



HPE ATP – Hybrid IT Solutions V2

Official Certification Study Guide (Exam HPE0-V14)

HPE Partner Ready Certification and Learning

Master all HPE0-V14 exam topics

Assess your knowledge with learning checks

Understand key concepts and objectives



Hewlett Packard
Enterprise

HPE ATP – Hybrid IT Solutions V2
OFFICIAL CERTIFICATION STUDY GUIDE
(EXAM HPE0-V14)

First Edition

Radek Zima

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

HPE ATP – Hybrid IT Solutions V2
Official Certification Study Guide (Exam HPE0-V14)

Radek Zima

© 2019 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: 978-1-7331277-5-2

WARNING AND DISCLAIMER

This book provides information about the topics covered in the HPE ATP – Hybrid IT Solutions V2 certification exam (HPE0-V14). 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.

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: Chris Powell, Chris Bradley

HPE Press Program Manager: Michael Bishop

About the Author

Radek Zima is an independent consultant who specializes in IT infrastructure design, implementation and maintenance for HPE servers, storage, networking, management and cloud software. He develops and delivers trainings, workshops, demonstrations, and conference presentations for HPE channel partners, customers and employees at training centers and events around the world. Radek has a Bachelor's degree and a Master of Science degree from the Faculty of Informatics and Statistics, University of Economics in Prague.

Introduction

This book is based on the Building HPE Hybrid IT Solutions, Rev. 19.41 course, which is designed to expose participants to the fundamental principles required to architect Hybrid IT solutions. The guide helps readers prepare for the Building HPE Hybrid IT Solutions (HPE0-V14) exam, which focuses on HPE networking, server, and storage solutions for SMB customers. The HPE ATP Hybrid IT V2 certification validates that a successful candidate has the knowledge and skills necessary to plan, design, recommend, and demonstrate HPE Hybrid IT solutions and deliver a proof-of-concept for a solution.

Areas of study include the ability to:

- Describe, differentiate, and apply industry standard, foundational SMB architectures and technologies.
- Differentiate the functions, features, and capabilities of HPE product and solution specific components and offerings.
- Recommend and position HPE SMB products, solutions, and appropriate services for customer use cases (positioning/use cases).
- Evaluate customer environment, and plan and design solutions using the HPE SMB portfolio to meet customer business requirements.
- Validate, install, configure, and upgrade HPE SMB solutions and their components.
- Troubleshoot, diagnose, and replace HPE SMB solution components.

- Manage, monitor, administer, and operate HPE SMB solution components.

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 accelerate your career.

We provide

- **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;
- **Continuous learning activities and job-role based learning plans** to help you keep pace with the demands of the dynamic, fast paced IT industry; and
- **Advanced training** to help you navigate and seize opportunities within the top IT transformation areas that enable business advantage today.

As a Partner Ready Certification and Learning certified member, your skills, knowledge, and realworld 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, trendmakers 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 book is designed for presales solution architects involved in supporting the sale of HPE SMB hybrid IT solutions, encompassing servers, storage, networking, and management. Typical candidates for the HPE ATP – Hybrid IT Solutions V2 certification are IT professionals who work in and around the data center and who have the responsibility to achieve and improve the

availability and manageability of the data center. Typical candidate job roles include but are not limited to presales architects, presales engineers, enterprise architects, solutions engineers, and technology architects.

Assumed Knowledge

HPE ATP Hybrid IT Solutions V2 Study Guide is an entry-level book. However, before reading this book, you should be familiar with HPE server and storage platforms as well as basic networking technologies such as IP addressing, broadcast domains, and the Open Systems Interconnection (OSI) model. It is assumed that you have an understanding of HPE ProLiant servers, HPE MSA storage systems, and HPE FlexFabric networking solutions. You should also have an understanding of HPE data center technology and an interest in learning about HPE StoreEasy, StoreOnce, 3PAR, and Nimble storage systems as well as Aruba Network devices to help you understand customers' business issues and to propose appropriate solutions.

Minimum Qualifications

To pass the Building HPE Hybrid IT Solutions (HPE0-V14) exam, you should have a minimum of 12 months hands-on experience or the equivalent in at least one of the core HPE areas (server, storage, and networking) and six months hands-on experience or the equivalent in other HPE SMB solutions and technologies.

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 Exam HPE0-V14

This self-study guide does not guarantee that you will have all the knowledge

you need to pass the exam. It is expected that you will also draw on real-world experience and would benefit from completing the hands-on lab activities provided in the instructor-led training.

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 https://certification-learning.hpe.com/tr/learn_more_about_exams.html

CONTENTS

1 Introduction to a Hybrid World

SMB Hybrid IT for Dummies

HPE Small Business Solutions

Small businesses today: Head in the cloud—feet on the ground

Are your customers trying to find the right balance?

HPE Small Business Solutions

What business outcomes are top-of-mind for your customers?

Start with complete, validated hardware and software stacks

Enable hybrid cloud by adding on to on-premises solutions

Learning check

HPE overview

Innovation is in our DNA

Creating a world where everything computes

Driven by a new generation of apps and data—and a new speed of business

The edge has arrived

Transform your customer's technology

Transform your customer's people and processes

Transform your customer's economics

Accelerating our strategy with partnerships

Learning check

HPE Hybrid IT

Hybrid IT is the new normal—and it is complex

How we make Hybrid IT simple

Define your customer's Hybrid IT strategy

Power your customer's Hybrid IT platforms

Start with a foundation that is fast, always-on, secure—ready for anything

Optimize your customer's Hybrid IT financing and delivery

Activity: Navigating the SMB Hybrid IT portfolio at the HPE website

Activity: Examples of a hybrid cloud solution

Learning check

Summary

Prelearning check

2 Recommending HPE compute solutions for SMB customers

HPE has it all

Why a single partner for compute, storage, and networking?

Assessing customer requirements

Discovering a customer's business and technical requirements

Issues to consider

Assessing the customer's applications

Solution design considerations

Learning check

Recommending HPE compute solutions for SMB customers

Introducing the customer scenario

Activity: Discovering a customer's business and technical requirements

Selecting an HPE server platform

Selecting a Gen10 DL or Gen10 ML server

HPE ProLiant tower servers

Selecting an HPE ML server platform

ProLiant ML110 Gen10 Server

Front system detail: LFF and SFF chassis shown

Interior system detail

Rear system detail

HPE StoreEver LTO Ultrium Tape Drives

Secure: Hardware-based data encryption

Activity: HPE QuickSpecs: HPE ProLiant servers

Activity: HPE Server Support and Certification Matrices

Learning check

Components of a compute solution

Processors in HPE servers

Unified Intel® Xeon® Scalable Platform

Intel® Xeon® Scalable Family processor numbering and features

Select the Intel® Xeon® processor

AMD EPYC 7000 Processors—A new building block

Learning check

Memory for HPE servers

Memory/Storage hierarchy

Gen9 vs. Gen10 memory configuration (Intel®-based servers)

Gen10 memory configuration (AMD EPYC-based servers)

HPE Persistent Memory Kit featuring Intel® Optane™ DC Persistent Memory

Activity: Server memory and Persistent Memory population rule

Learning check

Storage for HPE servers

Drive array basics

Hardware-based RAID

Activity: Features of HPE Smart Array Controller

Learning check

Networking features of ProLiant servers

Secure networking flexibility at the speed of compute

Activity: HPE QuickSpecs: Server components

Learning check

Rack and power infrastructure

HPE Rack and Power Infrastructure portfolio

G2 rack product portfolio

G2 Power Distribution Unit product portfolio

Elevated temperature support

HPE Adaptive Rack Cooling System

Installing solutions and subsystems to the rack

Installing server options

Power requirements

Cabling requirements

Learning check

Alternative products

HPE ProLiant ML350 Gen10 Server

[HPE ProLiant MicroServer Gen10—The compact server](#)

[HPE ProLiant ML30 Gen10 Server](#)

[Selecting an HPE DL server platform](#)

[HPE ProLiant Gen10 rack servers](#)

[ProLiant Gen10 rack server portfolio](#)

[HPE ProLiant DL20 Gen10 Server](#)

[HPE ProLiant DL325 Gen10 Server](#)

[HPE ProLiant DL360 Gen10 Server](#)

[HPE ProLiant DL380 Gen10 Server](#)

[HPE ProLiant DL385 Gen10 Server](#)

[HPE Apollo 2000 Gen10 System](#)

[HPE Synergy: Powering your customer's hybrid cloud transformation](#)

[HPE Composable Cloud for ProLiant DL](#)

[HPE Composable Cloud for ProLiant DL delivers business value](#)

[HPE ProLiant for Microsoft Azure Stack](#)

[Learning check](#)

[Configuring a single server solution](#)

[HPE One Config Simple](#)

[HPE OneConfig Advanced](#)

[HPE Power Advisor](#)

[HPE Server Memory Configurator](#)

[SSD Selector Tool](#)

[SSD Availability Matrix](#)

[Learning check](#)

[Selecting HPE Support Services](#)

[Building on our heritage of services leadership, HPE Pointnext:](#)

[HPE Operational Support Services](#)

[One partner to help business succeed through IT](#)

[HPE Foundation Care](#)

[HPE Proactive Care](#)

[Learning check](#)

[Managing a single server solution](#)

[HPE server management](#)

[HPE iLO 5 management technologies](#)

- Server boot process
- Intelligent Provisioning
- Smart Storage Administrator
- HPE Smart Update solution stack
 - HPE Smart Update Tools
 - HPE Smart Update Manager
 - HPE Service Pack for ProLiant
 - HPE Firmware update technology
- Learning check
- Customer scenario
 - Introducing the customer scenario
 - Customer requirements
- Summary
- Prelearning check

3 Recommending HPE networking solutions for SMB customers

- Recommending HPE networking solutions for an SMB
 - Introducing the customer scenario
- Activity: Discovering a customer's business and technical requirements
 - Selecting an HPE Networking platform
 - Selecting an HPE switch: Selection criteria
 - Selecting an HPE switch: FlexNetwork
 - Selecting an HPE switch: Aruba and FlexFabric
 - Selecting an HPE switch: Port speed—Downlinks
 - Selecting an HPE switch: Port speed—Uplinks
 - Selecting an HPE switch: Port count
 - HPE Switch Selector
 - HPE Networking Online Configurator
- Activity: Switch Selector for Data Center, Campus, and Small Business
- Learning check
- Networking architectures
 - Tiered network configurations
 - Leaf-spine topology
 - Data center architectures

HPE Intelligent Resilient Framework

Example of LACP configuration

Switch infrastructure recommendations—iSCSI SAN

Learning check

ArubaOS-Switch platform

Aruba switching portfolio: Campus and data center

Aruba 2530 and 2540 Switch Series

Aruba 2930F and 2930M Switch Series

Aruba 3810 and 5400R Switch Series

Learning check

ArubaOS-CX platform

Aruba 8400 Switch Series

Aruba 8320 Switch Series

Aruba 8325 Switch Series

Learning check

HPE FlexFabric

HPE FlexFabric 5710 Switch Series

HPE FlexFabric 5940 Switch Series

HPE FlexFabric 5945 Switch Series

Learning check

HPE Networking management

Command-line interface

HPE Networking software

Unified network management with HPE Intelligent Management Center

IMC for the data center

HPE IMC base platform options

Aruba NetEdit

Learning check

Customer scenario

Customer requirements

Summary

Prelearning check

4 Recommending HPE NAS Solutions for Customer use Cases

Recommending HPE NAS solutions for SMB customers

Introducing the customer scenario

Activity: Discovering a customer's business and technical requirements

Selecting an HPE NAS platform

Learning check

Network-attached storage

Finding the right fit for business and data can be challenging.

Discover a better solution for businesses

Network-attached storage

StoreEasy and Storage File Controllers provide the best of NAS and SAN

Ideal uses for HPE StoreEasy network-attached storage

HPE StoreEasy Storage family overview

Activity: Discovering the HPE NAS portfolio

HPE StoreEasy 1460 Storage

Raw, usable, and effective capacity

HPE StoreEasy 1560 Storage

HPE StoreEasy 1660 Storage

HPE StoreEasy 1860 Storage

Selecting an HPE StoreEasy 1000 platform

Learning check

HPE StoreEasy features and management

HPE StoreEasy is optimized to simplify file storage

Opt for a hybrid flash configuration

Maximize storage efficiency through wide protocol support

Multi-dimensional security keeps data protected

Easy-to-use, asynchronous replication

Guided configuration simplifies storage deployment

Day-to-day monitoring made even easier

Learning check

Storage-sharing technologies

Storage management

Folder and share management

iSCSI targets

Learning check

Sizing NAS products

- NAS sizing considerations

- Performance best practice for Server Message Block protocol (SMB 3)

- Virtualization best practices using Hyper-V on SMB or iSCSI with HPE StoreEasy products

- Configuration process steps

- Learning check

Customer scenario

Summary

Prelearning check

5 Recommending HPE Storage solutions for SMB customers

Recommending HPE Storage solutions for SMB

- Introducing the customer scenario

Activity: Discovering a customer's business and technical requirements

Direct-attached storage

- RAID levels

- Factors influencing performance

- Disk enclosures

- HPE D3000 Disk Enclosures: HPE DAS solutions

- HPE D3000 Disk Enclosures: Target customers

Learning check

Recommending HPE Storage solutions for SMB

- Introducing the customer scenario

Activity: Discovering a customer's business and technical requirements

- Storage area network

- What is a SAN?

- Definition of a SAN

Activity: Requirements for switches used in iSCSI SAN

Use cases for a SAN

High availability

- Redundant design

- Data mirroring

Clustering

Better utilization of backup and restore solutions

Business continuance

Server and storage consolidation

SAN components

HPE storage array positioning

HPE Nimble Storage

HPE Nimble Storage portfolio overview

All-Flash Array family

What is in the AF and HF series arrays?

AF and HF chassis front view and drive layout

AF-Series write operations: Steps 1 to 4

AF-Series write operations: Step 5

AF-Series read operations

AF-Series triple+ parity and integrated spare

SCM/NVMe-ready HPE Nimble Storage All-Flash platform

NVMe and SCM

Activity: SCM and NVMe for HPE 3PAR and Nimble Storage

Activity debrief

Learning check

Alternative components for different environments

Software-defined storage

HPE SimpliVity Data Virtualization Platform

HPE SimpliVity 380 Gen10

HPE SimpliVity 380 Gen10 G and H models

HPE SimpliVity 2600

HPE MSA Storage

HPE MSA-automated tiered storage

HPE MSA volume tier affinity feature

HPE MSA 1050 Storage

HPE MSA 2050 SAN Storage

HPE MSA 2050 SAN Storage features

HPE MSA 2052 SAN Storage

HPE MSA 1050/2050/2052 hardware building blocks

HPE Nimble HF Storage

- Adaptive Flash Array family

- HF-Series write operation: Steps 1 to 4

- HF-Series write operation: Step 5

- HF-Series read operations

- HF-Series triple+ parity

- HPE Cloud Volumes

- HPE Nimble Storage ProStack objectives

- Flexible platform and deployment options

- HPE Nimble Storage ProStack platform overview

Activity: Configuring a storage system

HPE 3PAR StoreServ

- HPE 3PAR StoreServ portfolio

- HPE 3PAR StoreServ manageability

- HPE 3PAR StoreServ 8000, 9000, and 20000 software details

- HPE 3PAR StoreServ 8000 hardware building blocks

- HPE 3PAR StoreServ 8000 host adapters

- HPE 3PAR StoreServ 8400 node pair

- HPE 3PAR StoreServ 8000 controller nodes

- HPE 3PAR StoreServ 8000 12 Gb SAS Drive Enclosures

- HPE 3PAR StoreServ 8000 hardware architecture

- HPE 3PAR StoreServ 8000 disk drives

- HPE 3PAR StoreServ 8000 models and system features

- HPE 3PAR StoreServ 8000—Key performance metrics

- Load balancing

- HPE 3PAR StoreServ Data at Rest Encryption

- Adaptive Flash Cache

- Express writes

- HPE 3PAR StoreServ Express Layout

- Dynamic Optimization introduction

- HPE 3PAR StoreServ Dynamic and Adaptive Optimization

Learning check

HPE Primera 600 Series

- HPE Primera technical specs

HPE Primera performance

Learning check

HPE backup and restore strategy

What is backup?

What is restore?

Backing up the environment

Backup strategies

Reasons why backups fail

Snapshot functionality

Snapshot: How it works—Create read-only snapshot

Snapshot: How it works—Update base VV

Clone introduction

Learning check

HPE StoreOnce backup systems

HPE StoreOnce Catalyst

Deduplication ratios

What is replication?

Bidirectional replication

StoreOnce VSA

StoreOnce 3640

StoreOnce 3620

StoreOnce 5200

StoreOnce 5250

What is Cloud Bank Storage?

Cloud Bank Storage setup

Cloud Bank Storage—Backup and copy

Cloud Bank Storage—Enhanced restore

Learning check

Configuration tools

HPE Storage Sizer

HPE NinjaSTARS for Nimble/StoreOnce/3PAR/Primera

Learning check

Management tools

HPE 3PAR StoreServ Management Console

- 3PAR Service Console
- NimbleOS WebUI
- StoreOnce Management Console
- HPE StoreOnce Recovery Manager Central
- InfoSight

Learning check

Summary

Prelearning check

6 Recommending HPE management and support solutions

Recommending HPE management and support solutions for SMB customers

- Introducing the customer scenario

Activity: Discovering a customer's business and technical requirements

- Hybrid IT management

Learning check

Remote IT support

- HPE Insight Online

- Choice of HPE Remote Support tools with Insight Online

Activity: Insight Remote Support: Support Matrix

- Embedded Remote Support

- Insight Remote Support Central Connect

- HPE Insight Remote Support

- Insight Online and Remote Support integration

- Target markets of direct connect for Remote Support

- Registering for Insight Remote Support direct connect

- Registering for Insight Remote Support through a centralized hosting device

Activity: HPE Get Connected

Learning check

HPE OneView

- Infrastructure automation made simple

- Simplify lifecycle operations

- HPE Composable Infrastructure partner ecosystem

- Deploy infrastructure at cloud-like speed

Develop more apps faster

HPE OneView addresses two approaches to IT management

Two methods for an enhanced user experience

Supported devices for monitoring by HPE OneView

Resource health monitoring operations

SNMP trap receiver processing

Alert Aggregation Service processing

System health monitoring

SAN health and diagnostics

Learning check

HPE OneView partner integrations

HPE plugins for VMware

HPE OneView for VMware vCenter: Brings the power of HPE OneView to VMware environments

HPE OneView for VMware vCenter: Overview

HPE server hardware: Overview

Integration with HPE OneView Remote Support

HPE OneView Integrations for Microsoft System Center

Seamless integration into Microsoft System Center

HPE OneView for Microsoft System Center

System Center Virtual Machine Manager integrations

Learning check

HPE OneSphere

Building on our momentum to simplify hybrid cloud

HPE OneSphere simplifies hybrid cloud environments

Key features

Learning check

HPE InfoSight: AI for hybrid cloud

HPE InfoSight: Improving efficiency through AI

HPE InfoSight—AI for hybrid cloud world

See once, and prevent for all

Only InfoSight goes beyond the obvious

InfoSight for servers—Putting it all together

Learning check

HPE RESTful API and PowerShell cmdlets

- REST communication

- Example of HPE RESTful API implementation—iLO

- Using PowerShell to manage HPE devices

- Example of PowerShell cmdlets—HPE iLO with PowerShell

Learning check

Customer scenario

- Introducing the customer scenario

Summary

Learning check Answers

Chapter 1

Chapter 2

Chapter 3

Chapter 4

Chapter 5

Chapter 6

Practice Test

Exam details

- HPE0-V14 testing objectives

Test preparation questions and answers

- Questions

- Answers

1 Introduction to a Hybrid World

LEARNING OBJECTIVES

After completing this chapter, you should be able to identify:

- ✓ HPE Small Business Solutions
 - ✓ HPE as a company built on the idea of innovation
 - ✓ The Hybrid IT vision
-

SMB Hybrid IT for Dummies

LEARNING MADE EASY



Hewlett Packard Enterprise
and Intel Special Edition

SMB Hybrid IT

for
dummies[®]
A Wiley Brand



See how hybrid IT
is the way of the future

Identify the best location to
run individual workloads

Build your on-premises
environment

Brought to
you by:

Hewlett Packard
Enterprise



Scott D. Lowe

Figure 1-1 SMB Hybrid IT for Dummies

HPE is providing an eBook to understand hybrid IT: *SMB Hybrid IT for dummies*. Use the book, referenced in [Figure 1-1](#), to learn how HPE helps with transforming IT operations and save time and money in the process. The book can be downloaded from: <https://www.hpe.com/info/smb>.

HPE Small Business Solutions

This chapter is wrapped around SMB customer scenarios. It does not start with complexities because building Hybrid IT solutions can be as easy as taking a validated SMB solution and building the customer solution based on best practices and guidelines outlined in this guide.

Small businesses today: Head in the cloud—feet on the ground

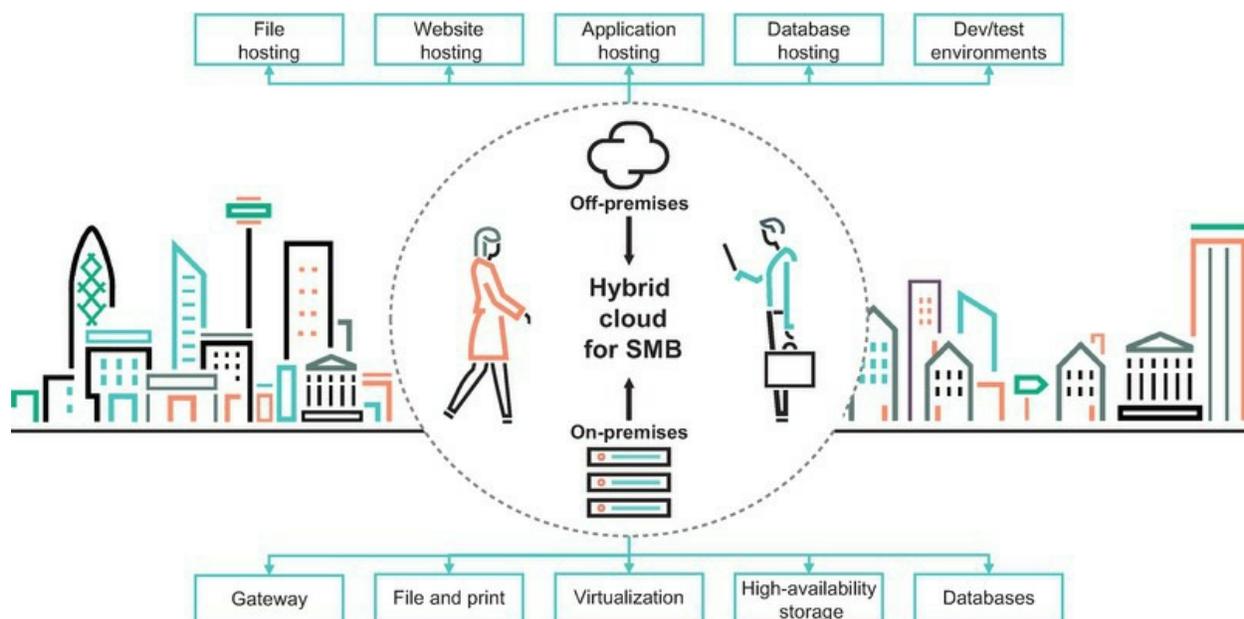


Figure 1-2 Small businesses today: Head in the cloud—feet on the ground

When it comes to IT purchase decisions, one of the biggest things small businesses struggle with is where to run your IT—on-premises or in the cloud.

Cloud has gained a lot of top-of-mind awareness due to perceptions of low cost of entry, ease of deployment, and flexible economics. However, cloud

tends to fall short when it comes to data protection, compliance and legal concerns, performance, and service levels, just to name a few.

Like a lot of other SMBs, you may need help striking the right balance of on-premises IT (one or more servers at your office) and off-premises IT (private or public cloud).

Hybrid cloud, as indicated in [Figure 1-2](#), is the term for mixing these two deployment options, and even SMBs can achieve a secure, simple, and affordable hybrid cloud environment that makes your customers more competitive.

Are your customers trying to find the right balance?

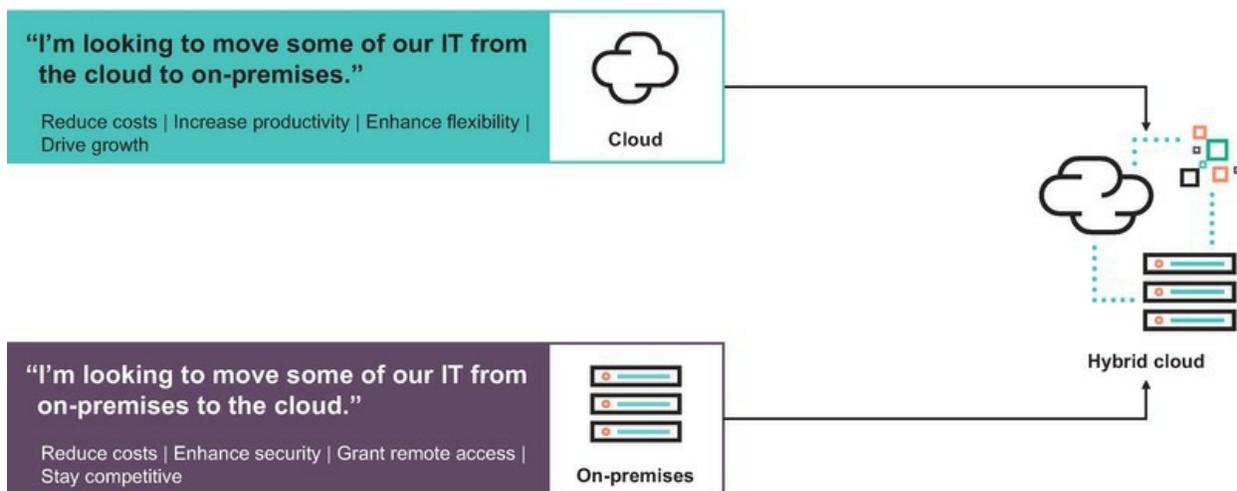


Figure 1-3 Steps to create a golden image with Image Streamer

When it comes to hybrid cloud, HPE can help your customers find the right balance.

[Figure 1-3](#) shows SMB IT tends to fall into two categories:

- Some small businesses have their IT running completely or primarily on public cloud. As the business grows and your customer builds a larger customer base, they need a more advanced means to protect customer data while improving performance and accessibility. This naturally leads to bringing some IT on-premises for greater security and control.
- Other SMBs already run IT primarily on-premises and are coming up for

a technology refresh, perhaps with end of support for Microsoft Windows Server 2008. If this is your customer, they might be looking to reduce their costs through server consolidation, while moving some nonbusiness-critical applications into the cloud. This also gives them an opportunity to build business continuity using cloud backups.

HPE Small Business Solutions

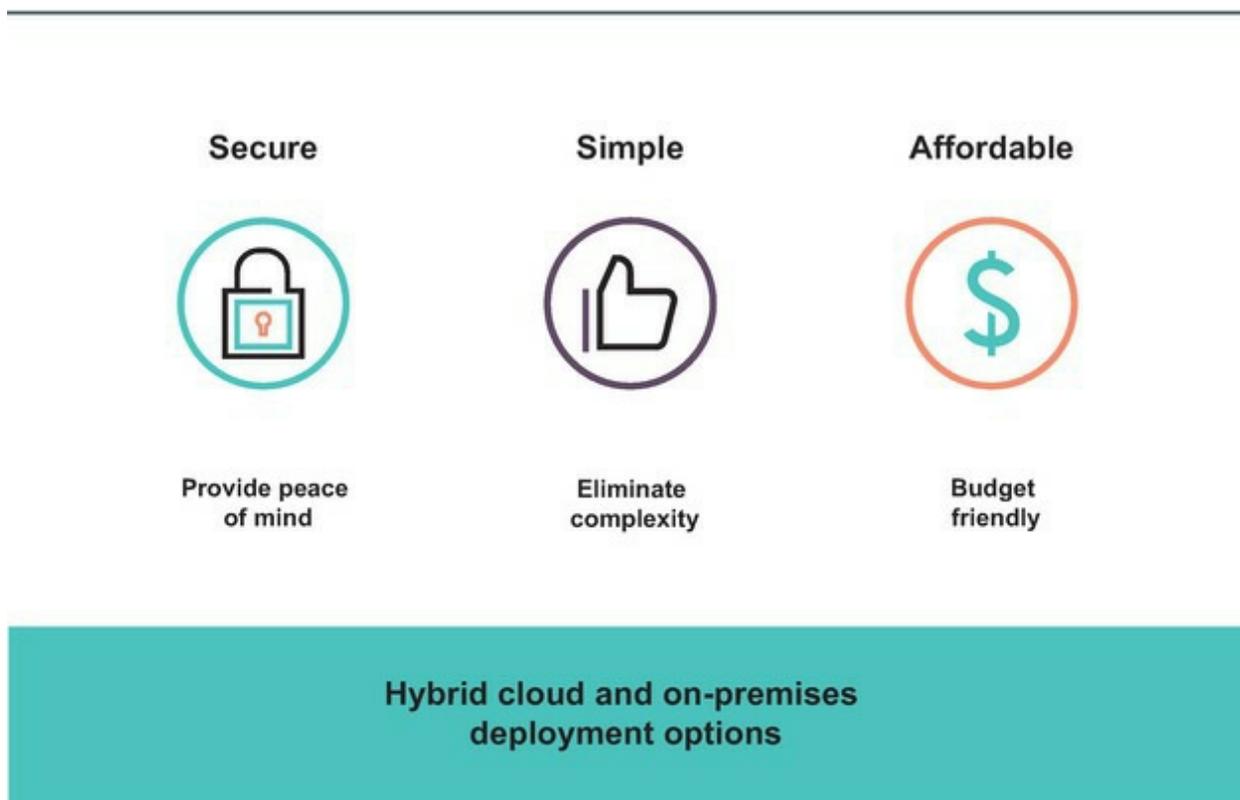


Figure 1-4 HPE Small Business Solutions—From a partner your customers can trust

Everything we have talked about so far leads us to why HPE offers HPE Small Business Solutions with on-premises and hybrid cloud options.

HPE Small Business Solutions capture best practices and incorporates the latest technologies optimized for foundational SMB use cases. These configured and validated solutions are prescriptive yet customizable, making

them easy to buy and deploy.

They are designed to be simple, secure, and affordable, which brings a host of business benefits. They also take the guesswork out of hybrid cloud for SMBs—even if customers have limited IT resources.

HPE Small Business Solutions are based on HPE ProLiant servers, and, as shown in [Figure 1-4](#), they are:

- **Secure**—Customers can protect against attacks and quickly recover from unplanned downtime with built-in security features that make HPE ProLiant Gen10 the World’s Most Secure Industry- Standard Servers (depending on the model).
- **Simple**—HPE Small Business Solutions are a one-stop shop and include preconfigured hardware and software, making them easy to deploy. In addition, HPE ProLiant servers include automated intelligence that reduces the time and skills required for deployment and management.
- **Affordable**—Low-cost HPE Small Business Solutions make hybrid cloud easy to consume. HPE Subscription offers a worry-free, flexible option that packages best-in-class hardware, software, accessories, and services into a predictable monthly payment.



Note:

HPE Small Business Solutions are regionally released under different names (such as “Smart Buy Express Offers” in the US and Canada, “Top Value Offers” in Europe, and “Intelligent Buy Offers” in Asia Pacific and Japan).

What business outcomes are top-of-mind for your customers?



Figure 1-5 HPE Small Business Solutions are built around business outcomes

Now, look at how easy it is to adopt hybrid cloud to help with each of these business outcomes. Let us look at those top goals—and HPE solutions—again, as indicated in [Figure 1-5](#).

Keep the business running

- **HPE Small Business Solutions for Hybrid File and Backup**—Protect files and meet regulatory compliance and security needs with a business continuity plan that includes cloud backup.
- **HPE Small Business Solutions for Hybrid Virtualization**—Minimize downtime with the ability to failover so files and data remain accessible.
- **HPE Small Business Solutions for Hybrid Database**—Maintain data availability with automatic replication to the cloud.

Increase IT efficiency

- **HPE Small Business Solutions for Hybrid Database**—Grow or shrink databases in size, performance, and cost as demand increases or decreases, and segment the database based on the sensitivity of the files.
- **HPE Small Business Solutions for Hybrid Virtualization**—Consolidate more workloads onto fewer servers and run legacy applications that are not cloud ready.

Improve productivity

- **HPE Small Business Solutions for Hybrid File and Backup**—Keep the team working at peak efficiency with automated hybrid cloud business continuity.
- **HPE Small Business Solutions for Hybrid Dev/Test**—Give developers the tools to create and collaborate without tying up on-premises resources.

Enhance the customer experience

- **HPE Small Business Solutions for Hybrid Web Hosting**—Boost sales by showcasing products and services with an always-open business website.
- **HPE Small Business Solutions for Hybrid Dev/Test**—Roll out new apps and updates faster than ever.

Because choice is so important, your customers can achieve these hybrid scenarios by tapping into cloud services from Microsoft Azure, public cloud services, or collocated cloud services from their IT Partner.

Start with complete, validated hardware and software stacks



Figure 1-6 Basic components of on-premises HPE Small Business Solutions

Now, let us take a closer look at how we build these solutions. [Figure 1-6](#) shows the basic components of on-premises HPE Small Business Solutions.

HPE Small Business Solutions for Small Office Deployment offers a choice between Aruba and HPE Office Connect switches and wireless access points.

For these solutions, customers have choice of operating system—either ClearOS (a high-value, lowcost OS platform and solution software based on Linux) or Microsoft Windows Server (today’s most popular SMB environment).

If the customer chooses a virtualization solution, then they have a choice of hypervisor—ClearVM, Microsoft Hyper-V, or VMware vSphere. If customers choose high-availability storage enabled by software-defined storage, they have a choice of Microsoft Storage Spaces Direct or VMware vSAN.

These solutions also include a variety of server options, service, and support choices, and they are ready to also be scaled to grow with the organization by adding memory and storage.



Note:

More information about ClearOS can be located at <http://clearos.com>

Enable hybrid cloud by adding on to on-premises solutions

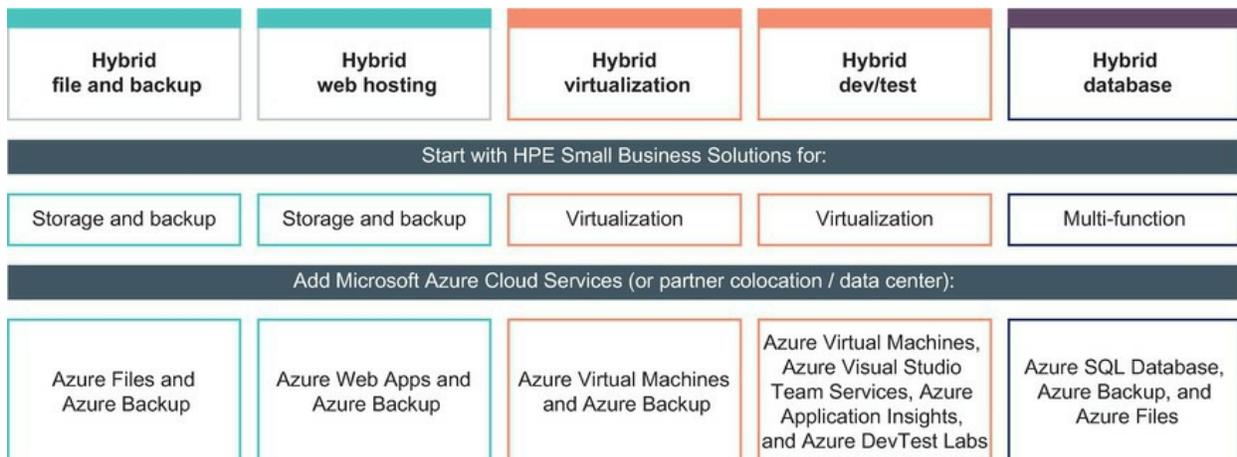


Figure 1-7 Choose hybrid cloud now, or later

These SMB solutions also form the basis of HPE hybrid cloud offerings. As we can see from [Figure 1-7](#), if the customer starts with an HPE Small Business Solution for storage and backup, for example, the customer can elect to add Microsoft Azure Cloud or partner private-cloud services to create an HPE Small Business Solution for hybrid file and backup with automated, cloud-based replication and backup.

Learning check

1. Which HPE server platform is used in SMB offers?

HPE overview

HPE is a company built on the idea of innovation. From the beginning, we have always been about helping customers get more value out of their apps and data. Our mission has always been to help our customers use technology to improve the way we live and work. We have done that through innovation, expertise, and most importantly through a trusted partnership with our customers. For over seven decades, as illustrated in [Figure 1-8](#), we have helped our customers through every major technology transition, including the shift from mainframe to industry-standard servers, the virtualization revolution, and now to cloud and tomorrow, to the edge, and artificial intelligence (AI).

Innovation is in our DNA

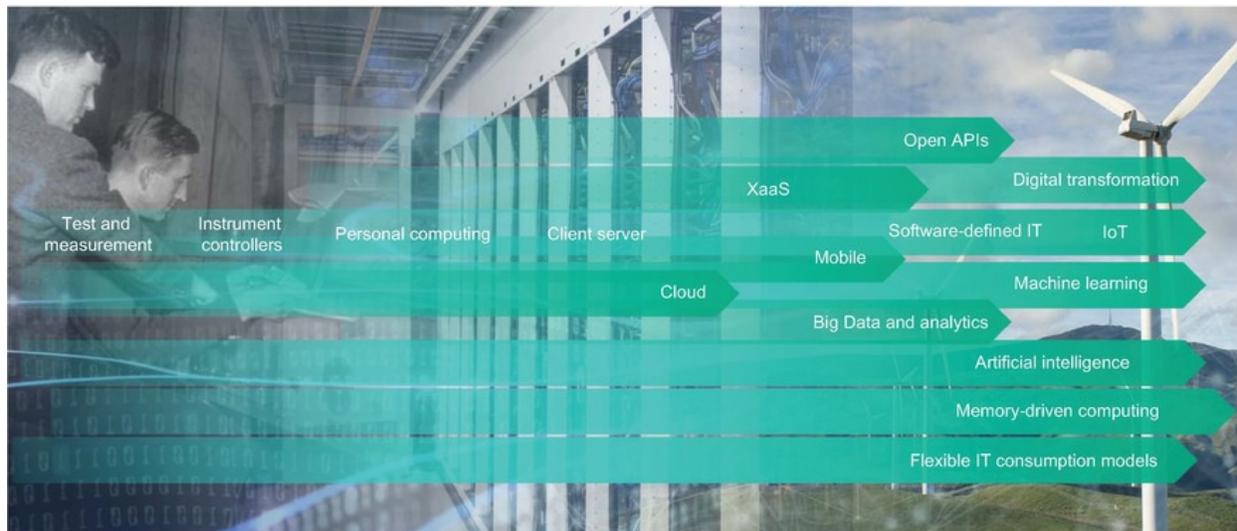


Figure 1-8 Powering the world's apps and data for over 75 years

Today we face the biggest shift yet, digital transformation. Fueled by an explosion of data and a desire to unlock insights through AI and advanced analytics, a new generation of technology, apps, and data that is increasingly accessible to all, digital transformation is disrupting every industry—creating massive opportunities and challenges for businesses of all sizes. Businesses everywhere want to grow, find new revenue streams, become more efficient, and gain a competitive edge.

These customers are starting to realize that the data generated by their business is not a burden, but a potent raw material in that data provides vital insights that can transform their businesses.

We have all been talking about digital transformation for a while, but we are far from done. Companies are not slowing down; in fact, their embracing of new technologies push businesses and governments to the edge of their capabilities and speed.

Creating a world where everything computes

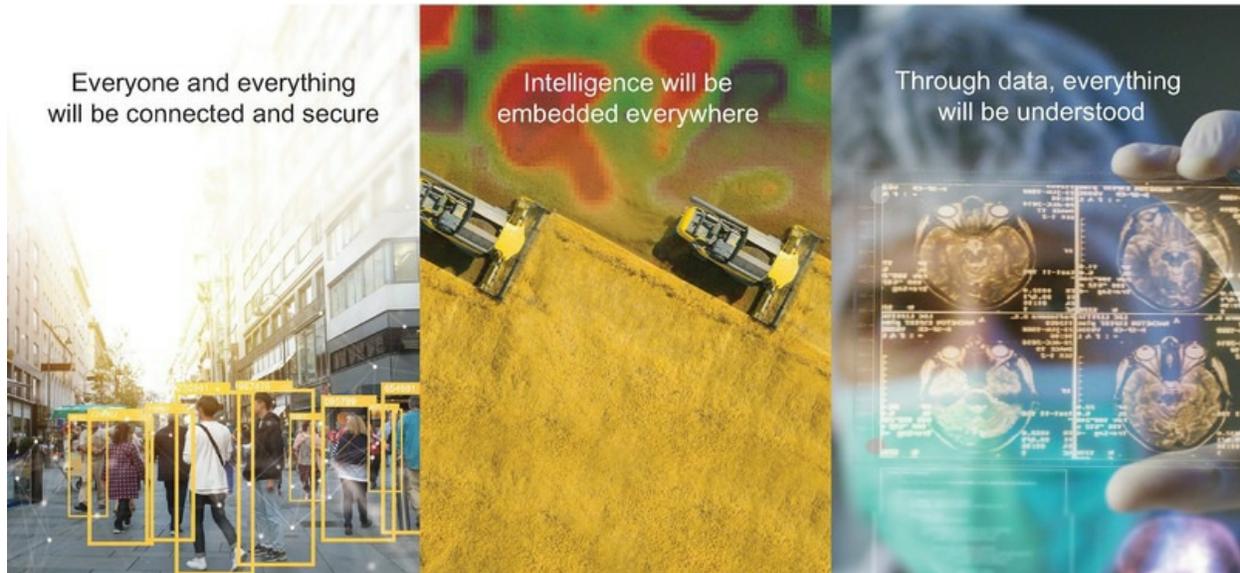


Figure 1-9 Creating a world where everything computes

We are creating a world where everything computes, a world that is hyperconnected, and where everyone and everything share data, as illustrated in [Figure 1-9](#).

It is a world where billions of connected people, devices, and things are sharing data from our cars and homes, workplaces, museums, stadiums, hospitals, factory floors, and the data center—creating petabytes of data that bring new insights and actions—and where AI is used to make decisions faster and autonomously—with speed measured in nanoseconds.

We want to help customers connect all their data—across all their edges and all their clouds.

Driven by a new generation of apps and data—and a new speed of business



Figure 1-10 A world driven by a new generation of apps and data—and a new speed of business

The world is driven by a new generation of apps and data—and a new speed of business. In our connected world, expectations continue for more engaging, personal, and seamless experiences across platforms and channels, reflective of our context in a given moment.

Businesses are capitalizing on new ways to engage and get closer to customers, thanks to deeper customer insights that drive smarter design and better experiences—with less friction at every step. Businesses are also tapping machine learning, Internet of Things (IoT), and AI to create new services and drive operational efficiencies, using predictive insights to reduce downtime and keep the business running.

All of this is fueled by a new generation of apps and data, existing on multiple platforms, from the data center to the cloud—and increasingly at the edge of the network. This is where people, places, and things converge to create new intelligent digital experiences, based on massive amounts of data being collected, analyzed, and often acted upon instantaneously.

As indicated in [Figure 1-10](#), the pace of business has never been faster—and accelerating time to value pays off. The possibilities are limitless—when optimally designed.

The edge has arrived



Figure 1-11 The edge has arrived

There is a shift underway—from data being highly centralized in large data centers to highly decentralized at the edge. It is shifting us to a whole new speed of business—measured in nanoseconds.

Let us take a look at [Figure 1-11](#), where we have come from. We began with mainframe and super computers that solved complex issues, but they were enormous and slow. In the 80s, we transitioned to the client-server and networked computing model, which was much more efficient, but still limited in capacity and speed. In the 2000s, we were introduced to a new way of computing with mobile devices with access to virtual/cloud data storage.

So, what do we really mean by the edge? The edge is where businesses interact with their customers; where products are manufactured; where employees work each day; where everyday people interact, purchase, and explore; and where technology gets put into action. The edge employs AI, machine learning, and automation to continuously learn, predict, and adapt to changes, needs, and threats in real-time—and enables us to act locally, in the moment, in context—creating new possibilities in every industry. Connecting these new insights from edge to cloud across the enterprise, integrating the data with existing business systems (ERP, CRM, and more) is what is driving value and growth.

We have been at the forefront of developing technologies for the intelligent edge to help enable realtime insights, personalized experiences, and the ability to take action instantaneously. Recognizing the opportunity at the intelligent edge is why we acquired Aruba Networks—to deliver mobile-first solutions. Through Aruba, we acquired IntroSpect to embed security with machine-learning capabilities that allow for rapid scale User and Entity Behavior Analytics (UEBA) at the edge and to indicate attacks among people—and things. Over the next four years, HPE will invest \$4 billion in intelligent edge innovation. We will accelerate our investments in R&D to continue to advance and innovate new products, services, and consumption models across a number of domains including security, edge computing, automation, AI, and machine learning.

When it comes to data processing, HPE's capabilities span every scale. When it comes to the kind of data processing involved in machine learning and AI, we again lead.

Accelerating time to value is the ultimate goal in a world where real-time is the new just-in time: time to market, time to insights, and time to action. Businesses need to be obsessive about accelerating time to value—and hyper-efficient, ready to capture new opportunities and defend against disruption, increasingly in real-time.

In a data-driven world, speed matters more than ever—we believe businesses that can act fast on a continuous stream of data—from everywhere—will win.

HPE: Accelerating your customer's digital transformation agenda

HPE was created to help your customers accelerate their digital transformation. We uniquely address the three biggest barriers to speed:

- **Transformation of technology for an edge-centric, cloud-enabled, data-driven world**—We help your customers create a software-defined, hybrid operating model that spans any cloud to multi-cloud, on- and off-premises, and the edge so they can create and deliver everything as a service dynamically, everywhere.
- **Transformation of people and processes, creating a culture for growth and innovation**—The hardest part of any digital transformation

involves people, their skills, processes, and culture.

- **Transformation of economics and technology investment strategy**— One of the most underappreciated aspects of any digital transformation is financial strategy. Capital trapped in underutilized fixed assets is a barrier to speed and innovation. With HPE Financial Services, we provide the right strategies and as-a-service consumption models to optimize IT investments.

Transform your customer's technology

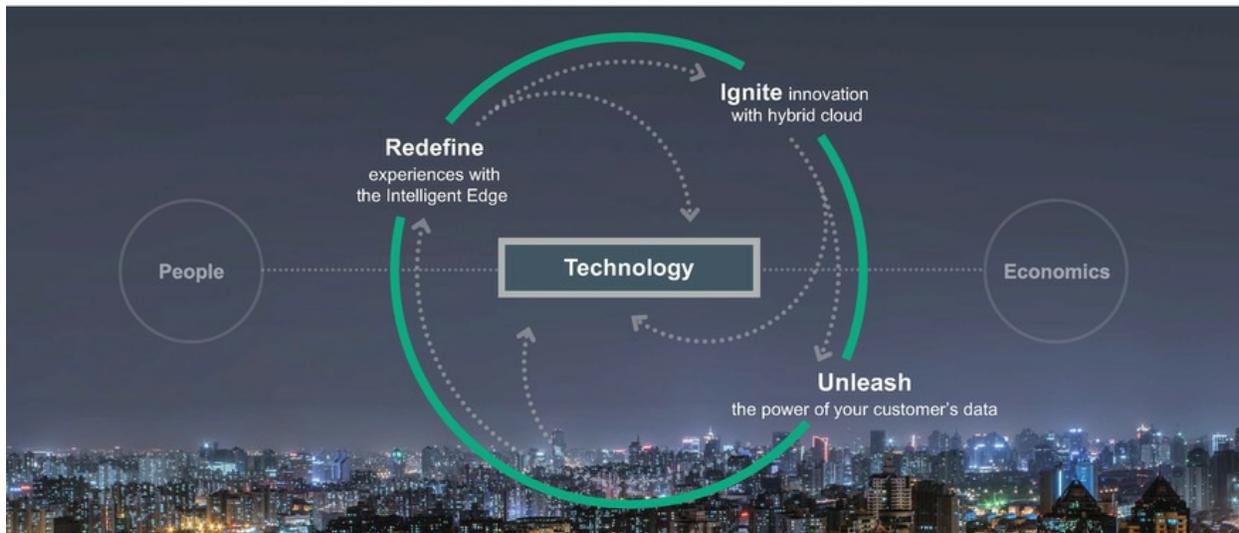


Figure 1-12 Accelerate with intelligent, software-defined solutions, edge to cloud

The explosion of connected things creates possibilities from the mobile, IoT edge to the cloud, and across the business. We see a world that is edge-centric, cloud-enabled, and data-driven—where everything is created and delivered as a service. As illustrated in [Figure 1-12](#), we are helping business transform to:

Redefine experiences with the intelligent edge

The revolution that is happening now is the explosion of data, powered by the edge. The explosion of data from connected things is driving data from large

data centers to many small centers of data. By 2020, Gartner estimates 75% of data generated will never see a data center or cloud (Gartner, *Top 10 Strategic Technology Trends for 2018: Cloud to the Edge*, Published: March 8, 2018). It will be created at the edge in an oil rig, a retail store, a medical device, a self-driving car, and so on.

The edge is where billions of people and places and trillions of things intersect—generating unimaginable amounts of data. Data that has the potential value to drive insights and actions is where your customers interact with their customers, manufacture their products, where their employees work, where experiences can be personalized, local, and in the moment—enriched by context.

Ignite innovation with hybrid cloud

Enterprises must now create and deliver services across hybrid and multi-cloud environments in a hybrid reality. To do this effectively and efficiently, businesses need to transform. It is not trivial. In fact, the biggest challenge is people adapting to the continuous change inherent in cloud.

HPE accelerates cloud transformation by addressing the toughest challenges from people and process, to economics and technology, with expertise and solutions to help enterprises get hybrid cloud right. We help your customers get the right mix of clouds for all of their apps and data—in public clouds and private clouds—connected to their on-premise environments and unified with a common software-defined operating model, delivered as a service with pay-per-use models that span public and private clouds. We help your customers create an IT culture for growth, with the control, security, compliance, and governance behind it, so they can focus on what is next for their business.

Unleash the power of customer data

Data is your intellectual property (IP)—the new currency. To stay ahead, your customers must capture, keep, and refine every bit of their data. They cannot afford to throw anything away because all their data is valuable. It is all about speed to insights, everywhere from edge to cloud, and “real time” is the new “just-in-time.”

Data is coming at us faster than ever before, and from new sources, like the

intelligent edge. Businesses can barely keep up with all this data—and the data explosion will continue.

We help businesses process and analyze data from every edge to any cloud and at every scale. From intelligent storage to high-performance computing—and with our breakthrough innovations like Memory-Driven Computing, we help businesses accelerate insights, action, and value—everywhere.

Transform your customer's people and processes

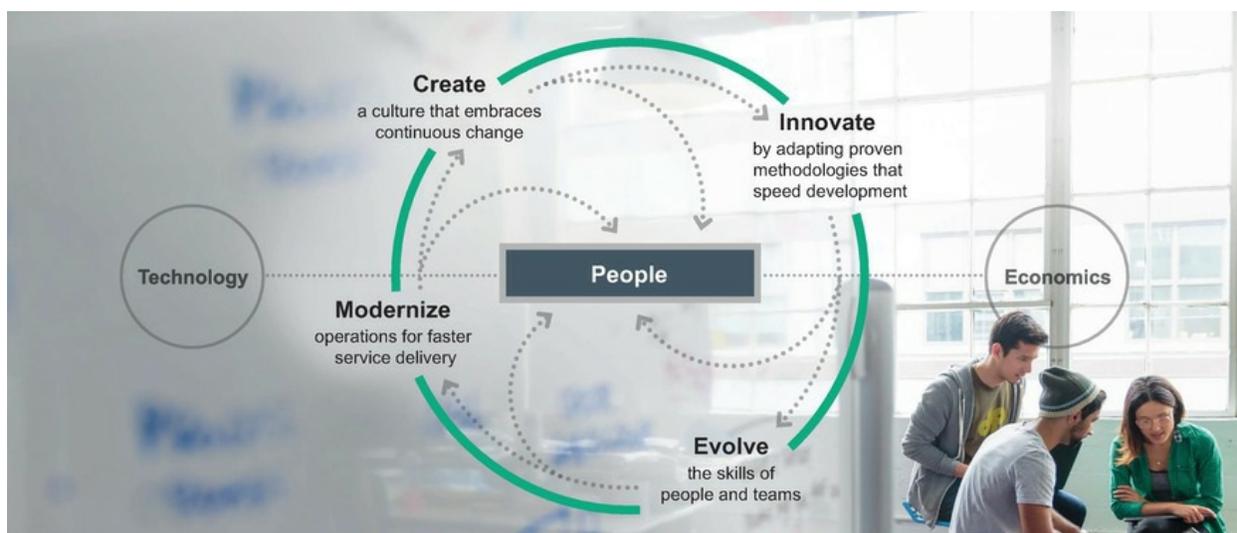


Figure 1-13 Create a culture for growth and innovation

Digital transformation comes down to how your customers drive business growth, ensuring they capitalize on their apps, data, people, and their capital to innovate and move faster. The hardest part of any digital transformation involves your customer's people—evolving skills, processes, and culture are the necessary components required to enable their teams to be successful through technology change. The challenge lies in helping people to understand how their roles change, why it matters, and equipping them with the skills that they need to be successful in their new roles. With HPE Pointnext we help your customer's people to be successful through their unique digital transformation journey, as illustrated in [Figure 1-13](#):

- **Innovate by adopting proven methodologies** that speed development. When DevOps is combined with Agile methodology and cloud-native

technologies, applications are built with a much shorter time to market, fewer issues, and improve quickly with frequent updates. This is frequently referred to as Continuous Integration/Continuous Delivery (CI/CD). HPE Pointnext helps businesses establish and implement a CI/CD strategy, accelerating cloud-native development and leveraging a developer-ready infrastructure to help them embrace these new techniques, tools, and practices.

- **Evolve the skills** of your customer's people and teams. Modern IT systems are comprised of new technologies, new methodologies, new partners, and a new generation of staff. Keeping your customer's IT organization skilled to design, build, and operate these new environments requires a new approach. HPE Pointnext skills-based training is most effective in a team dynamic and at the time of need. HPE Digital Learner is our on-demand learning platform that provides online access to micro-learning courses everywhere, with real-time video feeds and collaborative mentoring services.
- **Modernize operations** for faster service delivery. IT complexity has exploded over the past few years, and underfunded and overwhelmed IT operations teams are struggling to keep up. Public outages and security breaches are a rising trend. HPE Pointnext, with Datacenter Care, transforms your customer's IT operations. Working alongside their existing Ops team, we can help free them from the undifferentiated heavy lifting of infrastructure support, bringing a holistic view to service management and helping their team adopt new tools, processes, and approaches to manage a modern IT environment. By freeing up their people, we can help them shift these resources to other operational roles. We employ best practices from Information Technology Infrastructure Library (ITIL) and IT Service Management (ITSM).



Note:

ITSM is stands for IT Service Management, refers to the total activities—directed by policies, organized and structured in processes and supporting procedures. Enterprises dependent upon IT technology should ensure every aspect associated with IT technology is accomplished via proper planning and strategizing. Only then will

they be able to get the best out of it. There are specific tools in the market which can help enterprises accomplish the planning and strategizing of IT technology effectively in a seamless and hassle-free manner. These tools are usually known as ITSM tools (ITarian.com, September 2019).

Transform your customer's economics

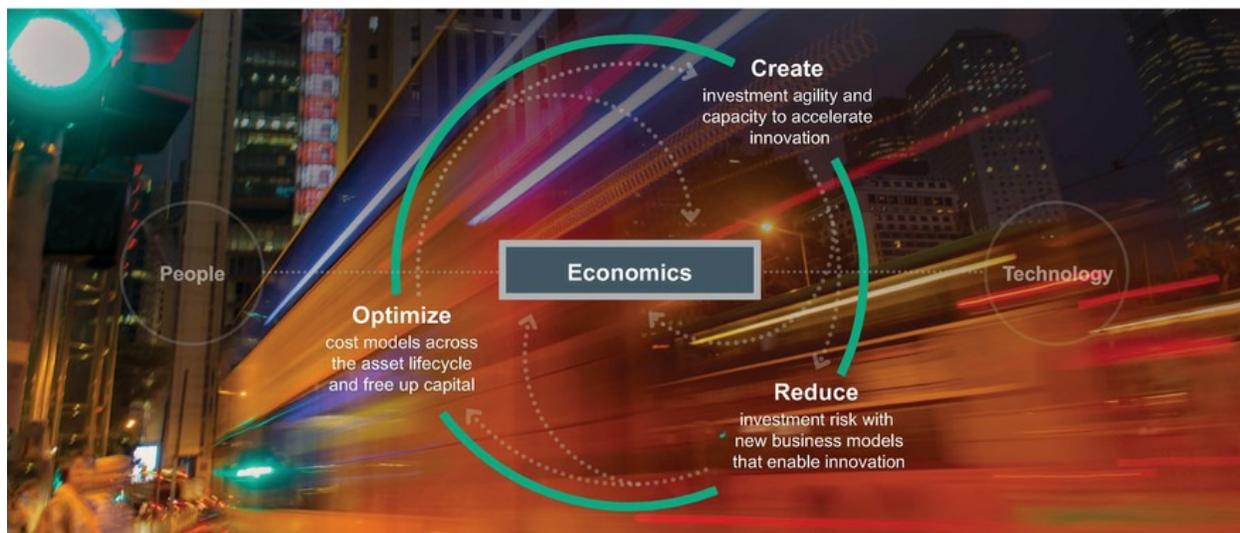


Figure 1-14 Align your customer's IT investment strategy for growth and innovation

Transforming IT investment strategies is key to successful digital transformation. As shown in [Figure 1-14](#), HPE helps businesses with three strategies to maximize your customer's IT investments— more power in their IT operational strategy today, future growth, and innovation.

Optimizing cost models and freeing up capital is focused on addressing existing operations. Today, 70 to 80% of most IT budgets is spent on simply maintaining the operational norm while only about 20% is available for investing in new areas of innovation or new applications. The challenge is how to reduce and optimize that maintenance spend.

HPE Financial Services can help by monetizing in-use and retired systems to fund new investments and by stabilizing legacy operations and spend aligned to in-use systems not ready to be retired. This increases profitability and

minimizes new investments in old systems.

Your customers can create investment agility and capacity to accelerate innovation. Investing in innovation that your customer's business needs now requires significant up-front capital. However, their existing budgets do not always cover everything they need. Companies can create additional investment capacity from current systems by selling off assets while continuing to use them. Oftentimes, IT leaders do not think about this, but in every IT asset they own, they have invested capital. That means the technology continues to have value for many years after its initial purchase.

Your customers can also look at better risk management. Risk has two forms. The first is that they have not invested in enough capacity or enough capability to handle the unexpected demand for a new market. It means turning away customers or giving them a bad experience. The other type of risk is equally bad, by which they build capacity and never use it. That ties up precious dollars that can be better used elsewhere.

HPE Financial Services can help with investment agility by creating an investment model that is designed for innovation, but right-sized to manage capacity risk. We can also remove infrastructure lock-in for systems that have not fully depreciated, but are holding up innovation initiatives. This frees up "trapped" capital, creating investment dollars that can be quickly aligned to new innovation investments.

Companies looking to implement new business models or run experimental deployments need to think about mitigating investment risk. Technology never stops changing, and most businesses must have a sharp eye on the future, looking at experimental technologies, rolling out products and services that are not ready for the mainstream. This can be the difference between being the disruptor or being disrupted, but it can be very risky. How do companies invest quickly while still managing the risk? If the innovation is successful, how do they scale quickly? If it fails, how do they make sure that the business is not impacted for years to come? There are many investment and asset lifecycle solutions that can help reduce risk, accelerate new model innovations, and help companies move into new spaces quickly.

HPE Financial Services can help in a number of ways. For example, in a phased deployment of new technology over a 12-month period, the customer can only pay when the systems are configured, tested, and turned on. This

allows for a more gradual introduction of new technology. With our aggressively tailored investment models that are perfect for untested or experimental deployments, we provide peace of mind. You can step away (all or part, with no penalty) if the deployment is not successful, and continue or expand, based on project and business success.

With HPE Pointnext, we help your customers assess where they are today, and we help them map their customer's apps and data to their optimal platforms to deliver speed and flexibility for their business. HPE Pointnext provides a full spectrum of advisory, design, integration, migration, and operational consulting services to help our customers deploy and manage workloads across their entire IT estate.

At HPE, we continue to strategically build a portfolio that helps our customers take advantage of digital transformation and the resulting explosion of apps and data. We are focused on helping our customers accelerate business outcomes by extracting critical insights from their data, wherever it lives, from edge to core to cloud. HPE Pointnext provides customers with a truly comprehensive Hybrid IT strategy that includes private, managed, and public clouds, as well as traditional IT.

Accelerating our strategy with partnerships

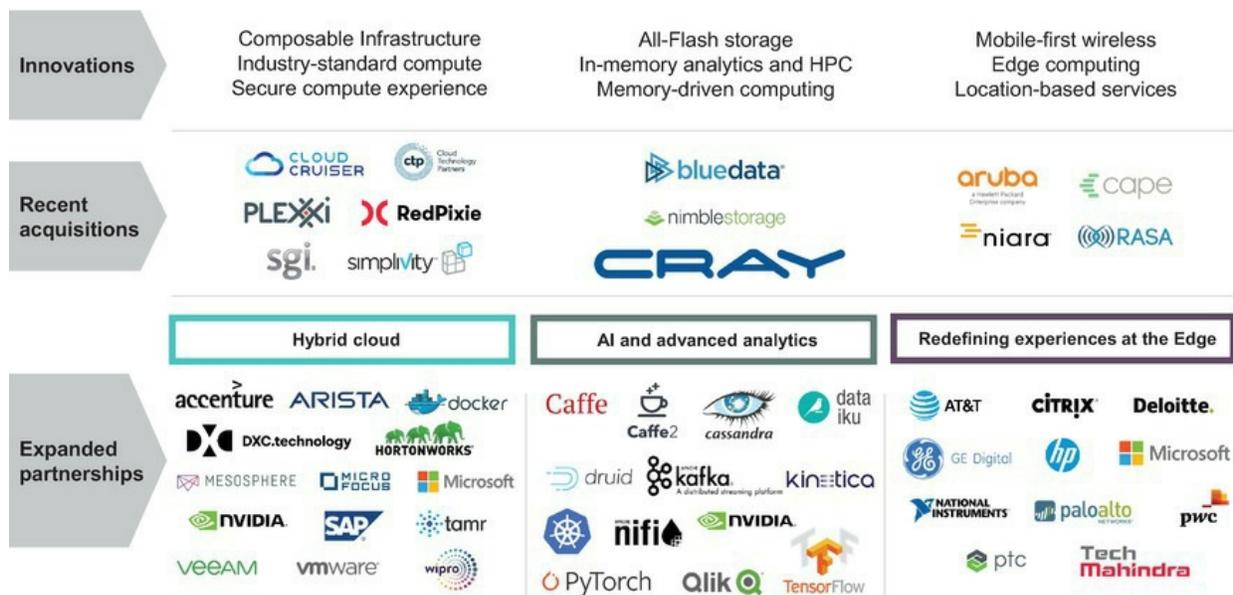


Figure 1-15 Accelerating our strategy with partnerships

It is important to be able to talk across hybrid/data/edge and be familiar with key technologies and recent acquisitions, as shown in [Figure 1-15](#). If you look at hybrid cloud, we started with an industry-leading base of innovations, and our portfolio has never been stronger.

With Synergy, we introduced the world's first Composable Infrastructure to simplify and accelerate DevOps and business innovation. Our 3PAR storage line has been leading the way for customers to the all-flash data center, and of course, we maintain decades of leadership in industry-standard computing.

For data and analytics, we also lead the market in High Performance Computing (HPC), in-memory analytics, and The Machine project from HPE Labs was out in front in planning the future of computing for when Moore's Law runs out of gas on current architectures.

Building on that foundation, we then put our cash to work to build on our strengths and to accelerate our roadmap and vision.

SimpliVity, Nimble, and Cloud Cruiser gave us an instant leadership position in hyperconverged, software-defined storage, and cloud and Hybrid IT management as well as radically simplifying onpremise environments in support of developers and business users, with a higher objective to help our customers accelerate insights and reduce time to value.

According to IDC, by 2022, the total addressable market for artificial intelligence/machine learning (AI/ML) and Big Data is expected to grow to approximately \$160 billion.



Note:

By 2022, IDC projects Big Data/Analytics software investments will reach \$90 billion (<https://www.idc.com/getdoc.jsp?containerId=prUS44215218>) and \$77.6 billion for cognitive and artificial intelligence-based systems (<https://www.idc.com/getdoc.jsp?containerId=prUS44291818>).

However, we have recognized that not all organizations will have the AI and data literacy skills needed to extract business value and actionable insights from their data, and demand is increasing for faster and more cost-effective solutions that can easily deploy AI/ML and Big Data analytics.

To accelerate this opportunity, HPE acquired BlueData in early 2019, a leading provider of software that transforms how enterprises deploy AI and Big Data analytics, expanding our offerings in these rapidly growing markets.

We also recently welcomed Plexxi and RedPixie into the fold, and HPE will continue to improve its hybrid cloud consulting capabilities that enable customers to deploy and manage workloads across environments seamlessly.

With Plexxi, we are delivering enhanced hyperconverged and composable solutions with a next-generation, software-defined data network fabric that can automatically create or rebalance bandwidth to workload needs, increasing agility, efficiency, and accelerating how quickly companies deploy applications and draw business value from their data. RedPixie, one of Europe's largest cloud specialists, expands our consulting capabilities to build and manage Microsoft Azure hybrid solutions for the financial sector.

We also acquired SGI, nearly doubling our market lead in HPC, which architecturally is forming the foundation for the future of HPE's approach to AI/ML and edge-to-core analytics. With SGI, we now have the perfect Big Data analytics solutions for industry, government, research, energy, and healthcare.

As we have said, innovation in today's world is a team sport, and our relationship with our ecosystem partners is growing from strength to strength. In many ways, our partnerships today are much deeper, and in ways that previously were not possible; our services and software business were seen to be in competition with them. As a result, we have made many moves to extend partnerships, and not just with traditional names, but with a new breed of players who really understand the future of apps and data.

Through our venture investment and partnership program, called Hewlett Packard Pathfinder, we are using our deep expertise to curate the best emerging start-ups and help customers innovate and develop solutions faster. One example is Mesosphere, a data center operating system that enables customers to run modern apps in a hybrid infrastructure. Another is Chef, an infrastructure automation technology that we have integrated into our Proactive Care Service.

In addition, rather than trying to compete head to head in public cloud, we understood that HPE was uniquely positioned to help our customers take

advantage of multi-cloud solutions, such as with our partnership with Microsoft Azure. We have also partnered with Docker and introduced the industry's first Docker-ready servers, which create the flexibility of a true hybrid solution. We also have a deep partnership with Arista for the software-defined data center (SDDC) networking that focuses on building application workloads for the cloud.

In the intelligent edge, we saw a huge new opportunity, but it was just emerging. We started with a core acquisition of Aruba that instantly put HPE at the front of the wireless first, mobility revolution. Because the edge is about data, and security is the biggest challenge, the next acquisitions we made were Rasa and Niara to add to the strengths of the Aruba mobile-first platform with more analytics and AI for managing and securing networks. We have also recently acquired Cape Networks, and the team have now joined Aruba to expand Aruba's AI-powered analytics and assurance capabilities with a sensor-based service assurance solution that gives a simple, proactive, and network-agnostic tool for measuring and monitoring Software as a Service (SaaS), application, and network services.

In IoT—we built on our own innovation in core-to-edge computing and storage—and extended our partnerships with Deloitte, GE Digital, National Instruments, PTC, Tech Mahindra, Veeam, and more. Our aim is to complete our core-to-edge solution stacks that merge IT and operational technology (OT) data to drive a variety of industrial IoT solutions—to put IoT into action across multiple regions. Recently, Forbes magazine called us the “quiet giant of IoT” because of the depth of capability and the fact that we do not just “talk about IoT”—we know how to implement it and get outcomes for our customers.

HPE Pointnext—bolstered by Cloud Technology Partners (CTP) consulting, design, and operational advisory services for cloud environments—has unparalleled ability to help your customers create a comprehensive IT strategy that successfully achieves the outcomes you seek.

Learning check

2. You are in a meeting with a customer, and they challenge you with a question: “What are the most important barriers of transforming IT

environment for an edge-centric, cloud-enabled, data-driven world?”
How should you respond?

HPE Hybrid IT

Although enterprises did not plan for a Hybrid IT world, it has become the answer to achieve speed and flexibility in our digital world.

Hybrid IT is the new normal—and it is complex

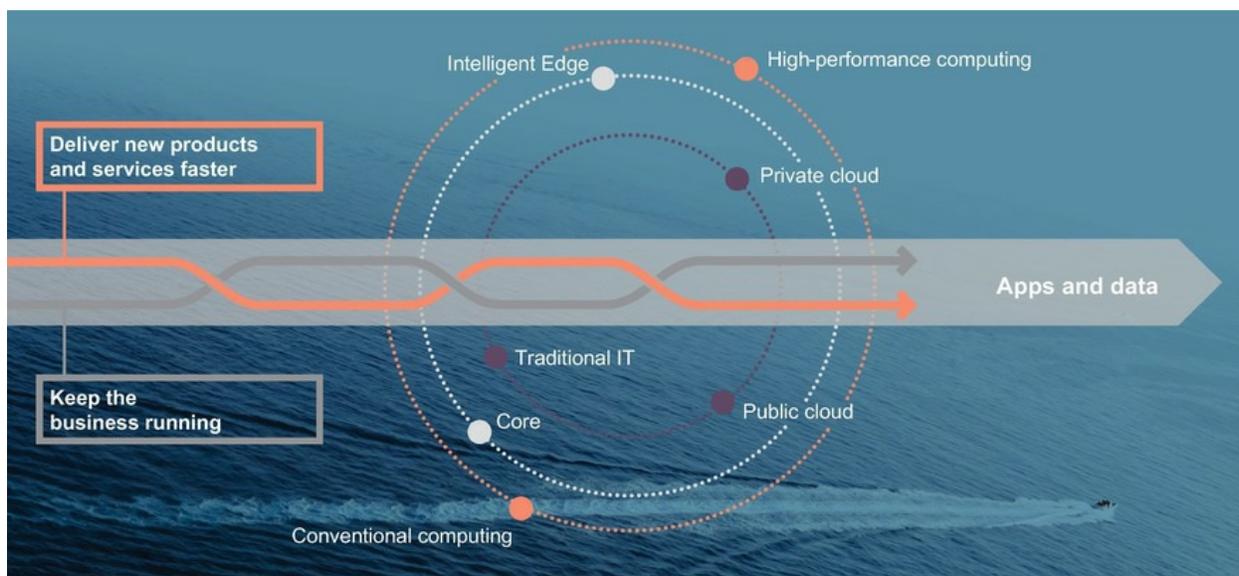


Figure 1-16 Operating customer apps and data across all technologies and platforms

To make Hybrid IT simple and optimal, as illustrated in [Figure 1-16](#), business must:

- Keep the business running and address the inefficiencies of **existing** environments (inflexible, legacy, siloes, overprovisioning, underutilization, capital-intensive) to free resources for innovation.
 - Enterprises continue to spend a disproportionate amount of budget keeping the lights on.
 - Rigid technology infrastructure was cited as the leading barrier preventing rapid innovation and speed to market (Source: HPE|Oxford Economics, January 2016).
- Deliver new products and services faster by speeding how you create **new** apps and services (agile, MVP strategies, containers, as a service) across mobile, IoT, AI, and more—amidst a growing ecosystem of partners.
 - HPE|Oxford Economics survey shows that companies leading their peers in pursuing rapid ideation also report superior financial results (Source: HPE Oxford Economics).
 - This demands a more experimental approach that is suited to the wealth of as-a-service models available, particularly pay-per-use models that help accelerate innovation and growth. Companies achieving growth do a lot of experimentation and build innovation into everyday operations. These companies take “lots of small steps” and readjust
- Create and deliver apps and data across multiple platforms and technologies, while understanding these critical new dimensions:
 - **Traditional, private, public cloud, multi-cloud**—Meeting control, cost, performance, and agility drivers
 - **Edge and core**—Aligned to responsiveness requirements (that is, a self-driving car cannot have decisions made in the cloud)
 - **Conventional and high-performance compute**—Systems must scale to new heights and at the right performance; a central learning engine is needed as today’s conventional computing architectures will not keep pace with the explosion of data

HPE believes achieving the optimal IT operating model for your customer’s apps and data will enable their business to move faster—with less risk. We make Hybrid IT simple for this reason:

- **As-a-service consumption, on-premises**
 - Pay-per-use consumption models and capacity planning strategies that limit risk and preserve capital
 - Pay for the resources you need, when you need them (whether you build or consume); this includes how you fund start-up ventures and new ideas—from experiment to strategic
- **Platforms from edge to core to cloud**
 - Modern Hybrid IT foundations that span the edge to cloud to core of the Hybrid IT environment
 - Software-defined infrastructure to speed development and delivery—and compose resources dynamically
 - Advanced data and analytics platforms that help turn data into insights and insights into action from all data
- **Invisible IT through automation and AI in the data center**—IT management that is AI driven from the edge to cloud to core
- **Ecosystem of partners**—Rich ecosystem of partners integrated across the Hybrid IT environment
- **Visibility across a multi-cloud environment**—Simplifying the complexity across the environment with unified management
- **Long-standing expertise to help bring it all together**—Helping customers bring it all together with HPE Pointnext, our services organization with the expertise to help them accelerate their Hybrid IT

How we make Hybrid IT simple

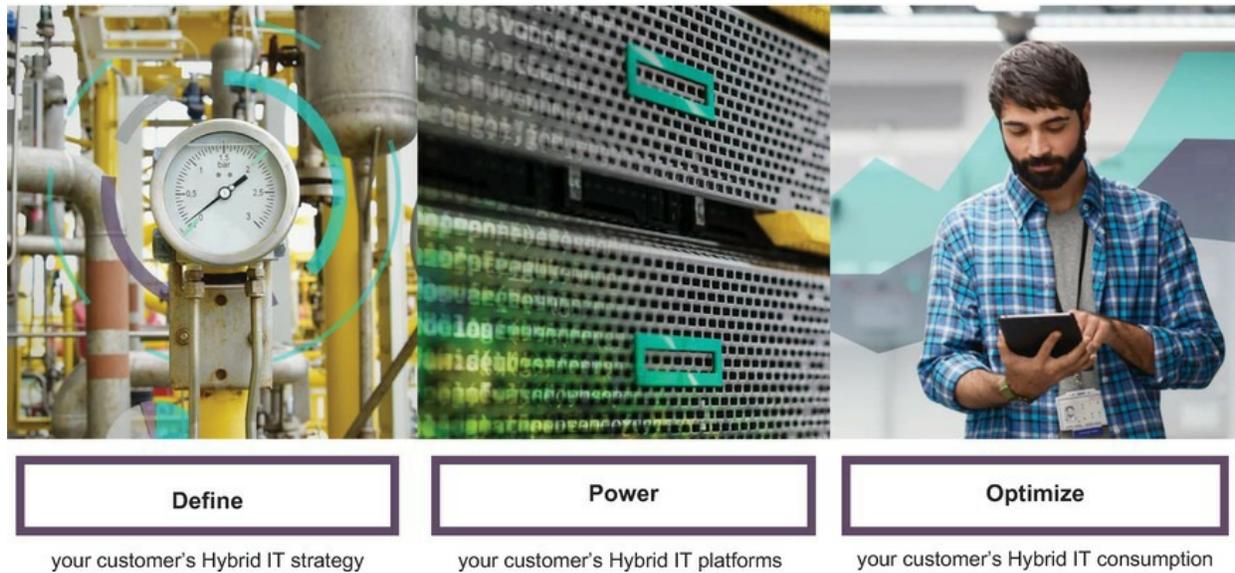


Figure 1-17 How we make Hybrid IT simple

HPE is making Hybrid IT simple with expertise, modern Hybrid IT foundations, software-defined infrastructure, advanced data and analytics platforms from core to edge. HPE is making Hybrid IT simple with innovative IT consumption and delivery models to help your customers get and pay for the resources they need, when they need them—with the control they need.

As shown in [Figure 1-17](#), HPE helps customers:

- **Define their Hybrid IT strategy** with expertise from HPE Pointnext and our partners to assess where customers are today, map their apps and data to their optimal platforms, migrate to their right mix of Hybrid IT, and continuously evolve.
- **Power their Hybrid IT** with a modern Hybrid IT foundation, the benefits of software-defined infrastructure, and turn insights into action with advanced data and analytics platforms.
- **Optimize Hybrid IT with consumption and delivery models** that give customers flexibility, speed, and on-premises control.

Define your customer's Hybrid IT strategy

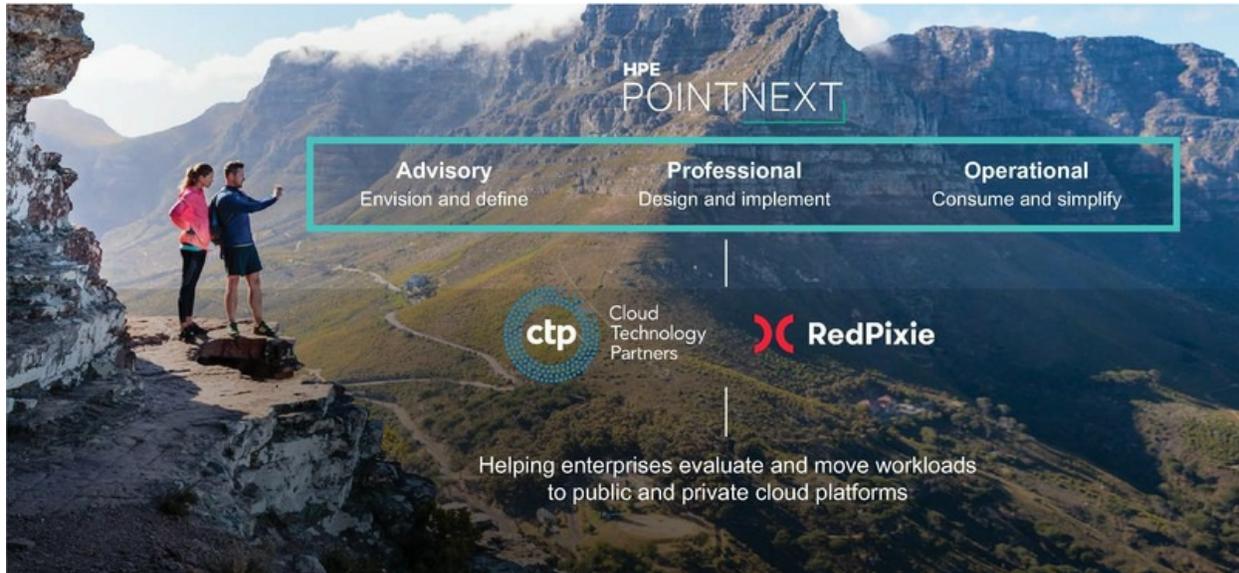


Figure 1-18 Expertise to help you accelerate your customer’s Hybrid IT strategy

With HPE Pointnext, we help you assess where you are today, help you map your customer’s apps and data to their optimal platforms—that deliver speed and flexibility for their business, as shown in [Figure 1-18](#).

Deliver speed and flexibility from apps and data

As we assess your customer’s current environment against their business goals, we help them drive efficiencies and free up resources. This includes modernizing, consolidating, automating IT (address siloes, sprawl) to drive better utilization and simplify and reduce their IT footprint. Driving efficiencies by modernizing IT helps them free resources for growth.

At the same time, we assess how customers can speed the creation and delivery of new services to their customers, developers, data scientists, and end users. This includes identifying ways to deliver app dev platforms with agility and speed; providing advanced data and analytics platforms to accelerate insights and action; and adopting emerging technologies and methodologies.

The Right Mix of Hybrid IT enables customers to pivot with

the business

HPE believes achieving the optimal IT operating model for your customer’s apps and data will enable their business to move faster—with less risk. We call it the Right Mix of Hybrid IT when you optimize:

- **Cost model**—Pay for the resources they need (whether you build or consume). This includes how they fund start-up ventures and new ideas—from experiment to strategic.
- **Platforms and as-a-service models (on- and off-premises)**—Give them the resources at the right cost, performance, control, and agility across the data center, private cloud, public cloud, and edge—securely.
- **Processes**—Simplify, automate, and accelerate the creation and delivery of your customer’s services, combined with the enabling technologies.
- **People**—Free up your customer’s resources for their most strategic initiatives, and tap new ways to drive collaboration in a digital workplace.

With HPE Pointnext, we do this with a 360-degree view across your customer’s cost model, platforms, processes, and people.

Power your customer’s Hybrid IT platforms



Figure 1-19 Power your customer’s Hybrid IT platforms

We help power your customer's apps and data with innovative HPE Hybrid IT platforms introduced in [Figure 1-19](#). To drive insights, speed, and action from their apps and data, enterprises must:

- Speed development and delivery with a software-defined infrastructure.
- Create a modern Hybrid IT foundation.
- Turn data into insights with advanced data and analytics platforms.

In a digital world, you must create and deliver services fast, and software-defined infrastructure helps your customers do just that.

Software-defined infrastructure is at the heart of HPE's strategy to make Hybrid IT simple, and we believe it is the foundation for Hybrid IT. HPE's software-defined infrastructure solutions enable customers to create and deliver services dynamically, with the control they need for their business. It enables them to:

- Compose, teardown, recompose resources as needed. For example, HPE IT provisions resources in three minutes, not weeks.
- Repurpose resources dynamically from one workload to another. For example, HudsonAlpha reprovisions resources to teams of researchers in two hours, instead of four days.
- Infinitely scale resources.

Let us look at how important your customer's Hybrid IT platforms are in driving insights, speed, and action from their apps and data.

The ability to pivot with the business and compose resources as they are needed is key to thriving in a digital world. HPE software-defined infrastructure solutions help businesses do just that.

HPE Synergy

HPE Synergy is the industry's first Composable Infrastructure and addresses challenges in your customer's existing environment **and** speeds the development and delivery of new services:

- Flexible resource pools of compute, storage, and fabric **in one infrastructure**

- Intelligently composes workloads across physical, virtual, and containers
- A single API that allows ops and developers to access the resources as code with developer environments and tools, including Docker, Mesosphere, Ansible, and Chef

HPE SimpliVity

The HPE SimpliVity hyperconverged infrastructure solution is designed from the ground up to meet the increased performance, scalability, and agility demands of today's data-intensive, highly virtualized IT environments. HPE SimpliVity technology transforms IT by virtualizing data and incorporating all IT infrastructure and services below the hypervisor into compact building blocks. With three times total cost of ownership (TCO) reduction, HPE SimpliVity delivers the best of both worlds— the enterprise-class performance, protection, and resiliency that today's organizations require, with the cloud economics businesses demand.

HPE OneView

Simplify your customer's Hybrid IT environment with the HPE OneView management platform that transforms servers, storage, and networking into a software-defined infrastructure. With HPE OneView, you can:

- Implement Infrastructure-as-code.
- Deploy infrastructure at cloud speed.
- Simplify operations with a unified application programming interface (API).

HPE ProLiant for Microsoft Azure Stack

With Azure Stack, your customer's business will experience lower costs, more flexibility, and the increased business agility that the cloud offers. The public Azure variant, already implemented by many, has more than proven itself in recent years. Stack is its extension, with a smaller footprint within an on-premise data center. Azure Stack has the same user interface and APIs as the Azure public cloud.

Start with a foundation that is fast, always-on, secure—ready for anything



Figure 1-20 Start with platforms that are efficient, always-on, and secure

Start with platforms that are efficient, always-on, and secure. HPE solutions do just that, as illustrated in [Figure 1-20](#):

- **Predictive analytics to anticipate and prevent issues with HPE All-Flash Storage and HPE InfoSight**—86% of issues automatically predicted and resolved
- **All-Flash storage for fast and reliable access to data with HPE All-Flash Storage**—99.9999% proven availability
- **Speed application performance with HPE Gen10 Servers—World’s Most Secure Industry- Standard Server**

Predictive analytics provide continuous digital operations by anticipating and preventing issues across the infrastructure stack with HPE InfoSight and HPE All-Flash Storage. HPE all-flash storage is the only storage smart enough to predict and prevent issues and close the app-data gap.

Fast and reliable access to data with HPE All-Flash Storage Predictive analytics, coupled with cloudready flash storage, deliver fast and reliable access to data both on- and off-premises.

Optimize your customer’s Hybrid IT financing and delivery



Figure 1-21 Optimize with as-a-service delivery and pay-per-use consumption

The fast pace of our digital world is driving change in how IT is delivered, paid for, and consumed. Today, too much capital is tied up in IT, making it difficult to get capital to fund new growth. Moreover, enterprises no longer want to (or have to) own IT assets.

Figure 1-21 shows HPE helping enterprises rethink the way they acquire, pay for, and consume IT to:

- Accelerate innovation with the right IT investment strategy (better align costs to revenues, monetize assets, smarter capital allocation).
- Start new projects faster, through pay-per-use and consumption-based models with on-premises control.
- Enhance the Hybrid IT service delivery with a global partner ecosystem.

Activity: Navigating the SMB Hybrid IT portfolio at the HPE website

1. Open the SMB Hybrid IT portfolio at: <https://www.hpe.com/info/smb>.

2. Answer the following questions:

a. What are the three important aspects of HPE solutions oriented to SMBs?

b. Name three families listed in the SMB portfolio.

Activity: Examples of a hybrid cloud solution

1. Open the HPE Community site at:

<https://community.hpe.com/t5/Shifting-to-Software-Defined/HPE-SimpliVity-integrationfor-Micro-Focus-Hybrid-Cloud/bap/7021447#.XQJJ9Rb7SHu>

2. Answer the following questions:

a. What are the two components in this example of a hybrid cloud solution?

b. What are the benefits of the integration provided by the HPE and Micro Focus product managers?

3. Open the HPE GreenLake Backup Brochure at:
<https://h20195.www2.hpe.com/v2/getpdf.aspx/a00005054ENW.pdf>

4. Answer the following questions:

a. What does HPE GreenLake Backup combine?

b. What is the metering unit HPE GreenLake Backup is using?

Learning check

3. How does HPE make Hybrid IT simple?

Summary

- HPE SMB offers capture best practices and incorporate the latest technologies optimized for foundational SMB use cases.
- HPE is a company built on the idea of innovation. HPE is accelerating the customer’s digital transformation agenda—transformation of technology, people and processes, and economics.
- Achieving the optimal IT operating model, Hybrid IT for apps and data will help enterprises to move faster and with less risk. HPE makes Hybrid IT simple for this reason.

Prelearning check

Before proceeding with this section, answer the following question to assess your existing knowledge of the topics covered in this chapter. Record your answers in the space provided.

- 1 You are in a meeting with a customer, and they challenge you with a statement: “I heard that the configuration process required to prepare a quote is time consuming, with a high of risk of errors and risk of ordering incompatible components.”

How should you respond?

2 Recommending HPE compute solutions for SMB customers

LEARNING OBJECTIVES

After completing this chapter, you should be able to recommend and position compute solutions for SMB customers.

HPE has it all

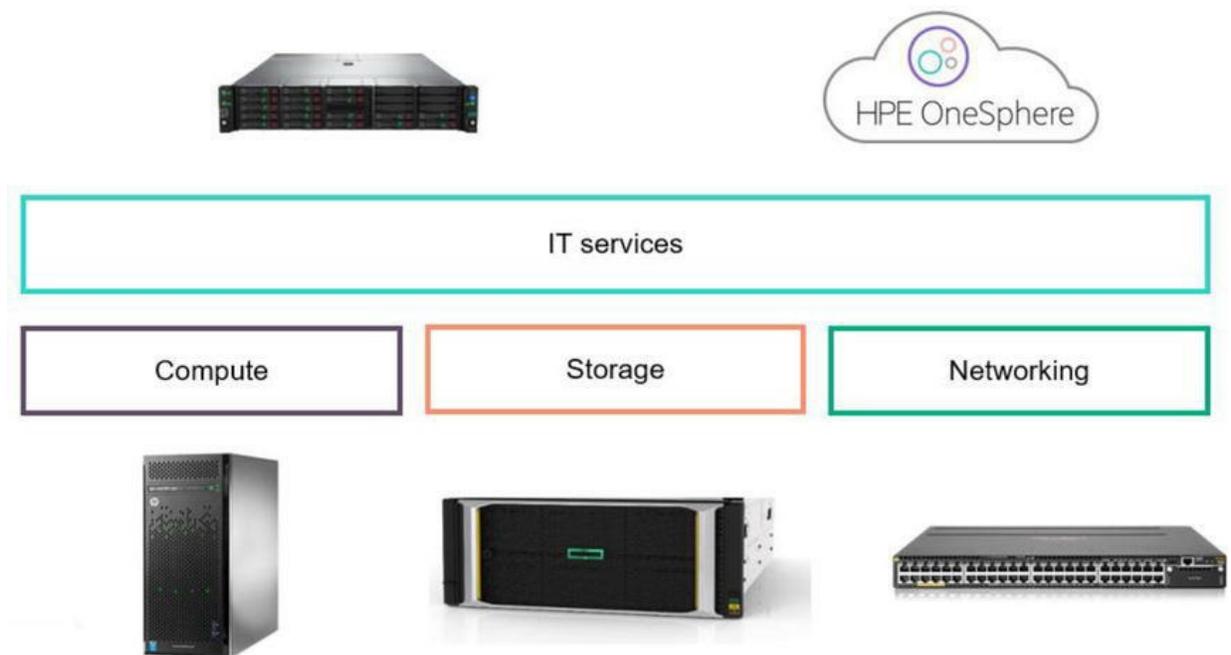


Figure 2-1 HPE has it all

Are you looking for industry-leading enterprise servers, storage, networking, software, or integrated systems products? Whether the IT service

consumption model includes public cloud, private cloud, hyperconverged solutions, or traditional systems, there are always compute, storage, and networking components in the data center that provide necessary resources, as shown in [Figure 2-1](#). In addition, HPE is providing any data center Hybrid IT component that is necessary for a successful deployment of IT service.

Why a single partner for compute, storage, and networking?

Using a single vendor at the rack level for compute, storage, and networking simplifies operations:

- **Plug-and-play integration**—No interoperability issues
- **Simplified logistics**—Single point of contact for procurement
- **Simplified support**—Single point of contact for server, OS, hypervisor, storage, and networking
- **Simplified management**—Unified administration

Using a single vendor for compute, storage, and networking improves business outcomes:

- **Lower operations expenses**—Unified management simplifies administration and reduces support burden.
- **Faster time-to-value**—Reference architectures and plug-and-play integration simplifies planning and accelerates deployment.
- **Lower risk/higher service levels/faster problem resolution**—No vendor squabbles or finger-pointing.

The solution design can be simplified by using validated HPE Reference Designs and Proactive Care Services.

Assessing customer requirements

We are presented with our protagonist, a solution architect (SA), who has just been hired at one of HPE's top partners. The SA gets a phone call at the desk from the sales team, and they are presenting a new opportunity to help a customer. They give the customer scenario and the problem. What would

they do first to help the customer?

This chapter will guide you through how and what to think about when gathering customer requirements, which tools to use, and how to make platform decisions.

Discovering a customer's business and technical requirements

When planning and developing an IT solution, the business and IT needs of the customer should be assessed. These assessment results should be used to guide the implementation planning process. When all customer information has been gathered, your experience and knowledge can be used to recommend the best possible solution. The IT recommendation is ultimately described in the statement of work (SOW). However, before a solid SOW can be developed, the specific needs, expectations, and environment for the solution must be understood. Thorough planning helps to avoid potential costly mistakes and prepares an upgrade path for the future.

Discovering a customer's business and technical requirements starts with a customer interview. Ask questions to determine current challenges, ways to address those challenges, and business goals. Consider asking questions in the following categories:

- Current environment
 - What applications are running and what are they expected to do?
 - How much storage is used currently?
 - Have storage needs grown over the last 12 months?
- Future plans
 - What are the business goals?
 - What is the projected role of the server and what operating system will it run?
 - To which kind of network will the server be connected?
 - Is there a long-term IT strategy in place?
- Resources

- Is the customer willing to commit resources to achieving these goals?
- Is the customer willing to let technical professionals help guide the way?
- Does the customer provide in-house training?
- Technical requirements
 - What is the expected availability of the server?
 - Will backups be performed?
 - Is power protection needed?
 - Does the server need to be configured or should it be ready to install out of the box?
 - What level of maintenance and support is desired?
- Obstacles
 - What is the biggest IT problem facing the business today?
 - What are the barriers to the solution?
 - Are there any open service calls or other customer sensitivities?

Based on answers to these questions, recommendations can be made about which server components are required and which are optional. For example, if file and print is the projected role for the server, storage capacity and transfer rate are important selection factors. Alternatively, if the server will be a database server, processor speed and memory are the primary considerations.

