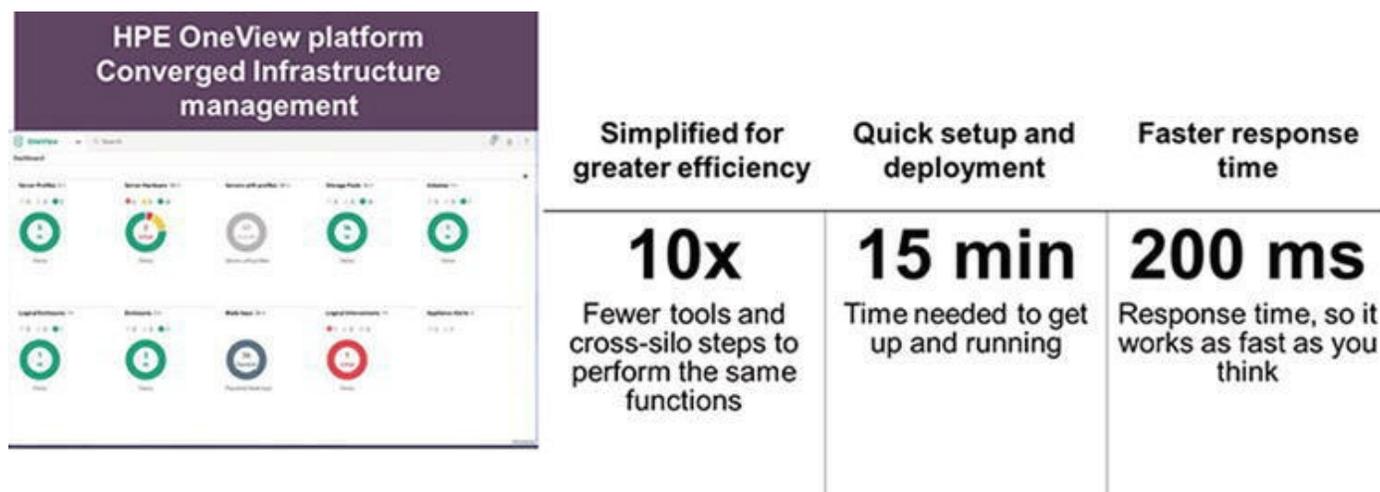


Figure 10-3 HPE OneView designed for simplicity

The software-based approach to lifecycle management in HPE OneView includes a fully programmable interface to create customized workflows, scripts, and configuration profiles for push-button builds that instantly deliver resources without mistakes or variations. Workflow templates capture best practices and policies to increase productivity and enable compliance and consistency. Built-in intelligence automates common BladeSystem management tasks usually performed by users, and connections to HPE SUM automate firmware and system software maintenance.

The HPE OneView GUI, shown on the left side of [Figure 10-3](#), reduces the number of tools needed for everyday management of server, storage, and network resources. One tool can perform numerous tasks within a single operating model that provides a consistent way of viewing resources and results.

HPE OneView increases productivity for the IT team by facilitating collaboration, removing friction, and accelerating time to value. Administrators responsible for different areas (including servers, storage, networks, and virtualization) are authorized to use the same tool to manage resources. This allows them to collaborate and make authorized decisions without stepping on each other's responsibilities. For example, virtualization administrators can automate control of HPE compute, storage, and networking resources without detailed knowledge of each device.

The HPE OneView appliance provides several software-defined resources, such as groups and server profiles, to enable an administrator to capture the best practices of an organization's experts across a variety of disciplines. An administrator can reduce cross-silo miscommunication by defining server profiles, networking objects, and other

resources.

HPE OneView has also been optimized for productivity with a response time target of 200 milliseconds for all user interface interactions. Tasks that take longer than 200 milliseconds to process are labeled **asynchronous** tasks, meaning they can finish independently in the background. This allows the administrator to move on to other activities.

HPE OneView: Your infrastructure automation engine

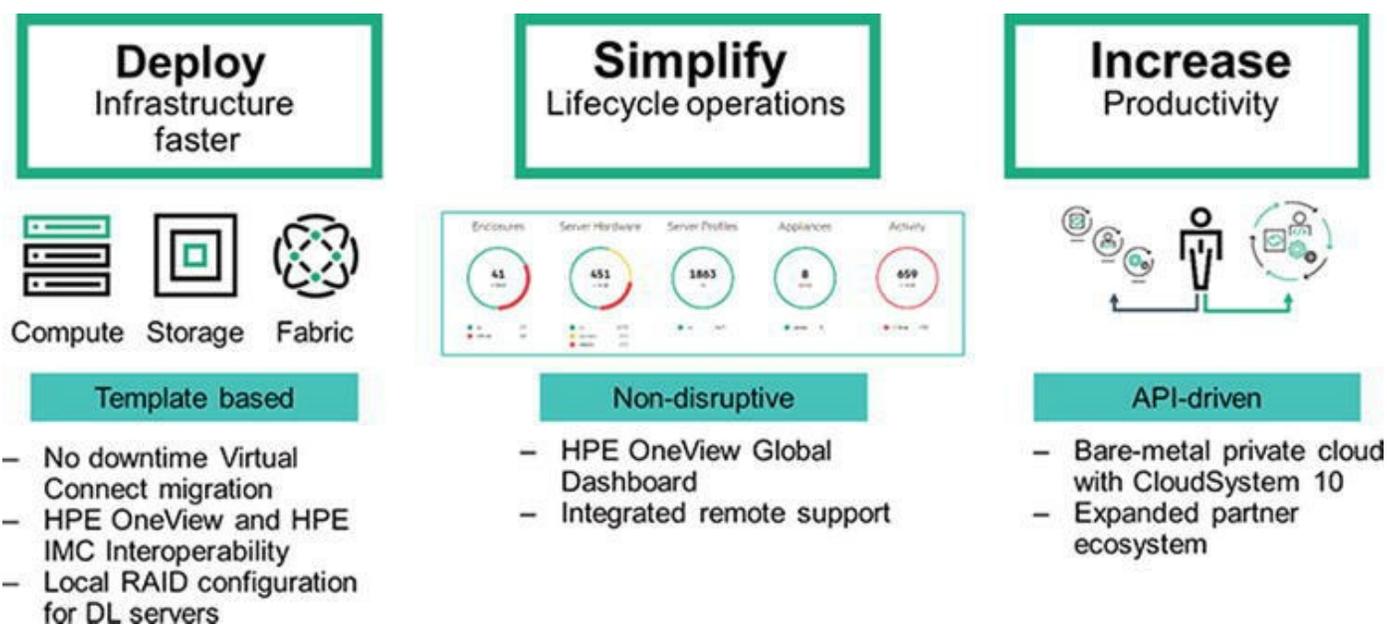


Figure 10-4 HPE OneView is an infrastructure automation engine

HPE OneView is your infrastructure automation engine to simplify operations, increasing the speed of IT delivery for new applications and services. Through software-defined intelligence, HPE OneView brings a new level of automation to infrastructure management by taking a template-driven approach to provisioning, updating, and integrating compute, storage, and networking infrastructure. Designed with a modern, standard-based API and supported by a large and growing partner ecosystem, HPE OneView also makes it easy to integrate powerful infrastructure automation into existing IT tools and processes.

The benefits of this approach are shown in [Figure 10-4](#):

Deploy infrastructure faster—Template-based automation enables IT generalists to rapidly and reliably provision resources in response to application owner requirements.

Simplify lifecycle operations—Agentless monitoring, online firmware updates, and a new Global Dashboard deliver streamlined lifecycle operations at scale.

Increase productivity—The unified API allows developers and independent software vendors (ISVs) to harness the power of HPE OneView and unify infrastructure automation with application and IT service delivery.

HPE OneView innovations provide you the industry's best infrastructure management experience, simplifying operations for HPE BladeSystem, HPE ProLiant servers, HPE StoreServ 3PAR storage, StoreVirtual, HPE Networking, HPE ConvergedSystem, and

powers management for the industry's first composable infrastructure, HPE Synergy. HPE OneView is integrated into the Synergy Composer for hassle-free infrastructure management and frictionless scaling.

HPE OneView architecture

The HPE OneView modern management architecture is designed to accelerate IT operations for managing both logical and physical resources. Logical resources are items such as network configurations, server profiles, and connections. Physical resources are tangible items such as server hardware, interconnects, and enclosures. This advanced architecture connects the resources with a common, domain-specific representation of the servers, networks, and storage. It also models the associations and interdependencies of the resources.

The HPE OneView resource-oriented architecture provides a uniform REST interface. Every resource has one Uniform Resource Identifier (URI) and represents a physical device or logical construct.

HPE OneView resource managers

✓ **Encl1, bay 2** | Overview ▾ | 🔍

Hardware >

Server name	<i>not set</i>
State	Profile Applied
Server profile	● <u>iscsi-bay-1</u>
Server power	Off
Model	ProLiant BL660c Gen9
Server hardware type	<u>BL660c Gen9 1</u>
Serial number	SGH101X7RN
License	HPE OneView Advanced w/o iLO
iLO address	<u>172.18.6.2</u> <u>and 1 additional</u>
Location	<u>Rack-221, Encl1, bay 2</u>
Resource type	----

Figure 10-5 Physical server resource manager

HPE OneView is designed as a set of cooperating components referred to as **resource managers** that implement the logical structures corresponding to the physical resources in the data center. HPE OneView resource managers run on the HPE OneView management appliance.

Resource managers focus on a specific type of resource, including servers, storage, and networking and provide the REST APIs for those resources, as well as publish state change messages to the message bus when their resources are changed in any way. Resource managers detect state changes both in response to user-initiated changes and also by monitoring the actual environment using device-level APIs and protocols such as Simple Network Management Protocol (SNMP). Both types of changes are

consistently reflected in the REST API and state change messages that are published to the SCMB to notify interested parties such as partner integrations or higher-level automation.

[Figure 10-5](#) shows the HPE OneView Hardware screen for a ProLiant Blade server and shows the types of entities managed by the physical server resource manager.

OneView has several resource managers that manage dedicated physical and logical components and interact with other components to make OneView work:

Environmental Resource Manager—Is responsible for presenting the management capabilities for power and cooling components of the physical systems. Through the Environmental Resource Manager, an administrator or application using the REST API can retrieve the current power status of a resource such as a server blade and then subsequently power the server on or off as needed.

Storage Resource Manager—Enables you to easily set up and consume storage resources. It sends information to the dashboard, which displays storage pools and volumes status. Communication issues to the storage system are seen as critical for all storage managed resources. Storage pool status is updated based on provisioning status of the common provisioning group (CPG) created in the HPE 3PAR StoreServ and StoreVirtual system.

Physical Server Resource Manager—Uses Core Foundation Services to access the physical components of the data center that are responsible for managing the rack servers and the server blades in enclosures.

Connectivity Resource Manager—Uses Core Foundation Services to access the physical components of the data center that it is responsible for managing—the edge switches for LAN and SAN connectivity that correspond to Virtual Connect (VC) modules.

REST API

HPE OneView administrators, application plug-ins, and converged systems interact with HPE OneView through either the GUI or through the REST API. The REST API provides a simple, stateless, and scalable approach, so users can easily integrate, automate, and customize on their own. Developers are also interested in the REST API because they are common to two-thirds of the top web environments.

You can use REST APIs to manipulate resources. A single, consistent resource model is embodied in an HTML5 user interface and the REST API. The REST architecture generally runs over HTTP, although other transports can be used. What the HPE

OneView user interface allows you to do graphically, the REST API enables you to do programmatically. For example, you can use a scripting language such as Microsoft PowerShell to perform tasks by using REST API calls that you might otherwise complete through the web-based UI.

REST support in HPE OneView is implemented as a web service. This web service allows basic create, read, update, and delete (CRUD) operations to be performed on resources using HTTP POST, GET, PUT, and DELETE messages.

SCMB

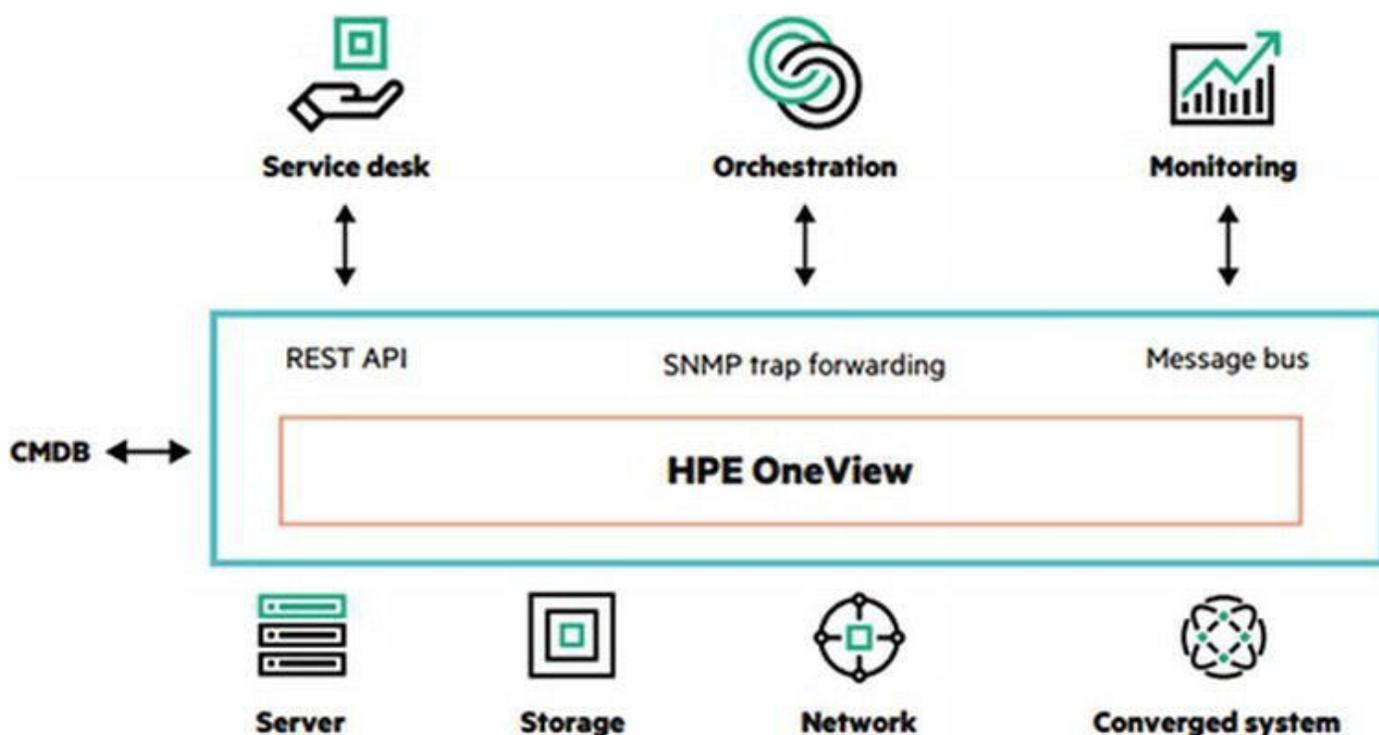


Figure 10-6 SCMB

The SCMB is a high-performance client request and server response mechanism used for retrieving status information about the underlying physical resources and requesting state changes. As shown in [Figure 10-6](#), the SCMB provides a communication channel between application processes and the underlying systems that it manages. It uses asynchronous messaging to notify subscribers of changes to both logical and physical resources. For example, you can program applications to receive notifications when new server hardware is added to the managed environment or when the health status of physical resources changes—without having to poll the appliance continuously for status using the REST API.

The integrated SCMB messaging platform in HPE OneView enables dynamic and responsive integrations with solutions such as HPE OneView for VMware vCenter and OpenStack for HPE Helion CloudSystem. It integrates with many commonly used tools, applications, and products, including service desk, reporting, monitoring, and configuration management database (CMDB) tools.



Note

A CMDB is a repository that acts as a data warehouse for an IT organization. It contains information describing IT assets such as software and hardware products and personnel.

Consumer-inspired user experience

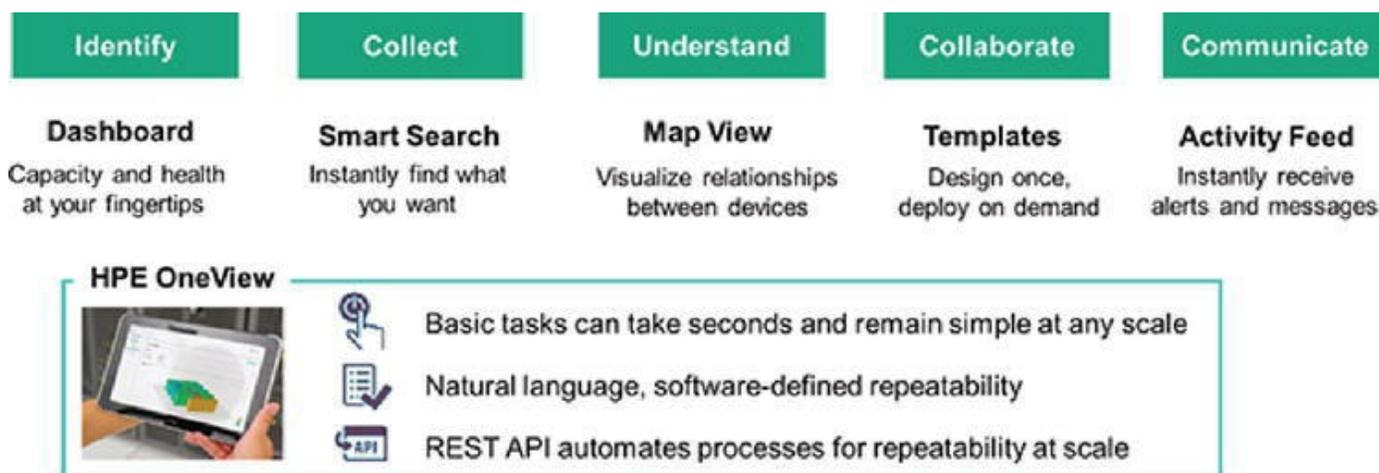


Figure 10-7 Five key data center tasks to which HPE OneView corresponds

HPE OneView simplifies common tasks to accelerate IT processes. [Figure 10-7](#) shows five key data center tasks: identify, collect, understand, collaborate, and communicate. HPE OneView key features correspond to those tasks:

Identify—Use the HPE OneView **Dashboard** to identify what needs to be changed, adjusted, or completed. Using the Dashboard, you can assess what is happening across the entire data center in seconds.

Collect—Use **Smart Search** to collect information and find what you want instantly. Smart Search is built into almost every task, so you have immediate access to the event device or task information you need.

Understand—**Map View** helps you visualize relationships between devices to understand device status and other dependencies. In a converged environment,

understanding relationships between devices and connections is critical, especially for change management or troubleshooting. Map View shows you how things are connected so you can quickly find, triage, and fix problems.

Collaborate—Traditionally, IT tools were built for silos and did nothing to foster collaboration, which is a critical requirement for convergence. Using **Templates**, IT teams can design once and deploy on demand. Teams can work together to create profiles of configurations and procedures. Any configuration or process can now be designed by the experts and captured in software for later use by everyone.

Communicate—Most communications between teams happen outside of management tools. **Activity Feed** sends alerts and messages about progress, completion, or failure. Activity Feed functions similarly to Twitter in terms of collaboration and communication. An administrator can add notations to events and assign them to an appropriate user. Activity Feed is a single place for IT administrators to communicate, work together, and stay on the same page with alerts, jobs, and projects in the data center.

OneView Dashboard

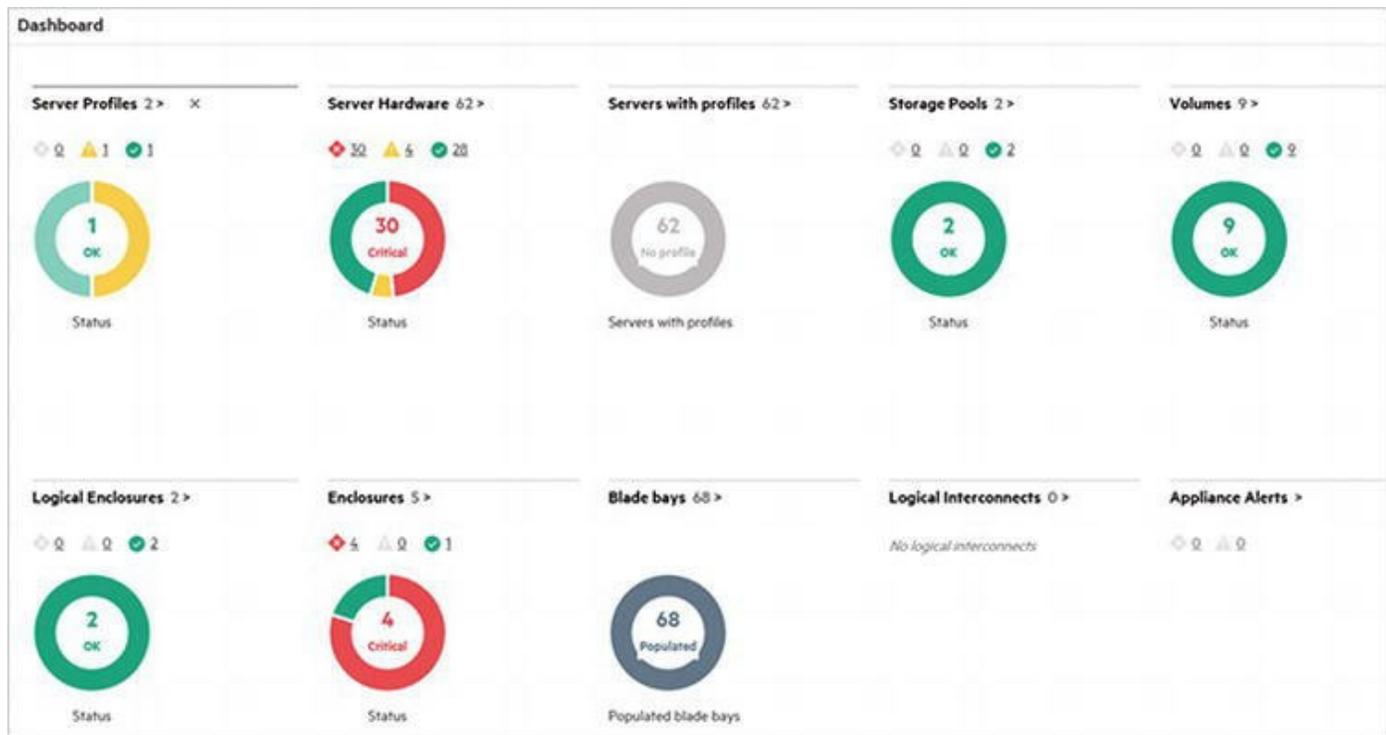


Figure 10-8 HPE OneView Dashboard

When you log in to the OneView appliance, you can view the Dashboard, which

provides an overview of system activity and alerts, as shown in [Figure 10-8](#). The Dashboard provides an at-a-glance visual health summary of the appliance resources you are authorized to view, such as:

Server profiles

Server hardware

Enclosures

Logical interconnects

Volumes

Storage pools

SAN management

The status of each resource is indicated by an icon:  OK,  Warning, or  Critical. You can link to the resource screens in the UI for more information by clicking these status icons.

Dashboard chart colors help you to quickly interpret the reported data:

Green—Indicates a healthy status

Yellow—Means that an event has occurred that might require attention

Red—Signifies that a critical condition requires immediate attention

Blue—Specifies the percentage of resource instances that match the data being measured (with a solid blue graph indicating 100%)

Light gray—Shows the remainder of resource instances that do not match the data being measured (used in combination with blue to total 100%)

Dark gray—Identifies which resource instances are not reporting OK, Warning, or Critical (with a status of Disabled or Unknown)

Map View

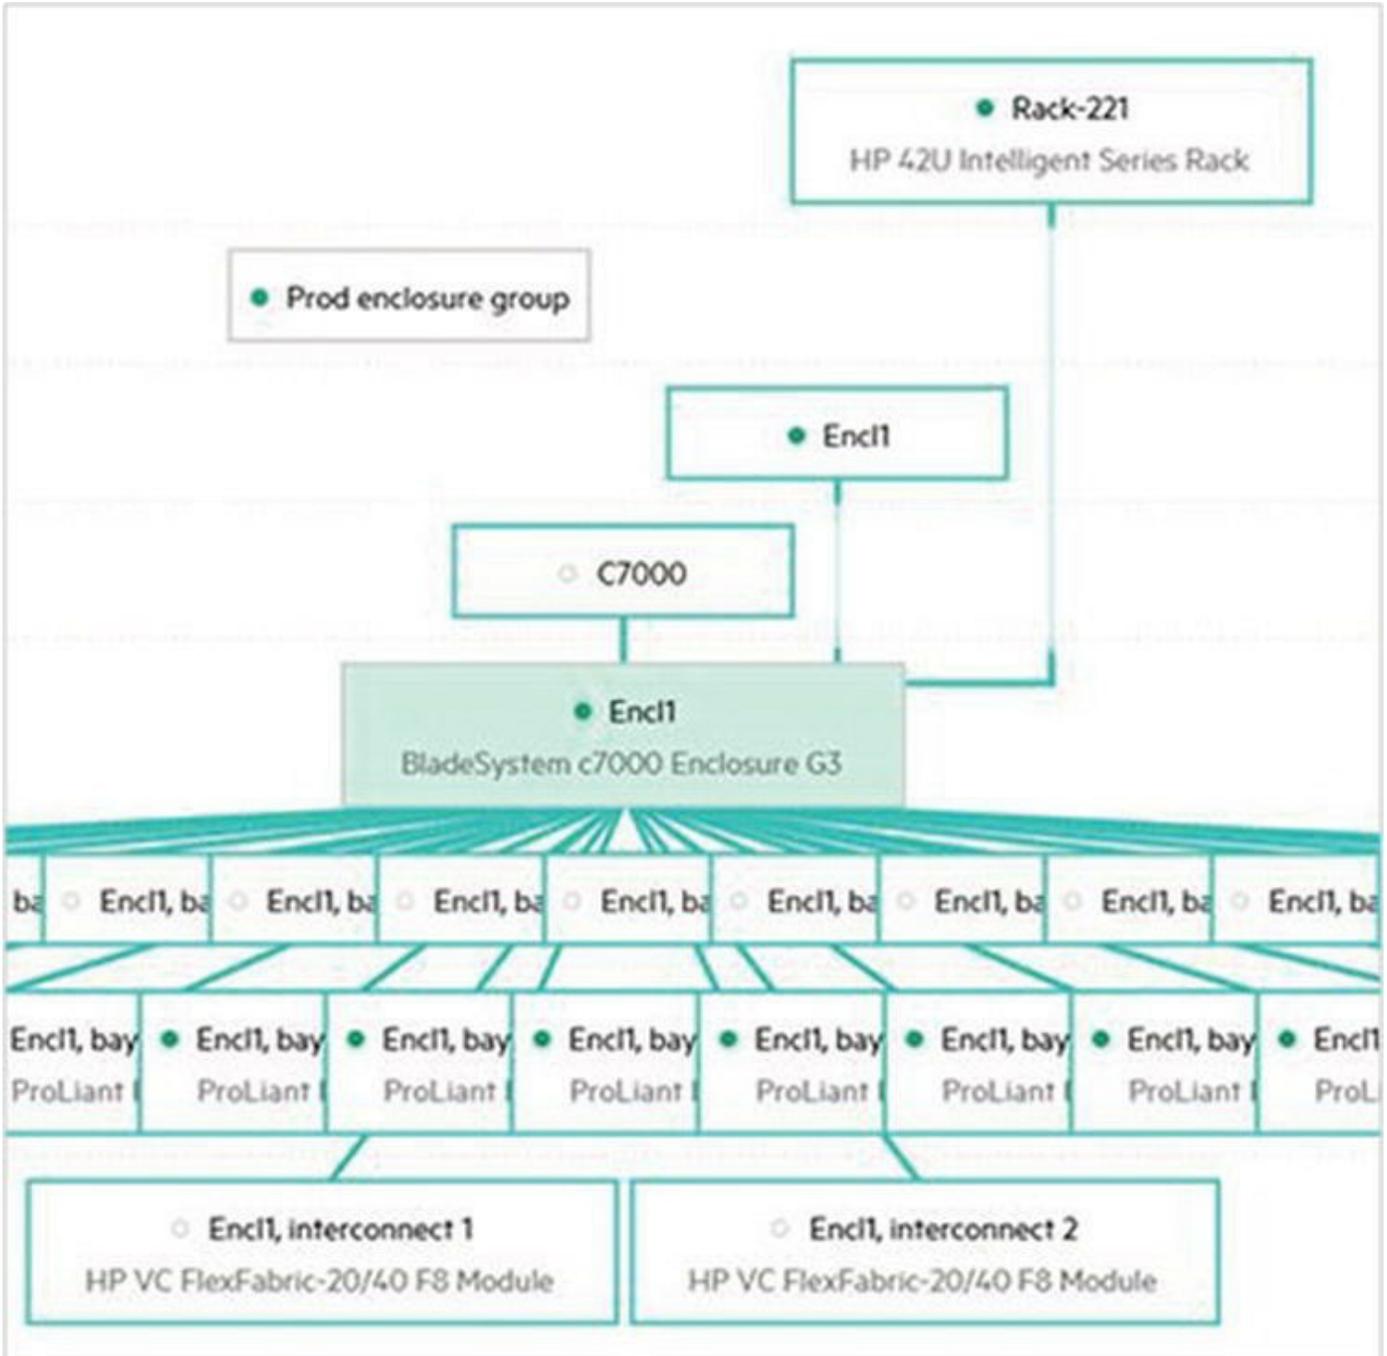


Figure 10-9 HPE OneView Map View

Map View, shown in [Figure 10-9](#), enables you to examine the configuration of and understand the relationships between logical and physical resources in the data center. This view gives you immediate visibility into resources, from individual Ethernet and Fibre Channel networks up to the enclosure, rack, and top-level physical data center. This view is designed to be highly interactive and useful even at scale.



Note

The Map View is resource centric. Everything above the selected resource is an ancestor. Everything below the selected resource is a descendant. A line connecting boxes indicates a direct relationship, such as server blades in an enclosure. Other items can be indirectly related to the resource, such as logical interconnect groups and server profiles.

Activity screen

Name	Resource	Date	State	Owner
Add	ThreePAR82200-6753 Storage Systems	2/22/16 6:37:17 am 36 minutes ago	Completed 4s	Admin
Subtasks 2 All statuses All states - 0 - 0 2				
Add	EST_CPG1	2/22/16 6:37:20 am	Completed 1s	
Add	EST_CPG2	2/22/16 6:37:21 am	Completed 1s	
Connect	ThreePAR82200-6753 Storage Systems	2/22/16 6:36:01 am 30 minutes ago	Completed 4s	Administrator

Figure 10-10 HPE OneView Activity screen

The Activity screen provides a log of health and status notifications, as shown in [Figure 10-10](#). The HPE OneView appliance verifies the current activity of resources in the environment, and it posts alerts to the Activity screen and to the associated resource screens for you to review.

The Activity screen also serves as a database of all tasks that have been run—either synchronously or asynchronously—and initiated by the user or system. It is similar to an audit log, but provides more detail and is easily accessed from the UI.

Introduction to server profiles

A server profile is the configuration for a server instance. Server profiles capture the entire server configuration in one place, enabling you to replicate new server profiles consistently and to modify them rapidly to reflect changes in the data center. In HPE OneView, a server profile enables you to implement key aspects of a ProLiant server blade or Synergy compute module configuration in one place, including:

Connectivity settings for Ethernet networks, network sets, and Fibre Channel and FCoE networks

Firmware levels

BIOS and UEFI boot settings

iLO settings

RAID configuration

Storage volume assignment, both boot disk and shared storage

HPE StoreVirtual and HPE StoreServ volumes

Network edge LAN and SAN connectivity

Primary and secondary PXE boot and boot from SAN

Virtual, physical, and user-defined assignment of media access control (MAC) addresses, World Wide Names (WWNs), and server serial numbers



Note

There are fewer server profile features available for ProLiant rack servers. For example, you cannot manage network connections.

Server profiles are one of the features that enable you to provision converged infrastructure hardware quickly and consistently according to best practices. Server profiles enable you to specify a server configuration before it is deployed so that you can quickly bring a new server under management when the hardware is actually installed. For example, you can create a server profile that is not assigned to a particular server but specifies all the configuration aspects to use for a server hardware type. After the server is installed in an enclosure bay, you can:

Directly assign the server profile to the enclosure bay.

Use the server profile template function.

Server profile templates

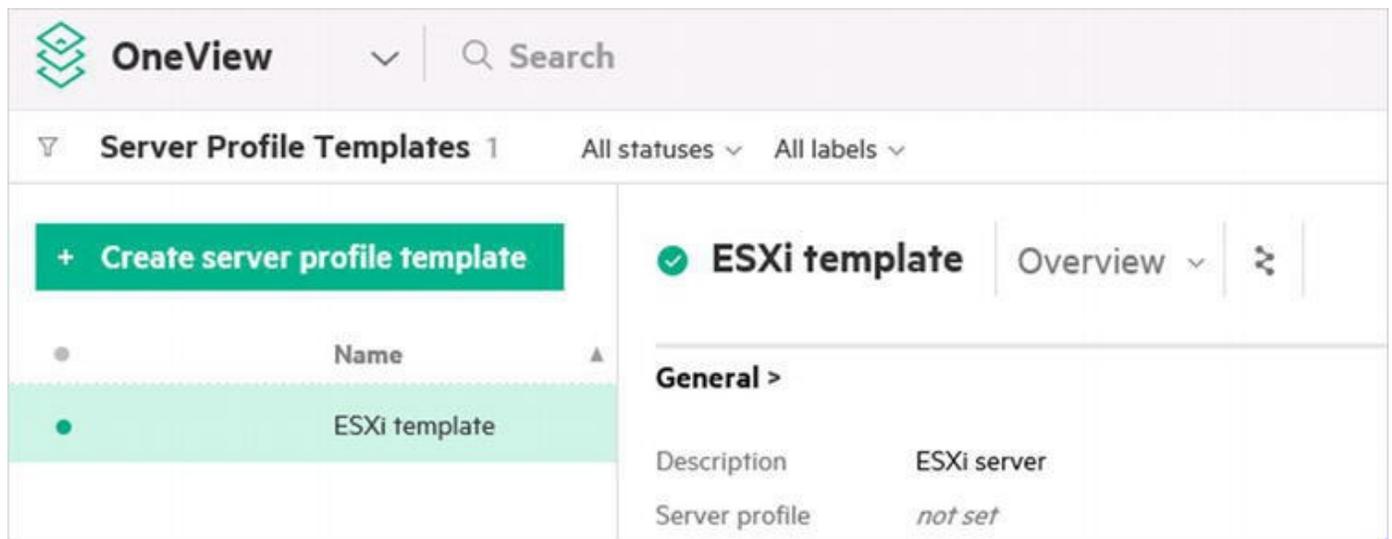


Figure 10-11 Create a server profile template in OneView

Server profile templates, an example of which is shown in [Figure 10-11](#), provide automated change management across multiple systems. You make the update once at the template level and automate updates to configurations or apply new system software baselines. Server profile templates are at the center of software-defined policies and solutions. HPE OneView provides a quick glance of server connections, SAN storage, direct-attached storage (DAS), and BIOS or UEFI compliance with the profile.

With HPE OneView, administrators can quickly drill down into dashboard panels to identify issues or troubleshoot problems. By clicking the connection panel, administrators can easily move from the high-level status down to the connections summary. You can expand individual connections to reveal complete connection configuration details. Conversely, you can collapse these details to hide them.

Server profile templates help to monitor, flag, and update server profiles in HPE OneView. A server profile template serves as a structural reference when creating a server profile. All of the configuration constructs of a server profile are present in the server profile template. This template type defines the centralized source for the configuration of firmware, connections, local storage, SAN storage, boot, BIOS, profile affinity, and hides unused FlexNICs.

Server profile mobility

The HPE OneView server profile mobility feature defines configurations once, in minutes, and then provisions or updates the configuration many times—consistently and reliably with no repetitive tasks—across compute, storage, and networking resources. This way, profile mobility is not limited to migrations across the same server hardware

type and enclosure groups. HPE OneView provides profile mobility across different:

Adapters

Generations

Server blade models

The HPE OneView appliance monitors both the server profile and server profile template. It compares both elements and ensures the server profile matches the configuration of its parent server profile template.