Issues to consider

In the process of discovering an SMB customer's business and technical requirements, you must consider several important issues:

- Determine the customer's business application and data requirements:
 - Latency requirements
 - Security
 - Determine dependencies

- Pay close attention to how much computing equipment is required.
- Understand the hardware components and sizing requirements.
- Consider nontechnical issues.
- Create specifications for the configuration options.
- Identify service levels, and determine how software and hardware will be used.
- Consider the customer's environmental requirements.
- Perform two key analyses:
 - Applications/workload analysis
 - Requirements analysis

Two analyses should be performed to thoroughly assess each customer:

- **Applications/workload analysis**—Analyzing the types of workloads the customer is running can provide guidance in selecting a server family. It is important to find out if the customer needs support for virtualization, cloud, web infrastructure, database, app development, and so on. For example, SMB customers requiring servers for a small IT infrastructure should consider the HPE ProLiant ML family.
- **Requirements analysis**—Using the answers to the business environment evaluation can guide the selection of a particular HPE server portfolio. For example, if the customer needs to start small with room to grow, HPE ProLiant DL servers should be considered. If the customer needs high server density, HPE Synergy solutions should be considered. Specific factors to evaluate include power and cooling requirements along with interconnectivity.

Assessing the customer's applications

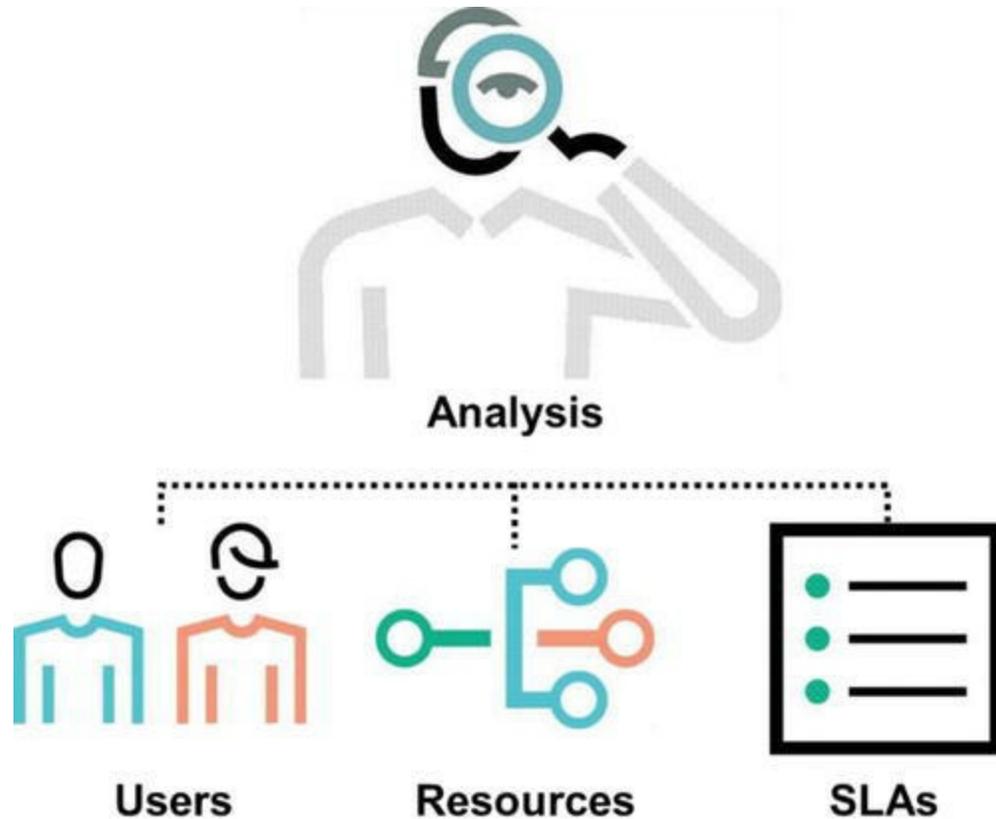


Figure 2-2 Examine the customer's existing resources, workloads, and overall operations

When you are planning an HPE solution, one of the first tasks to perform is an analysis of the customer's current infrastructure and the applications the customer plans to run on the servers. As we can see from [Figure 2-2](#), several factors should be considered:

- **The number of users for the application**—The number of users has a direct impact on the amount of CPU, memory, networking, and storage resources that will be required.
- **IT resources required by the application**—The application installation or user guide should offer recommendations regarding the amount of CPU, memory, networking, and storage resources that will be required for the application.
- **Applications that can be consolidated**—Are there any applications that no longer provide business value? If so, could they be retired, or could a different application deliver the desired results?

- **Service-level agreements (SLAs) in place for the various business organizations**—SLAs have a direct impact on the hardware and software that will be required. For example, a solution capable of achieving 99.999% availability will need to be designed with no single points of failure (SPOFs) and will require redundant components, duplicated systems, and clustering software. If the solution is not considered business critical and is only expected to achieve 99% availability, then redundant components, duplicated systems, and clustering software are not normally required.
- **Customer's current methodology**—Assess the customer's current infrastructure and processes to ensure that any new solution will integrate with their existing framework.

Solution design considerations

When applying the information you have gathered from the customer, consider the following areas:

- **Servers**—Will the solution consist only of physical servers or a mix of physical and virtual servers? If there will be a mix, which virtualization technology will be used? For physical servers, memory and processor technology components are important. Fault-tolerant memory and redundant processors are less crucial to a file and print server than they are to a database server, which performs computations and requires temporary storage. Determining the relative importance of server technologies and the server's projected role narrows the focus to a server with the required capabilities.
- **Storage**—Every server can have storage, but deciding on a RAID or non-RAID configuration depends on factors such as cost, storage availability, and fault tolerance. For a file and print server, a non-RAID configuration leaves data vulnerable to disk failure or data corruption. This would be less important to a network-centric firewall server.
- **Networking**—Depending on the existing network topology or the decision for a new topology, server networking capabilities, wireless devices, and edge devices must also be determined. Current data center networks are at least 10 Gb/s and can reach 25/50/100 Gb/s.

- **Operating system**—The choice of operating system directly affects the server components. As a general rule, the more recent (and thus more advanced) the operating system, the greater its demands on system hardware. Certain operating system features can also steer the decision. Potential server purchases should be made based on careful consideration for meeting or exceeding the highest minimum system requirements.
- **Applications**—Often referred to as workloads, the applications that the customer requires have a major impact on several design considerations. Understanding the nature of the workloads early in the process can reveal components that should be present in the solution. For example, if the main application is a business-critical database application, it is important to ensure that the server solution is designed with minimal points of failure and with sufficient storage, networking, memory, and processor resources to ensure the smooth operation of the database application.
- **Availability**—Workloads should be assessed for their level of business importance and housed on an appropriately available server solution. Availability clustering should be considered to ensure that if there is a catastrophic failure, applications can continue to run (by failing over to a standby or secondary system). Some applications might not be considered business critical and can, therefore, be unavailable without significant business impact. These might be located on nonclustered virtual or physical servers.
- **Security**—When you are planning where to place the server and how it should be configured, it is important to consider security. Be alert for physical and virtual security holes. When an employee leaves an organization, it is important to recover any keys and access cards. It might be necessary to change locks and codes. Disabling the user account and changing high-level passwords to which an employee had access are good practices.
 - **Physical**—Security measures also involve locks, codes, and location. Deciding to place a server in an interior room with a locked door sufficiently addresses most physical security needs. Because the temperature and humidity in a windowless interior room remain relatively constant, there should be no need to keep the door open to enhance airflow. A closed and locked door ensures that only

individuals with authority and access can enter.

- **Virtual**—Passwords, permissions, and access control lists should also be secured. Setting up users, groups, and permissions addresses virtual security needs. Each user needs a password to access project files stored on the server, and being a member of a particular group allows or denies access to other network resources. Grant each user only as much access as they need. All passwords should be changed regularly; meet the minimum character length (as defined by the operating system); and include letters, numbers, and special characters.
- **Scale-out vs. scale-up**—Scale-out approach means adding more servers to do the workload, while scale-up approach adds more resources (such as processors and memory) to already existing servers. Adding an additional two-socket system to the VMware cluster is considered scale-out. Adding two additional processors to the four-socket system currently equipped with only two processors is considered scale-up approach.

Learning check

1. What are typical storage considerations when assessing customer requirements?

Recommending HPE compute solutions for SMB customers

The wide HPE portfolio of data center products for SMB customers can be introduced using customer scenarios.

Introducing the customer scenario

The fictional customer LLP distribution company will be used as a storyline through this scenario. We will introduce the company using an interview:

- What is your primary business?
 - City-wide courier, parcel, and documents delivery
- How many employees do you currently have?
 - 25 employees + external workers
- What does your selling and delivery channel look like?
 - Contracts with dozens of local e-shops and expanding
- What does your server, storage, and network infrastructure look like?
 - SaaS (email, collaboration, CRM)
- Do you have an IT department?
 - No, external help on request
- What are your current plans?
 - Deploying a CRM and tracking onsite (speed, regulatory requirements —General Data Protection Regulation [GDPR])

Customer requirements

As an HPE presales consultant or partner, your primary goal is to find and convey value to your customers. You need to learn as much as you can about the customer's business situation and needs. To fully understanding their needs, you should be able to:

- Communicate value by providing business insights.
- Identify other ways to support the customer through cross-selling.
- Develop and maintain long-term relationships that can lead to more

business.

- Distinguish HPE from its competitors.
- Be the customer's first call when they need support.

It is so important that we make every single customer interaction as valuable as it can be. We can do this by knowing our customers, understanding our values, making ourselves relevant to our customers, and delivering meaningful business outcomes.

There are five stages to the consulting process. Keep these notes in mind as you progress through each stage:

1. Prepare

- Learn as much as possible before you begin to engage a customer.
- Research the company and the people, find out about areas of concern, and decide who is the best person to contact.

2. Interview

- Know what you want to learn before the initial interview with the customer.
- Consider the person's position at the company and anticipate their concerns accordingly.
- Listen to what the customer has to say.
- Use time effectively and address their concerns in a timely manner.

3. Plan

- Determine which solutions are closely aligned with the customer's business needs.
- Plan the next steps in the project, including who to add to your team when the time comes.

4. Propose

- Design the solution that meets the customer's needs and successfully conveys value.

- Do not propose products only; recommend a complete solution that demonstrates value at all levels of the organization.

5. Present

- Articulate the recommended solution, positioning it in terms that the customer understands.
- Ensure that the customer knows HPE has the solutions to meet their needs.

As a result of multiple interviews and gathering information about customer plans and customer's current infrastructure, the following requirements emerged for the new solution:

- Server for a custom-made CRM and tracking application
 - Requires RHEL or SLES
- Single server, on-site installation
- No local data center, sufficient space in a storeroom
- Single-phase power
- 1 Gb/s Ethernet connection to third-party router
- 20 GB of data anticipated per year, 100% grow rate annually
 - Design the solution for five years with sufficient room to grow
- Testing performed by supplier, and the following compute characteristics were verified as sufficient:
 - A single eight-core processor running at 2.1 GHz
 - 128 GB RAM
- Tape backup for long-term data protection and retention

Activity: Discovering a customer's business and technical requirements

1. Prepare a list of additional questions to ask the customer about:

Selecting an HPE server platform

Depending on workload, performance, capacity, and features required following decision trees, as seen in [Figure 2-3](#), can be used to select appropriate server, storage and networking platforms:

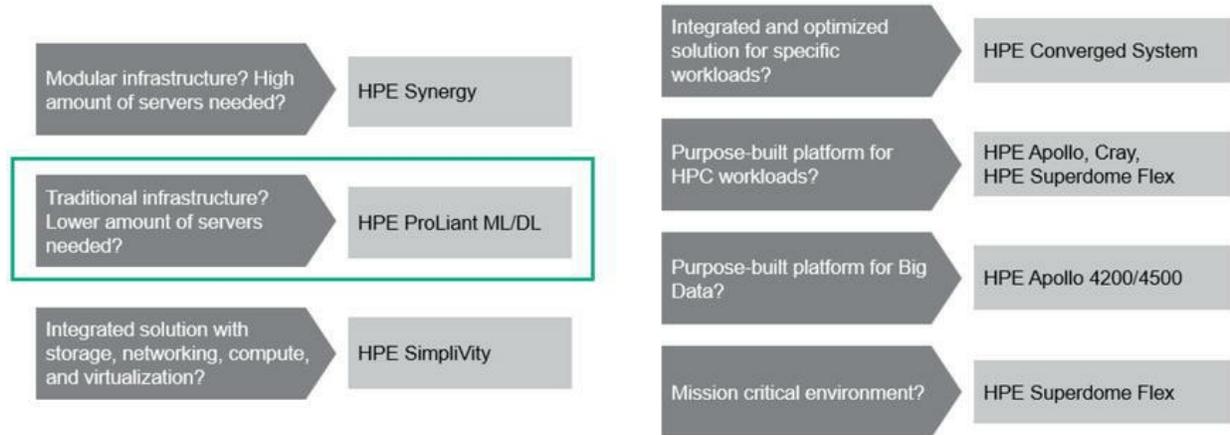


Figure 2-3 Selecting an HPE server platform



Note

Most platforms are out of the scope for this chapter; details focus on ProLiant ML/DL.

Selecting a Gen10 DL or Gen10 ML server

Depending on requirements shown in [Figure 2-4](#), an ML or DL platform can be selected:

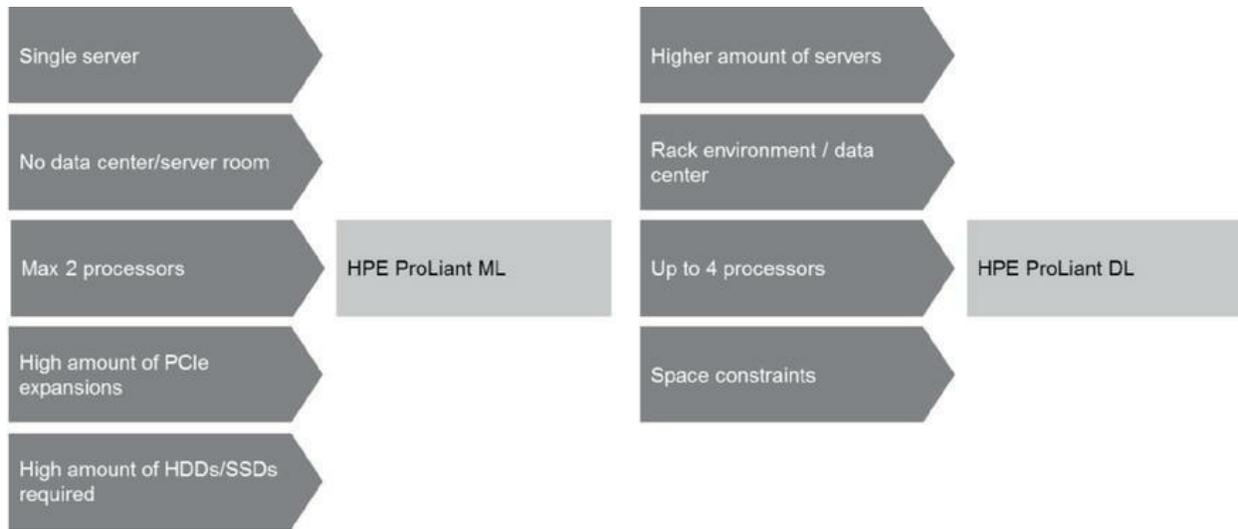


Figure 2-4 Selecting a Gen10 DL or Gen10 ML server

HPE ProLiant tower servers

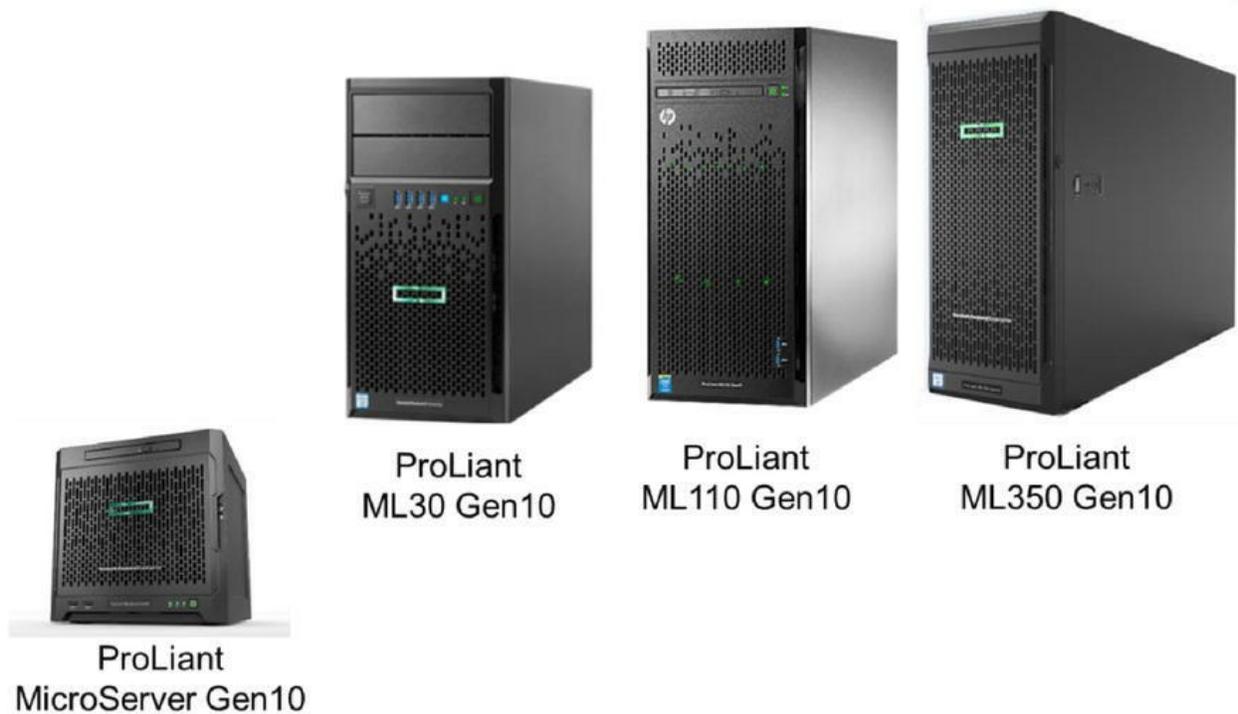


Figure 2-5 HPE ProLiant tower servers

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 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**—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, 100, 200, 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 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 portfolio delivers business value, helps 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 IO flexibility. They are also available in rack deployment options. 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
- Efficiency that office managers need to help improve employee productivity
- Affordability to increase business agility and help acquire and retain customers

As illustrated in [Figure 2-5](#), the ProLiant ML Gen10 family includes the following server models:

- ProLiant ML350 Gen10
- ProLiant ML110 Gen10
- ProLiant ML30 Gen10
- ProLiant MicroServer Gen10

Selecting an HPE ML server platform

Depending on workload, performance, memory and storage requirements, and requirements for expansion, the following decision trees can be used to select an appropriate Gen 10 ML server:

Performance, memory, and storage requirements



Figure 2-6 Selecting an HPE ML server platform, based on performance, memory, and storage requirements



Note

Other limitations and prerequisites apply. Consult QuickSpecs.

Availability and expansion



Figure 2-7 Selecting an HPE ML server platform, based availability and expansion requirements



Note

Other limitations and prerequisites apply. Consult QuickSpecs. Some

features require additional components.

ProLiant ML110 Gen10 Server



Figure 2-8 ProLiant ML110 Gen10 Server

The ProLiant ML110 Gen10 Server, as shown in [Figure 2-8](#), 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 Peripheral Component Interconnect express (PCIe) slots, eight large form factor (LFF) or 16 small form factor (SFF) disk drives, and eight Double Data Rate 4 dual in-line memory module (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.

ProLiant's most powerful and versatile 1P tower

Key selling points:

- Right-sized, with balance of performance to meet growing SMB requirements
- Expandability to fit the needs of growing, budget-conscious businesses
- Security innovations, industry-leading services, and ease of deployment

ML110 Gen10 is ideal for:

- The HPE ProLiant ML110 Gen10 Server is the perfect first server for growing small businesses and ROBOs. Workloads include email, file and print, web serving, networking, application development, and business processing workloads (enterprise resource planning [ERP], customer relationship management [CRM]).
- This affordable tower server delivers expansion, reliability, and performance with proven enterprise- class capabilities.

Price positioning:

- Standard warranty (3-3-3)
- US ILP starting at \$1490 for entry SKU and \$1815 for performance SKU (US ILP is preliminary and subject to change without notice)

Key competitive advantages:

- Redundant fan
- Supports up to eight LFF NHP HDDs

- Greater GPU support (supports up to two GPUs)
- Serial port
- Integrated Lights-Out (iLO) service port
- Flex Slot power supply unit (PSU)
- Industry-leading security

The ProLiant ML110 Gen10 features include:

- One Intel® Xeon® Processor Scalable Gold, Silver, Bronze Series, up to 14 cores, 105 W
- 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
- Standard HPE Dynamic Smart Array S100i, HPE Smart HBA Controllers for performance or additional features
- 16 SFF (8*2 cages) HP/8 LFF (4*2 cages) HP/NHP max, HDD/SSD m.2 enablement
- 2x1 GbE embedded + Standup
- Rear VGA and Optional Serial, 8 USB 3.0, 1 Micro-SD
- One Std ATX Multi-Output PSU 350 W / 550 W up to 85% efficiency
- Up-to two HPE Flex slot RPS 94 to 96%, RPS option
- A single-wide or double-wide graphics processing unit (GPU) support
- Front iLO Service Port, HPE iLO 5, HPE Smart Update Manager (SUM), HPE RESTful Interface Tool, Unified Extensible Firmware Interface (UEFI)
- HPE Secure Encryption, Trusted Platform Module (TPM) option
- Replaces ML150 Gen9 and ML110 Gen9



Note

For more information go to:

HPE.com/servers/ml110-gen10

Front system detail: LFF and SFF chassis shown

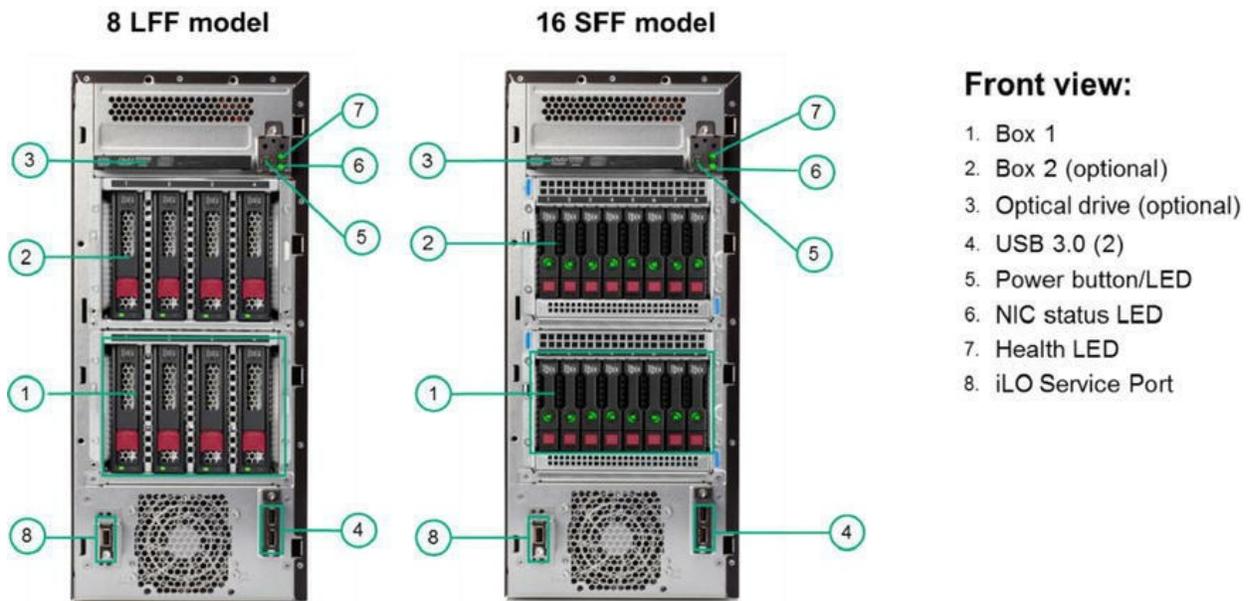
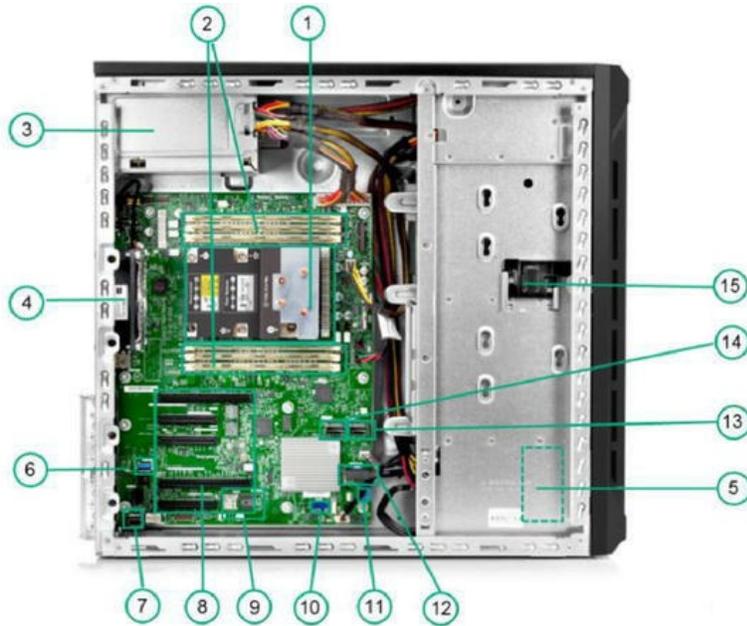


Figure 2-9 Front system detail: LFF and SFF chassis shown

Interior system detail



Interior view:

1. One processor and heatsink
2. Six DDR4 DIMM slots
3. ATX Power Supply
4. System fan (92 x 32 mm default)
5. PCIe fan (92 x 32 mm default)
6. Internal USB 3.0 connector
7. Internal USB 2.0 connector
8. Five PCIe3 expansion slots
9. MicroSD slot
10. SATA Port 10
11. SATA Port 9
12. Front USB 3.0 connector
13. X4 SATA Port 1
14. X4 SATA Port 2
15. Front Bezel Lock

Figure 2-10 Interior system detail

Rear system detail

Rear view:

1. PCIe3 slots (slots 1 to 5)
2. USB 3.0 (2)
3. USB 2.0 (2)
4. Network RJ-45 ports (2)
5. Video connector
6. UID LED/button
7. Power supply bay
8. Power supply connection
9. iLO management port
10. Serial port (optional)

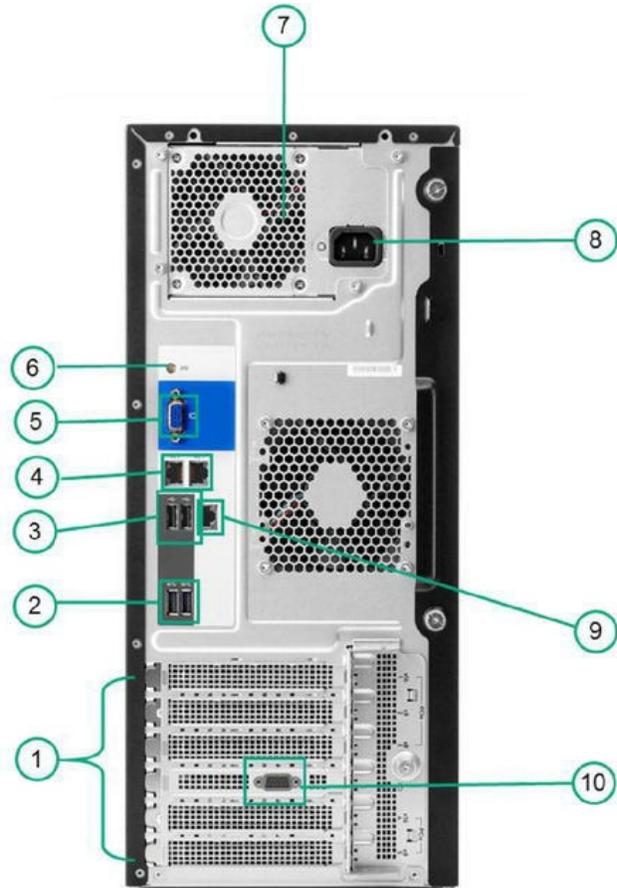


Figure 2-11 HPE ProLiant ML110 rear view

HPE StoreEver LTO Ultrium Tape Drives



Figure 2-12 HPE StoreEver LTO-8 Ultrium 30750 External Tape Drive

The family of HPE LTO Ultrium tape drives is HPE's premier line of backup and archive devices. Based on the Linear Tape-Open (LTO) Ultrium format, an open standard with a well-defined 12-generation roadmap, HPE LTO Ultrium tape drives offer customers the best choice for investment protection. With ultimate reliability and ease of use in mind, even at 100% duty cycles, HPE's rugged design builds on superior LTO technology and adds advanced features like hardware data encryption to create a new level of data protection. The tape drives employ dynamic data rate matching to adjust to the speed of the host, reducing wear and tear on both drive and media and increasing performance, even on slower hosts.

HPE StoreEver LTO-8 Ultrium 30750 Tape Drive

The HPE LTO-8 Ultrium 30750 tape drive, as referenced in [Figure 2-12](#), represents HPE's eighth generation of LTO tape drive technology capable of storing up to 30 TB (compressed 2.5:1) per cartridge while providing enterprise tape drive monitoring and management capabilities with HPE TapeAssure, ease of use with support for Linear Tape File System (LTFS) and Advanced Encryption Standard (AES) 256-bit hardware data encryption easy-to-enable security to protect the most sensitive data and prevent unauthorized access of tape cartridges. Capable of data transfer rates up to 300 MB/sec (native), its Speed Matching feature further optimizes performance by matching the speed of the host to keep drives streaming and increase the reliability of the drive and media. HPE LTO Ultrium half-height drives are designed for server customers in direct-attached storage (DAS) environments. For IT organizations with stringent, long-term data retention requirements, the Ultrium 30750's Write-Once, Read-Many (WORM) data protection feature offers a simple and secure method for archiving records as part of a compliance solution to meet the toughest of industry regulations. The Ultrium 30750 tape drives provide investment protection with full read and write backward support with LTO-7 media. By doubling the native capacity of previous generation Ultrium drives, HPE customers now require fewer data cartridges to meet their storage needs, significantly reducing their IT costs and increasing their return on investment (ROI).



Note

HPE StoreOpen Standalone and LTFS software are available for free downloads: <http://www.hpe.com/storage/ltfs>

Reliability: HPE StoreEver innovation



Figure 2-13 HPE StoreEver LTO-8 Ultrium 30750 Internal Tape Drive

HPE StoreEver is designed for reliability and long-term data protection and retention.

HPE proven reliability and performance features:

- Built on LTO-8 standard format, as illustrated in [Figure 2-13](#)
- Robust store and restore over extended periods
- Energy efficient
- Performance optimizing with data rate matching
- HPE Tape media reliability with a 30-year shelf-life

Secure: Hardware-based data encryption

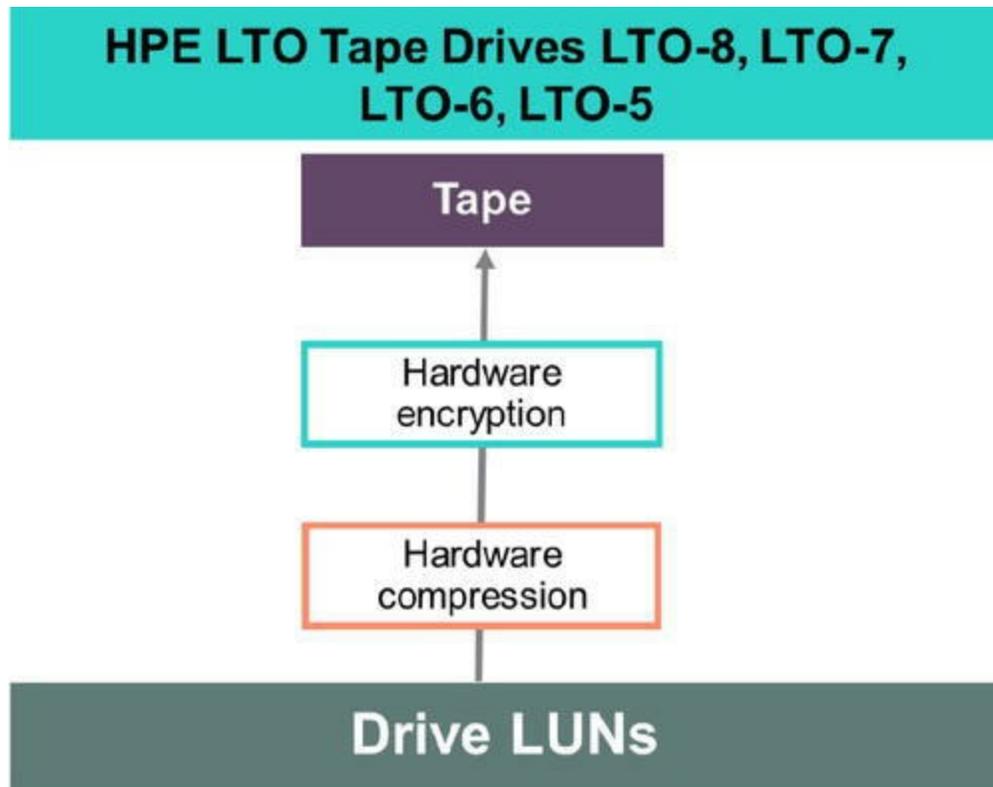


Figure 2-14 Compliant security for data at rest

Hardware-based data encryption means:

- No performance degradation for compression or encryption
- Compression occurs before encryption for optimum results
- Encryption keys are supplied by backup applications or key management appliances
- Comply with legislation—granted FIPS140-2 Level 1 certification (LTO-5, LTO-6, LTO-7), as indicated in [Figure 2-14](#)

Activity: HPE QuickSpecs: HPE ProLiant servers

1. Download this information resource:
<https://h20195.www2.hpe.com/v2/getdocument.aspx?docname=a00021852enw>

2. Answer the following questions:

a. What are the two differences between the entry LFF model and the base LFF model?

b. How much memory does the high-performance tower model with second generation of the Intel[®] Xeon[®] Scalable Processor have?

Activity: HPE Server Support and Certification Matrices

1. Access this information resource:

Red Hat Enterprise Linux Certification and Support:

<https://h17007.www1.hpe.com/us/en/enterprise/servers/supportmatrix>

SUSE SLES Certification and Support:

<https://h17007.www1.hpe.com/us/en/enterprise/servers/supportmatrix>

2. Answer the following questions:

a. Is the DL20 Gen10 supported and certified by SUSE Linux Enterprise Server 15?

b. Is a 64-bit version of Red Hat Enterprise Linux 8 certified on ML110 Gen10?

Learning check

- 2. A customer requires a tower server with redundant, hot-plug fans for an environment without a server room. Which server supports these requirements?
 - A. ML350
 - B. ML110
 - C. ML30
 - D. DL380

Components of a compute solution



Figure 2-15 Server subsystems

As we can see from [Figure 2-15](#), typical HPE x86 server contains several subsystems that we will cover in this material, such as:

- Processor
- Memory
- Storage
- Networking

Processors in HPE servers

HPE servers are equipped by Intel[®] Xeon[®] Processor Scalable Family, AMD EPYC, and Marvel ThunderX2 ARM.

Unified Intel[®] Xeon[®] Scalable Platform

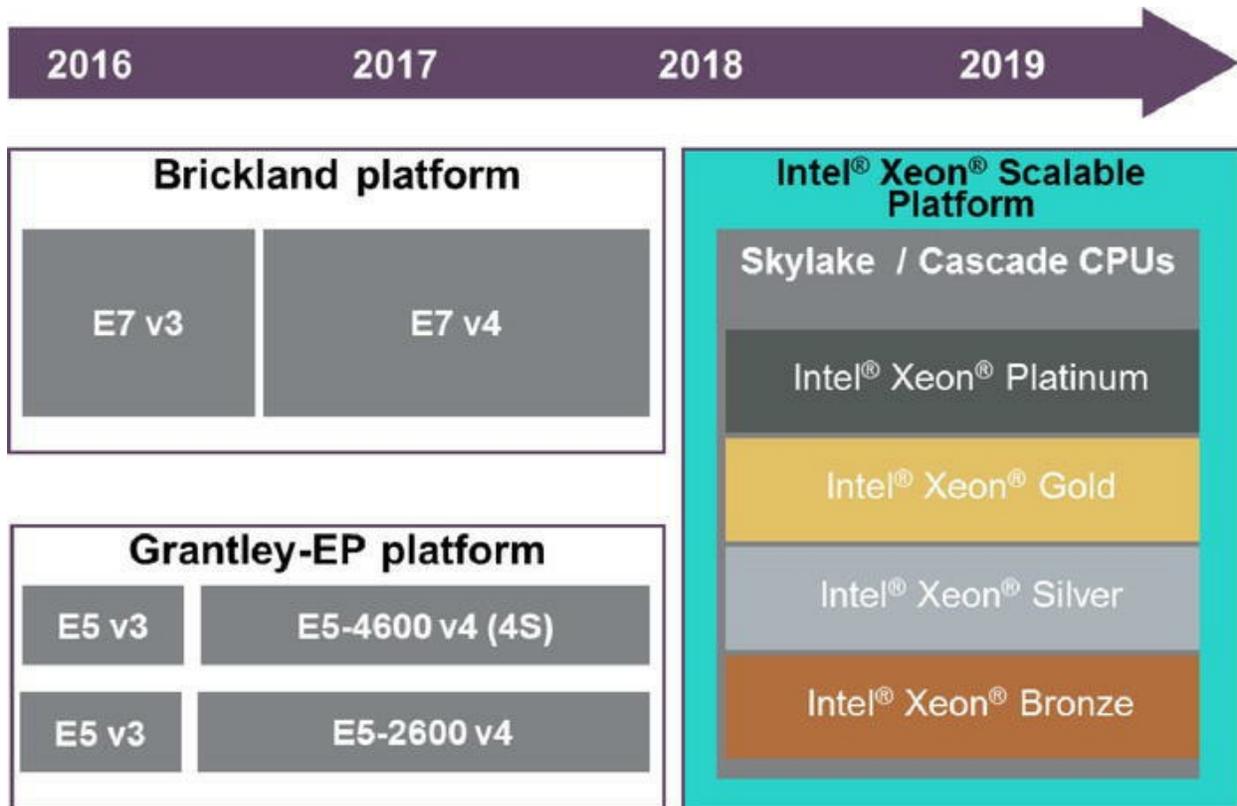


Figure 2-16 Intel® Xeon® Scalable Platform

Intel® Xeon® Processor E7 was targeted at mission-critical applications that value a scale-up system with leadership memory capacity and advanced reliability, availability, and serviceability (RAS). Intel® Xeon® Processor E5 was targeted at a wide variety of applications that value a balanced system with leadership performance/watt/\$.

As illustrated in [Figure 2-16](#), five model families are available with Intel® Xeon® Processor Scalable Family replacing former Intel® Xeon® E7 (4/8S+) and E5 (2S, 4S) families:

- Platinum
- Gold (6xxx)
- Gold (5xxx)
- Silver
- Bronze

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

Intel® Xeon® Processor Scalable Family overview

Features	Intel® Xeon® E5-2600 v4	Intel® Xeon® Processor Scalable Family (Skylake / Cascade 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 / 2933
TDP (W)	55 W to 145 W	70 W to 205 W

Figure 2-17 Intel® Xeon® Processor Scalable Family overview

Figure [Figure 2-17](#) offers comparison between Intel® Xeon® Processor Scalable Family overview and Intel® Xeon® E5-2600 v4. 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, we are moving to a metal-based SKU system: Platinum, Gold, Silver, and Bronze.

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 service ability.

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).

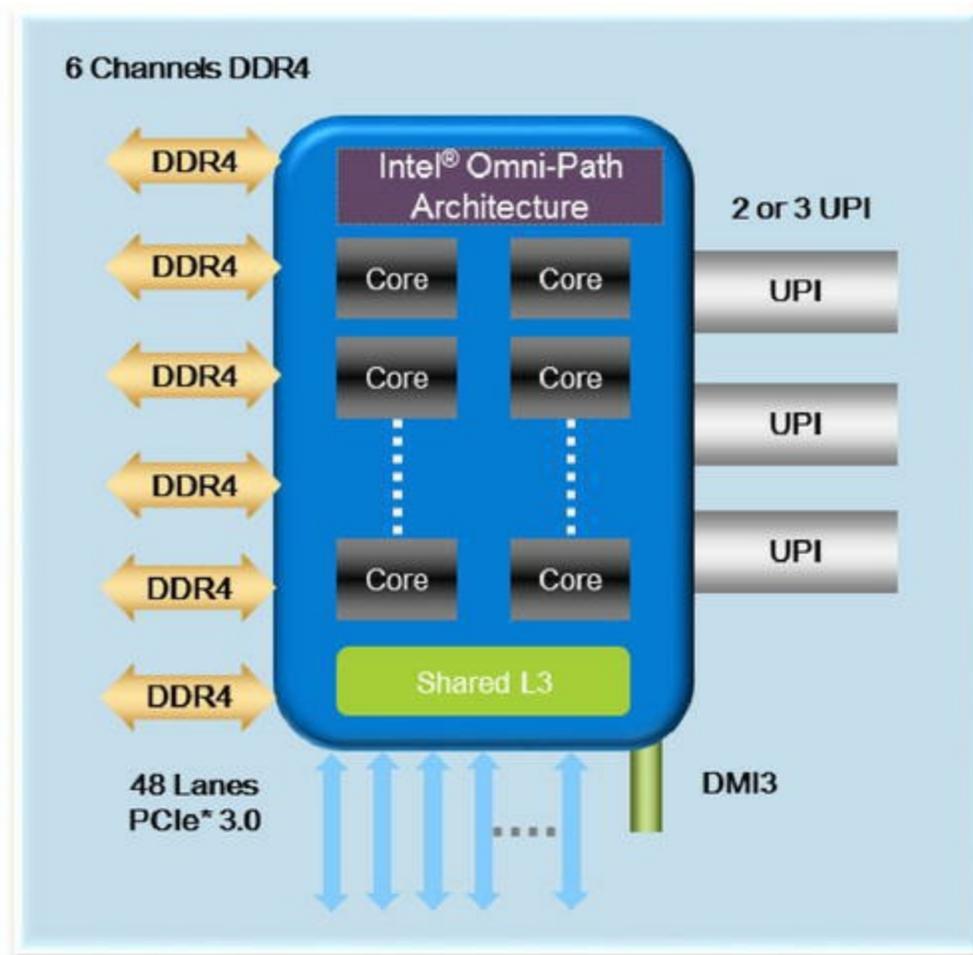


Figure 2-18 Intel® Xeon® Processor Scalable Family: Architecture

This generation introduces increased memory capacity and performance due to the six-memory channel design and the advanced inter-processor communication Ultra Path Interconnect (UPI), as indicated in [Figure 2-18](#).

Intel® Xeon® Scalable Family processor numbering and features

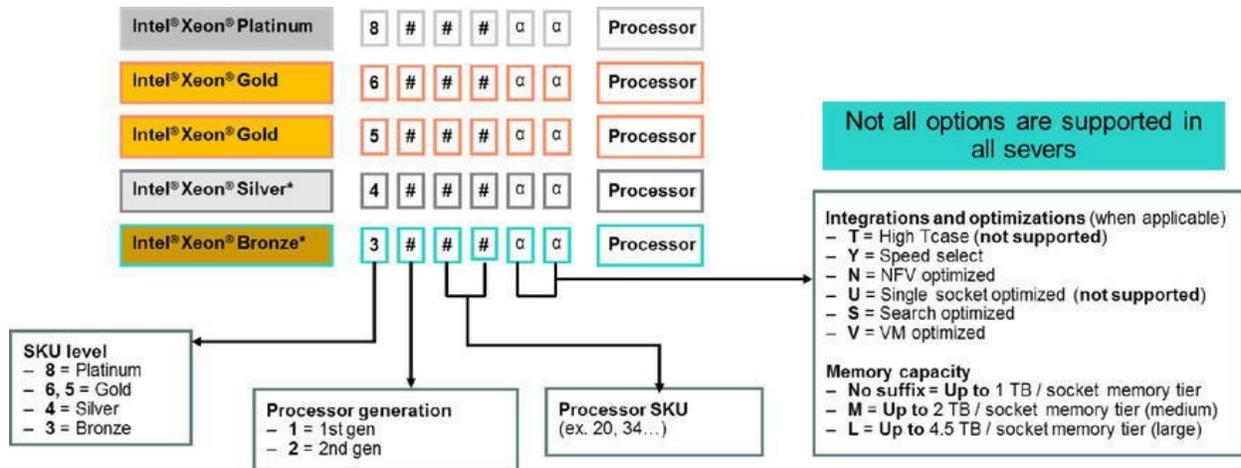


Figure 2-19 Intel® Xeon® Scalable Family processor numbering and features

Figure 2-19 illustrates the Intel® Xeon® Processor Scalable Family product numbering of Gen1 and Gen2 processors.

Select the Intel® Xeon® processor

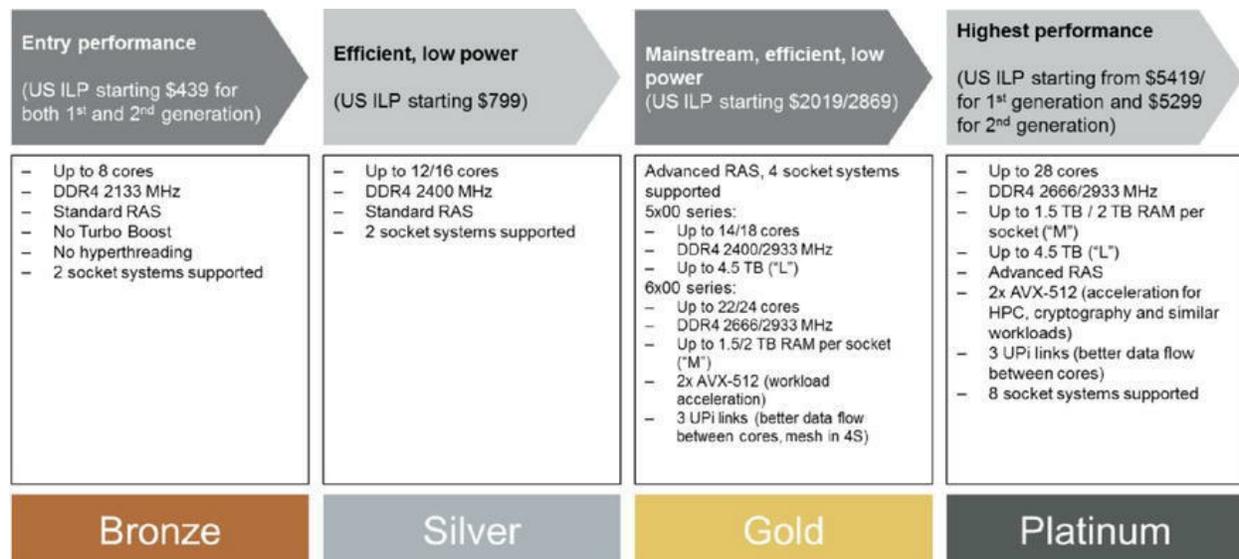


Figure 2-20 Gen1/Gen2 Intel® Xeon® Processor

Select the processor based on customer and business requirements using the information shown in Figure 2-20.



Note

Numbers above can change over time, so always verify latest information in QuickSpecs. US ILP are subject to change without notice. Support must be verified for specific models as well. Not all variants are available for all models.

AMD EPYC 7000 Processors—A new building block



Figure 2-21 Zen architecture

AMD is a key partner in the HPE compute portfolio, and HPE is offering key platforms that will best utilize the AMD processor technology and address a variety of workloads. HPE is further extending the ProLiant Gen10 portfolio with the introduction of single-socket and dual-socket servers based on the AMD EPYC processors, such as the HPE ProLiant DL325 Gen10 or HPE ProLiant DL385 Gen10.

This platform not only provides the upmost in security, but it also delivers 2P performance at 1P economics.

With today's application and data demands, mid-market organizations have been forced into buying up to a two-socket server to achieve necessary levels of performance. Single-socket servers have traditionally shipped with low-performance processors, limited memory, and IO.

The AMD Secure Processor is a dedicated security processor embedded in the AMD EPYC system on a chip (SoC). The security processor manages secure boot, memory encryption, and secure virtualization.

With respect to secure boot, the HPE Silicon Root of Trust ties to the AMD Secure Processor at the firmware level. The AMD Secure Processor also validates the HPE BIOS during the boot process.

With respect to Secure Memory Encryption (SME), you can encrypt all of the memory or a portion of memory. With partial memory encryption you can encrypt a subset of memory, for example, marking memory used by guest virtual machines (VMs) as encrypted.

With respect to Secure Encrypted Virtualization (SEV), there are encryption keys in the security processor that never leave the processor where they can be exposed to intruders. With SEV, VMs have separate encryption keys as does the hypervisor, isolating the VMs from other VMs, and even the hypervisor itself. SEV is supported by SUSE.

Important features of AMD EPYC platform, as illustrated in [Figure 2-21](#):

- Compute
 - Eight to 32 or 64 AMD “Zen” x86 cores (up to 128 threads)
 - 512 KB L2 cache per core (16 MB total L2 cache)
 - 64 MB shared L3 cache (8 MB per four cores)
 - Thermal design power (TDP) range: 120 W to 180 W
- Memory
 - Eight channel DDR4 with ECC up to 2667 MHz
 - Registered DIMM (RDIMM) and Load-Reduced DIMM (LRDIMM), 3DS, NVDIMM, Flash

- Two DIMMs/channel capacity of 2 TB/socket
- Integrated IO (no chipset)
 - 128 lanes PCIe Gen3
 - Used for PCIe, serial ATA attachment (SATA), Ethernet, and Coherent Interconnect
 - Up to 32 SATA or NVMe devices
 - Up to 16 10 GBASE-KR or 1 Gb Ethernet
 - Server controller hub (USB, UART, SPI, LPC, I2C, and so on)
- Security
 - Dedicated security subsystem
 - Hardware root of trust

Learning check

3. Which processor will support a four-socket system? (Select three.)
- A. Silver 4208 Processor
 - B. Gold 5215L Processor
 - C. Gold 6240Y Processor
 - D. Platinum 8280 Processor
 - D. Pentium Gold G5400

Memory for HPE servers

Recently, huge increases in the power of CPUs have left the rest of the server architecture lagging. 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

dynamic random access memory (DRAM). While solid state disks (SSDs) seek to eliminate the delay incurred during the mechanical operation of hard disk drives (HDDs), 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; 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 non-volatile dual in-line memory module (NVDIMM).

Memory/Storage hierarchy

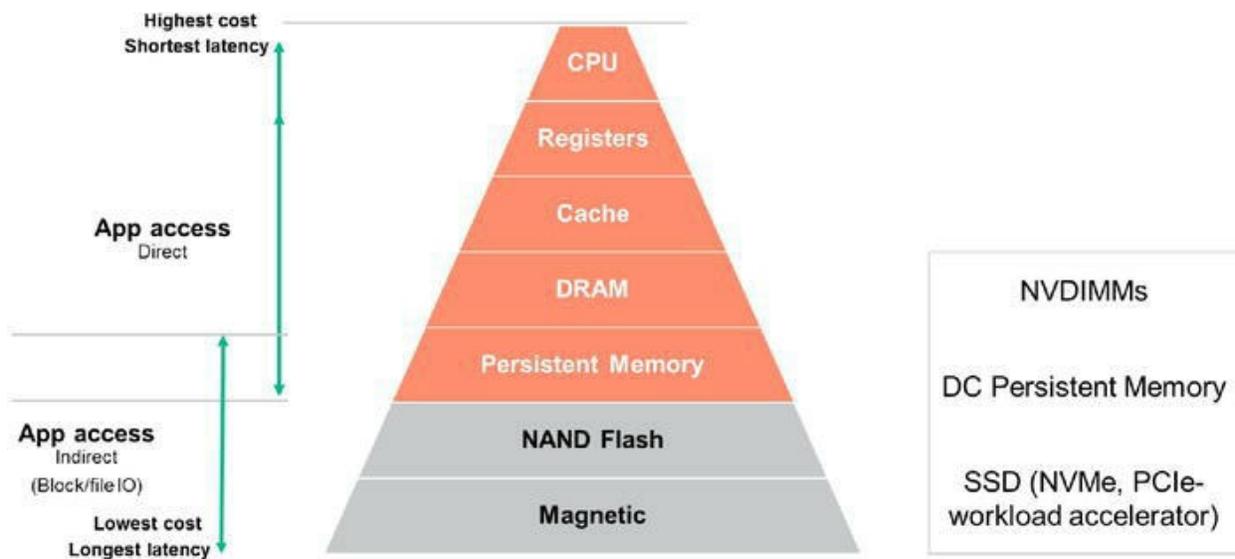


Figure 2-22 Memory/Storage hierarchy

Let us take a look at [Figure 2-22](#); technologies at the top of the pyramid have the shortest latency (best performance) but come at a higher cost relative to the items at the bottom of the pyramid. These layers are comprised of DRAM (memory), CPU cache(s), and CPU registers. All of these components are accessed directly by the application (CPU)—that is, byte-addressable access. These layers are also volatile in that their contents are lost when power is removed.

Technologies at the bottom of the pyramid—represented by magnetic media (HDDs and tape) and NAND flash (represented by SSDs and PCIe Workload Accelerators)—have longer latency and lower costs relative to the technologies at the top of the pyramid. Data stored on these technologies are non-volatile, even when power is removed. Applications access data on these layers indirectly typically using Block IO and/or File IO.

The new DC Persistent Memory technology layer sits between NAND flash and DRAM. It provides faster performance relative to NAND Flash while also providing the non-volatility not typically found in traditional memory. This technology layer provides the performance of memory with the persistence of traditional storage.

HPE SmartMemory

Gen10 DDR4-2666/2933 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 or 2933, depending on the processor.
- 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 Fast Fault Tolerance Memory

With HPE Gen10, HPE SmartMemory provides increased performance supporting speeds of up to 2666/2933 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. In the event of a memory failure, the memory subsystem automatically reorganizes the way data is stored in memory to create a protected region. It is just large enough to correct the DRAM failure while allowing the remaining portions of memory to continue to run at full performance.

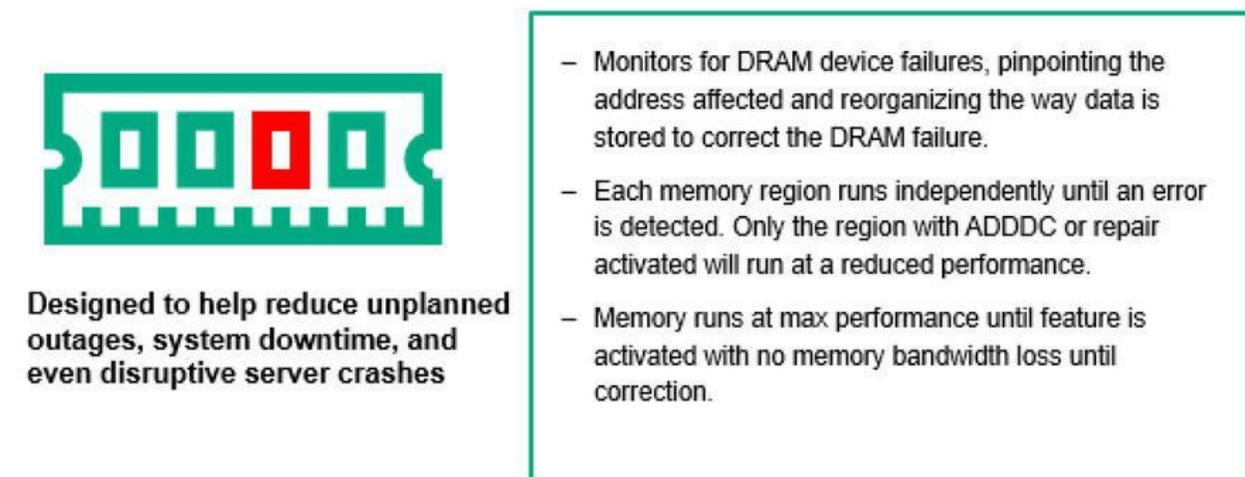


Figure 2-23 Seamless repair and recovery

Fast Fault Tolerance is a 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.

- **Features, as indicated in [Figure 2-23](#):**
 - 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

Device Data Correction (DDC) for random access memory patterns.

- **Improvement over Gen9 solutions**—The prior technology (Double Device Data Correction, or DDDC) is not used by most customers because of a severe performance penalty (all ranks must be configured so they are permanently in lockstep [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.
 - Only the region with Adaptive Double DRAM Device Correction (ADDDC) invoked will run at reduced performance.

Gen9 vs. Gen10 memory configuration (Intel[®]-based servers)

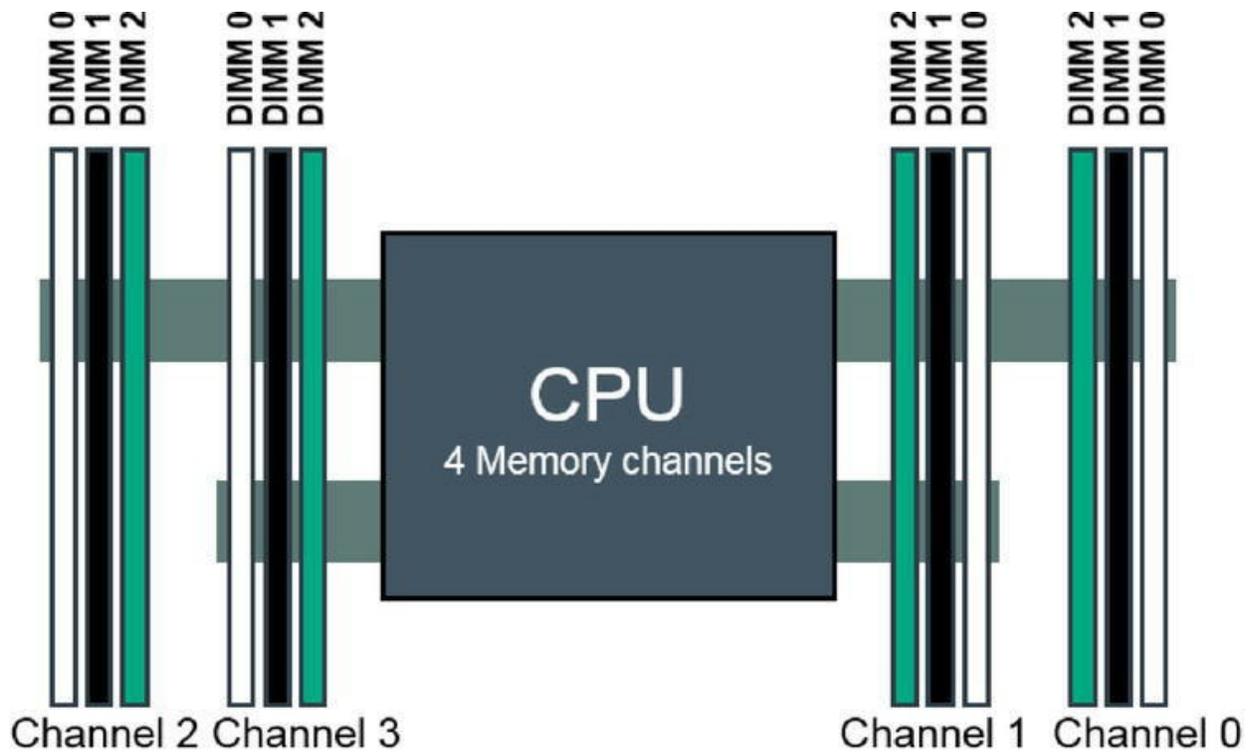


Figure 2-24 Gen9 with four memory channels per CPU

As illustrated in [Figure 2-24](#), Gen9 servers are using four memory channels per CPU (three slots per channel). This setup can be found on DL360/DL380/ML350/DL560, and so forth.

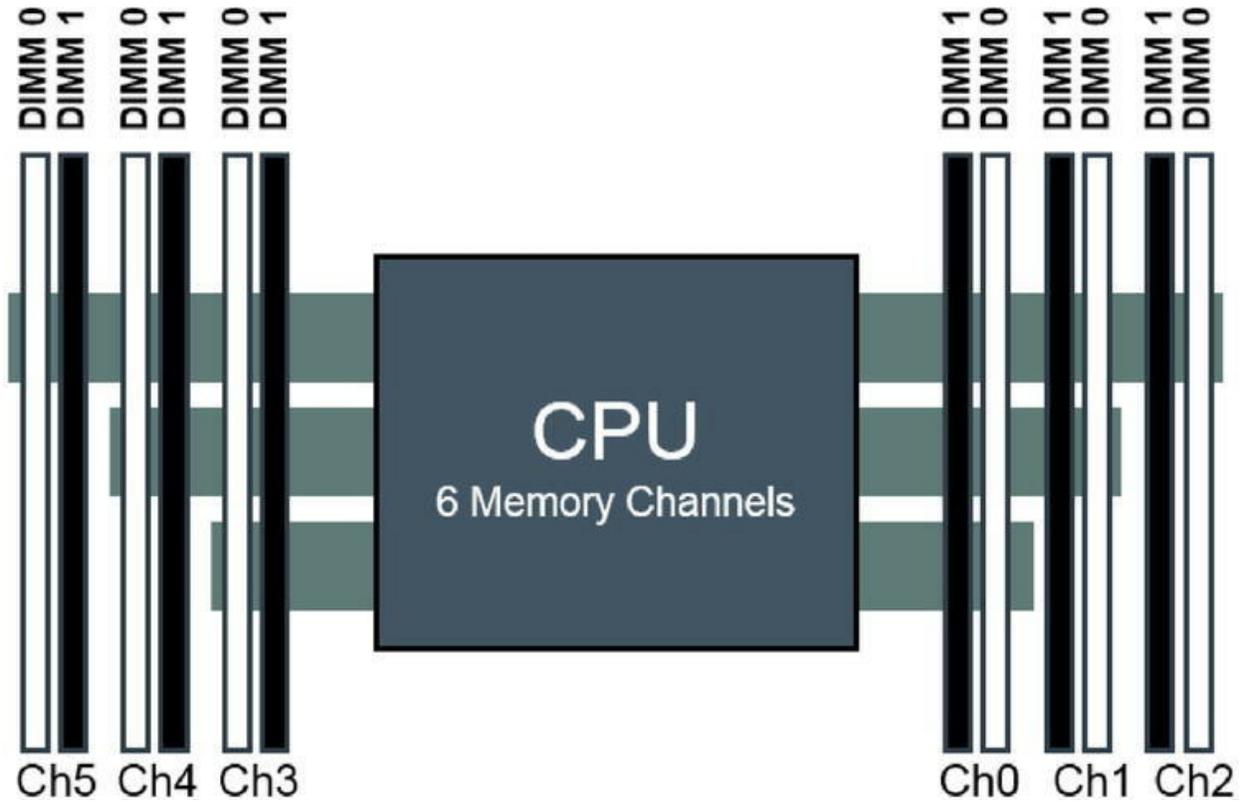


Figure 2-25 Gen10 with six memory channels per CPU

Gen10 servers with Intel[®] processors are equipped with six memory channels per CPU (two slots per channel), as shown in [Figure 2-25](#). Such architecture can be found on DL360/DL380/ML350/DL560, and so forth.

Gen10 memory configuration (AMD EPYC-based servers)

HPE ProLiant Gen10 servers—AMD 1 processor configuration																
DIMM population order																
Number of DIMM(s) to populate	Processor 1															
	CH D	CH C	CH B	CH A	CH E	CH F	CH G	CH H								
1	16															
2	16														1	
3	16			12											1	
4	16			12							5				1	
5	16		14	12							5				1	
6	16		14	12							5		3		1	
7	16		14	12		10					5		3		1	
8	16		14	12		10			7		5		3		1	
9	16	15	14	12		10			7		5		3		1	
10	16	15	14	12		10			7		5		3	2	1	
11	16	15	14	12	11	10			7		5		3	2	1	
12	16	15	14	12	11	10			7	6	5		3	2	1	
13	16	15	14	13	12	11	10		7	6	5		3	2	1	
14	16	15	14	13	12	11	10		7	6	5	4	3	2	1	
15	16	15	14	13	12	11	10	9	7	6	5	4	3	2	1	
16	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1

Figure 2-26 DIMM population order for HPE ProLiant DL385 Gen10 Servers with one processor installed (16 slots per processor)

The high-level memory system architecture for HPE ProLiant Gen10 servers using AMD EPYC processors is different from that of older AMD-based HPE ProLiant servers or Intel®-based servers. As indicated in [Figure 2-26](#), HPE ProLiant Gen10 servers using AMD EPYC processors include eight separate memory channels per CPU and up to 32 DIMM slots in two-socket servers—as opposed to the traditional 24 DIMM slot servers—allowing large memory configurations and delivering improved memory performance.

HPE ProLiant Gen10 systems support a variety of flexible server memory configurations, enabling the system to be configured and run in any valid memory controller configuration. For optimal performance and functionality, you should follow the rules when populating HPE ProLiant Gen10 servers with HPE DDR4 memory. Violating these rules may result in reduced memory capacity, performance, or error messages during boot.

DIMM population rules for HPE ProLiant Gen10 servers include the

following:

- Install DIMMs only if the corresponding processor is installed. If only one processor is installed in a two-processor system, only half of the DIMM slots are available.
- If a memory channel consists of more than one DIMM slot, the white DIMM slot will be located furthest from the CPU. White DIMM slots denote the first slot to be populated in a channel. For one DIMM per channel (DPC), populate white DIMM slots only.
- When mixing DIMMs of different ranks on the same channel, place the DIMMs with the heaviest electrical load (highest number of ranks) in the white DIMM slot. Within a given channel, populate DIMMs from the heaviest electrical load (dual rank) to the lightest load (single rank).
- If multiple CPUs are populated, split the DIMMs evenly across the CPUs and follow the corresponding CPU rule when populating DIMMs.
- To maximize performance in most of the potential applications, it is recommended to balance the total memory capacity across all installed processors, core group/channel pairs (P0A/B, P0C/D [or socket numbers 9 to 12, 13 to 16]), and channels. Populate memory local to each core group to reduce the latency. Load the channels similarly whenever possible to enable optimal interleaving. If the number of DIMMs does not spread evenly across the CPUs, populate as close to evenly as possible.
- The maximum memory capacity is a function of the number of DIMM connectors on the platform—the largest DIMM capacity qualified on the platform and the number and model of qualified processors installed on the platform.
- Do not mix RDIMMs and LRDIMMs in the same system. Do not mix 128 GB LRDIMMs with other capacity DIMMs. Do not mix x4 and x8 DRAM widths in the same system.
- Unbuffered DIMMs (UDIMMs) are not supported.
- The maximum memory speed is a function of the memory type, memory configuration, and processor model.
- DIMMs of different speeds may be mixed in any order; however, the server will select the highest common speed among all of the

DIMMs/CPU.

- HPE memory from previous generation servers is not compatible with the current generation. Certain HPE SmartMemory features such as memory authentication and enhanced performance may not be supported.



Note

To read more information about Server memory population rules for HPE ProLiant Gen10 servers with AMD processors, download the following PDF:

<https://h20195.www2.hpe.com/V2/GetPDF.aspx/a00038346enw.pdf>

HPE Persistent Memory Kit featuring Intel[®] Optane[™] DC Persistent Memory

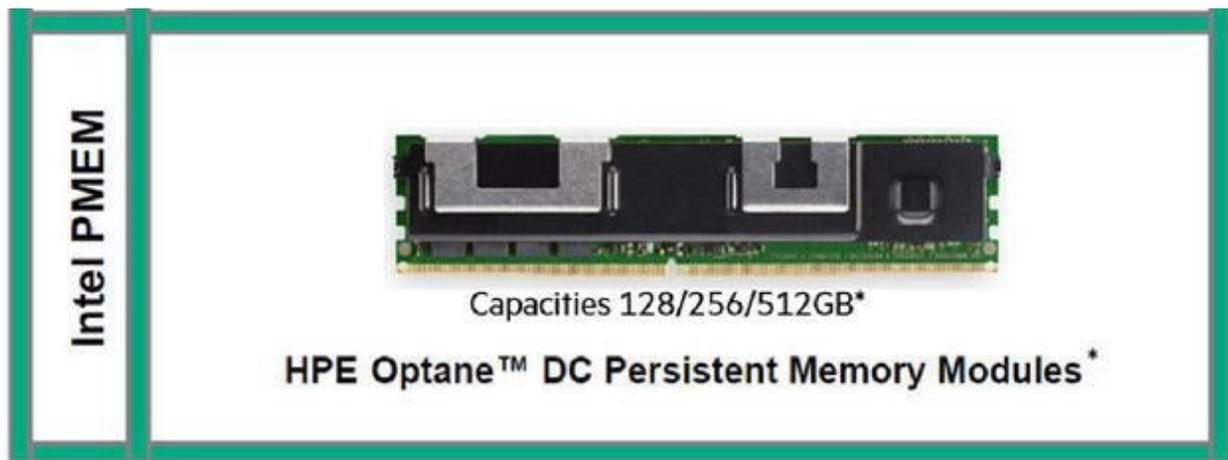


Figure 2-27 HPE Persistent Memory Kit featuring Intel[®] Optane[™] DC Persistent Memory

The new style of business requires the fastest access to customer data to put it to work more quickly for better business outcomes. Customers need solutions designed and built around the processes and applications that are at the core of the business. HPE Persistent Memory transforms critical data workloads—from cloud, databases, and high-performance computing to in-memory analytics and content-delivery networks. HPE Persistent Memory technology

is:

- Fast
- Reliable
- Designed around business workloads

As shown in [Figure 2-27](#), HPE Persistent Memory, available in 128, 256, and 512 GB kits, features Intel® Optane™ DC Persistent Memory to approach the speed of traditional DRAM with the persistence of storage, ensuring high capacity, high performance, and ongoing data safety—even in the event of a power interruption due to an unexpected power loss, system crash, or normal system shutdown.

HPE Persistent Memory provides large memory and fast storage and is available on select Intel® Xeon® Scalable Family Generation 2 only (select SKUs).

Software ecosystem enablement:

- **OS:** Windows, RHEL, SUSE, and VMware
- **Apps:** Databases and analytics

Activity: Server memory and Persistent Memory population rule

1. Open the technical white paper at:
<https://h20195.www2.hpe.com/v2/getmobile.aspx?docname=a00017079enw>
2. Answer the following questions:
 - a. Is the BL460c Gen10 with eight DIMMs considered a balanced configuration?

-
-
- b. If a channel has both an HPE SmartMemory DIMM and an HPE Persistent Memory module, what is the correct order of installation?

Learning check

4. What are the requirements for using HPE Persistent Memory featuring Intel® Optane™?

Storage for HPE servers

Storage is an important subsystem inside an HPE ProLiant server. HPE offers a wide portfolio of storage products for ProLiant servers, including:

- HPE Smart Array controllers
- HPE Hard Disk Drives (HDDs)

- HPE Solid-State Drives (SSDs)

Drive array basics

An array is a set of physical disk drives that can be combined into a single logical drive or subdivided into multiple logical drives that are distributed across all disks in the set.

Having several physical hard drives enables the controller to divide the data across multiple drives. A file is divided into a selected number of sectors, and then the file data is written concurrently across a series of drives in an array.

This approach of combining drives brings several advantages:

- **Performance**—Because multiple drives are accessed simultaneously, the process of writing (or reading) a file across multiple drives is much faster than writing to or reading from a single drive.
- **Redundancy**—If configured properly, one or more drives can fail without affecting the data accessibility on the array.
- **Capacity**—Combining multiple drives together allows you to create logical unit numbers (LUNs) that are bigger than individual drives.
- **Management**—An interface enables users to interact with the storage array to create Redundant Array of Independent Disk (RAID) sets and virtual volumes, to present virtual volumes to hosts, and to monitor the health and performance of the storage array.

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 is 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 is 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 functioning.

RAID level	Function/Application	Limitations
0 Minimum 1 disk	Data is distributed, striped, across separate disk drives.	Highly vulnerable to failure. The entire array will fail if one drive fails.
1 Minimum 2 disk	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 disk	Implemented as striped, mirrored disks. Good for database applications requiring high performance and fault tolerance.	50% of capacity dedicated to fault protection. Doubles the number of drives required.
5 Minimum 3 disk	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).

<p>6 Minimum 4 disk</p>	<p>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 24x7 operations that require higher levels of tolerance than RAID 5.</p>	<p>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.</p>
<p>50 Minimum 2 groups of 3 disks</p>	<p>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.</p>	<p>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</p>
<p>60 Minimum of 2 groups of 4 disks</p>	<p>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 requires at least 8 disks.</p>	<p>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 increase, this impacts the fault-recovery time correspondingly as the interval for rebuilding the RAID set increases.</p>
<p>JBOD</p>	<p>Just a Bunch Of Disks. Presented by a controller in host bus adapter (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 softwaredefined storage (SDS).</p>	<p>Performance is limited by the speed of the processor and the IO subsystem. Redundancy is handled by the software.</p>

Hardware-based RAID



Figure 2-28 HPE Smart Array Controller

Advantages of hardware-based RAID, as in [Figure 2-28](#), in contrast to software RAID include:

- Faster, more reliable performance
- Decreased processor load
- User-friendly configuration utilities
- No operating system interface for rebuild
- Safe write caching

HPE Smart Array controllers and storage

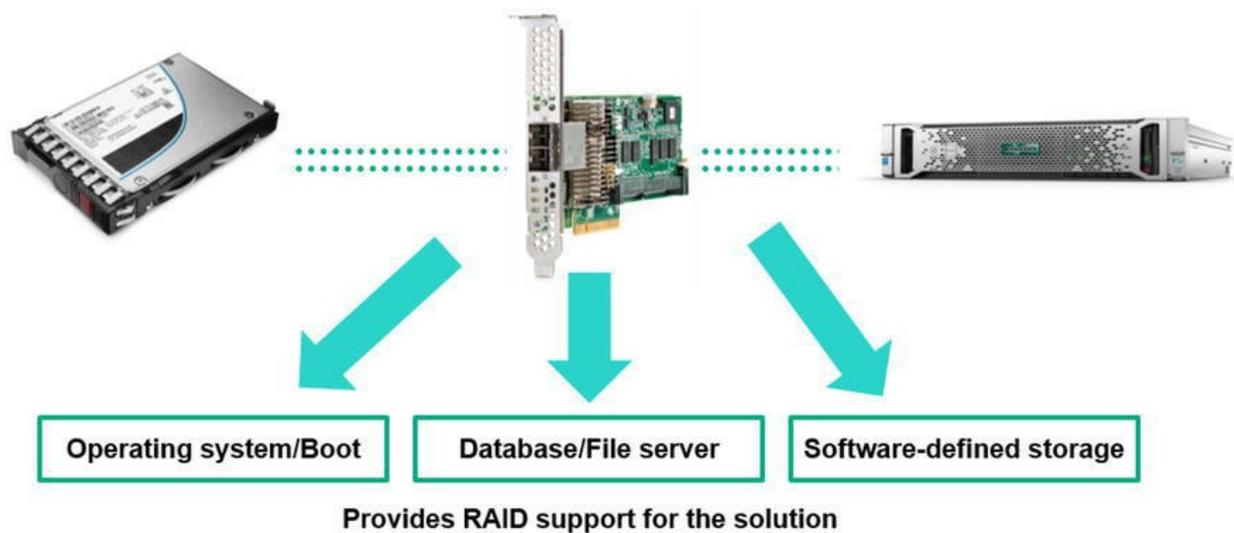


Figure 2-29 How a Smart Array controller works

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 indicated in [Figure 2-29](#). There are two types: Smart Host Bus Adapter (HBA) and Smart Array Controller.

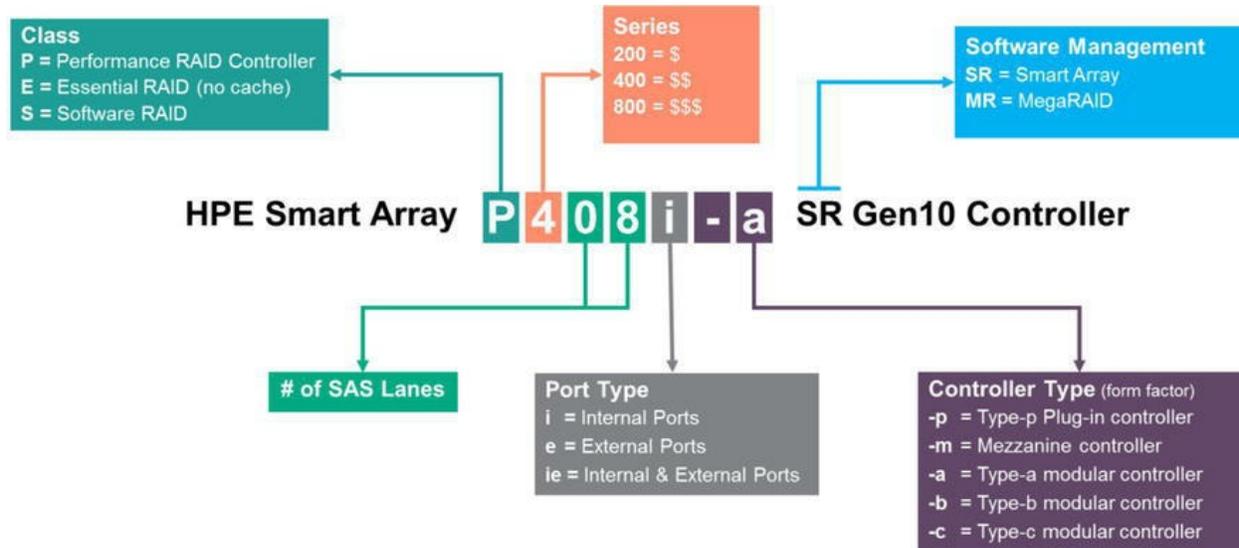


Figure 2-30 Gen10 Smart Array naming framework

An HBA is an expansion card that plugs into a slot (such as PCIe) 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 (RAID), 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.

HPE Smart Array controllers are ideal for maximizing performance, capacity, and resilience for direct-attached storage (DAS) connected to entry-level and enterprise servers. The HPE’s portfolio of Smart Array controllers includes several lines and series, as illustrated in [Figure 2-30](#).

External HPE Smart Array controllers use mini-SAS connectors for connecting to 6 Gb SAS external devices and mini SAS-HD connectors for 12 Gb SAS devices.

Features of HPE Smart Array controllers 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**—An UEFI Configuration Tool reduces time to configure simple RAID volumes.
- **Save time**—Newly created RAID 5 or RAID 6 volumes require parity initialization with Rapid Parity Initialization (RPI).
- **Enhanced protection**—Data at rest on all SAS/SATA drives and data security comply with regulations for sensitive data using HPE Smart Array SR Secure Encryption.
- **Better performance**—HPE Gen10 controllers deliver up to 1.6 M Input/Output Operations Per Second (IOPS) (4 KB random reads), 65% more performance over the previous generations of controllers.
- **Less power**—The 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 4x by caching the hot data on SSDs.

Example of tiered storage solutions: VMware vSAN

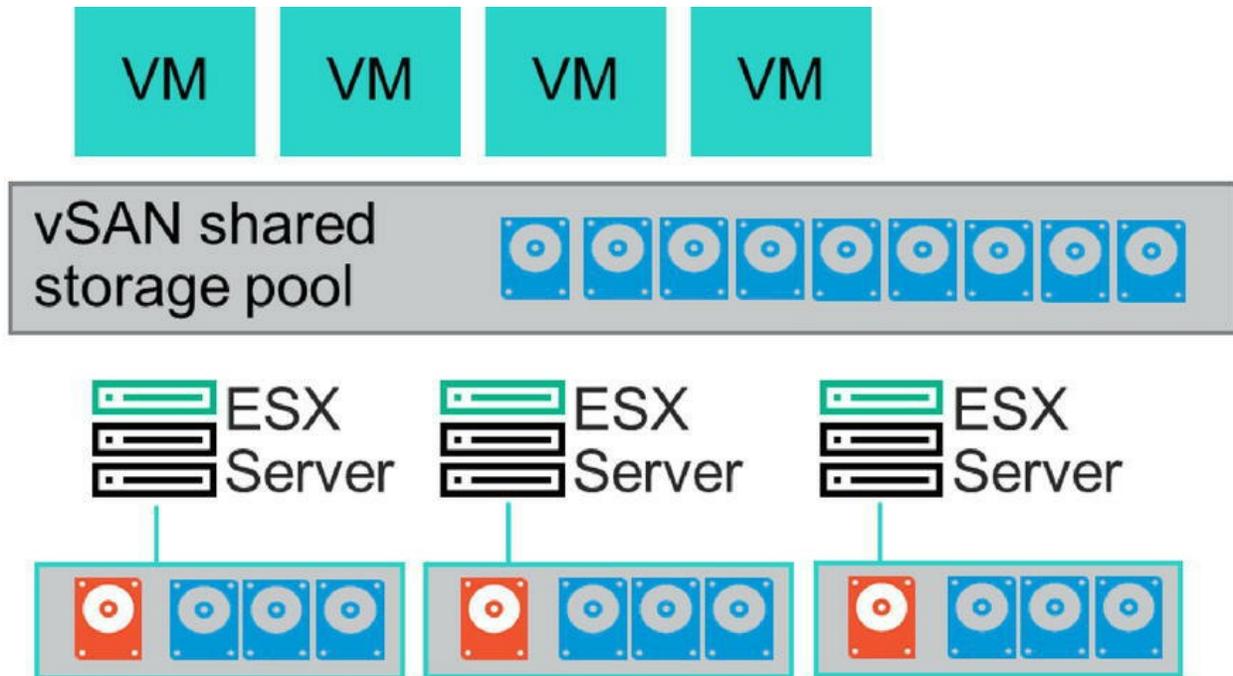


Figure 2-31 Architecture

VMware vSAN, as referenced in [Figure 2-31](#), is built into the hypervisor kernel, making it very simple and seamless with a zero-step installation process.

Product features:

- Deduplication, compression, and erasure coding on All-Flash systems
- Data at rest encryption (software based) with vSAN 6.5 or 6.6
- Smart Drive support (degraded device handling)
- One-click firmware and driver upgrades
- iSCSI target service for block storage
- Two-node cluster with a witness at a third site
- Min. 100 Mbps bandwidth, <300 ms RTT for witness

Architecture:

- Minimum of three VMware vSphere ESX servers, up to 64 nodes
- One to five disk groups per host:

- One SSD per disk group
 - One to seven HDDs per disk group
 - Software mirroring
 - Up to four copies of data can be stored in the vSAN pool
-



Note

vSAN is enterprise-class, storage virtualization software that, when combined with vSphere, allows you to manage compute and storage with a single platform. It seamlessly joins all storage devices across a vSphere cluster into a shared data pool, allowing you to easily scale up or down as your needs change. Using commodity x86 server components, vSAN-powered hyperconverged infrastructure lowers storage costs 40% or more versus traditional server and storage architectures (<https://www.vmware.com/products/vsan.html>, October 2019).

Activity: Features of HPE Smart Array Controller

1. This activity requires you to locate the QuickSpecs for the HPE Smart Array P408i-a SR Gen10 Controller:
<https://h20195.www2.hp.com/v2/GetDocument.aspx?docname=a00008200enw>

2. Answer the following questions:

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

- b. What is the maximum number of logical drives?

- c. What is the form factor of this controller?

d. What is required to support the Flash Back Write Cache?

e. Describe the feature known as “Mixed Mode” controller operation?

Learning check

5. Name at least three advantages of hardware-based RAID over software-based RAID.

Networking features of ProLiant servers

HPE server adapters, also known as network interface cards (NICs), are expansion cards that enable the server to be connected to the network. They enable you to increase server optimization for variable networking workloads. Gen10 server network adapters enable customers to:

- **Safeguard a customer’s data and business network**—Block installation of corrupted firmware and verify that the executing firmware

is trusted using silicon root of trust and authentication.

- **Boost IO bandwidth with lower latency**—Deliver single-root IO virtualization with the ability to run a large number of VMs per server, which reduces the need for hardware and the resulting costs of space and power required by hardware devices.
- **Reduce CPU utilization and help improve host VM density and server efficiency**—Use Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) for live migration and Microsoft Server Message Block (SMB) Direct environments.
- **Improve operating costs**—Off-load packet processing to lower power with tunnel off-loads. Speed packet processing to service more IO requests using data plane development packets.

Regarding network adapters, the right mix of performance and cost with reliability and security are essential for any given workload. HPE offers three series of NICs to choose from: Standard, Advanced, and Performance series.

Secure networking flexibility at the speed of compute

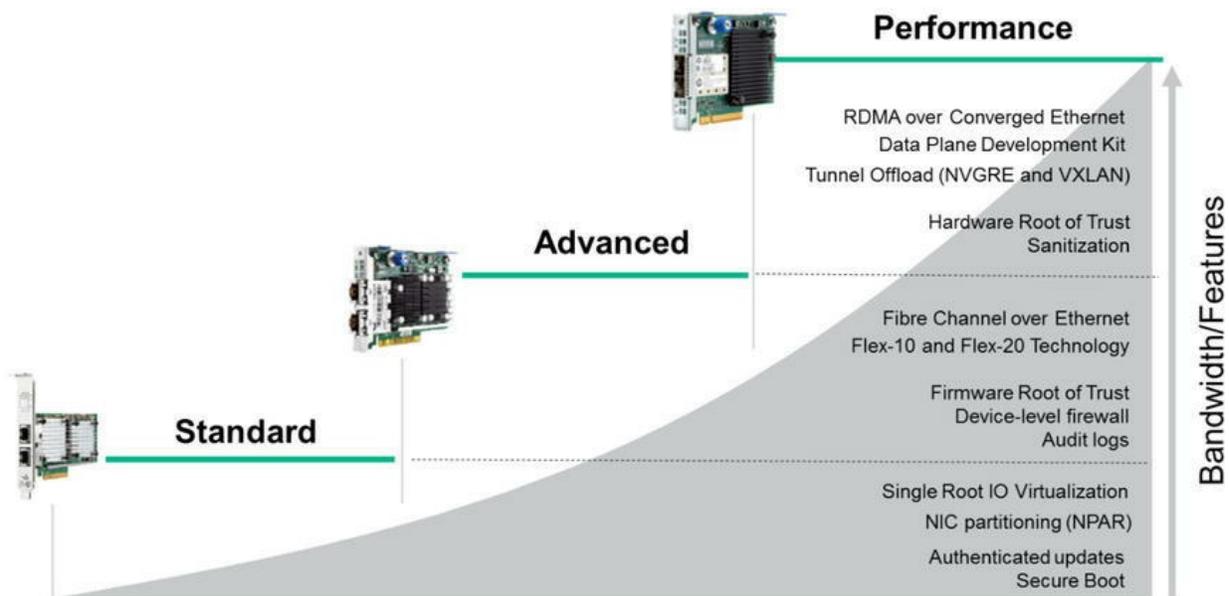


Figure 2-32 Gen10 Network Adapter product portfolio

As indicated in [Figure 2-32](#), the **Standard** series adapters enable a cost-

effective Ethernet solution for everyday applications. The economic scalability of these adapters feature 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, which can be seen as individual NICs (adapters) in the operating system. NPAR functionality is not available in all network adapters. The series offers two- and four-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 the following:

- 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 startup 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 simplify network and storage topology to build the new hybrid server infrastructure using the converged network adapter (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. One of the Advanced series features is Fibre Channel over Ethernet (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 Top of Rack (ToR) switch cable connection. Other features of the Advanced series include the following:

- Reduce cabling and maximize bandwidth using NPAR
- Increase performance via direct access to hardware from virtual

environment with Single-Root IO Virtualization

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

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%.

Other features of the Performance series include:

- **Multispeed**—Support for high speeds, such as 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**—Hardware root of trust to create a chain of trust for 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 HPE Network Adapter Options are based on Institute of Electrical and Electronics Engineers (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.

Activity: HPE QuickSpecs: Server components

1. Download the HPE ProLiant ML350 Gen10 Server QuickSpecs:

<https://h20195.www2.hp.com/v2/getdocument.aspx?docname=a00021852enw>

2. Answer the following questions:

- a. What are the typical speeds of Ethernet adapters available for ML350 Gen10?

- b. What is the name of the connector required for 100 Gb?

- c. Are there any other requirements for using HPE Ethernet 10/25 Gb 2-port 621SFP28 Adapter?

- d. What are the requirements for using additional network adapters in PCIe slots 5 to 8?

Learning check

6. Match the server networking component with its description.

CNA	System board module that provides network connectivity without occupying an expansion slot
FlexibleLOM	Expansion card that connects the server to the network and increases server optimization
NIC	Carries both Ethernet and storage traffic on a single adapter, thereby reducing costs

Rack and power infrastructure

HPE offers a variety of rack and power infrastructure products to provide a smarter, stronger, and simpler rack and power infrastructure.

Benefits include:

- They are compatible with proven industry-leading IT infrastructure products and software.
- HPE warranties and services extend to the rack and power product portfolio at no additional cost.
- Streamline operations with a single vendor for one-stop shopping.
- HPE Foundation Care Services support your customers to minimize the time to troubleshoot, monitor, and remediate issues, while innovating for the future.

HPE Rack and Power Infrastructure portfolio

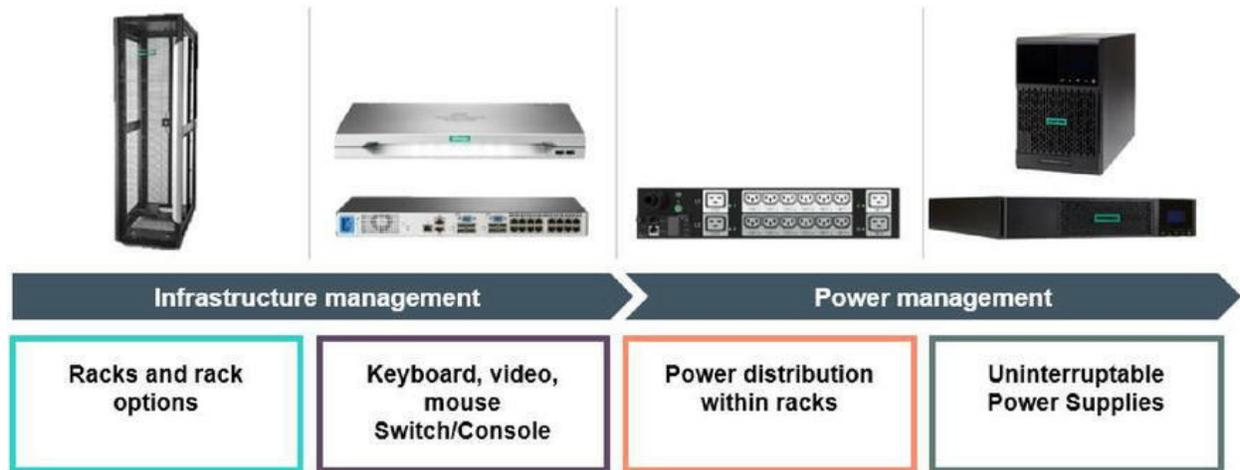


Figure 2-33 HPE Rack and Power Infrastructure is compatible, flexible, powerful, and trusted

As referenced in [Figure 2-33](#), HPE's rack and power products include:

- Racks and rack options
- Keyboard, Video, Mouse (KVM) Switch/Console
- Power Distribution within Racks (PDU)
- Uninterruptable Power Supplies (UPS) and Extended Runtime Modules (ERM)

G2 rack product portfolio

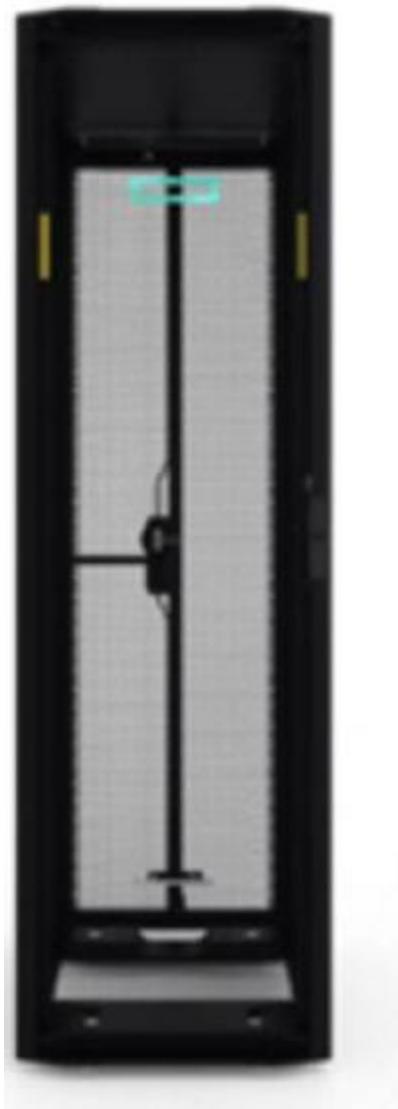


Figure 2-34 Flexible and robust infrastructure solutions

As illustrated in [Figure 2-34](#), G2 rack portfolio includes two lines: Advanced and Enterprise.

G2 Advanced Racks are suitable for any load—from a server room to a data center:

- 22U, 36U, 42U, and 48U
- Side panels and baying kit included
- Pallet and shock pallet models

- 10-year limited warranty

G2 Enterprise Racks are primarily targeted to enterprise data centers:

- 42U and 48U
- Highest load capacity
- Tall/Wide/Deep models
- Pallet and shock pallet models
- 10-year limited warranty

G2 Power Distribution Unit product portfolio

HPE PDUs, as shown in [Figure 2-35](#), provide power to multiple objects from a single source. In a rack, the PDU distributes power to the servers, storage units, and other peripherals.



Figure 2-35 G2 Power Distribution Unit product portfolio

PDU systems:

- Address issues of power distribution to components within the computer cabinet.
- Reduce the number of power cables coming into the cabinet.
- Provide a level of power protection through a series of circuit breakers.



Note

For more information on the HPE PDU portfolio, visit:
<https://www.hpe.com/us/en/product-catalog/servers/power-distribution-units.hits-12.html>

The G2 PDU portfolio includes four types of products:

- **G2 Standard**—Non-critical/SMB (G2 basic models)
 - High operating temperatures up to 60° C
 - IEC locking power cord support
 - Flexible mounting options, including true 0U configurations
 - Optional dual-locking power cords
 - No metering or switching
 - Up to 48 outlets
- **G2 Advanced**—Server room to data center (G2 switched and G2 metered models)
 - Metered and switched models
 - Load segment metering (metered models) and outlet switching (switched models)
 - High operating temperatures up to 60° C
 - IEC locking power cord support
 - Flexible mounting options including true 0U configurations
 - 1 GB Ethernet, up to 48 outlets
 - Single Phase only (Switched models)
- **G2 Enterprise**—Enterprise data centers (G2 metered and switched models, intelligent models)
 - Outlet metering and outlet switching
 - High operating temperatures up to 60° C
 - IEC locking power cord support
 - Flexible mounting options including true 0U configurations

- 1 GB Ethernet, up to 48 outlets
- **iPDU**—iPDU is a PDU with full remote outlet control, outlet-by-outlet power tracking, and automated documentation of power configuration.

Elevated temperature support

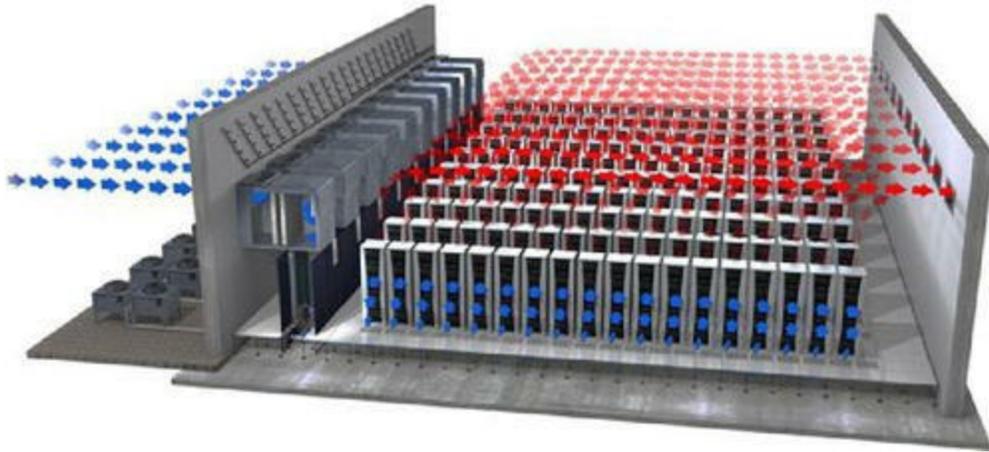


Figure 2-36 Elevated temperature support

The ASHRAE guidelines for data center operation suggest thermal ranges of operation for IT equipment that are commonly accepted by many data center administrators and facility managers. ProLiant Gen9 and Gen10 products support A3 and A4 guidelines, which means that ProLiant Gen9 and Gen10 servers can operate at higher temperatures than previous generations. This provides more choices for air-cooled solutions and the potential to reduce cooling costs.

The allowable ranges for classes A3 and A4 are intended to remove obstacles to data center cooling strategies such as free-cooling methods. Free-cooling takes advantage of a facility's local climate by using outside air to cool IT equipment directly, as [Figure 2-36](#) shows—without the use of mechanical refrigeration (chillers or air conditioners) whenever possible.

A variety of implementations for free-cooling are possible. For example, filtered outside air can be drawn directly into the data center. Other techniques keep the outside air isolated from the data center, but still transfer the data center heat directly to the outside air without refrigerating it. Careful

application of the new ASHRAE guidelines might enable free-cooling in more climates or allow for the data center to be cooled without refrigeration more days of the year. Reducing the use of refrigerated cooling lowers the operating expenses for the data center. In some cases, the refrigeration equipment can be eliminated or significantly reduced in size, saving capital investment expense as well.

The traditional maximum server inlet air temperature is 35°C. For ProLiant Gen9 and Gen10 servers, 40°C (ASHRAE A3) support is available on most platforms with configuration limitations. For select ProLiant Gen9 and Gen10 platforms, 45°C (ASHRAE A4) support is available with configuration limitations.



Note

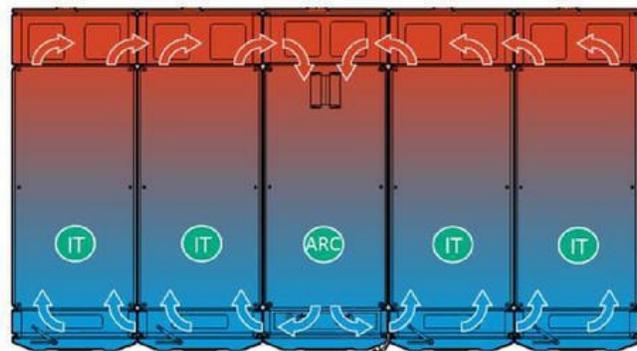
Check the QuickSpecs for specific servers to find ASHRAE support information.



Note

The American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE/^æʃreɪ/ ASH-ray) is an American professional association seeking to advance heating, ventilation, air conditioning and refrigeration (HVAC&R) systems design and construction ([Wikipedia.org](https://www.wikipedia.org), October 2019).

HPE Adaptive Rack Cooling System



Top down view

Figure 2-37 HPE Adaptive Rack Cooling System

The HPE Adaptive Rack Cooling System, as referenced in [Figure 2-37](#), is a closed-loop cooling system that removes the high levels of heat generated by today's servers, mass storage, and core networking hardware. The HPE Adaptive Rack Cooling System allows for simultaneous cooling of up to four racks and 150 kW of IT capacity when used in conjunction with the HPE Adaptive Rack Cooling System 42U or 48U racks.

HPE Adaptive Rack Cooling System is designed to complement any new or existing data center by enabling computing power growth without adding to the heat load in the data center. In addition, by allowing equivalent cooling capacity as a commodity direct liquid cooling (DLC), the HPE Adaptive Rack Cooling System will extend, considerably, the life of the data center. The horizontal air flow of the HPE Adaptive Rack Cooling System fully supports industry-standard front-to-back cooling designs and standard server dimensions. Additionally, the implementation of variable speed fans within the system enables improved energy efficiencies while providing all devices the right volume of airflow, regardless of the mounting position or workload.

Installing solutions and subsystems to the rack

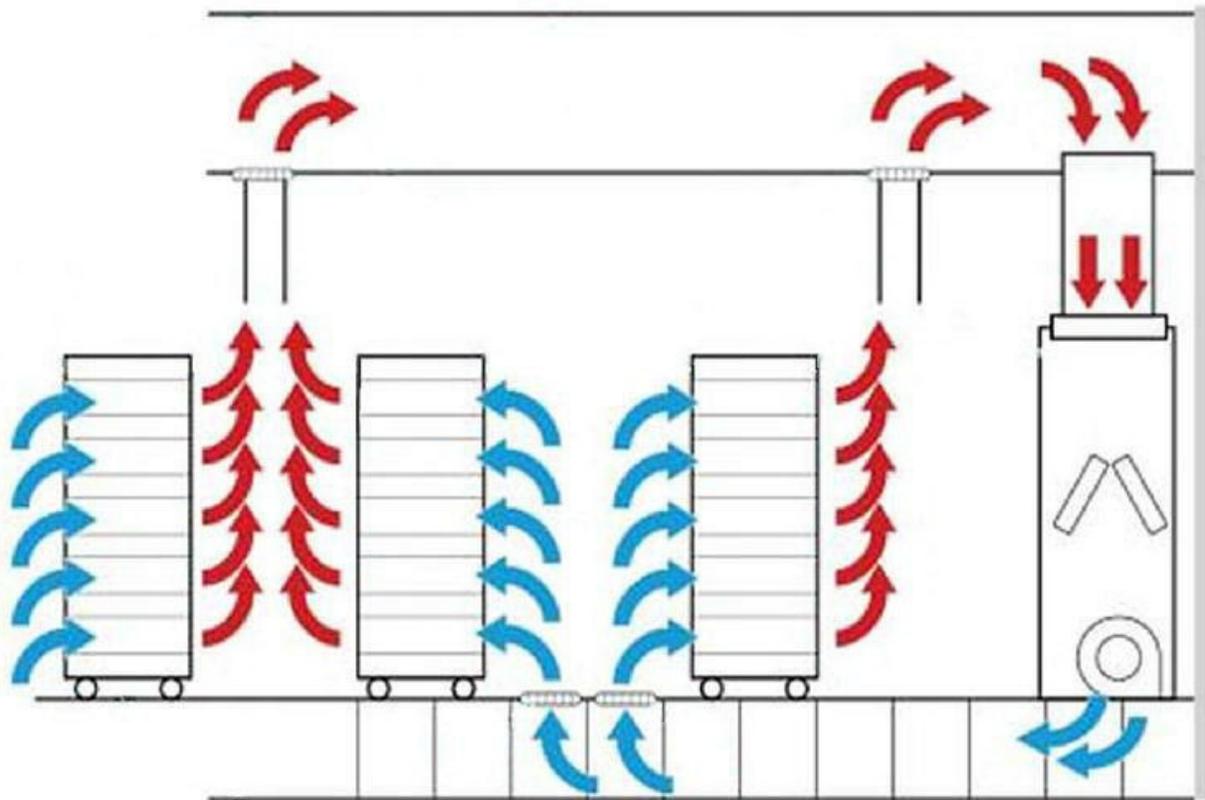


Figure 2-38 Rack placement for proper airflow in a data center with a suspended ceiling and raised floor

Before installing equipment in a data center, you should be aware of precautions and other requirements to ensure safety and proper handling of the equipment, as shown in [Figure 2-38](#). To prevent damaging the system, follow the precautionary guidelines that ship with the equipment when setting up the system or handling parts.

When installing a component, select a location that meets environmental requirements, including:

- A stable place where the component will not be bumped, scratched, or disturbed
- An area in which the component cannot easily be disconnected from its power source
- An area that is ideally locked or, at minimum, not accessible to

unauthorized personnel

- Within patching distance, directly or via cable management cross patches, of the location of the WAN connection and the switch that supplies the office/room floor network ports
- Low-to-medium relative humidity to protect against poor insulation, electricity creepage, mechanical property change of materials, and metal corrosion
- Room temperature not to exceed 35°C (95°F)

When installing an HPE network switch, it is important to keep the equipment clean. Dust buildup on the chassis may result in electrostatic adsorption, which causes poor contact of metal components and contact points, especially when indoor relative humidity is low. In the worst case, electrostatic adsorption can cause communication failure.

A discharge of static electricity from a finger or other conductor can damage system boards or other static-sensitive devices. This type of damage might reduce the life expectancy of the system or component.

All electrostatic interference from outside or inside the system adversely affect components in a conduction pattern of capacitance coupling, inductance coupling, electromagnetic wave radiation, or common impedance (including the grounding system) coupling. To prevent against electrostatic interference and electromagnetic interference, take the following actions:

- If AC power is used, use a single-phase, three-wire power receptacle with protective earth (PE) or ground to filter interference from the power grid.
- Keep switches far away from radio transmitting stations, radar stations, and high-frequency devices.
- Use electromagnetic shielding, such as shielded interface cables, when necessary.
- Route interface cables only indoors to prevent signal ports from getting damaged by overvoltage or overcurrent caused by lightning strikes.

The power input end of the switch has a noise filter, whose central ground is directly connected to the chassis to form the chassis ground (commonly known as PGND). You must securely connect this chassis ground to the earth

so the faradism and leakage electricity can be safely released to the earth to minimize electromagnetic interference (EMI) susceptibility of the switch. You can ground a switch by using a grounding strip at the installation site or the AC power cord connected to the switch.



Warning

Correctly connecting the switch grounding cable is crucial to lightning protection and EMI protection. Connect the grounding cable to the grounding system in the equipment room. Do not connect it to a fire main or a lightning rod.

Installing data center components

To view and navigate the initial server setup screens, you must first connect IO devices to the server. To do so, connect a keyboard and mouse to the USB ports, and then connect a monitor to the Video Graphics Array (VGA) or display port. Next, connect the network cable by connecting one end of the network cable to the NIC port on the back (or front, depending on the model) of the server. Then connect the other end of the network cable to a network jack or a network device such as router or LAN switch. The last step is to connect the power cord to the server power jack and to the AC power source.

As a best practice, HPE recommends that your customer install the latest firmware, drivers, and system software before using the server for the first time. To get the latest firmware and software updates, visit the download page on the HPE website for that server.

Heavy components, such as UPS or ERM, need to be installed at the bottom of the rack to prevent tipping of the rack.

Installing server options

The decision to select a new HPE server, storage, or networking system often can depend on the ability to configure or upgrade that server based on workload or performance requirements.

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 upgraded to meet business and technical requirements.

Supported hardware options such as drives and expansion cards depend on the server model in which the option is being installed. Be sure to visit the HPE website for product QuickSpecs or the installation guide for that option. If your customer is installing more than one option, read the installation instructions for all the hardware options and identify similar steps to streamline the installation process.

When installing an internal storage drive in a server, keep these considerations in mind:

- Non-hot-plug drives are not designed to be installed or removed from a server while the system is still powered on.
- Populate drive bays based on the drive numbering sequence. Start from the drive bay with the lowest device number.
- All drives grouped into the same drive array must meet the following criteria:
 - They must be either all HDDs or all SSDs.
 - Drives should be the same capacity to provide the greatest storage space efficiency when drives are grouped together into the same drive array.
- The system automatically sets all device numbers.

Before powering down the server for any upgrade or maintenance procedures, perform a backup of critical server data and programs.



Note

For complete information on servicing ProLiant servers, visit the Hewlett Packard Enterprise Support Center and search for the server in which your customer is interested. On the server page, click **Manuals** → **Service and maintenance information**. Then click the link to open the maintenance and service guide. You can access the Support Center at: <http://h20565.www2.hpe.com/portal/site/hpsc>.

Before installing and configuring IT solution components at a customer location, it is important to make sure that the customer understands the physical requirements for the installation and has planned for its location in the data center. Make sure the customer has a good understanding of:

- The racking requirements where the solution will be installed. Is the customer installing HPE racks or third-party racks? Ensure that the customer also understands the floor footprint requirements.



Note

Rack installations should be planned for operational efficiency as well as ease and safety of maintenance. HPE offers cable management solutions that can relieve the cord/cable congestion often created when populating a rack with servers, PDUs, and UPSs.

-
- The power requirements based on where the solution will be installed. Will it be a single-phase or three-phase power scenario? Confirm that the proper PDUs are configured. Make sure that the customer is aware of the circuit and power draw requirements per rack or floor tile along with data center cooling requirements.
 - The host port connectivity requirements. Clarify what the customer needs to provide or confirm that the proper cables are included in the configuration.

HPE offers a variety of options for racks to maximize the effectiveness of the racks:

- The HPE Baying Rack Option Kit enables you to connect multiple racks together.
- The Airflow Optimization Kit seals air gaps inside the rack, in between bayed racks, and the clearance from the floor to the rack to prevent hot and cold air from mixing, thus optimizing the temperature at which the chilled air enters the racked servers.
- Sidewall Panel Kits include lightweight, locking side panels in two sections for better handling.

- Anti-tip stabilizer kits provide stability and support when equipment is installed, removed, or accessed within the rack. Heavy Duty Stabilizer Kits should be used when a single-racked component exceeds 200 lbs.
- The Ballast Option Kit is designed to provide additional stability and support when used with the Stabilizer Kit. It is typically used when very heavy equipment is routinely installed, removed, or accessed within the rack.
- The Rack Tie Down Kit is used to secure the rack to the data center floor for seismic anchoring.
- HPE recommends the use of blanking panels to ensure proper airflow. The HPE Universal 1U Filler Panels are made from high-grade molded plastic and are designed to be installed without tools.
- HPE fixed and sliding rack shelves can be mounted at any height to place a monitor or another component within the rack.
- Adjustable rails allow HPE servers, storage, UPSs, and other products to be mounted in HPE or third-party racks.
- The Rack Light Kit is mounted under the roof of the rack and is used to illuminate the rear section of the rack. The light is turned on and off with the opening and closing of the rear doors.
- The locking drawer mounts on HPE G2 Advanced Series racks with sliding rails and is used to store optical discs, instruction manuals, tools, and any other items frequently used in the rack.
- The Rack Fan Kit mounted extractor fan kit enhances natural convection cooling by increasing airflow in the rack. The Rack Roof Mount Fan Kit is not compatible with 14U racks.
- Properly grounded racks help maintain system performance, protect network equipment, and maximize uptime. The HPE G2 Rack Grounding Kit is available for HPE G2 Advanced Series racks.
- The HPE Rack Branding Kit includes a blank that allows your customer to customize the rack with rack location information, branding, or other description information.

Power requirements



Figure 2-39 Modular 0U PDUs maximize rack space

When installing more than one server, you might have to use additional PDUs, as shown in [Figure 2-39](#), to safely provide power to all devices. Observe the following guidelines:

- Balance the server power load between available AC supply branch

circuits.

- Do not allow the overall system AC 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.

Install vertical PDUs on either side of the rack with outlets facing back for easy access and improved clearance. For higher power density configurations, install them side by side on both sides of the rack with the outlets facing the center. Horizontal/modular PDUs offer a modular architecture designed for data center customers who want to maximize power distribution and space efficiencies.

The control unit or core can be mounted in any 1U space or in 0U space on the side of HPE racks between the rails. Extension bars mount directly to the frame so your customers can configure access to power needed.

Utility AC power standards for connector types and voltage levels vary from country to country. HPE provides a variety of power cords, jumper cables, and accessories that facilitate using HPE servers and options around the world.

Most ProLiant servers feature auto-sensing power supplies with a power input range of 100 to 240 VAC (alternating current voltage), so they can be used globally. The appropriate power cables ship standard with all AC-powered products. HPE power cords meet the standards for the country for which they are intended. Rack-mountable HPE servers usually ship with IEC-to-IEC jumper cables. Some installations may require optional cable assemblies and cable retention devices.



Important

High-density systems can place significant current loads on PDUs and UPSs. HPE recommends using the HPE Power Advisor and the Rack and Power Infrastructure Architect Tool for calculating the total current (amperage) and apparent power (volt-amps) that a PDU or UPS will need to handle.

Cabling requirements

Cable options for a ProLiant ML/DL server include a mini-SAS cable, which is required depending on the combination of the server and storage controller. Some server models ship with the mini-SAS cables required to connect drives to the embedded SATA controller or HBA. You should always verify that a cable/transceiver is supported with a connecting device such as a switch or NIC/iSCSI HBA.

HPE recommends redundant power supply cables for high availability.

Some HPE rack components are accessed by sliding the chassis out on rails. This means the power cords or jumper cables connecting to the servers must have adequate length and slack to allow chassis movement while staying connected and powered up. High-density server systems employing 1U enclosures or server blades allow your customers to replace critical, hot-pluggable components without removing the chassis from the rack (and in some cases, without even extending the chassis from the rack). Power cords or jumper cables for these systems can therefore be short because cable movement is of little or no consideration.

HPE offers a cable management arm (CMA) that is used when installing a rack-mounted server to secure the server and network cables. The fully adjustable arm is available in 1U and 2U form factors and keeps the cables in the back of the rack in order. It also enables your customers to extend devices from the rack without powering them down or disconnecting cables.

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 most servers is 35°C (95°F). The temperature in the room where the rack is located must not exceed 35°C (95°F).
- **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

7. What is the maximum inlet ambient temperature for systems compliant with ASHRAE Class A4?
 - A. 30°C
 - B. 35°C
 - C. 40°C
 - D. 45°C

Alternative products

We presented solution components for the customer scenario. Some of the solution components, however, can be changed for alternative environments where different performance, capacity, costs, features, and similar

requirements can be satisfied with other HPE solutions. Not all the workloads are the same, and a single HPE product cannot match all the requirements.



Note

Decision trees, attributes, and requirements that can be used to select appropriate server are covered at the beginning of this chapter; please refer to the **Selecting an HPE server platform** section if needed.

HPE ProLiant ML350 Gen10 Server

<p>Performance with unmatched capacity and reliability—all in a much smaller form factor</p> 	<p>Availability, expandability, and serviceability—a winning combination</p> 	<p>Agile infrastructure management for accelerating IT service delivery</p> 
---	---	--

The ML350 Gen10 Server is a 2P server providing high performance, expandability, and manageability—ideal for expanding SMBs and enterprise businesses.



Figure 2-40 Efficient, high-performance server delivers expandability, manageability, and reliability

As illustrated in [Figure 2-40](#), the ProLiant ML350 Gen10 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 SFF drives (8 x 3 cages) or four to 12 LFF drives (4 x 3 cages) and increased IO expansion. It features one or two Intel[®] Xeon[®] Skylake/Cascade Lake Series[®], four to 28 cores with up to

70% performance gain and HPE DDR4 SmartMemory offering up to 14% performance increase.

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, database)**—Improve productivity of your customer’s organization. Best match for performance and best growth possibilities for growing medium businesses.
- **Data warehouse/analytics**—Your customers can find the information they need, when they 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/Cascade Lake 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 NVMe PCIe SSDs
- Choices of SFF x24 or LFF x12 HDD/SSD

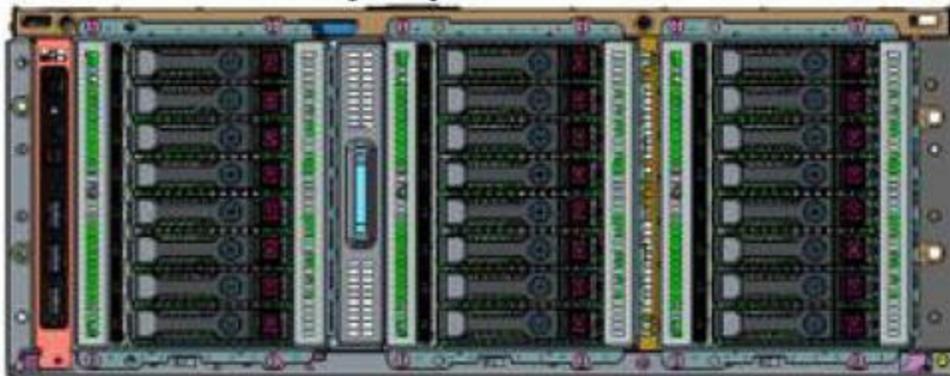
- Graphic support including NVIDIA Quadro RTX6000 or NVIDIA Tesla T4
- S100i Storage Controller and a choice of HPE Flexible Smart Array (AROC) or Smart HBA Controllers for performance or additional features
- Up to 24 DIMM slots of HPE DDR4 SmartMemory delivering 2666 Mt/s (3 TB Max) or 2933 MT/s (4.5 TB)
- 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
- 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 RPS 500 W/800 W/1600 W new Gen10 RPS

Storage options

Support up to **24SFF / 12LFF**



24 (8*3) SFF Model



16 (8*2) SFF HDD/SSD + 8 PCIe NVMe SSD



Figure 2-41 HPE ProLiant ML350 Gen10 storage options

As [Figure 2-41](#) shows, HPE ProLiant ML350 Gen10 Server can be used in several storage configurations, including:

- Four LFF (3.5") NHP (non-hot-plug)

- Eight SFF (2.5”)
- Four LFF
- Optional SFF or LFF cage kits
- Optional NVMe Express Bay

HPE ProLiant MicroServer Gen10—The compact server

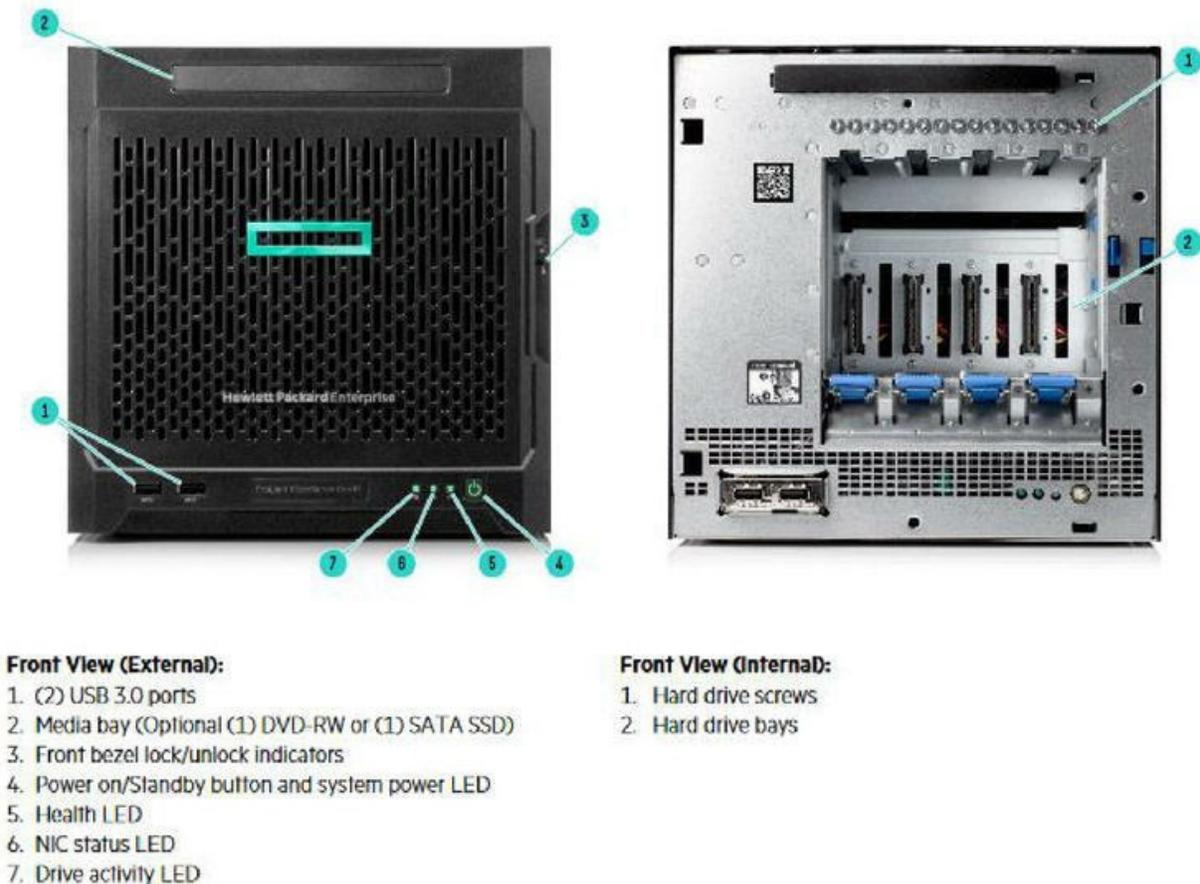


Figure 2-42 ProLiant MicroServer Gen10 (front view)

The HPE ProLiant MicroServer, as seen in [Figure 2-42](#), 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 SMBs, small office, home office, and enabled for graphics-related workloads.

ProLiant MicroServer Gen10 features include:

- Unique and serviceable design
 - Compact and stylish designed for small space
 - Simple installation and upgrades with easy access to drives, memory, and PCIe slots
- Graphic performance
 - Up to 48% better performance with Passmark 2008 measurement
 - Dual display ports capable of 4 K display resolution—4x 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 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



Rear View:

1. (2) USB 2.0 ports
2. (2) USB 3.0 ports
3. NIC 1 port
4. NIC 2 port
5. Power jack
6. PCIe Gen3 slots
7. VGA port
8. (2) Display Ports 1.2

Internal View:

1. DIMM slots
2. TPM connector
3. LFF/SFF SATA connector
4. SATA connector
5. System board power connector
6. System battery
7. Internal USB 2.0 port
8. Expansion slot 2, PCIe3 x4 (1)
9. Expansion slot 1, PCIe3 x8 (8, 4, 1)

Figure 2-43 ProLiant MicroServer Gen10 (rear and internal views)

As indicated in [Figure 2-43](#), ProLiant MicroServer Gen10 options include:

- AMD Opteron™ X3216 or X3421 or X3418 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 many HPE ProLiant servers, including the HPE ProLiant MicroServer Gen10. Servers with preloaded ClearOS are ready to be used 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.



Note

More information about ClearOS can be located at:
<http://www.hpe.com/servers/clearos>

HPE ProLiant ML30 Gen10 Server



Figure 2-44 ProLiant ML30 Gen10 Server

Referenced in 2-45, the HPE ProLiant ML30 Gen10 Server is a powerful yet affordable tower server designed for small offices, remote, and branch offices to run on-premises and hybrid cloud solutions, delivering enterprise-class performance, security, reliability, and expandability at a lower cost and standard 3-1-1 warranty.

Powerful yet affordable server for daily business workloads

Flexibility, reliability, and manageability for growing business needs

The industry leading security tower server at lowest TCO

Figure 2-45 Introducing the HPE ProLiant ML30 Gen10 Server

Ideal use cases:

- Small businesses, remote, and branch offices environment
- Traditional workloads such as IT, web, file, print, and mail servers
- Business applications or an original equipment manufacturer (OEM)-integrated solution
- Small to medium virtualization or data storage

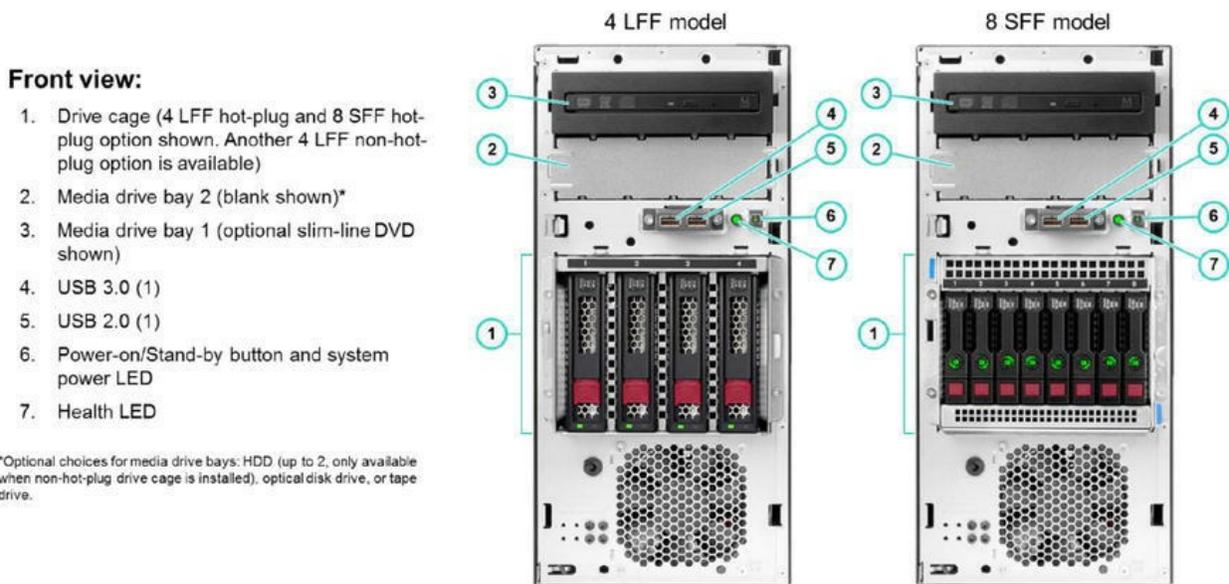


Figure 2-46 HPE ProLiant ML30 Gen10 Server front view

Rear view:

- 1. PCI slots (slots 1 to 4)
- 2. Dedicated iLO management port (optional)
- 3. USB 3.0 (4)
- 4. VGA port
- 5. Serial port (optional)
- 6. Power supply bay (one optional redundant power supply shown with enablement kit and blank shown)
- 7. Power supply power connector
- 8. Power cord clip
- 9. Padlock eye
- 10. Kensington security slot
- 11. NIC1/shared iLO port (upper), NIC2 port (lower)

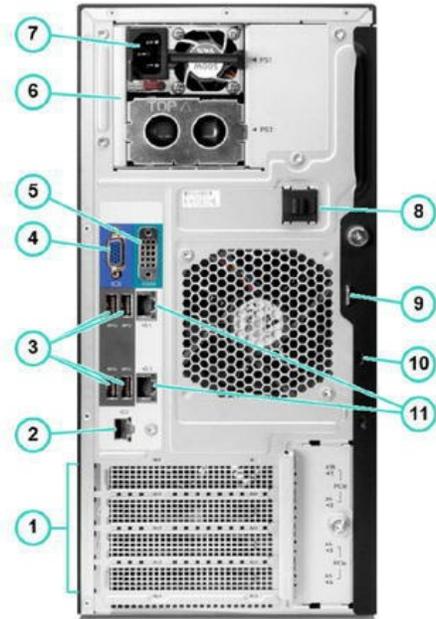


Figure 2-47 HPE ProLiant ML30 Gen10 Server rear view

Interior view:

- 1. System fan
- 2. Power supply
- 3. One processor and heatsink/fan
- 4. Four DDR4 DIMM slots
- 5. Bezel lock
- 6. Front PCI fan (optional or included, depending on models)
- 7. One internal USB 3.0 port
- 8. Six SATA ports
- 9. One M.2 slot (optional module shown)
- 10. TPM connector
- 11. Four PCIe expansion slots

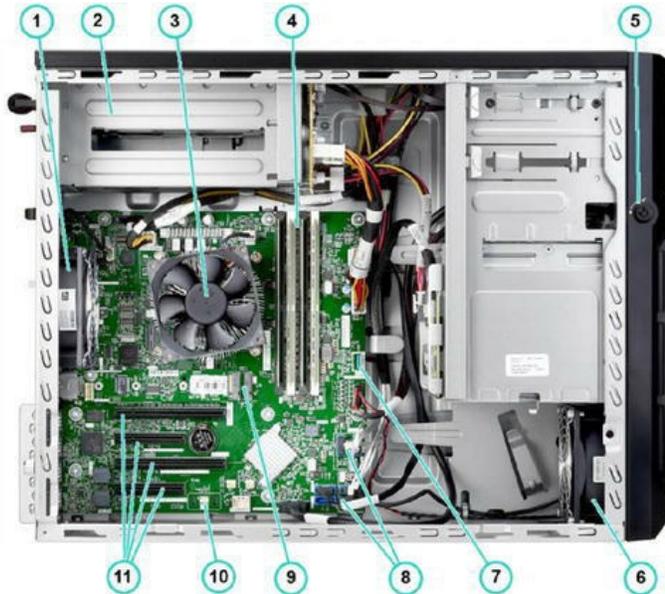


Figure 2-48 HPE ProLiant ML30 Gen10 Server interior view

Selecting an HPE DL server platform

Depending on workload, performance, memory and storage requirements, and requirements for expansion, the following decision trees, shown in 2-50,

can be used to select an appropriate Gen 10 DL server:

Performance, memory, and storage requirements



Figure 2-49 Selecting an HPE DL server platform, based on performance, memory, and storage requirements



Note

Other limitations and prerequisites apply. Consult QuickSpecs.

Size and expansion

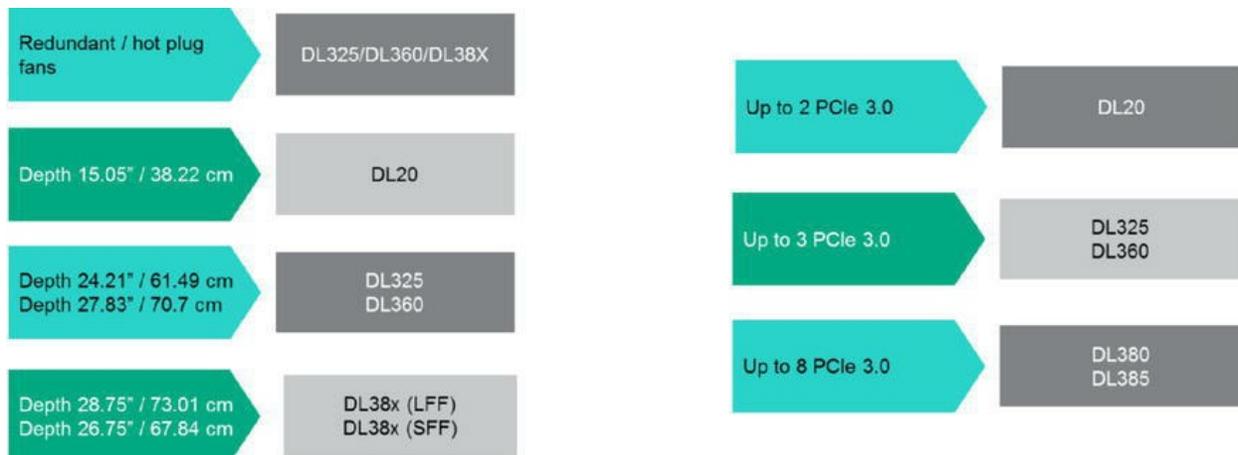


Figure 2-50 Selecting an HPE DL server platform, based on size and expansion requirements



Note

Other limitations and prerequisites apply. Consult QuickSpecs. Some features require additional components. Four-socket models are not covered.

HPE ProLiant Gen10 rack servers



Figure 2-51 HPE ProLiant Gen10 rack servers are based on industry-standard x86 architectures

As illustrated in [Figure 2-51](#), 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.

The ProLiant DL Gen10 family of optimized 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 SMB customers who want a smaller 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. The ProLiant DL Gen10 family includes a high-speed PCIe V3.0 interface and is supported by the latest in remote management, such as iLO 5.

The ProLiant DL server families offer rack-mounted servers based on Intel® or AMD EPYC processors. ProLiant DL servers are:

- Density optimized for flexibility and manageability in rack-mounted server environments
- Ideal for multi-server deployments
- Available in rack and cluster models

Compared to the ProLiant Gen9 line, 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 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 servers



Note

FlexibleLOM adapters for ProLiant Gen10 ML and DL series are designated with FLR (FlexibleLOM for rack servers) suffix.

- 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 2-52 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.

As indicated in [Figure 2-52](#), 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.

HPE ProLiant DL20 Gen10 Server



Figure 2-53 HPE ProLiant DL20 Gen10 Server

As illustrated in [Figure 2-53](#), the dense and compact HPE ProLiant DL20 Gen10 Server, powered by the Intel[®] Xeon[®] E Processor, Pentium, and Core i3 processors, provides a unique blend of enterprise-class capabilities and value. It offers outstanding configuration flexibility to cater to a wide variety of business requirements at an affordable price point. 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 iLO 5.

Key selling points include:

- Short-depth chassis of 15.05 inches provides mobility to deploy in a wide variety of environments.
- Affordable and versatile, from light workloads to the most demanding workloads.

- Offers outstanding configuration flexibility with six SFF or two LFF drives, 290 W, 500 W, or 800 W 48 VDC Redundant Power Supply, FlexibleLOM, and PCIe slots.

HPE ProLiant DL20 Gen10 Server is ideal for:

- Small and medium businesses and enterprises running general-purpose (IT, web, file, print, email) workloads
- OEM, transport, retail, and hospitality environments deploying custom applications (point of sale [PoS], ERP, billing)
- Service providers who want affordable, edge of network servers for scale-out deployments

Specifications of HPE ProLiant DL20 Gen10 Server include:

- Intel® Xeon® E-2XXX Series, six cores Pentium and i3 2/4 core processor options
- Four DDR4 (64 GB max)
- Up to six SFF or two LFF max, HDD/SSD, M.2 2280 NVMe SSDs (optional)
- Standard HPE Smart Array S100i Software RAID
- Choice of HPE Smart Array Essential and Performance RAID Controllers for performance or additional features
- 2x1 GbE embedded + choice of FlexibleLOM + Standup
- Two PCIe 3.0 slots
- Dedicated port for HPE iLO5 (option)
- HPE iLO 5, UEFI
- 290 W Standard NHP PSU (80Plus Silver) with up to 92% eff., 500 W Flexible slot Hot Plug Redundant PSU with 94% efficiency, 800 W 48VDC Flexible slot Hot Plug Redundant PSU with 94% efficiency
- Standard 3-3-3 warranty
- HPE Secure Encryption and TPM option, Intel® SGX
- Replaces DL20 Gen9



Note

More information about DL20 Gen10 can be located at:
<https://HPE.com/servers/dl20-gen10>

HPE ProLiant DL325 Gen10 Server



Figure 2-54 HPE ProLiant DL325 Gen10 Server

As shown in [Figure 2-54](#), the HPE ProLiant DL325 Gen10 Server is the secure and versatile single-socket server for virtualization and IO-intensive workloads. This 1P, 1U server delivers an exceptional balance of processor, memory, and IO for 2P performance at 1P TCO. Your customers will get the most out of their virtualized environment with this versatile server.

Key selling points include:

- Single-processor HPE ProLiant Gen10 server, based on the AMD EPYC™ processors
- Dual-processor performance at single-processor economics

HPE ProLiant DL325 Gen10 Server is ideal for:

- Virtualized environment
- VDI
- vSAN
- Big Data Analytics

Specifications of HPE ProLiant DL325 Gen10 Server include:

- One AMD EPYC 7000 series processor, 180 W 32 cores, PCIe 3.0, up to

three available slot(s)

- HPE Smart Memory 16 DDR4, up to 2666 MHz (2 TB max)
- Five to seven hot swappable redundant fans
- Standard HPE Dynamic Smart Array S100i. Choice of HPE Smart Array Essential or Performance Controllers for performance or additional features
- Ten SFF or four LFF Smart Drives, HDD/SSD, 10 SFF NVMe
- 4x1 GbE embedded + choice of FlexibleLOM (FLOM slot) + Standup
- Four USB 3.0 (one front standard, one interior, two rear)
- Up to one single width active only
- HPE iLO 5 4 GB NAND includes:
 - Embedded iLO physical serial port command-line interface (CLI)
 - Dedicated iLO NIC port
- 92% 550 W power supply optimized, 96% (Titanium), 500 W, 800 W, 1600 W Flex Slot Power Supply
- ASHRAE A3 and A4
- 1U, 24.2”
- Chassis intrusion detection option/digital signed firmware/TPM 2.0
- Easy install (standard) and CMA
- Standard 3-3-3 warranty
- HPE-qualified options only



Note

More information about the DL325 Gen10 can be located at:
<https://HPE.com/servers/dl325-gen10>

HPE ProLiant DL360 Gen10 Server



Figure 2-55 HPE ProLiant DL360 Gen10 Server

The HPE ProLiant DL360 Gen10 Server, as illustrated in [Figure 2-55](#), 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 2666/2933 MT/s HPE DDR4 SmartMemory supporting up to 3.0 TB max (or 4.5 TB with Gen2 processors), 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 x 128 GB 2666 MT/s LRDIMM, 768 GB when populated with 24 x 32 GB @ 2666 MT/s or 2933 MT/s RDIMM memory, and 192 GB 2 x 16 GB NVDIMM @ 2666 MT/s. Gen2 processors are capable of addressing up to 4.5 TB per processor.

Key selling points of the ProLiant DL360 Gen10 include:

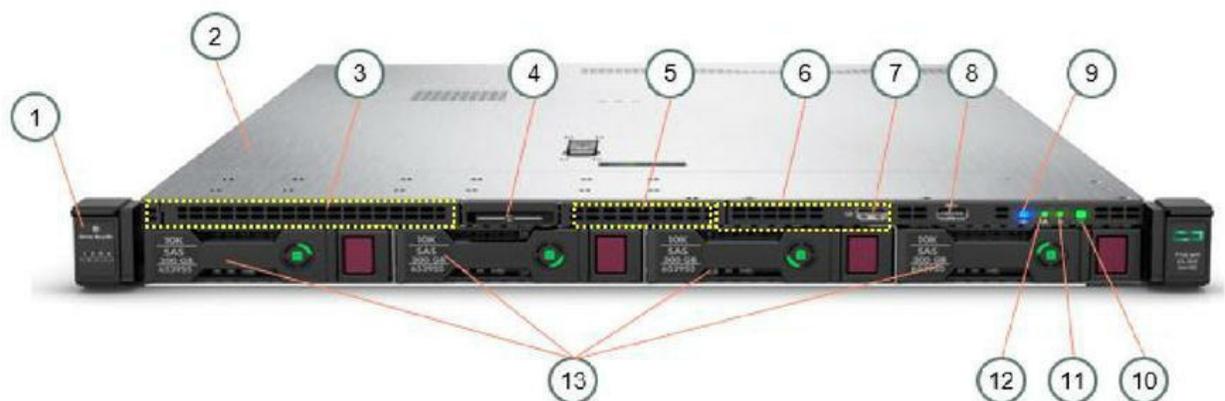
- Leading energy-efficient dense server for better ROI
- Higher performance with Intelligent System Tuning (IST), Persistent Memory, and greater NVMe capability
- Increased flexible storage options to support various workloads

- PCIe expansion with HPE FlexLOM and HPE Flexible Smart Array

The ProLiant DL360 Gen10 Server may be ideal for:

- Hypervisor
- Dynamic workloads in dense-virtualized environments
- Compute-intensive applications (web caching, data analytics)
- Low latency and transactional applications (warehouse and database)
- High-Performance Computing (HPC) focus on the financial services industry (FSI)

HPE ProLiant DL360 Gen10 Server: Four LFF front of system detail



Options:

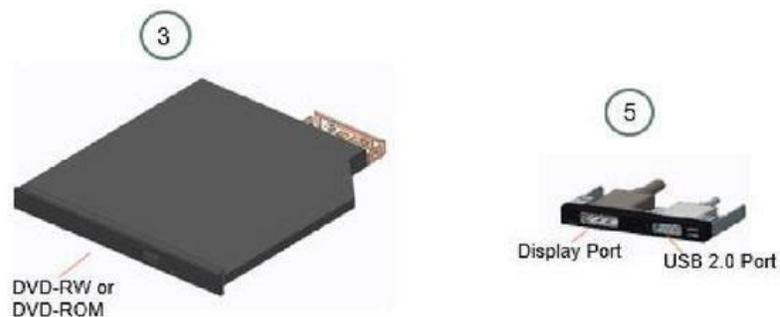


Figure 2-56 ProLiant DL360 Gen10 Server with four LFF, front view

Key features of the ProLiant DL360 Gen10 (front) as shown in [Figure 2-56](#)

are:

1. Drive support label
2. Quick removal access panel
3. Optional: DVD-RW or DVD-ROM (blank shown)
4. Serial no. label pull tab
5. Optional: Display port and USB 2.0 port Kit (blank power LED shown)
6. Optional: System Insight Display (SID)*—std. shown *This option will lose #7 iLO service port
7. iLO service port
8. USB 3.0 port
9. UID button/LED
10. Power On/Standby button and system
11. Health LED
12. NIC Status LED
13. SAS/SATA/SSD drive bays

ProLiant DL360 Gen10: Eight SFF front of system detail



Figure 2-57 ProLiant DL360 Gen10 Server with eight LFF, front view

Key features of the ProLiant DL360 Gen10 (front) as shown in [Figure 2-57](#) are:

1. Drive support label
2. Quick removal access panel
3. Serial no. label pull tab
4. Universal Media Bay options:
 - a. Option shown: +2 SFF SAS/SATA (total max 10SFF)
 - b. Optional: +2 SFF NVMe drives
 - c. Optional: DVD-RW or DVD-ROM
 - d. Optional: +2 Dual uFF (4x M.2 cartridges)
 - e. Optional: Display port + USB 3.0 Port
5. UID button/LED

6. Power On/Standby button and system power LED
7. Health LED
8. NIC status LED
9. USB 3.0 port
10. iLO Service Port
11. Eight SAS/SATA/SSD drive bays

HPE ProLiant DL360 Gen10: New premium 10 SFF NVMe

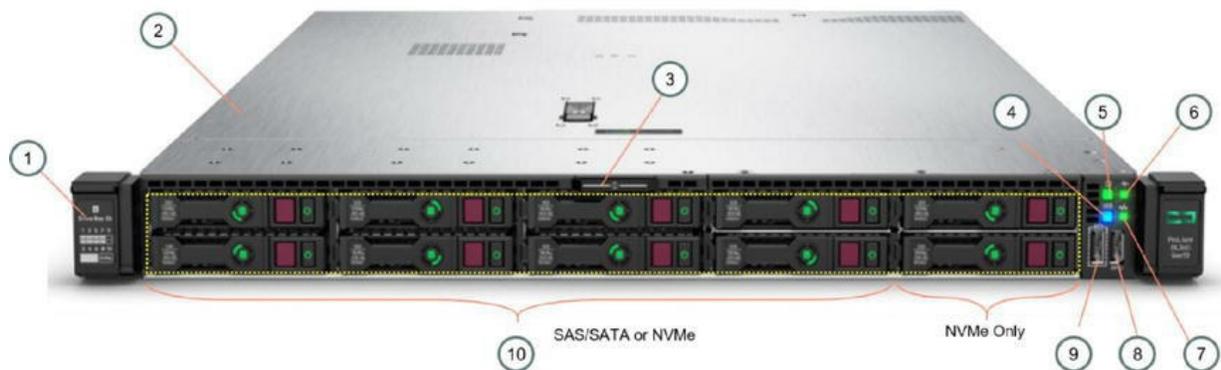


Figure 2-58 ProLiant DL360 Gen10 Server 10 SFF NVMe, front view

Key features of the ProLiant DL360 Gen10 (front) as shown in [Figure 2-58](#) are:

1. Drive support label
2. Quick removal access panel
3. Serial no. label pull tab
4. UID button/LED
5. Power On/Standby button and system power LED

- 6. Health LED
- 7. NIC status LED
- 8. USB 3.0 port
- 9. iLO service port
- 10. Max 10 NVMe drive bays (PCIe direct attached)

*Mix-n-match SAS/SATA or NVMe:

- Drive bays 1 to 8 support SAS/SATA and NVMe
- Drive bays 9 and 10 are NVMe ONLY

ProLiant DL360 Gen10: Interior system detail—SFF chassis

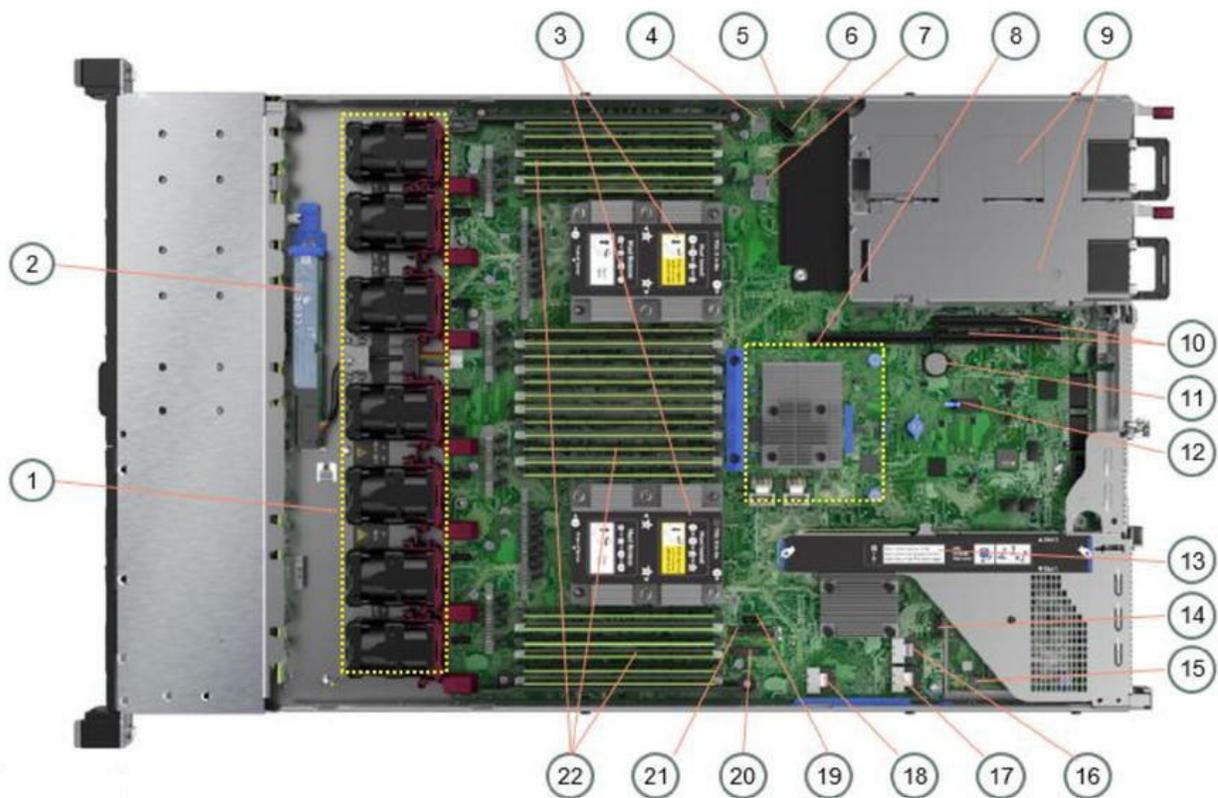


Figure 2-59 ProLiant DL360 Gen10 Server interior view

Key features of the ProLiant DL360 Gen10 (interior) as shown in [Figure 2-59](#) are:

1. Standard single rotor hot-plug fans (for eight SFF and four LFF chassis only, 10 NVMe chassis)
 - One CPU—five standard fans
 - Two CPUs—seven standard fans
 - Optional: High-performance fans
 - Will ship with seven high-performance fans as standard
2. Optional: Smart Storage Battery
3. Up to two processors (shown with standard heat sinks)
4. MicroSD card slot (Dual Micro-SD option available)
5. Optional: Chassis Intrusion Detection
6. Hard drive backplane power connector
7. Dual internal USB 3.0 connector
8. Smart Array Controller (type -a shown)
9. Up to two power supplies for redundant power
10. Secondary (CPU2) PCIe 3.0 riser (not available on Premium 10 SFF NVMe chassis due to dedicated 10 x4 NVMe riser for PCIe direct attached)
 - Optional: Low profile x16 or 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 1x8
- Optional: Two SATA M.2 + 2 x16
- Optional: Two 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

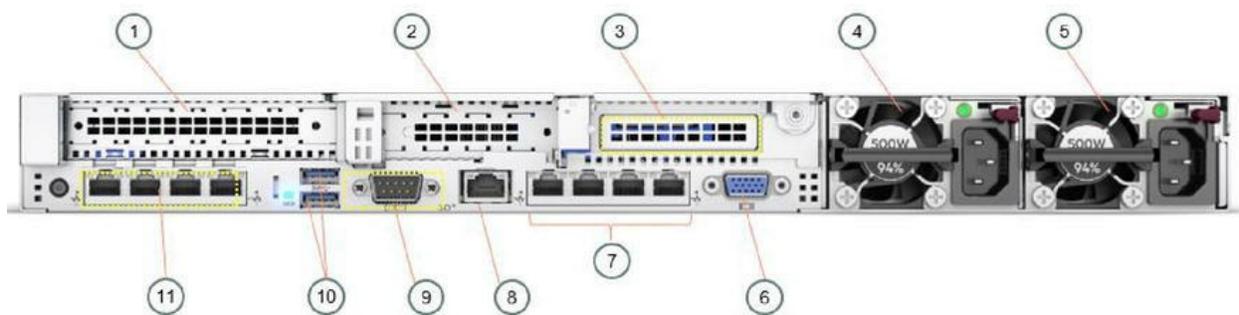


Figure 2-60 ProLiant DL360 Gen10 Server, rear system detail

Key features of the ProLiant DL360 Gen10 (rear) as shown in [Figure 2-60](#) are:

1. Slot 1 PCIe 3.0 Optional: Rear drive +1 SFF or 1 uFF (2x M.2 cartridges)

2. Slot 2 PCIe 3.0
3. Optional: Slot 3 PCIe 3.0 (Requires 2nd processor)
4. Power Supply 2
5. Power Supply 1
6. VGA port
7. Embedded 4x 1 GbE adapter
8. iLO management port
9. Optional: Serial port
10. USB 3.0 ports
11. Optional: FlexibleLOM* (shown 4x 1 GbE) *Supports various FlexibleLOM NICs up to 25 GbE

Specifications of HPE ProLiant DL360 Gen10 Server include:

- Up to two Intel® Xeon® Skylake/Cascade Lake Series, up to 28 cores
- 24 DDR4 2666/2933 MHz
- Persistent Memory up to 12 NVDIMM options
- 10 NVMe/eight +two SFF/four LFF max, HDD/SSD, M.2 SATA/PCIe Enablement Kit and NVMe PCIe SSD option
- Std. HPE Dynamic Smart Array S100i. Choice of HPE Flexible Smart Array or HPE Smart HBA Controllers for performance or additional features
- 4x1 GbE embedded + choice of FlexibleLOM or Standup
- Three PCIe 3.0 slots
- HPE iLO 5, UEFI, and legacy boot mode
- Up to 96% efficiency (Titanium), Flexible Slot PS, and Battery Backup Unit

- Standard 3-3-3 warranty
- HPE Secure Encryption, TPM option, and chassis intrusion
- Replaces DL360 Gen9 and DL360 Gen8



Note

More information about the DL360 Gen10 can be located at:
<https://HPE.com/servers/dl360-gen10>

HPE ProLiant DL380 Gen10 Server



Figure 2-61 HPE ProLiant DL380 Gen10 Server

The HPE ProLiant DL380 Gen10 Server as shown in [Figure 2-61](#) 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 make 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 2666/2933 MT/s HPE DDR4 SmartMemory supporting up to 3.0/4.5 TB max, 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 business needs.
- Huge storage footprint for large storage workloads.

- PCIe expansion with HPE FlexLOM and HPE Flexible Smart Array.

The ProLiant DL380 Gen10 Server may be ideal for SMB or enterprise customers with workloads/apps: 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.

The HPE ProLiant DL380 Gen10 Server includes:

- Up to two Intel[®] Xeon[®] Skylake/Cascade Lake Series, up to 28 cores
- 24 DDR4 2666/2933 MHz (1.5 TB max)
- Persistent Memory up to 12 NVDIMM options
- Up to 30 SFF or 19 LFF max, HDD/SSD, M.2 Enablement Kit and 20 NVMe PCIe SSD option
- Std. HPE Dynamic Smart Array S114i. Choice of HPE Flexible Smart Array or HPE Smart HBA Controllers for performance or additional features
- 4x1 GbE embedded + choice of FlexibleLOM + Standup
- Seven PCIe 3.0 slots (FH/FL)
- HPE iLO 5, UEFI, and legacy boot mode
- Up to 96% efficiency (Titanium), Flexible Slot PS, and Battery Backup Unit
- Standard 3-3-3 warranty
- HPE Secure Encryption, TPM option, and chassis intrusion
- Replaces DL380 Gen9 and DL380p Gen8

HPE ProLiant DL380 Gen10—Front system detail

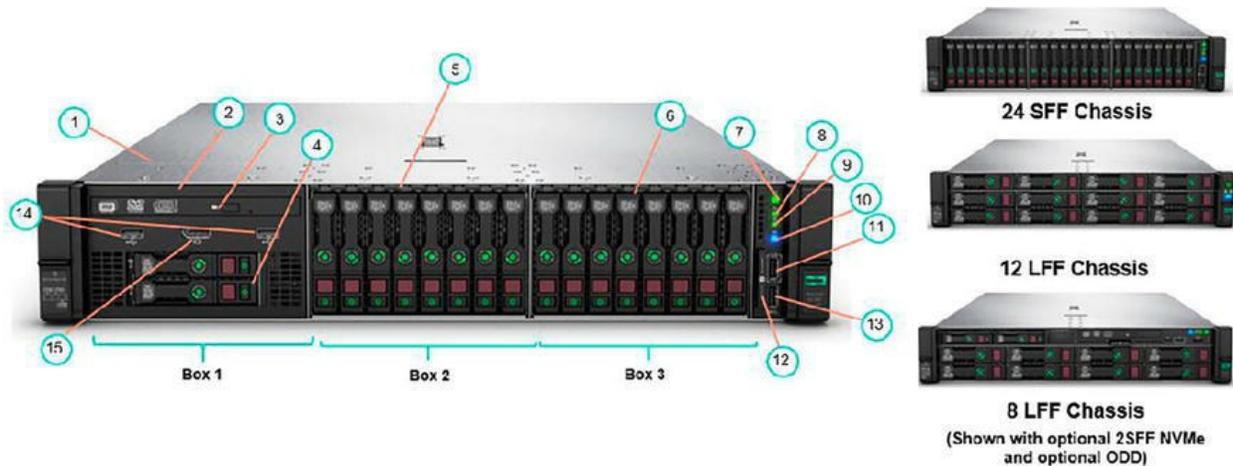


Figure 2-62 ProLiant DL380 Gen10 Server, front view and chassis

Key features of the ProLiant DL380 Gen10 (front) as shown in [Figure 2-62](#) are:

1. Quick removal access panel
2. Optional Universal Media Bay. Two USB 2.0 and display port (eight SFF bay or six SFF +two NVMe or eight NVMe optional)
3. Optional optical drive. Requires Universal Media Bay
4. Optional two SFF HDD, requires optional Universal Media Bay
5. Drive Bay 2. NVMe shown (eight SFF, six SFF + two NVMe or eight NVMe PCIe SSD optional)
6. Eight SFF drive cage bay
7. Power On/Standby button and system power LED button
8. Health LED
9. NIC status
10. UID 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 DL380 Gen10: Interior of system detail—Eight SFF

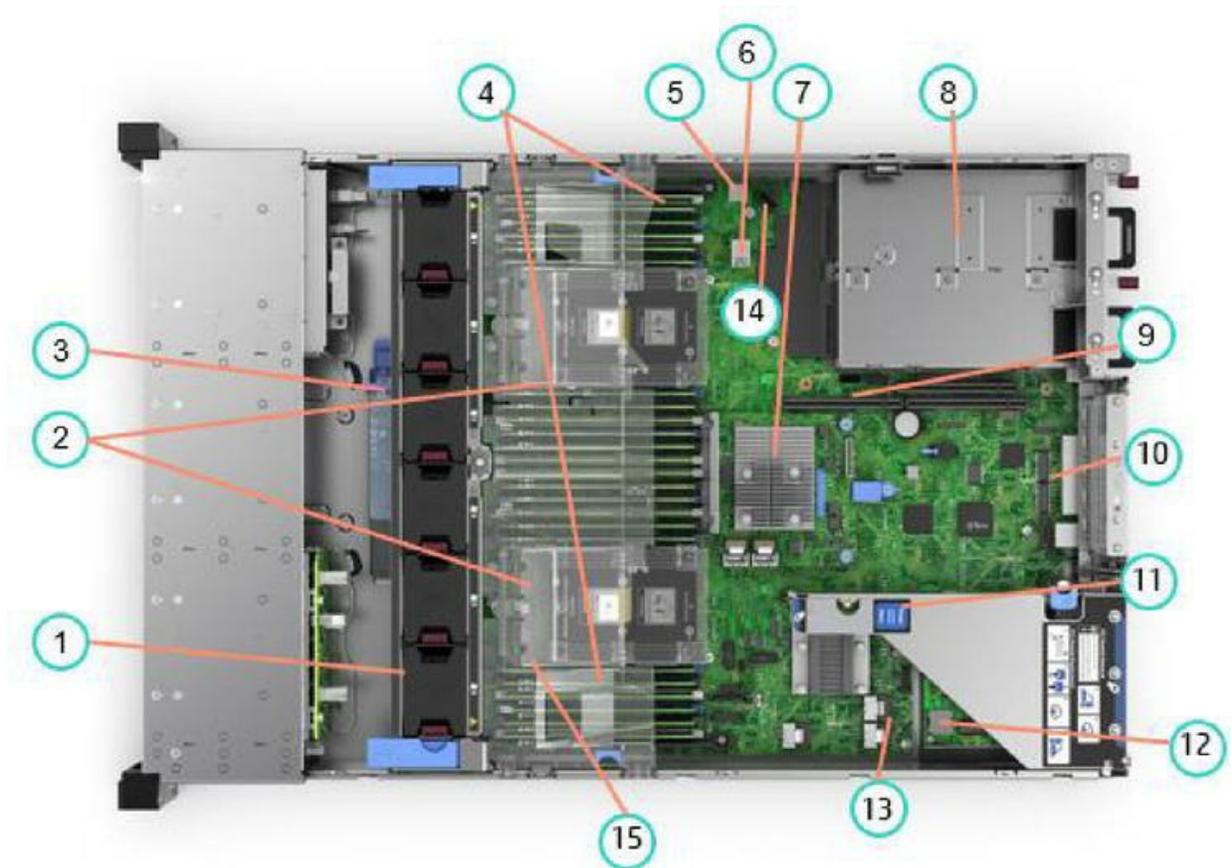


Figure 2-63 ProLiant DL380 Gen10, interior layout

Key features of the ProLiant DL380 Gen10 (interior) as shown in [Figure 2-63](#) are:

1. Fan cage shown with six standard hot-plug fans (High Ambient temperature fans optional)
2. Two processors, heatsink showing
3. Optional HPE Smart Storage Battery
4. DDR4 DIMM slots. Shown fully populated in 24 slots (12 per processor)
5. MicroSD card slot (optional dual Micro-SD option)
6. Internal USB 3.0 connector
7. 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 4x1 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

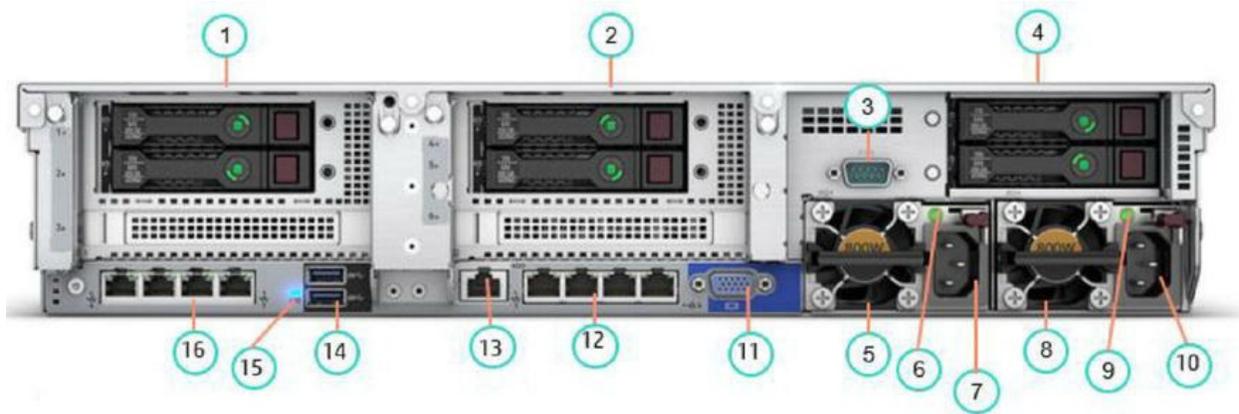


Figure 2-64 ProLiant DL380 Gen10, rear view

Key features of the ProLiant DL380 Gen10 (rear) as shown in [Figure 2-64](#) are:

1. Primary riser. PCI slots (slots 1 to 3 top to bottom, riser shipped standard, not shown), optional two
2. SFF rear drives
3. Secondary riser. PCI slots (slots 4 to 6 top to bottom, not shown, requires second riser card, and second processor). Showing optional 2 SFF rear
4. Optional serial port
5. Tertiary riser (slots 7 to 8). Optional rear 2 SFF HDD (supported in 24 SFF or 12 LFF front end)
6. HPE Flexible Slot Power Supply bay 1 (800 W shown)
7. Power supply Power LED
8. Power supply Power connection HPE Flexible Slot Power Supply bay 2 (800 W shown)
9. Power supply power LED
10. Power supply power connection

11. VGA connector
12. Embedded 4 x 1 GbE Network
13. Adapter
14. Dedicated iLO management port
15. Two USB connectors 3.0
16. Unit ID LED Optional FlexibleLOM ports (shown: 4 x 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 makes it an ideal server for many environments.

Other key features of the ProLiant DL380 Gen10 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 3xDW or 5xSW cards
- Additional boot/drive/rear options: SATA M.2; dual uFF SSD (2x 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
- Embedded Security features: Silicon Root of Trust; digitally signed FW
- Optional Security features: Chassis Intrusion Detection; TPM 2.0; TPM 1.2

There are several drive bay options:

- One to three SFF drive bays support eight SFF or NVMe drives for a total of 24 with additional six SFF rear drive bay options 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 options to total 24 SFF drives.
- Or the 12 LFF with optional four LFF mid-plane and optional three LFF + two SFF rear drive bays 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. The maximum memory capacity is limited by processor selection and DIMM type. Mixing of RDIMM and LRDIMM memory is not supported. Currently, the maximum memory capacity is limited to 1.5 TB, assuming 24 x 64 GB LRDIMM and 768 GB when populated with 24 x 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 eight-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 network technology, which offers customers a choice of HPE 1 GB Ethernet four-Port 331i Adapter plus optional HPE FlexibleLOM or stand up card, available in Entry and Base models. In the Performance models, the HPE 1 Gb Ethernet four-Port 331i Adapter plus HPE Ethernet 10/25 GB two-port 640FLR-SFP28 Adapter or the two-port 631FLR-SFP28 Adapters are available. A range of NIC cards are also available to enhance networking capabilities.

Flex slot 500 W, 800 W, and 1600 W power supplies are available as Platinum Low Halogen models. In addition, there are 800 W 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

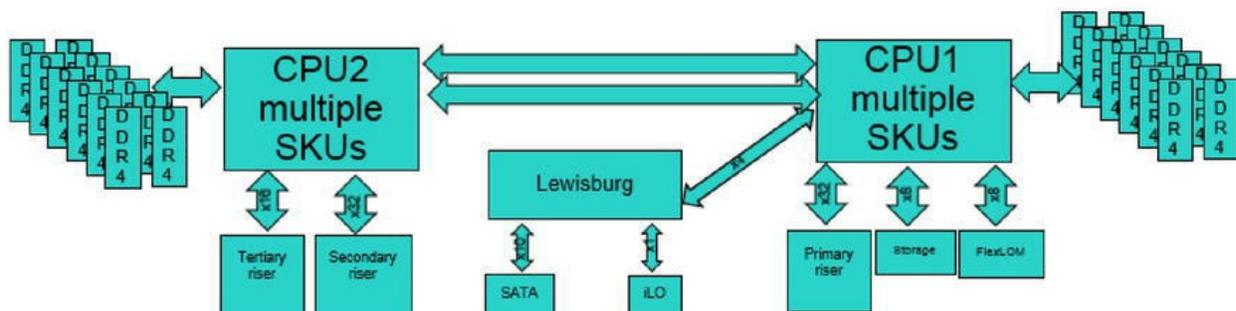


Figure 2-65 ProLiant DL380 Gen10, system/block diagram

The system block diagram illustrates in [Figure 2-65](#) 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 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.

HPE ProLiant DL385 Gen10 Server



Figure 2-66 The right-sized, enterprise-grade server ideal for virtualized workloads

As illustrated in [Figure 2-66](#), the HPE ProLiant DL385 Gen10 Server is an enterprise-grade, dualsocket 2U rack server that delivers enhanced performance, enterprise-class management, expandability, and scalable capabilities to meet the needs of general and heavy workloads. This is 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.

Key selling points include:

- The DL385 takes full advantage of the available compute to remove bottlenecks in memory and IO.
- The DL385 provides the right amount of compute to get the job done efficiently. Up to 64 cores and 32 DIMMs.
- Security is top of mind, leaving teams more at ease.

Ideal for virtualization, general purpose compute, memory-centric workloads,

high-core count workloads, SharePoint storage, and enterprise applications.

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 x 1 GbE, FlexibleLOM, Standup
Management	iLO 5
Power	Redundant Flex Slot
Additional	Embedded M.2, USB, and MicroSD

HPE ProLiant DL385 Gen10 options

DL385 Gen10	
Compute	Up to (2) AMD® EPYC® 7000 Series Processor family, (32C), PCIe 3.0, up to 8 available slot(s) + m.2 support embedded
Memory	HPE Smart Memory (32) DDR4 (8 channels per CPU), up to 2667MT/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 SmartDrives	24+6 SFF / 12+4+3 LFF+2SFF max, HDD/SSD, 24 NVMe PCIe SSD and m.2 connector embedded on system board.
Networking	4x1 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 (5) & Active/Passive up to 10.5" (3)
Management	
Converged Support	HPE OneView and HPE iLO Advanced
Embedded	HPE iLO 5, SUM, RESTful Interface Tool, UEFI
Power & Cooling	Up to 96% eff. (Titanium) with Flexible Slot PS
Industry Compliance	ASHRAE A3 & 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

Figure 2-67 ProLiant DL385 Gen10 options

Figure 2-67 shows options available for the ProLiant DL385 Gen10.

HPE ProLiant DL385 Gen10—Front system detail

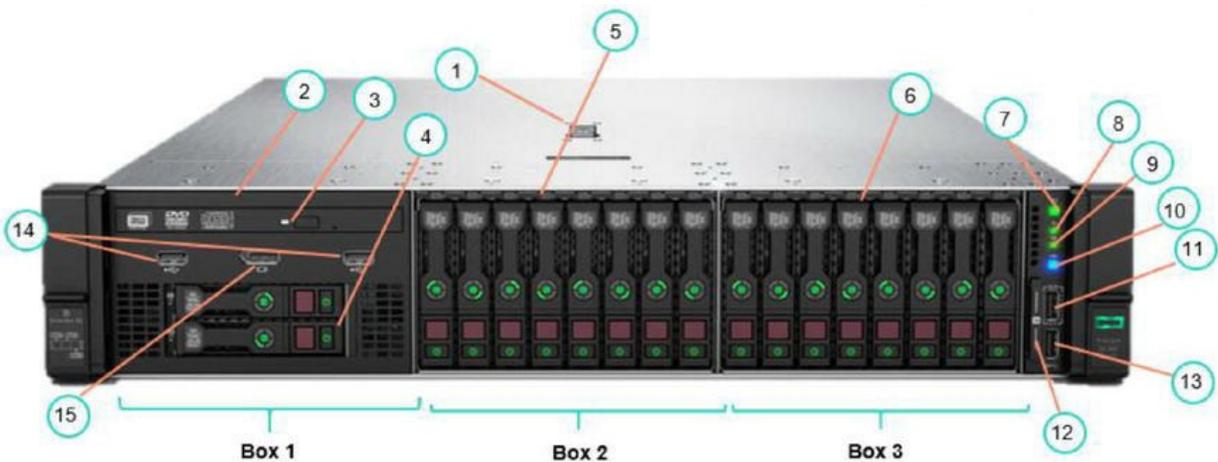


Figure 2-68 HPE ProLiant DL385 front view

Key features of the ProLiant DL385 Gen10 as shown in [Figure 2-68](#) are:

Front view:

1. 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 optional)
2. NVMe or eight NVMe optional)
3. Optional optical drive. Requires Universal Media Bay
4. Optional two SFF HDDs. Requires optional Universal Media Bay
5. Drive Bay 2. NVMe shown (eight SFF, six SFF + two NVMe or eight NVMe PCIe SSDs optional)
6. Drive Bay 3. NVMe shown (eight SFF, six SFF + two NVMe or eight NVMe PCIe SSDs optional)
7. 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

15. Universal Media Bay) Optional front display port (via Universal Media Bay)

HPE ProLiant DL385 Gen10: Interior of system detail (eight SFF)

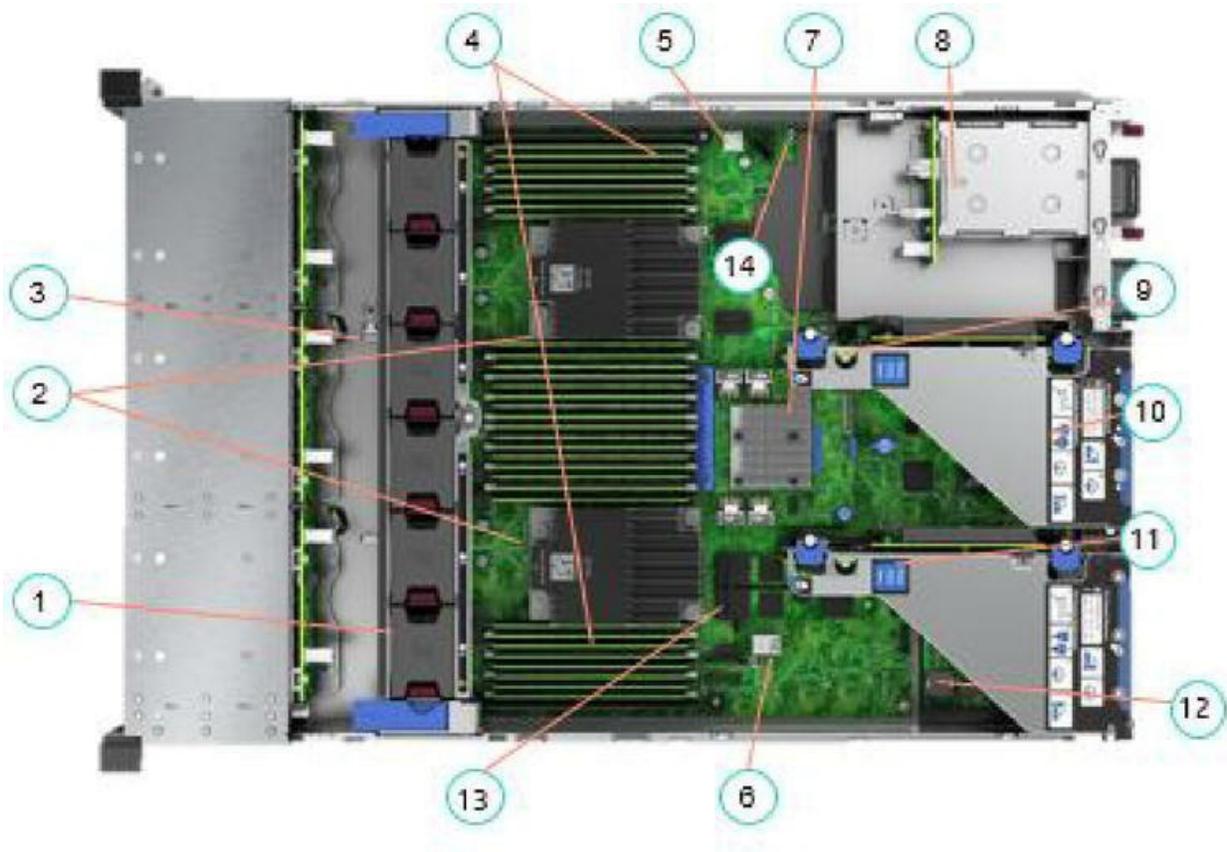


Figure 2-69 HPE ProLiant DL385 interior view

Key features of the ProLiant DL385 Gen10 (interior) as shown in [Figure 2-69](#) are:

1. Fan cage shown with six standard hot-plug fans (high-performance fans optional)
2. Two processors (heatsinks shown)
3. Optional: Smart Storage Battery location (battery not shown)
4. DDR4 DIMM slots. Shown fully populated in 32 slots (16 per processor)
5. MicroSD card slot (optional dual micro-SD option)
6. Internal USB 3.0 connector
7. 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 4x1 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

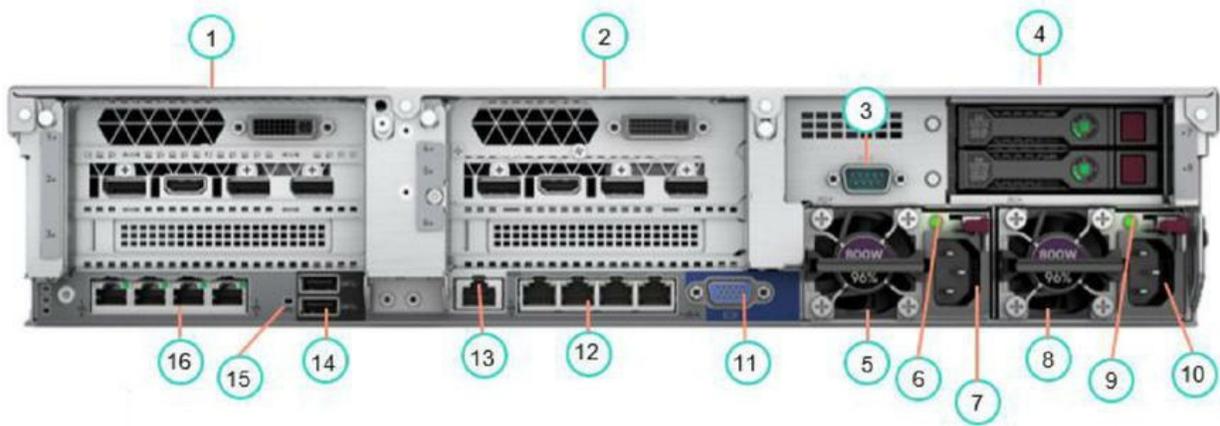


Figure 2-70 HPE ProLiant DL385 rear view

Key features of the ProLiant DL385 Gen10 (rear) as shown in [Figure 2-70](#) are:

1. Primary riser. PCI slots (slots 1 to 3 top to bottom, riser shipped standard, not shown). Optional two SFF rear drives
2. Secondary riser. PCI slots (slots 4 to 6 top to bottom, not shown, requires second riser card and second processor).
3. Optional two SFF rear drives
4. Optional serial port Tertiary riser (slots 7 to 8 top to bottom, not shown).
5. Optional rear two SFF HDDs
6. HPE Flexible Slot power supply bay 2
7. Power supply power LED
8. Power supply power connection HPE Flexible Slot Power Supply bay 1
9. Power supply power LED
10. Power supply power connection
11. VGA connector

12. Embedded 4 x 1 GbE Network
13. Adapter
14. Dedicated iLO management port
15. Two USB connectors 3.0
16. Unit ID LED Optional FlexibleLOM ports (shown: 4 x 1 GbE)



Note

More information about the DL360 Gen10 can be located at:
<https://HPE.com/servers/dl385-gen10>

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 CPU cores. It provides a dedicated secure space to run multiple security-related functions that require full isolation and can be off-loaded from the main CPU core. 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 below.

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)

Migrating VMs within the data center, to off-premises private cloud, or to the public

servers, APIs,
third-party key
management)

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 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 SSD) 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 are 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 Apollo 2000 Gen10 System



Figure 2-71 Enterprise bridge to HPC

As seen in [Figure 2-71](#), 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. 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 require 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/admins). 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 non-disruptive manner to become enterprise service providers.

The 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 Apollo 8000. Apollo 8000 is a water-cooled supercomputer which delivers 250 teraflops of compute per rack with game changing data center efficiency with a power usage effectiveness (PUE) of 1.1. It is ideal for very complex numerical simulations and mathematical modeling applications that require massively parallel computing capabilities. The 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.



Note

PUE is a ratio of how efficiently a computer data center uses energy—specifically, how much energy is used by the computing equipment (in contrast to cooling and other overhead).

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 fits 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 Apollo 2000—r2800 Chassis-24 SFF



Figure 2-72 HPE Apollo r2800 Chassis

The Apollo 2000 System brings HPE ProLiant Gen10 server technology, including iLO 5, into this 2U, multi-server chassis.

The HPE ProLiant XL170r Gen10 Server and the HPE ProLiant XL190r Gen10 Server offer more configuration choices that cover a much wider range of scale-out workloads. Storage and IO flexibility enable customers to optimize for performance or economy—the right compute for the right workload.

As illustrated in [Figure 2-72](#), the r2800 Chassis, with 24 SFF drives, allows customers 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 Cluster Management Utility (CMU) to monitor, manage, and optimize compute clusters of any size

Apollo 2000 System offerings

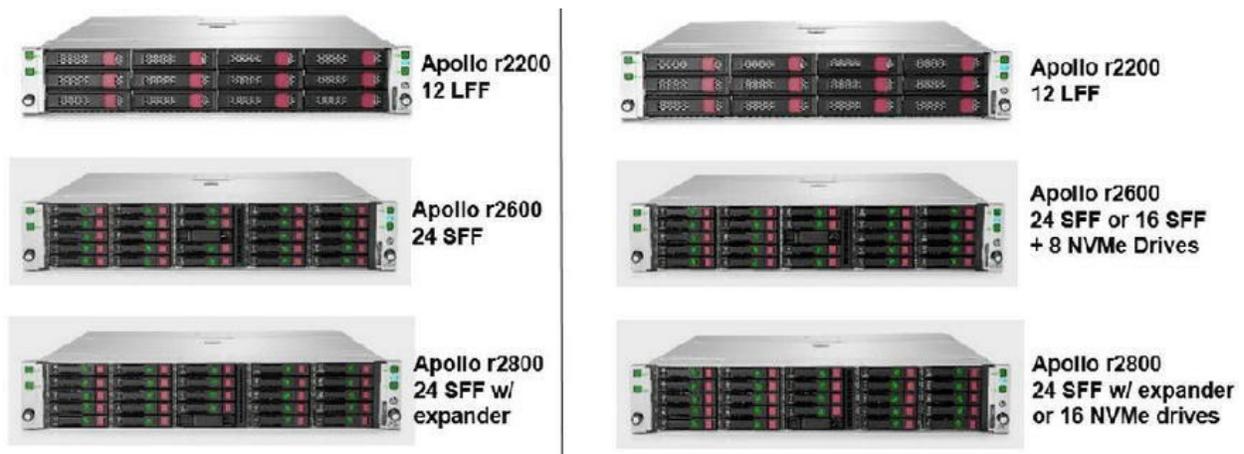


Figure 2-73 Apollo 2000 system offerings

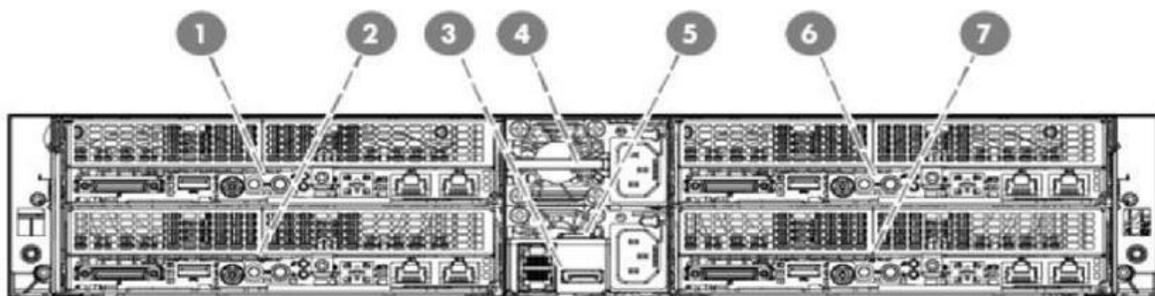
As seen in [Figure 2-73](#), Apollo 2000 can support up to four servers in 2U and offers three different enclosures to address different needs:

- **Apollo r2200 with 12 LFF**—If your customer’s focus is low cost, high capacity, and storage is the key, go with Apollo r2200 LFF with twelve 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 SFF with 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 + eight 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 0 to 24 drives allocated. For ultimate performance, the 16 NVMe drive option will 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 ToR management switch, eliminating the need of multiple iLO cables.

HPE Apollo 2000—r2800 Chassis four 1U nodes

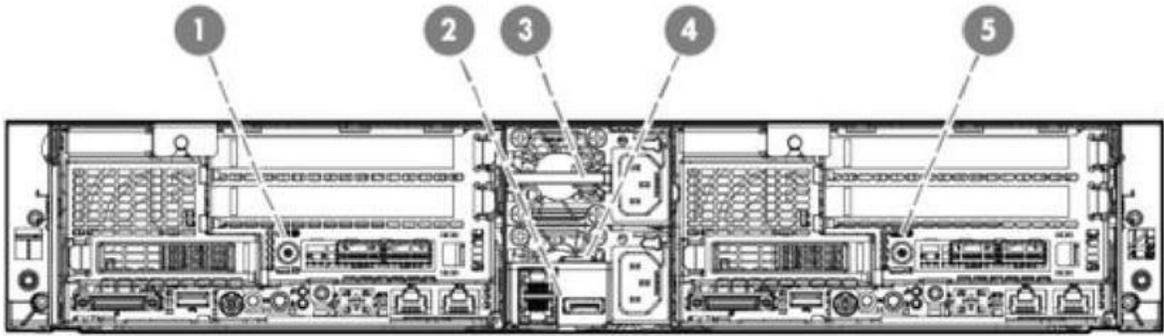


Item	Description
1	Node 4
2	Node 3
3	RCM module
4	Power Supply 2
5	Power Supply 1
6	Node 2
7	Node 1

Figure 2-74 Apollo 2000—r2800 Chassis four 1U nodes

Figure 2-74 shows the major components of the r2800 chassis with four 1U nodes.

HPE Apollo 2000—r2800 Chassis two 2U nodes



Item	Description
1	Node 3
2	RCM module
3	Power Supply 2
4	Power Supply 1
5	Node 1

Figure 2-75 Apollo 2000—r2800 Chassis two 2U nodes

Figure 2-75 shows the major components of the r2800 chassis with two 2U nodes.

HPE Apollo 2000—r2800 Chassis fans

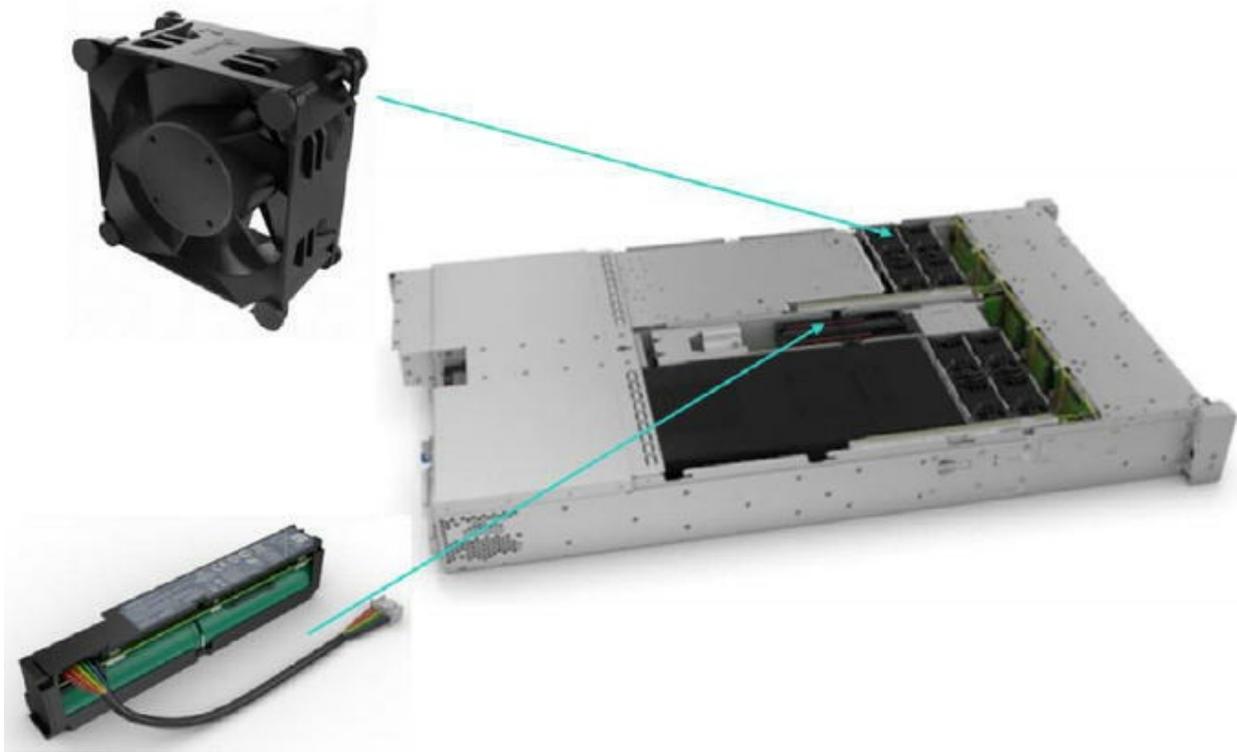


Figure 2-76 r2800 chassis fans and the HPE Smart Battery

[Figure 2-76](#) shows 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 condition, 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 the HPE Smart Array Px4x Controllers.