Adding enclosures to HPE OneView

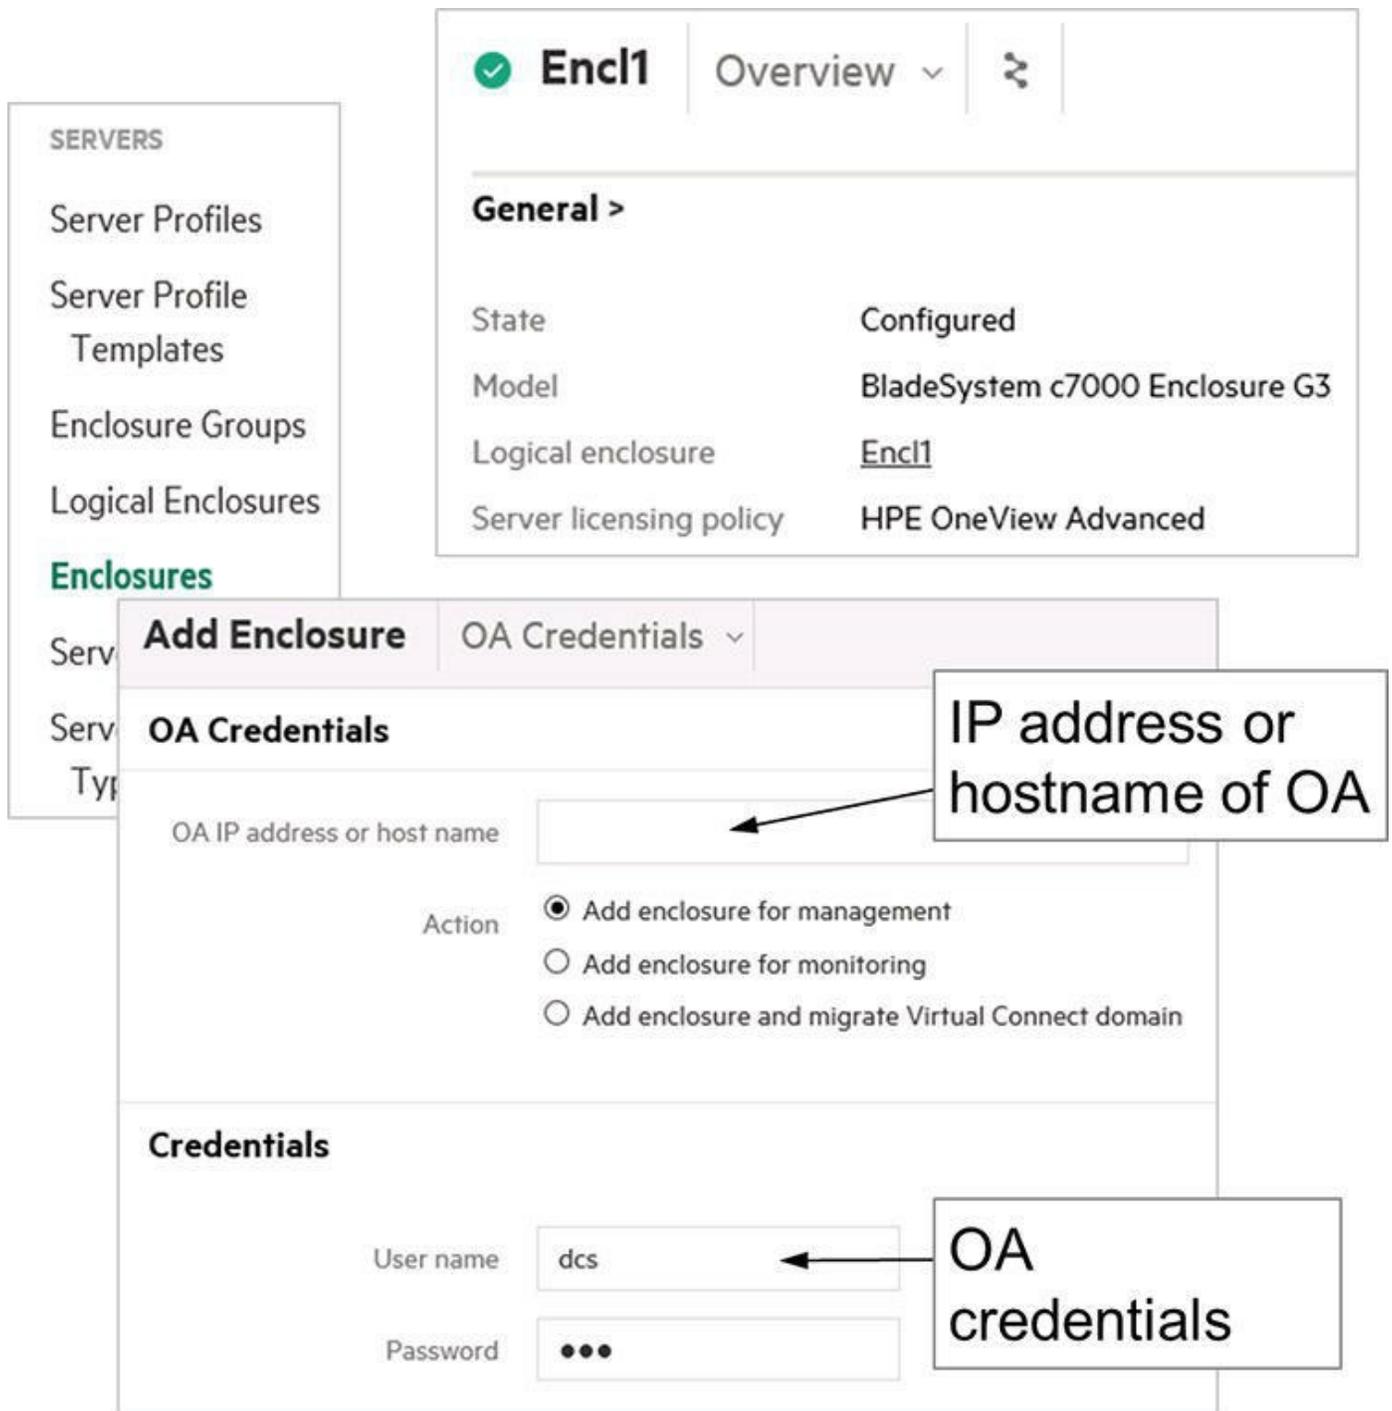


Figure 10-12 Adding enclosures to HPE OneView

As shown in [Figure 10-12](#), to add an enclosure, on the HPE OneView main menu select **Enclosures**. On the Enclosures page, click **Add enclosure**. The Add Enclosure window appears.

Monitored only

When adding an enclosure that will be monitored only, specify:

Onboard Administrator (OA) IP address or hostname—The IP address or Domain Name System (DNS) of the primary OA module in the enclosure.



Note

The name that will be associated with the enclosure in HPE OneView UI is the enclosure name set in the OA and is not the same as the DNS name associated with the IP address of the OA module.

Add an enclosure—If the server blades will be monitored only, no license is required for the servers. The selection applies to all server blades in the enclosure.

Credentials—Username and password of an OA user with administrator-level privileges.

Adding an enclosure for monitoring, like adding one for management, requires this basic information. When adding enclosures, you will see a sequence of messages appear on the Enclosures page indicating the progression of the operation. This includes:

Retrieve enclosure information

Configuring enclosure

Add servers

Because a monitored enclosure cannot be managed, HPE OneView does not create various configuration objects for a monitored enclosure. These include an enclosure group, a logical interconnect group, and logical interconnects.

You will find server hardware instances for each server blade on the Server Hardware page. You also will find a data center rack object populated with the enclosure because this is considered a monitoring feature.

Managed with or without a VC domain

The screenshot shows the 'Add Enclosure' configuration page with the following sections and callouts:

- Add Enclosure** (Header) with a dropdown menu set to 'OA Credentials'.
- OA Credentials** section:
 - Field: 'OA IP address or host name' (Callout: 'IP address or hostname of OA')
 - Action: Radio buttons for 'Add enclosure for management' (selected), 'Add enclosure for monitoring', and 'Add enclosure and migrate Virtual Connect domain' (Callout: 'To be managed or monitored only?')
- Credentials** section:
 - Field: 'User name' with value 'dcs' (Callout: 'OA credentials')
 - Field: 'Password' with masked characters '●●●'
- General** section:
 - Field: 'Enclosure group' with value 'Prod enclosure group' (Callout: 'Enclosure group name')
 - Licensing: Radio buttons for 'HPE OneView Advanced' (selected) and 'HPE OneView Advanced w/o iLO' (Callout: 'Server licensing'). A 'Learn more' link is also present.
- Firmware** section:
 - Field: 'Firmware baseline' with value 'Manage manually' (Callout: 'SPP firmware version or set to manual')

Figure 10-13 Enclosure that is managed and without a VC domain

The process of adding an enclosure that will be managed brings the rack, the enclosure, and the enclosure’s server hardware and interconnects under management. [Figure 10-13](#) shows the Add Enclosure screen. There are two main scenarios for adding an enclosure that will be managed.

Adding an enclosure—You can add an enclosure that has been newly installed, the OA Enclosure Bay IP Addressing (EBIPA) settings have been configured, but a VC domain does not exist.

Existing VC domain—You can also add an enclosure that is a member of a VC domain, which involves migrating the enclosure from a VC domain.



Note

Before you add the first enclosure to HPE OneView, you can configure various template objects. These include the Ethernet and Fibre Channel networks, the logical interconnect group, Ethernet and Fibre Channel uplink sets, and the enclosure group. With those configuration tasks complete, you can add the first enclosure. After the configuration objects are in place, you can potentially add other enclosures without further configuration work, if all enclosures are equivalent.

When adding an enclosure that will be managed, specify:

OA IP address or hostname—The IP address or DNS name of the primary OA module in the enclosure.

Add an enclosure as—Whether the server blades will be managed (which implies a license for each server is required) or monitored only (which implies no license is required for the servers). The selection applies to all server blades in the enclosure.



Note

If you select **Monitored**, the screen display changes to show only the Credentials section. If you select **Managed**, the screen displays as shown in the graphic—the Credentials, General, and Firmware sections appear.

Credentials—Username and password of an OA user with administrator-level privileges.

Enclosure group—The enclosure group to which the enclosure will be assigned along with an associated logical interconnect group.

Licensing—For a managed enclosure, whether the HPE OneView licenses for the server blades include iLO Advanced or not.

Firmware baseline—The firmware for the enclosure's components (OA and VC modules and servers) can be manually managed or updated based on a selected HPE SPP bundle installed on the appliance. An installation of the HPE OneView appliance includes one SPP firmware bundle. Additional SPPs, including custom bundles, can be uploaded to the appliance's repository subject to the maximum amount of storage space supported.

The logical interconnect is a collection of interconnect modules and typically includes

pairs of VC modules (FlexFabric, FlexFabric-20/40 F8, Flex-10/10D, Flex-10, and Fibre Channel). Interconnects in an enclosure share resources on their uplinks, share available networks, and manage the connectivity to the server blades on the downlinks. As a result of this type of design, an interconnect failure offers redundancy by failing over traffic to another interconnect, ensuring that a server maintains network connectivity. NIC teaming, in the case of LAN traffic, or multipath IO (MPIO), in the case of storage traffic, must be implemented by the server operating system to ensure failover can be accomplished.

When you add an enclosure, the existing OA configuration is maintained, as are the EBIPA and other related settings. The VC configuration, if any is present on the VC modules, will be erased unless the VC configuration is migrated.

If the firmware of OA modules is not at the minimum firmware level, or if you choose to update the firmware to a more recent version automatically by selecting an SPP, this task is performed during the Add Enclosure operation. This behavior also applies to the server iLO management processors.

If the firmware on the VC modules must be updated to a required minimum level, it is actually performed manually by the administrator after the enclosure is added. Optionally, you can update the VC firmware to any other version, provided the SPP firmware bundle is loaded on the HPE OneView appliance, just as you can for all enclosure components.



Note

If you add a second OA module to a managed enclosure that previously had only one OA module, you must refresh the enclosure in HPE OneView UI in order to discover the second module.

When adding an enclosure, you will see a sequence of messages appear on the Enclosures page indicating the progress of the operation. This includes:

Retrieving enclosure information

Configuring enclosure

Adding servers and interconnects

Adding an enclosure by migrating its VC domain configuration

The screenshot shows the 'Migrate Enclosure' page with the following content:

Note: It is highly recommended that you back up and download the HP Virtual Connect Manager configuration prior to initiating the migration.

The following settings will not be migrated:

- User account specific information such as certificates, user accounts, LDAP, Radius, TACACS+, session timeout, and user role configuration.
- Port Monitoring configuration. After migration, port monitoring can be configured separately on each Logical Interconnect.
- HP Virtual Connect domain pools of MACs, WWNs, and serial numbers.

Do not make changes to logical interconnect groups or networks while an enclosure is being migrated.

Server profile addresses and identifiers will be migrated.

OneView does not support migration of SR-IOV configuration.

Ensure all configured interconnect modules are present within the enclosure and all hardware is included in the OneView support matrix.

Ensure iLO3 Firmware is at least 1.61 and iLO4 firmware is at least 1.30 for all servers.

Ensure a OneView Advanced license is available for all servers.

Issues have been identified, which may affect or prevent a successful migration. Evaluate all warnings before proceeding. Resolve any issues by modifying either the Virtual Connect domain or the OneView configuration.

The enclosure SB-1 will be migrated using OneView logical interconnect group LIG_USE115932C. 2 Virtual Connect interconnect module(s) and 3 server(s) will be brought under OneView control. 6 Ethernet network(s), 4 fibre channel network(s) and 1 server profile(s) will be migrated. Unless specified otherwise, warnings indicate a capability will not be migrated to OneView if the migration is performed. Please ensure the detected capability is not critical to operations before proceeding.

3 error(s) have been detected which will block the migration.
8 warning(s) have been detected.

Proceed with the migration by selecting the "OK" button at the bottom of the page.

[View full report](#)

Callouts on the left side of the image point to specific parts of the page:

- 'Settings that are not migrated' points to the 'Note' and 'The following settings will not be migrated' section.
- 'Example of settings that are migrated' points to the 'Server profile addresses and identifiers will be migrated' section.
- 'Interconnect modules supported? Minimum iLO firmware installed?' points to the 'Ensure all configured interconnect modules...' and 'Ensure iLO3 Firmware...' sections.
- 'View report to see blocking issues that must be resolved' points to the '3 error(s) have been detected...' section.

Figure 10-14 VC domain migration scenario

When you choose to add an enclosure by migrating its VC domain configuration, you start the process by specifying:

OA IP address or hostname

Managed as opposed to monitored

OA credentials

Instead of selecting an existing name in the Enclosure group drop-down list, select **Migrate from Virtual Connect domain**. Doing so will refresh the screen and additional fields will display. Among several page display changes, notice that the Firmware field is inactive.

Specify the credentials of a VC domain user with administrator-level privileges.

Next, you need to enable HPE OneView to check (test) the compatibility of the VC domain relative to HPE OneView requirements for migrating an enclosure. Click **Test compatibility** to proceed to the potential step of having the enclosure's VC domain configuration migrated.

As shown in [Figure 10-14](#), when migrating an enclosure's VC domain configuration, several minimum requirements must be met, such as iLO firmware levels, and configuration settings that will be migrated and those that will not.

Smart Search

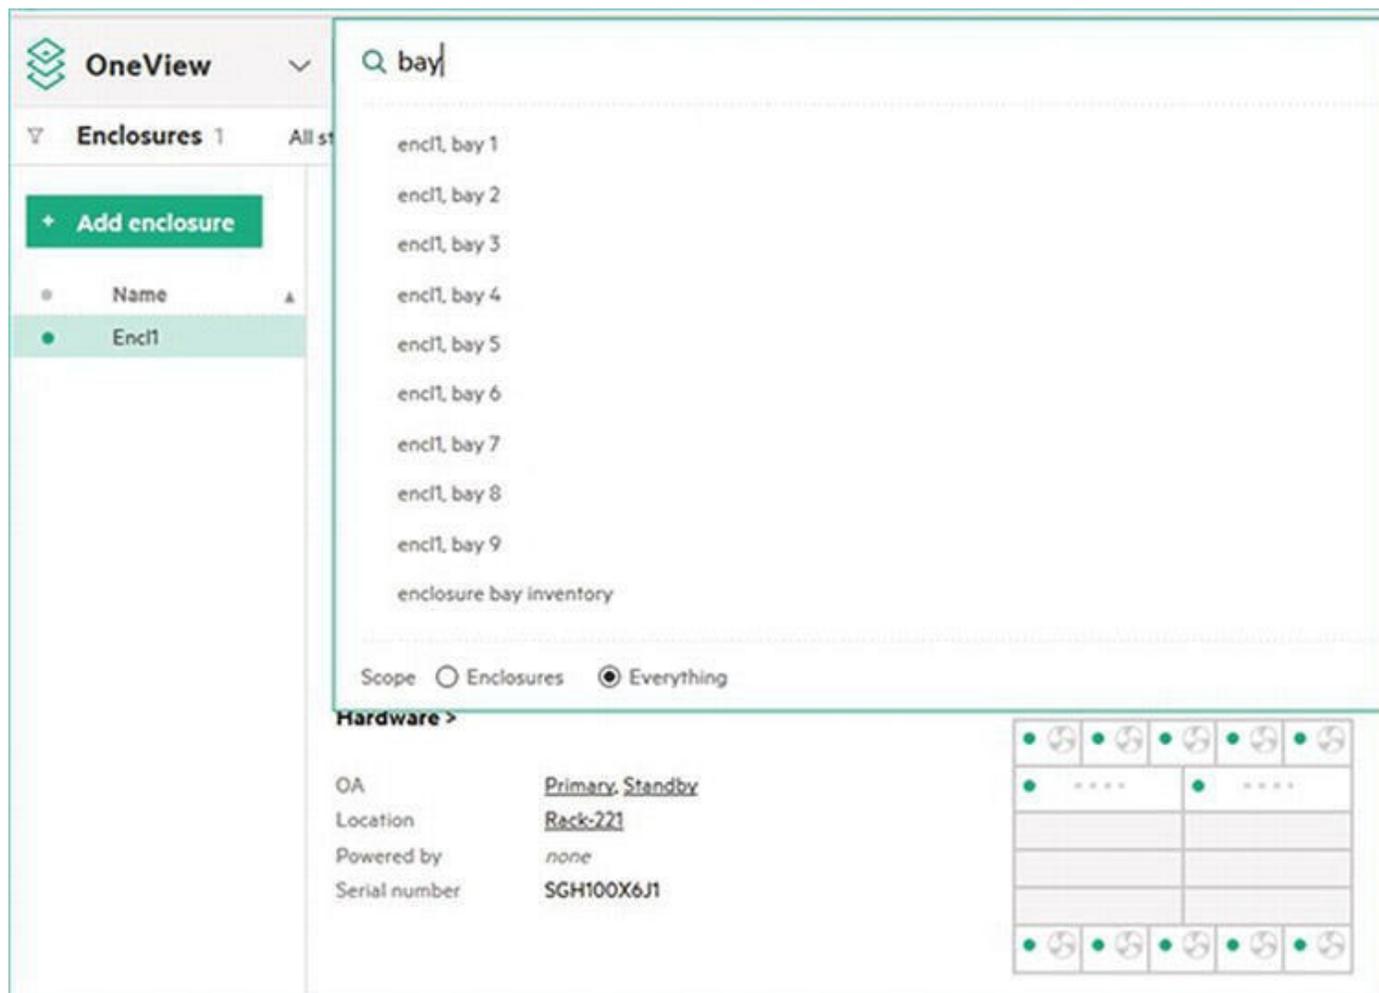


Figure 10-15 OneView Smart Search

Smart Search, shown in [Figure 10-15](#), is an HPE OneView feature that provides fast, sophisticated search capabilities to instantly find the information you are looking for, such as configured objects and device information, without having to search through long and complicated tree views. Smart Search provides an instant drop-down list when searching for devices based on physical MAC addresses, WWNs, and device names, among other criteria. When you start typing, search suggestions are provided based on pattern matching and previously entered search criteria. For example, Smart Search makes locating resources by model as simple as entering the model string (such as

BL660), so you can inventory or take action on a set of devices. To search:

Select a suggestion to change your filter to the suggestion and submit it (as if you had pressed **Enter**).

Press **Enter** to see the list of search matches.

When you find what you are looking for in the search results, which are organized by resource type, select the item to navigate to it.

- If you are doing a resource match, the master pane is filtered to match your search input.
- If you enter a multi-word search term, results show matches for all words you enter.



Note

Enter complete words or names as your search criteria. Partial words or names might not return the expected results.

The default search behavior is to focus on the resource you are currently viewing. However, to broaden the scope of your search across all resources, you can select **Everything**, which searches all resources. Certain resources do not include the option to choose between the current resource or everything, in which case the default search is for everything.



Note

The Smart Search feature does not search the help system.

HPE OneView provides proactive alert notifications by email (instead of using SNMP trap forwarding) and automated alert forwarding. You can view, filter, and search your alerts using Smart Search. Alerts can be assigned to specific users and annotated with notes from administrators. Notifications or traps can be automatically forwarded to enterprise monitoring consoles or centralized SNMP trap collectors.

Learning check

You are visiting a customer who tells you that they have recruited a new assistant system administrator who is tasked with streamlining their infrastructure management. They currently rely on Insight Manager, but they want to use scripting to perform server deployment. What should you advise?

Activity: Creating server profiles in HPE OneView

This activity requires you to use Internet Explorer or Chrome to browse to the emulator index page. On the emulator index page select **OneView Server Profiles**.

Follow the instructions as they appear on the screen.



Note

Access the emulator by entering the hyperlink into your Chrome or Internet Explorer browser.

<http://hpe.azureedge.net/gen10emulators/index.html>

Activity debrief

What is the purpose of changing the Workload Profile?

What is the state of the Server power after the profile is complete?