HPE Apollo 2000—r2800 Rack Consolidation Management Module



Figure 2-77 r2800 Rack Consolidation Management Module

As shown in [Figure 2-77](#), the RCM Module kit allows iLO aggregation at the chassis level and also has the port for the HPE Advanced Power Manager (APM). The RCM Module consists of two 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, 100Gv/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)

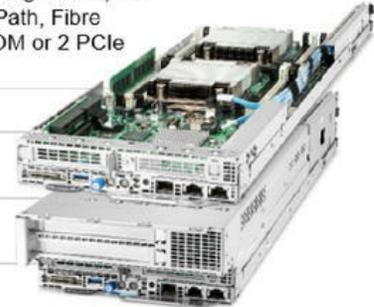


Figure 2-78 ProLiant XL170r and XL190r features

As illustrated in [Figure 2-78](#), HPE ProLiant XL170r and XL190r features include:

- Intel® Xeon® Processor Scalable Family up to 2x 145 W (four named shelves—Platinum, Gold, Silver, and Bronze) @ 35C Ambient with Roadmap processors
- Intel® Storm Lake (Omni-path) enabled motherboard (CPU1)
- 16 DDR4 Memory Sockets supporting RDIMM and LRDIMM up to 2666 MHz@1DPC and 2400 MHz@2DPC
- Type4 NVDIMM up to 2400 MHz, Up to four NVDIMMs (two NVDIMMs per processor)
- Type1 NVDIMM (Trumpkin) up to 2666 MHz@1DPC and 2400 MHz@2DPC

HPE Apollo 2000—XL170r Gen10



Density-optimized general purpose cloud server with all-in-one high-performance computing/storage

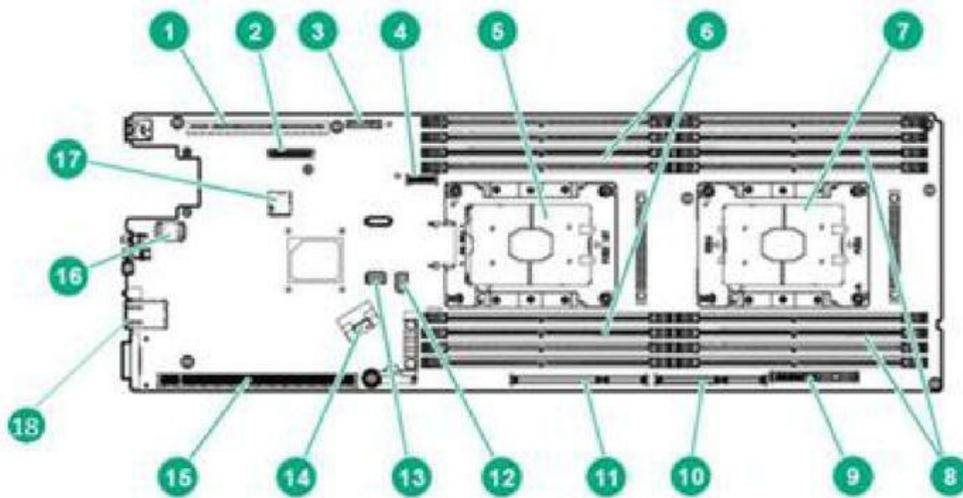
Figure 2-79 ProLiant XL170r Gen10 Server

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. As illustrated in [Figure 2-79](#), the ProLiant XL170r Gen10 Server is a 1U half-width, two-processor server. Rackscale power efficiency is easy with HPE Apollo Platform Manager for rack, chassis, server, and component level management.

The ProLiant XL170r Gen9 Server is a 1U half-width, two-processor server with configuration options for:

- Performance and efficient CPUs via 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



Item	Description	Item	Description
1	Primary riser slot 1	10	Secondary riser slot 4
2	mLOM connector	11	Secondary riser slot 3
3	System battery	12	Slimline connector (SATA x4)
4	M.2 SSD riser connector	13	System maintenance switch
5	Processor 1	14	Slimline connector (SATA x8)
6	DMMs for processor 1	15	Secondary riser slot 2
7	Processor 2	16	TPM connector
8	DIMMs for processor 2	17	microSD slot
9	Bayonet board slot	18	iLO dedicated port

Figure 2-80 HPE Apollo 2000—XL170r Gen10

Figure 2-80 shows the major components of the XL170r Gen10 Server.

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 2-81 ProLiant XL190r Gen10 Server

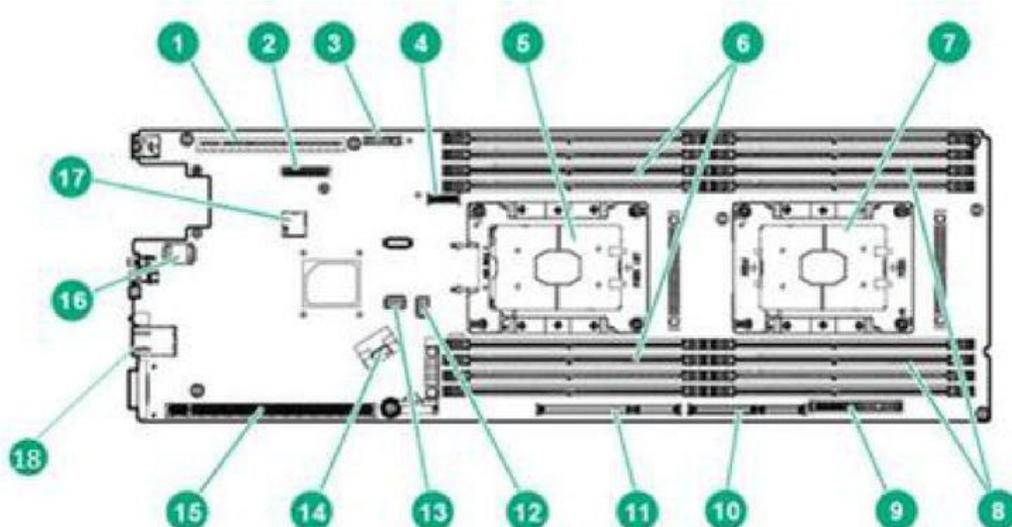
As shown in [Figure 2-81](#), the ProLiant XL190r Gen10 Server is a 2U half-width, two processor server with the same configuration options as the XL170r for CPU and memory, but has additional PCIe slots in multiple configurations providing support for additional expansion cards and future support for two integrated accelerators per server (GPUs or coprocessors).

This server, shown in [Figure 2-81](#), leverages Intel's® latest Xeon® Scalable Family series processors options ranging from Bronze to Platinum, increasing performance up to 30 to 40%, 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 fueled by the optional HPE Apollo Platform Manager for rack, chassis, server, and component level management.

HPE Apollo 2000—XL190r Gen10: Top view



Item	Description	Item	Description
1	Primary riser slot 1	10	Secondary riser slot 4
2	mLOM connector	11	Secondary riser slot 3
3	System battery	12	Slimline connector (SATA x4)
4	M.2 SSD riser connector	13	System maintenance switch
5	Processor 1	14	Slimline connector (SATA x8)
6	DIMMs for processor 1	15	Secondary riser slot 2
7	Processor 2	16	TPM connector
8	DIMMs for processor 2	17	microSD slot
9	Bayonet board slot	18	iLO dedicated port

Figure 2-82 HPE Apollo 2000—XL190r Gen10

The Apollo 2000 chassis can accommodate two of the XL190r trays; however, you can mix and match types if necessary.

HPE Synergy: Powering your customer's hybrid cloud transformation



Figure 2-83 HPE Synergy Composable frame

As illustrated in [Figure 2-83](#), HPE 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.

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

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 fabric pools into any configuration for any application. As an extensible platform, it easily enables a broad range of applications and operational models such as virtualization, hybrid cloud, and 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. 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 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 VMs and cloud.

The Synergy platform is unique because it enables customers to create a completely stateless infrastructure. Only HPE allows configuration and provisioning of compute, fabric, storage, and now 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.

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. The key components of Synergy are:

- **Composer**—Composer leverages integrated software-defined intelligence to accelerate operations through a single interface.
- **Streamer**—Image Streamer enables stateless computing and very rapid provisioning, updating, and reprovisioning of servers. 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 DAS, storage area network (SAN), software-defined storage (SDS), and file, block, and object storage.
- **Fabric**—Rack-scale multi-fabric connectivity eliminates stand-alone ToR switches.

Synergy Composer

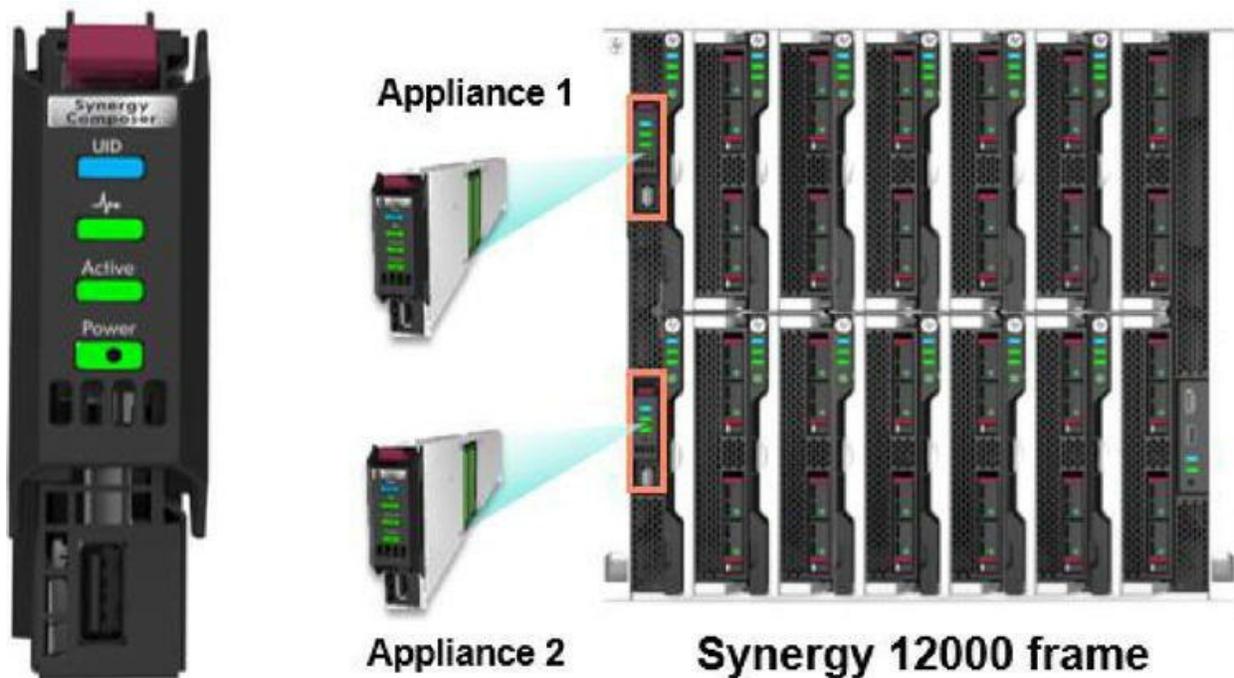


Figure 2-84 Synergy Composer plugs into an appliance bay in the Synergy frame

As seen in [Figure 2-84](#), Composer is a physical appliance integrated within the Synergy frame. It plugs into an appliance bay in the side of the Synergy frame.

 **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 embeds 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 new 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, 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.

Synergy Image Streamer



Figure 2-85 Synergy Image Streamer

To accelerate workload deployment, Synergy uses the Image Streamer, as shown in [Figure 2-85](#), 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 of infrastructure blocks.

Traditional server deployment is a sequential process of provisioning the physical hardware, followed by provisioning an operating system, and then by provisioning a hypervisor installation. 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 (preinstall) environment**.

Image Streamer enables a true stateless operation by integrating server profiles with golden images (operating environment and IO driver) 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 of 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, 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 implementation requires other components and affects the solution design.

HPE Synergy fabric and Interconnect Modules

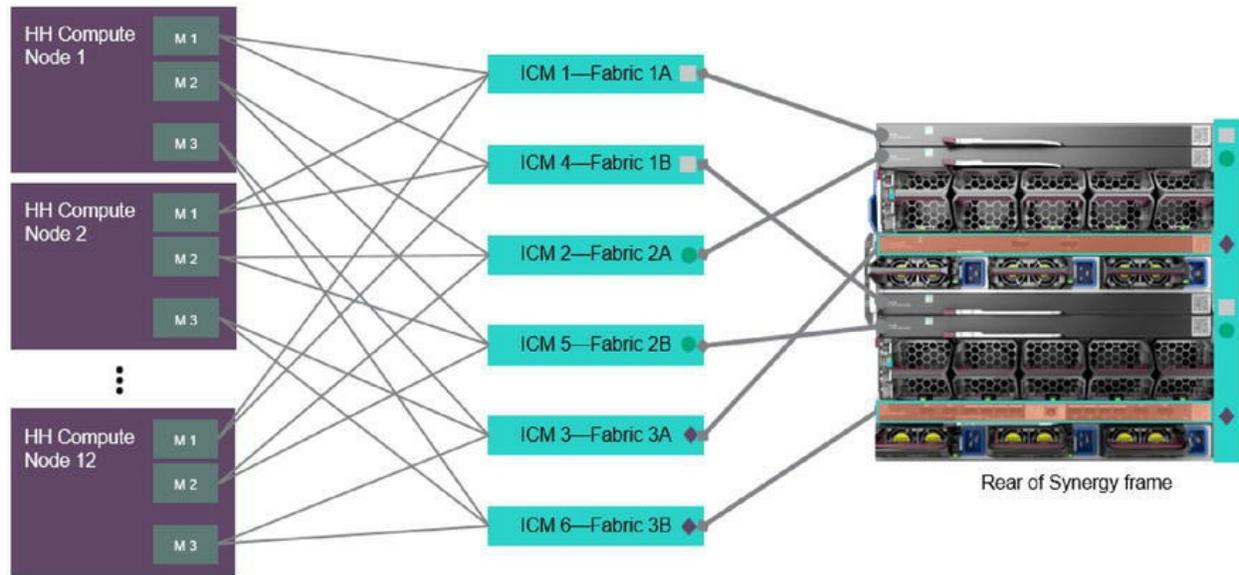


Figure 2-86 HPE Synergy frame interconnect module mapping

As indicated in [Figure 2-86](#), each compute module can be configured with one, or up to three, mezzanine card interface modules, occupying slots M1 to M3. The Synergy backplane supports two high throughput interface ports for each of the mezzanine card slots and they are connected to a pair of redundant interconnect module bays.

The interconnect bays are numbered one to six, from top to bottom. Fabric one is constructed from a redundant pair of interconnect modules 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 compute module portfolio



Figure 2-87 Synergy compute Gen10 and Gen9 modules

The flexible design Synergy composable compute module optimizes internal storage choices to match workload requirements. As shown in [Figure 2-87](#), 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 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. 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 data services and external-attached storage options.

The compute modules converge traffic over high-speed 10/20 Gb connections using HPE CNAs. When connected to an HPE Virtual Connect SE 40 Gb 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 SY480 Gen10 Compute Module



Figure 2-88 HPE Synergy SY480 front view

As referenced in [Figure 2-88](#), the HPE Synergy 480 Gen10 Compute Module delivers superior capacity, efficiency, and flexibility in a two-socket, half-height form factor to support demanding workloads. It also delivers superior capacity, efficiency, and flexibility to power more demanding workloads and increase 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 HPC in physical and virtualized environments while lowering costs within a Composable Infrastructure.

The HPE Synergy 480 Gen10 Compute Module has increased memory capacity (up to 1.5 TB per CPU socket) and 24 DIMM slots. It supports the entire Intel® new Intel® Xeon® Scalable Family of processors—3100, 4100, 5100, 6100, and 8100 series. Greater consolidation and efficiency are

achieved through an increase in VM density per compute module.

Key features of the Synergy SY480 as shown in the preceding graphic are:

1. Quick Access Panel
2. UID LED
3. Health Status LED
4. Mezzanine NIC status LED
5. Power On/Stand by button and system power LED
6. Compute Module handle release latch
7. Removable drive cage with two hot-plug drive bays
8. External USB 3.0 connectore & iLO USB connection (behind serial label pull tab)

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 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, VDI, SAP, and large storage capacity, such as Microsoft Exchange, and so forth.

HPE Synergy 660 Gen10 Compute Module



Figure 2-89 HPE Synergy SY660 front view

The HPE Synergy 660 Gen10 Compute Module delivers agility, control, and security in a foursocket, full-height form factor to support demanding workloads and virtualization density. It handles data-intensive workloads with uncompromised performance and exceptional value. It supports, highperformance, high-density Intel® 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 of the Synergy SY660 as shown in [Figure 2-89](#) are:

1. Quick Access Panel
2. UID LED
3. Health Status LED
4. Mezzanine NIC Status LED
5. Power On/Stand by Button and System Power LED
6. Zero (0) to Four (4) Hot-Plug Drive Bays (or up to 8 uFF with Smart Drive bay adapter)
7. Compute Module Handle Release Latch
8. External USB Connectore and iLO USB (behind serial label pull tab)

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 25/50 G 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
- 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

Synergy storage

Flexible capacity	High performance	Simplified management
Up to 200 drives No fixed ratios with compute SDS ready	12 GB SAS Full SSD performance Up to 2M total IOPS	Fluid resources Simplified lifecycle operations HPE OneView integrated
		



Figure 2-90 Fully integrated, performance capable, composable storage with Synergy

As illustrated in [Figure 2-90](#), Synergy composable storage's complete integration into the frame drives simplicity and reliability. Each storage module has 40 drives, and Synergy Composer can zone up to 200 storage drives to any compute module. 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 foursocket 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 is an end-to-end non-blocking 12 gigabit SAS fabric that delivers on the promise of full SSD performance with up to two million total IOPS. The resource pool is highly dense and yet fully flexible, offering up to 200 drives in a single frame that are composed with compute without any fixed ratios supporting file, block, and object data.

In addition to simplified deployment and provisioning as a composable resource deployed with templates and profiles, the full lifecycle of operations is simplified on this platform, from basic operational management to making storage highly portable and transformable as a managed resource within HPE OneView.

The 12 Gb SAS connection module, residing in interconnect module bays 1 and 4, creates a nonblocking 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 Synergy D3940 Storage Module

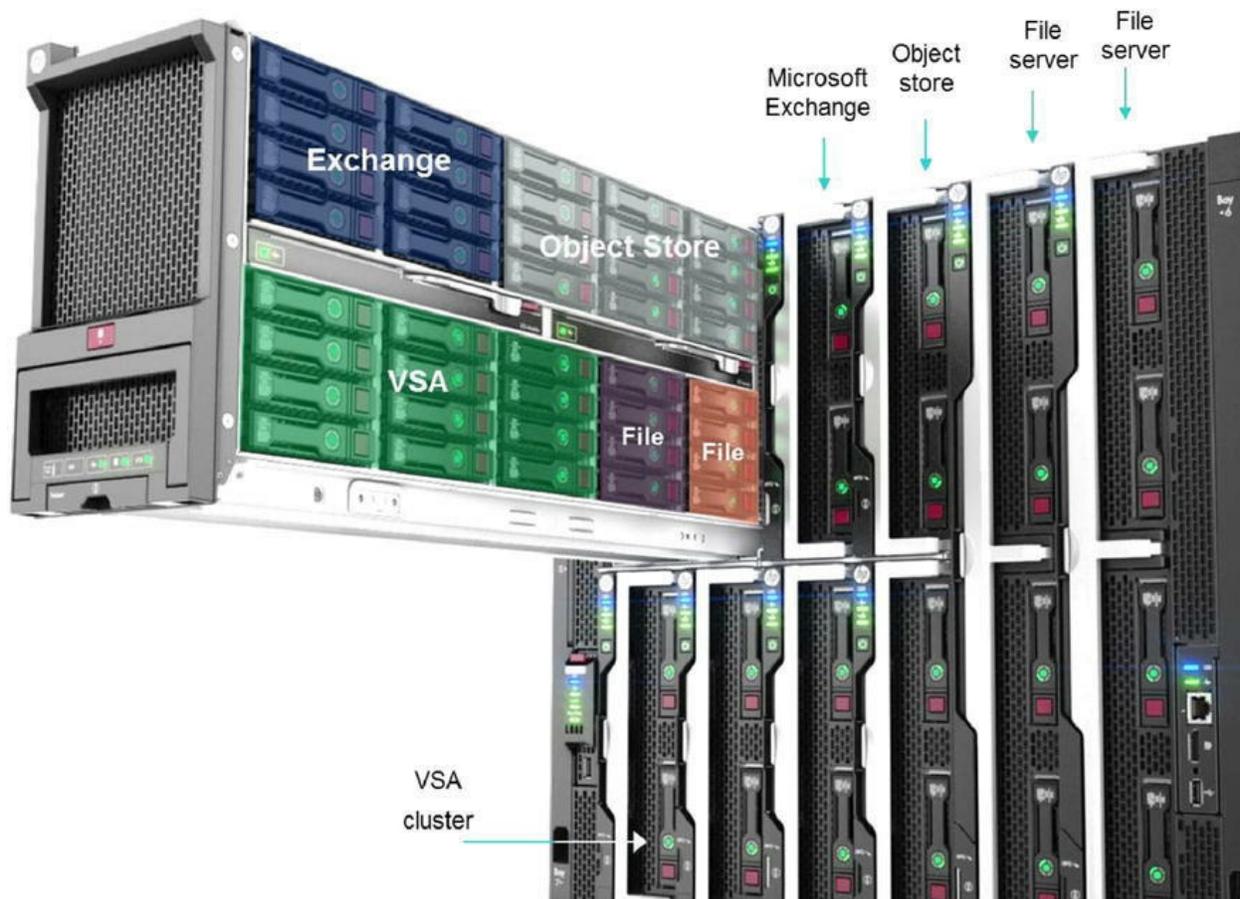


Figure 2-91 HPE Synergy D3940 Storage Module

As seen in [Figure 2-91](#), the HPE Synergy D3940 Storage Module is the foundational element of Synergy storage. Within a frame, an administrator can use 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 SDS solutions.

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 delivers an industry-leading density of 24 drives per rack U of space with 40 SFF 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:

- 4x controller bandwidth (96 Gb) compared with HPE BladeSystem c-Class (24 Gb)
- 2x drive enclosure bandwidth (192 Gb) compared with the HPE D3700 enclosure (96 Gb)
- HPE Dynamic Bandwidth Optimizer for 6 G SATA
- Synergy D3940 Storage Module as DAS = 2 M IOPs

HPE Composable Cloud for ProLiant DL



Figure 2-92 HPE Composable Cloud for ProLiant DL

As illustrated in [Figure 2-92](#), HPE Composable Cloud for ProLiant DL is an open private cloud platform enabled by composability that allows organizations to easily and efficiently deploy cloudnative applications at

rack-scale on virtual machines, containers, or bare-metal. The solution brings powerful composability to HPE ProLiant DL360 Gen10, DL380 Gen10, DL560 Gen10, and HPE SimpliVity 380 servers and delivers intent-based management through an enhanced HPE OneView enabling programmatic control of both the included HPE servers and HPE Composable Fabric.

HPE Composable Cloud for ProLiant DL is available as a complete solution or as stand-alone licenses for new and existing HPE servers that meet the solution support matrix requirements along with the HPE Composable Fabric. Other features include the following:

- It supports multiple, mixed use cases including HPE SimpliVity, VMware, Red Hat OpenShift, and bare-metal.
- Each HPE server sold with the solution includes HPE Composable Cloud for ProLiant DL software that enables leadership management and composability of the servers and fabric via enhanced HPE OneView features integrated with the HPE Composable Fabric.
- HPE Composable Fabric networking delivers dual 25 G links to each HPE server in the solution.
- It includes industry-leading HPE ProLiant DL360 Gen10, DL380 Gen10, and HPE SimpliVity 380 servers with broad configuration flexibility (CPUs, memory, and local storage).
- It scales from four up to 240 HPE ProLiant DL360, DL380 Gen10, DL560 Gen10, HPE SimpliVity 380, or HPE ProLiant DL380 Gen10 6126 2P 256 GB eight SFF Server for All Flash six VMware vSAN Certified ReadyNode servers per solution.
- When the full solution is purchased, a solution deployment is performed by HPE Pointnext for quality and faster time-to-value, with the flexibility to address specific data center needs.
- HPE Composable Cloud for ProLiant DL is now available with optional Fibre Channel (FC) storage connectivity. Customers can simply order a standard FC HBA with the HPE servers.
- HPE InfoSight for Servers may be used on all included HPE servers (no additional fee) to monitor, predict, and prevent problems across the server infrastructure.

- HPE Proactive Care support is included when purchased as full solution.
- It is available in select countries, including the United States, the United Kingdom, Ireland, France, Germany, and Australia.

HPE Composable Cloud for ProLiant DL delivers business value

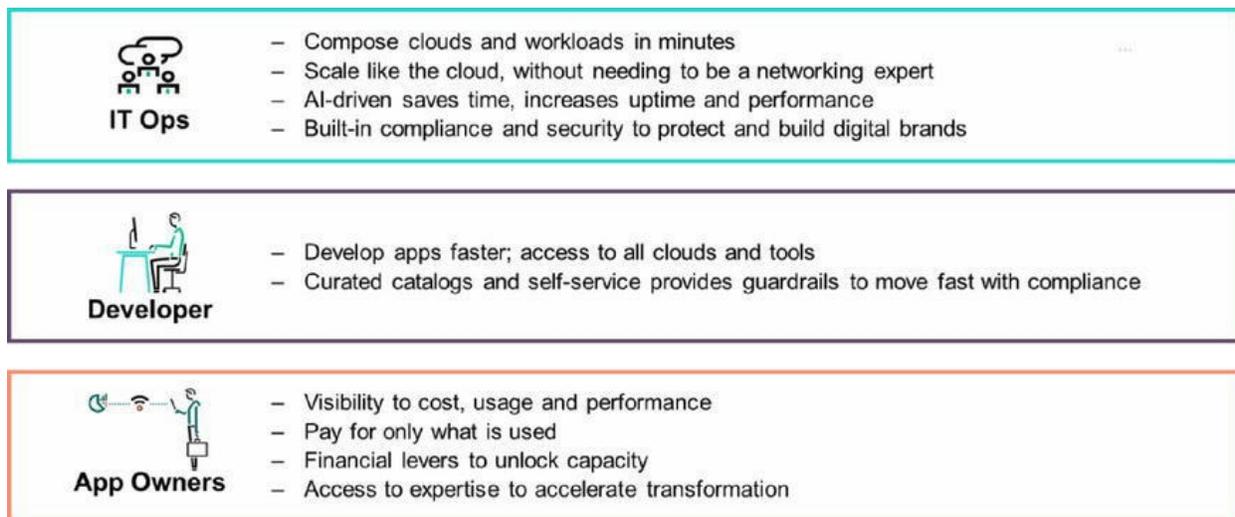


Figure 2-93 Business value of HPE Composable Cloud for ProLiant DL

HPE is extending the vision of composability to rack servers. As indicated in [Figure 2-93](#), with HPE Composable Cloud for ProLiant, IT professionals can:

- Scale the infrastructure in real-time without needing specialized networking skills.
- Compose clouds and workloads in minutes.
- Manage on-premises and public cloud resource utilization and costs in a uniform way.

The solution is optimized for cloud-native workloads and composable rack-scale environments with template-driven automation, which streamlines the deployment of cloud stacks and automates the lifecycle, initially supporting Red Hat OpenShift and VMware workload deployments.

HPE Composable Cloud for ProLiant DL is ideal for:

- Companies that have standardized on rack
- Cloud-native apps
- Storage flexibility (IP-based, SDS)
- Large enterprise deployments with up to 40 nodes

HPE ProLiant for Microsoft Azure Stack



Figure 2-94 HPE ProLiant for Microsoft Azure Stack

As seen in [Figure 2-94](#), the HPE ProLiant for Microsoft Azure Stack is a hybrid cloud solution that transforms on-premises data center resources into flexible hybrid cloud services that provide a simplified development, management, and security experience that is consistent with Azure public cloud services. The hybrid cloud solution is co-engineered by HPE and Microsoft to enable the easy movement and deployment of apps to meet security, compliance, cost, and performance needs. HPE ProLiant for Microsoft Azure Stack:

- Provides scaling from four nodes up to 16 hybrid nodes or up to eight All-Flash nodes per rack, with single-node scaling increments.
- Is based on industry-leading ProLiant servers with complete configuration flexibility (cores, memory, storage) and unmatched memory scale up.
- Is factory integrated for quality and faster time-to-value, with on-site deployment to address specific data center needs.
- Has a networking performance capability at 25 G.
- Uses HPE OneView with infrastructure automation engine for hardware monitoring and lifecycle management tasks, providing dashboards with status and real-time alerts for the health of the underlying infrastructure.

HPE ProLiant DL360 Gen10 Server as the Hardware Lifecycle Host (HLH) is used for solution deployment, monitoring, and management. The HLH is a Hyper-V host for the HPE OneView and HPE OneView Remote Support software management components.

While cloud architectures are growing fast with enterprises, not all workloads are suitable for public cloud. Data sovereignty, privacy, IP, compliance, performance, and cost make public cloud infeasible or impractical for many organizations and applications. Unfortunately, this has not stopped many IT users from seeking out public cloud services and putting their companies at risk. Private and hybrid cloud solutions can bridge this gap, but designing, deploying, and operating them can be complicated, risky, and expensive. If the new cloud does not meet the needs of the users, IT can find itself supporting something that no one uses, in addition to what they already

support.

Enter HPE ProLiant for Microsoft Azure Stack. HPE and Microsoft are working together to deliver a hybrid cloud solution that delivers on the promised speed and agility benefits of public cloud, in a package that can be quickly and easily deployed in enterprise or service provider data centers. This allows organizations to reap the benefits associated with cloud operating models, delivered in their own data centers, by two of the biggest names in the industry.

The HPE ProLiant for Microsoft Azure Stack allows organizations to run Azure-consistent cloud services in their data center. The key features of this solution include:

- **Common developer experience**—Applications can be deployed to either an Azure Stack private cloud or an Azure public cloud without modification. This is a powerful benefit to developers that makes it easier and faster to develop applications.
- **Workload portability**—Since Azure public cloud and Azure Stack services are API-compatible, applications can be deployed in Azure Stack and in Azure public cloud.
- **Integrated systems management**—Managing and maintaining a rack of servers, storage, and switching can be complex. HPE ProLiant for Microsoft Azure Stack includes integrated system management that makes it easier to perform software and firmware updates as well as monitor system status. HPE includes HPE OneView and HPE Insight Remote support running on a separate management server.
- **Single vendor support**—Solving problems for services running on multiple servers, network switches, and software products can be complex and frustrating. HPE offers single vendor support, so regardless of what component causes a problem, HPE manages the resolution with no finger pointing.
- **Multiple configuration options**—We know your customer's workload requirements are dependent on their specific needs. That is why we offer more configuration options. You can choose the number of nodes (from four to 12), the processor type, memory, and storage. We support larger memory configurations, up to 768 GB/node, than most competing

systems.

- **Multiple packaging options**—For most of our customers, we build the systems in our factory and ship them to the customer site. However, some of our customers would rather use their own racks. In those cases, we can build the systems **in our factories**, using your customer’s racks, or we can build the systems **in their data** center using their rack. The choice is theirs.
- **Multiple financing options**—We provide you with multiple ways to purchase HPE ProLiant for Microsoft Azure Stack. We have pay-as-you-go, leasing, and capital purchase options.
- **Multiple sourcing options**—We provide your customers with multiple ways to source the solution. They can purchase the entire solution from HPE, hardware and software; you can purchase the hardware from us and the software from Microsoft; or HPE can be an indirect Cloud Solution Provider (CSP) or distributor. Whatever fits your needs the best.

HPE ProLiant for Microsoft Azure Stack delivers the largest capacity and highest performance, at full speed, giving the ability to run even more workloads, faster.

Component	Default	Options
Environmentals	Cabinet, cables, POU, factory build	
Server	4 HPE ProLiant DL380 Gen10	4 to 16 nodes 4 to 8 DL380 Gen10 All-Flash
Processor ¹	Intel Skylake CPUs	10 to 28 cores 2.0 to 3.0 GHz
Memory ²	24 DIMM Slots 384 GB minimum	384 GB, 576 GB, 768 GB, 1.5 TB
Storage ³	12 HDD per server 6 SSD per server	4 TB, 6 TB, 8 TB, or 10 TB 800 GB, 1.6 TB, or 3.2 TB All flash: 800 GB, 1.6 TB, 3.2 TB, or 6.4 TB (total capacity up to 102 TB)
Network	25 GbE SFP28 1 GbE RJ45 (OOB Management)	FlexFabric 5950 48SFP28 FlexFabric 5900 AF-48G 4XG-2QSFP+ Arista 7160 / 7020RA Cisco Nexus 9300 / 3048(Roadmap)
Power	110V and 220V HPE Power Advisor	
OOB server	HPE ProLiant DL360 Gen10 8SFF CTO Server	Higher performance HLH

Figure 2-95 HPE ProLiant for Microsoft Azure Stack hardware overview

Figure 2-95 shows the defaults and options for the HPE ProLiant for Microsoft Azure Stack components.



Note

¹ Processors must be identical pairs. New generation of processors under investigation.

² Memory must follow population and configuration rules.

³ Storage must follow configuration rules.

Configurable to meet workload requirements

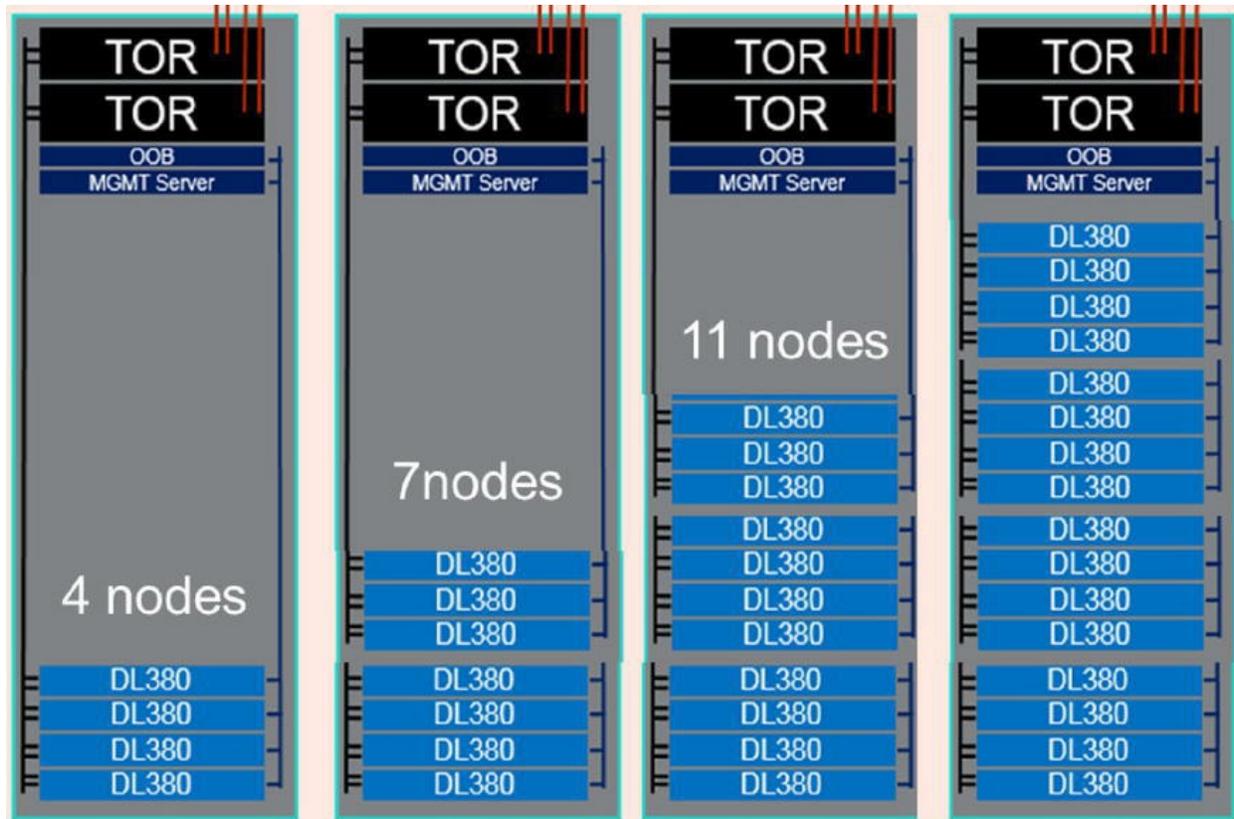


Figure 2-96 Any quantity between four and 16 nodes

As indicated in 2-97, HPE ProLiant for Microsoft Azure Stack is highly configurable—processor, memory, storage, single-node sizing, and rack/virtual rack.

Example entry configuration:

- Four HPE ProLiant DL380 Servers
- Dual Skylake 2.6 Hz processor, total 24 cores/server
- 384 GB RAM memory/server

Example maximum configuration:

- 16 HPE ProLiant DL380 Servers
- Dual Skylake 2.5 GHz processor, total 56 cores/server
- 1.5 TB memory/serve

Example configurations

Cores	Memory	Processor	Storage	RAW vs. usable capacity	VMs
4	952	Xeon 6126 12 c, 2.6 GHZ	72 TB	288 TB / ~96 TB	272 (A)
8	2488	Xeon 6138 20 c, 2.0 GHZ	7962 TB	768 TB / ~256 TB	710 (A)
12	4024	Xeon 8160 24 c, 2.1 GHZ	120 TB	1440 TB / ~480 TB	1150 (A)
16	22840	Xeon 8180 28 c, 2.5 GHZ	120 TB	1920 TB / ~640 TB	6525 (B)

(A) Based 384 GB RAM per node less HA and system overhead with standard D1 VMs supported
(B) Based 1.5 TB RAM per node less HA and system overhead with standard D1 VMs supported

Figure 2-97 Example configurations



Note

Estimate only. Use Microsoft Capacity Planning tool for final configuration.

Learning check

8. Recommend a single-socket server providing up to 32 cores and 4 TB in a 1U rack form.

9. What are the typical components of an HPE ProLiant for Microsoft Azure Stack solution?

Configuring a single server solution

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 Simple

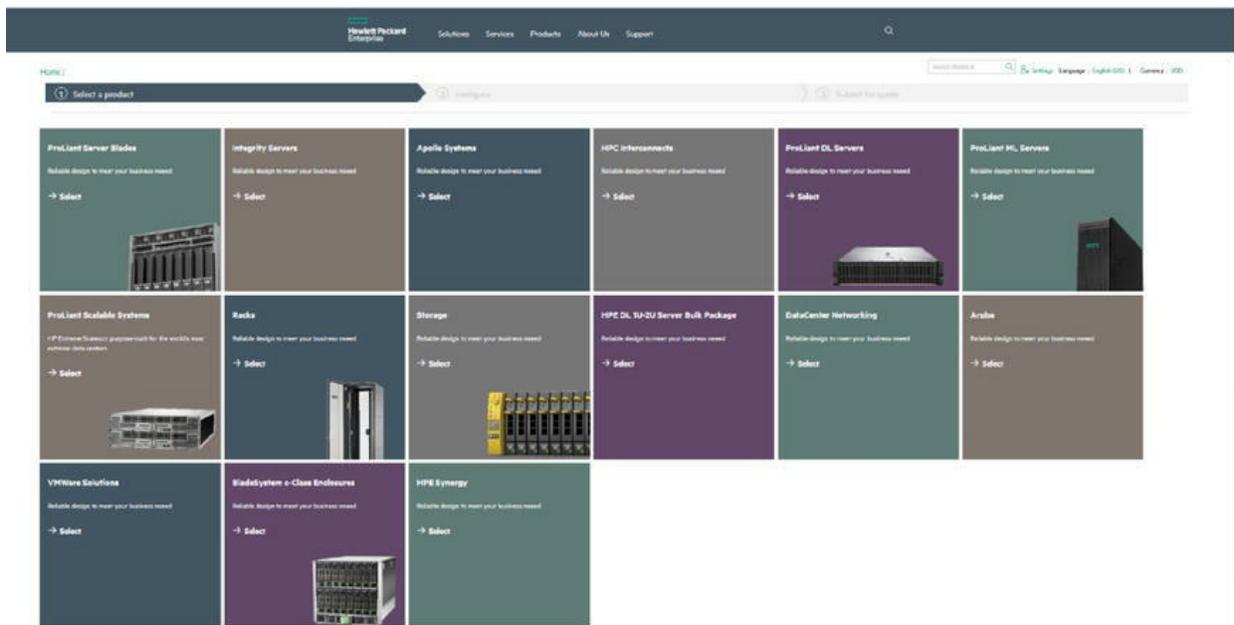


Figure 2-98 HPE One Config Simple

As illustrated in [Figure 2-98](#), the HPE One Config Simple (OCS) online configuration tool actively gathers customer requirements and maps them to a set of products or service options. The guided selling capability directs users to an optimal solution for the customer’s specific workloads or application

needs. This capability makes OCS ideal for non-technical salespeople, allowing them to quickly create simple configure to order (CTO) and build to Order (BTO) solutions. Final configuration files can be imported into partner quote and ordering systems for fast conversion to an order. It is most appropriate to use OCS early in the sales cycle for budgetary quotes and “what if” scenarios.

OCS is a web-based application that:

- Provides a self-service environment.
- Generates initial solution configurations in three to five minutes.
- Offers local list pricing.
- Uses solution wizards based on applications.
- Lists popular configurations that can be modified.
- Can export configurations to Microsoft Excel file format.

OCS includes a simple mode for creating configuration files. The user selects an HPE product family (ProLiant servers, storage, networking, and so on) and then completes a series of steps to add or confirm components (for example, disk, memory, and services).



Note

The HPE One Config Simple tool can be accessed at:
<https://h22174.www2.hpe.com/SimplifiedConfig/Welcome>

HPE OneConfig Advanced

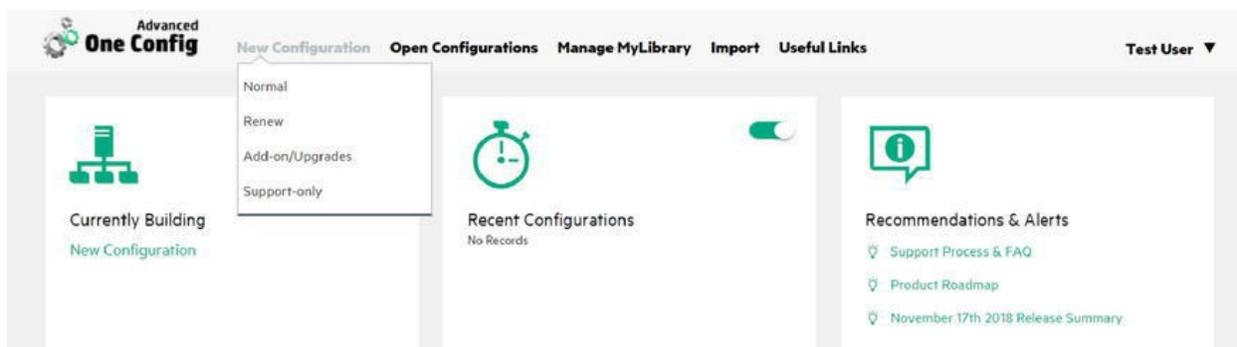


Figure 2-99 HPE OneConfig Advanced

HPE OneConfig Advanced is available both internally and for HPE Partners. The tool can be accessed from HPE Partner Ready portal (<https://partner.hpe.com/>) and is efficient in handling large configurations of multiples of rack/servers.

HPE OneConfig Advanced:

- Uses intuitive and web-based interface, as seen in 2-100.
- Eliminates redundant steps to recreate the configuration in other tools (such as Power Advisor).
- Includes latest products such as Nimble and SimpliVity.

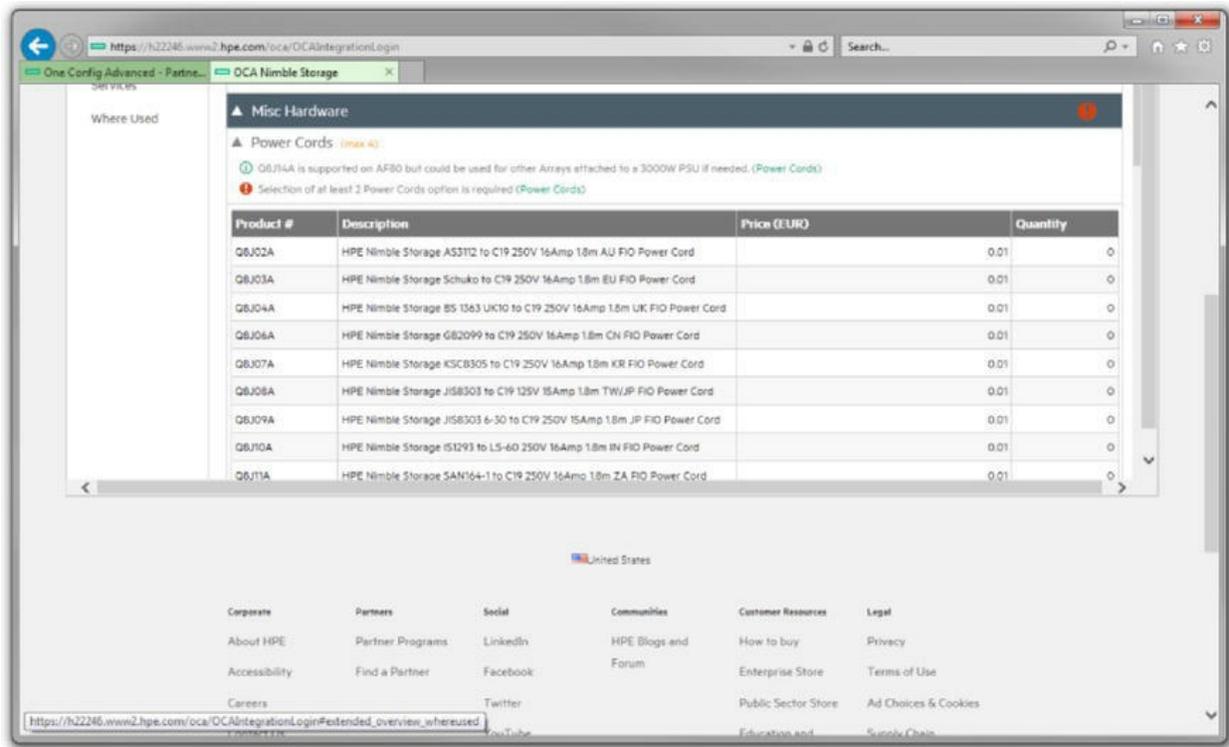


Figure 2-100 HPE OneConfig Advanced

Advantages of OCA:

- Allows a user to run multiple concurrent sessions.
- Reliable performance supported by Akamai accelerators around the

globe.

- Minor configuration edits can be made without having to start over.
- Ability to build configuration starting with any product (rack, server, blade, enclosure, switch, storage) and integrate them as required.
- Automatic updates to Knowledge Base, which allows for real-time issue resolution rather than relying on large weekly downloads.
- “Where Used” function finds instances where a Unique Configuration ID (UCID) or stock keeping unit (SKU) was used.
- Visibility to recommended SKUs enables business units (BUs) to steer demand and meet delivery commitments.
- Ability to configure on the go on a tablet (using supported browsers Chrome, Firefox, and IE).
- Configurations automatically save as you work.
- Context-sensitive messaging throughout the application.
- View into all saved configurations (not just orders) allows for detailed SKU demand analysis (by geography, by customer, and so forth).
- Comprehensive services view and edit capabilities for complete configuration or individual products therein.
- Reverse engineers a configuration solution from an imported bill of materials (BOM) or SKU list.
- Prevents unbuildable configurations by using “CLIC Check” functionality.
- Provides configuration wizards for complex products, ensuring accuracy of products and services.
- Power consumption report available for most configured products.
- Reduces factory cycle times for select complex solutions by sending Config-to-Build (C2B) instructions when needed.
- Eliminates memory and storage requirements for user’s machine.
- Eliminates timeouts and enables new business models like cloud services (for example: CS 500 appliances).

- Quickly scales solutions, from one rack to 55 racks in 3.5 minutes or less.
- Share configurations within OCA directly with a distributor.
- Share and reuse templates with “My Library”.
- Set up work groups to share configurations easily.
- Quote all existing HPE configuration types (Watson, Smart Quote, IQ on BMI, SalesBUILDER for Windows [SBW]).

Navigating the OCA user interface

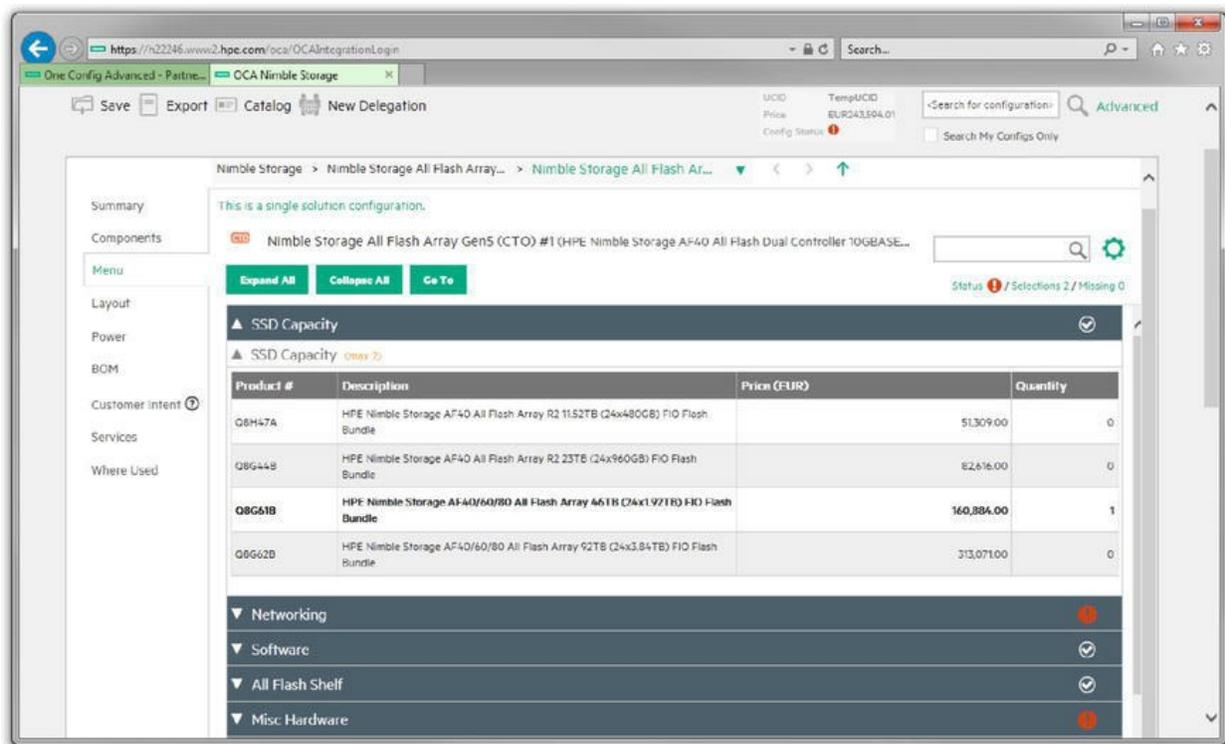


Figure 2-101 Navigating the user interface

Navigating the user interface, as seen in [Figure 2-101](#):

- **Home page**—Click the icon to go to OCA home page.
- **User settings**—Click the login name to access the user setting page, to manage the workgroup, and delegate your config.
- **Simple Search/Advanced Search**—Search section to allow quick search

and advanced searches.

- **New Configuration**—Tab allows you to create a new config. You can also click the New Configuration link to initiate the process.
- **Open Configurations**—Allows you to open local config file saved as “.oca” format and view recent configs that are also displayed in the Recent Configurations section.
- **MyLibrary**—Tab allows you to view and manage saved configurations, user-defined SKUs, and customer intent documents.
- **Currently Building**— OCA remembers your last configuration you saved and offers it in the “Currently Building” section.
- **Import**—Tab enables you to upload configs in a certain file format.
- **Useful Links**—Tab provides links to useful configuration-related information.
- **Recommendations & Alerts**—Tab provides links to trainings, support, and knowledge base roadmap.

HPE Power Advisor

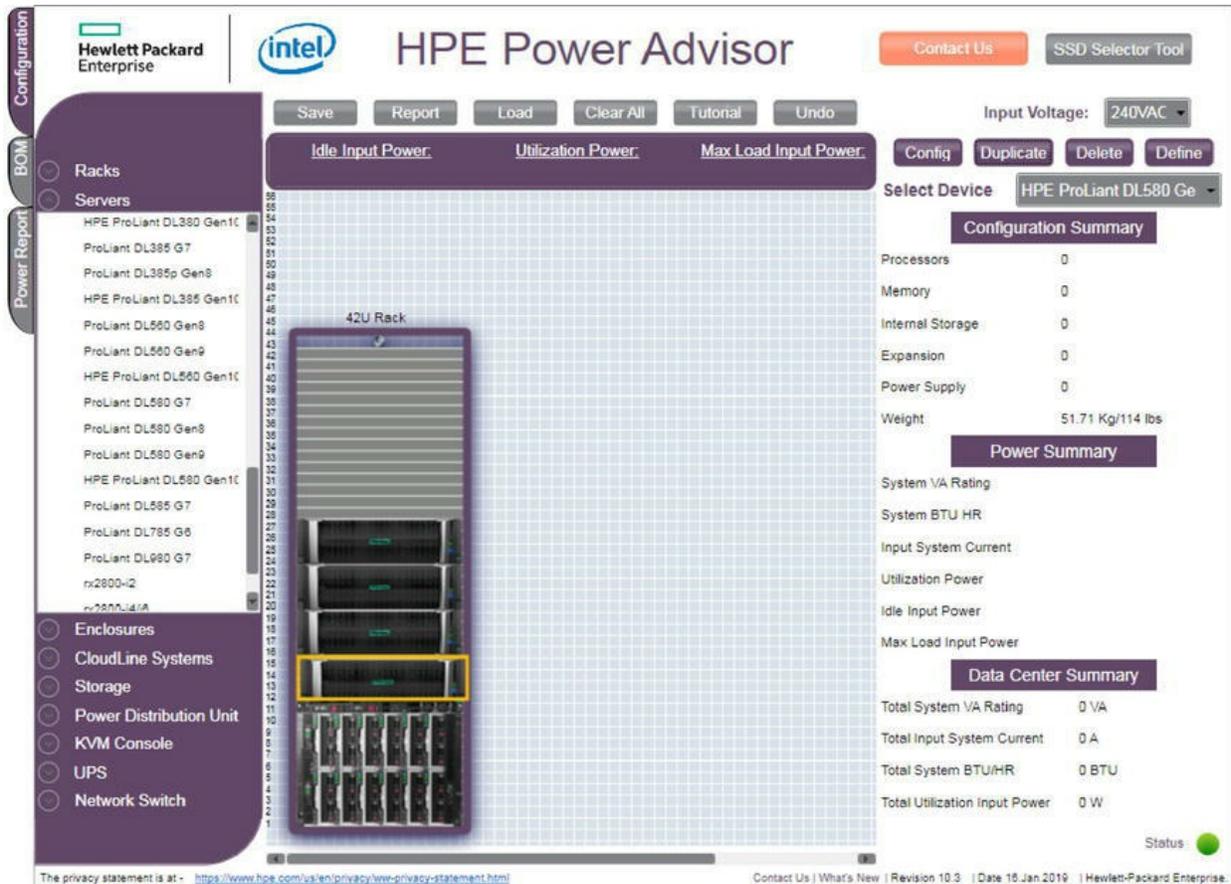


Figure 2-102 HPE Power Advisor

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.

As referenced in [Figure 2-102](#), HPE Power Advisor is an easy-to-use tool that estimates data center power requirements for server and storage configurations. A downloadable version and online Windows application are available. 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

HPE Power Advisor can be accessed at:
<https://paonline56.itcs.hpe.com>

HPE Server Memory Configurator

HPE DDR4 Server Memory Configurator

1 Select server > 2 ID existing memory > 3 Select memory capacity > 4 Select memory option > 5 Create parts list

Step 1a: Do you have a pre-configured BTO server model part number? No Yes

Step 1b: Select HPE ProLiant Server Series



Figure 2-103 HPE Server Memory Configurator

As shown in [Figure 2-103](#), the HPE Server Memory Configurator is a web-based tool used to assist with populating DDR3/DDR4 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:

1. Select the relevant ProLiant server.
2. Identify existing memory (auto-detect memory with HPE Insight Diagnostics or manually enter current memory).
3. Select memory capacity.
4. Select memory option.
5. Create parts list.



Note

HPE Server Memory Configurator can be accessed at:
<https://memoryconfigurator.hpe.com>

SSD Selector Tool

HPE Solid-State Drive Selector Tool

Workload | SSD type | Server type | Drive capacity | Interface type | Form factor | Category | Results

Select Your Workload(s) from the List Below
Check All Workloads That Apply
→ To Learn More About Workloads Click Here

<input type="checkbox"/> Boot/Swap <input type="checkbox"/> Read Caching <input type="checkbox"/> Web Servers <input type="checkbox"/> Social Media	<input type="checkbox"/> Virtualization (Medium Density) <input type="checkbox"/> Database (Mid Range) <input type="checkbox"/> Scientific and Engineering (Medium) <input type="checkbox"/> General Business Applications	<input type="checkbox"/> Financial Computing <input type="checkbox"/> OLTP <input type="checkbox"/> Mission Critical Applications <input type="checkbox"/> Business Intelligence
--	---	---

Note: The workloads in this column are typically addressed with lower priced SSDs, with a focus on READ performance. You will have the opportunity to choose READ Intensive as a "SSD Type" on the next screen.

Note: The workloads in this column typically require a balance of price and READ/WRITE performance. You will have the opportunity to choose Mixed Use as a "SSD Type" on the next screen.

Note: The workloads in this column are typically addressed with SSDs with high WRITE performance. You will have the opportunity to choose WRITE Intensive as a "SSD Type" on the next screen.

Select All / I Don't Know Yet

Back Next

Figure 2-104 SSD Selector Tool

As shown in [Figure 2-104](#), the HPE Solid-State Drive Selector Tool is an

online tool that helps 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 a fine-grained selection process.

Manual selection criteria include:

- Type of the workload
- SSD type (read-intensive, mixed use, write-intensive)
- Type of the server and the server model
- Capacity
- SSD Interface (SAS, SATA, NVMe)
- Form factor (SFF, LFF, M.2, Add-In Card, ...)
- Category (recommended vs. extended)



Note

HPE SSD Selector Tool can be accessed at: <https://ssd.hpe.com>

SSD Availability Matrix

Servers Storage

SATA SAS NVMe M.2 2019-03-19

Read Intensive (RI) Mixed Use (MU) Write Intensive (WI)

SFF/2.5" (Smart Carrier)

Part Number	GB	Product Description	Recommended	Life Cycle Status	Large Volume Quantity	Current Availability	Estimated Lead Time for Future Orders
875503-B21	240	HPE 240GB SATA RI SFF SC DS SSD	✓	SUST	25	Green	General availability
868814-B21	240	HPE 240GB SATA 6G RI SFF SC DS SSD		EOL	25	Yellow	Extended lead time
875511-B21	960	HPE 960GB SATA RI SFF SC DS SSD		DISC	25	Green	General availability
868822-B21	960	HPE 960GB SATA 6G RI SFF SC DS SSD		EOL	25	Yellow	Extended lead time
868826-B21	1920	HPE 1.92TB SATA 6G RI SFF SC DS SSD		EOL	15	Green	General availability
875507-B21	240	HPE 240GB SATA RI SFF RW DS SSD		OBS	25	Yellow	Extended lead time

Figure 2-105 SSD Availability Matrix

In addition to servers and storage, now more products industry-wide use NAND flash technology, while SSDs progressively decrease in availability.

Consequently, the suppliers that had been able to meet our demand are now no longer able to fulfill vendor needs. To prevent long lead times, steer your customers away from certain SKUs to ones with better availability. SSD Availability Matrix shows the latest SSD product availability.

As indicated in [Figure 2-105](#), the recommended options depicted in the chart have a more robust stocking strategy, which means orders containing these options will have more predictable fulfillment times.

Availability status for each SSD part number:

- **Green**—Supply available
- **Yellow**—Constrained supply with extended lead times
- **Red**—Little or no supply. Incoming SSDs will be allocated to existing orders



Note

The SSD Availability Matrix can be accessed from the partner ready portal at: <https://partner.hpe.com>

Learning check

10. Name at least three criteria that the SSD Selector Tool can use to filter down to the correct SSD drive for your customer.

11. You are in a meeting with a customer, and they challenge you with a statement: “I heard that the configuration process required to prepare a server configuration quote is time consuming, with a high risk of errors and risk of ordering incompatible components.”

How should you respond?

Selecting HPE Support Services

HPE Pointnext, the services organization for HPE, helps customers optimize IT and build the ideal technology, partner, and operational foundations needed to accelerate the digital journey of the business.

Building on our heritage of services leadership, HPE Pointnext:

- Builds for speed to help accelerating innovation and time to value.
- Curates and aggregates a best-in-class partner ecosystem to deliver the best solutions for customers.
- Removes complexity across the IT lifecycle.
- Optimizes infrastructure for current and new generation of apps and data.

HPE Operational Support Services



Figure 2-106 Simplify and improve your customer’s services experience

As referenced in [Figure 2-106](#), HPE Operational Support Services gives customers the support experience they need. Customers can choose the level of support that meets their IT and business needs.

HPE Foundation Care is cost effective, simplified support that helps get devices running when there is a problem. This is reactive support. Choose from hardware response times and save time with one place to call.

Foundation Care includes software support as well as collaborative support with Independent Software Vendors (ISVs). If we find that an issue is with the software and if it is not a known issue, we will transfer the call to the vendor, including Microsoft, SUSE, RedHat, VMware, and more.

HPE Proactive Care and Proactive Care Advanced (PCA) includes everything that Foundation Care offers and gives more personalized, proactive support with an enhanced call experience and tailored reports with recommendations. In addition, PCA includes assigned technical experts for a more personalized experience.

Finally, **Datacenter Care** offers a tailored experience to drive operational efficiency for the entire IT environment—even multivendor IT environments. Datacenter Care builds on what Foundation Care and Proactive Care offers, but offers solution support.

Whatever the support experience customers need, HPE can help.

One partner to help business succeed through IT

Single point of contact for service-related tasks offers many advantages:

- Save time
 - One place to call
 - Collaborative support with ISVs
- Cost effective
 - Choose the response that meets IT and business needs
 - Three-, four-, or five-year support duration
- Keeps devices running
 - Help when there is a problem
 - Access to technical resources
- Reduces complexity
 - Hardware and/or software support
 - Exchange service for select products

- Gain control
 - Connect to HPE
 - Access to IT anytime, anywhere from any device
 - 24x7 monitoring, automatic call logging
- Meet service-level expectations
 - Choose coverage to help meet IT and business needs
 - Call to Repair, 24x7, Next Business Day

HPE Foundation Care

All products come with a warranty, and it is important that customers understand exactly what the warranty gives. Warranty is an attribute of the product and is available 9x5, business hours only, not on weekends or holidays. Parts are shipped to customers and the warranty duration may vary from three months to three years with different components in an array (for example) having different warranty durations. There is no service-level response commitment and no software support.

HPE Foundation Care helps with saving time, with simplified support when there is a problem. We help customers get devices up and running. Customers will have access to our Standard Support Center via phone, chat, email, and the internet to help log cases. HPE Foundation Care extends the duration, response, and on-site features of the product warranty. In addition, your customer can choose to receive software support with a two-hour critical response window.

We recommend connecting products to HPE for near 100% diagnostic accuracy, and one consolidated view of your customer's connected HPE devices. They will receive 24x7 monitoring, near 100% diagnostic accuracy, prefailure alerts, automatic call logging, and parts dispatch. Connected devices will give more control and information to help manage IT.

All of this will help your customers save time with simplified support with one partner to call and give more control over their IT.

Hardware support includes:

- Support for HPE and non-HPE hardware
- On-site hardware support
- Exchange service for select devices
- Access to firmware updates
- Replacement parts
- Collaborative software support for x86 servers

Software support includes:

- Access to technical resources
- Software operational support
- Software documentation updates
- Where applicable, access to licenses and updates

Choose the response level that meets your IT and business needs

Choose from a set of reactive support levels to meet business and operational needs. Get connected to save time and resolve problems faster to ensure the level of coverage customers need to support the IT infrastructure with HPE Foundation Care services.

For server, storage, and networking products covered by Foundation Care, HPE offers three distinct service levels:

- **HPE Foundation Care Call to Repair Service**—Foundation Care Call to Repair Service offers 24 hours service per day, seven days per week, including HPE holidays with a six-hour call-to-repair time where our commitment is to have the hardware operational within six hours after a call is opened.
- **HPE Foundation Care 24x7 Service**— Foundation Care 24x7 Service offers 24 hours service per day, seven days per week, including HPE holidays with a four-hour on-site response time for hardware and a two-hour response time for software.

- **HPE Foundation Care Next Business Day Service**—Foundation Care Next Business Day Service offers a next business day on-site response with coverage on standard business hours, standard business days. Service is available nine hours per day between 8:00 a.m. and 5:00 p.m. local time, Monday through Friday, excluding HPE holidays.

For selected products, such as HPE Networking products, we provide two additional levels of exchange service:

- **HPE Foundation Care 4-Hour Exchange Service**— Foundation Care 4-Hour Exchange Service is available 24 hours per day, seven days a week, including HPE holidays. The networking component will be exchanged within four hours after opening the case, and networking software support is also included. This service is available for select HPE Networking products and is not available on all products.
- **HPE Foundation Care Next Business Day Exchange Service**— Foundation Care Next Business Day Exchange Service provides exchange of the networking component the next business day after the call is opened and includes networking software support. This service is available for select HPE Networking products and is not available on all products.

HPE Proactive Care



Figure 2-107 Get ahead of IT problems

Building on Foundation Care, we understand that it is not only important to fix problems when they occur, but it is also important to help reduce outages and prevent problems before they occur. As illustrated in [Figure 2-107](#), we can help customers to maximize performance and get the most out of their HPE devices with access to technical skills and best practice advice, receive a rapid resolution when a device is down, and connected to HPE for 24x7 real-time monitoring to help stay in control, to be informed, and to manage the IT.

Stay in control: Connect devices to HPE

These are the three key elements of HPE Proactive Care service:

- Solve problems faster
- Prevent problems
- Stay informed and in control

We help prevent problems before they occur using data analysis with a proactive scan, firmware/patch analysis, and incident/trend reports. Customers receive tailored reports with data analysis, recommendations, and advice from a Technical Account Manager.

We want to help with solving problems faster—using 24x7 monitoring, that will rapidly and accurately diagnose a problem that may be occurring, automatically dispatch parts, and if needed, to call in, to have an enhanced call experience, working with a Technical Solution Specialist who will manage the case from start to finish, engaging with any specialists as needed.

It is critical that customers stay informed and in control with access anytime, anywhere, and from any device to a personalized dashboard with all of the current information for monitored devices. HPE provides the access to a global knowledgebase and a community of experts.

Each of these elements comes together as HPE Proactive Care service.

Helping manage IT



Figure 2-108 Prevent problems and reduce outages with tailored reports

Connecting products to HPE is key for the proactive deliverables included in Proactive Care Support. HPE's goal is to help prevent problems and downtime that is costly and to help customers focus on the projects and initiatives that are important to business. A Technical Account Manager reviews tailored reports with data analysis and makes recommendations.

As indicated in [Figure 2-108](#), helping manage IT with:

- **Proactive scan**—Twice yearly proactive scan with recommendations and advice
- **Firmware release & software patch report**—Twice yearly firmware, patch, and software revision assessment and update recommendations
- **Incident report**—Quarterly incident trend reports to help identify configuration or environment issues



Note

Devices need to be connected to HPE.

Resolve problems with a rapid response

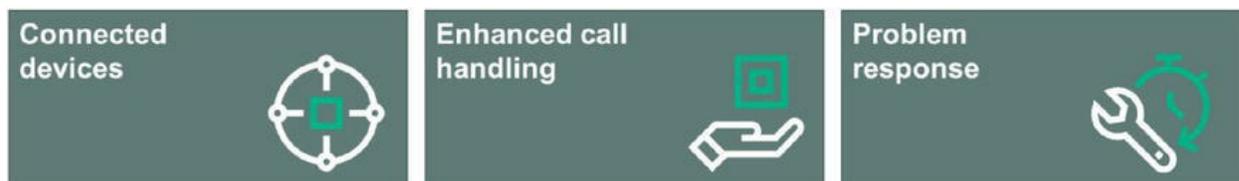


Figure 2-109 Resolve problems with a rapid response

Reactive support is part of Proactive Care support. Connecting products to

HPE will provide automatic call logging, generating a case and an email notification. Part of this is accurate and quick diagnostics and automatic parts dispatch.

If customers do need to call HPE for an incident, then they receive an enhanced call experience working with a Technical Solution Specialist (a TSS) who will manage/own the case from start to finish. Three levels of on-site support are available from which your customers can choose.

If customers have their own software support through the manufacturer, then we will work to solve basic, known issues, and if necessary open a call with the ISV, as shown in [Figure 2-109](#):

- Solve problems faster with enhanced call handling:
 - Start-to-finish case management
 - 24x7 monitoring
 - Prefailure alerts
 - Fast, accurate problem detection
 - Automated call logging and parts dispatch
- Prevent problems with tailored reports, expert reviews, and advice:
 - Proactive system scans
 - Firmware/patch revision management and analysis
 - Incident/Trend analysis
- Stay informed and in control when connecting devices to HPE for:
 - Personalized dashboard
 - Global knowledgebase
 - Community of experts (HPE, partners, peers)

Save time with one place to call

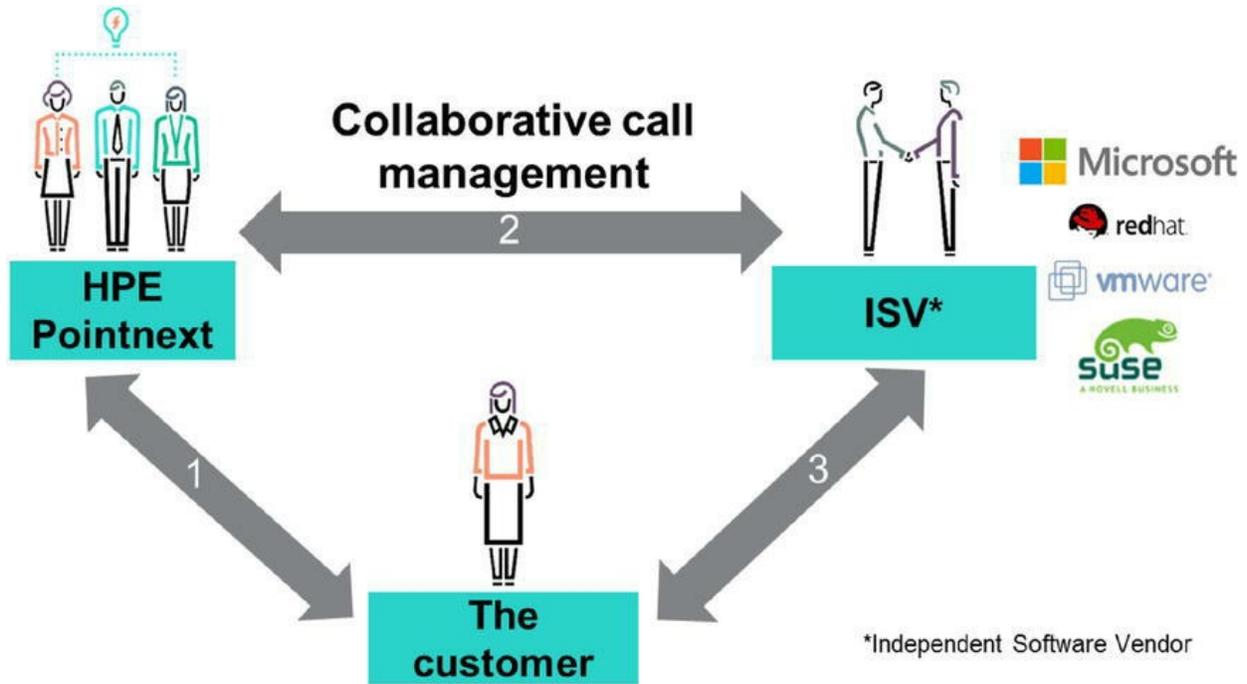


Figure 2-110 Collaboration with independent software vendors

As seen in [Figure 2-110](#), Proactive Care provides a collaborative customer support system. If the customer has software support with the ISV, we will diagnose the problem when the customer calls us. If it is not a known issue, then we will open a call with their ISV, transfer the case history, and the customer only has to place to one call and that is to HPE. This saves the time and gets the issue resolved faster.

Learning check

12. Name the available response levels with HPE Foundation Care services.

Managing a single server solution

HPE ProLiant management innovations target three segments to ensure that customers have complete lifecycle management—for their current environment and in the future as their business grows.

HPE server management

- **Managing single system (on-system)**—Built-in intelligence and automation for increased server admin productivity. It provides on-system management to provision, monitor, and troubleshoot servers as well as remote and out-of-band management. The on-system management tools available for ProLiant servers include:
 - Intelligence on every HPE server enabling setup, health and alerting, and firmware maintenance
 - UEFI
 - iLO 5
 - RESTful API, HPE RESTful Interface Tool, and other HPE scripting tools
 - Intelligent Provisioning
 - Smart Storage Administrator
 - HPE Smart Update and Service Pack for ProLiant
- **Managing multiple systems (on-premise)**—Installed at customer site and includes:
 - HPE OneView
 - HPE OneView Global Dashboard
 - iLO Amplifier Pack
- **Managing multiple systems (on cloud)**—Cloud-based management for service ticket and warranty tracking, health, and alerting, and includes:
 - Remote Support and Insight Online
 - InfoSight

HPE iLO 5 management technologies

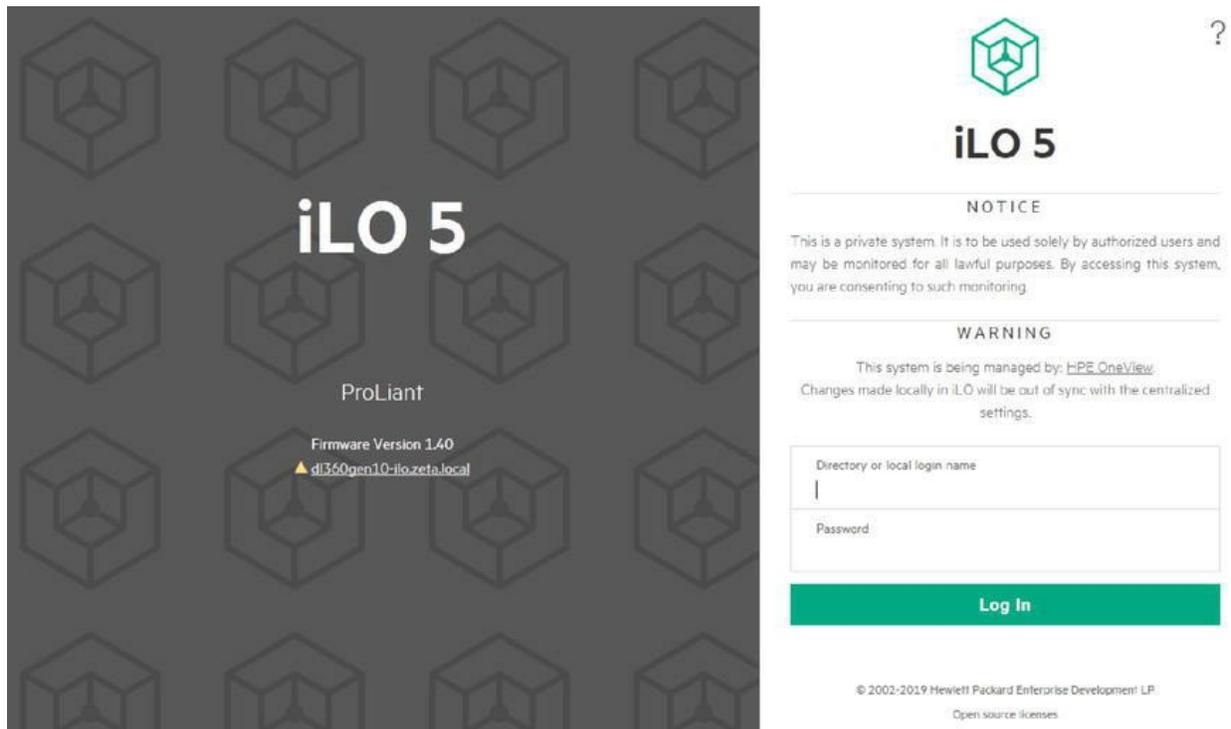


Figure 2-111 HPE iLO5

As shown in [Figure 2-111](#), 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 your customer 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. Your customer 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 Common Access Card (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 can be found at: https://support.hpe.com/hpsc/doc/public/display?docId=a00018324en_us

Using the iLO web interface

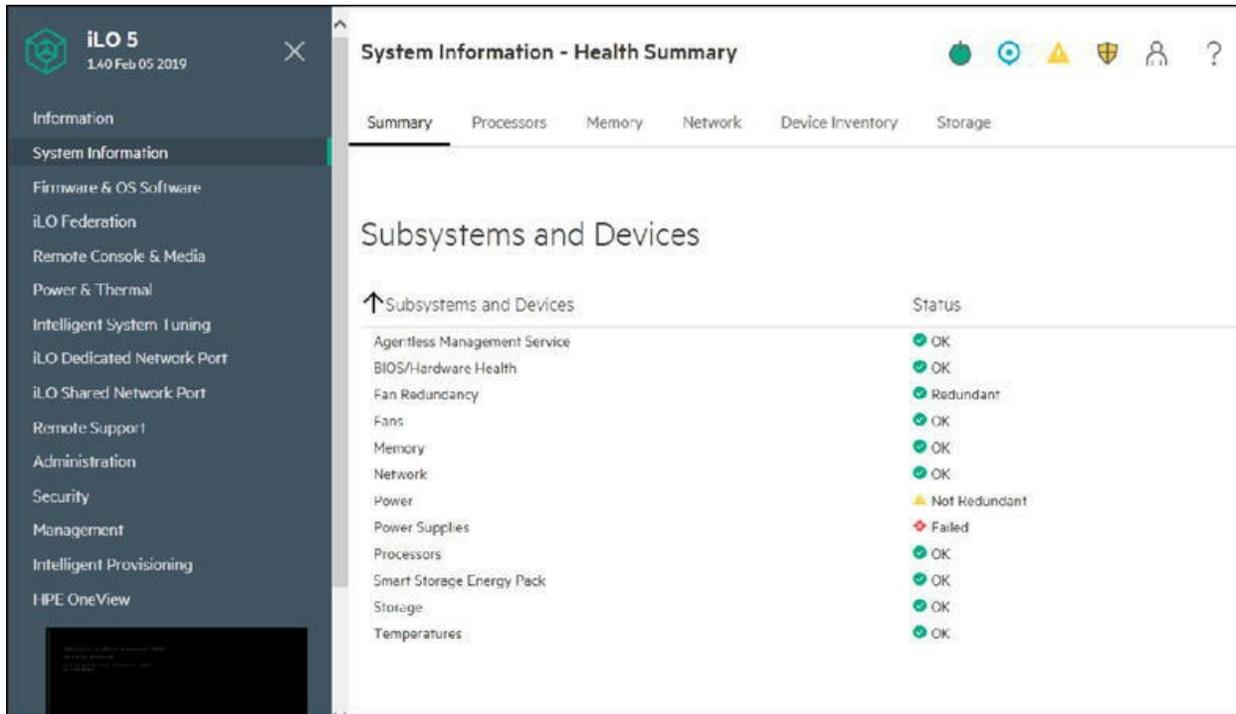


Figure 2-112 iLO web interface

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 seen in [Figure 2-112](#), navigational tree view displays include:

- Information
- System Information
- Firmware & OS Software
- iLO Federation
- Remote Console & Media
- Power & Thermal
- Intelligent System Tuning
- iLO Dedicated Network Port
- iLO Shared Network Port
- Remote Support

- Administration
- Security
- Management
- Intelligent Provisioning
- HPE OneView (optional)

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 web interface can be used to change the boot order for next restart, such as booting into Intelligent Provisioning. iLO can also generate SNMP traps or email alerts in case of hardware issues.

iLO Federation

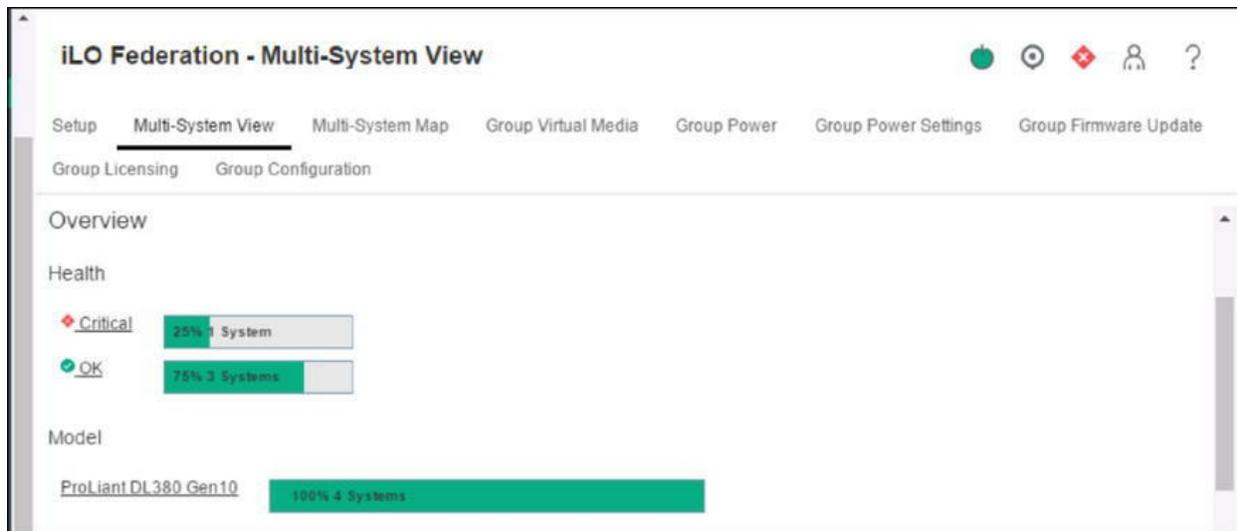


Figure 2-113 Viewing group iLO Federation Multi-System View

As illustrated in [Figure 2-113](#), 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 an 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.

iLO Federation discovery is supported with HPE SUM 6.3.0 and later.

Server boot process

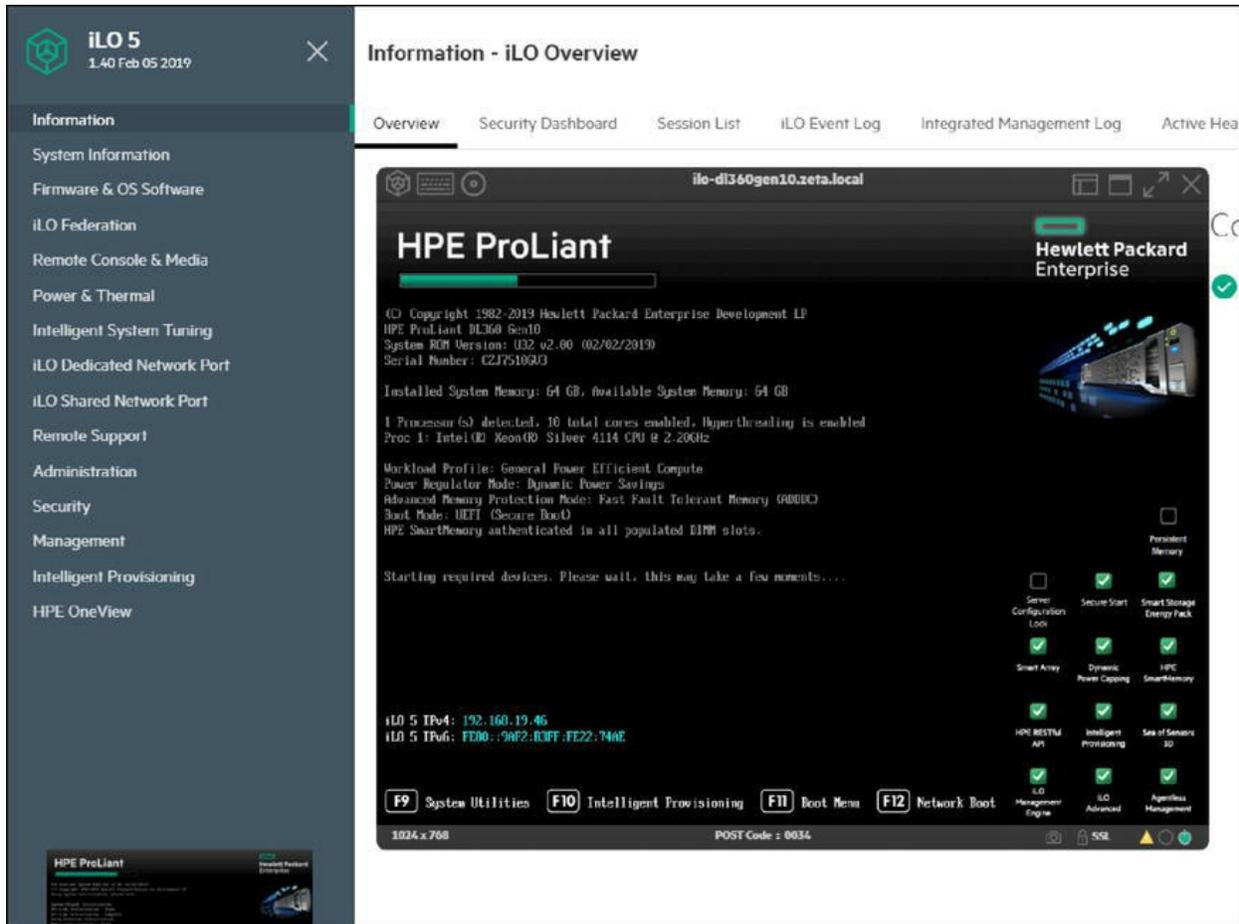


Figure 2-114 ProLiant POST screen

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.

The following steps are general instructions for booting up ProLiant Gen10 servers:

1. 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.

2. Press the **Power On/Standby** button.
3. Near the end of the boot process, the POST sequence runs and the POST screen displays, as shown in [Figure 2-114](#). 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 access the System Utilities screen.
 - If you do not need to modify the server configuration and are ready to install the operating system software, press **F10** to access Intelligent Provisioning.
 - Access the Boot menu by pressing the **F11** key.
 - Force a Preboot Execution Environment (PXE) network boot by pressing the **F12** key.
4. To exit the System Utilities screen and reboot the server, select one of the following options:
 - Press **F7** or click **F7: Load Defaults** at the bottom of the screen.
 - Press **F10** or click **F10: Save**; this saves any changes pending.
 - Press **F12** or click **F12: Save and Exit**.
 - Click **Exit**, in the pop-up select **Cancel** or **OK** to reboot the system, as seen in [Figure 2-115](#).

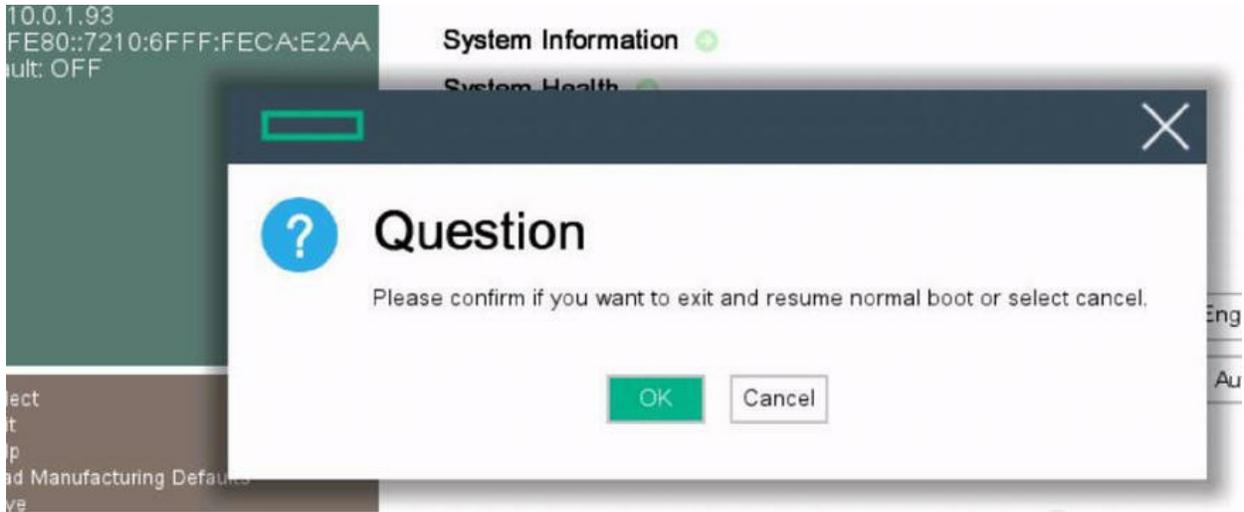


Figure 2-115 Click to confirm reboot

ProLiant Gen10 preboot environment

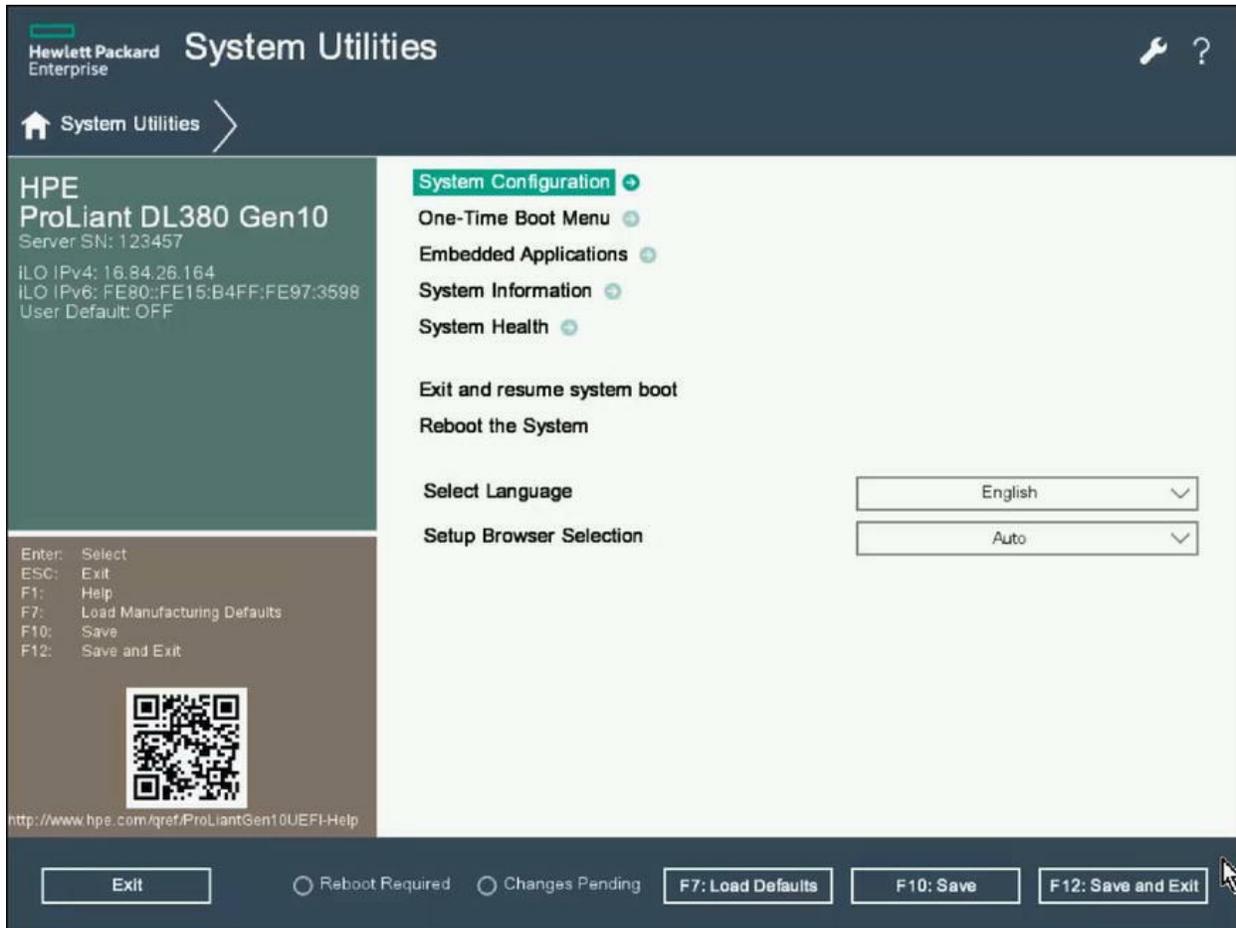


Figure 2-116 ProLiant System utilities

The System Utilities screen is the main screen in the UEFI menu-driven interface. As illustrated in [Figure 2-116](#), 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.
- **Setup Browser Selection**—Enables you to configure the appearance of UEFI System utilities.