HPE OneView Global Dashboard

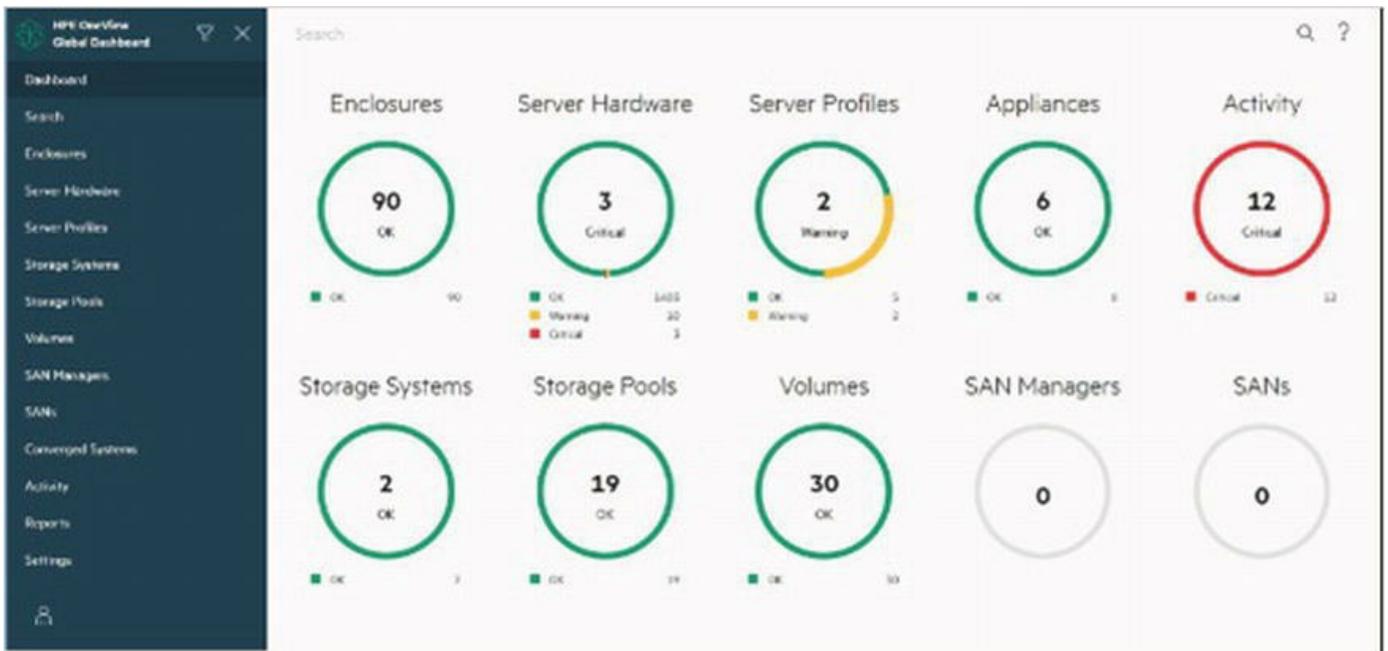


Figure 10-16 HPE OneView Global Dashboard

The HPE OneView Global Dashboard provides a unified view of the health of HPE servers, profiles, enclosures and HPE 3PAR storage systems across multiple appliances for ease of management.

The HPE OneView Global Dashboard, shown in [Figure 10-16](#), allows you to manage up to 12,800 nodes in the data center today and more in the future. It helps the IT staff troubleshoot alerts and view core inventory data across as many as 10 HPE OneView appliances. You get better infrastructure visibility which allows for more informed and faster decision-making. With single sign-on (SSO) places, you are one click away from powerful, device-level lifecycle management, and it removes the need to navigate between appliances. At no additional cost, customers can simply download and run the OneView Global Dashboard alongside HPE OneView, allowing them to manage their HPE BladeSystem c-Class, HPE ProLiant DL, HPE Synergy, and HPE Apollo systems as one infrastructure.

The HPE OneView Global Dashboard provides a unified view of the health of Hewlett Packard Enterprise servers, profiles, enclosures and HPE 3PAR storage systems across multiple appliances for ease of management. Some of the key features are listed below:

Create storage volumes with:

- SAN storage support—3PAR
- HPE StoreVirtual VSA storage

View server OS name in server hardware and server profile overview for faster access

to specific server information

Aggregates critical activates from multiple appliances into a single feed to quickly identify issues occurring on monitored hardware for prompt resolution

Generates reports for monitored assets to view inventory, including firmware versions, as well as compliance reports that allow you to verify that your equipment meets corporate standards

Integrated reporting for inventory allows you to see server hardware inventory, including associated firmware versions.

Create user-defined groups

Customizable dashboard

Increased scale

- Up to 10 appliances and up to 6400 nodes

New canned reports

VMware ESXi, Hyper-V and Kernel-based Virtual Machine (KVM) support

Localization—Simplified Chinese and Japanese

REST APIs

Easy access to reporting of inventory and health status

The HPE OneView Global Dashboard delivers compliance reports making it easy to identify what is not at your current firmware baseline to accelerate configuration compliance and consistency.

Integrated reporting for inventory allows you to see server hardware inventory, including associated firmware versions. Quickly verify that there are not duplicate resource identifications numbers (serial number, World Wide Identifier [WWID]), in your environment.

Near real-time search feature

The HPE OneView Global Dashboard has a Smart Search function that delivers access to devices of interest across data centers in seconds for better visibility of the systems enabling you to make faster decisions.

Search for attributes associated with a resource to quickly find and navigate to the desired item. Easily filter a view down to the set of resources that are having problems, enabling you to focus on only those items that need remediation.

Learning check

What is included in an HPE OneView server profile? (Select two.)

- A. Software license keys
- B. iLO settings
- C. Operating system images
- D. HPE StoreVirtual iSCSI storage volumes
- E. Error logs

Which process does HPE OneView automate when a profile is applied to the server?

- A. Operating system deployment
- B. Monitor data center infrastructure
- C. Assign MAC and WWID numbers
- D. Assign server hostname

HPE OneView partner integrations

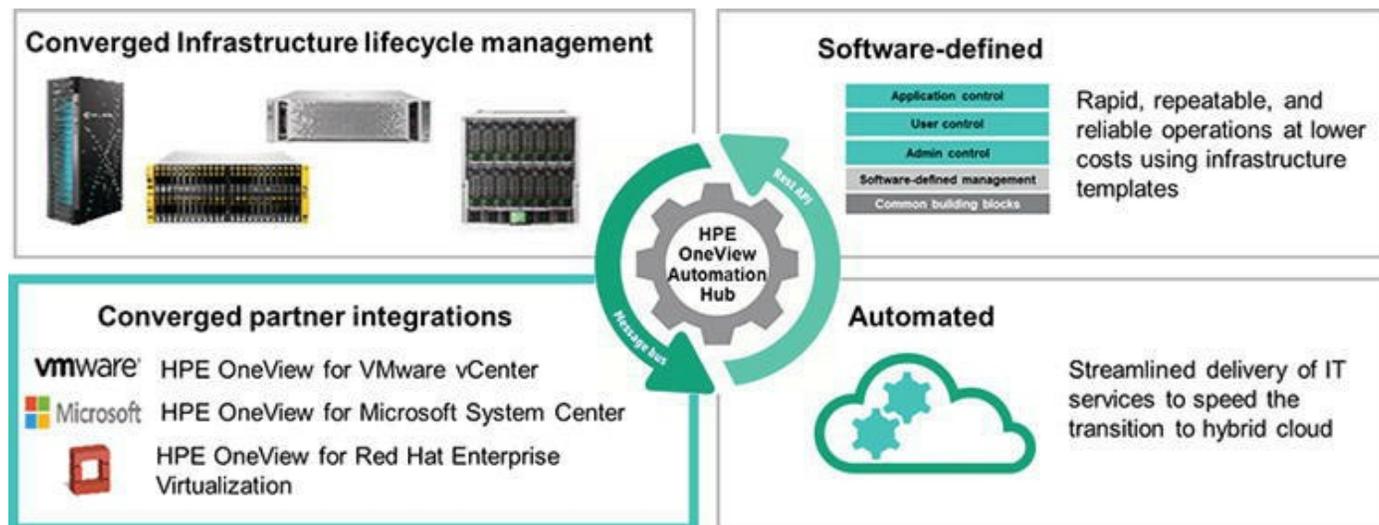


Figure 10-17 HPE OneView automation hub

As shown in [Figure 10-17](#), the concept of the automation hub is a key aspect of HPE OneView. The automation hub consists of an SCMB and the REST API. This means all of the information that HPE OneView collects, changes apparent to HPE OneView, various templates, and supported control functions can be accessed programmatically.

The automation hub supports two categories of converged management:

Infrastructure lifecycle management—Supports HPE ConvergedSystem, HPE 3PAR StoreServ storage systems, HPE BladeSystem c7000 enclosures with ProLiant BL server blades and VC modules, and ProLiant DL rack servers.

Partner integrations—Plug-ins for virtualization management products including:

- HPE OneView for VMware vCenter—Supports enhanced integration with VMware vCenter Operations Manager and vRealize Log Insight with HPE unique entitlement to use HPE vCenter Operations Manager integration on the standard edition. It also leverages HPE 3PAR StoreServ shared storage and boot from SAN for creating and growing clusters using an automated workflow.
- HPE OneView for Microsoft System Center—Supports health monitoring and alerting with server profile provisioning to create or grow a Hyper-V cluster.
- HPE OneView for RHEV—Supports network management and health monitoring for ProLiant servers and BladeSystem solutions.

The software-defined capabilities provide support for policy-driven infrastructure templates that drive automated activities. The templates capture best practices and

allow for rapid, repeatable, and reliable automated operations that reduce operational costs. This means that IT organizations can easily incorporate HPE OneView functions into workflows that extend beyond infrastructure management. It streamlines the delivery of IT services to help organizations speed the transition to hybrid cloud solutions.



Note

HPE OneView provides a right-to-use (RTU) agreement for these integrated partner products.

HPE OneView for Microsoft System Center

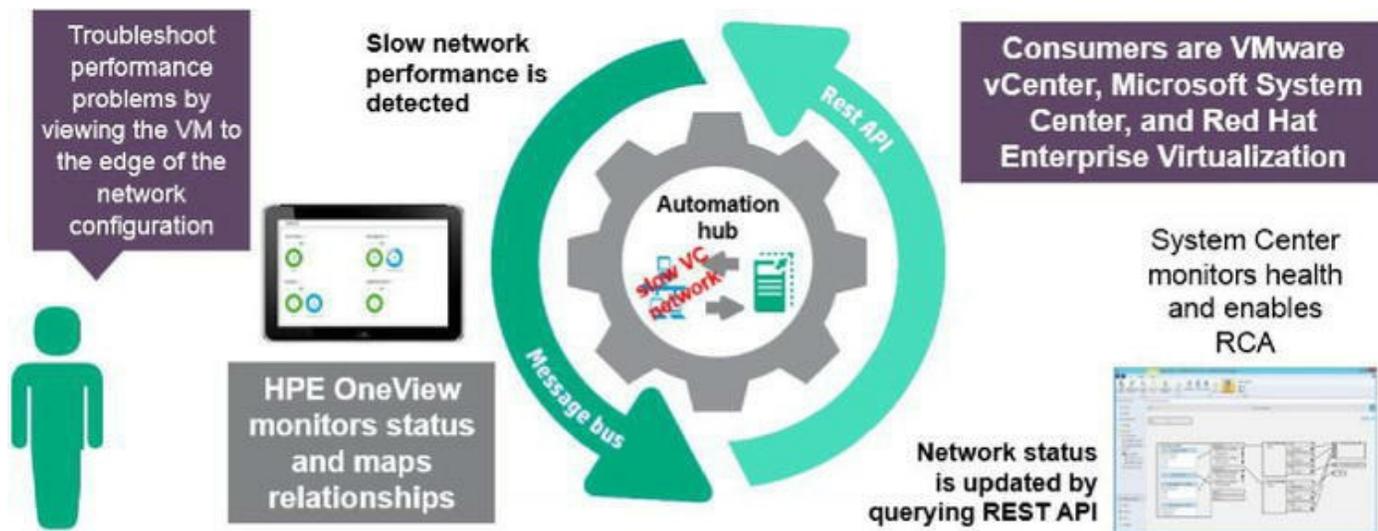


Figure 10-18 System health monitoring, configuration mapping, and remote support

One example of the HPE OneView automation hub being used for partner integrations is the HPE OneView for Microsoft System Center, which provides seamless integration of ProLiant and BladeSystem manageability features into Microsoft System Center.

As shown in [Figure 10-18](#), this integration provides comprehensive system health and alerting, driver and firmware updates, operating system deployment, detailed inventory, enhanced provisioning (integrated with HPE OneView server profiles), and HPE fabric visualization.

Using this plug-in for Microsoft System Center, an administrator can provision hosts using HPE OneView server profiles to create or grow a Hyper-V cluster consistently and reliably. In addition, the health monitoring and alerting information from an HPE

OneView infrastructure perspective shows relationships clearly.

This integration delivers server profile-based Grow Cluster deployment and automated HPE StoreVirtual or StoreOnce deployment for secure backup and recovery. In addition, Microsoft System Center displays the network mapping, storage connections, and hardware status.

HPE OneView for VMware vCenter

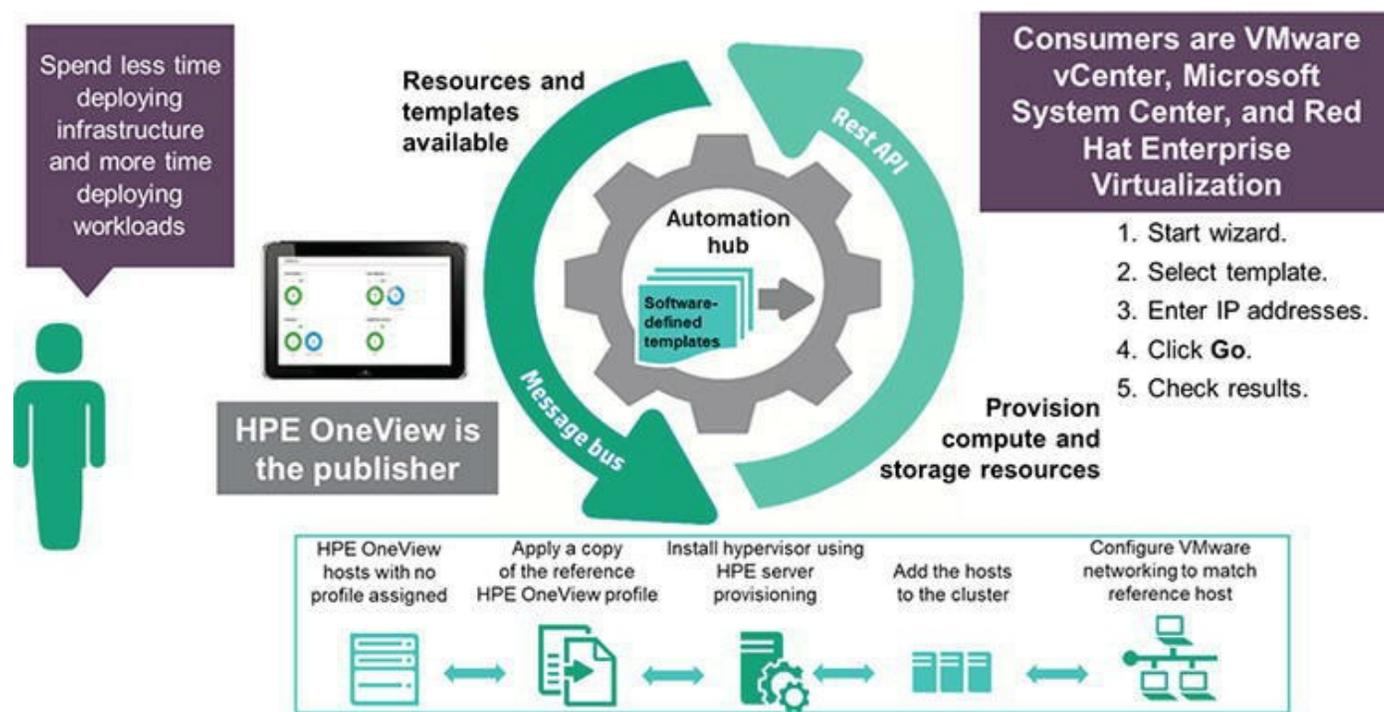


Figure 10-19 End-to-end cluster deployment

HPE OneView for VMware vCenter is an integrated application that can be used to manage ProLiant servers and HPE storage systems. It brings together management of VMware vSphere and HPE infrastructure within the context of the VMware vSphere console.

By leveraging HPE OneView software-defined templates and the HPE OneView REST API, the HPE OneView for VMware vCenter server simplifies the process of deploying a complete vSphere cluster. Processes that used to require 15 or more steps can now be accomplished in five steps—and with considerably less administrative time and effort.

[Figure 10-19](#) shows how organizations can streamline infrastructure management and dedicate a greater portion of their IT resources to defining and delivering the workloads

that will live on that virtual infrastructure.



Note

For more information, enter the following URL into your browser;

<https://www.hpe.com/us/en/product-catalog/detail/pip.4152978.html#>

Architecture overview

HPE OneView for VMware vCenter consists of these modules:

Core Module—This required module provides the framework required by the Server Module for vCenter and Storage Module for vCenter.

Server Module for vCenter—This optional module adds ProLiant and BladeSystem hardware monitoring into the vCenter Server console. This module provides server hardware management capabilities, including comprehensive monitoring, firmware updates, ESXi image deployment, remote control, end-to-end monitoring for VC, and power optimization for HPE servers in the VMware environment.

Storage Module for vCenter—This optional module enables monitoring of the HPE storage infrastructure from the vCenter Server console. It provides storage configuration and status information for mapping VMs, datastores, and hosts to logical unit numbers (LUNs) on HPE storage systems. It supports provisioning on HPE 3PAR StoreServ and HPE StoreVirtual systems. Provisioning tasks include creating, expanding, or deleting a datastore, along with creating or cloning a VM. This module also displays view-only information for HPE StoreOnce Backup systems.

3PAR Recovery Manager for VMware—This module provides Virtual Copy management of HPE 3PAR virtual copies and recovery of VMs and datastores.

Storage Administrator Portal—This module enables you to configure storage systems and management servers for management by the Storage Module for vCenter.

The plug-in software is installed on a separate server or a VM. During the installation process, you provide credentials for the plug-in to communicate with VMware vCenter. After installation is complete, you can access the vCenter Server to enable the installed plug-in and then configure other credentials to enable the plug-in to access HPE OneView.

You access the plug-in through the vSphere web client. The HPE OneView for VMware vCenter plug-in uses the published vCenter APIs to communicate with the vCenter Server. The server module supports consoles for connecting to HPE OneView, OA, and server iLOs. The storage module connects to storage systems such as HPE 3PAR

StoreServ to obtain information.