BIOS/Platform Configuration

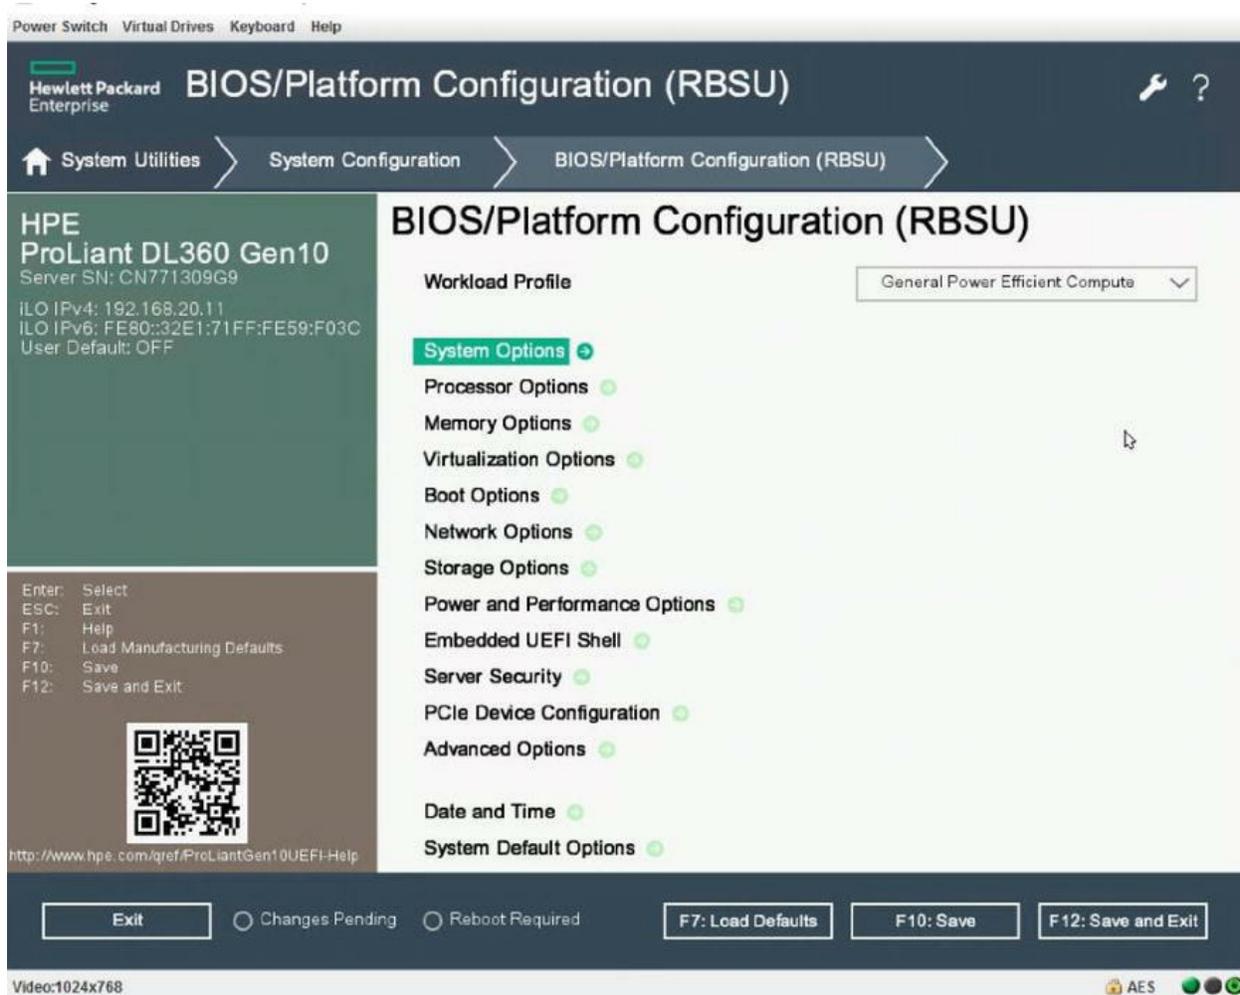


Figure 2-117 BIOS/Platform Configuration

As seen in [Figure 2-117](#), the BIOS/Platform Configuration menu 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**— Choose 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**—This option allows Intel® Hyper-Threading control, cores enabled, and distributed interrupt handling control.
- **Memory Options**—This option allows Advanced Memory Protection control and other advanced memory options to be set.
- **Virtualization Options**—This menu option 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 operating system installation. If not, changing the boot mode can impact the ability of the server to boot to the installed operating system. 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 operating systems 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 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 internet Small Computer System Interface (iSCSI) and virtual local area network (VLAN) settings.
- **Storage Options**—This option is used to configure the embedded SATA Controller options, the embedded storage boot policy, Fibre Channel/FCoE scan policy, and NVMe 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, and to enable or disable the Intelligent Provisioning F10 prompt, and 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 Peripheral Component Interconnect (PCI) devices. Disabling devices reallocates 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 options to control POST debug message reporting and advanced service options are found here.

cooling reducing setting, and many others. Therefore, servers can be tuned to match specific workloads.

There are a multitude of BIOS settings that can be adjusted to change performance, power, latency, memory, and other variables.

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

Here are descriptions of the preconfigured workload profiles:

- **General power-efficient compute (default)**—The default out-of-the-box profile. It uses the most common performance settings seen across most application workloads, but at the same time, turns on most power management settings that have minimal impact to 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 who value mission-critical 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 solve 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.
- **IO 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.

Intelligent Provisioning

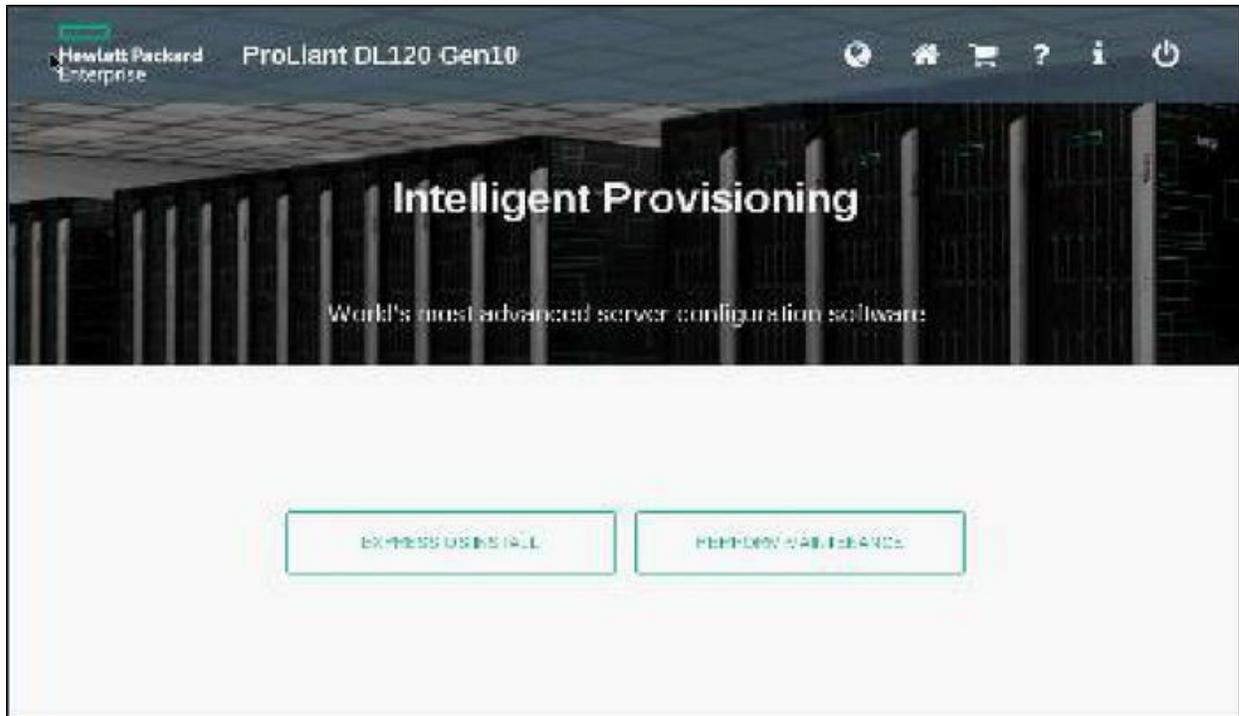


Figure 2-119 Server provisioning made easy

As seen in [Figure 2-119](#), Intelligent Provisioning (IP) 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 your customers can quickly and easily get their 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 your customer.

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 45% of server setup, configuration, and deployment steps, so you can deploy servers three times faster than before.

Use Intelligent Provisioning to:

- Deploy an operating system to Smart Array-based storage.
- Manage and update necessary drivers not included in base operating

system 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 operating system to two Secure Digital (SD) cards supported by HPE.

Launching Intelligent Provisioning

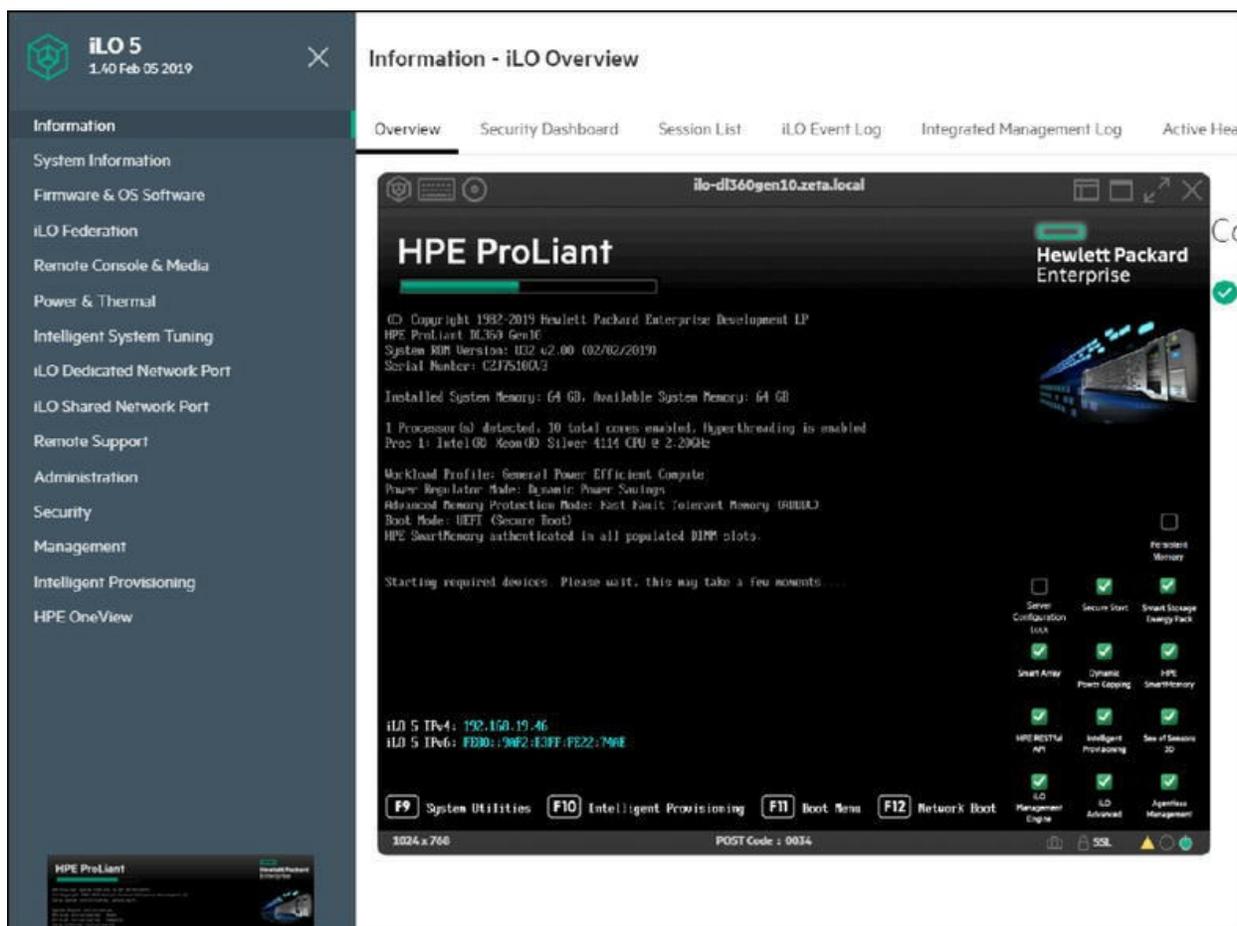


Figure 2-120 Launching Intelligent Provisioning

As seen in [Figure 2-120](#), to launch Intelligent Provisioning, power up the

server, and press **F10** during system boot. You can also access Intelligent Provisioning from the System Utilities screen. Select **Embedded Applications** and press **Enter**. Then select **Intelligent Provisioning**, and press **Enter**. Alternatively, IP can be launched from iLO and changing the next one-time boot option.

During the booting time, the startup menu options display:

- Intelligent Provisioning
- Smart Storage Administrator
- Scripting Toolkit Windows PE 64-bit mode

The highlighted option will start automatically in 15 seconds.

The Intelligent Provisioning home screen displays.

This screen shot 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.

- **Express OS Install**—Enables you to configure hardware and install an operating system.
- **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 entering the hyperlink into your web browser. https://h20564.www2.hpe.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 **F10** during start up and wait for Intelligent Provisioning to load.

Intelligent Provisioning operating system installation

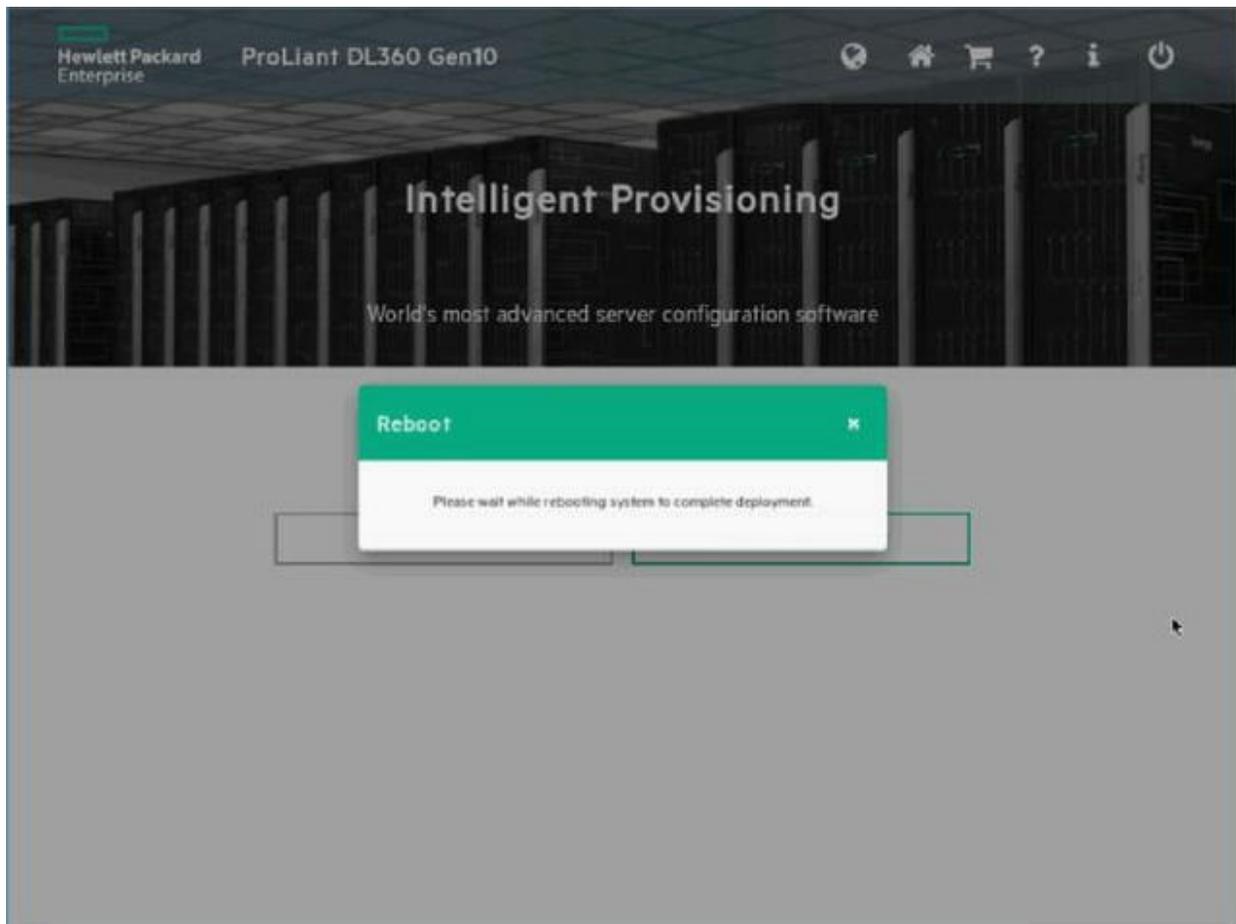


Figure 2-121 Intelligent Provisioning operating system installation

As referenced in [Figure 2-121](#), the steps to install the Intelligent Provisioning operating system are:

1. Select Install OS.
2. Mount the media in iLO.

3. Confirm selection.
4. Enter customization.
5. Confirm settings.
6. Launch installation.
7. Initiate reboot.

To use Intelligent Provisioning to configure the hardware and install an operating system 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 operating system 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 operating system you are installing, you might be prompted to enter or confirm information regarding the partition size, keyboard language, or product key.
- **Review**—The Review screen displays your hardware and operating system settings. Continuing past this screen installs the operating system and configures the server.



Notice

Continuing past this screen resets the drives to a newly installed state and installs the selected operating system. Any existing information on the server is erased. This does not affect a first-time setup, because there is no data present on the server.

Intelligent Provisioning Perform Maintenance options

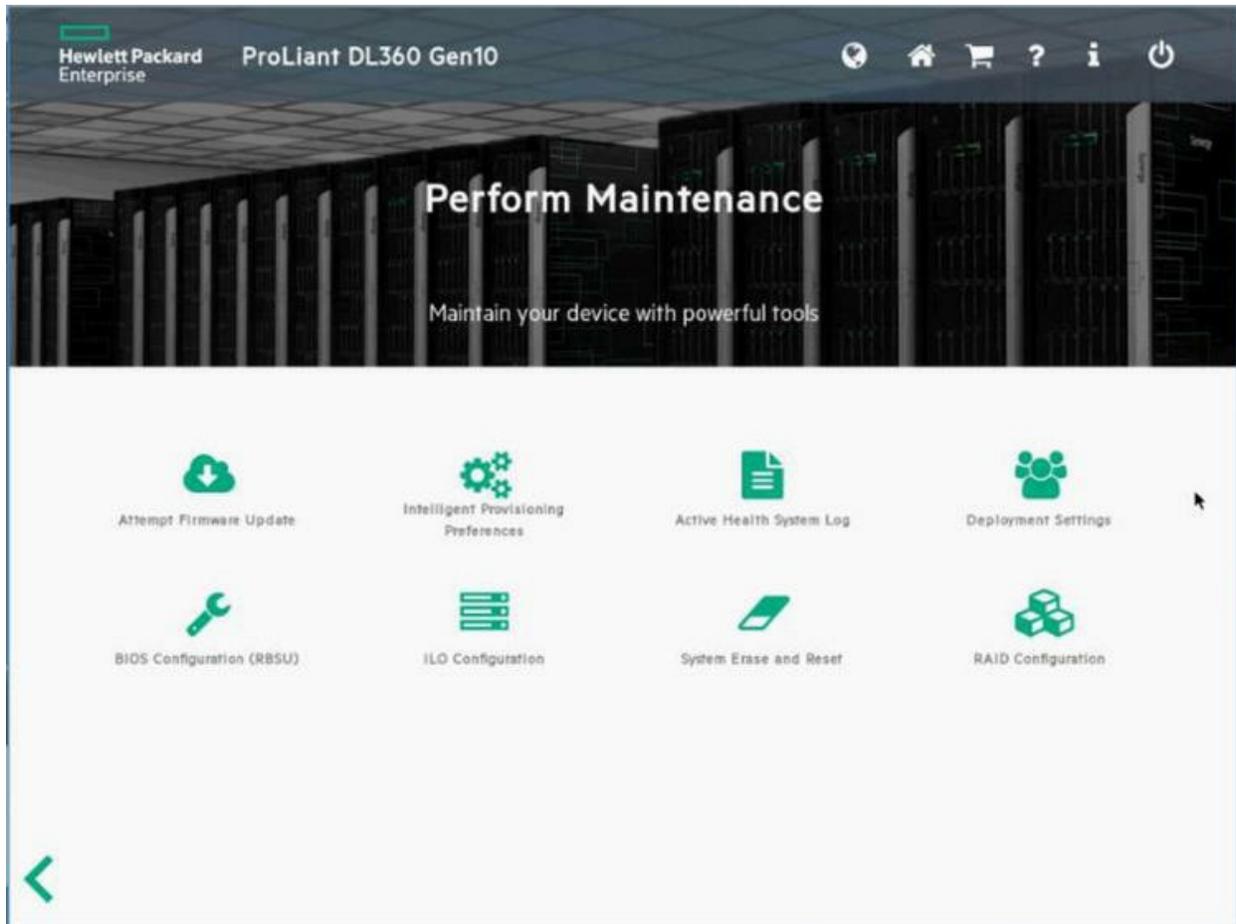


Figure 2-122 Intelligent Provisioning Perform Maintenance options

As shown in [Figure 2-122](#), clicking the **Perform Maintenance** link on the Intelligent Provisioning home screen displays a screen with several options for configuration utilities needed to maintain ProLiant servers, such as:

- **Attempt Firmware Update**—Intelligent Provisioning manages the firmware update process. The Service Pack for ProLiant (SPP) firmware repository URL is configured in the Provisioning preferences and requires network interface configuration.
- **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 End User License Agreement (EULA) is also accessible from this screen.

- **Active Health System Log**—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 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.
- **Deployment Settings**—Deployment Settings enable you to create a server configuration package that you can deploy to one or more ProLiant 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.
- **iLO Configuration**—Use the iLO Configuration Utility to change iLO settings through Intelligent Provisioning instead of through the iLO web interface. You must configure the 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.
- **System Erase and Reset**—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 the data wiping method 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, and 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 SAS, and Fibre Channel storage before using secure erase, unless they should also be erased.

- **RAID Configuration**—Starts Smart Storage Administrator to configure and diagnose local storage.

HPE SMB setup software

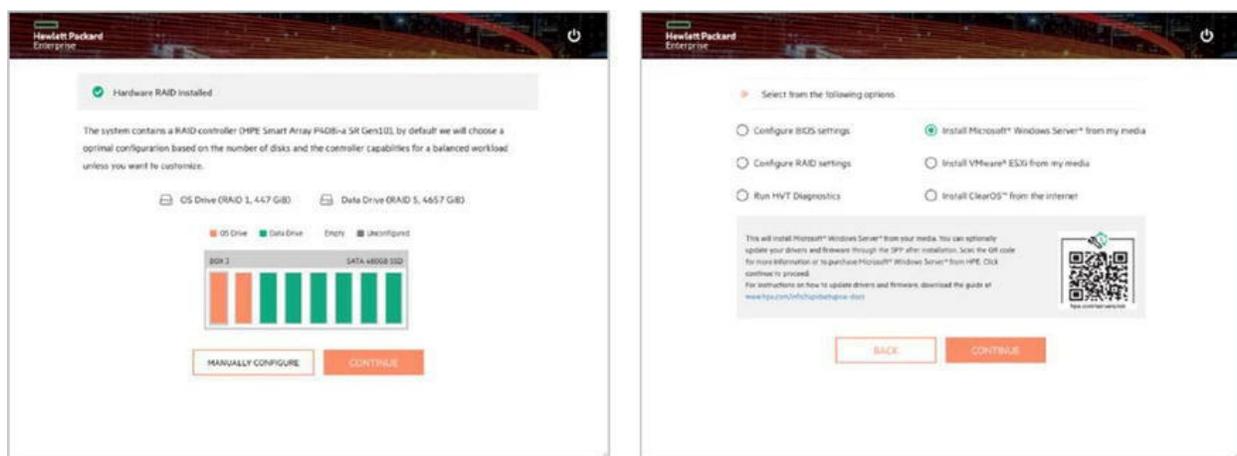


Figure 2-123 HPE SMB setup software

HPE SMB setup software, as referenced in [Figure 2-123](#), is embedded in ProLiant Gen10 servers and it is part of Intelligent Provisioning (ProLiant 10,100, and 300).

Simple three-step process:

- Set up hardware RAID.
- Select OS drive target.
- Select and install the OS or hypervisor.
- Optional hardware configuration (BIOS, hardware virtualization [HVT], RAID).

New or enhanced features:

- Automated RAID configuration for servers with five or more drives
- Enhanced manual RAID configuration option
- Enhanced server drive GUI
- ClearOS internet install
- Windows Server 2019 support

Smart Storage Administrator

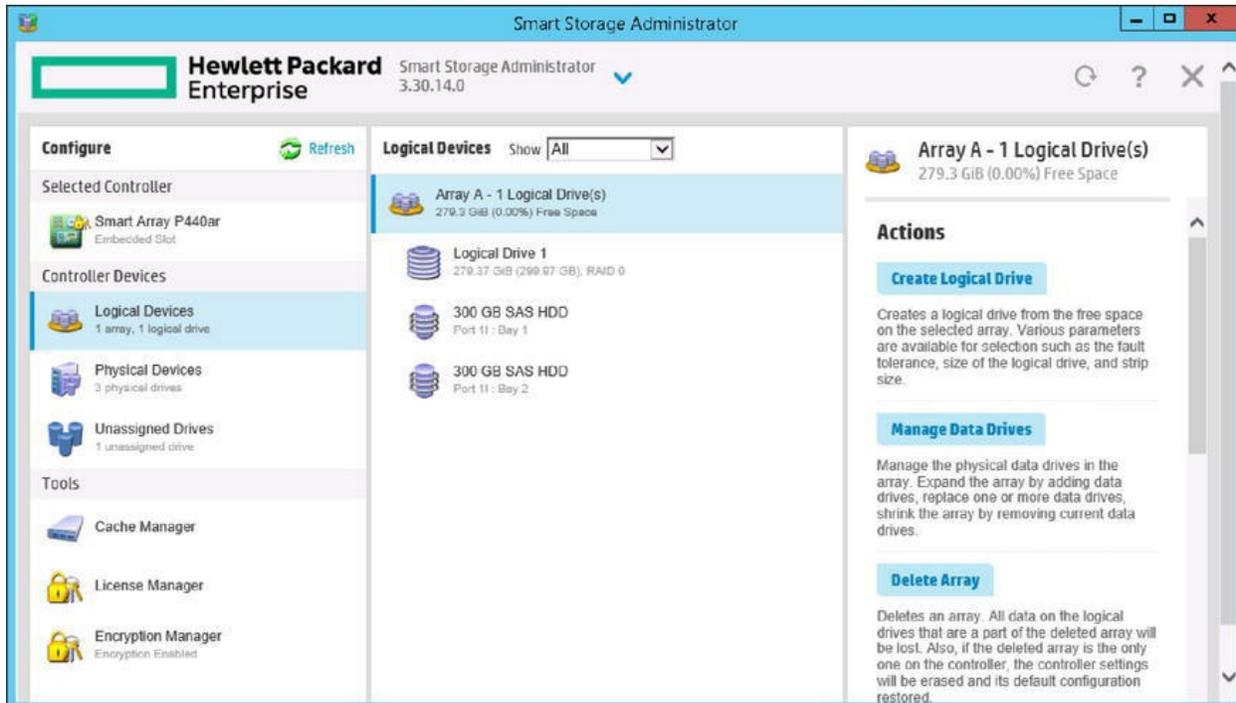


Figure 2-124 Smart Storage Administrator

SSA is an application that is accessed via Intelligent Provisioning, as illustrated in [Figure 2-124](#). 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 Smart Storage Administrator 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 operating system.



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.

Smart Storage Administrator 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

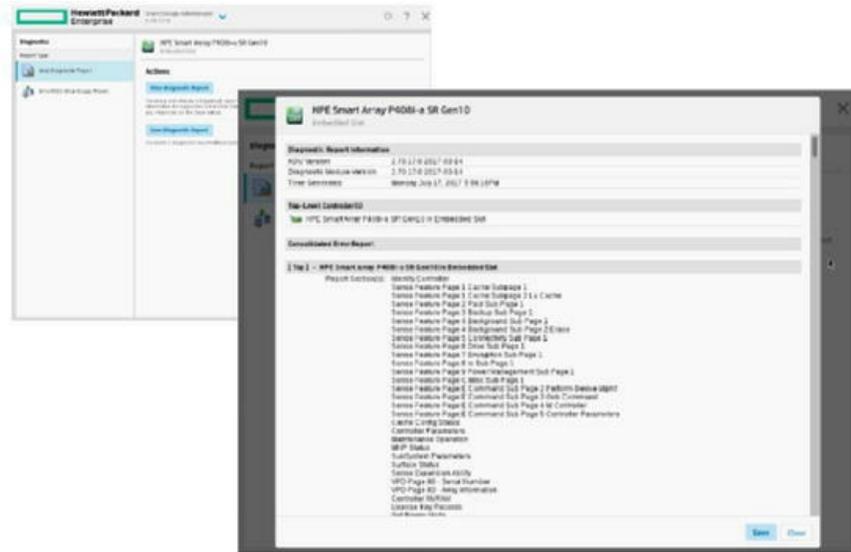


Figure 2-125 Smart Storage Administrator Diagnostic

As reference in [Figure 2-125](#), 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 SSDs, 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 SSDs 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

1. Open HPE SSA.
2. 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.
3. Select a report type. For this example, use the **Array Diagnostic Report** selection. 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
- Operating system device names or mount points

HPE Smart Update solution stack

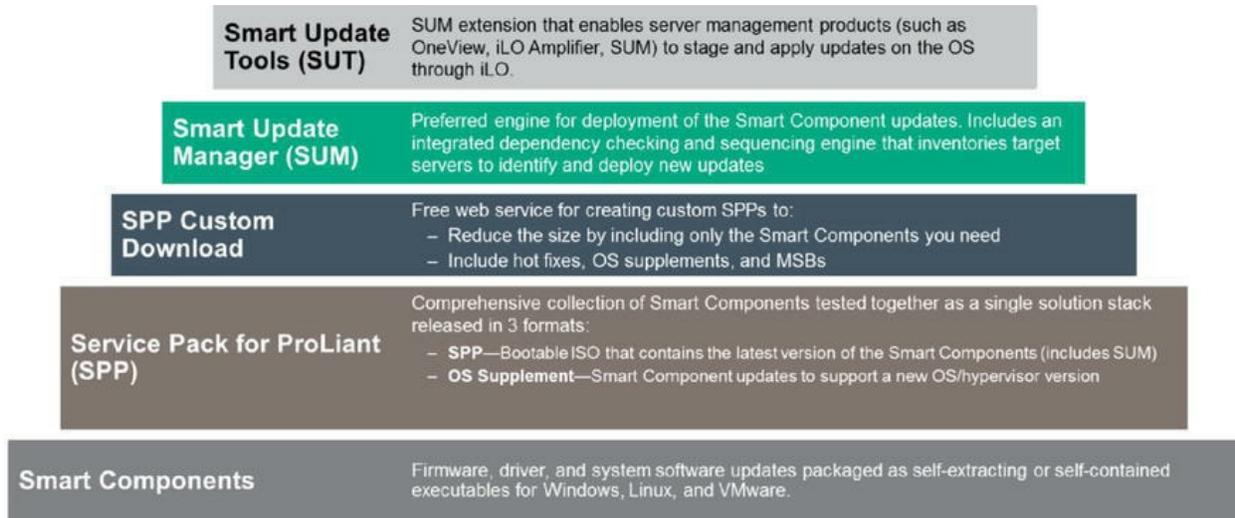


Figure 2-126 HPE Smart Update solution stack

As illustrated in [Figure 2-126](#), the HPE Smart Update solution stack 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

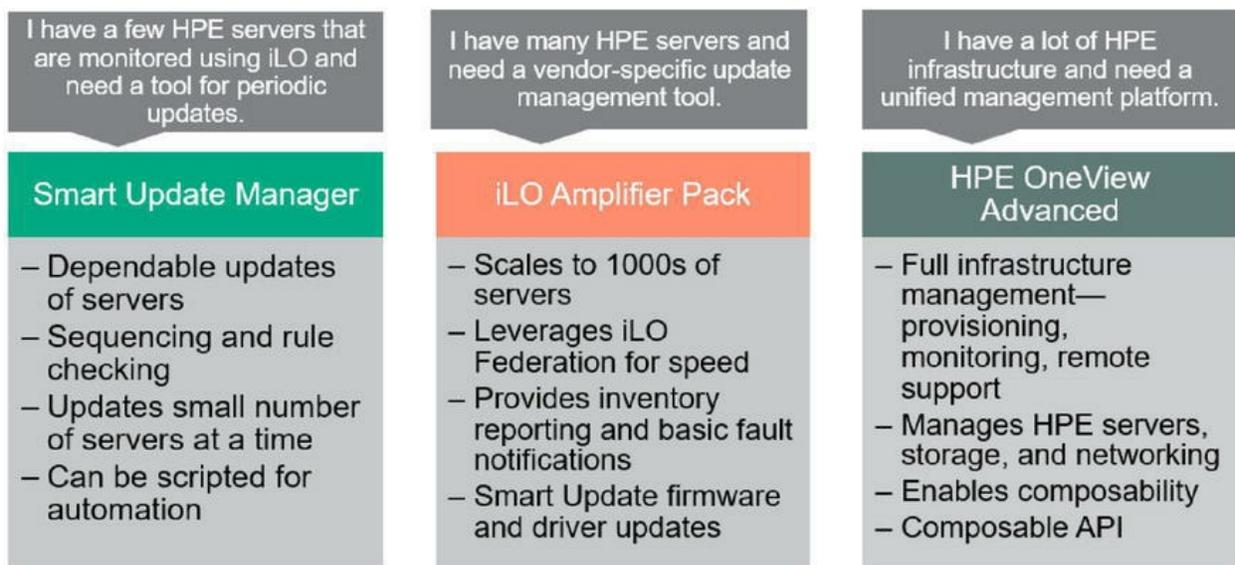


Figure 2-127 Smart Update Tools and how they help customers

As reference in [Figure 2-127](#), HPE Smart Update Tools (SUT/iSUT) are a SUM extension that enables HPE OneView and HPE iLO Amplifier Pack to stage, schedule, and apply updates automatically to reduce IT operations. SUT is an operating system (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.

Smart Update Tools 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.

Smart Update Tools support:

- HPE OneView external repository
- HPE OneView staging, scheduling, & deployment of updates
- 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 Smart Update Manager

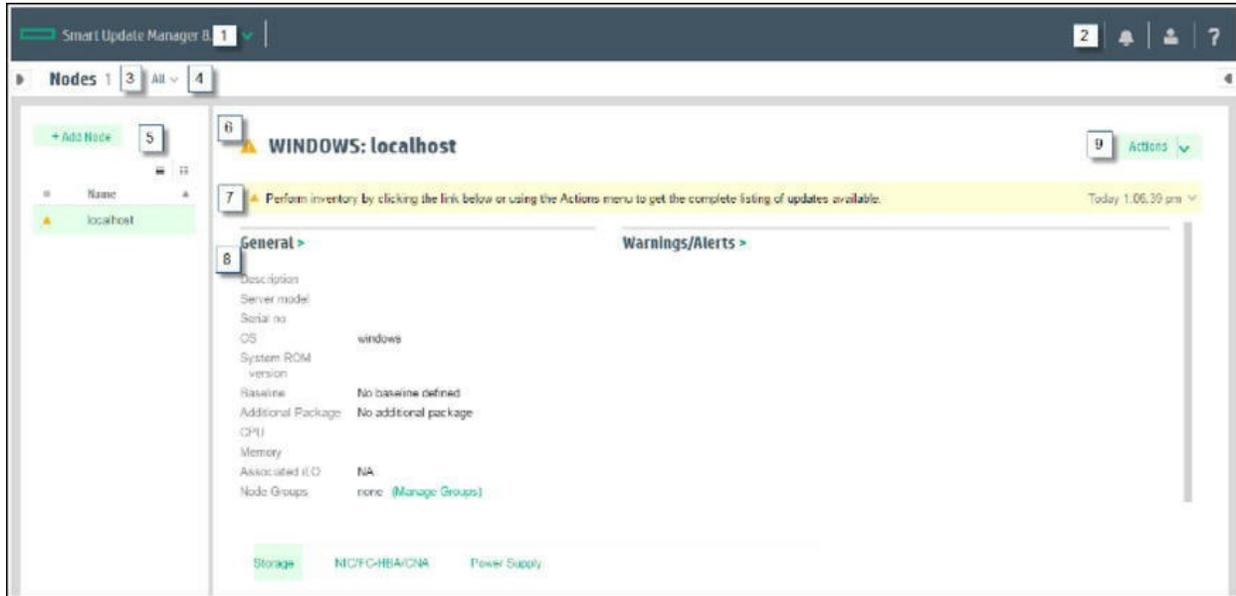


Figure 2-128 HPE Smart Update Manager GUI

HPE Smart Update Manager (SUM), as seen in [Figure 2-128](#), is a technology 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. The graphic above shows the major areas of the GUI. Screens options vary based on the node type, and the legend below describes them.

1. **Main menu**—Provides links to all HPE SUM screens
2. **Activity list**—Displays the status of recent activities
3. **Screen name**—Displays the name of the screen
4. **Screen filters**—Filters screen objects

5. **Screen list**—Displays a list of items on the screen
6. **Baseline or node name**—Lists the baseline or node that it selects
7. **Informational highlight**—Information you might need to perform or know about a selected item
8. **Item details**—Provides full details about a selected item
9. **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 local host 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:

1. Read/review all release notes.
2. Update drivers.
3. Update additional software, such as agents and utilities.
4. Update firmware.
5. Reboot, if necessary.

If you are installing a new server, update the firmware before installing the operating system.

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 Service Pack for ProLiant

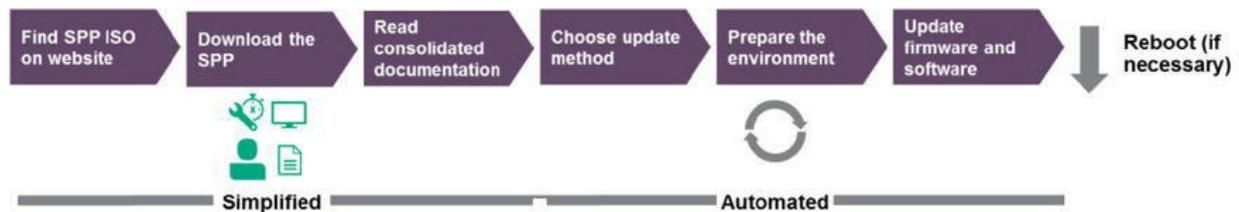


Figure 2-129 HPE Service Pack for ProLiant simplifies firmware and driver updates

The SPP includes components for HPE ProLiant server and enclosures systems, software, and firmware.

SUM deploys components to supported Windows, Linux, and VMware operating systems. 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, as referenced in [Figure 2-129](#).



Note

Each SPP release includes a version of SUM you can use to deploy the SPP components. Access the SUM website by right-clicking the hyperlink to open it.

<https://support.hpe.com/hpesc/public/home/productHome?sp4ts.oid=1008862656>

A single solution, consisting of SUM and SPP, 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 frames and 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 CNAs
- HPE SAS 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, visit:
<http://h17007.www1.hpe.com/us/en/enterprise/servers/solutions/inlibrary/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 operating system
- **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 operating systems.

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 applicable 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).

Learning check

13. Name three available tools to configure local storage in an HPE ProLiant server.

Introducing the customer scenario

This is the same scenario from the beginning of this chapter.

- What is your primary business?
 - City-wide courier, parcel, and documents delivery
- How many employees do you currently have?
 - 25 employees + external workers
- What does your selling and delivery channel look like?
 - Contracts with dozens of local e-shops and expanding
- What does your server, storage, and network infrastructure look like?
 - SaaS (email, collaboration, CRM)
- Do you have an IT department?
 - No, external help on request
- What are your current plans?
 - Deploying a CRM and tracking onsite (speed, regulatory requirements —General Data Protection Regulation [GDPR])

Customer requirements

As a result of multiple interviews and gathering information about customer plans and customer's current infrastructure, the following requirements emerged for the new solution:

- Server for a custom-made CRM and tracking application
 - Requires RHEL or SLES
- Single server, on-site installation
- No local data center, sufficient space in a storeroom
- Single-phase power
- 1 Gb/s Ethernet connection to third-party router
- 20 GB of data anticipated per year, 100% grow rate annually

- Design the solution for five years with sufficient room to grow
- Testing performed by supplier, and the following compute characteristics were verified as sufficient:
 - A single eight-core processor running at 2.1 GHz
 - 128 GB RAM
- Tape backup for long-term data protection and retention

Summary

- Assessing customer requirements includes analysis of business application and data requirements.
- HPE ProLiant server portfolio includes tower and rack servers suitable for SMB environment.
- Components of a compute solution include processor, memory, storage, and networking options.
- Cloud solutions from HPE include HPE Composable Cloud for ProLiant DL and HPE ProLiant for Microsoft Azure Stack.
- HPE provides a variety of presales configuration tools and management options for server portfolio.

Prelearning check

Before proceeding with this section, answer the following question to assess your existing knowledge of the topics covered in this chapter. Record your answers in the space provided.

1. You are in a meeting with a customer, and they challenge you with a statement: “I heard that the management of networking devices includes manual and repetitive tasks, with a high risk of errors.”

How should you respond?

3 Recommending HPE networking solutions for SMB customers

LEARNING OBJECTIVES

After completing this chapter, you should be able to recommend and position HPE networking products, solutions, tools, and appropriate services for customer use cases.

Recommending HPE networking solutions for an SMB

The wide HPE portfolio of data center products for SMB customers can be introduced using customer scenarios.

Introducing the customer scenario

The fictional customer LLP distribution company will be used as a storyline through this scenario. We will introduce the company using an interview:

- What is your primary business?
 - City-wide courier, parcel, and documents delivery
- How many employees do you currently have?
 - 25 employees + external workers
- What does your selling and delivery channel look like?
 - Contracts with dozens of local e-shops and expanding
- What does your server, storage, and network infrastructure look like?

- Several tower servers and older networking devices from multiple brands
- Do you have an IT department?
 - No, external help on request
- What are your current plans?
 - Replace multiple vendor networking devices

Customer requirements

As a result of multiple interviews and gathering information about customer plans and customer's current infrastructure, the following requirements emerged for the new solution:

- Single vendor networking infrastructure
- Port requirements:
 - Sixteen ports for servers
 - Two switches
- Speed requirements
 - 1 Gb (26 ports)
- Teaming/bonding required on server ports
- L3 functionality
 - Dynamic Host Configuration Protocol (DHCP)
 - Static routing

Activity: Discovering a customer's business and technical requirements

1. Prepare a list of additional questions to ask the customer about:
 - Current environment

Selecting an HPE Networking platform

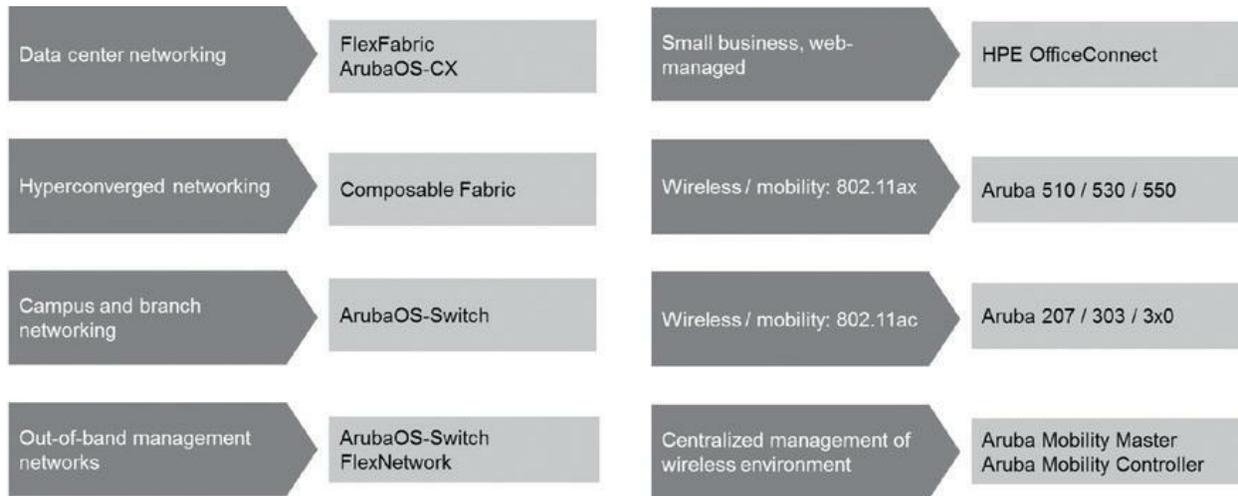


Figure 3-1 Selecting an HPE networking platform

Depending on the customer's current environment and requirements, decision trees can be used to select appropriate networking platform, as illustrated in [Figure 3-1](#).



Note

Other limitations and prerequisites apply. Consult QuickSpecs. Some features require additional components. Capacities are changing during time.



Note

Many HPE networking platforms are missing in the tree and are out of the scope for this book.



Note

Not all HPE networking platforms are suitable for data center and are out of the scope for this book.

Selecting an HPE switch: Selection criteria

Typical criteria used for selecting specific HPE networking device:

- Port count
 - For example: 5, 8, 16, 24, 28, 48
 - Modular vs. fixed ports
- Management
 - Unmanaged
 - Smart management
 - Full management (Simple Network Management Protocol [SNMP], Secure Shell [SSH], cloud-based, on-premise centralized management)



Note

Unmanaged or smart-managed devices are not typical products for data center.

- Layer 2/3 functionality
 - Switching only
 - Lite L3 (DHCP, static routing, Routing Information Protocol [RIP], ...)
 - Basic L3 (Open Shortest Path First [OSPF], policy-based routing, ...)
 - Advanced L3 (Border Gateway Protocol [BGP], Bidirectional Forwarding Detection [BFD], Virtual Router Redundancy Protocol [VRRP], and so forth)
- Port type and speed
 - Downlinks
 - Uplinks
 - 1/10, 10, 25, 40, 100

- Copper, Fiber
 - Extra features
 - Hardware redundancy (fans, power, fabric)
 - Logical stacking
 - Product line
 - ArubaOS-Switch
 - ArubaOS-CX
 - FlexFabric
 - FlexNetwork
 - OfficeConnect
 - Arista
 - Altoline
-



Note

Open Shortest Path First (OSPF) is a link state routing protocol (LSRP) that uses the Shortest Path First (SPF) network communication algorithm (Dijkstra's algorithm) to calculate the shortest connection path between known devices ([Techopedia.com](https://www.techopedia.com/entry/43147/ospf-protocol), October 2019).



Note

Border Gateway Protocol (BGP) is a routing protocol used to transfer data and information between different host gateways, the Internet, or autonomous systems. BGP is a Path Vector Protocol (PVP), which maintains paths to different hosts, networks, and gateway routers and determines the routing decision based on that. It does not use Interior Gateway Protocol (IGP) metrics for routing decisions, but only decides the route based on path, network policies, and rule sets ([Techopedia.com](https://www.techopedia.com/entry/43147/bgp-protocol), October 2019).



Note

Not all HPE networking platforms are suitable for data center or SMB. Some platforms are out of the scope for this book.

Depending on functionality, downlink port speeds, uplink port speeds, or port count, decision trees can be used to select an appropriate HPE switch, as illustrated in 3-6.

Selecting an HPE switch: FlexNetwork



Figure 3-2 Functionality of FlexNetwork



Note

Other limitations and prerequisites apply. Consult QuickSpecs. Not all platforms and models are covered.

Selecting an HPE switch: Aruba and FlexFabric



Figure 3-3 Functionality of Aruba and FlexFabric



Note

Other limitations and prerequisites apply. Consult QuickSpecs. Typical use case for ArubaOS-Switch devices (2530, 2540, 2930, 3810M, 5400R zl2) are out-of-band management networks.

Selecting an HPE switch: Port speed—Downlinks

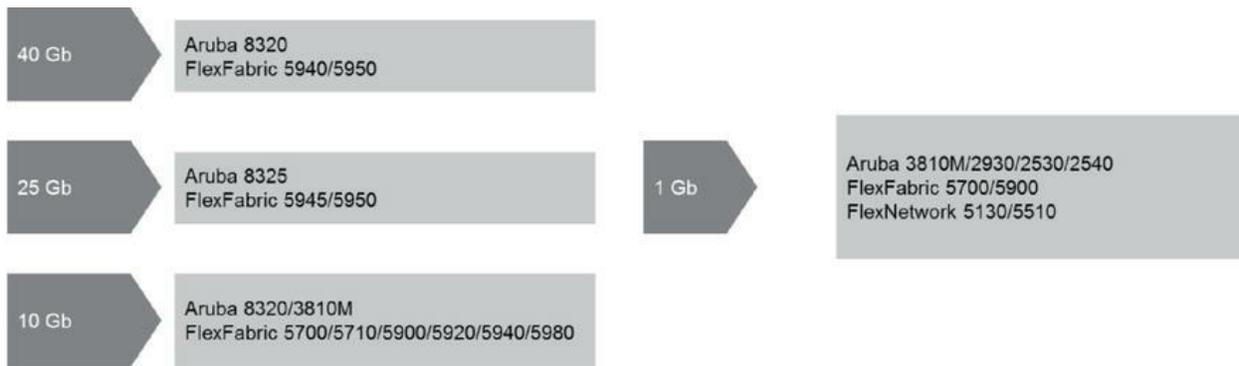


Figure 3-4 Port speed—Downlinks



Note

Other limitations and prerequisites apply. Consult QuickSpecs. Modular switches are not included.

Selecting an HPE switch: Port speed—Uplinks

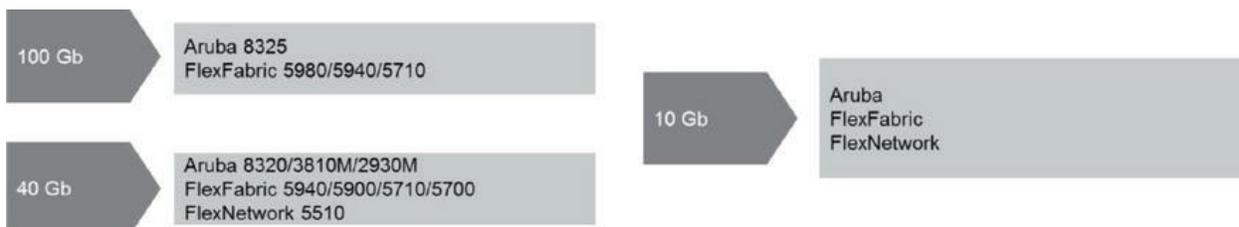


Figure 3-5 Port speed—Uplinks



Note

Other limitations and prerequisites apply. Consult QuickSpecs. Modular switches are not included.

Selecting an HPE switch: Port count

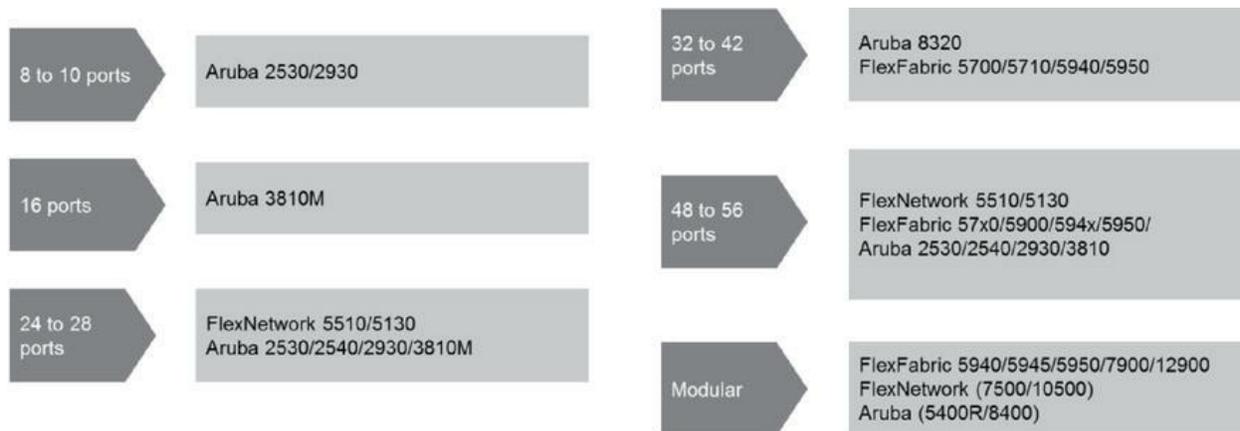


Figure 3-6 Port count



Note

Other limitations and prerequisites apply. Consult QuickSpecs.

HPE Switch Selector

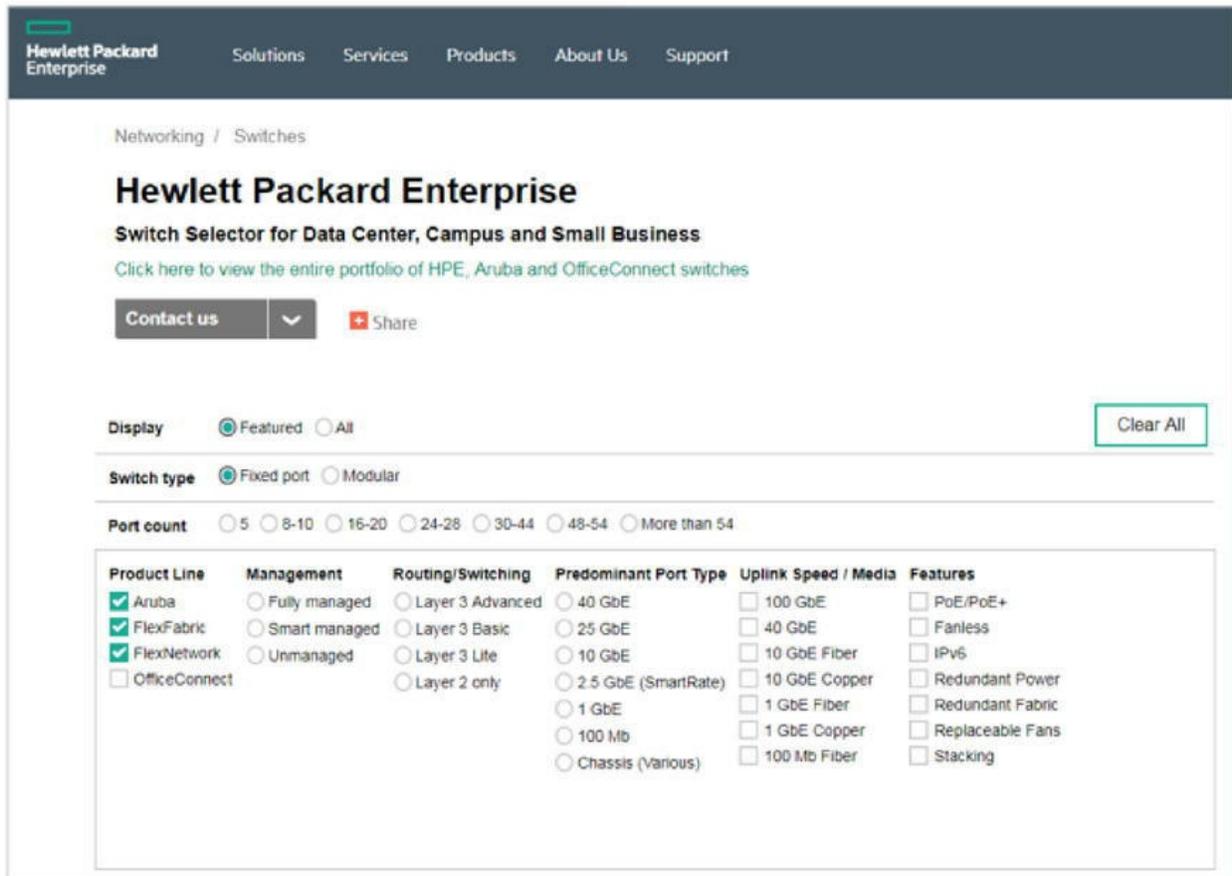


Figure 3-7 HPE Switch Selector

For Ethernet and TCP/IP-based storage solutions, an adequate switch might be needed. The HPE Switch Selector Tool can help architects narrow product family selections based on network needs, infrastructure requirements, or competitive equivalents. The HPE Switch Selector, as shown in Figure 3-7, allows you to view the complete HPE line of switches at a glance, and you can also sort the switches by port counts, speed, and other criteria.



Note

To access the HPE Switch Selector tool, visit:
<http://h17007.www1.hp.com/us/en/networking/products/switches-selector.aspx>

HPE Networking Online Configurator

Welcome to the HPE Networking Online Configurator which contains the most current HPE Networking product and pricing information. The online configurator streamlines the ability to select and configure our products and to create quotes for you and your customers.

Step 1. Select price list United States Step 2. Add products Step 3. Configure Step 4. Save quote Step 5. Export quote

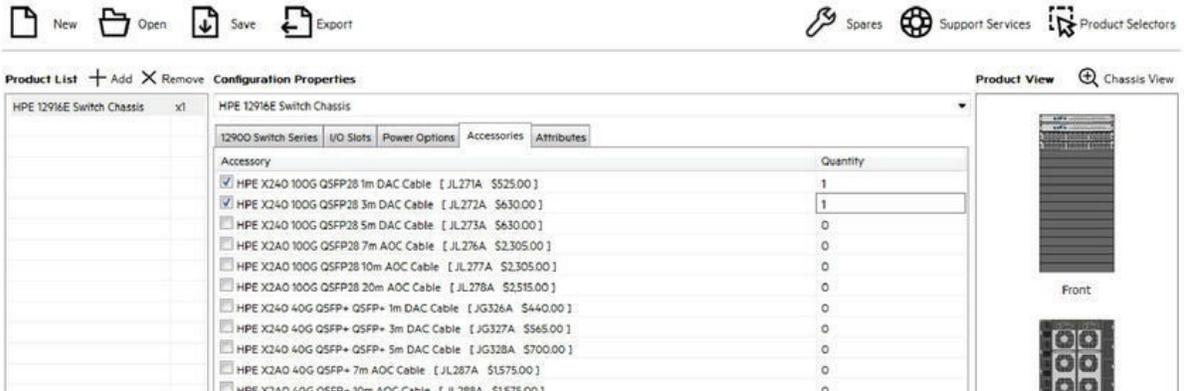


Figure 3-8 HPE Networking Online Configurator

The HPE Networking Online Configurator, as shown in [Figure 3-8](#), contains the most current HPE Networking product and pricing information. The online configurator streamlines the ability to select and configure our products and to create quotes.

The HPE Networking Online Configurator enables architects to quickly and easily create quotations of HPE products using a web browser (Internet Explorer and Firefox are supported). Quotation files can be saved locally and can be exported in several formats, including Excel.

Product families and products in the configurator include:

- Aruba
- FlexFabric
- Routers
- Mobility controllers and access points
- Management tools
- Services
- And much more (Arista, Altoline, network virtualization ...)



Note

To access the HPE Networking Online Configurator tool, visit:
<http://h17007.www1.hpe.com/us/en/networking/products/configur>

Activity: Switch Selector for Data Center, Campus, and Small Business

1. Open the switch selector at:
<http://h17007.www1.hpe.com/us/en/networking/products/switches/switch-selector.aspx>

2. Answer the following questions:

a. Which HPE platforms are supported by the tool?

b. How many total devices/products are available?

c. Which availability features can be used to filter results?

Learning check

1. Recommend an HPE networking platform for an out-of-band management network.

Networking architectures

Switches can be organized by where they are deployed in the environment. Traditional networks are organized into three tiers:

- **Core** switches establish the backbone of the network.
- **Distribution** switches are consolidation points for local area network (LAN) access or server access switches and connect to the core switches.
- **LAN** or **server access** switches support servers.

Tiered network configurations

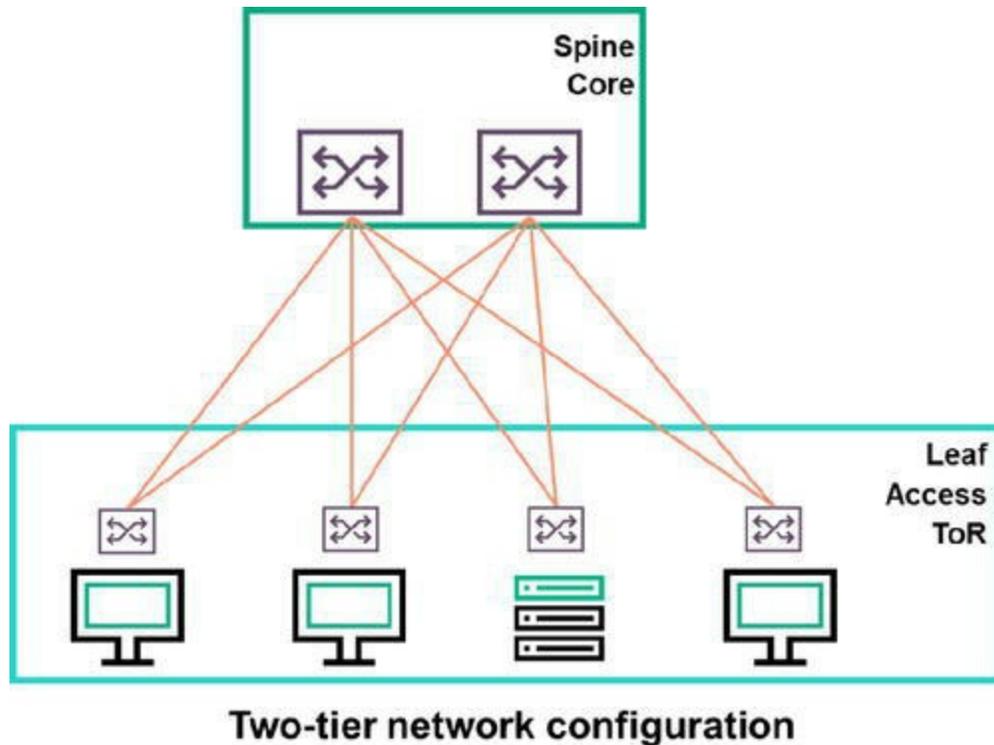


Figure 3-9 Two-tier network configuration

A traditional three-tier architecture includes the access layer, distribution layer, and core layer. Designed for network environments that were common 10 years ago, this architecture no longer meets the needs of many companies. Because data centers are highly virtualized, traffic patterns have changed. In the past, a majority of network traffic was exchanged between the clients and servers—a traffic flow called **north-south**. Now a lot of traffic is exchanged among virtual machines (VMs) in data centers and campus networks—a traffic flow called **east-west**. The traditional three-tier network is not designed to handle this traffic.

Users are demanding applications—such as multimedia, collaboration, and voice over IP (VoIP)—that require high-performance computing. A three-tier network introduces latency, disrupting these applications. A three-tier network relies on technologies such as Multiple Spanning Tree Protocol (MSTP) to help provide redundancy and reliability. Although MSTP is still widely in use, it fails over much too slowly to support today’s networks.

A two-tier network, as illustrated in [Figure 3-9](#), addresses those limitations. A two-tier network consists of the leaf or access layer, and the spine or core

layer.

In a two-tier network:

- The distribution layer is eliminated.
- LAN and server access switches connect directly to the core switches.
- Traffic flows directly from the edge to the core, reducing latency.

Leaf-spine topology

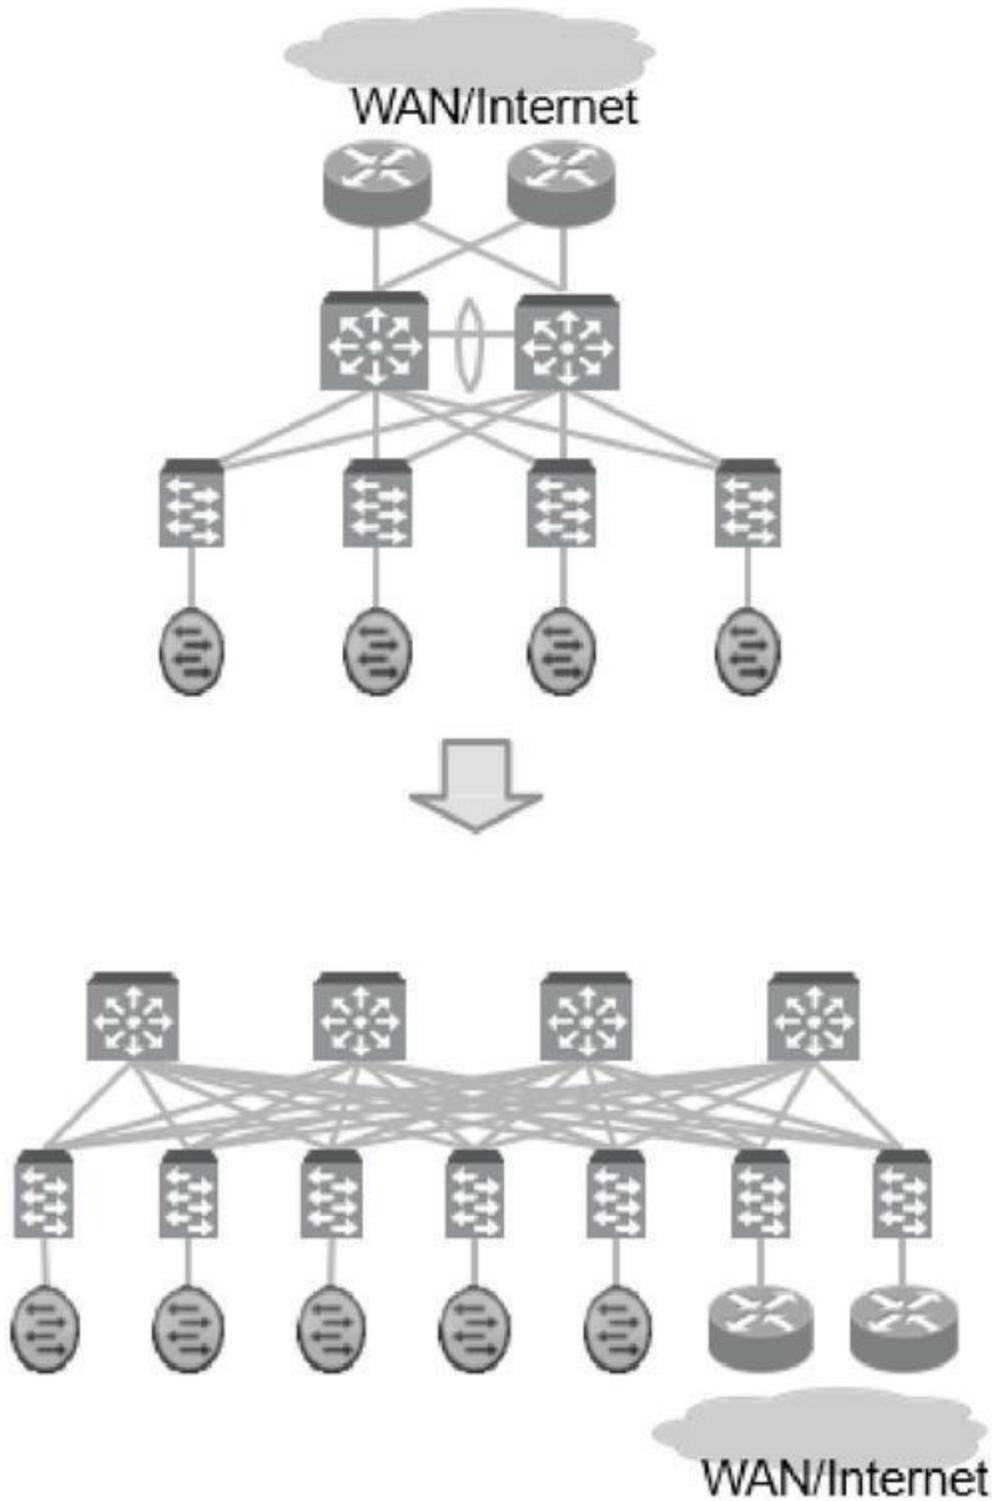


Figure 3-10 Three-tier traditional design and leaf-spine topology

Layer 2 networks were widely used in the early days of networking when

networks were small. As networks grew, they moved toward Layer 3 designs so they could scale and reduce broadcast domains. However, the rapid growth and deployment of virtualization saw a resurgence of L2 architectures and the deployment of new large-scale L2 architectures. These deployments are designed so that virtual LANs (VLANs) can extend from rack to rack and across data center boundaries, creating a large L2 environment optimized for the growing use of VM migration and disaster recovery.

L3 architectures that route packets at each device were widely deployed before the increase of virtualization. Virtualization drove the deployment of L2 architectures to satisfy the requirements of VM migration. However, within many environments, L2 extension is still not a priority. These types of environments can benefit from the scaling and efficiency advantages that L3 architectures provide.

An alternative to the core/aggregation/access layer network topology has emerged known as leafspine. In a **leaf-spine** architecture, as illustrated in [Figure 3-10](#), the access layer is formed by a series of leaf switches that are fully meshed to a series of spine switches. Leaf switches are often used in data centers to aggregate traffic from server nodes and then connect to the core of the network, consisting of spine switches. Next-generation multicore-server CPUs with dense VMs and storage make the use of leaf-spine topology critical. In addition, uplink, downlink, and peer ports that are not oversubscribed and are all switched at wire speed are important.

The leaf-spine topology solves the IP/Ethernet storage challenge of massive east to west traffic in a solution with deterministic latency characteristics, any-to-any non-blocking host communication, and deep buffers capable of absorbing the largest of bursts and Transmission Control Protocol (TCP) incast traffic patterns.

Leaf-spine architectures, as illustrated in [Figure 3-11](#), can be L2 or L3, meaning that the links between the leaf and spine layer could be either switched or routed. You can run L2 or L3 for connectivity between the leaves and spines.

Data center architectures

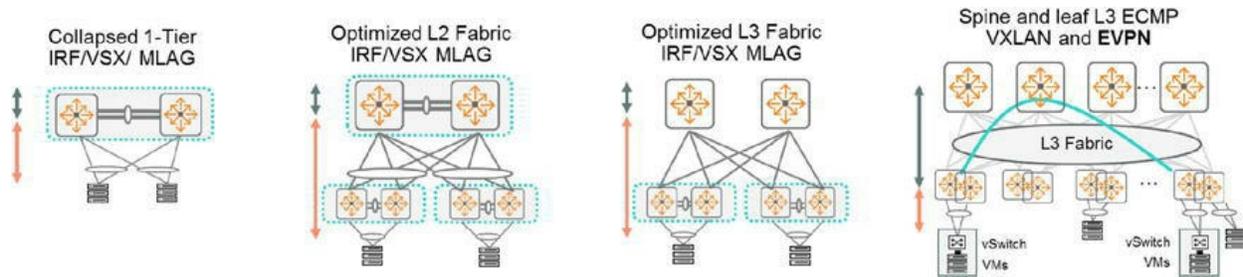


Figure 3-11 Data center architectures

Comparison of various data center architectures:

Architecture	Customer persona	Number of servers	VM scale	Features
Collapsed 1-Tier IRF/VSX	Small server rooms K to 12 school districts	50 to 100	5000+	L2, MLAG, VSX, DCB, API integration
Optimized L2 Fabric IRF/VSX	Small to medium data centers Education, local gov, retail	~ 100 to 500	25,000 to 50,000+	L2, VSX (MLAG + Config Sync), DCB, API integration
Optimized L3 Fabric IRF/VSX	Medium to large Enterprises, universities	~500 to 2000	100,000 to 500,000 +	ECMP, L3 Routing, IPv6, VSX, DCB, NSX, API integration
Spine and leaf L3 ECMP VXLAN and EVPN	Large and complex data centers Financial services, large enterprises	~ 2000+	750,000 +	VXLAN with MP-BGP EVPN, ECMP, L3, VSX, DCB, NSX, API integration



Note

Acronyms are defined below:

- Application programming interface (API)
- Data center bridging (DCB)
- Equal-cost multi-path routing (ECMP)

- Ethernet virtual private network (EVPN)
 - Intelligent Resilient Framework (IRF)
 - Multichassis Link Aggregation (MLAG/MCLAG)
 - Multiprotocol Extensions for BGP (MP-BGP)
 - Virtual Switching Extension (VSX)
 - Virtual Extensible LAN (VXLAN)
-

HPE Intelligent Resilient Framework

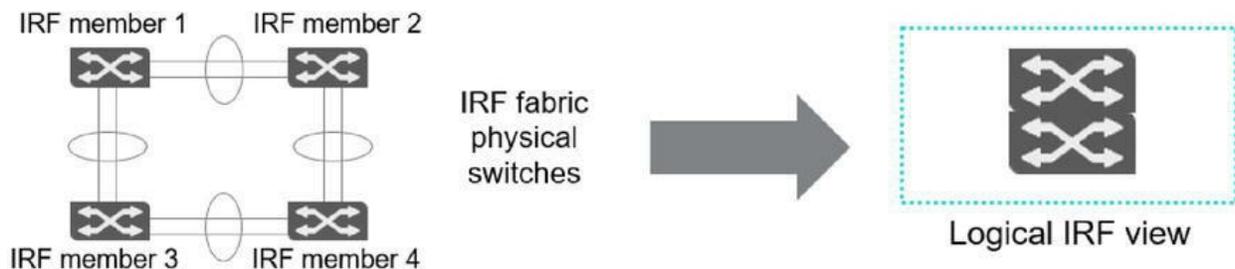


Figure 3-12 IRF provides a simple, resilient, and versatile pay-as-you-grow business model

HPE Intelligent Resilient Framework (IRF), as indicated in [Figure 3-12](#), is an HPE technology that extends network control over multiple active switches. The two-tier topology is connected with aggregated links—increasing bandwidth and providing redundancy. IRF can load balance traffic across all the links.

IRF virtual devices operate efficiently and eliminate the need to run outdated protocols that slow down the network and waste valuable network bandwidth. This can provide resiliency and other enhancements at any layer of the network.

At the core, IRF virtual devices eliminate a single point of failure (SPOF) while simplifying router redundancy using a single IP address rather than complex router redundancy protocols. At the edge of the network, IRF virtual devices can interconnect hundreds of virtual machines, reducing latency, optimizing performance for high-volume server-to-server traffic flows, and eliminating unnecessary network hops.

Devices that form an IRF virtual device are called **IRF member devices**. One IRF member operates as the primary system switch, maintaining the control plane and updating forwarding and routing tables for other devices. If the primary switch fails, IRF instantly selects a new primary switch, preventing service interruption and helping to deliver network, application, and business continuity for business-critical applications.

An IRF fabric uses member IDs to uniquely identify and manage its members. This member ID information is included as the first part of interface numbers and file paths to uniquely identify interfaces and files in an IRF fabric. If two devices have the same IRF member ID, they cannot form an IRF fabric. If the IRF member ID of a device has been used in an IRF fabric, the device cannot join the fabric. One IRF fabric forms one IRF domain. IRF uses IRF domain IDs to uniquely identify IRF fabrics and prevent IRF fabrics from interfering with one another.

Routing protocols calculate routes based on a single logical domain rather than the multiple switches it represents. Within an IRF domain, the geographic location of switches does not matter. Switches can be extended horizontally, and they continue to function as a single logical unit whether they are installed locally, distributed regionally, or even situated at distant sites. In addition, employing IRF can enhance disaster recovery by linking installations up to 70 kilometers apart and giving them the same fast failover as if they were sitting side by side within the data center. Such location independence is extremely important to support the global on-demand application access and dynamic traffic flows of technology-oriented businesses. IRF virtual devices provide the high performance and rapid failover required for delay-sensitive, business-critical applications.

You can assign several physical links between neighboring members to their IRF ports to create a load-balanced aggregate IRF connection with redundancy. You can use the Ethernet multichassis link aggregation (MLAG) feature to aggregate the physical links between the IRF fabric and its upstream or downstream devices across the IRF members.

With IRF, switches can be clustered across tiers to collapse and dramatically simplify the architecture. You can eliminate the distribution layer and unnecessary network hops. This two-tier architecture provides more direct, higher capacity connections between users and networks services.

IRF delivers a network-based, in-service-software-upgrade capability that allows an individual IRF-enabled switch to be taken offline for servicing or software upgrades without affecting traffic going to other switches in the IRF domain.

IRF and meshed stacking

In addition to IRF, HPE offers meshed stacking technology. Meshed stacking allows you to aggregate up to five switches to form a fully meshed stack for resiliency and management in a single interface. Direct links run to and from each switch in the stack, forming a single logical switch. IRF and meshed stacking offer many benefits over traditional stacking:

- **Unified management**—You can manage the stack through a single master switch.
- **High availability**—IRF and meshed stacking provide N:1 failover and redundant links.
- **Increased performance**—All available links remain active and provide load balancing, which increases efficiency in switching and routing.
- **Scalability**—You can increase network bandwidth by adding switches to the meshed stack or IRF system.
- **Flattened architecture**—By enabling access layer switches to share highly available links to the core, meshed stacking and IRF help customers create low-latency, two-tier architectures in the data center.



Note

Mesh stacking is only applicable on a very small subset of the ArubaOS-Switch product line.

Forming link aggregation connections

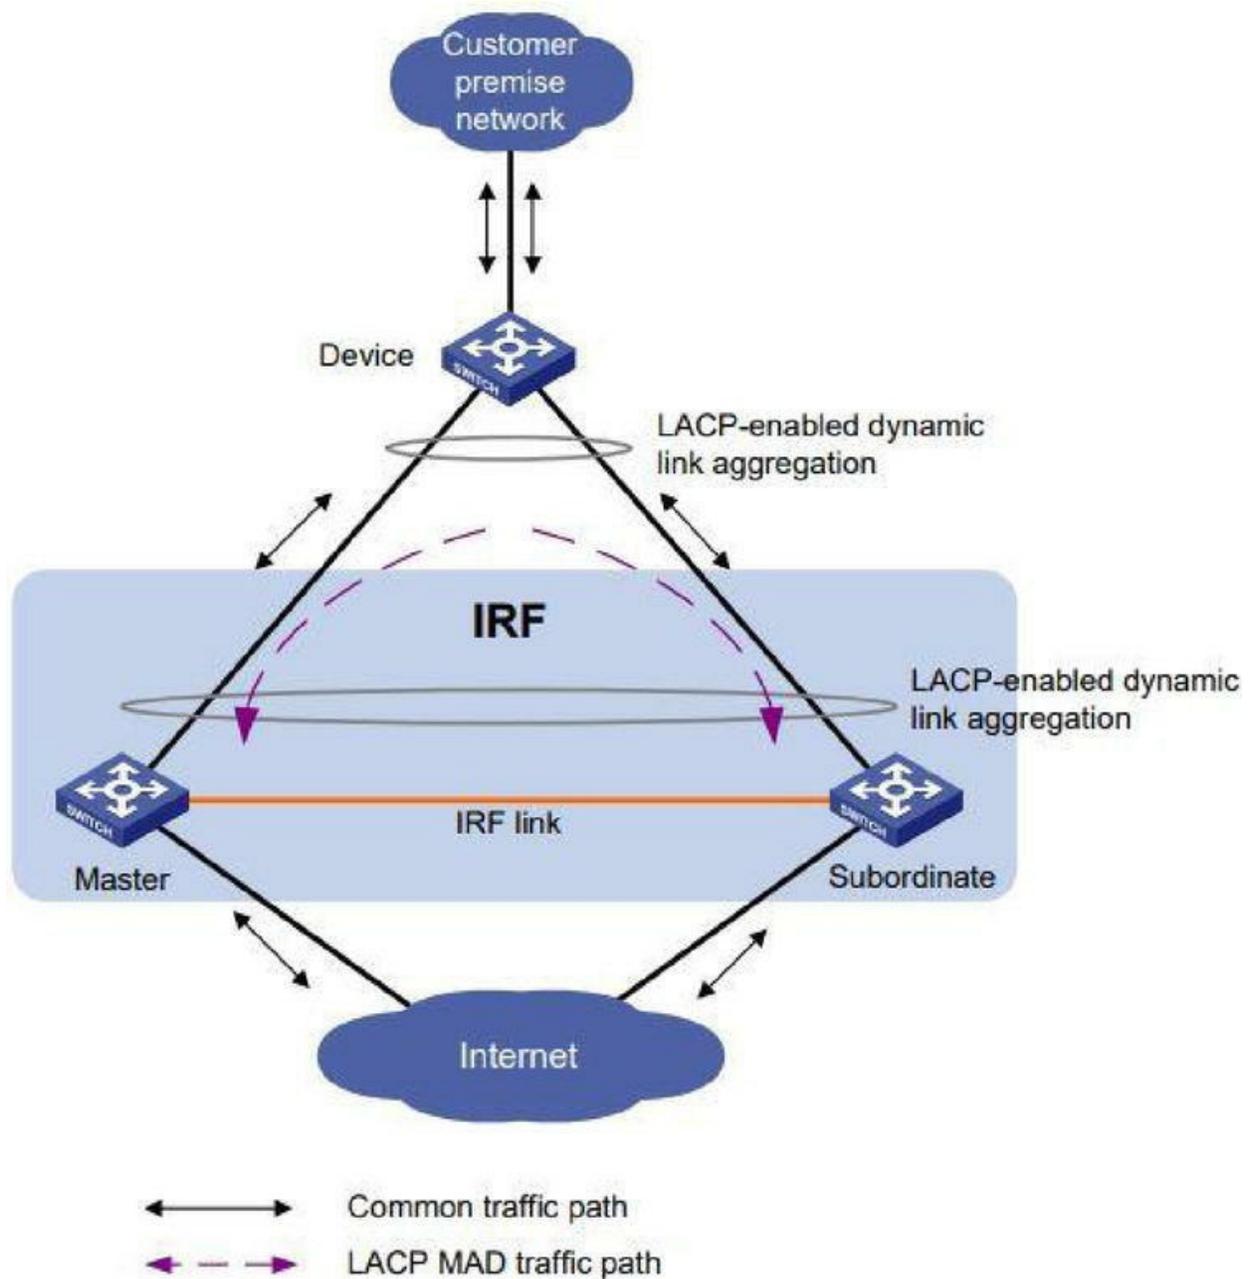


Figure 3-13 LACP multi-active detection (MAD) application scenario

IRF and LACP used together, as shown in [Figure 3-13](#), can further boost performance by bundling several parallel links between switches and servers, allowing scalable on-demand performance and capacity to support critical business applications. Should a network failure occur, IRF can deliver rapid recovery and network reconvergence in less than 50 milliseconds.



Important

When you are establishing a link aggregation to a server, remember that not all server operating systems support LACP. For example, VMware ESX/ESXi 4.0, 4.1, and 5.0 only support static link aggregation (no LACP). LACP support was added in 5.1.



Note

LACP MAD is supported only on some devices with appropriate firmware to support the passing of the LACP MAD packets.

Link aggregation connections can be formed manually or dynamically. A manual aggregated link is configured and maintained manually by the administrator. It recognizes only those ports that are configured as part of the link. When formed manually, it is up to the administrator to define the ports that will be part of the bundle of aggregated links on two connected devices; no special messages are exchanged between the switches concerning the aggregated links.

One advantage of the manual approach is that it is simple to configure. However, the downside of this approach is that you can only configure eight physical interfaces in the aggregation link. If one fails, you lose one-eighth of the bandwidth. An even greater disadvantage is that the manual link aggregation has no protections against misconfigurations. If you add the wrong links to the aggregation—for example, links that connect to different switches—you can cause dangerous loops and connectivity issues.

A key advantage of LACP is that LACP verifies that the ports in common between the two switches have compatible settings, removing the likelihood of a misconfigured aggregated link. LACP also verifies that the links all connect to the same link aggregation on the same peer switch, removing the possibility of loops resulting from misconfiguration. The main disadvantage of LACP is that it is more complex to configure and typically harder to troubleshoot when two devices do not successfully establish an aggregated link.

Dynamic LACP supports standby physical links, which provide additional

failover if a functioning physical link within the aggregated link becomes unavailable. Standby physical links are typically not counted in the maximum allowed number of physical links for an aggregated link. This allows you to set up an aggregated link with the maximum number of physical links allowed by the switch and designate standby links. Therefore, you can ensure maximum bandwidth even when links fail.

However, a disadvantage of dynamic LACP is that the established trunk joins only one VLAN (the untagged VLAN). Most aggregated links are also links on which you might want to carry traffic from multiple VLANs. This cannot be done on dynamic LACP links without also implementing a complicated GARP VLAN Registration Protocol (GVRP) configuration, which can open security holes.

Static LACP does not have this disadvantage. Therefore, in practice, network admins almost always prefer static LACP configurations.

Link Layer Discovery Protocol (LLDP) is used at Layer 2 to send periodic announcements about:

- Major device functions (such as bridging and routing)
- Management IP address
- Device identifier (such as sysname, MAC address, and chassis ID)
- Port identifier (port ID and description)

To display LLDP information from a neighborhood device, use the command:

```
display lldp neighbor-information
```

To see more information about a switch connected to a specific interface, use:

```
display lldp neighbor-information int hundred 1/0/1 verbose
```

LACP operational modes

LACP manages the dynamic link negotiation of an aggregated connection: Switches use LACP data units (LACPDU) to exchange information and establish a dynamic aggregated link. These LACPDU include information such as the source media access control (MAC) address of the device in

addition to identifiers and priorities:

- The **system identifier** has two parts: the LACP system priority and the switch's MAC address.
- The **LACP system priority** determines which switch will select the ports that are active in the aggregated link. Smaller numbers have a higher priority.
- The **port priority** is used in the process of determining which ports are active in the aggregated link. Smaller numbers have a higher priority.

Exchanging LACPDU allows devices to determine if the links can be aggregated. For example, devices can determine whether all the links are the same media type and speed. LACPDUs also allow the devices to manage the logical aggregated link, including adding or removing physical links and handling failovers.

An LACP port can operate in one of two states:

- **Active**—Active ports transmit LACPDUs to advertise that they can create aggregated links.
- **Passive**—Passive ports listen for LACPDUs. If passive ports receive an LACPDU from an active port, they will respond with their own LACPDU.

For an aggregated link to be established, either both sides need to be in an active state, or one side is in an active state and the other in a passive state. If both sides are in a passive state, then neither will initiate the exchange process.

In static LACP, both sides operate in active mode—they both send LACPDUs to each other. In dynamic LACP, one or both sides actively send LACPDUs. One side can be in a passive mode, waiting for the active side to initiate the process.

In dynamic aggregation mode, both switches need to be configured to use dynamic link aggregation mode; in static aggregation mode, both switches need to be configured as static.



Important

Do not confuse static LACP with a manual link aggregation. The first uses LACP; the second does not use a protocol.

Dynamic bridge-aggregation group

To ensure ports are automatically configured with the applicable bridge-aggregation interface configuration on FlexFabric switch, the following steps need to be performed when creating a dynamic bridge-aggregation group:

1. Create the bridge-aggregation interface.
2. Set the link-aggregation mode within the bridge-aggregation interface.
3. Add the desired physical interfaces to the link-aggregation group.
4. Configure the bridge-aggregation interface with the desired VLANs within the bridge-aggregation interface.

Link aggregation load sharing

HPE FlexFabric switches provide options for load sharing, such as MAC and IP addresses and TCP or User Datagram Protocol (UDP) ports. By default, these switches use the Layer 3 option for load balancing. This option uses a hash of the source and destination IP address to select the physical interface used for IP traffic. It uses a hash of source and destination MAC addresses for non-IP traffic.

You can also choose Layer 4 load sharing, which determines load sharing as follows:

- If the packet protocol is an IP packet and has Layer 4 port information, use the source and destination IP address as well as the source and destination UDP and TCP port. Only nonfragmented packets will have their TCP/UDP port number used by load sharing. This ensures that all frames associated with a fragmented IP packet are sent through the same trunk on the same physical link.
- If the packet protocol is an IP packet and does not have Layer 4 information, use the source and destination IP address.

- If the packet is not an IP packet, use the source and destination MAC address.

The load sharing process is local to a switch. You can implement different load sharing solutions on two devices that are connected. In addition, even when two connected devices use the same type of load sharing, traffic that travels across one link is not guaranteed to return on that link. It depends on the information that is being hashed as well as the load sharing algorithm implemented on the local switch.

Configuring IRF devices

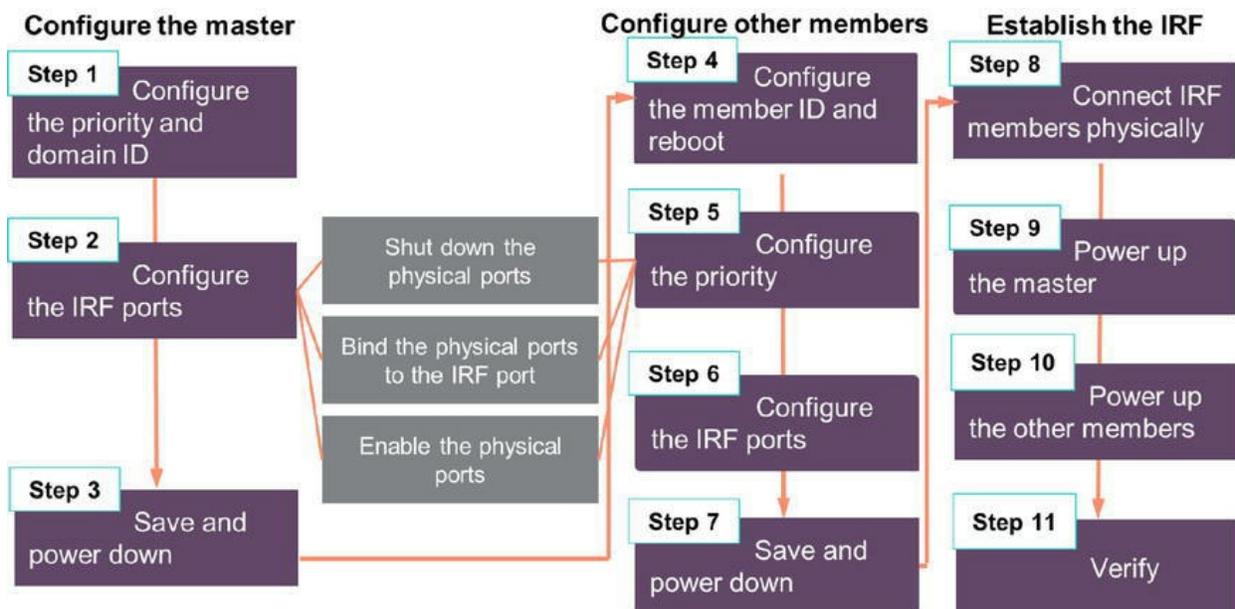


Figure 3-14 Eleven steps are involved in configuring IRFs

As shown in [Figure 3-14](#), IRF virtual devices are configured in a high-level three-stage process:

1. Configure the master.
2. Configure the other members.
3. Establish the IRF device, and activate the IRF port configuration.

A logical IRF port is a logical port dedicated to the internal connection of an

IRF virtual device. These ports cannot act as access, trunk, or hybrid ports. An IRF port is effective only when it is bound to a physical IRF port. Typically, an Ethernet port or optical port forwards frames to the network. When a physical port is bound to an IRF port, it acts as a physical IRF port and forwards data traffic such as IRF-related negotiation frames and data traffic among members.

An IRF stack can have a daisy chain topology or a ring topology. A ring connection is more reliable than the daisy chain connection. In a daisy chain topology, the failure of one link can cause the IRF virtual device to partition into two independent IRF virtual devices, which can disrupt connectivity as well as IRF functioning. The failure of a link in a ring connection results in a daisy chain connection and does not affect IRF services.

Example of LACP configuration

Aruba virtualization (and HA) solutions for access and core

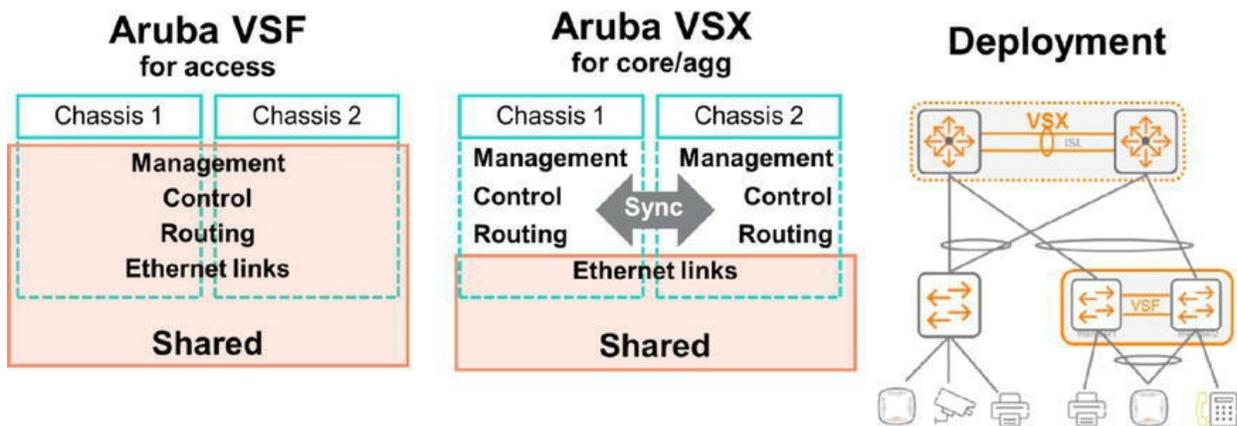


Figure 3-15 Aruba virtualization and HA solutions for access and core

Virtual Switching Framework (VSF) lets two Aruba switches function as a single virtual switch called a VSF fabric. One logical/virtual switch is distributed physically in two chassis, and two switches can be managed as a single entity. VSF, therefore, introduces many benefits in management simplicity, redundancy at both Layers 2 and 3, operational efficiency, and design simplicity. The VSF fabric can also provide high availability and fast

failover at both Layers 2 and 3 for servers connected with multiple network interface cards (NICs). VSF is used in many ArubaOS-Switch devices.

VSF, as illustrated in [Figure 3-15](#), provides all the advantages of a completely link-aggregation-based Layer 2 design:

- Faster failover
- Load-balancing
- Simple set up and design
- Better stability and less need for troubleshooting

VSF fabric looks like a modular switch that has been extended over two chassis, where all switches have:

- A management plane, which provides administrators and management solutions access to management interfaces such as the command-line interface (CLI)
- A control plane, which runs protocols and builds routing and MAC forwarding tables
- A forwarding plane, which processes and forwards packets

A modular switch has a management module that supports the management and control plane while its interface modules provide the forwarding plane. These interface modules also proxy portions of the control plane. If the modular switch has two management modules, one management module is active while the other is on standby.

A VSF fabric also has only one active management plane and one active control plane, which are provided by the management module in the commander switch. The VSF commander proxies the control plane to the standby member. Both members of the fabric participate in the forwarding plane. Interface modules from both members are combined as if over one large switch connected by the VSF link.

Aruba Virtual Switching Extension (VSX) is a virtualization technology for aggregation/core switches running the ArubaOS-CX operating system. This solution lets the switches present as one virtualized switch in critical areas. Configuration synchronization is one aspect of this VSX solution where the primary switch configuration is synced to the secondary switch. This solution

allows for a pseudo single pane of glass configuration and helps keep key configuration pieces in synchronization as operational changes are made. Since the solution is primarily for high availability, it is expected that most of the configuration policy is the same across both peers.

VSX virtualizes the control plane of two aggregation switches to function as one device at Layer 2 and as independent devices at Layer 3. From a datapath perspective, each device does an independent forwarding lookup to decide how to handle traffic. Some of the forwarding databases, such as the MAC forwarding database and neighbor tables, are synchronized between the two devices using a proprietary VSX control plane. Some of the forwarding databases are built independently by each switch.

Benefits of VSX:

- Control plane:
 - Dual control plane for better resiliency
 - Unified management (synchronized configuration and easy troubleshooting)
 - Independently software upgradable with near zero downtime
 - In-chassis redundancy for the 8400 series switches and device-level redundancy for all other platforms, such as for the 8320 series switches
- Layer 2 distributed link aggregations, or LAGs (aggregation switches to access switches):
 - Loop-free L2 multipathing (active-active)
 - Rapid failover
 - Simple configuration
 - No spanning tree required
- Layer 3 distributed LAGs (core switches to aggregate switches)
 - Distributed Layer 3 over VSX pair (various options: Routed Only Ports [ROPs], Switched Virtual Interfaces [SVIs], or LAG SVIs)
 - Unified datapath (active-active first hop gateway)

- Layer 3 equal-cost multi-path routing (ECMP) and Layer 2 VSX (highly fault tolerant)
- Active gateway:
 - Active-active first hop gateway (Virtual IP)
 - Simple configuration (one command)
 - No gateway protocol overhead
 - DHCP relay redundancy

Switch infrastructure recommendations—iSCSI SAN

There are a set of minimum switch capabilities that make building a high-performance, fault-tolerant storage network a relatively easy and cost-effective task. Any enterprise-class managed switch typically has the necessary capabilities most IP-SAN customers require. HPE recommends the HPE FlexFabric 5700, 5900 (which includes the 5920 and 5930 family switches for ToR convergence), and the FlexFabric 7900 and 12900 Switch Series for core converged switches. In an iSCSI packet-switching network, a hop is the trip a data packet takes from the initiator (host) to the target (storage system). A hop is one portion of the path between the source and the destination.

Each time packets are passed to the next device, a hop occurs. On an iSCSI network that uses TCP/IP, the number of hops a packet has taken toward its destination (called the hop count) is kept in the packet header. A packet with an exceedingly large hop count is discarded, and this should be avoided. This hop count value is a measure of distance and is also referred to as the diameter of the IP SAN fabric.

Known as time to live (TTL) in IPv4, and hop limit in IPv6, the field specifies a limit on the number of hops a packet is allowed before being discarded. A fixed TTL value must be at least big enough for the network “diameter,” that is, the longest possible path. The purpose of the TTL field is to avoid a situation in which an undeliverable datagram keeps circulating on a system and such a system eventually becoming swamped by such “immortals.”

A hop count of n means that n switches separate the source host from the destination. By itself, this metric is not useful for determining the optimum network path, as it does not take into consideration the speed, load, reliability, or latency of any particular hop. Since store and forward and other latencies are incurred through each hop, a large number of hops between source and destination imply higher latency and lower real-time iSCSI performance.

A FlexFabric iSCSI SAN does not have a hard hop limit, but keep in mind that latency can be a cause of slow IO performance and iSCSI disconnects. Of course, many other things can add to latency, including hops, congestion, and distance. HPE recommends a best practice iSCSI SAN network diameter (or hop count) below seven and as small as possible. Consider edge-core or edge-core-edge topologies instead of a many hop network topology. A small iSCSI SAN diameter, or low hop count, enables the spanning-tree algorithm to have enough time to converge. Also, the bridge protocol data unit (BPDU) needs time to be propagated across the network. BPDUs are data messages that are exchanged across the switches within an extended network that uses a spanning tree protocol topology. BPDU packets contain information on ports, addresses, priorities, and costs and ensure that the data ends up where it was intended.

The spanning-tree default timers account for seven hops:

- **Max-age time:** 20 seconds
- **Forward delay time:** 15 seconds
- **Hello time:** Two seconds

Minimum recommended switch capabilities for HPE storage-based iSCSI SAN

Although there are still many 1 GbE servers operating in the field, 10 GbE server connections are being widely deployed and will keep increasing. Additionally, 40 GbE links provide good switch-to-switch bandwidth and performance needed in converged networks. Carefully consider the fan-in ratio used (the number and speed of server ports versus the number and speed of uplink ports) to ensure that switch-to-switch links are not the cause of excessive congestion.

To achieve maximum performance on the SAN, it is important to select a switch that has a fully subscribed backplane. This means that the switch should be capable of operating at line rate for all packet sizes.

The nature of iSCSI and TCP means that the network switches in an iSCSI SAN receive bursts of packets. Network switches should support adequate buffers that are able to store the bursts of traffic while flow control has paused transmissions.

For optimal switch performance, HPE recommends that the switch have at least 512 KB of buffer cache per port. Consult switch specifications for the total buffer cache. If the switch aggregates cache among a group of ports (that is, 1 MB of cache per eight ports), space the utilized ports appropriately to avoid cache oversubscription.

Configuration of end-to-end data center bridging (DCB) iSCSI mitigates the need to utilize the deep buffer 5920, but does not preclude it. Deep buffering is useful when flow control is not implemented in the network, but it adds latency that limits throughput. The FlexFabric family of switches utilize smartbuffer technology to reduce the need for shared buffer memory. For example, without smart-buffer technology a 48-port switch should have 24 MB of buffering. With smart-buffer technology, a 48-port switch with 12 MB of buffering and lossless flow-control enabled performs as well as if it had 24 MB of buffering.

IP storage networks are unique in the amount of sustained bandwidth that is required to maintain adequate performance levels under heavy workloads. When a network port becomes saturated, excess frames may be dropped because the port cannot physically handle the amount of traffic it is receiving. This causes the packets to be resent, and the overhead of resending the packets can add to the decreased performance.

Ethernet Pause (flow control) technology can eliminate this problem by controlling the speed at which data is sent to the port. IEEE 802.1bb priority-based flow control enhances a converged iSCSI/LAN network by pausing storage traffic that should remain lossless and allowing the converged LAN traffic to buffer, queue, pass, or drop. This ensures storage packets do not get dropped while LAN traffic still follows the inherent nature expected of TCP-based traffic. Some switch manufacturers do not recommend configuring pause flow control when using jumbo frames or vice versa. HPE recommends

implementing flow control over jumbo frames for optimal performance. Flow control is required when using the HPE Device Specific Module (DSM) and multipath IO (MPIO).

All ports on the switch, servers, and storage nodes should be configured to auto-negotiate duplex and speed settings. Although most switches and NICs auto-negotiate the optimal performance setting, if a single port on the IP storage network negotiates a sub-optimal (100 megabit or less and/or halfduplex) setting, the entire SAN performance can be impacted negatively. Make sure to check each switch and NIC port to make sure the auto-negotiation resolved to be 1 Gb/s or 10 Gb/s with full duplex. Additionally, carefully consider uplink bandwidth in a network.

Link aggregation and/or trunking support is important to enable when building a high-performance, fault-tolerant IP storage network. HPE recommends implementing link aggregation and/or trunking technology when using switch-to-switch trunking and iSCSI target NIC balancing and aggregation.

Implementing a separate subnet or VLAN for the IP storage network is a best practice. If implementing VLAN technology within the switch infrastructure, typically you need to enable VLAN tagging (802.1q) and/or VLAN trunking (802.1q). VLANs are used to limit the scope of broadcast domains in an Ethernet network. Access ports are ports that carry untagged frames inside the network, while trunk ports carry multiple VLANs. Trunk ports can contain one VLAN, which will be carried without VLAN tag. Such a VLAN is usually called Port VLAN ID (PVID). To carry a VLAN, trunk port needs to be explicitly configured to permit a VLAN. The following configuration, as shown in [Figure 3-16](#), only allows VLAN 10 to be allowed through the port as tagged, and VLAN 1 is configured as untagged (PVID):

```
Ten-GigabitEthernet1/0/1
Current state: UP
Line protocol state: UP
IP packet frame type: Ethernet II, hardware address: d894-03fb-793d
Description: Ten-GigabitEthernet1/0/1 Interface
Bandwidth: 10000000 kbps
Loopback is not set
Media type is stack wire, port hardware type is STACK_SFP_PLUS
10Gbps-speed mode, full-duplex mode
Link speed type is autonegotiation, link duplex type is autonegotiation
Flow-control is not enabled
Maximum frame length: 10000
Allow jumbo frames to pass
Broadcast max-ratio: 100%
Multicast max-ratio: 100%
Unicast max-ratio: 100%
PVID: 1
MDI type: Automdix
Port link-type: Trunk
VLAN Passing: 10
VLAN permitted: 10
Trunk port encapsulation: IEEE 802.1q
Port priority: 0
```

Figure 3-16 Trunk port configuration example

To build a fault-tolerant IP storage network, multiple switches are typically connected into a single Layer 2 (OSI model) broadcast domain using multiple interconnects. To avoid Layer 2 loops, the Spanning Tree protocol (802.1D) or Rapid Spanning Tree protocol (802.1w) must be implemented in the switch infrastructure. Failing to do so can cause numerous issues on the IP storage networks, including performance degradation or even traffic storms. Use of Intelligent Resilient Fabric (IRF) technology allows users to create large logical switches using multiple physical switches. Implementing IRF can eliminate the need for loops in a network, allowing for all active/active paths.

Sequential read and write, or streaming workloads can benefit from frames larger than 1514 bytes. In fact, for the majority of 10 G or higher network infrastructure situations, jumbo frames are beneficial. A best practice is to always configure jumbo frames on the switch infrastructure in an iSCSI SAN. For example, it is recommended to use the default 10 k byte jumbo frames on the 5900 Switch Series, especially when using Ethernet-based storage access with vMotion and with other high throughput applications.

After jumbo frames are enabled on the switch infrastructure, you can leave it

to the initiator and target configurations to actually utilize jumbo frames, and if so, what maximum size to use. On the 3PAR StoreServ 8000, 20000, with the new iSCSI cards, it is recommended to enable jumbo frames on the storage arrays in addition to the switch infrastructure. Multi-initiator environments using 9 k jumbo frames can actually starve out devices. Better performance in those multi-initiator environments can be realized when the iSCSI initiators and targets are configured for 4 K byte (maximum frame size of 4088 bytes) jumbo frames. However, there is really no one-size-fits-all setting for jumbo frames. Depending on the hardware, a frame size that is too large can actually hurt network throughput. In the case of single large file and multiple small files-based transactions, a 4 K frame size gives optimal results. For multiple large files, a 9 K frame size is better suited.

It is important for you to evaluate the effect that different frame sizes have on different file transfers on a case-by-case basis. Some switch manufacturers do not recommend configuring jumbo frames when using flow control, or vice versa. Consult the switch manufacturer documentation for guidance on this issue.

High network latency can be the primary cause of slow IO performance, or worse, iSCSI drive disconnects. It is important to keep network latency on the IP-SAN subnet below two milliseconds. Many factors can contribute to increasing network latency—the most common factors are congestion, distance, and router hops. Keeping deployment topologies simple while configuring the IP-SAN on a single IP subnet with Layer 2 switching helps to lower the network latency.

Learning check

2. Match the virtualization and HA solution with the networking platform.



ArubaOS-Switch platform



Figure 3-17 Aruba logo

The Aruba portfolio includes both wired and wireless solutions. Aruba wireless access points deliver superb Wi-Fi performance to meet the density and performance needs of customer environment. They can be deployed as controller-managed (ArubaOS) or controllerless (InstantOS) access points (APs).

Aruba's modern, programmable switches easily integrate with network management solutions. These switches come with built-in security features and can integrate with Aruba ClearPass for advanced policy management.

REST APIs and OpenFlow support enables automation of network operations, monitoring, and troubleshooting. Aruba's Layer 3 switches are also capable of leveraging user and port-based traffic tunneling to apply policies, advance services, and encrypt traffic to secure the network.

Aruba's campus switching portfolio spans the access layer to the core and creates the foundation of the new digital workplace. Aruba's switches integrate seamlessly with Aruba's industry-leading security and network management solutions.

Aruba networking devices offer:

- Unified policy and security
- Wireless optimization
- Manageability
- Cloud-based management

- Automation
- Investment protection



Note

Not all Aruba devices and their features are suitable for data center.

Aruba switching portfolio: Campus and data center

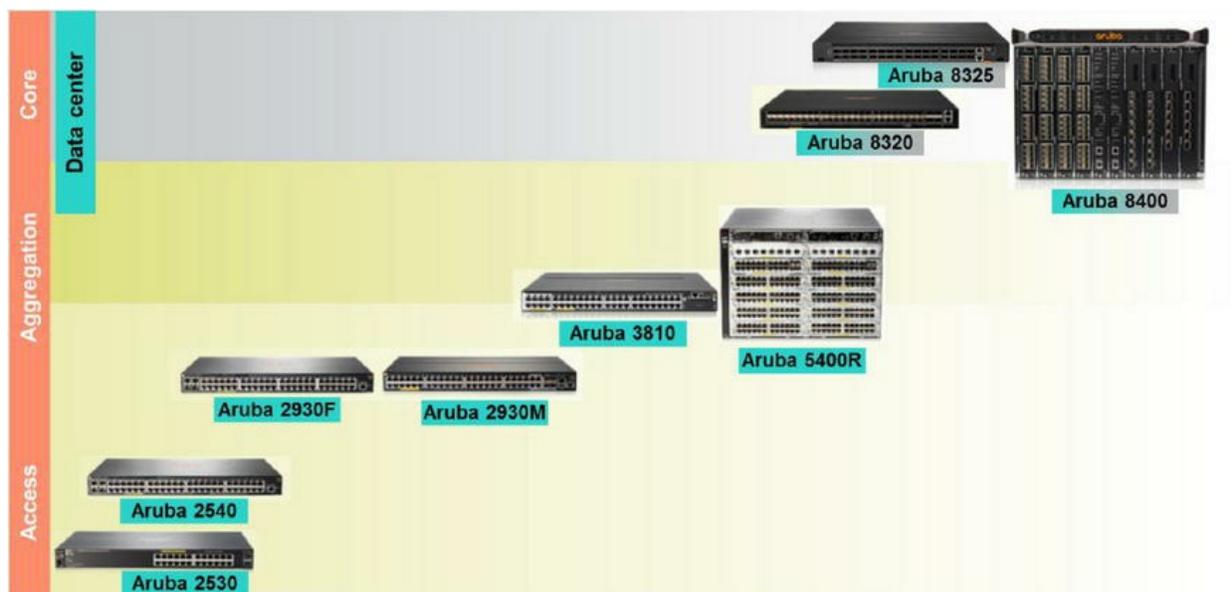


Figure 3-18 Aruba switching portfolio

Aruba switching portfolio, as illustrated in [Figure 3-18](#), includes access, aggregation, core devices, and data center installations:

- Aruba 2530
- Aruba 2540
- Aruba 2930F and Aruba 2930M
- Aruba 3810
- Aruba 5400R
- Aruba 8320 and Aruba 8325

- Aruba 8400

Aruba 2530 and 2540 Switch Series



Figure 3-19 Aruba 2530 and Aruba 2540 Switch Series

The Aruba 2530 Switch Series, as shown in [Figure 3-19](#), provides cost-effective, reliable, and secure access layer connectivity for enterprises, branch offices, and small and midsize businesses. The switches include a limited lifetime warranty.

These fully managed switches deliver Layer 2 capabilities with enhanced access security, traffic prioritization, sFlow, and IPv6 host support. Right-size deployment is available with a range of Gigabit and Fast Ethernet models including compact and fanless models that are ideal for use in quiet work spaces.

The 2540 has wire-speed backhaul bandwidth capacity with built-in 10 GbE uplinks, robust QoS, static/RIP routing, and IPv6. It also includes a limited lifetime warranty with no software licensing required.

Other features include:

- Access control lists (ACLs), sFlow, and IPv6
- Choice of eight, 24, 48 ports 10/100 and gigabit switches
- AirWave and ClearPass Policy Manager integration
- Ready for the software-defined network with REST APIs
- Energy savings with EEE, quiet fanless models, and low power mode settings

- Cloud manageable
- Zero Touch Provisioning

Aruba 2930F and 2930M Switch Series



Figure 3-20 Aruba 2930F and 2930M Switch Series

The Aruba 2930F and 2930M Switch Series, as shown in [Figure 3-20](#), are designed for customers creating smart digital workplaces that are optimized for mobile users with an integrated wired and wireless approach. These convenient Layer 3 network switches include built-in uplinks and power. They are simple to deploy and manage with advanced security and network management tools.

A powerful Aruba ProVision ASIC delivers performance, robust feature support, and value with programmability for the latest applications. Stacking with VSF provides simplicity and scalability. The 2930F supports built-in 1 GbE or 10 GbE uplinks, PoE+, Access OSPF routing, dynamic segmentation, robust QoS, RIP routing, and IPv6 with no software licensing required.

The flexible 2930M supports wire speed 10 GbE and 40 GbE uplinks, redundant dual modular power supplies for up to 1440 W of PoE, and new models with industry-standard IEEE 802.3bt Class 6 that provide up to 60 W of PoE per port.

The feature rich 2930M supports a robust QoS, RIP, access OSPF routing, Protocol-Independent Multicast (PIM), Virtual Router Redundancy Protocol (VRRP), IPv6, and dynamic segmentation for unified and secure access.

The robust basic Layer 3 feature set includes a limited lifetime warranty.

Aruba 3810 and 5400R Switch Series



Figure 3-21 Aruba 3810 and 5400R Switch Series

The Aruba 3810 Switch Series, as shown in [Figure 3-21](#), provides performance and resiliency for enterprises, SMBs, and branch office networks. With HPE Smart Rate multi-gigabit ports for high-speed access points and IoT devices, this advanced Layer 3 network switch delivers a better application experience with low latency, virtualization with resilient stacking technology, and line rate 40 GbE for plenty of back haul capacity.

A powerful Aruba ProVision ASIC delivers performance, robust feature support, and value with flexible programmability for the latest applications. The 3810 delivers resiliency and scalability via innovative backplane stacking technology and redundant, hot-swappable power supplies all in a convenient 1U form factor. The advanced Layer 2 and 3 feature set includes OSPF, IPv6, IPv4 BGP, dynamic segmentation, robust QoS, and policy-based routing with no software licensing.

The 3810M is easy to deploy, use, and manage using Aruba AirWave or Aruba Central. Aruba ClearPass offers centralized security and external captive portal support. The switches offer a limited lifetime warranty.

The Aruba 5400R zl2 Switch Series, as shown in [Figure 3-21](#), delivers enterprise-class resiliency with innovative flexibility and scalability for customers creating smart digital workplaces that are optimized for mobile users with an integrated wired and wireless approach. This modular series brings scalable aggregation with VSF stacking technology, hitless failover, and a fast software upgrade for 5400R VSF stacks. The advanced Layer 2 and 3 feature set includes OSPF, IPv6, IPv4 BGP, dynamic segmentation, robust QoS, and policy-based routing with no software licensing required.

Based on a powerful ProVision ASIC, the Aruba 5400R zl2 Switch Series has a high-speed, highcapacity architecture with 2 Tbps crossbar switching fabric with low 2.1μ robust feature support, and value with flexible programmability for the latest applications. This series offers flexible connectivity options with a six- or 12-slot compact chassis, line rate 40 GbE, up to 96 line rate Smart Rate multigigabit or 10 GbE ports and up to 288 ports of PoE+ for powering access points, cameras, and IoT devices. The 5400R is easy to deploy, use, and manage using Aruba AirWave or Aruba Central. Aruba ClearPass offers centralized security and external captive portal support. The switches include a limited lifetime warranty.

Learning check

3. Which ArubaOS-Switch series supports IPv4 BGP?
 - A. Aruba 2530
 - B. Aruba 2540
 - C. Aruba 2930F
 - D. Aruba 5400R zl2

ArubaOS-CX platform

ArubaOS-CX is a modern software system for the core that automates and simplifies network tasks using built-in Python interpreter and REST APIs. These industry-leading switches bring intelligence and automation to the core with full programmability and a built-in Aruba Network Analytics Engine, giving network operators the ability to see more, know more, and act faster.

Aruba 8400 Switch Series



Figure 3-22 Aruba 8400 Switch Series

The Aruba 8400 campus core and aggregation switch series, as shown in [Figure 3-22](#), provide a flexible and innovative approach to dealing with the new application, security, and scalability demands of the mobile-cloud and Internet of Things (IoT) era.

It combines a modern, fully programmable OS with carrier-grade hardware and leading performance. It incorporates the industry-first Network Analytics Engine to monitor and troubleshoot network, system, application, and security-related issues easily.

By enabling faster automation and network insights, the operating system reduces the time spent on manual tasks and addresses current and future demands driven by Mobility and IoT.

The 8400 contains a high-speed fully distributed architecture and provides up to 19.2 Tbps switching capacity with up to 8571 billion packets per second (BPPS) for throughput. All switching and routing is performed in the IO modules and meets the demands of bandwidth-intensive applications today

and in the future.

Features of Aruba 8400 include:

- High-speed 32-port 10 GbE module, eight-port 40 GbE module, or six-port 40/100 G module
- Convenient bundles that simplify ordering
- Redundant management, fabric, and power supplies
- Energy-efficient, hot-swappable, power supplies with N+N redundancy with 80 Plus Gold certification

Management features include:

- Out-of-Band management (OOBM)
- Console management ports
- Status LEDs for fans, power supplies, and modules



Note

More information about the Aruba 8400 Switch Series can be found here:

<https://h20195.www2.hp.com/V2/GetPDF.aspx/a00017758enw.pdf>

Aruba 8320 Switch Series



Figure 3-23 Aruba 8320 Switch Series

Aruba 8320, as shown in [Figure 3-23](#), is a high-performance switch with 2.3 Tbps switching capacity. It has 48 ports of 10 Gig and six ports of 40 Gbe—so plenty of capacity for campus core and aggregation. This switch supports two hot-swappable, redundant power supplies that help to ensure high availability that is needed at the aggregation and core layers of a campus network.

It supports the same advanced Layer 3 software features as the 8400: OSPF, Border Gateway Protocol (BGP), virtual routing and forwarding (VRF), Multi-Chassis Link Aggregation (M-LAG) for switch virtualization, and it is designed for usability with easy to replace fans and five fan modules. 8320 provides a five-year hardware warranty with no software licensing required.

There are three models with high-speed connectivity:

- 48p 10 G SFP/SFP+ and 6p 40 G QSFP+ Switch (JL479A)
- 32p 40 G QSFP+ (JL579A)
- 48p 1 G/10 GBASE-T and 6p 40 G QSFP+ (JL581A)



Note

More information about the Aruba 8320 Switch Series can be found here: <https://h20195.www2.hp.com/v2/getdocument.aspx?>

Aruba 8325 Switch Series



Figure 3-24 Aruba 8325-32C and 8325-48Y8C

The Aruba 8325 Switch Series, as shown in [Figure 3-24](#), offers a flexible and innovative approach to addressing the application, security, and scalability demands of the mobile, cloud, and IoT era. These switches serve the needs of the next-generation core and aggregation layer as well as emerging data center requirements at the ToR and End of Row (EoR). They provide over 6.4 Tbps of capacity, with line-rate gigabit Ethernet interfaces including 10 Gbps, 25 Gbps, 40 Gbps, and 100 Gbps.

The 8325 series includes industry-leading line rate ports 1/10/25 GbE (SFP/SFP+/SFP28) and 40/100 GbE (QSFP+/QSFP28) with connectivity in a compact 1U form factor. These switches offer a solid investment for customers wanting to migrate from older 1 GbE/10 GbE to faster 25 GbE, or 10 GbE/40 GbE to 100 GbE ports.

Features of Aruba 8325 include:

- 6.4 Tbps switching capacity in compact 1U form factor
- VSX for high availability
- Advanced Layer 3, including OSPF and BGP
- High hardware scale, including full internet routing
- N+1 redundant, hot-swappable power supplies
- All bundles include hot-swappable, removable fan and power supplies
- REST for distributed or centralized orchestration
- Database-driven ArubaOS-CX architecture for HA and fault tolerance

- 48 ports of 10/25 G with eight ports of 40/100 G, and 32 ports 40/100 G

Most Aruba switches offer several different fan trays depending on the model and airflow direction preferences. Airflow directions include:

- Back (power side) to front (port side) airflow (blue handles)
- Front (port side) to back (power side) airflow (red handles)



Note

More information about the Aruba 8325 Switch Series can be found here: <https://h20195.www2.hp.com/v2/getdocument.aspx?docname=a00056519enw>

Aruba Virtual Switching Extension

Aruba Virtual Switching Extension (VSX) is virtualization technology for aggregation/core switches running the ArubaOS-CX operating system. This solution lets the switches present as one virtualized switch in critical areas. Configuration synchronization is one aspect of this VSX solution where the primary switch configuration is synced to the secondary switch. This solution allows for a pseudo single pane of glass configuration and helps keep key configuration pieces in synchronization as operational changes are made. Since the solution is primarily for high availability, it is expected that most of the configuration policy is the same across both peers.

VSX virtualizes the control plane of two aggregation switches to function as one device at Layer 2 and as independent devices at Layer 3. From a datapath perspective, each device does an independent forwarding lookup to decide how to handle traffic. Some of the forwarding databases, such as the MAC forwarding database and neighbor tables, are synchronized between the two devices using a proprietary VSX control plane. Some of the forwarding databases are built independently by each switch.

VSX components include:

- **Inter-Switch Link (ISL)**—This is a Layer 2 interface between two VSX peer switches.
- **Keepalive**—This is a Layer 3 interface that is used to exchange

heartbeats between VSX peer switches.

- **Switch roles**—Each VSX switch must be configured with a role (primary or secondary).
- **IGMP snooping**—VSX switches can be configured for IGMP snooping on downstream VLANs facing the access switches.
- **DHCP relay redundancy**—Only the primary switch actively relays DHCP client requests to the server; the secondary switch acts as a backup.



Note

More information about VSX can be found here:
https://support.hpe.com/hpsc/doc/public/display?docId=emr_na-a00052286en_us

The following is an example of a VSX configuration, after the necessary VLAN/VLANs, VLAN interfaces, and ISL are created:

```
switch(config)# vsx
switch(config-vsx)# role primary
switch(config-vsx)# inter-switch-link 1/1/1
switch(config-vsx)# keepalive peer 10.0.0.1 source 10.0.0.2
switch(config-vsx)# exit
switch(config)# interface vlan 10
switch(config-if-vlan)# active-gateway ip 10.1.1.1 mac 00:11:22:33:44:55
```

The example is not complete; additional steps are usually needed to configure routing and other configuration.

Learning check

4. Recommend a ToR switch with 25 GbE downlinks.

- A. Aruba 2530
- B. Aruba 2540
- C. Aruba 2930F
- D. Aruba 5400R zl2

HPE FlexFabric

HPE FlexFabric is designed to address more traditional data center networking architectures and requirements. It offers a full portfolio of data center switches offering a complete range of fixed port and modular edge switches and resilient modular core switches. HPE FlexFabric addresses more traditional customer requirements such as LAN and storage area network (SAN) convergence, Multiprotocol Label Switching (MPLS), and support for a wide variety of networking protocols. The automated, programmable FlexFabric data center architecture gives instant access to cloud-based apps and services, so a data center can grow and adapt to the business.

HPE FlexFabric portfolio provides an open approach to data center network fabric. The FlexFabric enables you to create a more converged, flexible, and scalable data center network architecture.

A converged data center network runs both network and storage traffic over the same infrastructure. HPE converged infrastructure solutions tie virtualization, automation, and unified infrastructure management software together into prebuilt, tested, and workload-optimized systems. These systems are software-defined for easy integration into existing infrastructure and quick transition to hybrid cloud delivery models.

FlexFabric is a highly scalable data center fabric architecture that enables you to provision network resources efficiently and securely to accelerate deployment of virtualized workloads. With FlexFabric, customers are never bound to a specific operating configuration, proprietary architecture, or network fabric. The HPE automated, programmable data center fabric architecture gives instant access to cloud-based apps and services, so the data center can grow and adapt to the business.

With highly scalable platforms and advanced networking and management technologies, FlexFabric network designs are simpler and easier to manage and grow over time. This open architecture uses industry standards to simplify server and storage network connections. It provides seamless interoperability with existing core data center networks. FlexFabric combines intelligence at the server edge with a focus on centrally managed connection policy management to enable virtualization-aware networking and security; predictable performance; and rapid, business-driven provisioning of data center resources.

Most HPE FlexFabric switches offer several different fan trays depending on the model and airflow direction preferences. Airflow directions include:

- Back (power side) to front (port side) airflow
- Front (port side) to back (power side) airflow

HPE FlexFabric 5710 Switch Series



Figure 3-25 HPE FlexFabric 5710 Switch Series

HPE FlexFabric 5710 Switch Series, as shown in [Figure 3-25](#), is a family of high-performance, lowlatency access switches aimed at providing superior edge device connectivity in modern spine leaf data centers.

HPE FlexFabric 5710 Switch Series is ideally suited for deployment at the server access layer. It delivers lower total cost of ownership (TCO) while enhancing networking performance to support demanding virtualized applications and server-to-server traffic. Resilience and ease of management

come hand in hand with the FlexFabric 5710.

HPE FlexFabric 5710 Switch Series is a great choice when proving 10 Gb connectivity for small iSCSI storage solutions.

	HPE 5710 48SFP+ 6QS+/2QS28 Switch (JL585A)	HPE 5710 48XGT 6QS+/2QS28 Switch (JL586A)	HPE 5710 24SFP+ 6QS+/2QS28 Switch (JL587A)
IO ports and slots	48 x 1/10 GB SFP+ ports 6 x QSFP+ ports (or a maximum of 2 x QSFP28 ports)	48 x 1/10 GBASE-T ports 6 x QSFP+ ports (or a maximum of 2 x QSFP28 ports)	24 x 1/10 GB SFP+ ports 6 x QSFP+ ports (or a maximum of 2 x QSFP28 ports)
Additional ports and slots	Management ports 1 x 10 M/100 M/1000 MBASE-T copper port 1 x SFP port Console ports 1 x mini USB console port 1 x serial console port	Management ports 1 x 10 M/100 M/1000 MBASE-T copper port 1 x SFP port Console ports 1 x mini USB console port 1 x serial console port	Management ports 1 x 10 M/100 M/1000 MBASE-T copper port 1 x SFP port Console ports 1 x mini USB console port 1 x serial console port

HPE FlexFabric 5710 technical specifications

HPE FlexFabric 5940 Switch Series



Figure 3-26 HPE FlexFabric 5940 Switch series

The HPE FlexFabric 5940 Switch Series, as shown in [Figure 3-26](#), is a family of high-performance and low-latency 10 GbE, 40 GbE ToR data center switches. The switch series also includes 100G uplink technology and a two-slot and four-slot modular form factor, providing ultimate flexibility for ever-changing data center requirements. This entire series is part of the HPE FlexFabric data center solution, which is a cornerstone of the FlexNetwork architecture.

The FlexFabric 5940 Switch Series is ideally suited for deployment at the aggregation or server access layer.

With the increase pace of deploying virtualized applications, adopting software-defined networking, and the server-to-server traffic, many data centers now require spine and ToR switch innovations that will meet their requirements. The HPE FlexFabric 5940 is optimized to meet the increasing requirements for higher-performance server connectivity, convergence of Ethernet and storage traffic, the capability to handle virtual environments, and low latency.



Note

More information about the HPE FlexFabric 5940 Switch Series can be found here:

<https://h20195.www2.hp.com/v2/GetDocument.aspx?docname=c05158726>

HPE FlexFabric 5945 Switch Series



Figure 3-27 HPE FlexFabric 5945 48SFP28 8QSFP28 Switch and HPE 5945 four-slot Switch

HPE FlexFabric 5945 Switch Series, as shown in [Figure 3-27](#), is a family of high-density, ultra-lowlatency, ToR switches that is part of the HPE FlexFabric solution (from the HPE Cloud-First Reference Architecture). HPE FlexFabric 5945 Switch Series provides choices that fit your customer's IT budget via unique modular options.

The HPE FlexFabric 5945 Switch Series enables customers to scale their server-edge 10/25/100 GbE ToR deployments with high-density ports

delivered in a 1U design:

- 48 x 10/25 GbE (SFP or BASE-T) with 6 x 40 GbE ports
- 48 x 10 GbE (SFP or BASE-T) with 6 x 100 GbE ports
- 32 x 40 GbE ports

Ideally suited for deployment at the aggregation or server access layer of large enterprise data centers, the HPE FlexFabric 5945 Switch Series is also powerful enough for deployment at the core layer of small and medium-sized enterprises. It supports modern spine and leaf data center solutions that require demanding throughput and modern feature support needed for today's virtual environments.

With the increase in virtualized applications and server-to-server traffic, customers require spine and ToR switches that can meet their throughput requirements. With the HPE FlexFabric 5945, data centers can now support up to 100 Gb per port using QSFP28 transceiver, allowing high-performance server connectivity and the capabilities to handle virtual environments. This is available in the low-latency HPE FlexFabric 5945 Switch Series.



Note

More information about the HPE FlexFabric 5945 Switch Series can be found here:

<https://h20195.www2.hpe.com/v2/getdocument.aspx?docname=a00047323enw>

Learning check

5. Which transceiver type supports 100 Gb?
 - A. QSFP28
 - B. QSFP28
 - C. QSFP+
 - D. SFP

HPE Networking management

You can evaluate switches based on their level of manageability:

- Managed switches support SNMP and allow you to configure each port's communication parameters and many other aspects of the switch through a CLI and a GUI, such as a web browser interface.
- Smart-managed switches provide basic capabilities and can be managed through a web browser interface. The web browser interface is designed to be intuitive, making it easy to configure and manage switch features. These switches also support SNMP, so you can manage them through a centralized SNMP console.
- Unmanaged switches provide basic Layer 2 switching and are not configurable. These are plug-and-play switches.

HPE Networking devices can be managed via CLI tools, APIs, or GUIs. Management interfaces can be embedded, installed on-premise, or accessible from cloud.



Note

Unmanaged or smart-managed devices are not typical products for a data center.

Command-line interface

```
[CDC01-Ten-GigabitEthernet1/0/1]port access vlan 5
[CDC01-Ten-GigabitEthernet1/0/1]description Management Lab A
[CDC01-Ten-GigabitEthernet1/0/1]display this
#
interface Ten-GigabitEthernet1/0/1
 port link-mode bridge
 description Management Lab A
 port access vlan 5
#
return
[CDC01-Ten-GigabitEthernet1/0/1]

[CDC01-luser-manage-student]service-type ?
  ftp      FTP service
  http     HTTP service type
  https    HTTPS service type
  pad      X.25 PAD service
  ssh      Secure Shell service
  telnet   Telnet service
  terminal  Terminal access service

#
local-user student class manage
password hash $h$6$o7nQDx0p5prsk$Sc$Pu8q4R7yBU7HjE9L9Px
authorization-attribute user-role network-admin
authorization-attribute user-role network-operator
#
return
```

Figure 3-28 Examples of CLI

CLI management is a traditional and powerful interface for networking devices. CLI can be accessed either locally, via management cable, or remotely using telnet or SSH. Commands differ for various products, so a reference guide is usually needed.

Examples of commands for collecting the diagnostic information from the network switches:

- On the **Comware** platform, use display diagnostic-information to display or save operating statistics for multiple feature modules in the system and to collect log messages.
- On the **ArubaOS-CX** platform, use command show tech to show detailed information about switch features.
- On the **ArubaOS-Switch** platform, use command show tech all to show detailed information about switch features.

These commands display only a subset of available information.

HPE Networking software

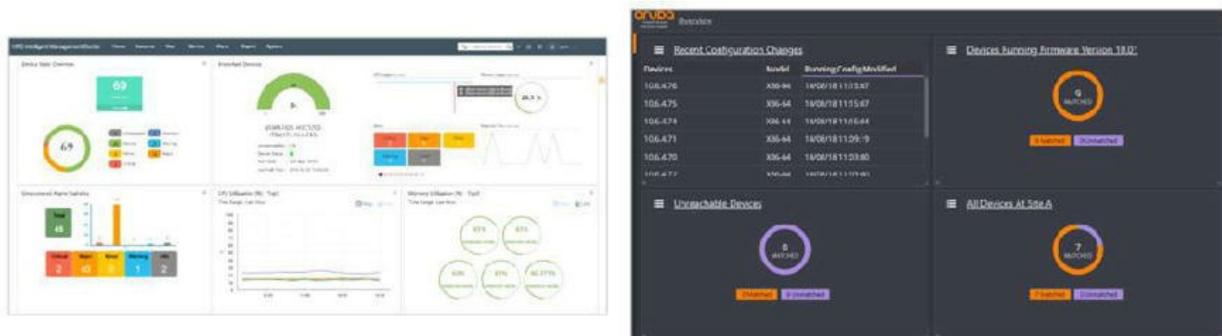


Figure 3-29 Networking software

Alternatively, HPE networking devices, as indicated in [Figure 3-29](#), can be managed using networking software, such as:

- **IMC management**—Campus and data center SW management platform
- **NetEdit**—Empowers IT teams to orchestrate multiple switch configurations with intelligent capabilities, including search, edit, validation (including conformance checking), deployment, and audit

Unified network management with HPE Intelligent Management Center

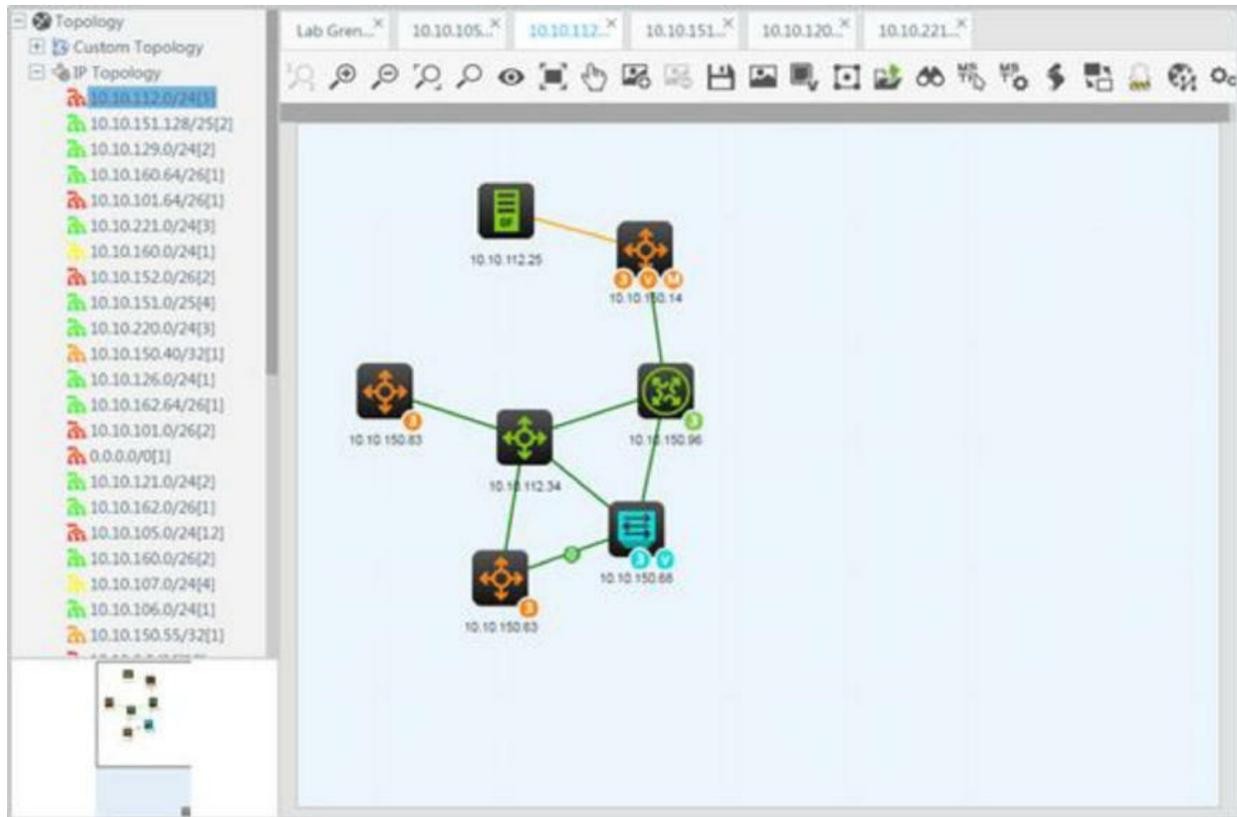


Figure 3-30 IMC

HPE Intelligent Management Center (IMC), as shown in [Figure 3-30](#), is a centralized network management platform that allows you to manage physical networks (both wired and wireless) and virtual networks. You can monitor and manage network traffic and devices from a single interface. Because IMC supports both HPE and third-party network devices, including Cisco, you can manage a heterogeneous environment.

Customers can manage the entire HPE networking solution (except for ArubaOS-CX or Composable Fabric)—from the data center to the main campus to each branch—from HPE IMC. It uses SNMP among other protocols to discover, monitor, and manage devices. IMC integrates fault management, element configuration, and network monitoring from a central vantage point.

IMC supports physical and virtual device management along with thousands of network devices from dozens of vendors, including Cisco. IMC also integrates with Aruba AirWave, Aruba ClearPass, and HPE OneView.

IMC has a modular architecture, which extends capabilities in areas such as BYOD management, wireless and wired user management, server health monitoring, and many others. HPE BYOD solutions enable customers to find the right balance between users' convenience and the ability to control access to the network. The solutions enable BYOD in three easy steps—onboarding, provisioning, and monitoring.

With the IMC auto-discovery feature, you can locate all devices on the network. IMC categorizes discovered devices into types of network devices (such as switches, routers, servers, access points, and desktops) and maps them on a network topology. You can view the network devices based on IP address or device type. You can also create custom views to make it easier to view and manage devices. In addition, IMC allows you to establish baseline configurations and software images. You can compare configurations, track versions, and establish alerts if configuration changes are made.

The modular architecture of IMC enables you to add management capabilities as needed. These modules are fully integrated into the IMC platform, allowing them to share information and functionalities. Examples of the management capabilities you can add through IMC modules include the following:

- Network Traffic Analyzer (NTA) is a graphical network-monitoring tool that provides real-time information about users and applications consuming network bandwidth. You can use NTA to plan, monitor, enhance, and troubleshoot networks, as well as identify bottlenecks and apply corrective measures for enhanced throughput.
- The Wireless Services Manager unifies the management of wired and wireless networks on the IMC platform. The Wireless Services Manager adds wireless devices to the IMC network topology and allows you to configure and apply policies to these devices. You can configure wireless LANs (WLANs) and use radio frequency heat mapping to plan and adjust wireless coverage.
- To complete some functions, IMC logs into a network device's CLI and

sends commands over a terminal session, just like you do as an admin. Similar to how SNMP GET and WRITE messages occur in the background, the IMC process happens in the background when you use a service such as VLAN management. You simply see that you make a change through the IMC management interface and the change is applied.

- To help you streamline, automate, and orchestrate processes, HPE provides several Virtual Application Networking (VAN) modules. VAN Connection Manager orchestrates VMs and network connections. It also allows you to apply policies with templates and manage VM migrations.
- VAN Resource Manager provides drag-and-drop tools, which you can use to design network services. It helps you provision the network and accelerate the delivery of services and applications.
- VAN Fabric Manager helps you manage large Layer 2 data centers, converged LAN/SAN, and connectivity across multiple data centers.
- When a solution uses SNMP v2c, the SNMP servers and agents include a community string in the messages. SNMP agents only accept GET messages if the community matches their readonly or read-write community string. They only accept WRITE messages if the community matches their read-write community string. SNMPv2 is easy to set up, but it is not secure because the community is included in messages in plain text. Therefore, the community string cannot function as a true password.

IMC for the data center

IMC brings critical functionality for the data center:

- Orchestrates network fabrics.
 - Virtual Extensible LAN (VXLAN) manageability
 - Unified Shortest Path Bridging (SPB), Transparent Interconnection of Lots of Links (TRILL), IRF fabric management
 - Manages across geographically dispersed locations
 - VMware vMotion playback
 - Unified data center bridging (DCB), Fibre Channel over Ethernet (FCoE) management

- Accelerates deployment of services and applications.
 - “Just right” network services tuned to business requirements
 - Simplifies provisioning, monitoring, and troubleshooting of applications
 - Eliminates manual provisioning of network service parameters
 - Easy-to-use service modeling tool with drag and drop user interface (UI)

HPE IMC base platform options

HPE IMC base platform options for HPE and third-party devices are:

- **IMC Basic**—Fixed-device limit of 50 nodes
- **IMC Standard:**
 - Base license for 50
 - Up to 15,000 devices
- **IMC Enterprise:**
 - Large network systems that scale to tens of thousands
 - Base license for 50 managed devices
 - Fifty APs for the included Wireless Services Module
 - Five nodes of Network Traffic Analyzer Software
 - Additional licenses for purchase

Aruba NetEdit



Figure 3-31 Coordinated Network Switch Configuration

Aruba NetEdit empowers IT teams to orchestrate multiple switch configurations with intelligent capabilities, as indicated in Figure 3-31, including search, edit, validation (including conformance checking), deployment, and audit.

NetEdit, as shown in Figure 3-32, runs as an Open Virtualization Application (OVA) virtual machine (for example, VMware’s ESXi) on a server. NetEdit supports switches running ArubaOS-CX, including Aruba 8400, 8325, and 8320. To use NetEdit, Aruba OS-CX network device configurations need to be imported first. When NetEdit is made aware of a new switch, it will interrogate the switch configuration, hardware inventory and neighbor information. A user-friendly web-based UI allows network admins to easily interpret the status of network devices. A customizable application dashboard provides quick visibility into information and metrics that matter most to you. NetEdit ensures that there is a complete history of changes for all devices, targeting all configuration and hardware changes.

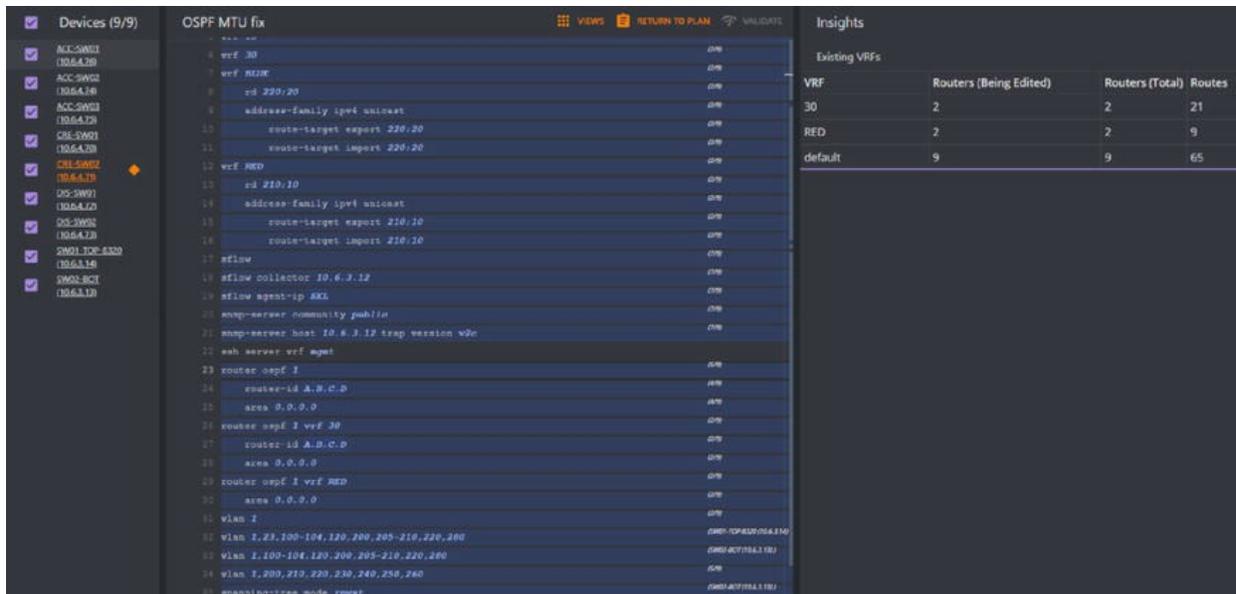


Figure 3-32 Orchestration using the Aruba NetEdit

Key features of Aruba NetEdit include:

- **Simpler device orchestration**—Configuration changes can be pushed to

multiple devices at the same time without any knowledge of API's.

- **Improved configuration consistency**—Configurations can be easily unified using conformance rules.
- **Reduced outage window**—Large configuration changes can be easily rolled back with the click of a button.
- **Better firmware management**—Firmware management is centralized in one place within NetEdit.

Learning check

6. How many network devices can be managed by the HPE Intelligent Management Center Basic Software Platform?
-

Customer scenario

The fictional customer LLP distribution company will be used as a storyline through this scenario. We will introduce the company using an interview:

- What is your primary business?
 - City-wide courier, parcel, and documents delivery
- How many employees do you currently have?
 - 25 employees + external workers
- What does your selling and delivery channel look like?
 - Contracts with dozens of local e-shops and expanding
- What does your server, storage, and network infrastructure look like?
 - Several tower servers and older networking devices from multiple brands
- Do you have an IT department?
 - No, external help on request

- What are your current plans?
 - Replace multiple vendor networking devices

Customer requirements

As a result of multiple interviews and gathering information about customer plans and customer's current infrastructure, the following requirements emerged for the new solution:

- Single vendor networking infrastructure
- Port requirements:
 - Sixteen ports for servers
 - Two switches
- Speed requirements
 - 1 Gb (26 ports)
- Teaming/bonding required on server ports
- L3 functionality
 - Dynamic Host Configuration Protocol (DHCP)
 - Static routing

Summary

- Switches can be organized by where they are deployed in the environment. Traditional networks are organized into tiers or leaf-spine topology.
- The Aruba portfolio for data centers includes ArubaOS-Switch and ArubaOS-CX platforms.
- HPE FlexFabric is designed to address traditional data center networking architectures and requirements.
- HPE Networking devices can be managed via CLI tools, APIs, or GUIs. Management interfaces can be embedded, installed on-premise, or

4 Recommending HPE NAS Solutions for Customer use Cases

LEARNING OBJECTIVES

After completing this chapter, you should be able to recommend and position HPE NAS products.

Recommending HPE NAS solutions for SMB customers

The wide HPE portfolio of data center products for SMB customers can be introduced using customer scenarios.

Introducing the customer scenario

The fictional customer LLP distribution company will be used as a storyline through this scenario. We will introduce the company using an interview:

- What is your primary business?
 - City-wide courier, parcel, and documents delivery
- How many employees do you currently have?
 - 25 employees + external workers
- What does your selling and delivery channel look like?
 - Contracts with dozens of local e-shops and expanding
- What does your server, storage, and network infrastructure look like?

- Tower servers, multiple switches including L3
- Do you have an IT department?
 - No, external help on request
- What are your current plans?
 - Deploying a centralized file-oriented solution
 - Consolidation of user data

Customer requirements

As a result of multiple interviews and gathering information about customer plans and the customer's current infrastructure, the following requirements emerged for the new solution:

- Centralized location for users' data consolidation
- 50 TB
- Simplified storage configuration tools
- Efficient deduplication
- Encryption for data at rest
- Possibility for replication and hybrid flash in the future

Activity: Discovering a customer's business and technical requirements

1. Prepare a list of additional questions to ask the customer about:
 - Current environment
 - Technical requirements
 - Room for growth
 - Resources
 - Management/monitoring



Figure 4-1 Selecting an HPE NAS platform

Depending on capacity requirements, the decision tree shown in [Figure 4-1](#) can be used to select appropriate NAS platform.



Note

Other limitations and prerequisites apply. Consult QuickSpecs. Some features require additional components. Capacities are changing during time.



Note

Most platforms are out of the scope of this guide; bigger details are coming for StoreEasy platform in this chapter.

Learning check

1. Recommend a NAS platform for a customer currently using an HPE 3PAR StoreServ Array.

Network-attached storage

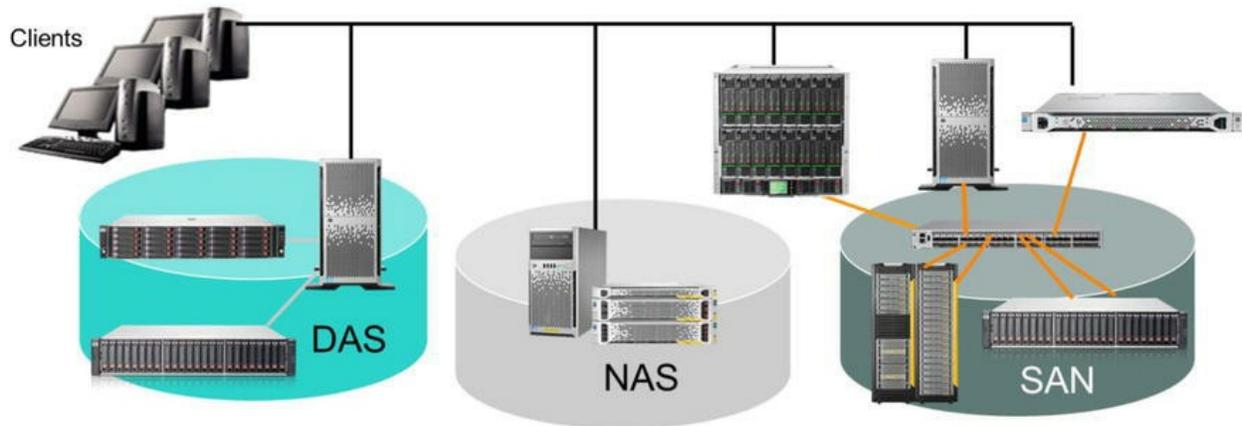


Figure 4-2 Implementing DAS-NAS-SAN

As illustrated in [Figure 4-2](#), there are three principal and complementary ways to implement storage solutions:

- **Direct-attached storage (DAS) storage solutions**—The most straightforward of storage solutions, either internal in a single server (usually using Serial-attached SCSI [SAS] protocol) or directly attached external storage. However, DAS storage cannot be shared among multiple servers, and the one-to-one connection between device and server can add management complexity.
- **Network-attached storage (NAS) devices**—Primarily used for unstructured data and file sharing, NAS storage adds increased storage management functionality, increased data protection, performance enhancements, and the ability to share data using industry-standard protocols across an Ethernet network. NAS devices can usually be easily and quickly installed and can be accessed by both servers and nonserver devices.
- **Storage area network (SAN) storage**—SAN is a solution rather than a product, using shared network components. SAN brings a dedicated network of storage systems accessible by multiple servers. In contrast to NAS delivering data using file systems, SAN allows block-level access and delivers even better levels of performance, management, and resilience. SANs can provide business continuity with redundant

components, automated failover, and centralized management. Common protocols used in SANs are Fibre Channel, internet Small Computer Systems Interface (iSCSI), Fibre Channel over Ethernet (FCoE), and InfiniBand.

This chapter covers NAS storage.



Note

InfiniBand (IB) is a computer-networking communications standard used in high-performance computing that features very high throughput and very low latency. It is used for data interconnect both among and within computers ([Wikipedia.org](https://en.wikipedia.org/wiki/InfiniBand), September 2019).

Finding the right fit for business and data can be challenging

Many customers are faced with limited resources—whether it is their IT staff or their budget—but to remain competitive, companies must take advantage of changing tides in the sea of technology to disrupt competitors and deliver superior results to customers.

Some of the challenges faced by SMBs and distributed organizations with limited on-site IT resources include:

- **Time**—The notion that managing on-premise storage is too time consuming and complex, and to combat that, they will need dedicated admins and storage specialists familiar with the platform they select.
- **Cost**—On-premise storage is an expensive upfront investment, and companies cannot afford to outpace and outgrow their storage.
- **Protection**—Data needs to be accessible to approved users, and protected in case of any “worst case scenario,” particularly if they are storing customer data and proprietary information. Many customers cannot compromise on the certainty of knowing their data is secure and inaccessible to unapproved parties and protected in case of tampering or disaster—something that is hard to do with cloud storage and front-end web apps for cloud storage like Dropbox and Box.

Discover a better solution for businesses



Figure 4-3 HPE network-attached storage

With the right technology powering your business, you can better serve your customers and sustainably grow without being held back by resource limitations.

HPE NAS, as referenced in [Figure 4-3](#), provides better storage solutions for businesses:

- **Storage dedicated to file**—Servers can be used for applications and high-performance and lowlatency work-loads.
- **Simple management**—For IT generalists or anyone with limited resources. Customer can spend less time on storage with simple management and a centralized repository for files and documents.
- **Access control**—Control who has access to what without inhibiting work and scale capacity as needed.

Network-attached storage

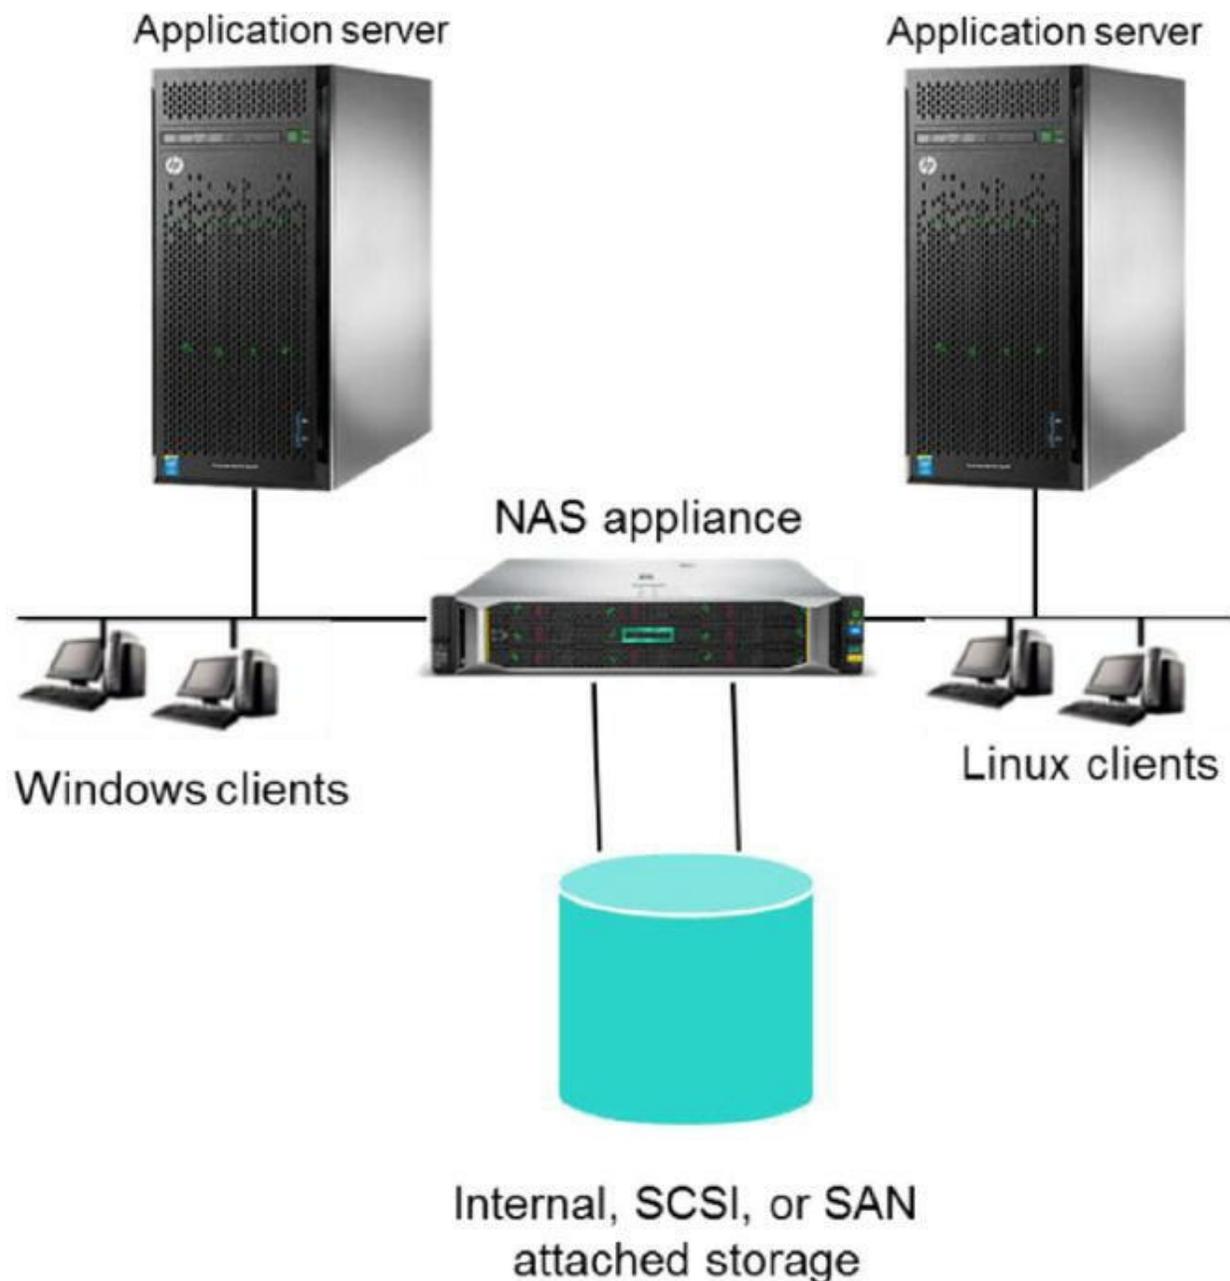


Figure 4-4 Network-attached storage

As illustrated in [Figure 4-4](#), NAS solutions consist of a specialized server-storage device that connects directly to the network. A file system is located and managed on the NAS device. Data is transferred to servers and other devices through the local area network (LAN) and to clients using industry-standard file-sharing protocols such as Common Internet File System (CIFS), Server Message Block (SMB), and Network File System (NFS).

NAS storage devices require storage cabinets providing specialized file access, security, and network connectivity. A network interface card (NIC) on the server is a requirement to access the storage. NAS provides file-to-disk block mapping and client access at the file level using network protocols.

NAS technology simplifies manageability and improves data access to clients and applications. A NAS solution is accessed over the Transmission Control Protocol/Internet Protocol (TCP/IP) Ethernet-based networks and supports multiple users connecting to it simultaneously. A NAS solution generally works with a mix of clients and servers running different operating systems.

NAS systems are fully optimized, secure, and highly available, utilizing Microsoft Windows Storage Server with HPE enhanced value and unique features.

The dedicated NAS appliance can provide shared storage between heterogeneous clients. NAS appliances connect to the network through a traditional LAN interface such as Ethernet, and they attach to network hubs similarly to the way servers and other network devices do. All network users have equal access to the stored data and do not have to go through the application server. NAS makes storage resources more readily available and helps to alleviate the server bottlenecks commonly associated with accessing DAS devices.

Advantages of NAS include:

- Is easy to install.
- Is affordable.
- Incorporates mature technologies.
- Offers a scalable solution.
- Increases network performance.
- Supports remote management.

Disadvantages of NAS include:

- Limited performance and storage capacity.
- Increased management costs.
- Data stored on a central NAS server is unavailable if the system shuts

down.

- Increased network traffic can create bottlenecks.

StoreEasy and Storage File Controllers provide the best of NAS and SAN

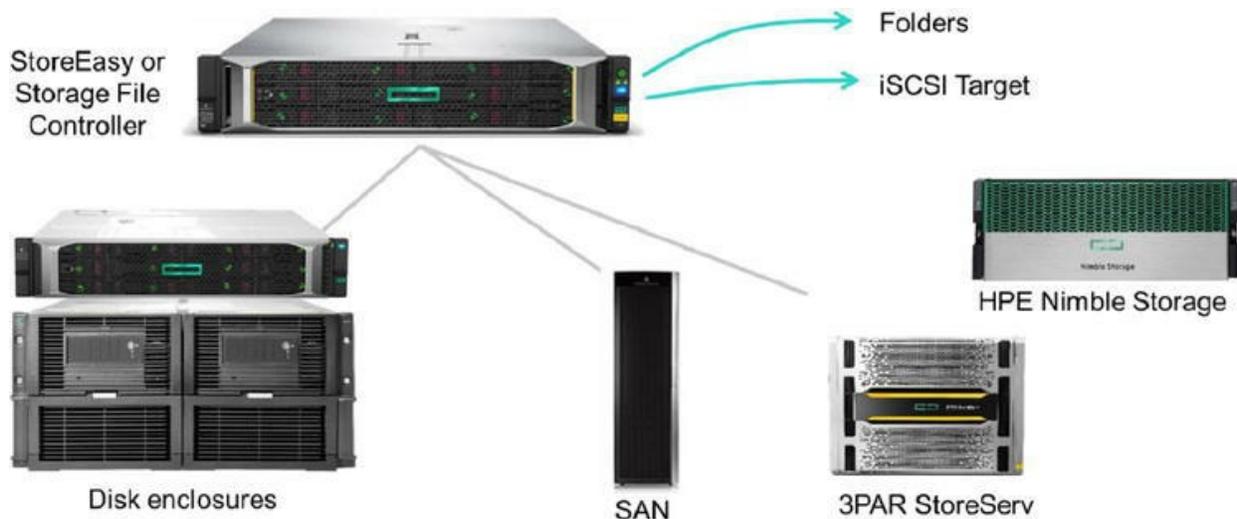


Figure 4-5 Centralized management interface to consolidate storage from DAS, NAS, and SAN

Figure 4-5 illustrates centralized management interface to consolidate storage from DAS, NAS, and SAN.

The major difference between SAN and NAS is that a SAN is a network separated from a company LAN and is configured to allow servers to communicate with storage arrays, typically using Fibre Channel. NAS requires a dedicated storage device, typically an optimized, network-attached server with several RAID storage drives. Both options have their strengths and weaknesses, with the primary advantages of a SAN being the major weakness of a NAS solution, and vice versa.

The benefits of SANs include network speed, reliability, centralization, and data protection. The main strengths of NAS are interoperability, low total cost of ownership (TCO), and relative simplicity. In a SAN architecture, administrators present storage to the server as a logical unit number (LUN) that appears to the server as a local disk that can be partitioned, formatted

with a file system, and used just like any other disk. It can also be used as raw storage.

HPE StoreEasy can take direct-attached disks (internal/external) as well as storage from SANs and make it available to client systems as either file shares or iSCSI targets. The StoreEasy administrator can use the same set of tools for working with the folders and LUNs.

HPE StoreEasy products are typical NAS products belonging to the primary storage products category. StoreEasy is a single, consolidated platform that delivers efficient, secure, and highly available application and file storage. It improves clients' TCO with standard functionalities typically found only in high-end systems, such as file and iSCSI access, granular deduplication, secure data encryption, and continuous availability. As a Windows-based solution, HPE StoreEasy also eliminates the customers' learning curve and enables them to maintain productivity levels without disruption based on their existing familiarity with the operating system.

StoreEasy is efficient, secure, and highly available storage to address file and application storage challenges for organizations large and small. It is powered by HPE ProLiant technology and Windows Storage Server, providing the tools and capabilities to help get the most out of your customers' capacity and spending less time managing storage. StoreEasy can densely scale capacity if required and seamlessly backup data to Microsoft Azure.

Ideal uses for HPE StoreEasy network-attached storage

While StoreEasy has an incredible number of features that can be implemented, typical use cases include:

- User, department, and corporate file shares
- Home directory consolidation
- File sync and share for desktop and mobile
- Connected remote office/branch office (ROBO)
- Bulk data preservation of unstructured file data
- Compliant archives using iTernity Compliant Archive Software
- Dedicated backup system for Microsoft apps, Hyper-V virtual machines

(VMs), and Windows PCs

- Simple data store for SQL Server and Hyper-V over SMB or iSCSI

Group, department, and corporate shares

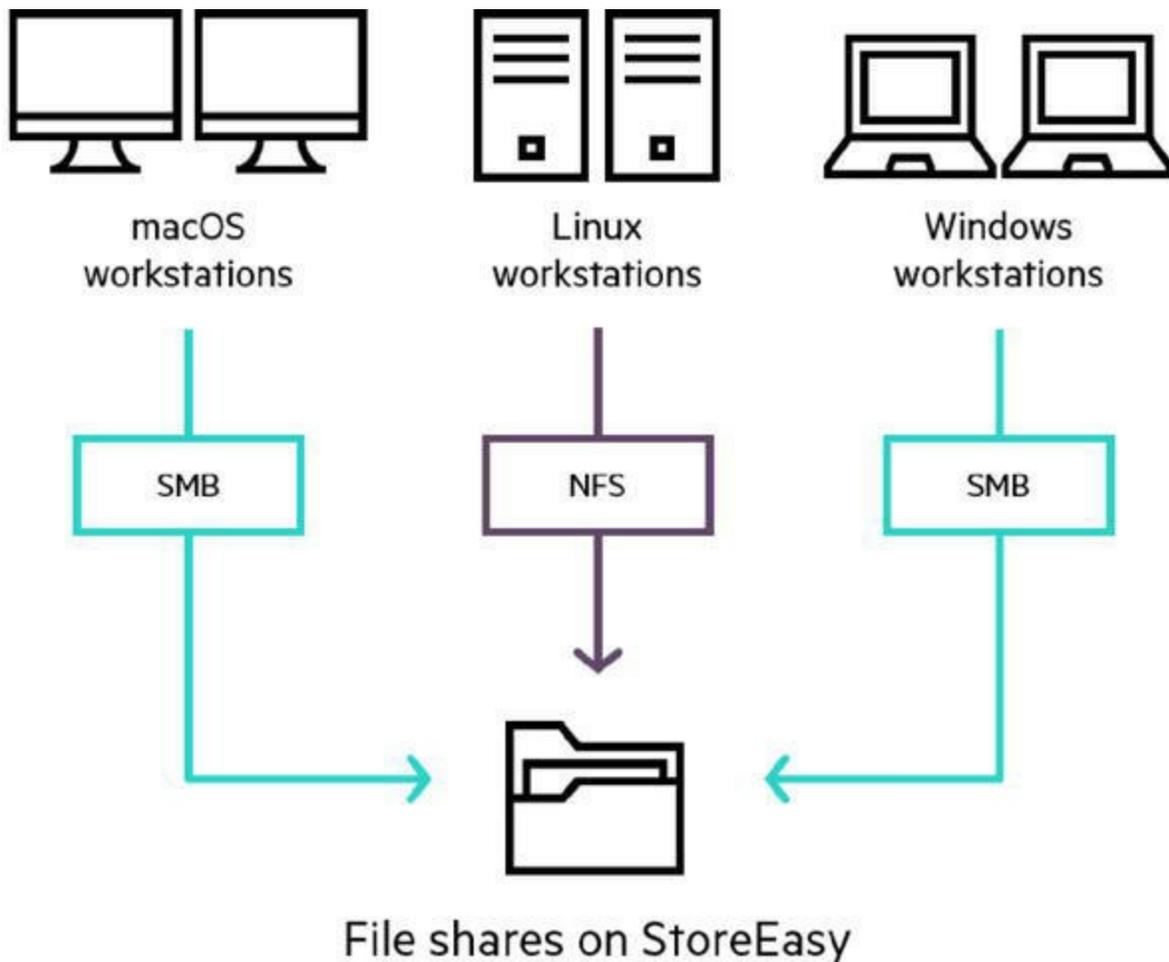


Figure 4-6 Group, department, and corporate shares

HPE StoreEasy products, as indicated in [Figure 4-6](#), support heterogeneous file sharing through multiprotocol access for both Windows clients and non-Windows clients. Centralized storage sharing enables efficient collaboration and data distribution with group, department, and corporate file shares that enhance organizational security and control of data.

HPE StoreEasy provides:

- **Multiple platform support**—Configured file shares, available via SMB and NFS protocols, provide seamless access from any client OS.
- **Reduced local storage footprint and redundancy**—Storing files on file shares reduces the volume of duplicate files and improves client performance with fewer files stored locally.
- **Improved productivity**—Access to mobile and distributed workforce and centralized file data ensures users access the current version of data.

Requirements for file sharing to Windows clients include:

- A failover cluster running Windows Server with at least two nodes, and the configuration of servers, storage, and networking must pass all tests performed in the “Validate a Configuration” wizard.
- File server role installed on all cluster nodes.
- Clustered file server configured with one or more file shares created with the continuously available property.
- SMB client computers running Windows 8, Windows Server 2012, or above.

SMB transparent failover enables administrators to configure Windows file shares in Windows Failover Clustering configurations to be continuously available. In case of a hardware or software failure, the server application nodes will transparently reconnect to another cluster node without interrupting the server applications.

SMB 3 provides active-active file sharing by allowing access to the same folders from multiple systems concurrently. This is achieved through Cluster Shared Volumes (CSVs).

SMB Encryption ensures that data transfers are secure by encrypting data as it transfers. The biggest benefit of using SMB Encryption over more general solutions (such as IPsec) is that there are no deployment requirements or costs beyond changing the SMB server settings. SMB Encryption is AES based.

The NFS 4.1 protocol is a significant evolution of the NFS protocol, and Microsoft delivers a standards-compliant, server-side implementation in Windows Server 2012.

Some of the Windows Server 2012 NFS 4.1 server features include a flexible, single-server namespace for easier share management, full Kerberos v5 support (including authentication, integrity, and privacy) for enhanced security, VSS snapshot integration for backup, and Unmapped UNIX User Access to enable easier user account integration.

Windows Server 2012 supports simultaneous SMB 3 and NFS access to the same share, identity mapping using stores based on RFC-2307 for easier and more secure identity integration, and highly available cluster deployments.

Windows Server 2012 includes a new, flat file-based identity mapping store. Windows PowerShell commandlets (cmdlets) also replace cumbersome manual steps to provision Active Directory Lightweight Directory Services (AD LDS) as an identity mapping store and to manage mapped identities.

Home directory consolidation

HPE StoreEasy can be used to centralize user data storage by consolidating local home directories with simple management, reduced client upgrade costs, and greater IT control.

Users in many organizations store data, such as application data files or documents, either locally on individual laptops or on desktop PCs. This practice has several disadvantages:

- **Reduced mobility**—Users cannot access the files or documents from other network locations or other PCs in the organization.
- **Risk of data loss**—Malfunction of the local disk can result in data loss.
- **Inefficient storage usage**—Often, different versions of the same file are stored in several places, which increase total storage-related costs.
- **Difficulty in performing regular data backups**—Individual backups for every PC are cumbersome.
- **Time-consuming maintenance or upgrades**—PC maintenance or upgrade activity requires more time to back up and restore the data.
- **Interferes with common user experience**—Even if a user's data is configured to be stored on a centralized network file share, the desktop, operating system, and application settings are not the same between

different desktop PCs or laptops. This requires additional learning time to understand the PC settings and to locate the data, thus decreasing productivity.

- **Network file accessibility**—User data stored on a network file share becomes inaccessible in the event of a network or server outage.

HPE StoreEasy systems can address these challenges, enabling an organization to store user data with greater efficiency and consistency in centralized storage instead of individual user PCs through the following key features:

- **Folder Redirection** enables administrators to redirect the path of a known folder to a new location. The new location is a file share on an HPE StoreEasy system. Users interact with files in the redirected folder as if it still existed on the local drive. For example, the administrator can redirect the “Documents” folder, which is usually stored on a local drive, to a shared location of a StoreEasy system.
- **Offline Files** makes network files available, regardless of a slow or unavailable connection to the server. When working online, file access performance is at the speed of the network and server. When working offline, files are retrieved from the “Offline Files” folder at local access speeds. For example, users can continue accessing their data in the event of a corporate network or server outage. It is even possible to take laptops to a remote location, such as a home office, and access data without connectivity to the corporate network.
- **Roaming User Profile** redirects user profiles to a file share so that users receive the same operating system and application settings on multiple computers.
- **Work Folder** provides users flexibility to work on files online or offline, and data will auto sync with a centralized file server when it is connected to internet. The user should have no worries about internet connection, login, or password to access their corporate network, and files are up to date whenever files are accessed.

Compliant data preservation

Business applications



WORM, retention, self-healing,
encryption & compression



StoreEasy 1000 with iCAS
installed and licensed

DR site A

DR Site B

Figure 4-7 Long-term data integrity with HPE StoreEasy

Another use case for HPE NAS platform, as presented in [Figure 4-7](#), is to ensure long-term data integrity via tamper-proof, legally compliant archive and data preservation solution with Compliant Archive Software (iCAS) and StoreEasy:

- Perfect fit for regulated industries, public sector, or organizations with product liability requirements
- Use as an archive backend for a variety of business applications, including enterprise resource planning (ERP)/SAP, enterprise content management (ECM), email, and so on
- Archive documents and digital media, including audio and videos, for long-term retention on cost-optimized storage
- Bulk file archiving up to 192 TB in a single system with StoreEasy 1660, with archives up to PB scale supported

iTernity is one of the leading enterprise software companies exclusively focused on managing and protecting enterprise data and compliance-relevant information. It enables customers of all sizes to realize a future-proof and hardware-independent “enterprise archive” where information can be accessed and migrated freely without hardware boundaries. iTernity customers gain a competitive edge by securing business continuity and data integrity for long-term relevant information in a most flexible and cost-efficient way.

Together, HPE and iTernity deliver unmatched price:performance in data management, archiving, and protection solutions for crucial enterprise data. For many years, HPE and iTernity have worked together to design and deliver advanced archiving and data protection solutions to customers. Our complementary technologies and strategies work together to deliver comprehensive, cost-effective solutions for enterprises, and:

- Ensure regulatory compliance and integrity of records.
- Maximize flexibility so your customers can adapt as business needs change.
- Increase security with constant integrity checks of important information.

- Drive down cost.

iTernity iCAS is an enterprise archive and data protection software solution that helps organizations store data flexibly and securely. Built upon industry standards and Windows platforms, and supporting SAN and NAS storage devices, iCAS offers the adaptability that today's IT departments need.

A huge, installed base of highly satisfied iTernity customers benefits from regulatory compliance solutions that can mitigate risk, simplify data management, and improve overall business success. iCAS complements applications and disk-based storage platforms with compliant archiving and data management features.

File data can also be backed up to Microsoft Azure using an Azure subscription to protect against disasters and catastrophic failures.



Note

Although it can be used in SMB segment, iCAS is primarily an enterprise archive and data protection software solution.

Flexible deployment options for Hyper-V

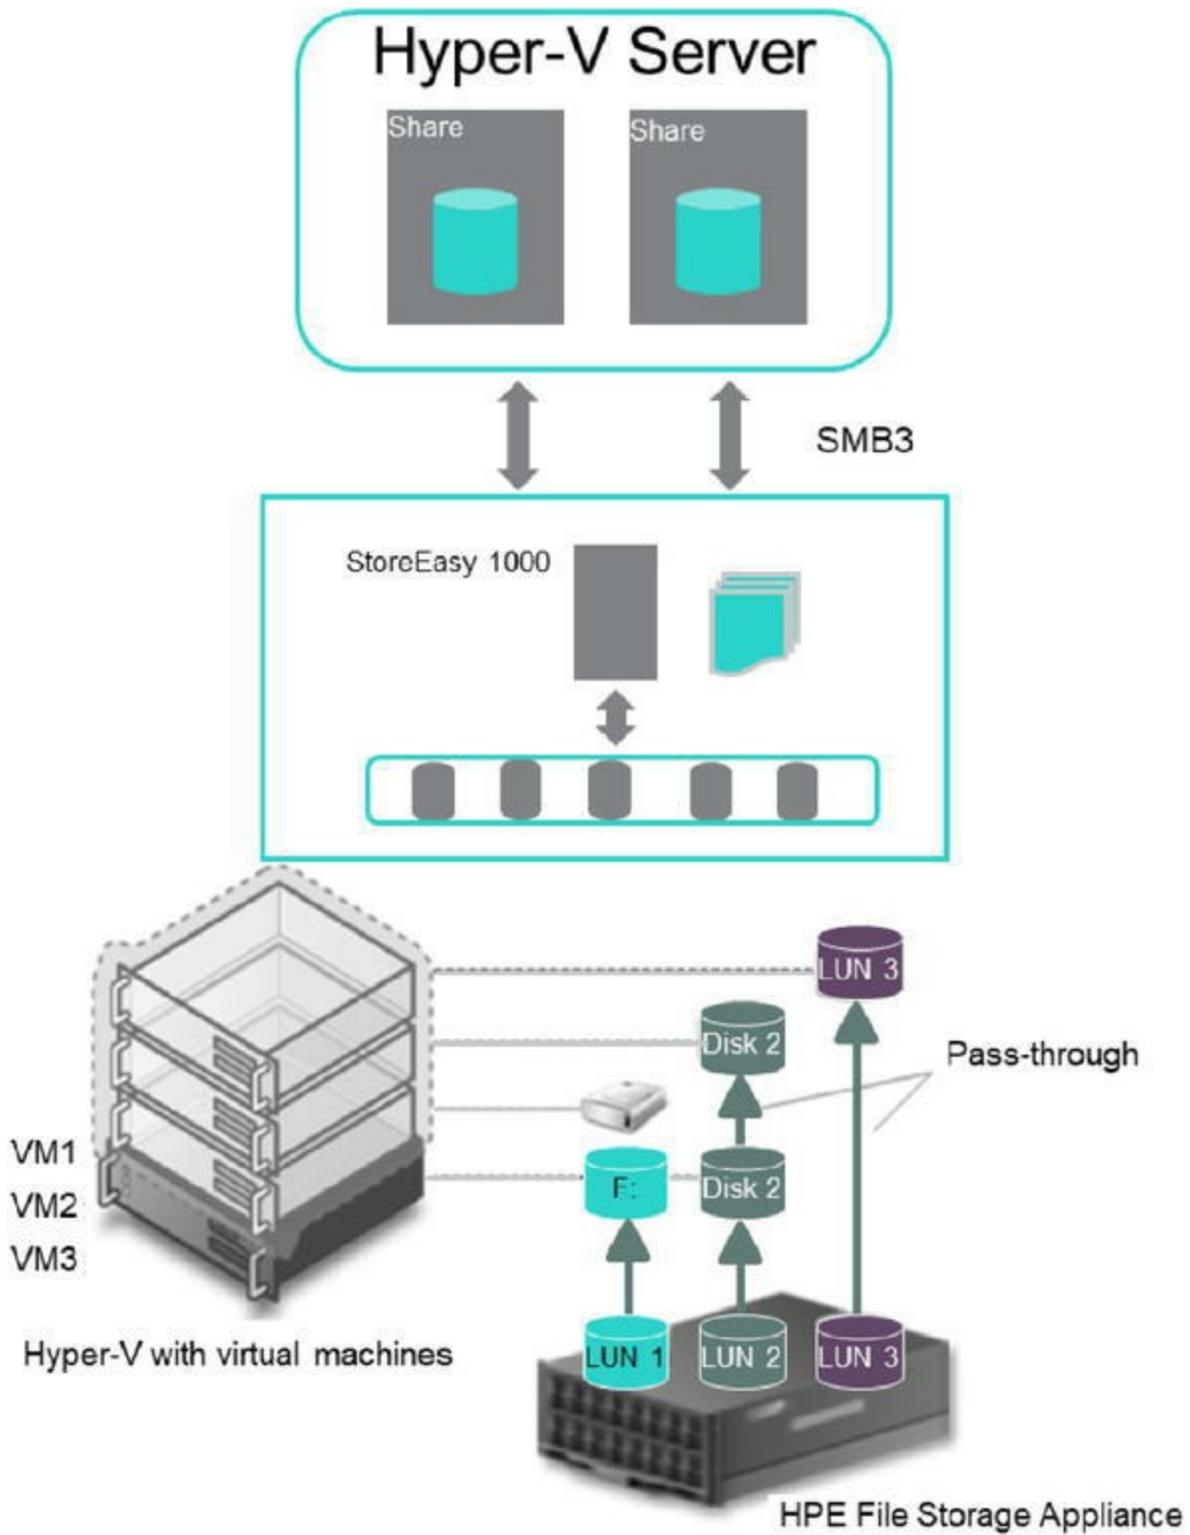


Figure 4-8 SMB 3 shared data store for Hyper-V

Before Windows Server 2012, remote storage options for Hyper-V were

limited to SAN solutions that can be difficult to provision for Hyper-V guests or other more inexpensive options that do not offer many features. By enabling Hyper-V to use SMB file shares for virtual storage, administrators have an option that is not only simple to provision with support for CSVs and inexpensive to deploy, but also offers performance capabilities and features that rival those available with Fibre Channel SANs.

SMB 3 shared data store for Hyper-V, as illustrated in [Figure 4-8](#), provides:

- Simplified provisioning and management (manage file shares instead of storage fabric and LUNs)
- Fibre Channel-like capabilities
- Fast data transfer and network fault tolerance

Traditional SAN capabilities with iSCSI are possible with:

- Host Hyper-V virtual machines on HPE StoreEasy
 - As a virtual hard disk on a parent partition volume
 - As pass-through a parent partition disk
- Add as additional storage to guest OS
 - iSCSI initiator in child partition directly connects to LUN
 - Reduce typical virtualization costs by utilizing SAN-like capabilities

Benefits include:

- **Fast data transfers and network fault tolerance with SMB Multichannel**—With Windows Server 2012, customers can store application data (such as Hyper-V and Microsoft SQL Server) on remote SMB file shares. SMB Multichannel provides better throughput and multiple redundant paths from the server—for example, Hyper-V or Microsoft SQL Server—to the storage on a remote SMB share.
- **Transparent failover and node fault tolerance with SMB—Supporting business-critical server application workloads** requires the connection to the storage back end to be continuously available. The new SMB server and client cooperate to provide transparent failover to an alternative cluster node for all SMB operations for planned moves and unplanned failures.

Traditional SAN capabilities with iSCSI are available, too. The administrator can host Hyper-V VMs on HPE StoreEasy either as a virtual hard disk (VHD) on the parent partition volume or as a pass-through parent partition disk. Another use of iSCSI is to add it as additional storage to the guest OS, with the iSCSI initiator in the child partition directly connecting to the LUN.

File storage optimized for SQL Server workloads

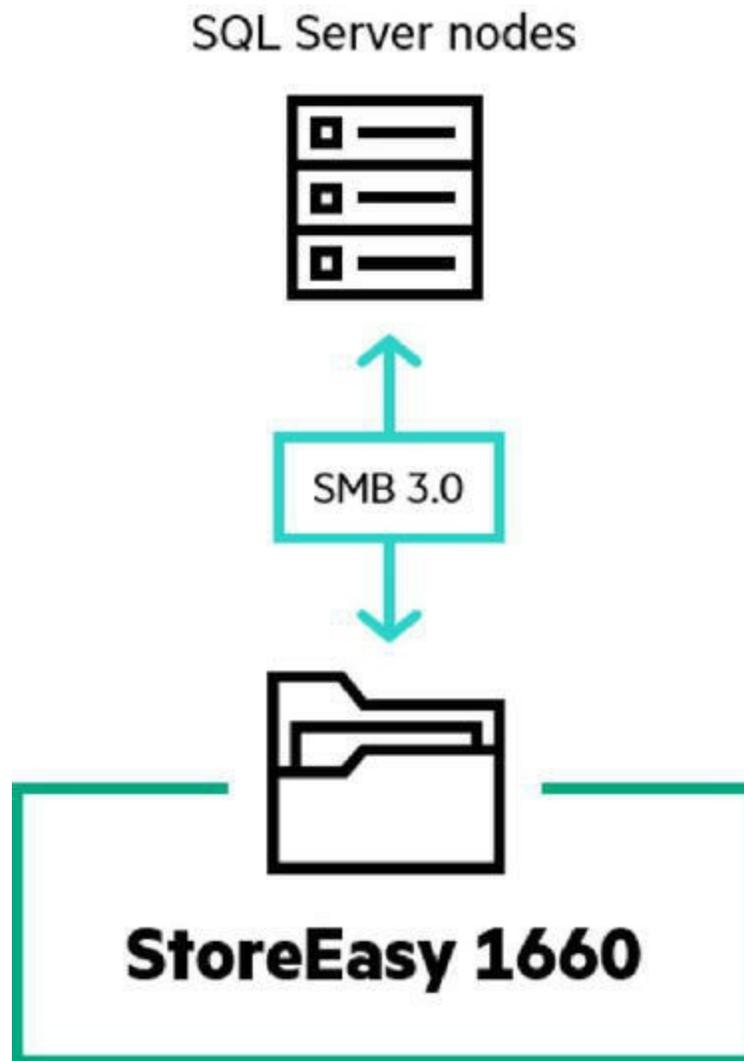


Figure 4-9 File storage optimized for SQL Server workloads

The SMB client and SMB server, as illustrated in [Figure 4-9](#), are both optimized for small random read/write IO, to increase performance for

common server application workloads, such as Microsoft SQL Server online transaction processing (OLTP). SMB also uses a large maximum transmission unit (MTU) feature, enabled by default, to significantly improve performance in large sequential transfers, such as those for Microsoft SQL Server data warehouse, in addition to database backup and restore operations.

Customer benefits include:

- **Fast data transfers and network fault tolerance with SMB Multichannel**—With Windows Server, customers can store application data (such as Hyper-V and Microsoft SQL Server) on remote SMB file shares. SMB Multichannel provides better throughput and multiple redundant paths from the server—for example, Hyper-V or Microsoft SQL Server—to the storage on a remote SMB share.
- **Transparent failover and node fault tolerance with SMB**—Supporting business-critical server application workloads requires the connection to the storage back end to be continuously available. The new SMB server and client cooperate to provide transparent failover to an alternative cluster node for all SMB operations for planned moves and unplanned failures.

HPE StoreEasy Storage family overview



Figure 4-10 Leading NAS product family

All StoreEasy servers, as illustrated in [Figure 4-10](#), can act as a NAS device

to provide storage for other clients. They also all support the same protocols. Where that storage actually comes from, though, depends on the StoreEasy model.

A gateway allows you to present LUNs from the SAN or array to HPE StoreEasy, and you can then offer LUNs as SMB and NFS file shares to clients. This also enables other Windows-based features such as snapshots (Volume Shadow Copies), data deduplication, quotas, and much more.

The HPE StoreEasy portfolio includes:

- **StoreEasy 1000**—Uses internal or attached storage
- **Storage File Controller**—Acts as a gateway for LUNs from a SAN array

The HPE StoreEasy 1000 Storage family is designed for small businesses, branch offices, and workgroup environments. The HPE StoreEasy Storage family integrates easily into new and existing environments, offering storage administrators and IT generalists a straightforward, consistent management experience.

With built-in security, data remains well protected and available to users when and where they want to access it. It also enables installation and running end-point protection such as antivirus on the system itself, reducing the cost and complexity of connecting an external, end-point protection server. With features including constant health monitoring, mirrored OS disks, advanced memory protection, reliable software support, and more, users can rely on a solution that keeps their data ready and highly available.

Available models:

- HPE StoreEasy 1460 Storage
- HPE StoreEasy 1560 Storage
- HPE StoreEasy 1660 Storage
- HPE StoreEasy 1860 Storage

StoreEasy 1000 solutions are available with:

- Capacity: None (additional drives needed), 8 TB, 9.6 TB, 14.4 TB, 16 TB, 32 TB, 64 TB

- Form factor
 - Rack
 - Tower
- Drive type
 - SAS, SATA
 - LFF, SFF
- Windows Storage Server 2016 preinstalled

Activity: Discovering the HPE NAS portfolio

1. Open the StoreEasy family overview at:
<https://www.hpe.com/info/storeeasy>
 - a. Review the first case study.
 - b. Expand the Related Case Studies section, select one case study, and review.
2. Answer the following questions:
 - a. What was the business driver to move the data to a new NAS solution?

- b. Why did the customer select an HPE NAS product?

HPE StoreEasy 1460 Storage



Figure 4-11 HPE StoreEasy 1460 Storage

HPE StoreEasy 1460 Storage products, as referenced in [Figure 4-11](#), are storage solutions that deliver multi-protocol file serving and application storage in a compact and affordable 1U rack-mount form factor. They are easy to install and manage, making them perfect to consolidate users' documents, images, audio, and videos. HPE StoreEasy 1460 is perfect for a small workgroup, small businesses, or remote offices. HPE StoreEasy 1460 models come with Windows Storage Server 2016 preinstalled from the factory on the four internal large form factor (LFF) drives. The HPE StoreEasy 1460 is built on a ProLiant DL360 Gen 10 Server and includes a rail kit.