Coordinated and scheduled firmware management

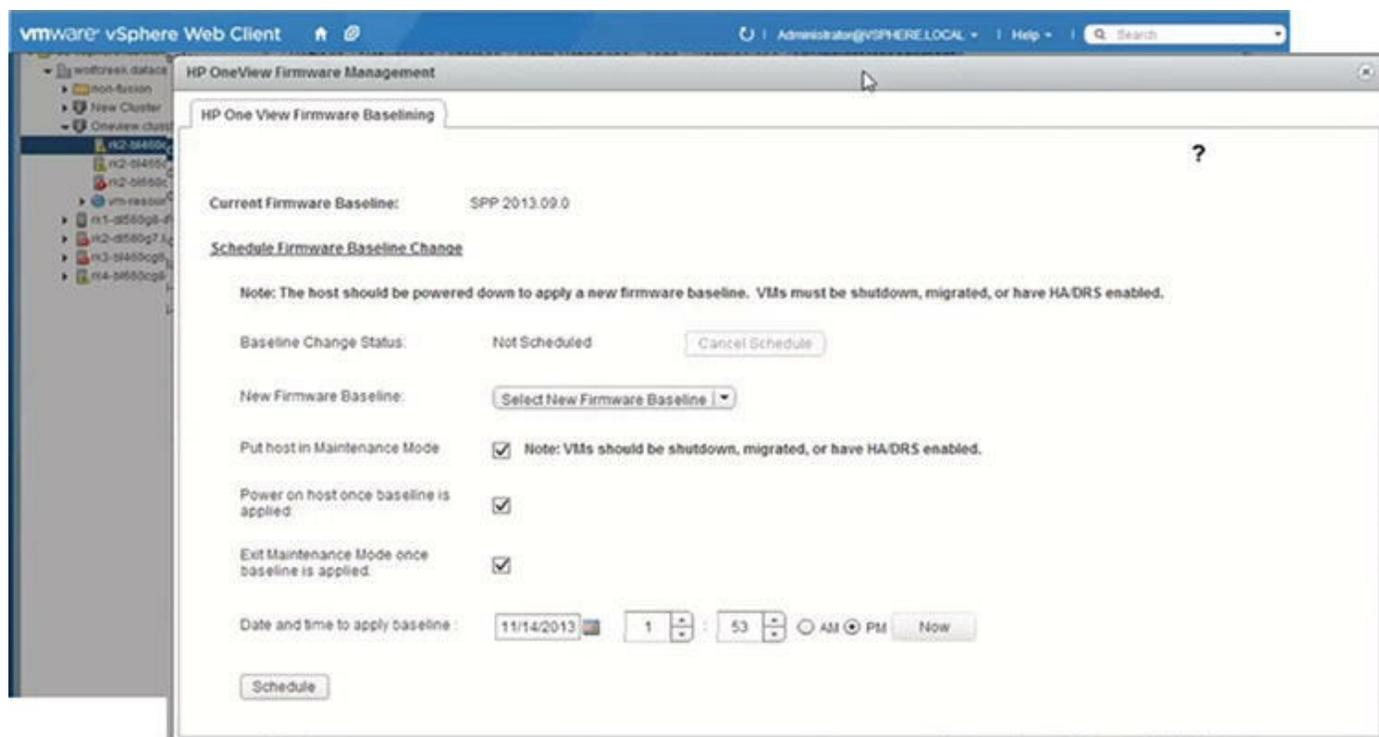


Figure 10-20 HPE OneView Firmware Management page

The HPE OneView for VMware vCenter plug-in allows you to view firmware information and update firmware for an entire cluster or selected servers in the cluster, as shown in [Figure 10-20](#).

The HPE OneView Firmware Management page is used with the vSphere web client to schedule firmware updates for hosts managed by HPE OneView. This feature is available only if the selected host has firmware managed by HPE OneView.

HPE OneView for RHEV

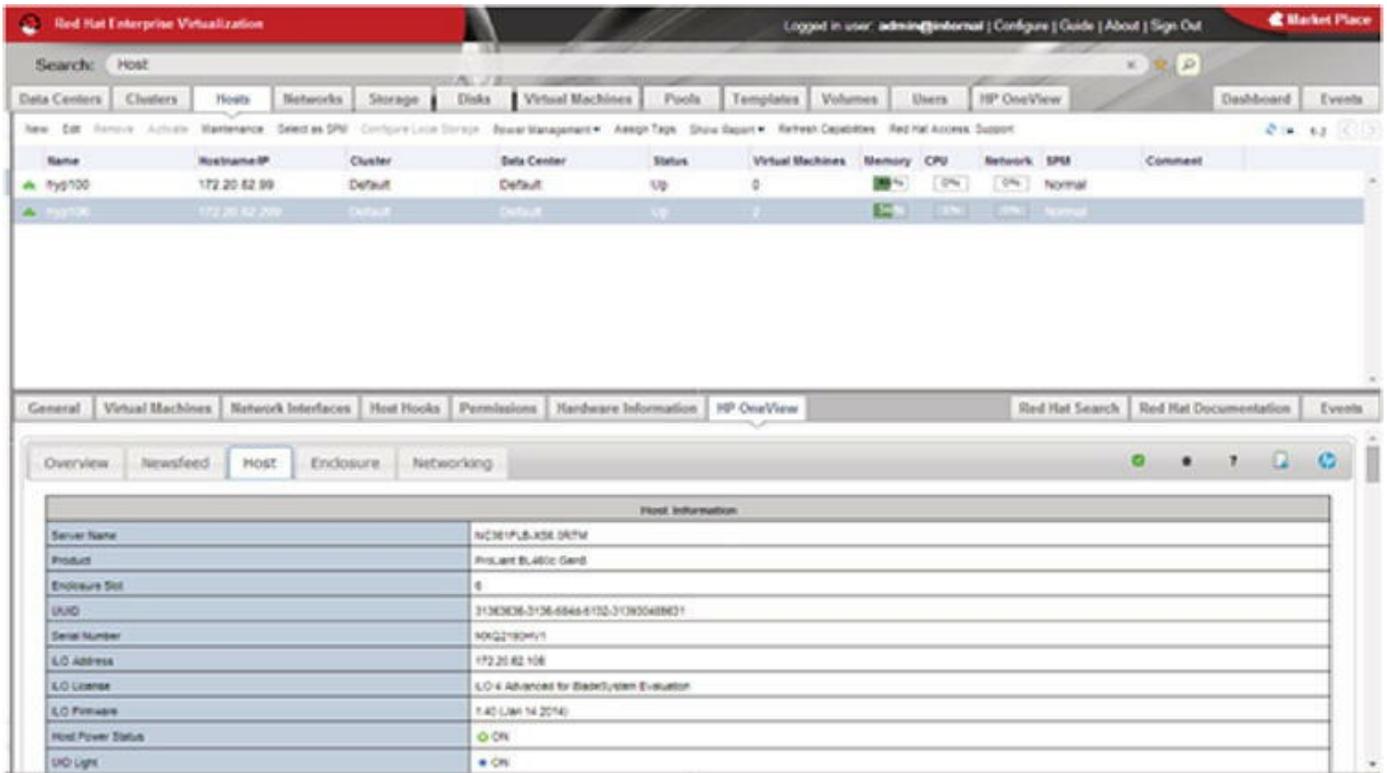


Figure 10-21 HPE OneView for RHEV

The HPE OneView integration for RHEV is a user interface plug-in that seamlessly integrates the manageability features of ProLiant, BladeSystem, and VC within the RHEV management console. As shown in [Figure 10-21](#), this integration gives administrators control over their HPE infrastructure and their Red Hat virtualized environment from a single screen.

HPE OneView for RHEV provides the following capabilities from within the RHEV management console:

Displays host and infrastructure inventory and health information on a single screen, reducing IT administration time and effort

Displays network relationships from the RHEV-defined network to the physical switch

Provides SSO capability to trusted HPE infrastructure tools, including OA, iLO, and VC

Offers a single screen display of host firmware, including:

- System ROM
- iLO
- Smart Array controllers and power management controller
- OA

- VC module

Logs events and delivers notifications in a dedicated newsfeed

Remote management

HPE iLO, Insight Online, and related remote support tools are available as part of an HPE warranty or support agreement. HPE iLO is an embedded management technology that supports the complete lifecycle of all ProLiant Gen8 and Gen9 servers, from initial deployment to remote management and service alerting.

Remote management with HPE iLO Advanced

All HPE OneView SKUs currently ship with an iLO Advanced license (except for the upgrade SKU). iLO Advanced provides remote access to server power control and event logs.

Features available to help solve complex IT problems include:

Graphical remote console turns a supported browser into a virtual desktop, giving the user full control over the display, keyboard, and mouse of the host server. The console is independent from the operating system, the console displays remote host server activities (such as shutdown or startup operations), and it can be launched from the HPE OneView server profile page.

Shared console and console replay functions allow up to six team members to view and share control of a single virtual KVM session. They can also capture and save screen video for later review.

Serial record and playback functions save the text-based output data for later access.

Remote System keeps a log of activity for later troubleshooting or reference.

HPE Insight Online

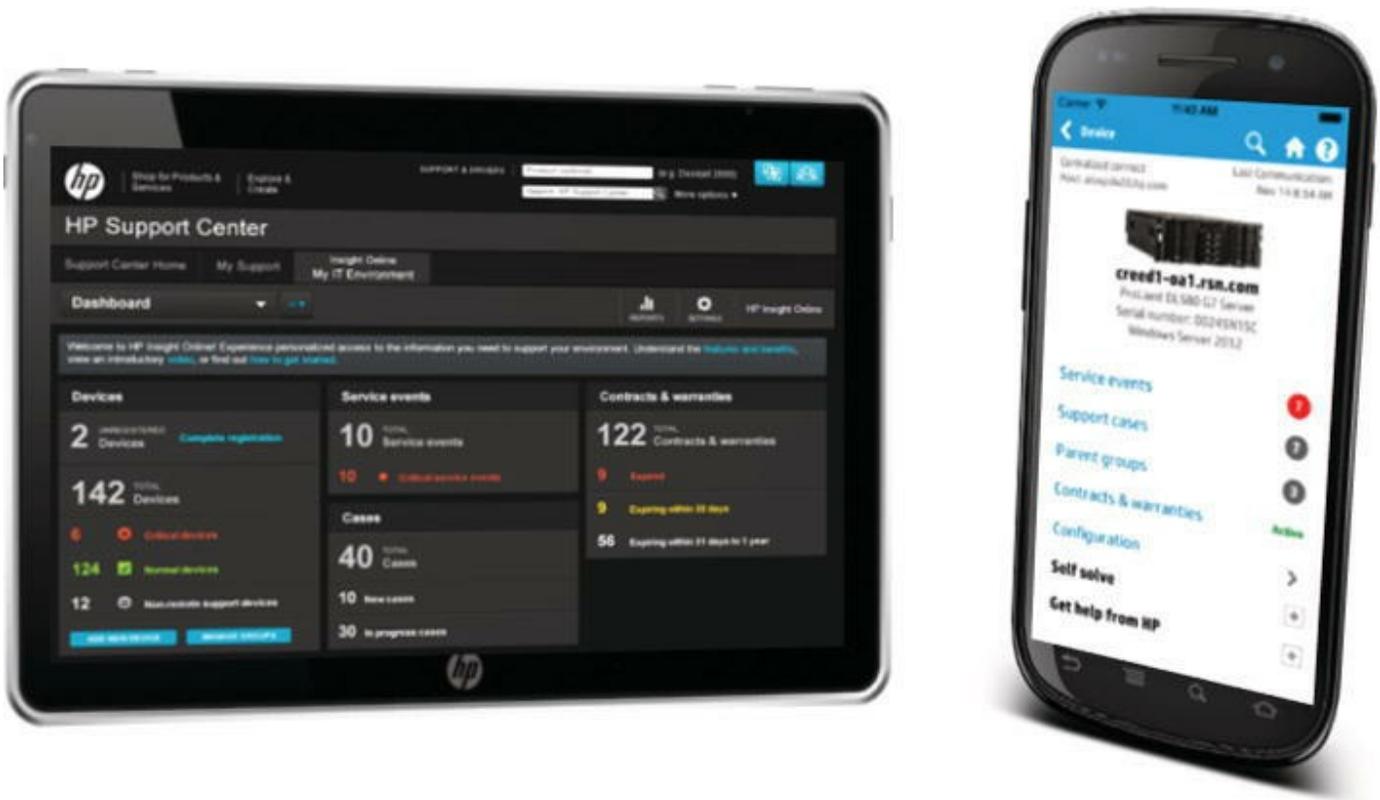


Figure 10-22 HPE Insight Online

HPE Insight Online is a cloud-based infrastructure management and support portal available through the HPE Support Center. Powered by HPE remote support technology, it provides a personalized dashboard to simplify tracking of IT operations and to view support information from anywhere at any time. Use the Insight Online dashboard to track service events and support cases, view device configurations, and proactively monitor HPE contracts and warranties as well as HPE Proactive service credit balances. It is an addition to the HPE Support Center portal for IT staff who deploy, manage, and support systems, plus HPE Authorized Channel partners who support an IT Infrastructure.

HPE Insight Online is complementary to HPE OneView. Insight Online provides access to device and support information in a cloud-based personalized dashboard, so you can stay informed while in the office or on the go, as shown in [Figure 10-22](#). Use the Insight Online dashboard to track service events and support cases, view device configurations, and proactively monitor HPE contracts and warranties for all devices monitored by HPE remote support tools. No installation is required to use Insight Online.

HPE Insight Remote Support

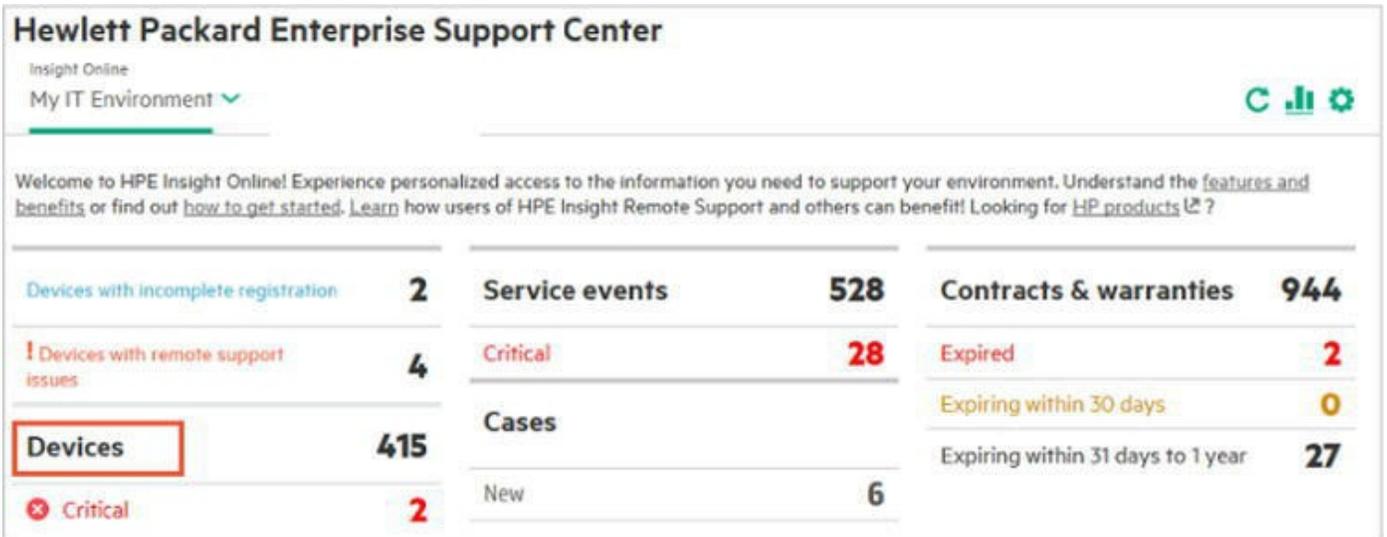


Figure 10-23 Insight Online dashboard

Figure 10-23 shows the Insight Online dashboard, which provides a quick overview of the current status of Devices, Service events, Support Cases, and Contracts.

Choice of HPE remote support tools with Insight Online

Insight Online direct connect

- ProLiant Gen8 and later servers
- SMB or remote IT environment
- No software installation or onsite console
- HPE Insight Online personalized dashboard and partner view
- Warranty and Service Contract entitlement management
- Detailed fault analysis of environment hardware

Insight Remote Support central connect

- Supports most HPE servers, storage, network devices
- Monitors up to 3500 devices per host server
- Onsite console
- Optional Insight Online dashboard and partner view
- Warranty and Service Contract entitlement management

Figure 10-24 Insight Online is provided by two connection methods

To take full advantage of Insight Online, you need to install HPE remote support

technologies. Remote support with Insight Online is provided by two connection methods as shown in [Figure 10-24](#):

Direct connect—Available for ProLiant Gen10, Gen9, and Gen8 servers and BladeSystem enclosures. Direct connect enables these devices to automatically submit hardware failure and diagnostic information directly to HPE for analysis, case generation, and automated parts replacement. No centralized hosting device is required; instead, you use Insight Online as the online console. ProLiant servers benefit from agentless remote support monitoring provided with HPE iLO 4 and iLO 5 management. This method is ideal for SMB and remote sites with ProLiant Gen10, Gen9, and Gen8 servers.

Insight Remote Support central connect—Available for servers, storage, and networks. Using central connect, you register the device to communicate with HPE through an Insight Remote Support centralized hosting device in your local environment. This method is ideal for HPE Converged Infrastructure IT environments with multiple device types.

These solutions automatically send hardware failures and configuration information to HPE for fast, accurate diagnosis and repair. With Insight Online, all devices monitored by Insight Remote Support central connect or Insight Online direct connect can be auto-populated to the Insight Online personalized dashboard to provide 24 × 7 access to product and support information regardless of location. HPE Authorized Channel partners can view asset information shared to them by their customer in a separate “My customers” section in Insight Online.



Note

To download the Insight Remote Support software, enter the following URL into your browser:

<https://www.hpe.com/uk/en/services/remote-it-support.html>

Embedded Remote Support

Embedded Remote Support enables you to send hardware issues to HPE Support, as part of an HPE warranty or Care Pack. You can easily and securely export the Active Health file to an HPE Support professional to help resolve issues faster and more accurately. When you use the Embedded Remote Support feature, choose from the two configuration options: Insight Online direct connect and Insight Remote Support central connect.

Device support

The following Insight Online direct connect device types support direct connect registration:

- ProLiant Gen10, Gen9, and Gen8 servers

The following Insight Remote Support central connect types support central connect registration:

- ProLiant Gen10, Gen9, and Gen8 servers

Insight Online direct connect

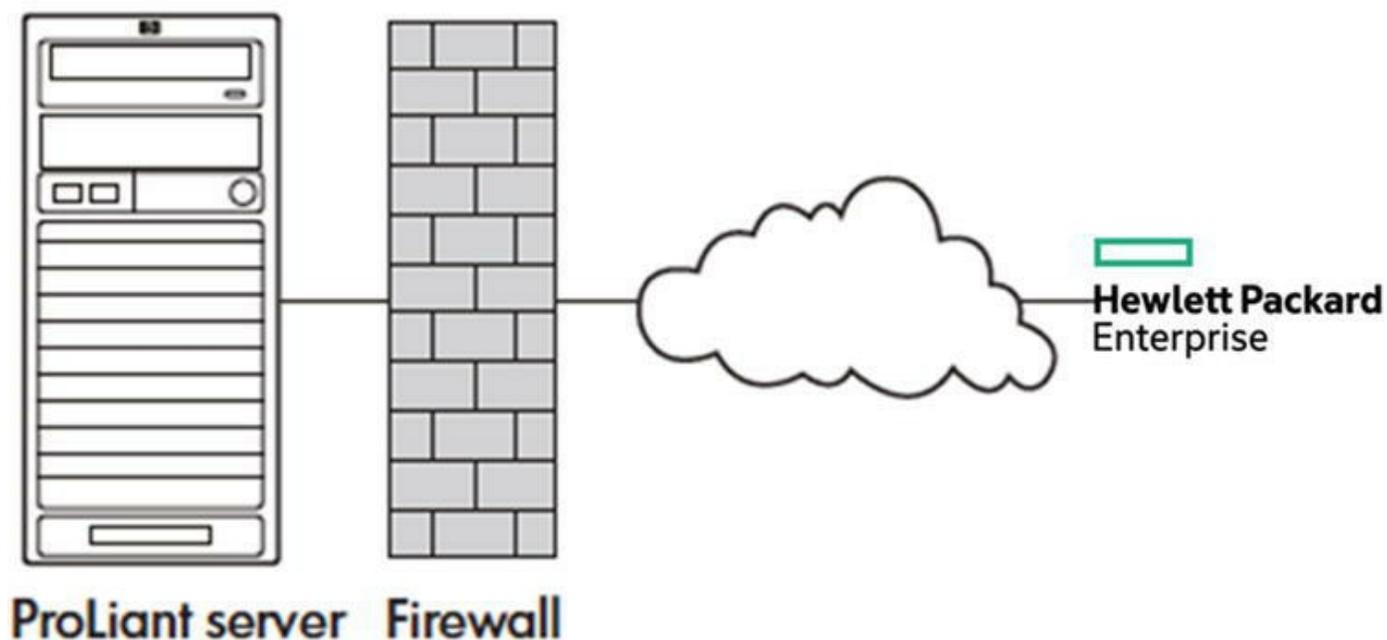


Figure 10-25 Insight Online direct connect

As shown in [Figure 10-25](#), you can register a server to communicate directly with Insight Online without the need to set up an Insight Remote Support centralized hosting device in your local environment. Insight Online is your primary interface for remote support information.

Insight Remote Support central connect

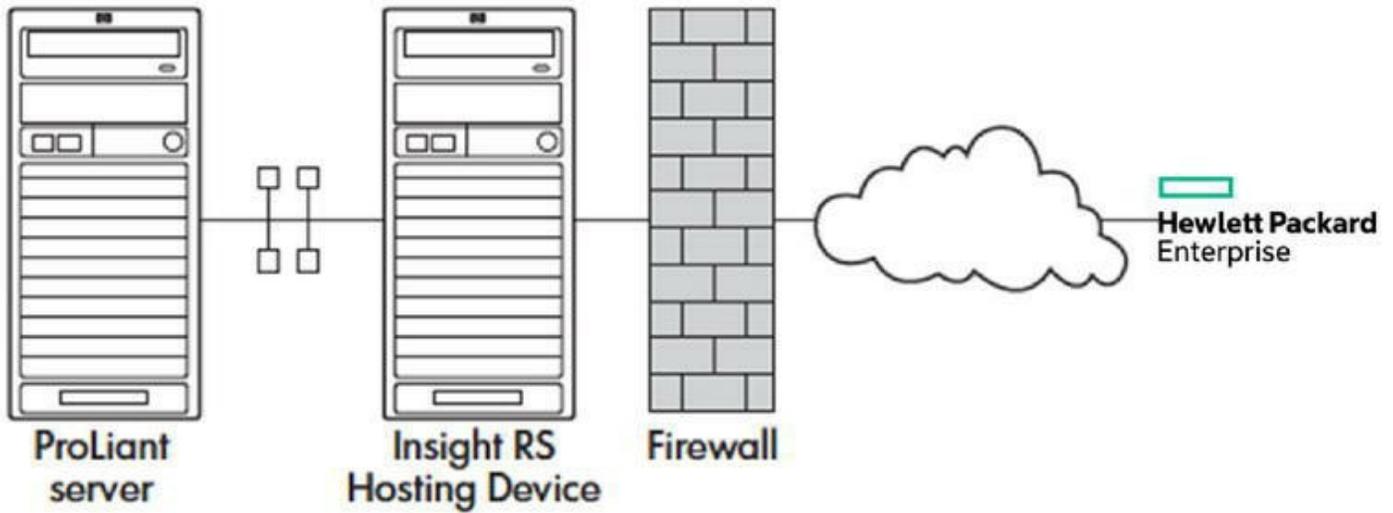


Figure 10-26 Insight Remote Support central connect

You can register a server to communicate with HPE through an Insight Remote Support centralized hosting device in your local environment, as shown in [Figure 10-26](#). All configuration and service event information is routed through the hosting device. This information can be viewed using the local Insight Remote Support Console or the web-based view in Insight Online (if it is enabled in Insight Remote Support).

HPE Insight Remote Support

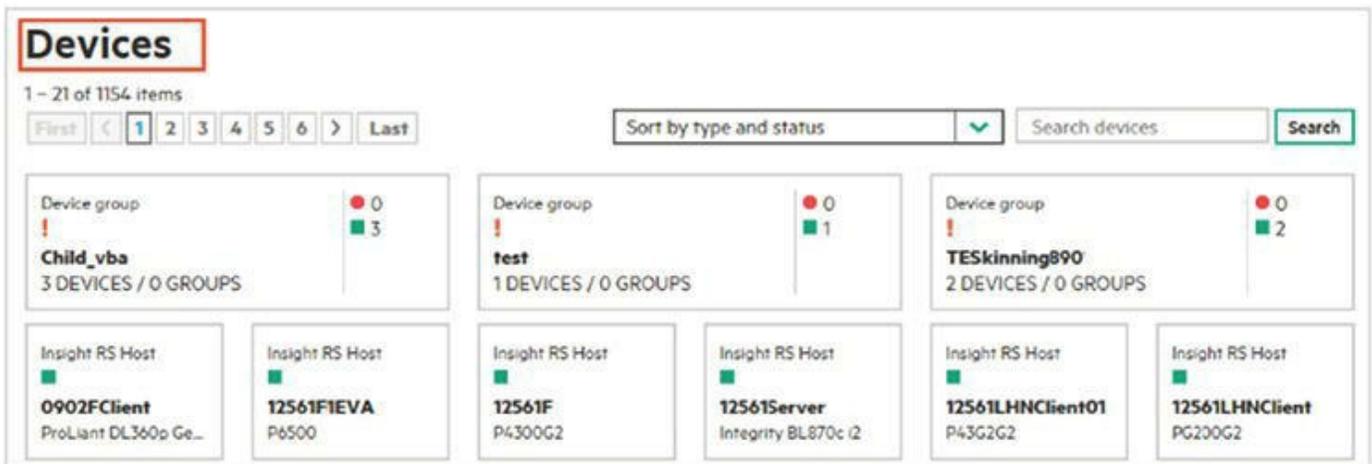


Figure 10-27 HPE Insight Remote Support Devices screen

Figure 10-27 shows the Insight Remote Support Devices screen, which is optimized to support up to 3500 monitored devices per hosting device and can be installed on a Windows ProLiant hosting device or a Windows virtual guest.

Intelligent Provisioning includes the embedded Insight Remote Support feature, which allows you to register supported servers for HPE remote support. HPE Insight Remote Support provides automatic submission of hardware events to HPE to prevent downtime and enables faster issue resolution.

There are two ways to register a ProLiant server for remote support—directly with HPE or through an Insight Remote Support hosting device.

When you register a server to communicate directly with Insight Online, you do not need to set up an Insight Remote Support centralized hosting device in the local environment. Insight Online will be the primary interface for remote support information.

When you register a server to communicate with HPE through an Insight Remote Support centralized hosting device in a local environment, all configuration and service event information is routed through the hosting device. This information can be viewed by using the local Insight Remote Support console or the web-based view in Insight Online (if it is enabled in Insight Remote Support). Insight Remote Support central connect shows the central connect configuration with a ProLiant server.



Note

Additional Insight Remote Support information is available from Insight Remote Support 7.x documentation. To access the material, enter the following URL into your browser;

<http://www.hpe.com/info/insightremotesupport/docs>

Insight Online and Remote Support integration



Figure 10-28 Insight Online and Remote Support integration for customers and partners

Insight Online provides a one stop, secure web access to product and HPE support information specific to IT environments.

As shown in [Figure 10-28](#), Insight Online adds two sections to HPE Support Center.

My IT Environment—A custom view of the IT environment. Within the My IT Environment, users can view the following:

- Personalized dashboard
- Device status, configurations
- Contracts & Warranty status
- Auto-generated events tracking
- Support cases

- Check service credit balance
- Proactive reports

My Customers—Where one or multiple HPE authorized partners may be allowed to see or manage designated remote support devices, as the customer chooses from Insight Remote Support. In this view, they will see:

- Customer grid
- Customer dashboard
- Contracts & Warranty status
- Monitor service events
- Proactive reports

Insight Online and Insight Remote Support are available at no additional cost as part of a warranty, HPE Care Pack Service, or contractual support agreement with HPE.

Target markets of direct connect for remote support

Small business	IT data center and remote sites	High Performance Computing
<p>Needing:</p> <ul style="list-style-type: none"> - Simple deployment and automated support - Little or no maintenance (no host server required) - Partner or non-tech IT support - ProLiant servers only (Gen8 and later) <p>Solution: Insight Online direct connect</p>	<p>Needing:</p> <ul style="list-style-type: none"> - Automated support for remote sites and data center in heterogeneous environment - Consolidated anywhere, anytime view for all sites - Easy warranty and contract management <p>Solution: Use Insight Remote Support and direct connect methods with one Insight Online dashboard</p>	<p>Needing:</p> <ul style="list-style-type: none"> - Automated parts replacement for ProLiant servers - Low overhead for performance-sensitive customers - Operating system independent solution <p>Solution: Insight Online direct connect</p>

Figure 10-29 Direct connect for remote support users

Direct connect for remote support provides automated support, faster resolution, simplified contract, and warranty management. As shown in [Figure 10-29](#), it is ideal for the following types of customers:

Small business—Insight Online direct connect

- Simple deployment and automated support
- Little or no maintenance (no host server required)
- Partner or non-tech IT support
- ProLiant Gen10 and Gen9 servers

IT data center and remote sites—Insight Remote Support and direct connect methods with one Insight Online dashboard

- Automated support for remote sites and data centers in heterogeneous environment
- Consolidated anywhere, anytime view for all sites
- Easy warranty and contract management

High-Performance Computing—Insight Online direct connect

- Automated parts replacement for ProLiant Gen10 and Gen9 servers
- Low overhead for performance-sensitive customers
- Operating system independent solution

Registering for Insight Remote Support direct connect

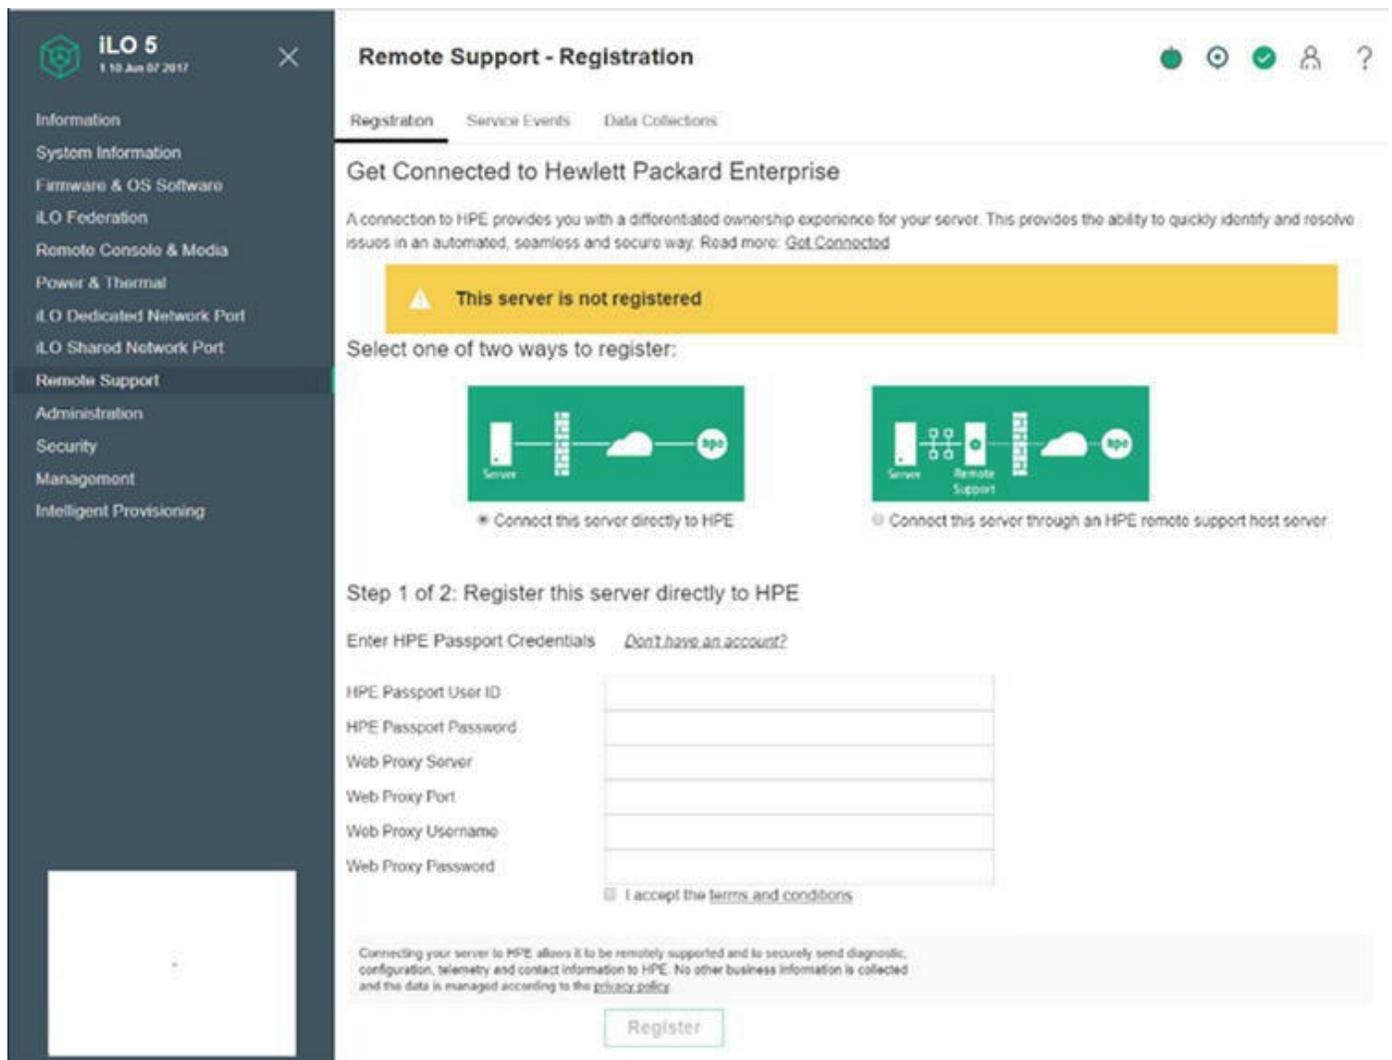


Figure 10-30 Remote Support registration

[Figure 10-30](#) shows the Remote Support registration screen. Use the following procedure to register a ProLiant server for direct connect remote support. By registering, you agree to send registration, service events, configuration, and Active Health System data to HPE. When you register for Insight Remote Support direct connect, you must complete these steps in both Intelligent Provisioning and the Insight Online portal.

First, verify that the server meets the prerequisites for using the Insight Remote Support solution.

Select **Register this server directly to HPE**. The page refreshes to show the Insight Online direct connect registration options.

Enter your HPE Passport credentials in the HPE Passport User ID and HPE Passport Password boxes. In most cases, your HPE Passport user ID is the same as the email

address you used during the Passport registration process.

Optional: Enter information about the web proxy server if the ProLiant server uses a web proxy server to access the internet.

Click **Register**. Clicking **Register** is Step 1 of a two-step registration process. Step 2 is completed in Insight Online. Allow up to five minutes for your registration request to be fully processed.

Navigate to the Insight Online website at <http://www.hpe.com/info/InsightOnline> and log in with your HPE Passport credentials.

Follow the instructions in Insight Online and provide your site, contact, and partner information so that HPE can deliver service for the server. If you have multiple servers to register, complete Step 1 for all of the servers, and then complete Step 2 for all of the servers during one Insight Online session.

Return to the Insight Remote Support page in Intelligent Provisioning, and then click **Confirm**.

Click the **Continue** right arrow to proceed to the Intelligent Provisioning home page (Overview).

Registering for Insight Remote Support through a centralized hosting device

Use this procedure to register a ProLiant server for Insight Remote Support central connect.

Verify that the server meets the prerequisites for using Insight Remote Support.

Select **Register this server through an HPE Insight Remote Support centralized hosting device**. The page refreshes to show the Insight Remote Support central connect registration options.

Enter the Insight Remote Support hosting device hostname or IP address and port number. The default port is 7906. Click **Register**.

Click the **Continue** right arrow to proceed to the Intelligent Provisioning home page (Overview).

Agentless Management Service

HPE recommends installing Agentless Management Service (AMS). AMS is one way in which iLO can obtain the server name. If iLO cannot obtain the server name, the

displayed server name in Insight Online and Insight Remote Support is derived from the server serial number.

AMS is installed automatically if you use the Intelligent Provisioning Recommended installation method for Windows installation.

If you did not install AMS, ensure that the server name is displayed correctly in Insight Online and Insight Remote Support.

Data collected by Insight Remote Support

By registering for Insight Remote Support, you agree to send registration, service events, configuration, and Active Health System data to HPE.

During server registration, iLO collects data (including server model, serial number, and iLO NIC address) to uniquely identify the server hardware. When service events are recorded, iLO collects data to uniquely identify the relevant hardware component and to enable proactive advice and consulting.

This data is sent to the Insight Remote Support hosting device (Insight Remote Support central connect) or directly to HPE (Insight Online direct connect). iLO or the Insight Remote Support hosting device sends Active Health System information to HPE every seven days and sends configuration information every 30 days. All data collected and sent to HPE is used to provide remote support and quality improvement.



Note

For more information about the data that is collected, refer to the HPE iLO 5 User Guide, which is available by entering the following URL into your browser:

http://h20564.www2.hp.com/hpsc/doc/public/display?docId=a00026409en_us

Learning check

Which environment is ideal for Remote Support direct connect?

- A. Installations where HPE MicroServer Gen10 are advised
- B. IT data centers and remote sites
- C. Simple deployment and automated support
- D. HPE Converged Infrastructure environments with multiple device types

Write a summary of the key concepts presented in this chapter.

Summary

HPE on-premises management tools are installed and run on computers in an organization's own data center rather than at a remote facility, such as in the cloud. HPE iLO Advanced provides remote management across the data center.

HPE OneView is an on-premises solution that focuses on the entire hardware lifecycle from initial configuration to ongoing monitoring and maintenance. The HPE OneView resource-oriented architecture provides a uniform REST interface. HPE OneView integrates with partner products such as:

- HPE OneView for VMware vCenter
- HPE OneView for Microsoft System Center
- HPE OneView for RHEV

Insight Control server provisioning (ICsp) is a complete provisioning solution for HPE ProLiant servers. It provides server provisioning capabilities designed to help customers to simplify and reduce the time associated with server provisioning tasks. ICsp has capabilities for multi-server operating system and firmware provisioning to rack-mount and BladeSystem servers.

HPE on-cloud solutions remove the need for the user to install or maintain software on-premises. HPE remote management solutions include:

- HPE Insight Online—Is a cloud-based infrastructure management and support portal available through the HPE Support Center. Complementary to HPE OneView, Insight Online provides access to device and support information in a cloud-based personalized dashboard.
- HPE Insight Remote Support—Is embedded in Intelligent Provisioning and optimized to support up to 3500 monitored devices per hosting device. It allows you to register supported servers for HPE remote support either directly with HPE or through an Insight Remote Support hosting device.