Additional features include:

- Scale up to 32 TB raw internal capacity
- Intel® Xeon® Bronze 3104 processor (6 core/85 W/1.7 GHz)
- 8 GB DDR4-2666 MT/s memory with room to expand
- Four-port 1 GbE NIC embedded with 10/25 GbE options
- Windows Storage Server 2016 Standard Edition

Raw, usable, and effective capacity

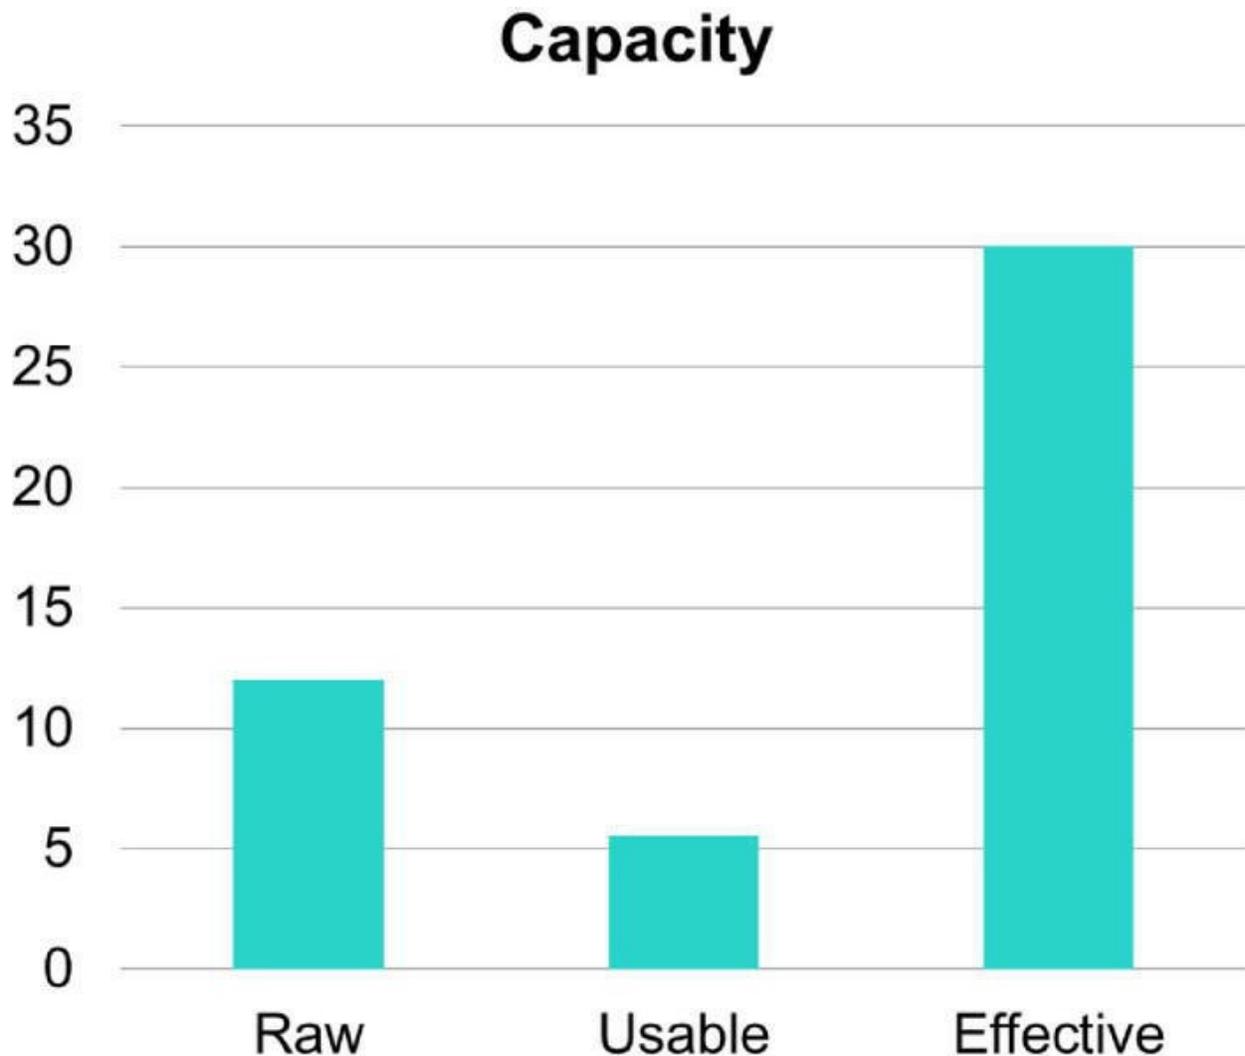


Figure 4-12 Example of raw, usable, and effective capacity

A typical attribute of a storage solution is the capacity. Three typical types of capacities, as shown in [Figure 4-12](#), are used when describing storage products:

- **Raw**—Sum of drive capacities.
- **Usable**—Raw capacity corrected of RAID overhead and system requirements.
- **Effective**—Usable capacity corrected of deduplication and compression. Sometimes external, cloud-based capacity is calculated.

Example:

- 4x 3 TB HDD = 12 TB Raw
- RAID 1 + 0, OS requires 500 GB = 5.5 TB Usable
- Deduplication 2:1, compression 1.5:1, 14.5 TB in cloud = 30 TB effective

HPE StoreEasy 1560 Storage



Figure 4-13 HPE StoreEasy 1560 Storage

HPE StoreEasy 1560 Storage products, as shown in [Figure 4-13](#), are storage

solutions that deliver multi-protocol file serving and application storage in a compact and affordable tower form factor package. They are easy to install and manage, making them perfect for a small workgroup, small businesses, or remote offices. HPE StoreEasy 1560 models come with Windows Storage Server 2016 Standard Edition OS preinstalled from the factory on the four internal LFF drives. The HPE StoreEasy 1560 is built on a ProLiant ML110 Gen 10 Server.

Additional features include:

- Scale up to eight LFF drives with optional drive cage (four LFF drive slots standard; optional drive cage adds four more slots)
- Intel® Xeon® Bronze 3104 processor (6 core/85 W/1.7 GHz)
- 8 GB DDR4-2666 MT/s memory with room to expand
- Four-port 1 GbE NIC embedded with 10/25 GbE options
- Windows Storage Server 2016 Standard Edition

HPE StoreEasy 1660 Storage



Figure 4-14 HPE StoreEasy 1660 Storage

HPE StoreEasy 1660 Storage products, as indicated in [Figure 4-14](#), deliver high capacity and drive down your cost per GB. Four- or six-core processing (depends on the model selected), up to 16 LFF hot-plug drive slots, and ready for external expandability enhance this ultra-dense 2U shared storage solution for small, medium, or large IT environments. HPE StoreEasy 1660 models come with Windows Storage Server 2016 Standard Edition OS preinstalled from the factory on two M.2 solid-state drives (SSDs). The HPE StoreEasy 1660 is built on a ProLiant DL380 Gen 10 Server and includes a rail kit.

Additional features include:

- Scale up to 192 TB raw internal capacity with new 12 TB HDDs
- Intel® Xeon® Bronze 3104 processor (6 core/85 W/1.7 GHz)
- 16 GB DDR4-2666 MT/s memory with room to expand
- Four-port 1 GbE NIC embedded with 10/25 GbE options
- HPE Smart Array P816i-a SR Gen10 controller
- Windows Storage Server 2016 Standard Edition

HPE StoreEasy 1860 Storage



Figure 4-15 HPE StoreEasy 1860 Storage

HPE StoreEasy 1860 Storage products, as illustrated in [Figure 4-15](#), deliver enterprise-class performance, expandability, and flexibility. Four- or six-core processing (depending on the model selected), up to 28 internal small form factor (SFF) data drives for performance, and ready for external expandability enhance this 2U shared storage solution for small, medium, or large IT environments. HPE StoreEasy 1860 models come with Windows Storage Server 2016 Standard Edition OS preinstalled from the factory on two M.2 SSDs. The HPE StoreEasy 1860 is built on a ProLiant DL380 Gen10 server and includes a rail kit.

Additional features include:

- Scale up to 67.2 TB raw internal capacity with new 2.4 TB HDDs
- Intel® Xeon® Bronze 3104 processor (6 core/85 W/1.7 GHz)
- 16 GB DDR4-2666 MT/s memory with room to expand
- Four-port 1 GbE NIC embedded with 10/25 GbE options
- HPE Smart Array P816i-a SR Gen10 controller

- Windows Storage Server 2016 Standard Edition

Selecting an HPE StoreEasy 1000 platform



Figure 4-16 Selecting an HPE StoreEasy 1000 platform

Depending on capacity and fabric form requirements the decision tree shown in [Figure 4-16](#) can be used to select appropriate NAS platforms.



Note

Other limitations and prerequisites apply. Consult QuickSpecs. Some features require additional components. Capacities are changing during time.

Learning check

2. Name three typical use cases for HPE StoreEasy platform.

HPE StoreEasy features and management

HPE StoreEasy helps regain control over file data. Benefits of HPE StoreEasy products are as follows:

- **Optimized**—Tailored for file storage
- **Efficient**—Save money and time
- **Secure**—Protect data always
- **Highly available**—Prevent business and user disruption

Optimized: Tailored for file storage

The HPE StoreEasy Storage comes with preconfigured hardware with Windows Storage Server 2016, preinstalled for purpose-built file storage. StoreEasy installation and network configuration wizards simplify installation and prevent configuration errors for faster deployment.

HPE StoreEasy Dashboard consolidates all key status information such as capacity utilization, performance, and health in a single interface and reduces the need to navigate multiple tools and interfaces for system management. Low-capacity warning email alerts are now available to minimize disruption.

The best practices-driven storage provisioning tool discovers system storage and recommends the most efficient setup for optimum performance, optimum capacity, or a balanced configuration based on the available storage.

Hard disk drive (HDD) bundles make it easier for you to buy, deploy, and provision additional storage capacity over the life of the product.

Efficient: Save you money and time

The file sync and share feature enables mobile workers to access their files from anywhere. It provides corporate data controls including data encryption, data wipe, user lock-screen password enforcement, file classification, and storage quotas.

Customers can:

- Use advanced data management to automatically improve use of resources, comply with data retention policies, and enhance protection of sensitive files.
- Consolidate multiple workloads and IT services in a single platform with included support for a broad range of file access protocols, Microsoft Hyper-V and SQL Server application storage, iSCSI block access, and IT infrastructure services such as print, Domain Name System (DNS), and Dynamic Host Configuration Protocol (DHCP).
- Reduce branch office wide area network (WAN) usage with hosted BranchCache, automatically caching frequently accessed corporate data center files.
- Achieve an average of 50 to 60% space savings with granular data deduplication.

Secure: Protect your data always

HPE StoreEasy Storage protects data while at rest with file system encryption and BitLocker Drive Encryption, and while it is being transferred with SMB encryption and signing.

Customers can:

- Use file classification infrastructure to dynamically identify files based on sensitivity, and implement sophisticated access controls using Active Directory Rights Management Services.
- Install and run end-point protection, such as an antivirus, on the system itself, thereby reducing the cost and complexity of connecting to an external end-point protection server.
- Prevent against data loss with Volume Shadow Copy Service online snapshots and support for agent-based backup software.
- Enable simplified and cost-efficient compliance archiving with iTernity Compliant Archive Solution (iCAS), now validated to meet the United States Securities and Exchange Commission (SEC) 17a-4(f) requirements.



Note

SEC Rule 17a-4 is a regulation issued by the U.S. Securities and Exchange Commission pursuant to its regulatory authority under the U.S. Securities Exchange Act of 1934 (Known simply as the “Exchange Act”), which outlines requirements for data retention, indexing, and accessibility for companies that deal in the trade or brokering of financial securities such as stocks, bonds, and futures. According to the rule, records of numerous types of transactions must be retained and indexed on indelible media with immediate accessibility for a period of two years, and with non-immediate access for a period of at least six years. Duplicate records must also be kept within the same time frame at an offsite location ([Wikipedia.org](https://en.wikipedia.org/wiki/SEC_Rule_17a-4), September 2019).

Highly available: Prevent business and user disruption

HPE StoreEasy Storage increases uptime with such standard features as near continuous health monitoring with HPE Active Health System, mirrored OS drives, redundant fans (not available on all models) and power, and advanced memory protection.

Customers can:

- Confidently deploy large file systems leveraging file system online self-healing, online verification, and precise and rapid repair that massively reduces downtime from file corruption.
- Take advantage of real-time byte-level replication and disaster protection with Double-Take Availability by Carbonite. Installation software is included on fifth-generation HPE StoreEasy.
- Get support from file storage experts with three-year parts, three-year labor, and three-year on-site support with next business day response (StoreEasy 1650/1850) or three-year parts, one-year labor, and one-year on-site support with next business day response (StoreEasy 1450/1550). All StoreEasy products also include one year of 24x7 software telephone support for the Windows Storage Server operating system.

HPE StoreEasy is optimized to simplify file storage

HPE StoreEasy is a purpose-built system for your customer to get StoreEasy

out of the box and online fast and includes:

- Preconfigured hardware and software
- Task-driven installation wizard
- Simple network configuration and validation

HPE StoreEasy contains a simplified provisioning and management, which helps your customer spend less time on storage:

- Brand new management experience
- Guided storage provisioning and centralized monitoring
- Low capacity alerts with drive bundles to order

StoreEasy file system types

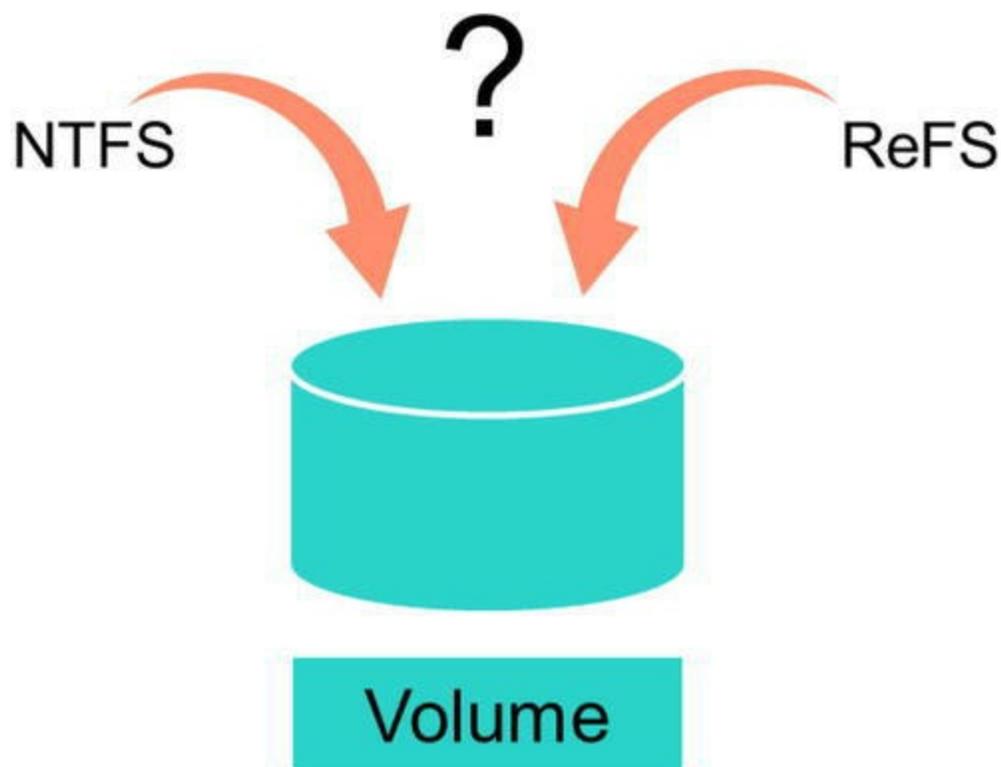


Figure 4-17 NTFS or ReFS

When creating a volume, you need to select between two different file systems: New Technology File System (NTFS) and Resilient File System

(ReFS), as illustrated in [Figure 4-17](#).

NTFS is the default file system when creating a new volume and should typically be selected. It is the traditional file system for Windows Servers and applications that run on them. NTFS also provides numerous features, including file compression, quotas, encryption, deduplication, and more.

In addition, very large volumes (for example, 64 TB) can now be confidently deployed without fear of service downtime associated with file system check/repair operations. Windows 2012/2016 performs online scanning for corruptions, online identification of error conditions for on-disk file system data structures, and online repair. Windows now incurs file system downtime only when absolutely needed and such downtime is proportional only to the specific number of (online) detected corruptions. All of these enhancements together deliver significantly improved NTFS uptime.

Currently, ReFS is rarely used with StoreEasy since typically it is deployed in combination with Storage Spaces. Storage Spaces is not supported with StoreEasy.

ReFS is Microsoft's newest file system, designed to maximize data availability, scale efficiently to large datasets across diverse workloads, and provide data integrity by means of resiliency to corruption. It seeks to address an expanding set of storage scenarios and establish a foundation for future innovations; it is not a replacement for NTFS.

Key features are:

- Built-in resiliency and availability.
- Maximizes availability in combination with Storage Spaces.
- Metadata corruption is always detected, and user data corruption is also detected.
- Scales to PB+ datasets comprising very large files or directories.
- Compatibility with existing system application programming interfaces (APIs) and technologies.
- Some NTFS features are not supported, including boot, disk quota, deduplication, file compression, File Server Resource Manager (FSRM), and so on.

Opt for a hybrid flash configuration

Up to **4x** faster access to data versus HDD-only configurations*



Figure 4-18 Smart Cache kit

While HDDs provide the best \$/GB for storing data, adding SSDs for use as a cache is an economical way to boost performance and lower latency, as indicated in [Figure 4-18](#). Some StoreEasy models allow adding SSD layers, such as The Smart Cache kit, contain two x HPE 800 GB SAS 12 G Mixed Use SFF (2.5 in) SC, a three-year warranty, and Digitally Signed Firmware SSD (SKU# 872376-B21). The Smart Cache kit is used with the SmartCache feature included with the HPE Smart Array P816i-a SR Gen10 Storage Controller, which provides a Flash Acceleration data tier to the HPE StoreEasy 1860.

Smart Cache features:

- Cache hot data on lower latency SSDs, with configuration and management through HPE Smart Storage Administrator
- Read and write-back caching supported
- Ideal for application acceleration or latency-sensitive use cases
- 1.6 TB SSD bundles available in LFF and SFF



Note

*Based on HPE internal testing with 90% cache hit rate

Maximize storage efficiency through wide protocol support

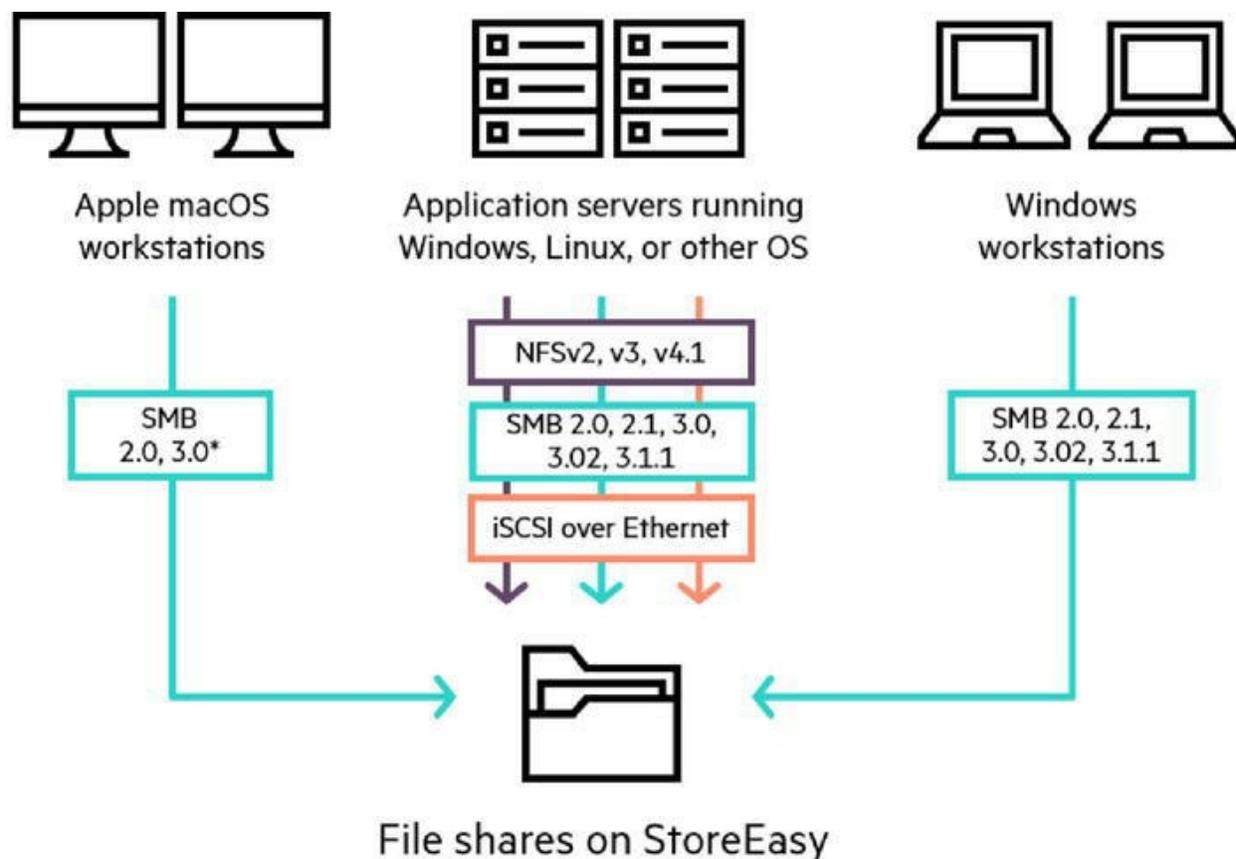


Figure 4-19 StoreEasy has wide protocol support

As referenced in [Figure 4-19](#), StoreEasy offers a variety of protocols:

- **Full-fidelity SMB protocol support** provides continuous availability, secure data transfer, increased network performance, and fault tolerance.
- **Broad NFS version support** provides continuous availability, easier share management, and simple user identity mapping.
- Use as **iSCSI target server** for virtualization and application workloads.

Specific version or generation of protocol may rely on support in the operating system.

Multi-dimensional security keeps data protected

StoreEasy Storage delivers many physical and digital protection against simple user errors, corrupted files, unauthorized access, physical tampering,

and malicious attacks:

- A **lockable bezel** blocks access to hardware components.
- **Drive encryption** ensures no one without the encryption key can read the data written to StoreEasy drives.
- **File system encryption, file classification, and dynamic access controls** ensure only the users your customers authorize can access files.
- **Endpoint assurance and encryption** secures data in transit.
- Automatically enabled **anti-malware** protects against malicious attacks.

Easy-to-use, asynchronous replication

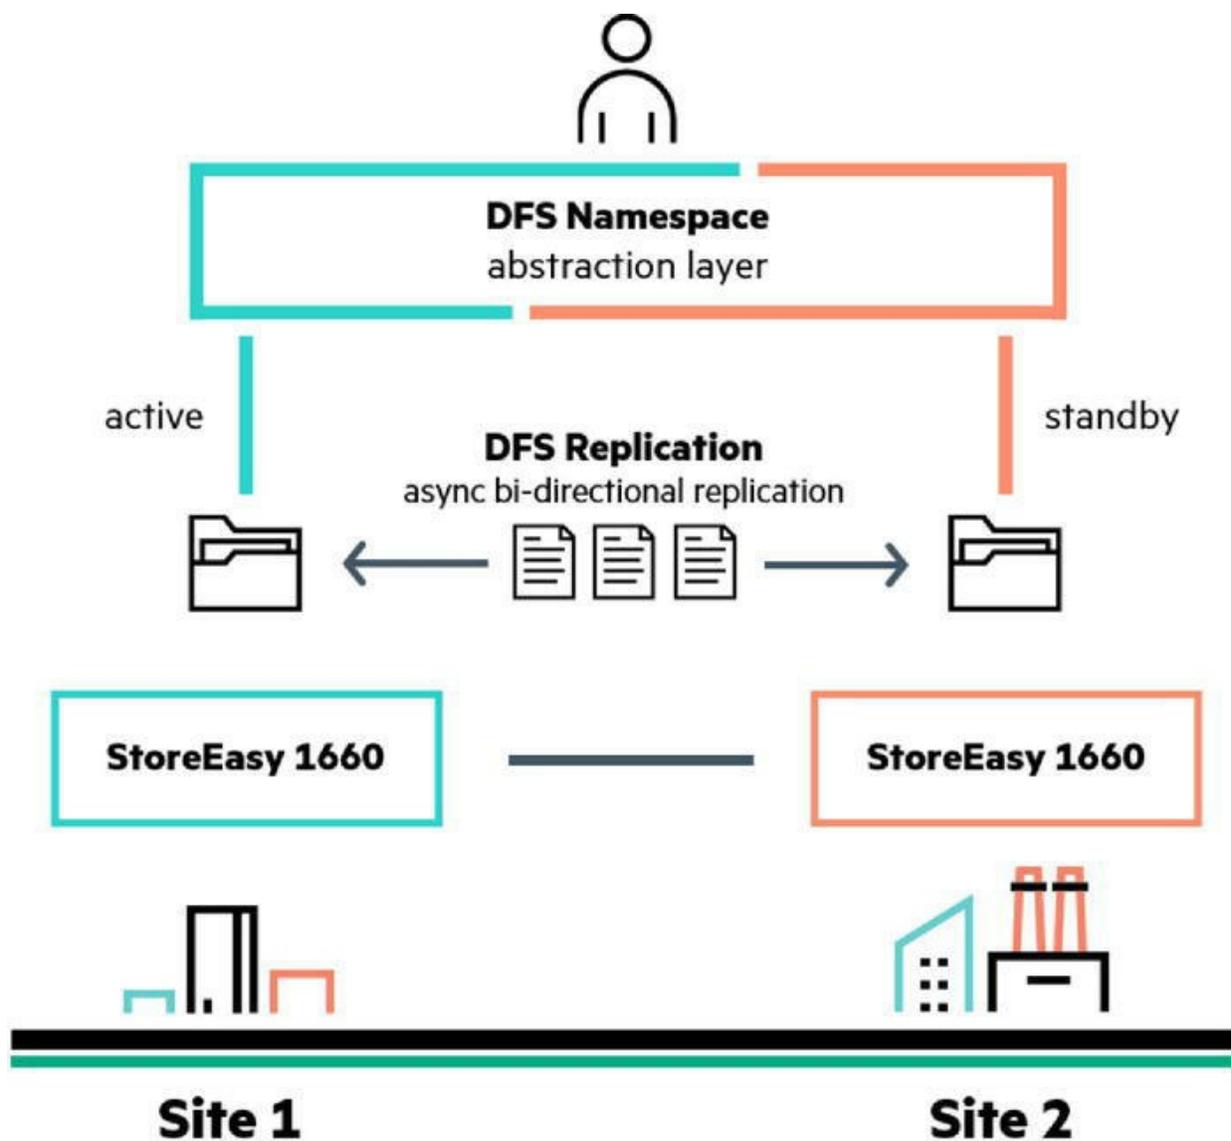


Figure 4-20 DFS-R and DFS-N

BranchCache enables data and user failover with active-active data access on both sites, data failover with Distributed File System Replication (DFS-R) and user access failover with Distributed File System Namespaces (DFS-N).

As illustrated in [Figure 4-20](#), DFS-N and DFS-R are role services in Windows Server in the File and Storage Services role:

- **DFS Namespaces** enables grouping of shared folders that are located on different servers into one or more logically structured namespaces. Each namespace appears to users as a single shared folder with a series of

subfolders. However, the underlying structure of the namespace can consist of numerous file shares that are located on different servers and in multiple sites.

- **DFS Replication** enables efficient replication of folders (including those referred to by a DFS namespace path) across multiple servers and sites. DFS Replication uses a compression algorithm known as remote differential compression (RDC). RDC detects changes to the data in a file, and it enables DFS Replication to replicate only the changed file blocks instead of the entire file.

Built-in asynchronous replication provided by **DFS-R** continuously mirrors file sets:

- Up to 100 TB per server
- Up to 70 million files per server
- Maximum file size of 250 GB
- Scheduled bandwidth limiting
- Bidirectional replication
- Replication over high-latency IP networks
- Capacity-efficient, deduplication aware replication

DFS-N provides an orchestration layer between replication targets for a seamless failover experience for SMB file-sharing clients.

Guided configuration simplifies storage deployment

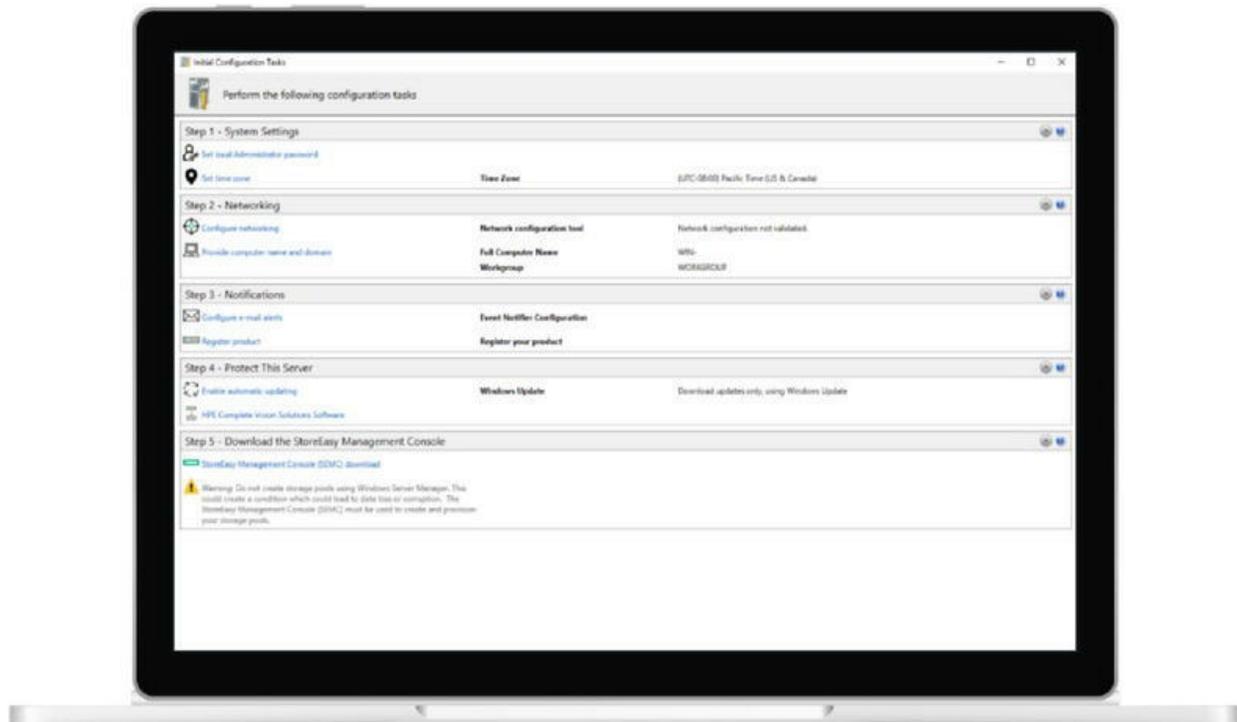


Figure 4-21 Initial setup of HPE StoreEasy system

Guided configuration, as indicated in [Figure 4-21](#), helps with initial setup through key tasks, includes network configuration validation tool, and minimizes disruption by configuring email alerts. Email alerts include:

- Low capacity notifications include drive bundles and SKUs to order when ready to expand capacity.
- Optional HPE Proactive Notifications deliver product updates.

The setup wizard also includes a link to download the new management console.

Day-to-day monitoring made even easier

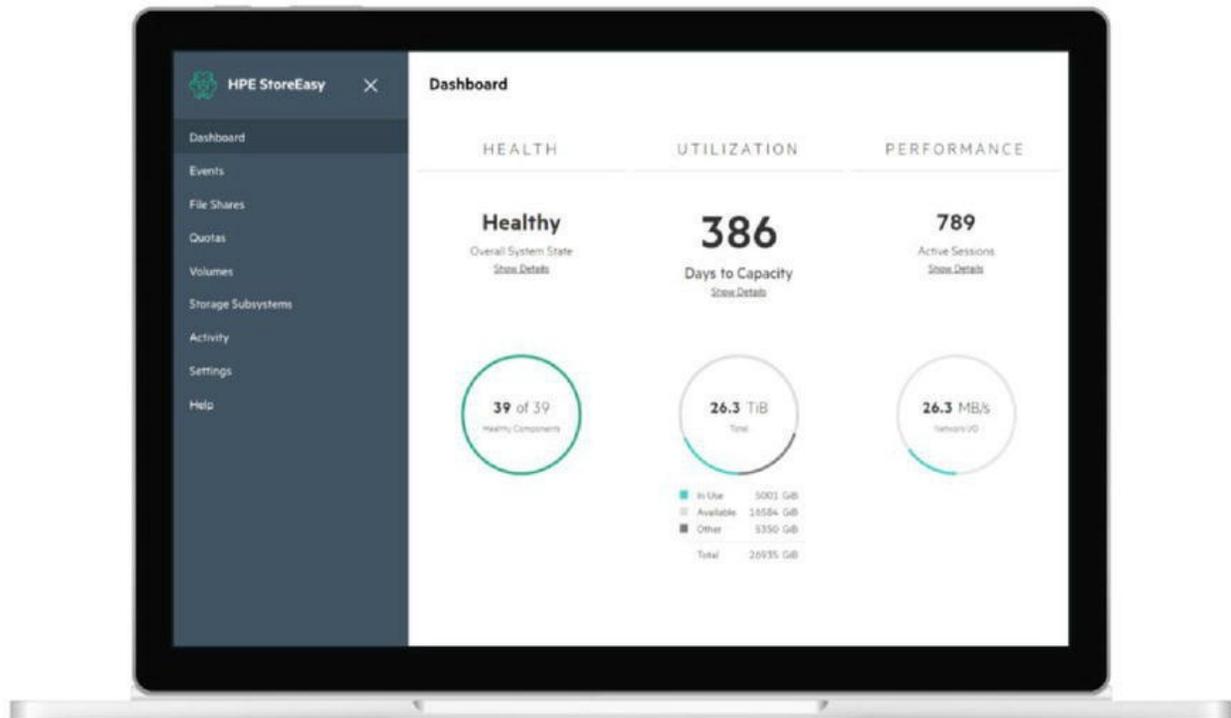


Figure 4-22 HPE StoreEasy management console, backed by a powerful REST API

HPE StoreEasy management console, as illustrated in [Figure 4-22](#), is designed to improve your customer’s experience by saving them time with streamlined management workflows. The console provides all-in-one file storage management that eliminates complexity, avoids legacy tools, and lightens the IT administrative burden. Simplified exposure of advanced features enables your customers to take advantage of compelling capabilities with less effort.

Hassle-free setup and monitoring of file shares

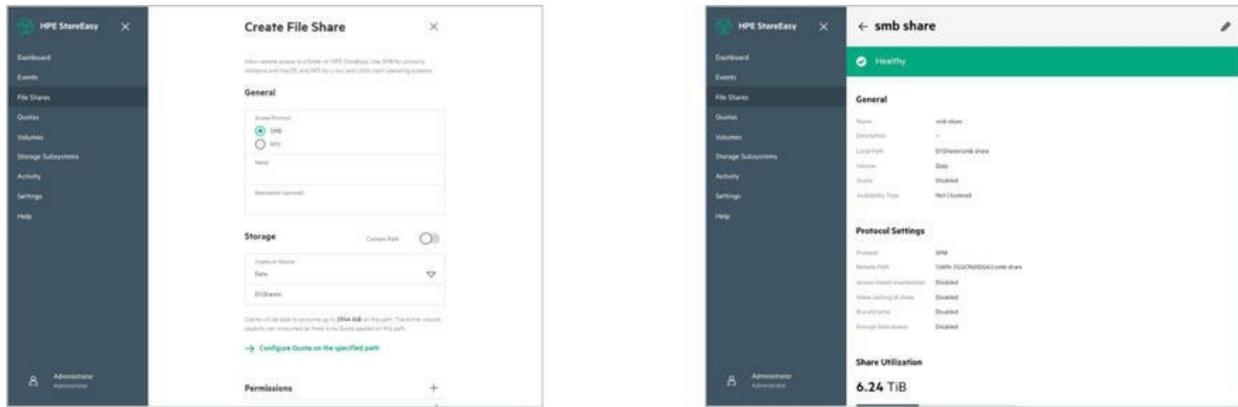


Figure 4-23 Hassle-free setup and monitoring of file shares

As seen in [Figure 4-23](#), HPE StoreEasy console allows administrators to quickly define and create new file shares, and monitor all relevant status and configuration details about shares with just a few clicks.

Managing sixth-generation StoreEasy 1000

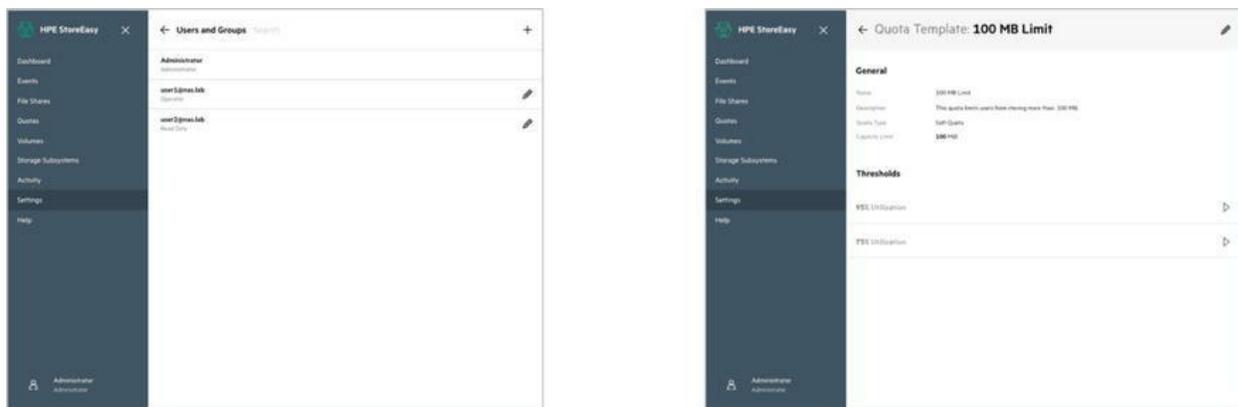


Figure 4-24 Managing sixth-generation StoreEasy 1000

The management console can also be used to define and apply a wide range of file storage administrative tasks, including user permissions to control access to sensitive data and quotas to control capacity consumption, as illustrated in [Figure 4-24](#).

Integrated management in Server Manager

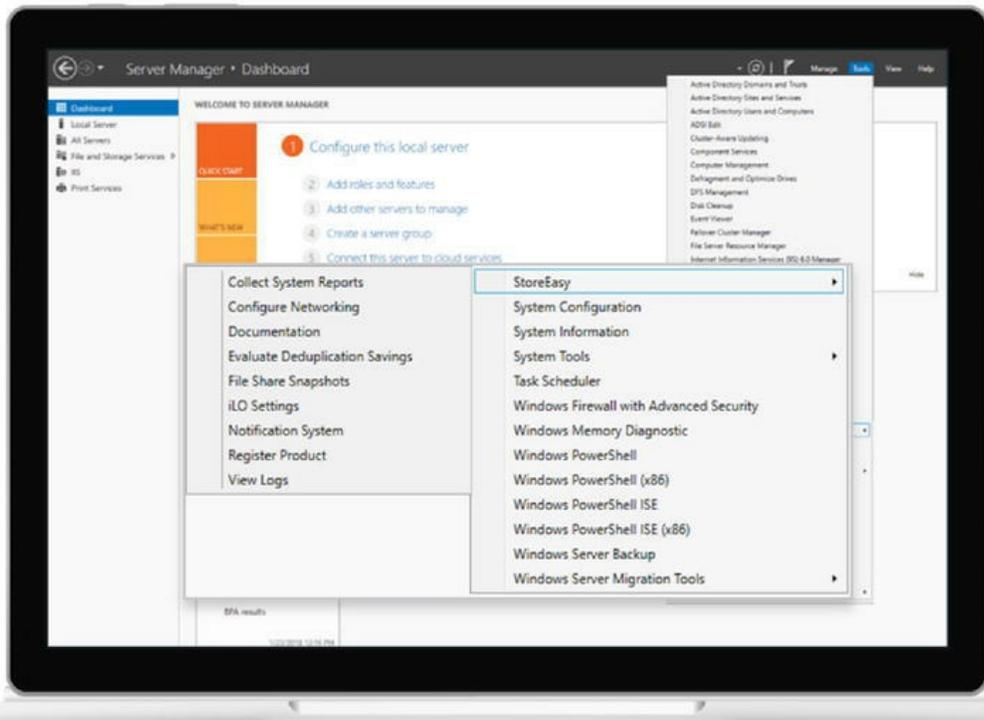


Figure 4-25 Integrated management in Server Manager

As illustrated in [Figure 4-25](#), alternatively, administrators who prefer traditional Microsoft tools can access StoreEasy tools and documentation through Server Manager.

Learning check

3. Which technology allows efficient replication of folders in an HPE StoreEasy solution?
 - A. DFS-N
 - B. NTFS
 - C. DFS-R
 - D. ReFS

Storage-sharing technologies

Multiple storage-sharing technologies are available with HPE StoreEasy.

Storage management

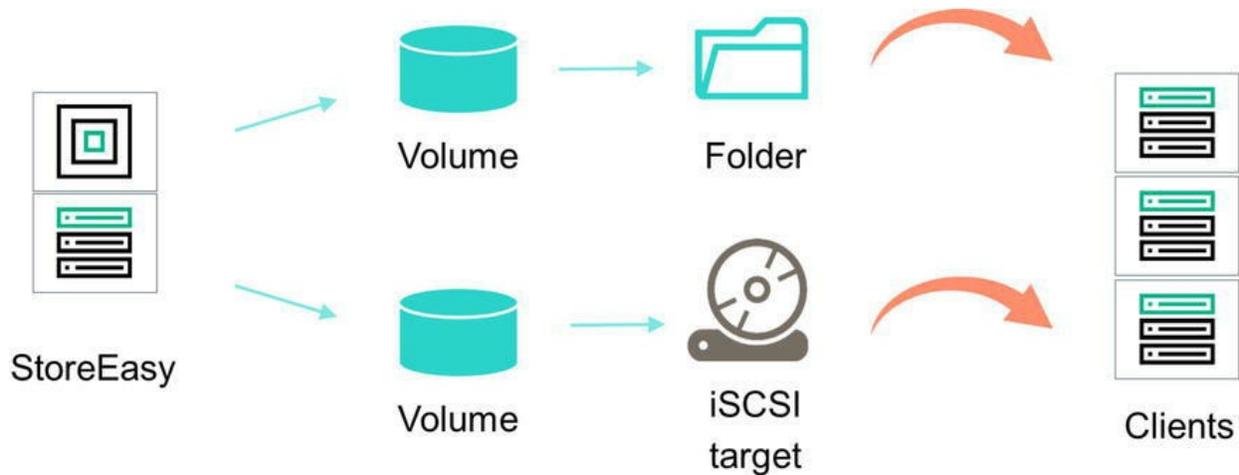


Figure 4-26 Process of sharing folders and/or creating iSCSI targets

After StoreEasy has volumes, you can perform storage management tasks such as sharing files and folders and/or creating iSCSI targets. These can subsequently be made available to clients. The whole process is shown in [Figure 4-26](#).

Folder and share management

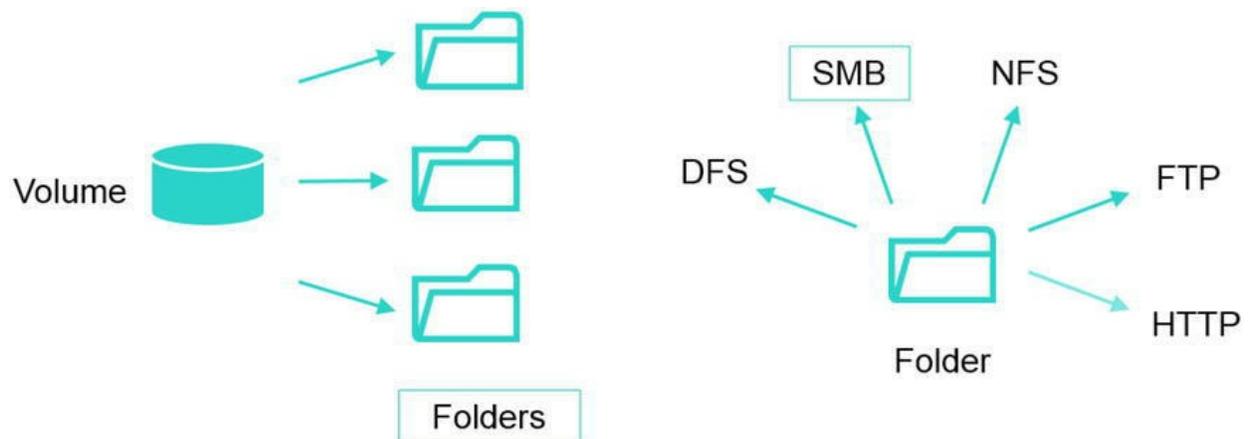


Figure 4-27 Sharing files across different protocols

Volumes and folders on any system are used to organize data. Regardless of system size, systematic structuring and naming conventions of volumes and folders ease the administrative burden. Moving from volumes to folders to shares increases the level of granularity of the types of data stored in the unit and the level of security access allowed.

Folders can be managed using the Server Manager or StoreEasy management console. Tasks include:

- Accessing a specific volume or folder
- Creating a new folder
- Deleting a folder
- Modifying folder properties
- Creating a new share
- Managing shares

There are several ways to set up and manage shares. Methods include using Windows Explorer, the command line interface (CLI), or Server Manager.

The file-sharing security model of the storage system is based on the NTFS file-level security model. Share security seamlessly integrates with file security.



Note

File-sharing protocols, as referenced in [Figure 4-27](#):

- **SMB**—Server Message Block
 - **NFS**—Network File System
 - **DFS**—Distributed File System
 - **HTTP**—HyperText Transfer Protocol
 - **FTP**—File Transfer Protocol
-

iSCSI targets



Figure 4-28 End-to-end connectivity

As illustrated in [Figure 4-28](#), you can also make StoreEasy storage available to clients as an iSCSI target. These iSCSI targets appear to clients as local disks that can be partitioned, formatted with a filesystem, and used just like any other disk. It may also be used as raw storage. The iSCSI target could be made from DAS (StoreEasy 1000) or from a SAN array if StoreEasy is acting as a SAN Gateway (Storage File Controller). One potential advantage of using a SAN in this way is if the SAN is using Fibre Channel, StoreEasy could allow non-Fibre Channel connected systems to access SAN storage over the IP-based network.



Important

Generally, StoreEasy should not be used to serve storage to systems that are already connected to the Fibre Channel fabric as this can reduce performance due to the extra processing overhead. Use the SAN management software to present the SAN LUN directly to the client over the Fibre Channel fabric.

Learning check

4. How can security be integrated at the file level in HPE StoreEasy?

-
-
5. You are in a meeting with a customer, and they challenge you with a statement: “I heard that HPE StoreEasy products are only supported in a Windows environment.”

How should you respond?

Sizing NAS products

HPE StoreEasy storage systems can be ordered using various configurations, so the solution will fit customer needs. HPE provides a variety of tools and resources to help size and configure a customer solution. One of the best places to find online information about HPE products and solutions is the HPE website and HPE Product Bulletin.



Note

HPE Product Bulletin can be downloaded from <http://hpe.com/info/qs>.

NAS sizing considerations

To correctly select a StoreEasy system, evaluate the following considerations:

- Total capacity and room for growth
- Number of users accessing storage (type of protocol used by clients)
- Required availability (RAID-level selection, cluster)

- Performance requirements (protocol, rotational speed, SSD, SAS, SATA, cache in the controller)
- Requirements for installation support and enhancements to warranty services

Performance best practice for Server Message Block protocol (SMB 3)

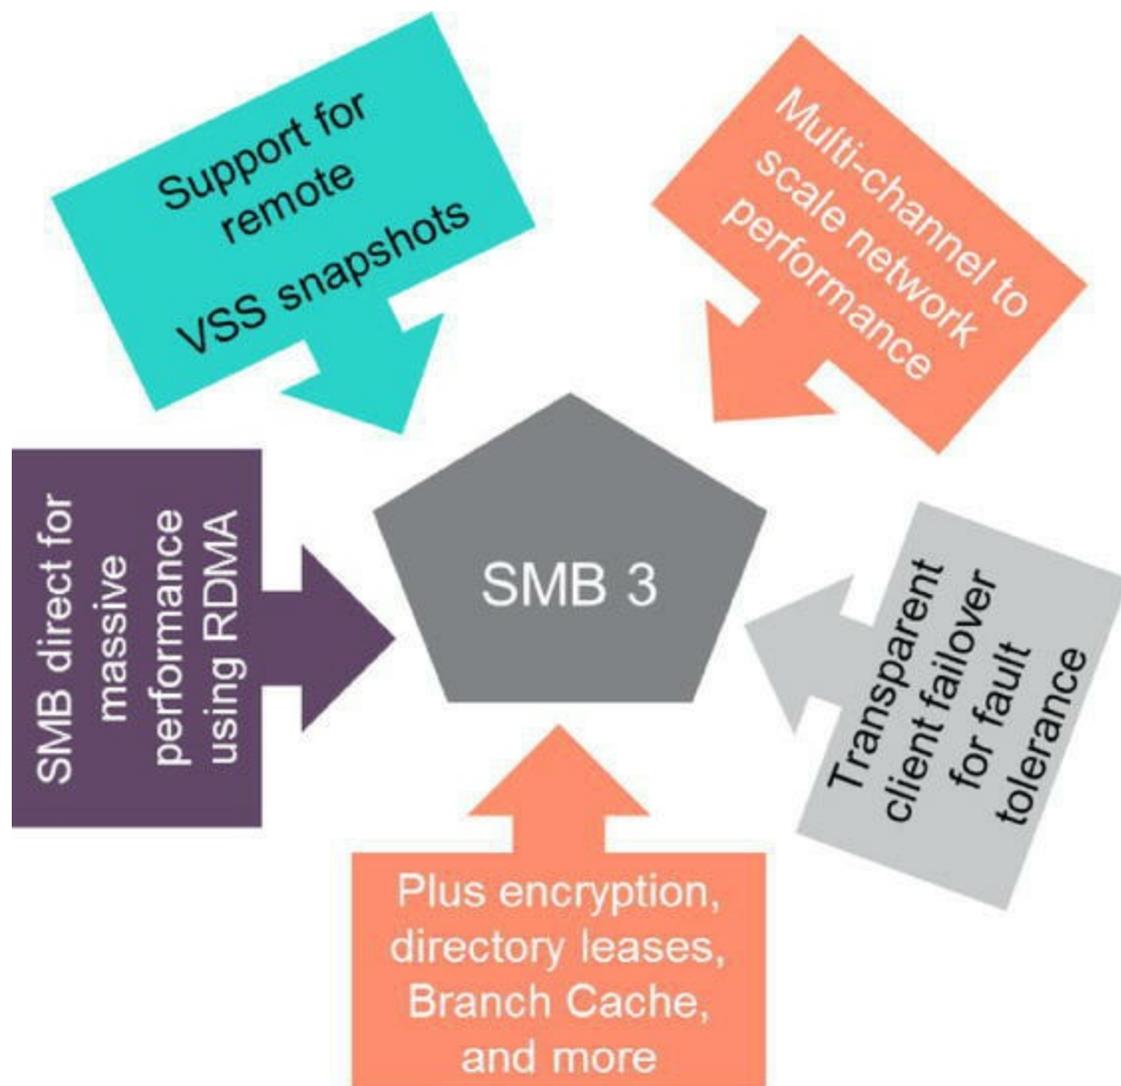


Figure 4-29 Server Message Block protocol SMB 3

HPE StoreEasy 1000 and Storage File Controller are built on the Windows

Storage Server 2016 operating system. Windows Storage Server 2012 and 2012 R2 introduced major new features in the SMB 3 network protocol stack, as referenced in [Figure 4-29](#), which can improve the performance and reliability of SQL Server with databases stored on SMB 3 File Share on a StoreEasy Storage product.

SMB multichannel uses multiple network interfaces, if possible (all StoreEasy models).

Both the performance and reliability of the system can be increased as follows:

- Improve performance by providing the combined throughput of all network adapters and cluster nodes together.
- Improve reliability by transparently falling back to the remaining network adapters or cluster nodes if one component fails.

SMB 3 protocol compatibility

SMB 3 is compatible with clients using version 1 or 2 of the SMB protocol, but the new features mentioned above are available only to client operating systems also using SMB version 3. Only operating system variants, such as Windows Server 2012 or Windows 8.1 and newer, support the SMB 3.0+ protocol.

Using SMB 3 multichannel

SMB 3 multichannel is a feature that enables SMB clients to automatically identify all the network interfaces on a target server and to simultaneously use all the network interfaces on the client that can connect to those on the server. This is similar in concept to network teaming, except that it requires no configuration in the network and sharing center. All available network devices are automatically used. If one of the links fails, SMB automatically reroutes requests through the remaining links. The application will experience only a slight delay on completion of one of its IOs. To use SMB 3 multichannel, connect multiple network interfaces both on the Windows 2016, 2012, or 2012 R2 server hosting the SQL Server application, and on the StoreEasy 1000 or Storage File Controller.

Using SMB 3 scale-out

SMB 3 scale-out enables multiple clustered servers to share the same file system. A client accesses that share through all clustered servers simultaneously, provided the file share is on a Cluster Shared Volume (CSV).

CSVs are suitable for SQL data files or for similar applications making many IOs into a few large files kept open for a long time. They are not suitable for workloads generating a large number of metadata operations, such as opening files, closing files, creating new files, or renaming existing files.

Configuring SQL Server to use an SMB share

Configure the default database locations using UNC paths (\\server\share) instead of a local drive (D:).

SMB 3.1.1

Encryption has been enhanced in SMB 3.1.1. The cipher can now be negotiated during connection establishment. In addition to AES-128-CCM for SMB 3.0.x compatibility, Windows 10 (and Windows Server 2016) added AES-128-GCM in SMB 3.1.1. The Galois/Counter Mode (GCM) offers a significant performance gain.

Virtualization best practices using Hyper-V on SMB or iSCSI with HPE StoreEasy products

Features	RAID 1	RAID 10	RAID 5	RAID 50	RAID 60
Minimum no. of drives	2	4	3	6	8
Data protection	Single-drive failure	Up to 1 disk failure in each pair	Single-drive failure	Up to 1 disk failure for every 3 disks	Up to 2 disk failures for every 4 disks
Read performance	High	High	High	High	High
Write performance	Medium	Medium	Low	Medium	Medium
Read performance (while in degraded state)	Medium	High	Low	Medium	Medium
Write performance (while in degraded state)	High	High	Low	Medium	Low

Figure 4-30 RAID features

Figure 4-30 highlights the RAID configurations and their performance with different RAID configurations.

When planning for hosting Hyper-V on an HPE StoreEasy 1000 or Storage File Controller system, the storage should be configured to give both the level of reliability that the customer expects and the performance the clients need.

With a lower number of disks and a high workload, customers and clients might experience slow VM performance. Adding more disks to a RAID drive set might increase the performance of the VMs.

When provisioning the storage for use by Hyper-V, the actual process might be different between the different HPE StoreEasy products. For more information on how to connect, provision, and monitor storage using the HPE StoreEasy system, see the appropriate HPE StoreEasy User Guide or external SAN array user guide.

Configuration process steps

1. Choose the HPE StoreEasy model (1460, 1560, 1660, 1860).
2. Select the specific Build to Order (BTO) or Configure to Order (CTO) (some models only) configuration.
3. Select required internal storage drives for CTO (some models only: 1660,

1860).

4. Choose an optional external storage controller and disk enclosure (D3X10, D6020).
5. Choose an optional second network adapter.
6. Choose a Peripheral Component Interconnect express (PCIe) riser kit (some models only: 1460, 1660, 1860).
7. Choose an optional second LFF drive cage (1560)/mid-chassis cage (1660, 1860) and additional drives.
8. Choose an optional second processor or performance kit.
9. Choose an optional extra memory kit.
10. Choose an optional second power supply (some models only: 1460, 1560). Some models already contain two power supplies.
11. Choose extras (cable management arm, recovery software media kit).

Learning check

6. Describe the SMB three multichannel features:

- Contracts with dozens of local e-shops and expanding
- What does your server, storage, and network infrastructure look like?
 - Tower servers, multiple switches, including L3
- Do you have an IT department?
 - No, external help on request
- What are your current plans?
 - Deploying a centralized file-oriented solution
 - Consolidation of user data

Customer requirements

As a result of multiple interviews and gathering information about customer plans and the customer's current infrastructure, the following requirements emerged for the new solution:

- Centralized location for users' data consolidation
- 50 TB
- Simplified storage configuration tools
- Efficient deduplication
- Encryption for data at rest
- Possibility for replication and hybrid flash in the future

Summary

- A NAS device is a server that is dedicated only to file sharing. NAS appliances can provide shared storage between heterogeneous clients. Disk arrays and other storage devices connect to the network through a traditional LAN interface such as Ethernet.
- HPE StoreEasy products are optimized to simplify file storage, offer advanced features, and hassle-free management.
- HPE StoreEasy platform is sharing folders or iSCSI targets.

5 Recommending HPE Storage solutions for SMB customers

LEARNING OBJECTIVES

After completing this chapter, you should be able to recommend and position HPE Storage products.

Recommending HPE Storage solutions for SMB

The wide HPE portfolio of data center products for SMB customers can be introduced using customer scenarios. We will introduce two scenarios in this chapter. This is the first scenario.

Introducing the customer scenario

The fictional customer LLP distribution company will be used as a storyline through this scenario. We will introduce the company using an interview:

- What is your primary business?
 - City-wide courier, parcel, and documents delivery
- How many employees do you currently have?
 - 25 employees + external workers
- What does your selling and delivery channel look like?
 - Contracts with dozens of local e-shops and expanding
- What does your server, storage, and network infrastructure look like?
 - Several tower servers, NAS, and new networking devices

- Do you have an IT department?
 - No, external help on request
- What are your current plans?
 - Add local storage

Customer requirements

As a result of multiple interviews and gathering information about customer plans and customer's current infrastructure, the following requirements emerged for the new solution:

- Local application, in house developed
- Data growth higher than anticipated
- Increase the local storage for the server hosting the application
- No local data center, sufficient space in a storeroom
- 16 more spindles to satisfy input/output operations per second (IOPS) requirements
- Serial-attached SCSI (SAS) solid-state drive (SSD) (endurance: mixed used)
- Attached to the server over SAS
- 50 GB of data anticipated per year, 100% growth rate annually
- Design the solution for five years

Activity: Discovering a customer's business and technical requirements

1. Prepare a list of additional questions to ask the customer about:
 - Current environment
 - Technical requirements
 - Room for growth

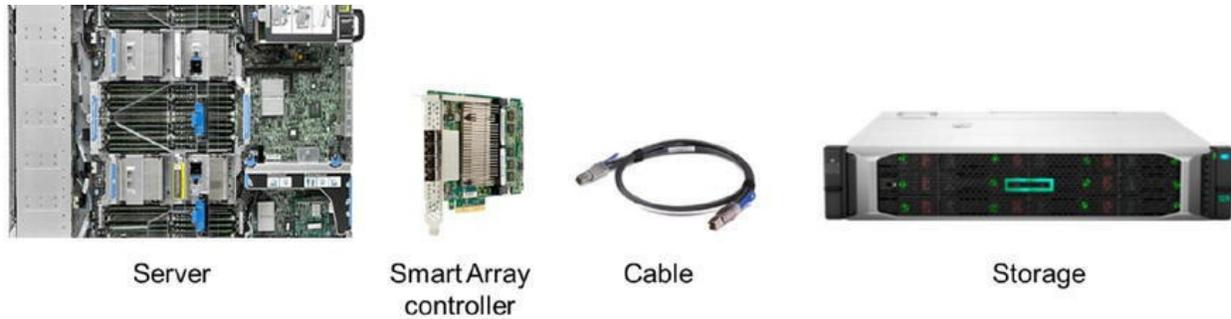


Figure 5-1 Components of Direct-attached storage

Direct-attached storage (DAS) provides the most straightforward HPE storage solutions for small-to-medium businesses (SMBs). However, DAS storage cannot be shared among multiple servers, and the one-to-one connection between device and server can add management complexity.

DAS solutions typically are connected directly to the attached server using the SAS protocol and can be either an internal part of the server or implemented as external enclosures.

RAID functionality is provided by an array controller installed inside the server.

Typical components of a DAS solution, as shown in [Figure 5-1](#), include:

- Server
- Host bus adapter (HBA) or array controller installed in server
- Connectivity, cables
- Storage enclosure

RAID levels

Redundant array of independent disks (RAID) refers to a means of configuring groups of disks in different ways to provide increased performance and/or resilience to drive failure. Common RAID levels are:

- RAID 0 (striping) provides increased capacity and performance but no extra data protection. If one drive is lost, all the data are lost.
- RAID 1 (mirroring) protects against failure of one drive.

- RAID 1+0 (mirroring and striping) protects against failure of one drive (and failure of particular multiple drives).
- RAID 5 (distributed data guarding) protects against failure of one drive.
- RAID 6 with ADG provides the highest level of fault tolerance.
- RAID 50 (RAID 5+0) protects against failure of one drive (and failure of particular multiple drives).
- RAID 60 (RAID 6+0) allows administrators to split the RAID 6 storage across multiple external boxes.
- RAID 1 ADM allows customers to create mirrored RAID set using three identical hard drives.

Factors influencing performance

HPE drive	IOPS (70% read/30% write)
15 K RPM Enterprise HDD	220
10 K RPM Enterprise HDD	165 to 180
7.2 K RPM Midline HDD	100 to 125
Read-intensive SSD	Random read from 54000, random writes from 9800
Mixed use SSD	Random read from 56000, random writes from 20000
Write-intensive SSD	Random read from 100000, random writes from 100000
NVMe high-performance, low-latency, write intensive	Random read from 575000, random writes from 550000

Figure 5-2 HDD and SSD performance comparison

The table shown in [Figure 5-2](#) compares the IOPS of several HPE HDD and SSD devices.

Two factors primarily characterize drive performance: sequential data transfer rate of the media and random IO operations per second. Sequential data transfer rate is typically specified when reading or writing relatively large blocks (≥ 64 kilobytes) of data to sequential disk sectors on the outer most tracks of the disk, called the outer diameter (OD). This maximum sequential data rate is valid only for the OD tracks on the drive. The rate can be 40 to 50% lower on the inner diameter (ID) tracks because they have a smaller diameter with fewer sectors per track. Random IO operations occur when reading or writing relatively small blocks (≤ 8 kilobytes) of data to sectors that are scattered across the whole capacity of the disk. The speed of the actuator seeks, and the disk rotation speed determines, random performance.

The realized performance of a disk drive depends heavily on the nature of the workload. Performance varies when accessing large blocks of sequential data or small blocks of random data. All rotating media disk drives have varying levels of IO request reordering, called optimization. This optimization reduces the combined seek and rotation distance of several outstanding IO requests that the drive has in its request queue. The more requests in the drive's queue, the better the ability to optimize. From a system perspective, the seek distance to random data can also be minimized by maintaining contiguous files on the disk by using appropriate system utilities. Disk file fragmentation can significantly degrade both random and sequential performance.

SSDs and hard disk drives (HDDs) use SAS or Serial Advanced Technology Attachment (SATA) protocols to interface with the host system, but SSDs store and retrieve data in flash memory arrays rather than on spinning media. SSDs have no seek or rotational latency time. They address any sector of the NAND flash directly in 0.1 millisecond. SSD latency includes the time for memory access and transfer combined with controller overhead.

SSDs excel at random read operations, where their performance can be more than 100 times better than that of spinning media drives. SSDs perform random writes at least 25 times faster than a comparable 15 K rpm HDD. This means SSDs provide improved application performance.

Processor and memory performance have grown in step with Moore's Law, getting faster and smaller; however, storage performance has lagged far

behind, creating a significant bottleneck to system performance. Today, more IT dollars are spent on storage-centric applications such as database, data warehousing, and virtualized workloads than on compute-centric applications.

Choosing the best server storage options for a computing environment requires an understanding of storage devices and storage needs. HPE offers enterprise storage options based on two technologies: HDDs and SSDs.

HPE provides server storage solutions using either spinning magnetic media (HDDs) or solid-state NAND flash-based technology (SSDs). HPE enterprise HDDs are designed for use in unconstrained (24 x 7 x 365 up to 100% write) IO workloads.

Solid-state drives are used for mission-critical applications such as large databases, email servers, and back-office applications. They provide maximum reliability, high performance, and error management under the most demanding conditions.

HPE also has classes of SSD solutions:

- Read-intensive SSDs provide relatively large storage capacities at low costs. They are best suited for high read (for example, boot) environments where the workload is high read/low write.
- Mixed use SSDs are suited for high IO applications with a workload that is balanced between read and write.
- Write-intensive SSDs have similar capacities as mainstream SSDs, but they are suited for mission-critical environments with workloads high in both read and write applications.
- Non-Volatile Memory Express (NVMe) SSDs offer high performance and low latency.

The interconnect type used for accessing drives and enclosures is SAS or SATA with a bandwidth of 3 Gb/s, 6 Gb/s, or 12 Gb/s. Drive IO performance is usually saturated first; with more drives in the enclosure, interconnect bandwidth can be saturated as well.

Disk enclosures



Figure 5-3 HPE D3610/D3710, HPE D6020, and HPE D8000

HPE has products to satisfy SMB customers with a range of DAS, network-attached storage (NAS), and storage area network (SAN) solutions. Ideal for company-wide deployment and mission-critical applications, HPE Storage solutions offer maximum scalability, industry-leading performance, a fully integrated suite of centralized management tools, and unmatched data protection and disaster tolerance features.

Disk enclosures are:

- **Affordable**—Cost effectively extend and grow the existing storage capacity.
- **Scalable**—Leverage HPE Smart Array technology with the latest SAS/SATA hard drives and SSDs.
- **Simple**—Enhance productivity with intuitive setup and day-to-day interaction through ProLiant management tools.

Disk enclosures can be used when the internal drive slots of an HPE ProLiant server or HPE BladeSystem are filled. Disk enclosures address growing needs—flexible, modular solutions that simplify capacity expansion.

The disk enclosure portfolio, as illustrated in [Figure 5-3](#), includes the following product families:

- HPE D3000 Disk Enclosures
- HPE D6000 Disk Enclosures
- HPE D8000 Disk Enclosures

Unique features of HPE disk enclosures

HPE provides a complete, end-to-end DAS solution including enclosures, controllers, and servers. Advanced Data Guarding (RAID 6) is the highest level of fault tolerance. It allocates two sets of parity data across drives and allows simultaneous write operations. This level of fault tolerance can withstand two simultaneous drive failures without downtime or data loss. The most important components allow hot-plug operations (remove, replace, and add) while the system is running without service disruption. Redundant power and cooling provide increased reliability because failure of a power supply or fans does not interrupt system operation.

Self-Monitoring Analysis and Reporting Technology (S.M.A.R.T.), first developed at HPE, detects a possible hard disk failure before it occurs, allowing replacement of the component before failure occurs.

HPE D3000 Disk Enclosures: HPE DAS solutions



Figure 5-4 HPE D3710 Disk Enclosures

The HPE D3000 Disk Enclosures, the next-generation 12 Gb SAS low-cost, high-capacity, tiered external storage system, as shown in [Figure 5-4](#), are ideal for small application environments, remote offices, and departmental locations. The 12 Gb SAS enclosures—large form factor (LFF) D3600/D3610 with 12 drive bays and small form factor (SFF) D3700/D3710 with 25 drive bays—offer flexible, modular solutions to simplify capacity expansion of HPE ProLiant server and HPE BladeSystem environments to external storage without having to make the full move to SAN or NAS. This allows your customers to buy what is needed today and purchase additional capacity as data storage needs to grow. D3000 Enclosures support HPE SmartDrive Carriers.

Designed to give customers room to expand as their storage demands increase, the D3600 and the Gen10 branded D3610 LFF can grow to 96

drives. Customers can also add up to eight D3600/D3610 LFF enclosures from a single controller port.

HPE D3000 Disk Enclosures: Target customers

HPE D3000 Disk Enclosures are ideal for small application environments in SMBs, remote offices, departmental locations, Tier 2 or 3 storage for enterprise customers, customers without SAN, and storage option for backups, such as Veeam Backup.

D3000 enclosures are:

- **Affordable and low cost**—Customers can purchase only what they need today and avoid upfront costs. The modular platform provides investment protection, and by purchasing only what is needed today allows for growth as storage requirements evolve. Buying storage only when needed simplifies planning and relieves budget pressures.
- **Flexible and scalable**—Storage capacity can grow easily by attaching enclosures as storage demands increase.
- **Easy to manage**—Little IT expertise is required to expand. Enclosures are compatible with the latest HPE Smart Array SAS controllers and HBAs. Enclosures can be configured using familiar ProLiant management tools provided through the Smart Controller.



Note

For technical and configuration details, see the D3000 QuickSpecs at <https://h20195.www2.hp.com/v2/getdocument.aspx?docname=c04227611>.

Learning check

1. What is the difference between D3610 and D3710?

Recommending HPE Storage solutions for SMB

The wide HPE portfolio of data center products for SMB customers can be introduced using customer scenarios. We will introduce two scenarios in this chapter. This is the second scenario.

Introducing the customer scenario

The fictional customer LLP distribution company will be used as a storyline through this scenario. We will introduce the company using an interview:

- What is your primary business?
 - City-wide courier, parcel, and documents delivery
- How many employees do you currently have?
 - 25 employees + external workers
- What does your selling and delivery channel look like?
 - Contracts with dozens of local e-shops and expanding
- What does your server, storage, and network infrastructure look like?
 - Several tower servers, NAS, and new networking devices
- Do you have an IT department?
 - No, external help on request
- What are your current plans?
 - Adding a shared storage for highly resilient application

Customer requirements

As a result of multiple interviews and gathering information about customer

plans and customer's current infrastructure, the following requirements emerged for the new solution:

- Local application, in house developed
- Application modified for active/standby mode of operation
- Shared storage required for high availability
- Two controllers required
- No local data center, sufficient space in a storeroom
- 16 more spindles to satisfy IOPS requirements
- SAS SSD (endurance: mixed use)
- Attached to the server over internet Small Computer Systems Interface (iSCSI)
- 50 GB of data anticipated per year, 100% growth rate annually
- Design the solution for five years

Activity: Discovering a customer's business and technical requirements

1. Prepare a list of additional questions to ask the customer about:

- Current environment
- Technical requirements
- Room for growth
- Resources
- Management/monitoring
- How many servers are connected to the storage?
- Location of the storage and servers
- Identical rack?
- Current IOPS read/write ratio

2. Write down other topics you might need to discuss with your customer.

Storage area network



Figure 5-5 SAN devices

A storage area network (SAN) is a dedicated high-speed network that interconnects and presents shared pools of storage devices, shown in [Figure 5-5](#), to multiple servers. With the correct configurations a SAN can provide reliability, performance, and resiliency for the storage of data. SANs reduce the cost of managing storage by consolidating storage and sharing the resources across multiple hosts.

Most Fibre Channel (FC)–based storage supports changes made to storage configurations through a host operating system. This support allows storage resources to be added or deleted without interrupting production environments, thereby reducing downtime for storage changes.

Scaling storage can be as simple as plugging a storage array into a switch or adding drives into existing storage systems. This capability allows storage to be added on an as-needed basis.

The benefits of a SAN include:

- **Centralized storage**—Consolidating storage (primary and secondary) in a SAN and sharing the resource across multiple servers reduces the cost of storage management.
- **Elimination of server downtime while adding storage**—Using Fibre Channel-based storage, storage resources can be added or deleted without interrupting the production environment.
- **High availability**—Redundant fabric designs, storage replication, dynamic failover protection, traffic rerouting, and server clustering enable SANs to provide enterprise-class availability to open systems servers.
- **Modular scalability**—With modular scalability providing support for an unpredictable environment, the infrastructure can be changed as business needs evolve. Bandwidth, availability, redundancy, and capacity can be dynamically scaled on demand, providing maximum flexibility to accommodate business growth.

What is a SAN?

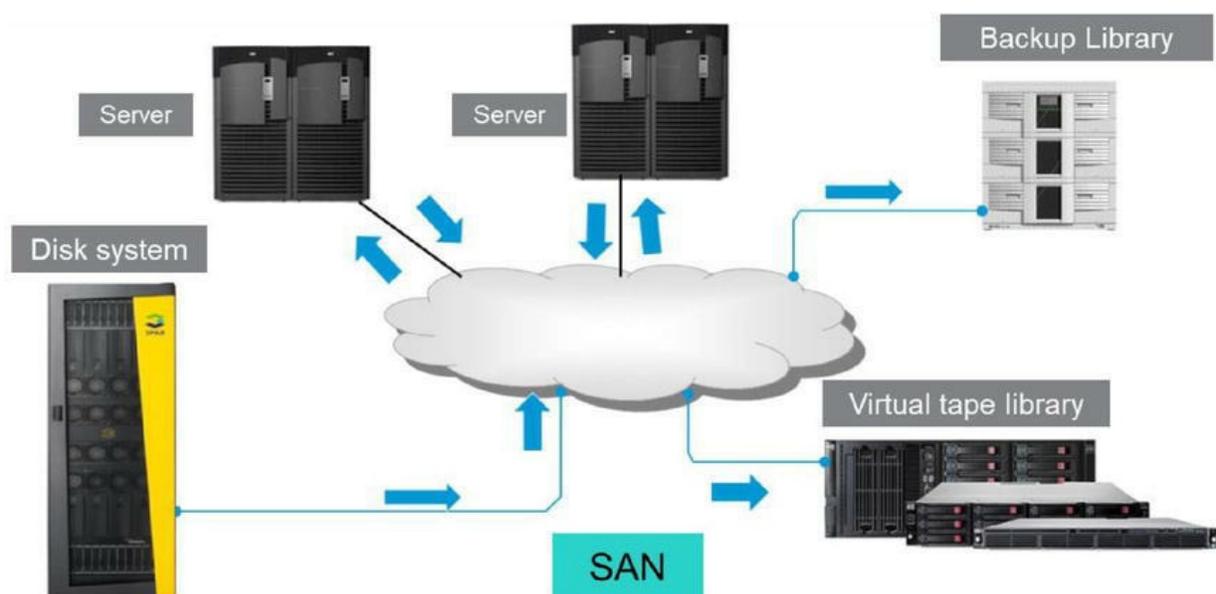


Figure 5-6 A SAN shares resources but not files

A SAN, as illustrated in [Figure 5-6](#), is an intelligent infrastructure that interconnects heterogeneous servers with shared, heterogeneous storage systems. It is a dedicated storage network, designed specifically to connect storage, backup devices, and servers.

SANs remove heavy storage server traffic from the local area network (LAN), freeing up bandwidth for client server traffic. They also provide the data communication infrastructure required by advanced, cost-efficient mass storage systems. They provide scalability of network and storage devices and serve as a platform for clustering and high availability.

Fibre Channel is the typical standard transport layer for SANs, with SCSI-3 providing the upper layer protocol.

SANs are maintained separately from parallel general-purpose networks and are isolated from the messaging network. They are optimized for movement of data from server to disk and tape.

SANs use multiple paths to connect different storage devices with associated servers and can provide backup and archival storage for multiple or remote locations. HPE StorageWorks SANs deliver value-added storage applications, management tools, storage arrays, and virtualization technology.

General-purpose networks, such as LANs and wide area networks (WANs), carry heavy user communications traffic involving printers, email, and so forth. A SAN is the back-end network that carries storage traffic, which provides a clear separation of storage devices from processing and presentation activities and enables the front-end LAN to carry normal Transmission Control Protocol/Internet Protocol (TCP/IP) traffic.

SANs have become almost synonymous with Fibre Channel; however, Fibre Channel is not a required component, because almost any networking or serial Small Computer System Interface (SCSI) technology can be used to create a SAN. In addition, the Fibre Channel protocol is designed to carry not just SCSI traffic, but also TCP/IP traffic and other protocols.

Definition of a SAN

The Storage Network Industry Association (SNIA) defines a SAN as “A network whose primary purpose is the transfer of data between computer systems and storage elements, and among storage elements. A SAN consists of a communication infrastructure which provides physical connections, and a management layer which organizes the connections, storage elements, and computer systems so that data transfer is secure and robust.”

Within this definition, there is no mention of Fibre Channel. SNIA recommends using the term Fibre Channel SAN when the network is based on Fibre Channel technology. Although a SAN typically references Fibre Channel, it can be based on other technologies such as Enterprise System Connection (ESCON) and iSCSI, which enables a server to access storage through an Ethernet network adapter.



Note

For more information about SNIA, refer to the website at:
<http://www.snia.org>.

Switch Infrastructure Recommendations—iSCSI SAN

There is a set of minimum switch capabilities that make building a high-performance, fault-tolerant storage network relatively easy and cost-effective task. Any enterprise-class managed switch typically has the necessary capabilities most IP-SAN customers require. HPE recommends the HPE FlexFabric 5700, 5900, which includes the 5920, and 5930 family switches for ToR convergence, and the FlexFabric 7900 and 12900 Switch Series for core converged switches.

In an iSCSI packet-switching network, a hop is the trip a data packet takes from the initiator (host) to the target (storage system). A hop is one portion of the path between source and destination. Each time packets are passed to the next device, a hop occurs. On an iSCSI network that uses TCP/IP, the number of hops a packet has taken toward its destination (called the hop count) is kept in the packet header. A packet with an exceedingly large hop count is discarded, and this should be avoided. This hop count value is a measure of distance and is also referred to as the diameter of the IP SAN fabric.

Known as time to live (TTL) in IPv4, and Hop Limit in IPv6, the field specifies a limit on the number of hops a packet is allowed before being discarded. A fixed TTL value must be at least big enough for the network “diameter,” that is, the longest possible path. The purpose of the TTL field is to avoid a situation in which an undeliverable datagram keeps circulating on a system, and such a system eventually becoming swamped by such “immortals.” A hop count of n means that n switches separate the source host from the destination. By itself, this metric is not useful for determining the optimum network path, as it does not take into consideration the speed, load, reliability, or latency of any particular hop. Since store and forward and other latencies are incurred through each hop, a large number of hops between source and destination imply higher latency and lower real-time iSCSI performance.

A FlexFabric iSCSI SAN does not have a hard hop limit, but keep in mind that latency can be a cause of slow IO performance and iSCSI disconnects. Of course, many other things can add to latency, including hops, congestion, and distance. HPE recommends a best practice iSCSI SAN network diameter (or hop count) below seven, and as small as possible. Consider edge-core or edge-core-edge topologies instead of a many hop network topology. A small iSCSI SAN diameter, or low hop count, enables the spanning-tree algorithm to have enough time to converge. Also, the bridge protocol data unit (BPDU) needs time to be propagated across the network. BPDUs are data messages that are exchanged across the switches within an extended network that uses a spanning-tree protocol topology. BPDU packets contain information on ports, addresses, priorities, and costs and ensure that the data ends up where it was intended.

Activity: Requirements for switches used in iSCSI SAN

1. Download HPE 3PAR StoreServ iSCSI best practices from:
<https://support.hpe.com/hpsc/doc/public/display?docId=c04873990>
2. Answer the following questions:
 - a. What is a priority-based flow (PFC) protocol?

b. What does Data Center Bridging Exchange (DCBX) protocol do?

Use cases for a SAN

There are many drivers to migrate current storage environment to SAN, including:

- Requirement for higher availability
- Better utilization of backup and restore solutions
- Business continuance
- Server and storage consolidation

High availability

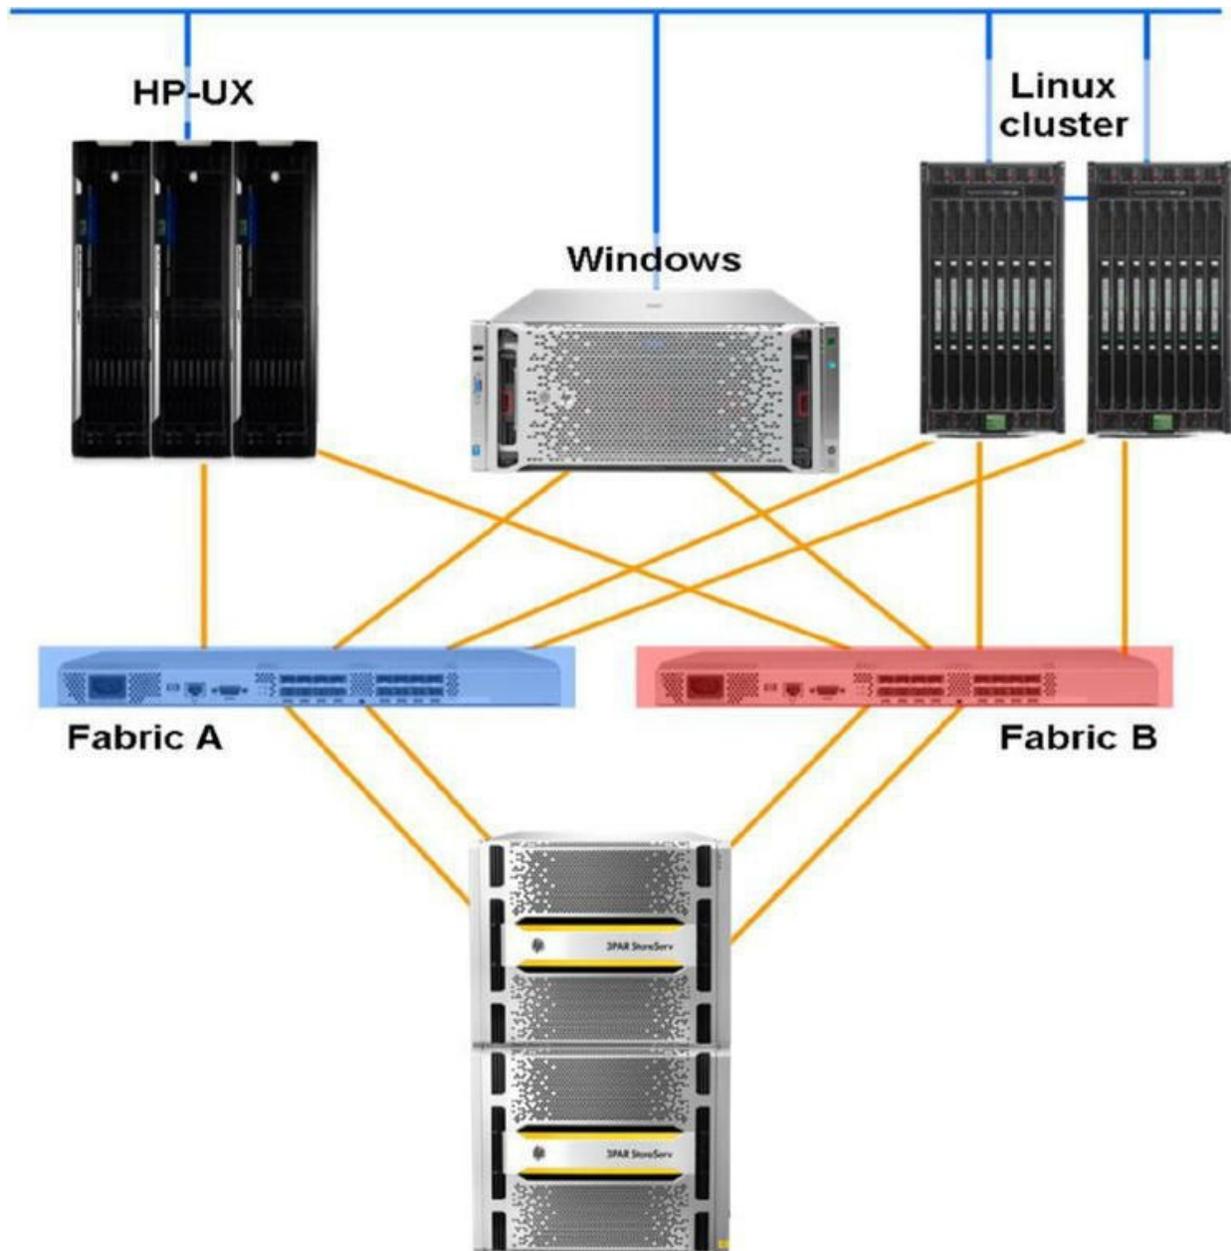


Figure 5-7 A SAN should always be redundant and resilient in design

Well-designed SANs are fault-tolerant with no single point of failure, as indicated in [Figure 5-7](#). Redundant fabric designs, server clustering, storage replication, and traffic rerouting provide enterprise-class availability. Because the long-distance capabilities of Fibre Channel can be used to mirror data to a remote location, SANs also can help facilitate disaster recovery.

SANs employ fiber optic and copper connections to create dedicated

networks for servers and their storage systems. The SAN enables several servers to access a common data set for server clustering arrangements.

- If an application fails, it can be restarted on another server without having to move the data.
- If a server fails, SANs enable fast failover. This means that operations automatically transfer to another node in the cluster. Again, data does not have to be moved as part of the failover.

SANs provide an automatic method for transferring operations from a failed or down system to a secondary system.

Redundant design

Robust high-availability architecture requires redundant components throughout: servers, host adapters, switches, and storage units.

In the extended architecture in the preceding graphic, both servers are equipped with dual HBAs, so that failure of the adapter alone does not disable the server. Both servers are connected to both switches, so the loss of one still leaves the servers with an intact path to storage. If an entire storage unit fails, the redundant unit enables the servers to continue to access the application data.

Data mirroring

Data mirroring is the process of copying data from one location to another storage device in real time, resulting in an exact copy, or mirror, of the data at that point in time. SAN data mirroring capabilities take advantage of Fibre Channel distances to allow the replication of key systems at a remote site.

Clustering

The emergence of SANs is enabling new system configurations that can leverage shared storage. One common system configuration is the shared nothing or availability cluster, in which multiple servers (nodes) share common storage and client access and can take over the processes of a node in the cluster that fails. The primary advantages that a SAN provides a cluster

are connectivity, scalability, and reliability.

Shared nothing clusters can be implemented on Microsoft Windows or UNIX systems. The servers in these clusters do not access the same data concurrently. Instead, one server owns and accesses data in the common storage. That capability can be passed to another server if the first server should fail.

A shared data cluster takes a different approach, allowing nodes in a cluster to access the same data at the same time. Although sharing data introduces the problem of maintaining data integrity, it also leverages the capabilities of the SAN to consolidate data, improve availability, and support scalable solutions.

High-availability clusters are used to support critical business applications. They provide a redundant, fail-safe installation that can tolerate equipment, software, and network failures, and continue running with as little impact on business as possible. High-availability clusters have been in use for several years; however, until the advent of Fibre Channel, they were limited in size and reliability. Clusters require shared storage, and sharing SCSI storage subsystems is difficult and unreliable.

Clustering technology has been improved by the network architecture of SANs. SANs provide the ability to interconnect a large number of devices. SANs can support as few as two hosts in a failover configuration and can be expanded to support “many-to-one” configurations.

This design can be used in combination with dual-attached hosts and storage devices to keep an application running even if one fabric fails because of operator error, catastrophe, or other issues. Upgrades can be performed on one fabric without path downtime. This is the best design approach for high-availability environments.

Both resilient and non-resilient dual fabrics can be referred to as redundant fabric SANs. Redundant designs are always recommended for high-availability systems and any large SAN deployment where downtime for the entire SAN could affect hundreds of servers.

Better utilization of backup and restore solutions

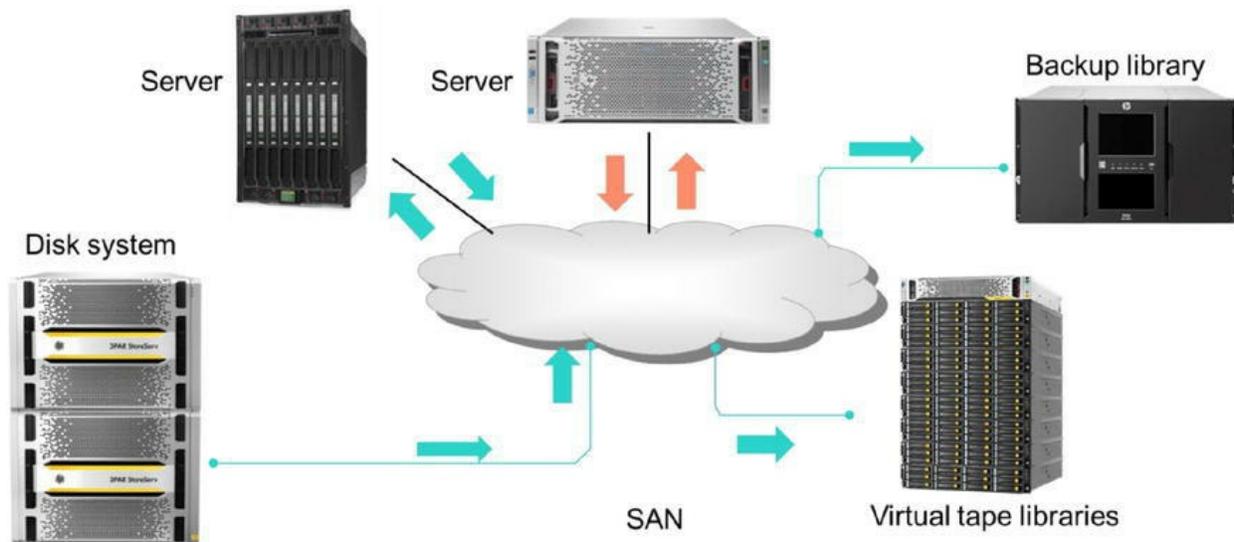


Figure 5-8 Utilizing the SAN for backups and restore operations

The amount of data being stored, the speed of access to that data, and the amount of time that data is kept online are all growing exponentially. Backup and recovery operations are coming under increasing pressure from collapsing backup windows and growing storage requirements, as illustrated in [Figure 5-8](#).

Because SANs remove backup and recovery traffic from the LAN, congestion is reduced, and backup windows are improved. System performance is dramatically increased because data and communications traffic no longer compete for the limited bandwidth on a standard LAN.

Centrally managed, high-performance storage resources such as tape libraries reduce backup time and overhead.

SANs also can be configured for serverless (active fabric) backup. Serverless backups use the Network Data Management Protocol (NDMP), the standard for backing up NAS, and SCSI Extended Copy to move the data directly from disk to tape. This enables continuous, uninterrupted access to data and applications during the backup and restore processes.

This technology also eliminates the read/write processes through the application and backup hosts, resulting in dramatically reducing processing power. By removing the LAN, CPU, and IO resources from the data path, network bottlenecks are eliminated, and high-application performance is maintained. Serverless backup is ideal for large databases or file servers for

which there is no backup window.

Business continuance

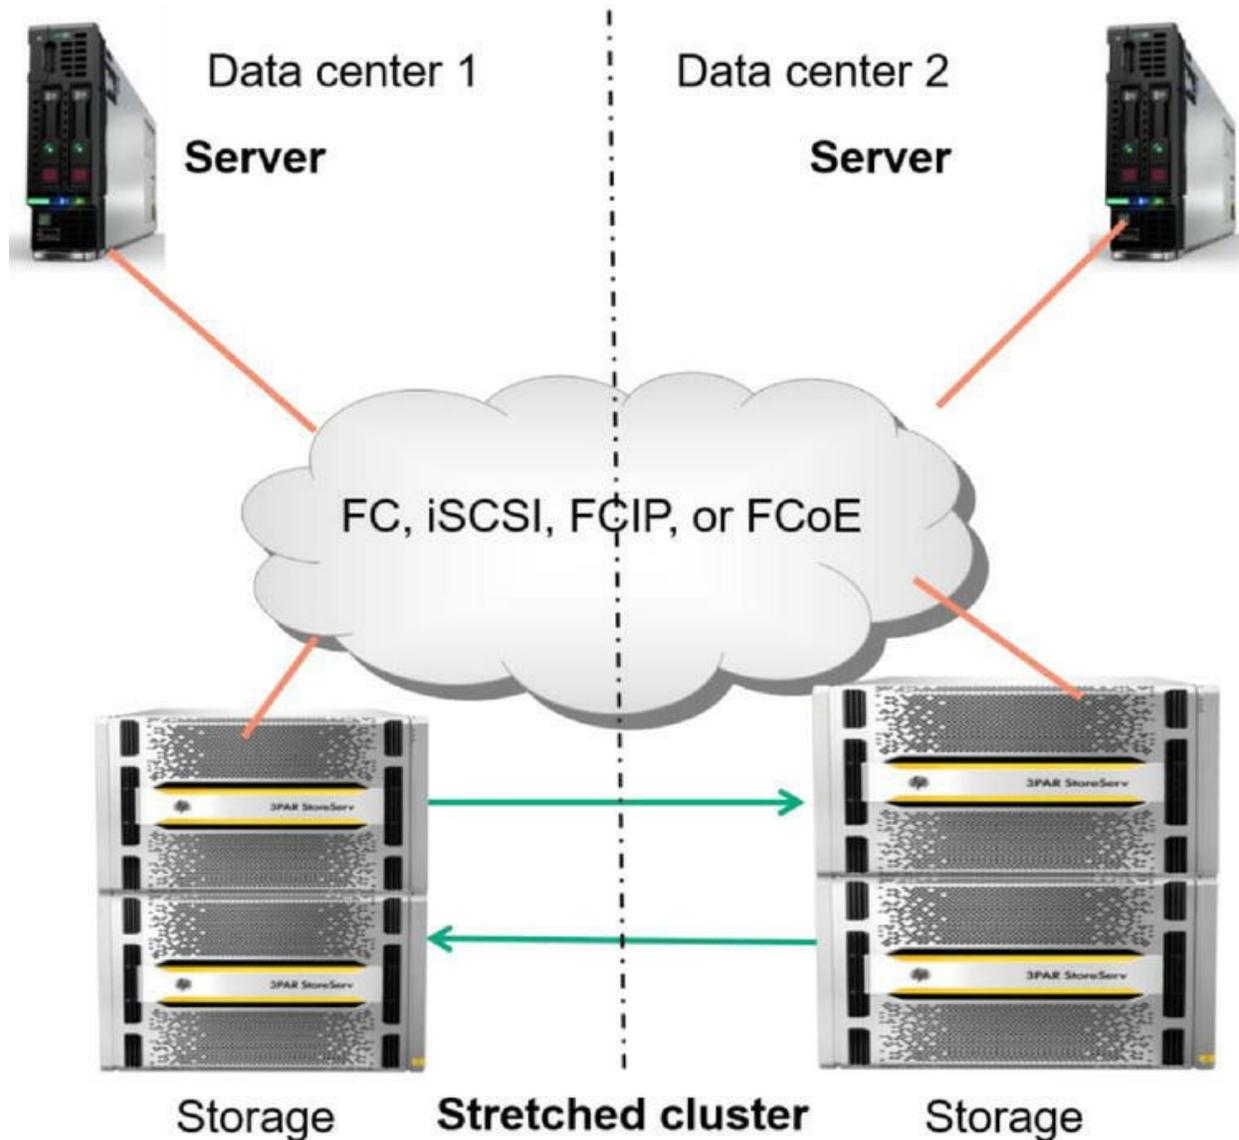


Figure 5-9 Multi-site replication to eliminate single points of failure

SANs can eliminate single points of failure, incorporate failover software, and support mirroring at geographically dispersed data centers for disaster recovery, as indicated in [Figure 5-9](#). Businesses can quickly restore productivity after a power failure or other cause of downtime.