11 Practice Test

Introduction

The Minimally Qualified Candidate (MQC) at the Accredited Technical Professional (ATP) level has foundational knowledge of HPE server solutions and underlying technologies, which include CPU, memory, disk, power, networking, and management subsystems. The typical MQC will have worked in the IT industry in a server administrator, architect, or integrator role for 6 to 18 months.

Ideal candidate for this exam

The MQC is able to articulate the HPE Server solutions strategy, including the purpose, benefits, and components; ask appropriate questions and conduct customer site surveys to be able to make judgments on current, high-level customer server needs; generate server configurations; and demonstrate the key features of HPE server solutions. The MQC is able to install and configure server operating systems.

Exam details

The following are details about the exam:

Exam ID: HPE0-S52

Exam Type: Proctored

Number of items: 60

Item types: Input text, input numbers, matching, multiple choice (single-response), multiple choice (multiple-response), point and click

Exam time: 1 hour 30 minutes

Passing score: 66%

Reference material: No online or hard copy reference material will be allowed at the testing site.

HPE0-S52 testing objectives

34%—Plan and design solutions

List information to gather from the customer.

Given a scenario, determine an appropriate tool to size, specify, or build a solution.

Given a customer requirement, determine the appropriate product/solution based on factors such as price, performance, product availability, and workload.

Identify an upgrade path or solution expansion based on the customer's plan for growth.

Given a customer's long-term vs short-term needs, determine a solution.

Describe key HPE differentiators.

Provide a rationale for a product line within a given solution.

Describe industry-standard server, storage, network architectures, and technologies, and their benefits.

Describe available server hardware and software options, their benefits, and use cases.

32%—Install, configure, and set up solutions

Evaluate the customer environment to determine solution suitability.

Analyze the success criteria and proposed solution to determine the gap between them.

Determine valid use cases for a proposed solution.

Describe the necessary tools, steps, and procedures for implementing the design.

19%—Troubleshoot solutions

Given a customer problem, determine an appropriate resolution path.

Determine the appropriate tools for troubleshooting.

15%—Manage, monitor, and maintain solutions

Given a customer situation, identify the appropriate management tool to use.

Describe the steps for performing common management tasks.

Test preparation questions and answers

The following questions will help you measure your understanding of the material presented in this book. Read all of the choices carefully, as there may be more than one correct answer. Choose all correct answers for each question.

Questions

What are the core principles of Composable Infrastructure? (Select three.)

- Ⓐ Unified API
- Ⓑ Differentiated API
- Ⓒ Hardware-defined intelligence
- Ⓓ Software-defined intelligence
- Ⓔ Static resource pools
- Ⓕ Fluid resource pools

Which ProLiant Gen10 technology ensures that a server will never boot with compromised firmware?

- Ⓐ Secure Recovery
- Ⓑ Secure Encryption
- Ⓒ iLO Advanced
- Ⓓ Silicon Root of Trust

A customer wishes to take advantage of the directory services integration capabilities of iLO 5. What is the minimum iLO license necessary to provide this feature?

- Ⓐ No license is necessary
- Ⓑ iLO Standard
- Ⓒ iLO Advanced
- Ⓓ iLO Advanced Premium Security Edition

Which security state is only available with the iLO Advanced Premium Security Edition license?

- Ⓐ Production mode
- Ⓑ HighSecurity mode

- ⋄. FIPS 140-2 mode
- 1. CNSA/Suite B mode

What are the features of HPE Smart Memory? (Select three.)

- 1. 16 banks of memory per rank
- ⋄. Rank multiplication
- ⋄. Maximum data rate of 1866 MT/s
- 1. Encoded rank selection
- ⋄. 1.2-volt operation

How does NVDIMM technology protect data in the event of unexpected power loss?

- 1. It enables recovery from two sequential DRAM failures on memory DIMMs.
- ⋄. It provides reliable RAID fault tolerance to maximize system uptime.
- ⋄. It uses NAND Flash storage to provide a backup copy of DRAM data.
- 1. It uses server memory identified in the BIOS as Persistent Memory.

What does iLO 5 workload matching do?

- 1. It configures BIOS settings to change performance, power, latency, memory, and other variables.
- ⋄. It engages processor turbo boost to cause frequency fluctuations or jitter.
- ⋄. It configures operating system settings to change performance, power, latency, memory, and other variables.
- 1. It eliminates the delay incurred during the mechanical operation of HDDs.

What is the purpose of core boosting?

- 1. It allows processors to attain higher frequencies, and as more CPU cores are activated, the highest turbo frequencies are achieved.
- ⋄. It minimizes processor frequency fluctuation to reduce latency and deliver deterministic and reliable performance.
- ⋄. It automatically configures BIOS settings to optimize the performance of specific workloads.
- 1. It enables higher processor frequencies on active cores, resulting in reduced costs for customers using core-based licensing software.

A customer is experiencing poor performance with their storage latency sensitive

application. They believe that the issue is being caused by their ProLiant Gen10 storage and they would like to upgrade to persistent storage with the lowest possible latency. What type of storage should you recommend?

- 1. 10K RPM HDD
- 2. SSD
- 3. NVDIMM
- 4. L2 cache

A customer has a virtualized server environment and wants to maximize the performance of the network connections to their virtual machines. Which technology should their NICs support to deliver maximum performance?

- 1. Jitter Smoothing
- 2. SR-IOV
- 3. SmartCache
- 4. UEFI

Which management technology do these features describe?

- 1. Provides rapid discovery and remote access features that allow customers to inventory and deploy servers using virtual media
- 2. Uses agentless management to monitor core hardware and related alerts
- 3. Allows users to view the Integrated Management Log and download Active Health System logs
 - a. SUM
 - b. UEFI
 - c. Insight Online
 - d. iLO 5

Which features are provided by HPE Insight Online? (Select three.)

- 1. Track service events and support cases
- 2. View the relationships between logical and physical resources in the data center
- 3. Create templates for automated infrastructure configuration and provisioning
- 4. View device configurations
- 5. Control server power remotely

3. Monitor HPE contracts and warranties

Which server models are included in the ProLiant ML Gen10 family? (Select two.)

- a. ML10
- b. ML30
- c. ML110
- d. ML350
- e. ML380

Which statement most accurately describes the ProLiant ML server family?

- a. Ideal for maximum internal storage and IO flexibility
- b. Density optimized for flexibility and manageability
- c. Ideal for multi-server deployments
- d. Available in rack and cluster models

A customer is looking for a low-cost compute solution that is optimized for small offices, is very quiet and capable of running graphics-related workloads. Which ProLiant Gen10 server should you recommend?

- a. MicroServer
- b. SY480
- c. XL170r
- d. BL460c

A customer requires a server with up to 24 SFF HDDs in the chassis. Which ProLiant server should you recommend?

- a. ML110 Gen10
- b. BL460c Gen10
- c. DL380 Gen10
- d. DL360 Gen10

What is the maximum number of Intel Xeon Scalable processors available in ProLiant DL Gen10 servers?

- a. 1
- b. 2

⋮. 4

1. 8

A customer is looking for a solution that consolidates servers, shared storage, storage switches, backup, deduplication, and WAN optimization in a small footprint. Which server should you recommend?

⋮. MicroServer Gen10

⋮. SimpliVity 380

⋮. Synergy 12000

1. DL580 Gen10

A customer is running an HPC application on a small cluster of commodity x86 servers and is experiencing performance and availability issues. They are looking for a low-cost, scale-out solution that is highly optimized for HPC workloads. Which solution should you recommend?

⋮. SGI 8600

⋮. ProLiant DL580

⋮. SimpliVity 380

1. Apollo 2000

A customer needs flexible drive mapping to enable custom drive allocations to match workloads in their new Apollo 2000 HPC solution. Which enclosure should you recommend?

⋮. r2200

⋮. r2600 24 SFF

⋮. r2600 16 SFF + 8 NVMe

1. r2800

What are components of HPE BladeSystem? (Select three.)

⋮. OA

⋮. UPS

⋮. Composer

1. Server blades

⋮. Interconnects

3. Location discovery engine

What are the options for populating a single c7000 enclosure?

- a. up to four full-height and 16 half-height blades
- b. up to eight full-height or 16 half-height blades
- c. up to eight full-height and 16 half-height blades
- d. up to 16 full-height or eight half-height blades

Which server model is included in the ProLiant BL Gen10 family?

- a. BL420c
- b. BL460c
- c. BL480c
- d. BL660c

What is a FlexNIC?

- a. A flexible LOM that provides bandwidth configuration in 10 Mb increments from 10 Mb to 10 Gb
- b. A physical PCIe function that presents itself to the operating system as a discrete physical NIC
- c. A way of dividing a 10 Gb or 20 Gb port into up to 16 partitions
- d. A virtual NIC contained in a software layer that supports full-speed protocols such as 10 GbE and 8 Gb Fibre Channel simultaneously

Which architecture does HPE Synergy use to consolidate data center network connections?

- a. Master/satellite
- b. Top of rack
- c. Multi-tier
- d. North/South

Which HPE Synergy component provides a single interface for assembling and re-assembling flexible compute, storage, and network fabric resources?

- a. Composer
- b. Frame Link Module
- c. Image Streamer

1. Onboard Administrator

Which Synergy component provides a repository of bootable images that can be accessed by multiple compute modules in a matter of seconds?

- 1. Image Streamer
- 2. Frame Link Module
- 3. Composer
- 4. Frame

What is Intelligent Provisioning used for? (Select two.)

- 1. creating server profiles
- 2. installing operating systems
- 3. updating firmware
- 4. providing a repository of bootable golden images
- 5. adjusting FlexNIC bandwidth

What information must be entered when adding an enclosure to be managed by HPE OneView?

- 1. iLO IP address of first blade in the enclosure
- 2. Primary and secondary OA IP address
- 3. Primary OA IP address or hostname
- 4. Virtual Connect MAC addresses

Which environment is ideal for Insight Remote Support central connect?

- 1. Small- to mid-sized businesses
- 2. Remote sites with Gen9 and Gen10 servers
- 3. Customers who do not want a centralized hosting device in the local environment
- 4. HPE Converged Infrastructure environments with multiple device types

Answers

1. **A**, **D**, and **F** are correct. The three core principles of Composable Infrastructure are unified API, software-defined intelligence, and fluid

resource pools.

- ⊗ **B, C, and E** are incorrect. Differentiated API, hardware-defined intelligence, and static resource pools are not core principles of Composable Infrastructure.

For more information, see [Chapter 1: Transforming IT Strategy](#).

2. **D** is correct. In ProLiant Gen10 servers, an immutable silicon-based fingerprint is installed in the server, thus ensuring the server will never boot with compromised firmware. This is known as the Silicon Root of Trust.

- ⊗ **A, B, and C** are incorrect. Secure Recovery initiates recovery from a known good version of the firmware in the event of the iLO firmware validation failure. Secure Encryption protects data at rest on bulk storage hard drives and SSDs attached to a compatible HPE Smart Array controller. iLO Advanced is a license, not a technology.

For more information, see [Chapter 2: Security](#).

3. **C** is correct. The iLO Advanced license offers advanced security features, including directory integration.

- ⊗ **A, B, and D** are incorrect. With no license, or the iLO Standard license installed, directoryservices integration is not available. The iLO Advanced Premium Security Edition includes all of the functionality of iLO Advanced, but it is not the minimum license to provide directory services.

For more information, see [Chapter 2: Security](#).

4. **D** is correct. With the iLO Advanced Premium Security Edition license, customers that need the highest-level encryption capabilities have a fourth security state available to them: CNSA/SuiteB.

- ⊗ **A, B, and C** are incorrect. The capabilities of HPE iLO Standard that comes with every ProLiant Gen10 server gives customers the ability to configure their ProLiant Gen10 servers in one of three security states: production, high security, or FIPS.

For more information, see [Chapter 2: Security](#).

5. **A, D, and E** are correct. DDR4 memory has 16 banks of memory in a DRAM chip compared to the eight banks in DDR3, and this helps to reduce memory latency. When more than eight ranks are installed, DDR4 uses a four-bit encoded chip select value for rank selection allowing DDR4 memory to address up to 24 memory ranks on a memory channel. All DDR4

memory operates at 1.2 volts, compared to 1.35 or 1.5 volts for DDR3 memory, leading to significant system power savings.

- ☒ **B** and **C** are incorrect. Rank multiplication is used by DDR3, not by DDR4 memory. DDR4 memory operates at up to 2166 MT/s bandwidth.

For more information, see [Chapter 3: HPE Server Technologies and Tools](#).

6. ☑ **C** is correct. NVDIMM technology uses NAND Flash storage to provide a backup copy of DRAM data.

- ☒ **A**, **B**, and **D** are incorrect. Adaptive Double Device Data Correction (ADDDC) enables recovery from two sequential DRAM failures on memory DIMMs. HPE Smart Array controllers provide reliable RAID fault tolerance for ProLiant storage to maximize system uptime. HPE Scalable Persistent Memory writes data to actual server memory marked as Persistent Memory.

For more information, see [Chapter 3: HPE Server Technologies and Tools](#).

7. ☑ **A** is correct. Workload matching configures BIOS settings to change performance, power, latency, memory, and other variables.

- ☒ **B**, **C**, and **D** are incorrect. Workload matching does not engage turbo boost to cause jitter, configure operating system settings, or eliminate the delay incurred during the mechanical operation of HDDs.

For more information, see [Chapter 3: HPE Server Technologies and Tools](#).

8. ☑ **D** is correct. Core boosting enables higher processor frequencies on active cores, resulting in reduced costs for customers using core-based licensing software.

- ☒ **A**, **B**, and **C** are incorrect. Core boosting does not allow the highest turbo frequencies to be achieved as more cores are activated, minimize processor frequency fluctuation (jitter smoothing does this), or automatically configure BIOS settings (workload matching does this).

For more information, see [Chapter 3: HPE Server Technologies and Tools](#).

9. ☑ **C** is correct. NVDIMM is the highest performing Persistent Memory on the market running at DRAM speeds.

- ☒ **A**, **B**, and **D** are incorrect. 10K RPM HDD and SSD do not perform as well as NVDIMM. L2 cache is faster than NVDIMM, but is not persistent.

For more information, see [Chapter 3: HPE Server Technologies and Tools](#).

10. **B** is correct. SR-IOV increases network performance via direct access to hardware in virtual environments.
- A**, **C**, and **D** are incorrect. Jitter smoothing mitigates processor frequency fluctuation and does not directly impact network performance. SmartCache accelerates access to data on HDD. UEFI is an on-system management technology and does not directly affect network performance in virtualized environments.

For more information, see [Chapter 3: HPE Server Technologies and Tools](#).

11. **D** is correct. iLO 5 provides rapid discovery and remote access features that allow customers to inventory and deploy servers using virtual media, uses agentless management to monitor core hardware and related alerts, and allows users to view the Integrated Management Log and download Active Health System logs.
- A**, **B**, and **C** are incorrect. SUM, UEFI, and Insight Online do not provide all of these features.

For more information, see [Chapter 3: HPE Server Technologies and Tools](#).

12. **A**, **D**, and **F** are correct. Insight Online allows you to track service events and support cases, view device configurations, and monitor HPE contracts and warranties.
- B**, **C**, and **E** are incorrect. Insight Online does not allow you to view the relationships between logical and physical resources in the data center or create templates for automated infrastructure configuration and provisioning (HPE OneView does this). Insight Online does not allow you to control server power remotely (iLO 5 does this).

For more information, see [Chapter 3: HPE Server Technologies and Tools](#).

13. **C** and **D** are correct. The ProLiant ML Gen10 family includes ML110 and ML350 server models.
- A**, **B**, and **E** are incorrect. The ProLiant ML Gen10 family does not include ML10, ML30, or ML380 server models.

For more information, see [Chapter 4: HPE ProLiant Tower Servers](#).

14. **A** is correct. The ProLiant ML server family is ideal for maximum internal storage and IO flexibility.
- B**, **C**, and **D** are incorrect. The ProLiant ML server family is not density optimized, ideal for multi-server deployments, or available in rack and

cluster models.

For more information, see [Chapter 4: HPE ProLiant Tower Servers](#).

15. **A** is correct. The ProLiant MicroServer Gen10 is a low-cost compute solution optimized for small office, home office, and enabled for graphics-related workloads.
- B**, **C**, and **D** are incorrect. SY480 is a Synergy compute module, and requires the Synergy 12000 frame and management infrastructure. XL170r is an Apollo compute server, and requires the Apollo r2200, r2600, or r2800 enclosure. BL460c is a Blade server and requires the BladeSystem c7000 chassis.

For more information, see [Chapter 4: HPE ProLiant Tower Servers](#).

16. **C** is correct. The ProLiant DL380 Gen10 supports a maximum of 24 SFF drives in the chassis.
- A**, **B**, and **D** are incorrect. The ProLiant ML110 Gen10 supports a maximum of 16 SFF drives in the chassis. The BL460c supports a maximum of two SFF drives. The ProLiant DL360 Gen10 supports a maximum of 10 drives in the chassis.

For more information, see [Chapter 5: HPE Rack Server Solutions](#).

17. **C** is correct. The DL580 Gen10 server supports a maximum of four Intel Xeon Scalable processors.
- A**, **B**, and **D** are incorrect. The maximum number of Intel Xeon Scalable processors available in ProLiant DL servers is four.

For more information, see [Chapter 5: HPE Rack Server Solutions](#).

18. **B** is correct. The SimpliVity 380 server consolidates servers, shared storage, storage switches, backup, deduplication, and WAN optimization in a small footprint.
- A**, **C**, and **D** are incorrect. The MicroServer Gen10, Synergy 12000, and DL580 Gen10 do not consolidate servers, shared storage, storage switches, backup, deduplication, and WAN optimization in a small footprint.

For more information, see [Chapter 5: HPE Rack Server Solutions](#).

19. **D** is correct. Apollo 2000 is the ideal solution for customers who are embarking on a scale-out journey and need a solution for small deployments of HPC.

- A, B, and C** are incorrect. SGI 8600 is a liquid-cooled supercomputer that can grow to tens of thousands of nodes. ProLiant DL580 and SimpliVity 380 are not highly optimized for scale-out HPC workloads.
For more information, see [Chapter 6: HPE Performance Computing Solutions](#).
- 20. **D** is correct. The r2800 enclosure provides control over the number of hard drives that are allocated to each server, enabling custom drive allocations to match workloads.
 - A, B, and C** are incorrect. The r2200 and r2600 enclosures do not support custom drive allocations. Drives are allocated equally across server nodes.
For more information, see [Chapter 6: HPE Performance Computing Solutions](#).
- 21. **A, D, and E** are correct. OA, server blades, and interconnects are components of HPE BladeSystem.
 - B, C, and F** are incorrect. UPS, Composer, and location discovery engine are not components of HPE BladeSystem.
For more information, see [Chapter 7: HPE BladeSystem Solutions](#).
- 22. **B** is correct. A single BladeSystem enclosure can be populated with up to eight full-height or 16 half-height blades.
 - A, C, and D** are incorrect. A single BladeSystem enclosure can be populated with up to eight full-height or 16 half-height blades.
For more information, see [Chapter 7: HPE BladeSystem Solutions](#).
- 23. **B** is correct. The BL460c is the only member of the ProLiant BL Gen10 family.
 - A, C, and D** are incorrect. There are no BL420c, BL480c, or BL660c Gen10 servers.
For more information, see [Chapter 7: HPE BladeSystem Solutions](#).
- 24. **B** is correct. A FlexNIC is a physical PCIe function that presents itself to the operating system as a discrete physical NIC.
 - A, C, and D** are incorrect. A FlexNIC is not a flexible LOM that provides bandwidth configuration in 10 Mb increments, a way of dividing a 10 Gb or 20 Gb port into up to 16 partitions, or a virtual NIC contained in a software layer.

For more information, see [Chapter 7: HPE BladeSystem Solutions](#).

25. **A** is correct. HPE Synergy uses a master/satellite architecture to consolidate data center network connections.
- B**, **C**, and **D** are incorrect. HPE Synergy does not use top of rack, multi-tier, or North/South architectures to consolidate data center network connections.

For more information, see [Chapter 8: HPE Synergy](#).

26. **A** is correct. The HPE Synergy Composer provides a single interface for assembling and re-assembling flexible compute, storage, and fabric resources.
- B**, **C**, and **D** are incorrect. Synergy Frame Link Module and Image Streamer do not provide a single interface for assembling and re-assembling flexible compute, storage, and network fabric resources. Onboard Administrator is a BladeSystem, not a Synergy component.

For more information, see [Chapter 8: HPE Synergy](#).

27. **A** is correct. The Synergy Image Streamer is a physical appliance repository of bootable (golden) images that can be streamed across multiple compute modules in a matter of seconds.
- B**, **C**, and **D** are incorrect. The Synergy Frame Link Module, Composer, and Frame do not provide a repository of bootable images that can be accessed by multiple compute modules in a matter of seconds.

For more information, see [Chapter 8: HPE Synergy](#).

28. **B** and **C** are correct. Intelligent Provisioning is used for installing operating systems and performing maintenance-related tasks such as updating firmware.
- A**, **D**, and **E** are incorrect. Intelligent Provisioning does not create server profiles, provide a repository of bootable images, or adjust FlexNIC bandwidth.

For more information, see [Chapter 9: HPE On-system Management](#).

29. **C** is correct. The IP address or hostname of the primary OA must be entered when adding an enclosure to OneView.
- A**, **B**, and **D** are incorrect. Adding an enclosure to OneView does not require the iLO IP address of first blade in the enclosure, secondary OA IP

address, or Virtual Connect MAC addresses to be entered.

For more information, see [Chapter 10: HPE On-premises and On-cloud Management](#).

30. **D** is correct. Insight Remote Support central connect is ideal for HPE Converged Infrastructure environments with multiple device types.
- A**, **B**, and **C** are incorrect. Insight Remote Support central connect is not ideal for SMB, remote sites, or customers who do not want a centralized hosting device in the local environment

For more information, see [Chapter 10: HPE On-premises and On-cloud Management](#).

Learning Check Answers

Chapter 1

A customer is considering migrating their data center to a Cloud Service Provider. They are concerned about the risks involved and would rather have the ability to host some of their application on-premises and some in the cloud. How should you respond?

The world is going to be hybrid and HPE makes hybrid IT simple. This is achieved through offerings in the traditional data center, software-defined infrastructure, systems software, private cloud, and through public cloud partnerships.

Which concept is central to the operation of HPE Composable Infrastructure?

- A. A single unified API allows operators and developers to access IT resources.**
- B. Multiple differentiated APIs allow operators and developers to access IT resources.
- C. A single portal allows operators and developers to have seamless access to systems management functions such as iLO 5.
- D. Multiple differentiated portals allow operators and developers to have seamless access to systems management functions such as iLO 5.

Which of the following represent customer requirements? (Select four.)

- A. Velocity**
- B. Legacy
- C. Efficiency**
- D. Flexibility**
- E. Dynamic
- F. Portability
- G. Security**

Chapter 2

You are attending a meeting with a customer to discuss their future data center modernization plans. The subject of security is high on their list of concerns. While they are comfortable with their progress in protecting against application-layer attack, they are aware of the possibility of attacks that are directed at the hardware level. They raise the question about what HPE offers in the Gen10 platform around protecting against firmware-level attack. How should you respond?

HPE has achieved the concept of a Silicon Root of Trust. The fact that this feature is anchored in the silicon is unique in the industry, giving HPE an impenetrable protection right through the entire supply chain. Therefore, at all stages of the manufacturing process, authenticity, and integrity is assured for the server platform.

What is the name that HPE has given to the process that includes the “Silicon Root of Trust”?

- A. Fast Start
- B. New Start
- C. **Secure Start**
- D. Up Start
- E. Valid Start

Which iLO license is required to support “Runtime Firmware Validation”?

- A. iLO Advanced license
- B. iLO Standard license
- C. **iLO Advanced Premium Security Edition license**
- D. iLO Foundation license

A customer is hosting solutions on behalf of other organizations and is interested in providing rack and server physical access security. For cost reasons, they would rather not have to build secure cages to provide this level of security. How should you respond?

ProLiant Gen10 servers include an option for a chassis Intrusion Detection switch, which detects if the chassis access cover is opened or closed. HPE provides enhanced security features to ensure only authorized personnel have access to owned racks and infrastructure. These features include: CAC, three-factor

authentication, and biometric locks.

Chapter 3

A customer is planning to implement a dense Virtual Machine solution where large memory capacity is key to success. Your customer asks for confirmation that Gen10 servers can provide the capacities they need. How should you respond?

Gen10 capacity increase over the previous generation is achieved as a result of Intel® Skylake supporting six memory channels plus support for 128 GB DIMMs. Potential maximum memory capacity of 1.5 TB per processor socket.

You are attending a meeting with a customer to discuss their future data center modernization plans. The conversation turns to Gen10 performance enhancements. You are asked to explain how HPE IST may provide benefit for their VDI application. How should you respond?

HPE ProLiant Gen10 servers are capable of supporting a very wide range of workloads. To achieve this degree of flexibility requires accurate selection of BIOS configuration parameters with the aim to provide the best match to workload. HPE Jitter Smoothing technology mitigates processor frequency fluctuation to reduce latency and deliver deterministic and reliable performance.

What are functions of the HPE iPDU? Select two.

- A. **Determine the exact power consumption for every component in the rack.**
- B. Support up to 24 devices through a battery backup unit.
- C. **Support for Location Discovery Services.**
- D. Use NAND Flash memory technology to consume less memory.
- E. Integrate seamlessly with ProLiant servers via the PCIe bus.

Which HPE standalone application enables you to determine the power consumption and total weight of a configuration?

- A. Server Memory Configurator
- B. Converged Infrastructure Sizing Suite
- C. **Power Advisor**
- D. HPE Product Bulletin

Chapter 4

Which description is most accurate regarding the ProLiant ML server family?

- A. **Ideal for maximum internal storage and IO flexibility**
- B. Density-optimized for flexibility and manageability
- C. Ideal for multiserver deployments
- D. Available in rack and cluster models
- E. Ideal for large data center deployments

You are attending a meeting with a customer to discuss their remote office server requirements. Four out of the 10 offices have fewer than 10 employees. They need a low-cost server to provide local print services and a network gateway with simple to use management. How should you respond?

The HPE ProLiant MicroServer offers a unique form factor that is compact and great for small spaces. ClearOS preloaded gives customers an out-of-box ready SMB solution for cloud, gateway, network, and server capability.

A customer approaches your stand at a trade fair and asks you to suggest the best server solution to provide a low cost, highly available server platform for a stand-alone virtual machine environment. How should you respond?

The ML350 Gen10 Server is ideal for enterprise IT infrastructure to mission-critical applications. With the HPE ProLiant ML350 Gen10 Server, you can deploy a single platform to handle a wide variety of enterprise workloads.

Chapter 5

A customer has expressed a need for a secure two-socket ProLiant rack server to host one of their low-density virtualized applications. Which server should you recommend?

- ⓐ. DL380 Gen10
- ⓑ. DL180 Gen9
- ⓒ. DL560 Gen10
- ⓓ. DL360 Gen10**

You are attending a meeting with a customer to discuss their future data center modernization plans. During the discussion they raise the subject of hyperconverged solutions and that they are considering Microsoft Storage Spaces Direct. When they considered the HPE DL380 Rack Server solution, they were concerned that the platform may not offer sufficient HDD and SSD storage to meet their needs. How should you respond?

The HPE ProLiant DL380 offers a massive storage expansion capability, supporting SAS SSD/HDD as well as NVMe drives. Depending on the chosen options the DL380 can support up to 30 drives.

A customer has expressed a need for a four-socket ProLiant rack server to host one of their business-critical applications? Which server should you recommend?

- ⓐ. DL380 Gen10
- ⓑ. DL180 Gen9
- ⓒ. DL560 Gen10**
- ⓓ. DL360 Gen9

You are at a customer meeting and one of the IT team comments that there is a belief that hyperconverged solutions performance is poor due to the CPU bound deduplication and compression functionality. They are considering the HPE SimpliVity 380 solution. How would you address their performance concerns?

The HPE SimpliVity 380 solution offloads the task of performing deduplication and compression to the Tiered Index Accelerator (TIA). The TIA performs deduplication and compression on primary data, in real-time, without impact to performance or latency.

Chapter 6

In the HPE Apollo r2600 24 SFF chassis, how are the 24 disks allocated to each of the XL170r and the XL190r Servers?

The disks are allocated equally across servers.

You are attending a meeting with the IT procurement manager of a large research university. Attending the meeting is the head of a small department performing medical research for the next generation of antibiotics. They already have an HPE Apollo 2000 solution, and they are looking to upgrade from Gen9 servers and scale to a larger solution. How should you respond?

The Apollo 2000 System is a density-optimized, 2U shared infrastructure chassis for up to four ProLiant Gen9 or Gen10 independent, hot-plug servers. A 42U rack fits up to 20 Apollo r2000 series chassis accommodating up to 80 servers per rack.

An HPE ProLiant XL170r provides two IO slots for a choice of networking and clustering options which includes Ethernet 1 Gb, 10 Gb, and 100 Gb. What are the alternatives? (Select two.)

- A. 100 Gb/s EDR InfiniBand
- B. Myrinet cluster interconnect
- C. Quadrics QsNet network
- D. Scalable Coherent Interconnect
- E. 100Gb/s Omni-Path

At a trade fair, a customer approaches you to explain that they are a small startup company doing research in developing organic battery technologies for the motor car industry. They have a limited budget to invest in the next generation of HPC. They are interested in technologies supplied by HPE. How should you respond?

HPE Apollo 10 series makes HPC and DL affordable. Low entry costs combine with a flexible industry-standard platform, custom solutions, and a choice of components. Apollo 10 series servers readily integrate into standard cluster environments.

Chapter 7

What is the maximum number of server blades that can populate a single c7000 enclosure?

- A. Up to four full-height and 16 half-height blades
- B. Up to eight full-height or 16 half-height blades**
- C. Up to eight full-height and 16 half-height blades
- D. Up to 16 full-height and eight half-height blades

Name components of HPE BladeSystem. (Select three.)

- A. Uninterruptable Power Supply
- B. OA**
- C. Insight Control server provisioning
- D. Server blades**
- E. Interconnect modules**
- F. Location Discovery Engine

A conversation with a customer reveals that they are considering an upgrade to their BladeSystem solution they had installed eight years ago. The data center edge has been upgraded to 20 Gb Ethernet and they are investing in a 16 Gb FC SAN. They ask you if they can upgrade the BladeSystem to support these new technologies. How should you respond?

You need to ask questions about their existing BladeSystem infrastructure. They may need to invest in a Platinum c7000 enclosure before they consider upgrading to higher bandwidth interconnects and newer Gen10 blades.

What is a FlexNIC?

- A. A flexible LOM that provides RoCE blades
- B. A physical PCIe function that presents itself to the operating system as a discrete physical NIC**
- C. A way of dividing a 10 Gb or 20 Gb port into up to eight partitions
- D. A virtual NIC contained in a software layer

A customer is rolling out a production environment that will consist of up to 30 c7000 enclosures. They are looking for a central management solution that will allow them to

perform template-driven server deployment and enclosure lifecycle management. How should you respond?

BladeSystem and HPE OneView form a federated system. This means that each BladeSystem chassis becomes part of a single management environment. Any workload can be moved to any chassis as long as the receiving blade has the right configuration.

Chapter 8

What are the three key elements of a composable architecture?

Fluid resource pools, software-defined intelligence, and a third element that uses all of these capabilities and exposes them through a unified API, which allows infrastructure to be programmed like code so it can become infrastructure as a service.

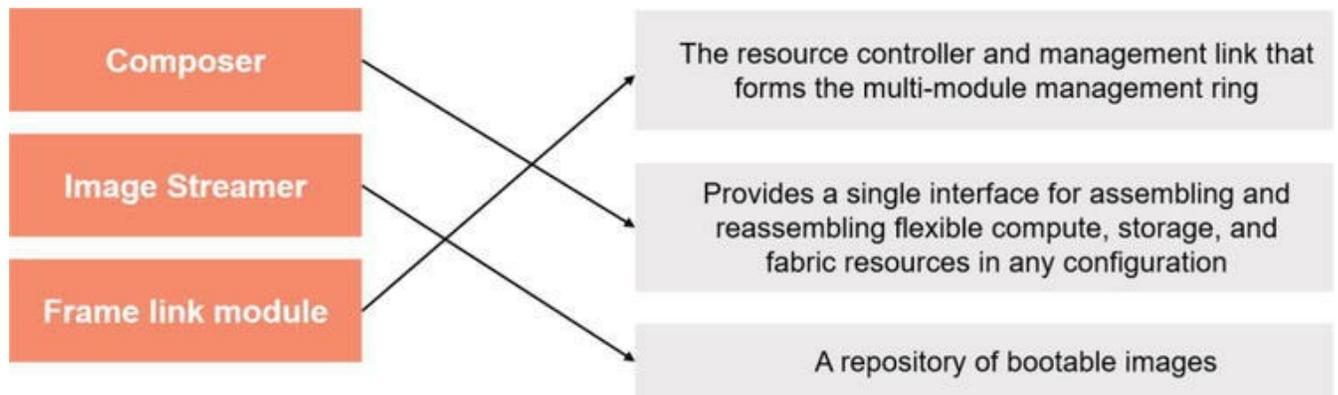
A customer is evaluating a competitor’s solution and has raised concerns that Synergy does not appear to have a solution for the rapid deployment of the host operating system to the Synergy compute nodes. How should you respond?

The Synergy Image Streamer represents a new approach to deployment and updates for Composable Infrastructure. This management appliance works with HPE Synergy Composer to provide fast, software-defined control over physical compute modules with operating system and application provisioning. Image Streamer enables truly stateless computing combined with instant-on capability for quick deployment and updates.

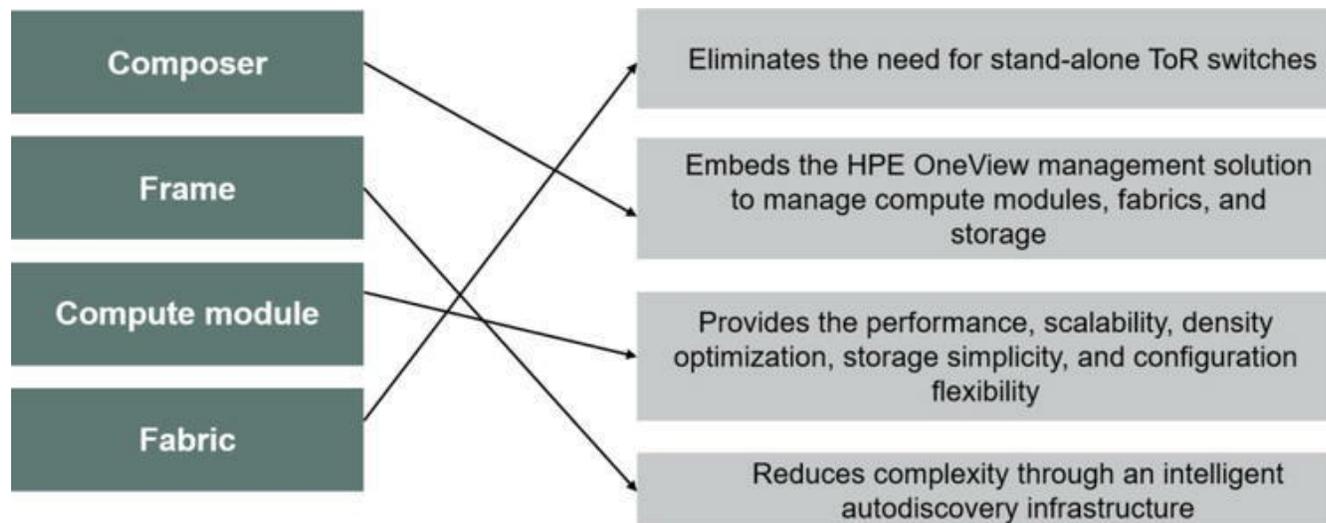
A client is evaluating the HPE Synergy solution as a replacement for their multi-stacked c7000 enclosure solution. Among many questions about the Synergy product, they specifically are interested in the management solution redundancy capability. How should you respond?

The HPE Synergy management solution is based around a redundant pair of Composer modules, redundant FLMs, and the management ring. Composer embeds the HPE OneView management solution to manage compute modules, fabrics, and storage, which is the essence of software-defined intelligence in Synergy.

Match the Synergy management subsystem component with its function.



Match the Synergy component with its function.



Chapter 9

What must you do if you have only one processor installed in a two-socket server?

A. Install fan blanks

- 3. Install a blanking panel in the second socket
- C. Balance the DIMMs across both processor sockets
- D. Designate one socket as active and the blank as passive

You are designing a solution for a client that will involve installing servers in a rack. The client needs local access to iLO but due to power availability problems in the data center the support staff cannot use KVM. Which Gen10 feature would be useful to satisfy this need?

The HPE ProLiant Gen10 Server provides a USB port accessible from the front of the server. This is a full-service iLO console port. You just need a USB to Ethernet converter then connect a laptop. Then browse to the default IP address 169.254.1.2.

Which tool should you use to configure the boot options on a ProLiant Gen10 server?

- A. HPE OneView
- 3. RBSU
- C. Secure Boot mode
- D. BIOS/Platform Configuration menu**

A customer approaches your stand at a trade fair and informs you that they are looking to upgrade their G7 servers to Gen10 servers. They ask if they still need to use the Smart Start CD to deploy Windows? How should you respond?

A Smart Start CD is no longer required since Gen8 servers. It has been replaced by Intelligent Provisioning which is embedded in the flash ROM. To run this deployment environment you simply press the F10 key during startup and wait for Intelligent Provisioning to load.

Chapter 10

You are visiting a customer who tells you that they have recruited a new assistant system administrator who is tasked with streamlining their infrastructure management. They currently rely on Insight Manager, but they want to use scripting to perform server deployment. What should you advise?

They need to evaluate their infrastructure to determine if HPE OneView would be the suitable solution to monitor their infrastructure. If their infrastructure is up to date and compatible, they could implement OneView to monitor their estate in parallel with Insight Manager while they develop scripts to interface with the OneView REST APIs.

What is included in an HPE OneView server profile? (Select two.)

- A. Software license keys
- B. iLO settings**
- C. Operating system images
- D. HPE StoreVirtual iSCSI storage volumes**
- E. Error logs

Which process does HPE OneView automate when a profile is applied to the server?

- A. Operating system deployment
- B. Monitor data center infrastructure
- C. Assign MAC and WWID numbers**
- D. Assign server hostname

Which environment is ideal for Remote Support direct connect?

- A. Installations where HPE MicroServer Gen10 are advised
- B. IT data centers and remote sites
- C. Simple deployment and automated support**
- D. HPE Converged Infrastructure environments with multiple device types