Business continuance (sometimes referred to as business continuity) describes the processes and procedures an organization puts in place to ensure that essential functions can continue during and after a disaster. Business continuance planning seeks to prevent interruption of mission-critical services and to re-establish full functioning as swiftly and smoothly as possible. Although business continuance is important for any enterprise, it might not be practical for any but the largest to maintain full functioning throughout a disaster crisis. According to many experts, the first step in business continuity planning is deciding which of the organization's functions are essential, and apportioning the available budget accordingly. After the crucial components are identified, failover mechanisms can be put in place. New technologies, such as disk mirroring over the internet, make it feasible for an organization to maintain up-to-date copies of data in geographically dispersed locations, so that data access can continue uninterrupted if one location is disabled.

A business continuance plan should include a:

- Disaster recovery plan, which specifies an organization's planned strategies for post-failure procedures
- Business resumption plan, which specifies a means of maintaining essential services at the crisis location
- Business recovery plan, which specifies a means of recovering business functions at an alternate location
- Contingency plan, which specifies a means of dealing with external events that can seriously impact the organization

Server and storage consolidation

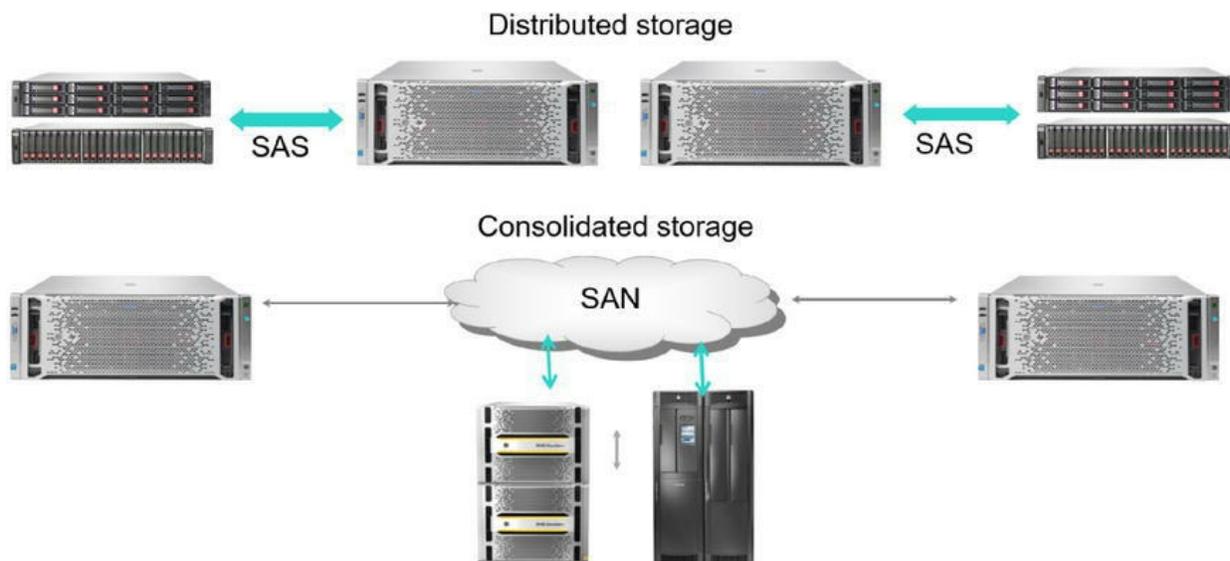


Figure 5-10 Comparing distributed to consolidated storage solution

Often SAN implementations begin as storage consolidation efforts and evolve into storage expansion mechanisms as needs grow, as illustrated in [Figure 5-10](#).

High connectivity enables storage for many servers to be consolidated on a small number of shared storage devices, reducing costs and easing management of capital assets.

When storage is managed as a pooled resource that can be allocated to servers in large or small amounts, administrators can deploy the most cost-effective storage products without concern about what fits into the limited slots available in the servers.

Adding more storage to a SAN is easy and does not normally result in loss of service for the application server, other than what it takes to mount the new storage. Servers can be added and removed, and storage reassigned, while the SAN is online. You also can increase the number of switches in a fabric to provide the number of ports required, or you can use routing technology to connect multiple SAN fabrics.

SAN components

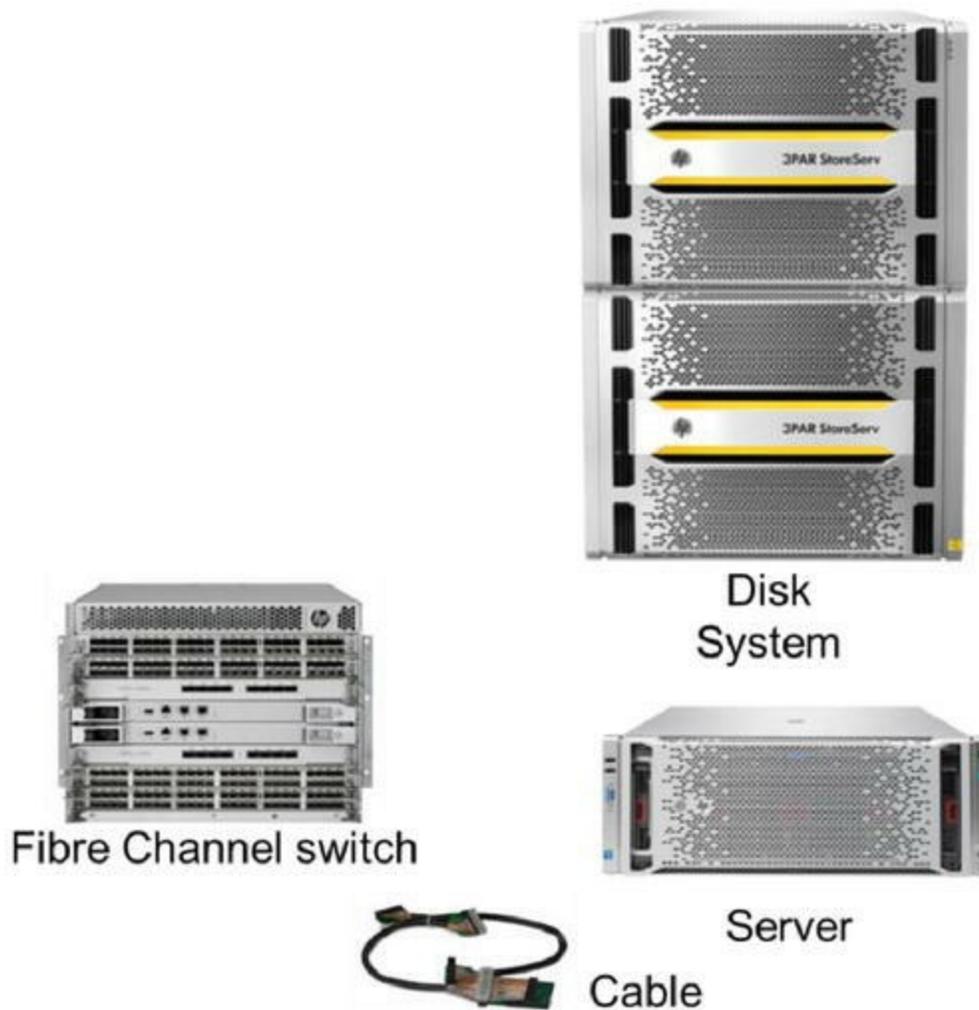


Figure 5-11 The basic components for creating a SAN

A SAN consists of both hardware and software components including, as indicated in [Figure 5-11](#):

- **Switches**—Switches create the fabric of the SAN. Fabric means a network of one or more Fibre Channel switches that transmit data between any two N_Ports on any of the switches. Switches can be added together to increase the size of the fabric and increase the number of supported devices. iSCSI-based SANs are using Ethernet or converged switches.
- **Routers, bridges, and gateways (not shown in graphic)**—The key role of routers is to enable device sharing, utilize multi-protocol technologies, and provide fault isolation. Routers can extend the SAN over long

distances, providing greater disaster recovery from site failures.

- **Storage devices**—A SAN storage solution will have an HBA or built in SFF pluggable transceiver (SFP) to connect to a switch. It will have dual controllers for resiliency and performance and RAID technology to recover from disk failures.
- **Servers**—Servers connect to the SAN with HBAs or Converged Network Adapters (CNAs). The server can be running any operating system provided it supported on the SAN infrastructure.



Note

Any device that is using the SAN for intelligent purposes is referred to as a node. Therefore, storage devices, a server or backup devices, are all nodes, and all connect to the SAN as N_Ports.

- **Cabling and connectors**—These can be fiber optic or copper, and care has to be taken in being aware of the limitations on distance and speed supported. All SFPs are tri-speed. For example, if an SFP is a 16 GB, it can run at speeds of 16 GB, 8 GB, and 4 GB. iSCSI-based SANs are using Ethernet connectors and cables, such as 10BaseT or 10GbE SFP+.
- **SAN management applications**—The aims in developing a SAN are to maximize resources to combine management tasks and to streamline the IT processes.

HPE offers an ideal range of Fibre Channel and multiprotocol or converged protocol switches for SAN and LAN that deliver the scalability, performance, and broad interoperability that missioncritical data and applications demand. Combining ease of use, flexibility, and improved security with cost-effective pricing, these switches are key enablers of a dynamic and adaptive storage networking environment. They also provide an ideal platform for adaptation as business needs grow and change— both in terms of basic connectivity and advanced SAN and/or converged fabric functionality. The SAN switch portfolio includes devices from multiple vendors, as indicated in 5-13.

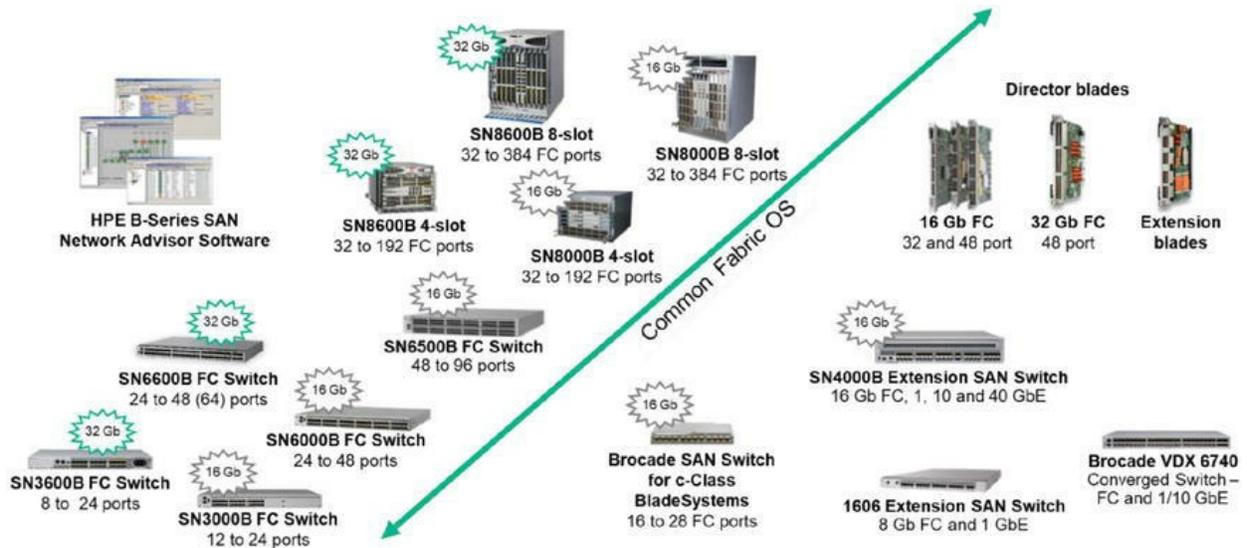


Figure 5-12 HPE B-Series SAN portfolio

In a switched fabric Fibre Channel SAN topology the host server and target storage devices are typically restricted from communicating with each other by default; in these cases, Fibre Channel zones are required to allow specific members contained in the zone to communicate.



Figure 5-13 HPE SN3000B 16 Gb Fibre Channel Switch

Zoning can be applied to either the WWN of the N-Port connected to the SAN fabric or to the specific switch port to which the end device is connected. WWN-based zoning typically allows the end device to be moved to another port without having to reconfigure the existing zones.

It is recommended to create two physically redundant SAN fabrics (such as “Fabric A” and “Fabric B”) to protect against hardware failures and to provide highly available connections to redundant systems.



Note

More information can be found in the HPE SAN Design Reference Guide: <https://support.hpe.com/hpsc/doc/public/display?docId=c00403562>

HPE storage array positioning



Figure 5-14 HPE's storage array portfolio

As shown in [Figure 5-14](#), the HPE storage array portfolio includes the following:

- **XP7**—HPE XP7 Storage is designed for mission-critical applications. Its array-based virtualization increases availability, simplifies disaster recovery (DR), and improves resource utilization by helping eliminate storage silos.
- **HPE Primera**—Intelligent storage for mission-critical apps that delivers extreme resiliency and performance with the agility of the cloud.



Note

The HPE XP7 and HPE Primera storage arrays are enterprise-level solutions and will not be covered in this guide.

- **3PAR StoreServ**—HPE 3PAR StoreServ was built to meet the extreme requirements of massively consolidated cloud service providers. With over 3 M IOPS and consistent sub-ms latency, 3PAR transforms midrange and enterprise deployments with solutions from a few TBs to more than 20 PB scale.
- **Nimble**—HPE Nimble Storage leverages flash storage and predictive analytics to deliver greater than 99.9999% guaranteed availability.
- **SimpliVity 380**—SimpliVity is a powerful, simple, and efficient hyperconverged infrastructure (HCI) platform that joins best-in-class data services with the DL380 Gen10.
- **MSA**—MSA storage includes flash-enabled arrays that raise the entry storage bar, making application acceleration possible for a wide range of budgets.

HPE Nimble Storage

Nimble Storage Arrays, illustrated in [Figure 5-15](#), are available from HPE and HPE partners. Nimble arrays are ideal for organizations seeking a new approach to storage that combines affordable flash performance with simplicity for mixed, mainstream workloads.

HPE Nimble Storage portfolio overview



Figure 5-15 HPE Nimble Storage

Nimble secondary storage arrays enable your customers to put their backup data to work with flash performance to support development, test, and analytics on copy data.

A Nimble array tolerates three simultaneous drive failures and provides additional protection from intra-drive parity. Application-consistent snapshots and replication as well as integration with leading backup software solutions provide an ideal foundation for comprehensive data protection. Backed by the Timeless Storage guarantee that ships with all Nimble Storage arrays, there is no need to pay for optional software, and upgrades are easy.

Nimble is built on the patented Cache Accelerated Sequential Layout architecture. One of the key points of CASL is how it populates flash. Traditional storage vendors typically wait for a cache miss to populate the cache. For example, if a user requests something that was not already in cache and then wants to pull it into cache, traditional vendors wait for the first miss, which unfortunately results in the first users not reaping the benefit of the flash. Nimble uses 22 patented algorithms that watch the data as it

comes between the storage processor and the bus down to the disc and inspects it to determine how random it is, and at that moment one copy gets laid down on disc and one on flash. This allows the first person that requests the data to immediately realize the benefit of flash.

Essentially, Nimble CASL evaluates how random the data is and makes a decision whether to put it in flash. Because of this, Nimble is able to get very high cache hits with only a small percentage of cache. This is how they have been able to get all-flash-like performance since 2008 and will serve as the foundation for all SCM-like performance in the future.

All-Flash Array family

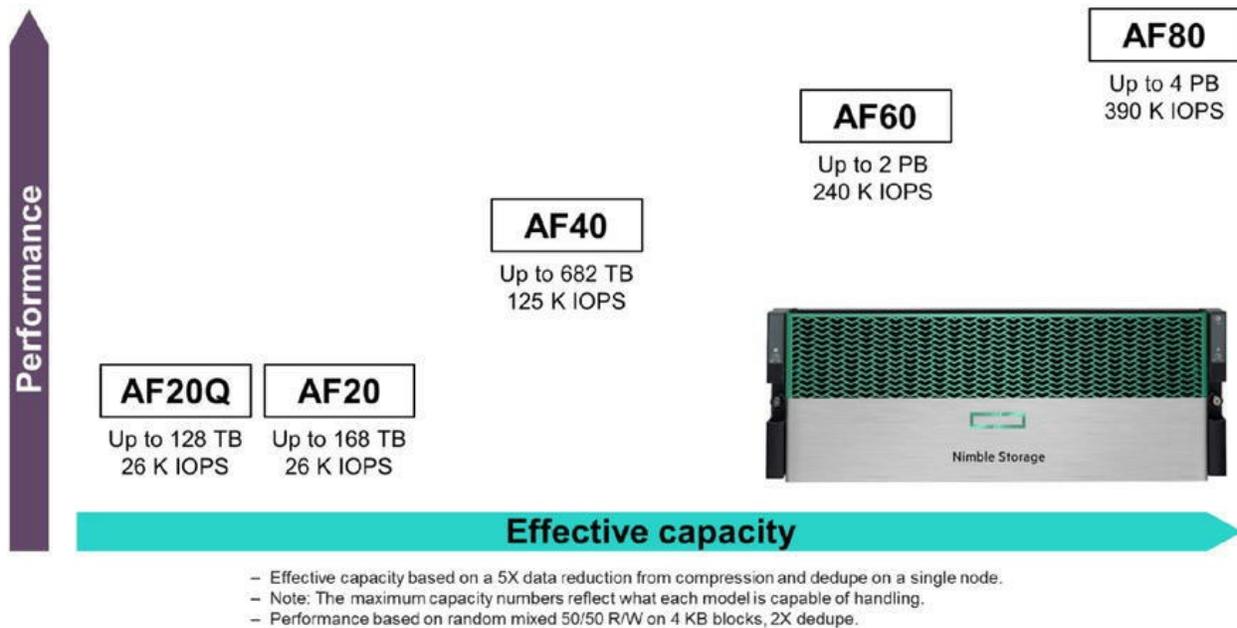


Figure 5-16 Nimble All-Flash storage family

As shown in [Figure 5-16](#), the Nimble All-Flash storage family. The Nimble All-Flash storage family includes several models that vary in both performance and maximum supported capacity:

- The AF20Q and AF20 are the perfect entry points for all IT organizations that require speed and economy for performance-sensitive workloads.
- The AF40 and AF60 offer high performance and attractive economics for

performance-sensitive workloads that need the best blend of price, performance, and scalability.

- The AF80 is designed for consolidating multiple large-scale performance-sensitive applications with aggressive performance and high scalability demands.



Note

HPE recommends HPE racks with a depth of 1200 mm to best accommodate the length of the Nimble Storage chassis; the HPE 1200 mm rack provides ample room for cabling and ease of serviceability.

What is in the AF and HF series arrays?

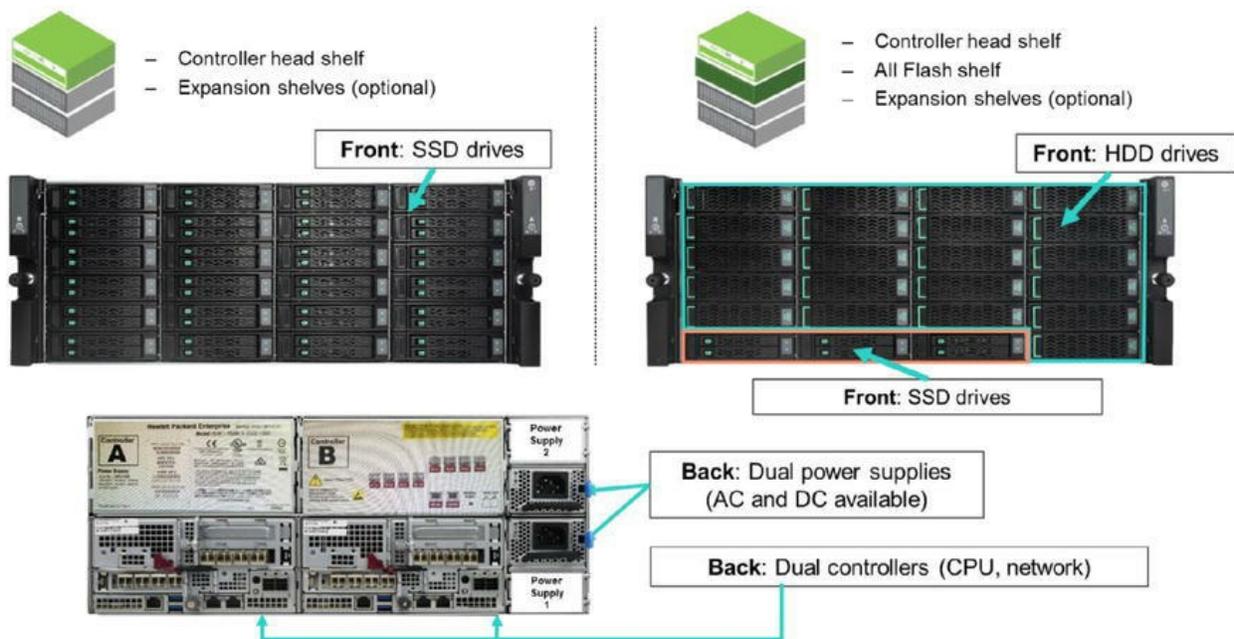


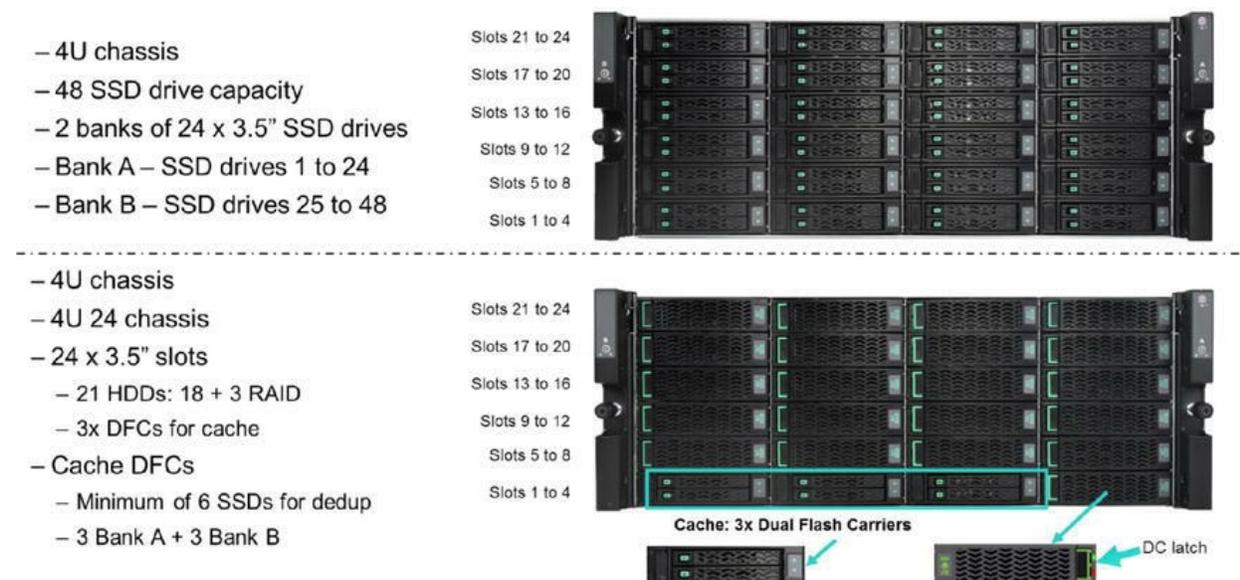
Figure 5-17 AF and HF Arrays

As shown in [Figure 5-17](#), a high-level overview of what is in a Nimble Storage AF (on the left side of the graphic) and HF-Series array (on the right side of the graphic) and what can be connected to a controller shelf. There are three primary components in each Controller Head Shelf:

- SSDs (and HDDs in the HF), located in the front of the chassis
- Dual power supplies, located in the back of the chassis
- Dual controllers, located in the back of the chassis

The Controller Head Shelf can then have the All-Flash Shelf and the Disk Expansion Shelves connected to it. Both the All-Flash Shelf and the Disk Expansion Shelves are optional.

AF and HF chassis front view and drive layout



which can be populated with 21 HDDs and three DFCs (Dual Flash Carriers) for a total potential of six SSDs. The minimum number of SSDs that can be configured with these arrays is three, which populates Bank A in the DFC.

The chassis slots are numbered from left to right and top to bottom. Slots 1 through 3 contain the DFCs, and slots 4 through 24 are populated with disk carriers.

For workloads that require peak performance, Nimble HF storage systems allow a pinning of the entire volume in the cache. When a volume is pinned in cache, the HF-Series arrays will write data to the hard disk drives and maintain a copy of the full volume in cache.

Specifications include:

- 4U 24 chassis
- 24 × 3.5” slots carry 21x HDDs + 3x DFCs
- HDDs: 18 + three RAID (new Nimble-branded HDD carriers)
- DFCs: A minimum of three SSDs in Bank A (Bank B is available for cache upgrades)

AF-Series write operations: Steps 1 to 4

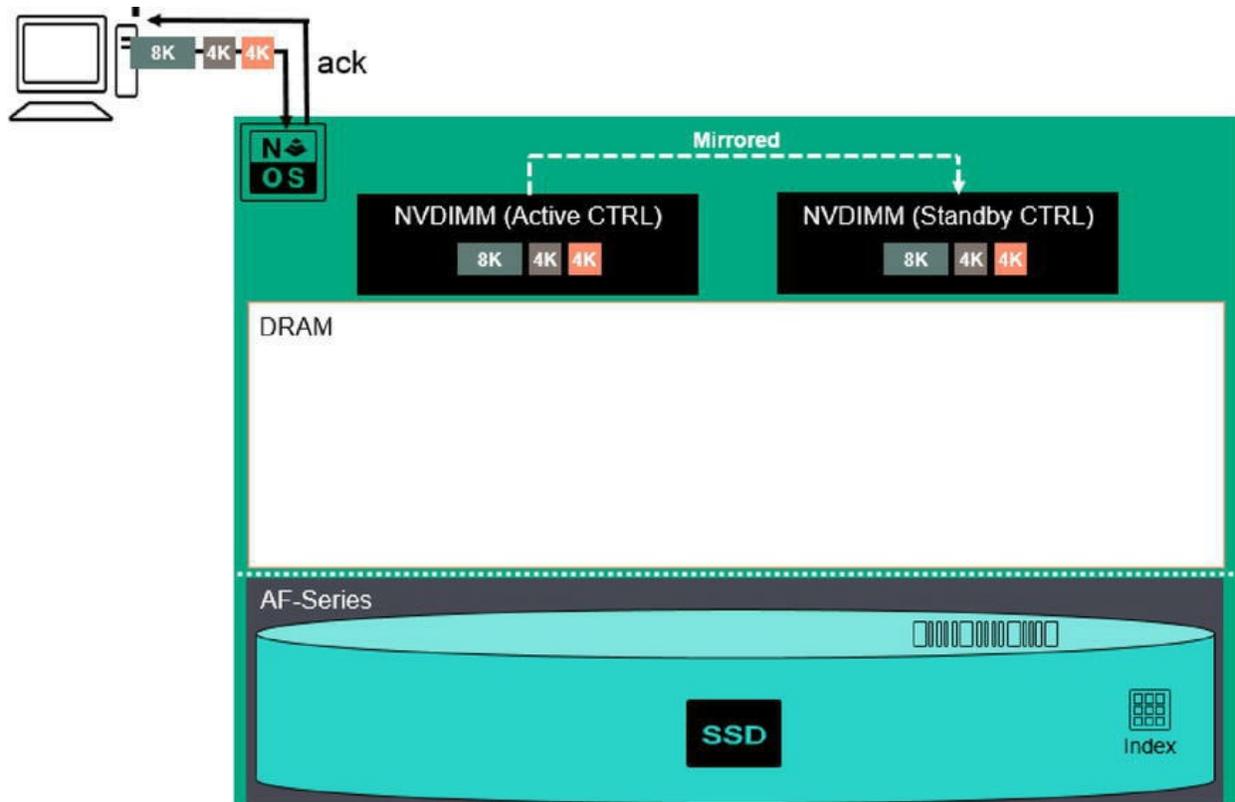


Figure 5-19 Data mirroring

NimbleOS is a unique architecture that has been designed to accommodate a variety of platforms. Currently, Nimble Storage offers two: the Hybrid Flash Array family (HF-Series) and the All-Flash Array family (AF-Series). You see a representation of NimbleOS as well as some of the hardware components found in Nimble controllers, including non-volatile dual in-line memory module (NVDIMM) and dynamic random-access memory (DRAM). At the bottom of the graphic, the AF-Series primary storage components are represented.

To help you better understand NimbleOS, we are going to first review the write process, illustrated in [Figure 5-19](#). When a write is received by the controllers, it is first accepted by the NVDIMM. From there, it is mirrored to the standby controllers NVDIMM before it is acknowledged. At this point, the array has two copies of the data in UltraCap secured NVDIMM. NimbleOS then acknowledges the write.

Steps:

1. Writes are sent by a variety of applications in variable block sizes.
2. NimbleOS places incoming writes into the active controllers NVDIMM, which is backed up via an UltraCap.
3. NimbleOS mirrors the active controllers NVDIMM to the standby controllers NVDIMM.
4. NimbleOS acknowledges the write.



Note

A supercapacitor (SC), also called an ultracapacitor, is a high-capacity capacitor with a capacitance value much higher than other capacitors, but with lower voltage limits, that bridges the gap between electrolytic capacitors and rechargeable batteries. It typically stores 10 to 100 times more energy per unit volume or mass than electrolytic capacitors, can accept and deliver charge much faster than batteries, and tolerates many more charge and discharge cycles than rechargeable batteries ([Wikipedia.org](https://en.wikipedia.org), October 2019).

AF-Series write operations: Step 5

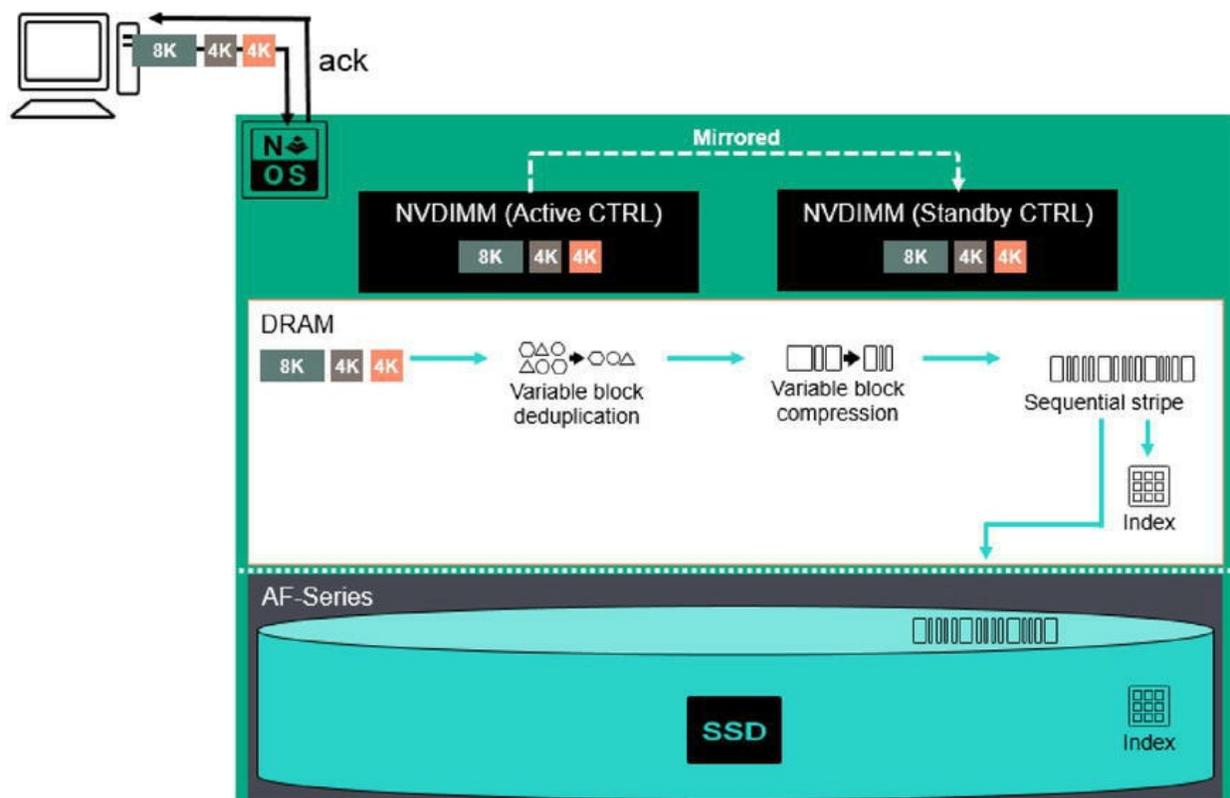


Figure 5-20 Write operation

Blocks are then copied into the DRAM. This memory space is used to perform a variety of functions. What occurs at this level depends on the product family. In this section, we are going to focus on the All-Flash Array (AF-Series).

After the blocks arrive in the NVDIMM, NimbleOS uses variable block deduplication which pairs down the total number of blocks that will be written to media, as illustrated in [Figure 5-20](#). NimbleOS then compresses remaining data blocks to further pair down the volume of data written to media. Next, the blocks are placed into a 10 MB sequential stripe write. This stripe is then sequentially written to the SSDs. Simultaneously, an index of the block locations is maintained in the NVDIMM, with backup copies maintained on SSDs. The index allows for extremely quick lookups when a read request is made.

Steps (continued from the previous section):

5. Blocks are copied into DRAM.

- Variable block deduplication is applied.
- Variable block compression is applied.
- Blocks are formed into a sequential stripe write.
- Blocks are written to SSD and are indexed in DRAM.

Compression engines will examine the contents of a block of data looking for opportunities to shrink that data block by recognizing sub-blocks of repeating zeros and will remove the superfluous data. Most compression algorithms, like the LZ4 compression algorithm, have a look up table of known patterns. The block will be inspected, looking for those known patterns, and eliminate them.

AF-Series read operations

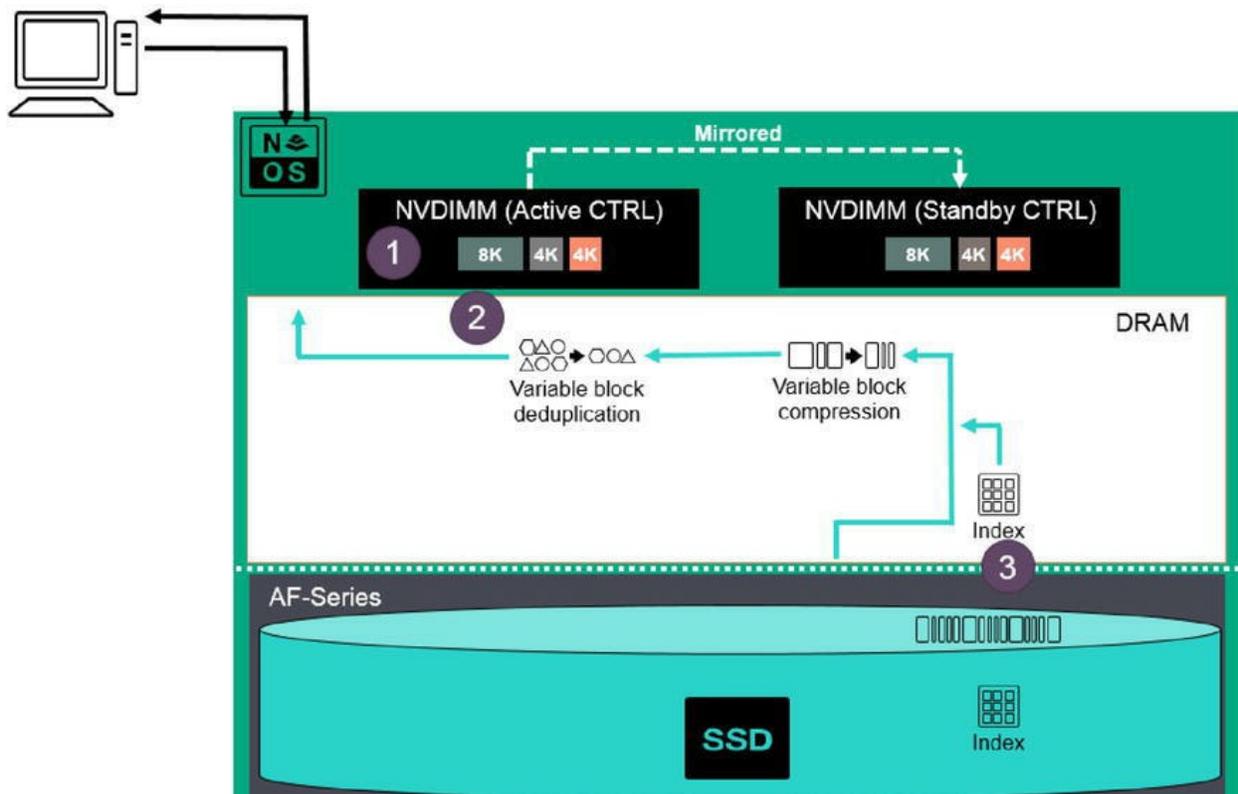


Figure 5-21 AF-Series read operation example

Now let us look at how reads are processed in the Nimble AF-Series Array,

as indicated in [Figure 5-21](#):

1. Read from NVDIMM.
2. If not found, check DRAM
3. If not found, read from SSD using the index for a quick lookup. If found, validate checksum, decompress, dedupe, and return data.

When a read request is sent to the array, the NimbleOS architecture will retrieve the data from the fastest area possible.

NimbleOS starts by checking in NVDIMM. If a write was just sent to the NVDIMM and is still there, it will be retrieved very quickly. If it is not in NVDIMM, NimbleOS will check DRAM, because that is the next fastest location from which the read can be retrieved. If the data is not located in the DRAM, it will pull the read from the SSD. It will then validate the checksum, decompress the data, dedupe the data, and return the data.

It is important to note that a checksum happens on both reads and writes to help ensure data integrity.

AF-Series triple+ parity and integrated spare

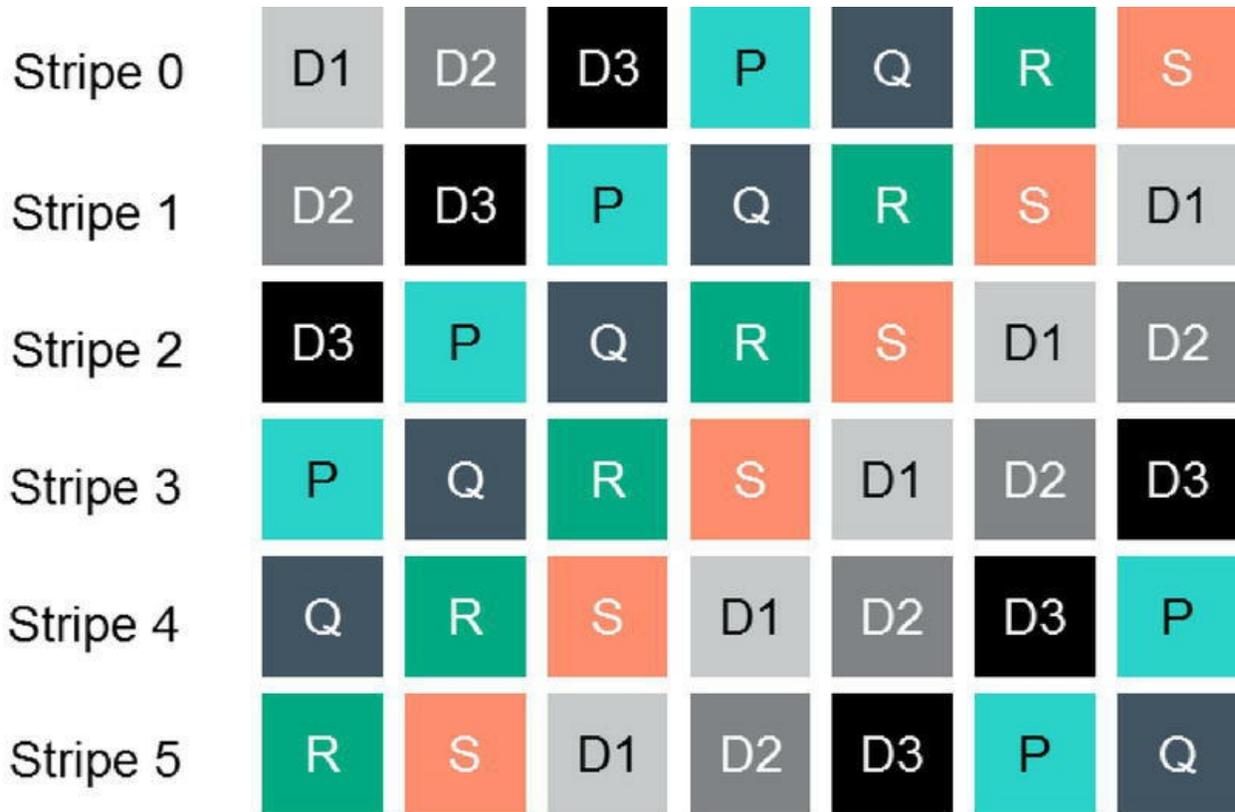


Figure 5-22 Triple+ parity layout

What is triple+ parity RAID? The first is the use of triple parity. Triple parity means that three drives' worth of parity information is written, as shown in [Figure 5-22](#).

Every chunk in a stripe consists of data, metadata (checksums, sequence number, and so forth), and parity (XOR of data). The system tolerates a simultaneous failure of any three SSDs, and it has a built-in virtual spare that allows a fourth failure.

Intra-drive parity fixes sector loss in single read.



Note

Triple+ parity RAID with integrated spare is used with All-Flash Arrays only.

SCM/NVMe-ready HPE Nimble Storage All-Flash platform

HPE Timeless program—Our commitment to future-proofed storage platforms

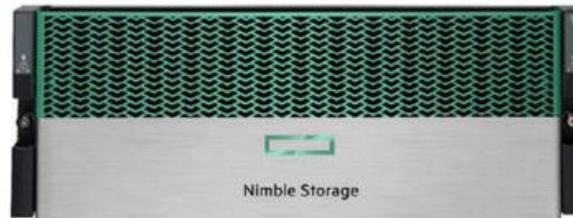


Figure 5-23 HPE Timeless program

HPE is committed to future-proofed technologies and storage platforms. All Nimble AFAs are Storage Class Memory (SCM)/NVMe ready and have a technology roadmap that unlocks the full potential of NVMe. Non-disruptive, seamless upgrades to next-generation technology will be facilitated, as indicated in [Figure 5-23](#).



Note

For more information, visit:
<http://hpe.com/storage/MemoryDrivenFlash>

NVMe and SCM

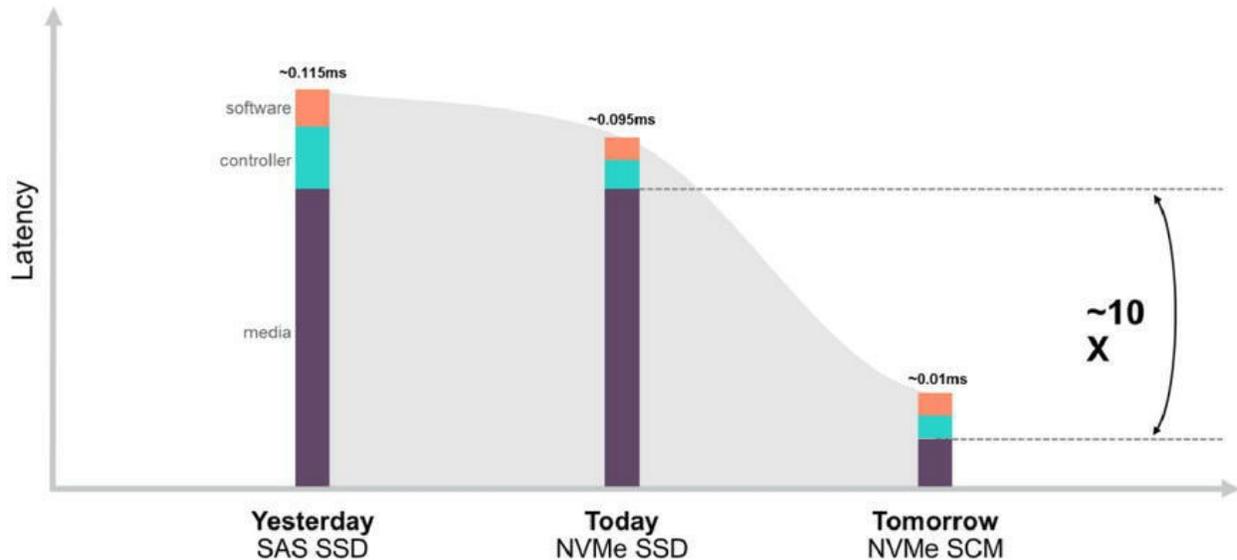


Figure 5-24 NVMe and SCM

There are several emerging technologies that are being implemented at the product level these days, as indicated in [Figure 5-24](#). These include:

- **NVM Express**
 - Is a high-performance, scalable host controller interface.
 - Is designed for enterprises, data centers, and client systems that use Peripheral Component Interconnect express (PCIe) SSDs.
 - Can be adopted across multiple platforms.
 - Improves random and sequential performance.
 - Reduces latency.
- **SCM**
 - Hybrid storage and memory access
 - Typically used as a cache/tier layer
 - Performance between DRAM and SSD

Activity: SCM and NVMe for HPE 3PAR and Nimble Storage

1. Watch this video: <https://www.youtube.com/watch?v=hKDcNWpipes>.
2. Take notes about SCM performance and typical workloads.

Activity debrief

Answer the following questions:

1. What is an example of a workload that may benefit from SCM?

2. Is the latency typically consistent over different workloads (IOPS) with current storage technologies?

3. What is the latency difference between SSDs using NVMe and SAS protocols?

4. What is the latency difference between NVMe SSD and NVMe SCM technologies?

- What is the expected latency reduction when used in an array?

5. Name three customer benefits of SCM.

Learning check

2. Name at least three typical use cases for SAN.

3. How many controllers does AF20Q have?

- A. 0 (external controllers are needed)
- B. 1
- C. 2
- D. 4

Alternative components for different environments



Figure 5-25 HPE MSA Storage, Apollo 4510, and HPE SimpliVity

We presented a validated solution for customer scenario. Some of the solution components, however, can be changed for alternative environments, where different performance, capacity, costs, features, and similar requirements can be satisfied with other HPE solutions. Not all the workloads are the same and a single HPE product cannot match all the requirements.

Alternative components, as indicated in [Figure 5-25](#), for other SMB segments include:

- Software-defined storage
 - HPE SimpliVity
 - HPE StoreOnce VSA
 - HPE ProLiant servers with partner solutions
- HPE MSA Storage

Software-defined storage



Figure 5-26 Apollo 4510 and HPE SimpliVity

Software-defined storage (SDS) is an alternative to traditional arrays, where

storage functionality is provided as a software. Shared storage can easily coexist with applications on servers and capacity is provided by storage connected to the server, usually using internal drives or external disk enclosures.

Examples of SDS, as shown in [Figure 5-26](#), include:

- HPE SimpliVity
- HPE StoreOnce VSA
- HPE ProLiant servers support:
 - VMware virtual SAN (vSAN)
 - Microsoft Storage Spaces Direct
 - Hadoop/Hadoop Distributed File System (HDFS)
 - Ceph, OpenStack Swift, Qumulo, Scality
 - Cohesity, Datera, Hedvig
 - And many other solutions



Note

Consult QuickSpecs and SMB-validated solutions for more details and limitations.

HPE SimpliVity Data Virtualization Platform

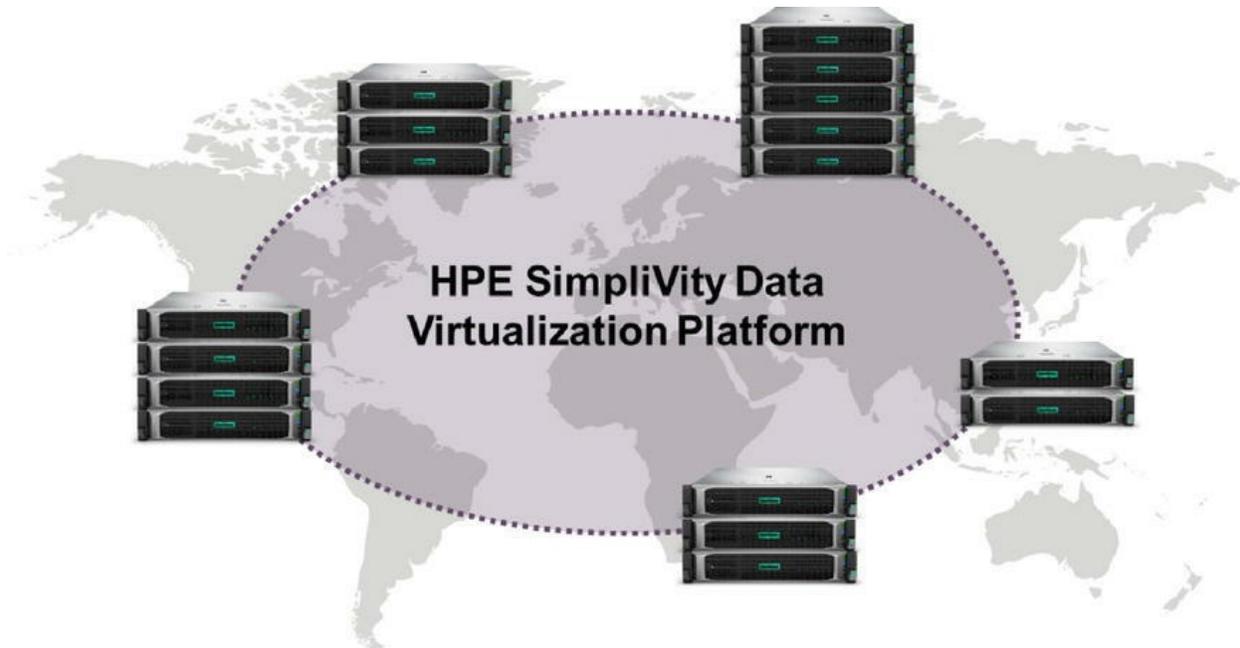


Figure 5-27 HPE SimpliVity Data Virtualization Platform

HPE SimpliVity Data Virtualization Platform HPE SimpliVity infrastructure's core technology is the HPE SimpliVity Data Virtualization Platform, a globally aware file system and object store that facilitates globally shared resource pools across multiple sites. The Data Virtualization Platform, as illustrated in [Figure 5-27](#), is characterized by the advanced data services it delivers guaranteed data efficiency, built-in resiliency, backup, and disaster recovery, and global VM-centric management and mobility.

HPE SimpliVity 380 Gen10



Figure 5-28 HPE SimpliVity 380 Gen10

HPE SimpliVity 380 (available on HPE ProLiant DL380 Gen10 servers), as shown in [Figure 5-28](#), is a compact, scalable 2U rack-mounted building block that delivers server, storage, and storage networking services. Adaptable for diverse virtualized workloads, the secure 2U HPE ProLiant DL380 Gen10 delivers world-class performance with the right balance of expandability and scalability. It also provides a complete set of advanced functionalities that enables dramatic improvements to the efficiency, management, protection, and performance—at a fraction of the cost and complexity of today’s traditional infrastructure stack.

HPE SimpliVity 380 Gen10 is scalable, configurable, and expandable in one-node increments to up to 32 total nodes. The OmniStack Virtual Controller (OVC) is a VM-based controller running on a hypervisor. It is the foundational software of the HPE SimpliVity HyperConverged infrastructure. The OmniStack Accelerator Card (OAC) is a purpose-built PCIe-based accelerator card designed by the HPE SimpliVity team. The OAC performs inline data efficiency processing, deduplication, compression, optimization, and write acceleration, as VMs write to their virtual disks.



Note

It is strongly recommended that an HPE SimpliVity 380 Gen10 cluster have a minimum of two nodes to realize the full benefit of the high-availability hyperconverged solution. A single cluster supports up to eight nodes, while a single Federation supports up to 32 nodes.

HPE SimpliVity 380 Gen10 offers:

- A choice of Intel® Xeon® Scalable Processors
- From 144 GB to 1536 GB of memory (per node)
- 10 Gb or 1 Gb network
- Redundant power
- Virtualization software and licensing
- A compact 2U form factor

The HPE SimpliVity 380 Gen10 software features include:

- VMware vSphere or Hyper-V
- HPE SimpliVity management plugin for VMware vCenter or Microsoft System Center
- Cluster expansion without downtime
- Hot-pluggable SSD
- Integrated storage controller with battery-backed cache
- HPE ProLiant iLO Advanced Remote Management

HPE SimpliVity 380 Gen10 G and H models



Figure 5-29 HPE SimpliVity 380 Gen10 G

HPE SimpliVity 380 Gen10 G, as shown in [Figure 5-29](#), is a software-optimized model:

- 6 × 1.92 TB SSD Kit only
- Memory: 128 GB to 3072 GB
- Without hardware accelerator
- Expanded scale: Up to 16 nodes per cluster and up to 48 clusters per federation
- Starting from \$3655 (US ILP is preliminary and subject to change without notice)

HPE SimpliVity 380 H is a storage optimized node for backup and recovery with LFF hard drives:

- 4x 1.92 TB SSD

- 8x 4 TB HDD to 25 TB usable

Hardware-optimized and software-optimized nodes are supported in the same cluster.

HPE SimpliVity 2600



Figure 5-30 HPE SimpliVity 2600

HPE SimpliVity 2600, as shown in [Figure 5-30](#), is based on Apollo 2600 and provides a different HPE SimpliVity platform choice for dense solutions for compute centric workloads.

The solution does not use the HPE OAC; it is 100% software optimized with always-on dedupe and compression. HPE is the first vendor to deliver consistent performance with always-on dedupe and compression enabled by software—tested by HPE Labs and Login VSI.

VDI use case includes:

- Architected and optimized for best VDI configuration performance
- HPE SimpliVity makes VDI simple to implement and manage
- VMware offering

HPE SimpliVity 2600 offers:

- Four times denser for space-constrained environments
- 100% software-defined (dedup/compression)
- Optimized for compute-intensive workloads

- Use cases:
 - General virtualization
 - VDI
 - ROBO
 - Retail locations
 - Internet-of-Things (IoT) systems
 - Other edge computing environments
- Global, unified scalability
- Enhanced operational efficiency for HPE SimpliVity

HPE MSA Storage

The screenshot displays the 'VOLUMES' section of the Storage Management Utility V3. On the left, an 'Action...' menu is open, listing various operations such as 'Create Virtual Volumes', 'Modify Volume', 'Add to Volume Group', 'Delete Volumes', and 'Create Snapshot'. The main area shows a table of four volumes, all with a health status of 'OK'. Below the volume table, there are sections for 'Related Snapshots' and 'Related Maps', both currently showing 'No snapshots for selected volumes'.

Name	Health	Class	Size	Allocated	Group	Pool	Type	Snapshots	Maps	Schedules
Vol0001	OK	Virtual	29.9GB	29.9GB	-ungrouped-	A	base	0	1	0
Vol0002	OK	Virtual	99.9GB	100.6MB	-ungrouped-	A	base	0	1	0
Vol0003	OK	Virtual	74.9GB	96.4MB	-ungrouped-	B	base	0	1	0
Vol0004	OK	Virtual	1199.9GB	150.9MB	-ungrouped-	A	base	0	1	0

Figure 5-31 Storage Management Utility V3

Fourth- and fifth-generation MSA arrays support a powerful set of data services enabled with a new virtualized array system architecture. MSA customers have access to a number of these new features through a firmware upgrade and paid license feature:

- **Thin provisioning**—This feature allows storage allocation of physical storage resources only after they are consumed by an application.
- **Automated tiering (archive tiering)**—Uses a real-time IO engine to intelligently position data on the most appropriate “tier” of storage. MSA customers leveraging the archive tiering functionality will be able to

better use storage resources by moving data that has not been accessed for long periods of time to more cost-effective midline SAS HDDs.

- **Virtualized Snapshot services**—Created specifically for the virtualized storage architecture of the MSA, the MSA will be able to use redirect-on-write technology to provide better snapshot performance and enable customers to create more complex snapshots with simplified snapshot management.
- **Storage Management Utility (SMU) V3 web user interface**—The new virtualized array architecture also comes with a powerful, embedded SMU, as indicated in [Figure 5-31](#), that will simplify setup and management of the MSA and these new features. Users familiar with the previous MSA web user interface (UI) will have the option to continue with it or take advantage of the new UI.
- **Wide striping**—By leveraging virtualized storage “pools,” MSA users can leverage wide striping technology to simplify volume expansion and to facilitate full utilization of all resources allocated to a specific volume.
- **SSD read cache**—Feature that extends the MSA controller cache. Read cache is most effective for workloads that are high in random reads. A maximum of two SSDs per pool can be added for read cache.

HPE MSA-automated tiered storage

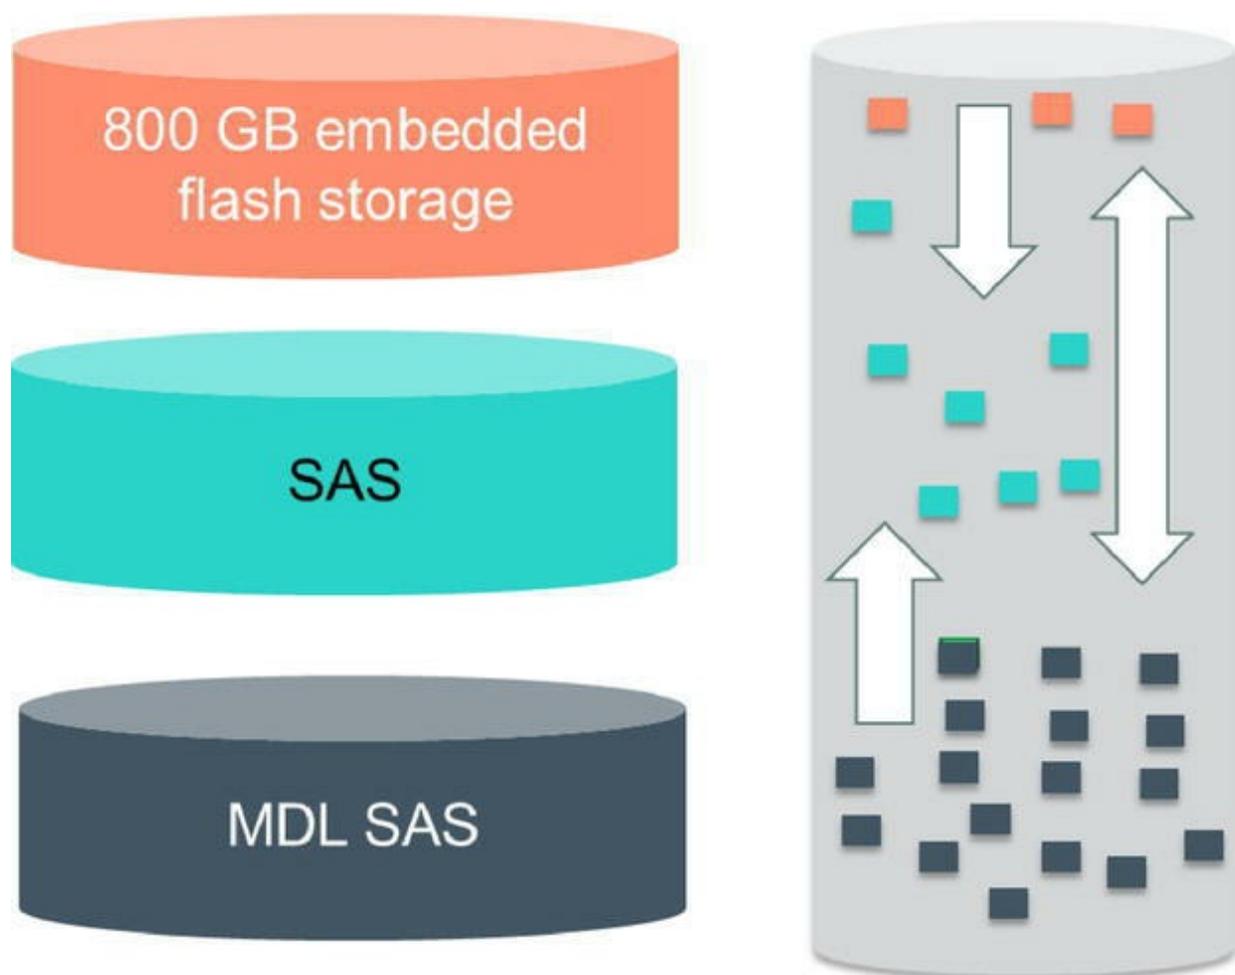


Figure 5-32 Automated tiered storage

Automated tiered storage, as shown in [Figure 5-32](#), is a virtual storage feature that automatically moves data residing in one class of disks to a more appropriate class of disks based on data access patterns:

- Frequently accessed, “hot” data can move to disks with higher performance, lower capacity, and higher costs.
- Infrequently accessed, “cool” data can move to disks with higher capacity, lower performance, and lower costs.

Each virtual disk group, depending on the type of disks it uses, is automatically assigned to one of the following tiers:

- **Performance**—This highest tier uses SAS SSDs, which provide the best performance but also the highest cost.

- **Standard**—This middle tier uses enterprise-class spinning SAS disks, which provide good performance with mid-level cost and capacity.
- **Archive**—This lowest tier uses midline spinning SAS disks, which provide the lowest performance with the lowest cost and highest capacity.

Some advantages of using automated tiered storage are:

- Because a virtual pool can have multiple disk groups, each belonging to a different tier, a virtual pool can provide multiple tiers of storage.
- The IO load is automatically balanced between components in a tier.
- Virtual disk groups can be added or removed without disrupting IO. Data in virtual disk groups that are being removed is automatically migrated to other disk groups if the other disk groups have enough storage space for it. If they do not have the space, the system will not delete the disk groups until enough data is removed.

HPE MSA volume tier affinity feature

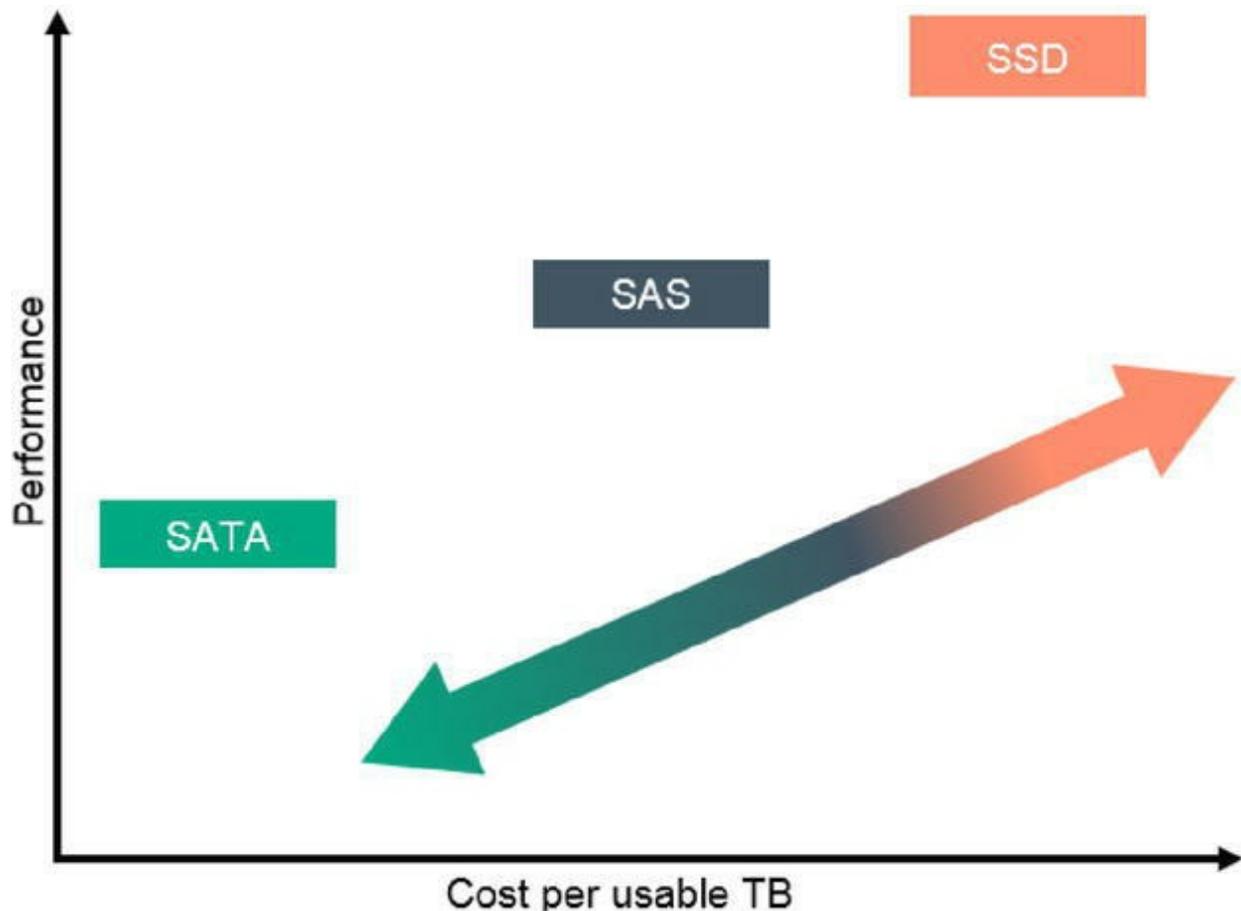


Figure 5-33 There are three volume tier affinity settings: no affinity, performance, and archive

The volume tier affinity feature enables tuning the tier-migration algorithm for a virtual volume when creating or modifying the volume so that the volume data automatically moves to a specific tier, if possible. If space is not available in a volume's preferred tier, another tier will be used.

There are three volume tier affinity settings, as illustrated in [Figure 5-33](#):

- **No affinity**—This setting uses the highest available performing tiers first and only uses the archive tier when space is exhausted in the other tiers. Volume data will swap into higher performing tiers based on frequency of access and tier space availability. This is the default.
- **Archive**—This setting prioritizes the volume data to the least-performing tier available. Volume data can move to higher performing tiers based on frequency of access and available space in the tiers.

- **Performance**—This setting prioritizes volume data to the higher performing tiers. If no space is available, lower performing tier space is used. Performance affinity volume data will swap into higher tiers based upon frequency of access or when space is made available.

HPE MSA 1050 Storage



Figure 5-34 The MSA 1050 is an entry consolidation and virtualization storage solution

The HPE MSA 1050 Storage, as shown in [Figure 5-34](#), brings affordable flash storage down to ever lower customer segments. The MSA 1050 is designed to meet entry-level storage requirements and budgetary constraints. With the lowest storage array price points in the HPE Storage portfolio and field-proven ProLiant compatibility, the MSA 1050 is the platform of choice for smaller SAS, iSCSI, and FC deployments.

The HPE MSA 1050 features 8 Gb Fibre Channel, 12 Gb SAS, and 1 GbE and 10 GbE iSCSI at previously unattainable entry price points. The array allows users to take advantage of the latest storage technologies in simple and efficient ways by providing a good balance between performance and budget resulting in a highly favorable \$/GB return on their investment.

The MSA 1050 provides affordable application acceleration. It is seriously simple and affordable flash-ready storage to help you get the most performance for the lowest cost.

Features include:

- Factory-configured, dual controller array with four host ports of either 8 Gb FC, 1 GbE, 10 GbE, or 12 Gb SAS protocols

- SFF or LFF models available
- Simple, proven Gen10 ProLiant compatible shared storage
 - Affordable, easy to use shared storage without a big learning curve
 - Intuitive setup and management for ProLiant administrators.
 - Look-and-feel will fit right in with Gen10 ProLiant servers
- Access to MSA 2050 advanced data services with simplified licensing approach
 - Leverage flash storage with easy to use, easy to maintain tiering and read cache technologies
 - No storage expertise necessary, system dynamically responds to workload changes
- Expandable and upgradable to meet new demands
 - Add new HDDs or disk expansions to MSA 1050
 - Upgrade system to the MSA 2050 through data-in-place upgrade—a simple 15-minute upgrade
 - No time-consuming migrations required

HPE MSA 2050 SAN Storage



Figure 5-35 MSA 2050—Eight models are supported

The HPE MSA 2050 SAN NEBS Certified DC-Power SFF Storage system is designed for network equipment providers (NEPs) and communication service providers. The MSA 2050, as shown in [Figure 5-35](#), is suited for those customers who need a robust telecom infrastructure. The NEBS compliant MSA 2050 Storage system supports configurations with up to seven compliant disk enclosures for a maximum of 192 SFF HDDs or SSDs.

The HPE MSA 2050 SAN DC-power LFF Storage system includes two (DC) power supplies, but is not NEBS certified. The two power supplies are designed to operate over the input range of -40 VDC to -75 VDC.

The HPE MSA 2050 SAN DC-power Carrier Grade SFF Disk Enclosure is a special model disk enclosure designed for use with NEBS compliant MSA 2050 configurations. This drive enclosure has 24 drive bays and has dual -48 VDC-power supplies. It is only sold with carrier grade arrays.

When used in conjunction with specific storage SFF SAS drives, the solution is NEBS certified (GR-63-Core and GR-1089-Core) and Seismic Zone 4 rated. NEBS level-3 certification provides the assurance that the equipment is safe to operate and sturdy enough to withstand certain physical and environmental conditions (for example, fire and earthquakes). For Seismic Zone 4 rating, the MSA 2050 must be mounted in an HPE Seismic Rack.



Note

NEBS is the abbreviation for the Network Equipment Building System, a set of technical requirements aimed at increasing the availability of network switches. The standard was developed internally at Bell Labs.

HPE MSA 2050 SAN Storage features

The MSA 2050 supports a maximum of seven disk enclosures (either LFF and/or SFF). Add-on enclosures can either be HPE MSA 2050 LFF Disk Enclosure or HPE MSA 2050 SFF Disk Enclosure. The MSA 2050 can grow incrementally to a maximum of 96 LFF, 192 SFF drives, or a combination of SFF and LFF enclosures up to the maximum of eight total enclosures. Virtual storage disks groups can be spanned across multiple enclosures. The Virtual Storage RAID levels supported are RAID 1, RAID 5, RAID 6, and RAID 10.

The maximum hard drive counts vary by RAID levels:

- **RAID 1**—Two drives maximum
- **RAID 5, RAID 6, and RAID 10**—A maximum of 16 drives can be used

The maximum logical unit number (LUN) size is 140 TB (128 TiB).



Note

SFPs are not included in MSA SAN solutions, and SFP are not required with SAS controllers.

Eight models are supported:

- HPE MSA 2050 SAN DC Power LFF Storage
- HPE MSA 2050 SAN NEBS Certified DC Power SFF
- HPE MSA 2050 SAN Dual Controller LFF
- HPE MSA 2050 SAN Dual Controller SFF
- HPE MSA 2050 SAS Dual Controller LFF Storage

- HPE MSA 2050 SAS Dual Controller SFF Storage
- HPE MSA 2050 SAS NEBS SFF Storage
- HPE MSA 2050 SAS DC-Power LFF Storage

Other features include:

- Each controller has the following:
 - 8 GB read/write cache
 - Battery-free cache backup with super capacitors and compact flash
 - Storage Management Utility V3
 - Thin Provisioning support
 - Multiple disk groups that can be aggregated into a single storage pool
 - Storage pools allow data on a given LUN to span across all drives in a pool
 - Owners of MSA 2040, MSA 2042, and MSA 1040 arrays can do data-in-place upgrades to the MSA 2050 array
 - Impressive performance figures:
 - 200,000+ random read
 - 100,000+ random write
 - >5 GB/Sec

HPE MSA 2052 SAN Storage

All MSA 2052 models come standard with the Advanced Data Services Software Suite LTU, which includes automatic tiering, 512 snapshots, and remote snaps for remote replication.

Multiple disk groups can be aggregated into a single storage pool. Storage pools allow data on a given LUN to span across all drives in a pool. When capacity is added to a system, the user is also getting a performance benefit of the additional spindles.

The two 800 GB SSDs can be optionally deployed either as read cache or as an SSD virtual disk group for tiered storage. Customers can deploy the two

embedded SSDs as Read Cache to improve random read performance. A maximum of two SSDs are supported per pool providing a maximum of 4 TB of read cache per controller and a maximum of 8 TB of read cache per array. Customers can also deploy the two embedded SSDs as tiered storage.

The MSA 2052 can manage up to three tiers of storage: performance tier, standard tier, and archive tier. The Performance Auto Tiering License necessary to create an SSD virtual disk group for both read and write capabilities comes standard with all MSA 2052 models. No additional software license purchase is necessary.

HPE MSA 205X highlights

MSA 2050 Array Performance ¹	HPE MSA 2050 Converged SAN Controller with HDDs	HPE MSA 2050 Converged SAN Controller with Mixed Use SSDs
Protocol (host connect)	16 Gb Fibre Channel	16 Gb Fibre Channel
MSA 2050 RAID 1 SSD Performance Results ²		
Random Reads (IOPs)		220,600
Random Writes (IOPs)		102,800
MSA 2050 RAID 5 Performance Results ^{3,4}		
Segmented Sequential Reads (MB/s)	5,280	
Segmented Sequential Writes (MB/s)	4,650	

Figure 5-36 MSA 2050 end-to-end performance results

The performance figures provided, as indicated in [Figure 5-36](#), are for reference as many variables exist between array configurations, workloads, hard drive types, disk group setup parameters, and host system setup. HPE has traditionally published a set of end-to-end MSA performance specifications that are fed into HPE Sizer tools, which are based on conservative real-world configurations. For consistency, the MSA

performance numbers have been documented in both benchmark and end-to-end performance tables.

HPE MSA 1050/2050/2052 hardware building blocks

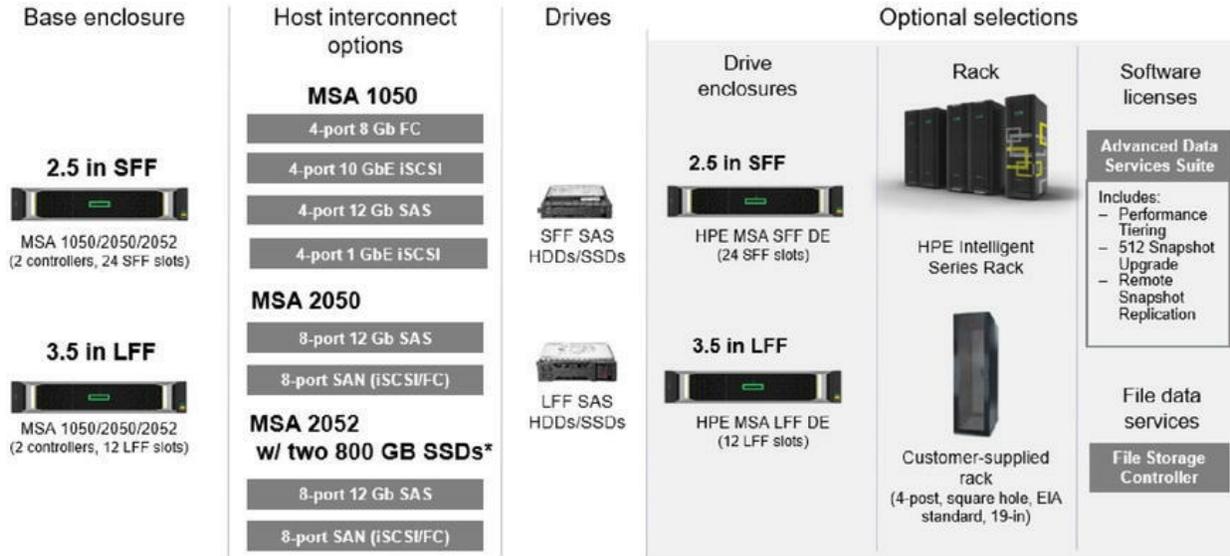


Figure 5-37 HPE MSA 1050/2050/2052 hardware building blocks

As shown in [Figure 5-37](#), the HPE MSA 1050, 2050, and 2052 consist of hardware building blocks, including optional selections.

HPE Nimble HF Storage



Figure 5-38 HPE Nimble HF array

Most of what we have talked about All-Flash Arrays is also in common to adaptive flash—or “hybrid” arrays. Let us highlight what is unique and where it is different.

When compared to other hybrid arrays, Nimble Adaptive Flash Arrays, as shown in [Figure 5-38](#):

- Have a five times smaller footprint.
- Are more efficient than other hybrids.
- Are unconstrained by disk.

Adaptive Flash Arrays can provide all-flash–like performance, but at a hybrid price point. The way it does this is through what is called Adaptable Service Levels.

What this means is that you can, on the fly, change the performance and responsiveness of the system for certain applications. For most applications, you really only need a small percentage of flash to accelerate performance, typically in the range of 5 to 10% of the capacity of the system. The normal operational mode of the array, which is called Auto Flash mode, would provide fast, all-flash–like performance for most applications. However, let us say you have one application running that requires 100% flash-level latency. Adaptive Service levels let you pin, or place, all of that application’s data into flash so that there is no chance of a cache miss that could impact latency to that app (this is what is called **All-Flash mode**). Again, this would

be for a small number of applications.

Conversely, there may be apps that need to optimize for cost per GB; in that case, you can even throttle down to use Minimal Flash mode. Just to stress again, all of this happens dynamically, on the fly, with absolutely no impact to applications or their users. This capability lets customers realize the value of all-flash-like performance for certain apps, but within a better price point for consolidated, mixed workloads.

The first fundamental architectural advantage we have is that Adaptive Flash Arrays are not limited by the performance of disk and the number of disk spindles. This is because when writing to disk, it uses mirrored NVRAM to rapidly respond to the host, and then compress the data in DRAM, and coalesces thousands of small random blocks into large sequential stripes on disk.

This allows us to sustain flash-like speed although we are actually writing to disk.

From a cost perspective, because we are not limited by the disks themselves, this suits the densest and cost-effective enterprise Nearline SAS drives (sometimes called capacity drives).

In addition, Nimble Adaptive Flash Arrays were designed from the ground up to be use cost-effective multi-level cell (MLC) and 3D-NAND flash in **the most efficient way possible**.

Rather than performing storage tiering, which is inefficient and does not rapidly adapt to changing workloads, it uses flash as a dynamic read cache. As a result, it achieves a 95% or better **cache hit rate** with only 5 to 10% flash capacity.

It also implemented advanced flash endurance management software, and write coalescing avoids write amplification on flash. This allows it to use cost-optimized 3D-NAND and still guarantee SSD lifespans.

Finally, all data is protected on disk using triple+ parity RAID. This means that Nimble arrays do not require RAID overhead on flash, making 100% of the flash capacity available for use. Compared to legacy hybrids, Nimble delivers about 50% more usable flash capacity for a given amount of raw flash.

Another way that Nimble offers radical simplicity and massive cost efficiency is through the ability to replicate between a primary array (typically an All-Flash Array) and a secondary array (either a Nimble Adaptive Flash or a new Secondary Flash Array). In fact, this is one of the most common ways that Nimble All-Flash Arrays are deployed. By deploying two Nimble arrays and taking advantage of simple, no-cost, replication capabilities, you can achieve backup and disaster recovery at about a one-third the cost of an all-flash backup or DR target.

Given the power of Nimble Secondary Flash Arrays, the same storage can be used for cost-effective dev/test and cloning—and allow data to be retained for months at a low cost.

Adaptive Flash Array family

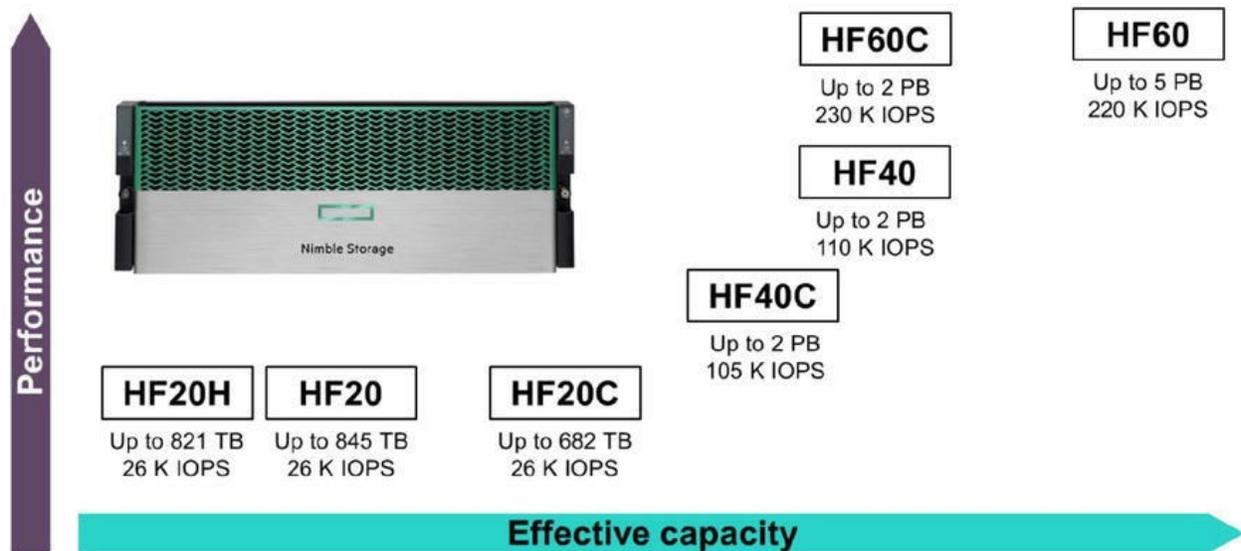


Figure 5-39 Adaptive Flash Array family

HPE Nimble Storage Adaptive Flash Platform, as shown in [Figure 5-39](#), is engineered to meet the needs of the small and mid-sized business (SMB) and smaller enterprise customers. Each consecutive member of the HF-Series family provides more performance along with a large effective capacity.

Variants with H are only half-populated, and variants with C are not using deduplication.

HF-Series write operation: Steps 1 to 4

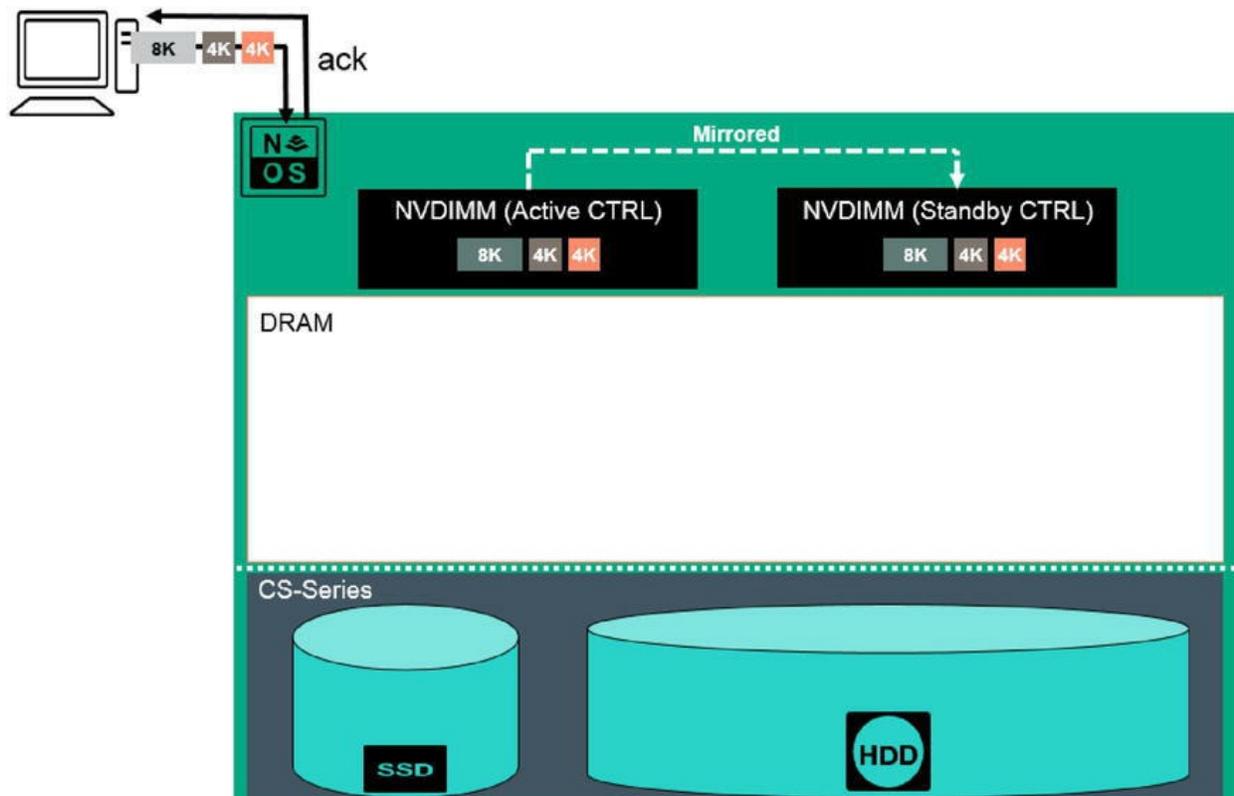


Figure 5-40 HF-Series write operation beginning

NimbleOS is a unique architecture that has been designed to accommodate a variety of platforms. Currently Nimble Storage offers two: the Hybrid Flash Array Family and the All-Flash Array family. As shown in [Figure 5-40](#), you see a representation of NimbleOS as well as some of the hardware components found in Nimble controllers including NVDIMM and DRAM. At the bottom of the graphic, both array families' primary storage components are represented.

To help you better understand NimbleOS, we are going to first review the write process. When a write is received by the controllers, it is first accepted by the NVDIMM. From there, it is mirrored to the standby controllers, NVDIMM, before it is acknowledged. At this point, the array has two copies of the data.

Process steps:

1. Writes are sent by a variety of applications in variable block sizes.
2. NimbleOS places incoming writes into the active controllers NVDIMM.
3. NimbleOS mirrors the active controllers NVDIMM to the standby controllers NVDIMM.
4. NimbleOS acknowledges the write.

HF-Series write operation: Step 5

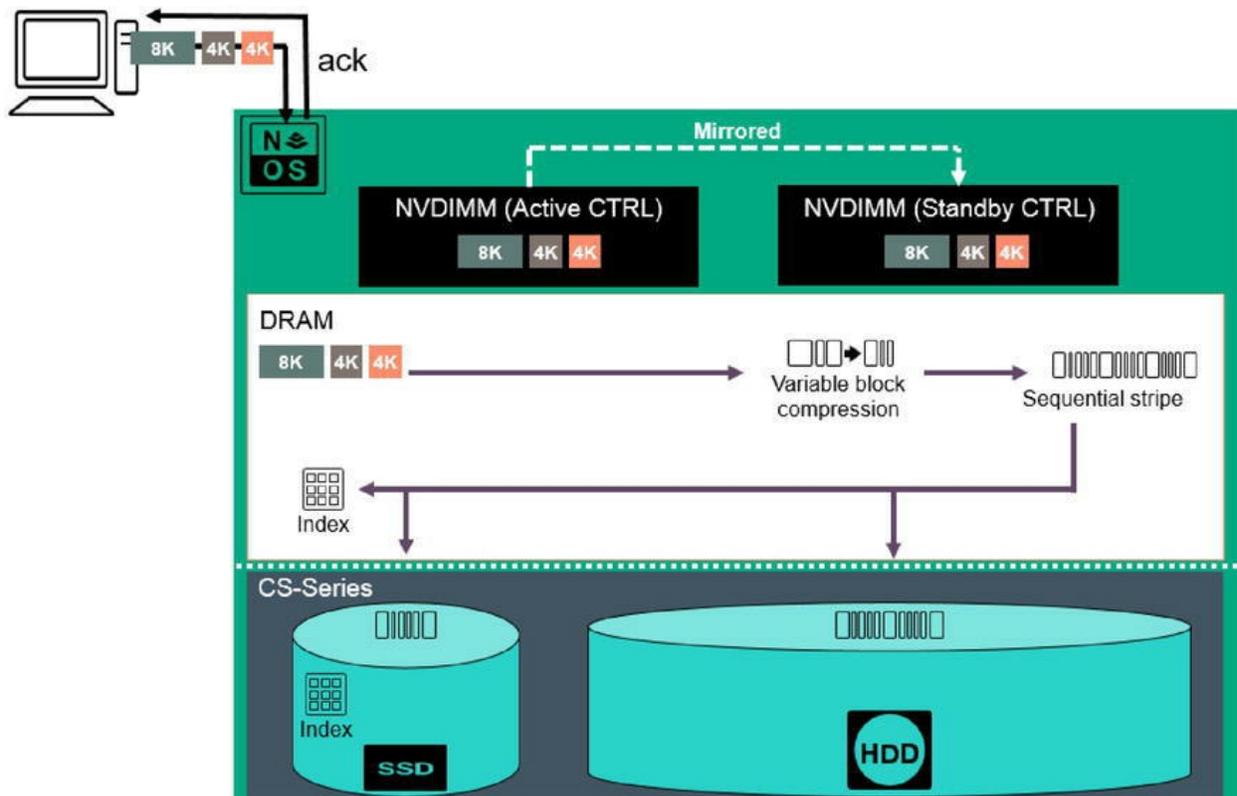


Figure 5-41 HF-Series write operation details

The same basic concepts are applied to the HF-Series as you have seen in the AF-Series, as shown in [Figure 5-41](#). When an HF-Series array receives blocks in the DRAM/NVDIMM, NimbleOS performs variable block compression and then builds a sequential stripe write; the stripe size is dependent on the Array model:

- **AF** = 10 MB
- **HFx0** = 18 MB
- **Half pop** = 8 MB

The compression engine limits the total number of IOPS sent to disk, while the sequential stripe write ensures both fast write and read performance. The stripe write is then sent to disk and an index of the blocks locations is maintained in the SSD flash cache. This ensures a fast lookup should a read operation need to pull data from the disks. At the same time data is written to hard disk drives, all cache-worthy data is written to flash cache to accelerate read operations.

Steps(continued from the previous section):

5. Blocks are copied into DRAM.
 - Variable block deduplication is applied (*dedupe is not supported on the HF20C).
 - Variable block compression is applied.
 - Blocks are formed into a sequential stripe write.
 - Sequential stripe is written to hard disk drives.
 - Cache-worthy data and any data destined for pinned volumes is also sent to SSD.
 - Blocks are indexed.

The NimbleOS architecture allows Nimble to use less expensive MLC SSDs while still ensuring high performance and a long device lifespan by:

- Converting random writes into sequential writes, which minimizes write amplification, resulting in less wear.
- By not treating an SSD like a hard drive and using RAID, which wears out SSDs faster.
- By compressing the data on the flash saving space and the number of writes made to the SSDs.
- By keeping the metadata in cache, which accelerates all the reads. Since the array does not need to access the disks for metadata, latency is

reduced.

NimbleOS considers all random writes as cache worthy. Other cache-worthy data includes data pinned in a volume or data that is considered “hot.” The caching of read data allows NimbleOS to use SSDs to provide extremely fast reads while using the HDDs for resilient storage of the data.

If you look at caching from a hardware prospective, the benefits of the Nimble adaptive flash architecture become evident:

- A 7200 revolutions per minute (RPM) SATA disk can support about 70 IOPS.
- A 10,000 RPM FC disk can support about 130 IOPS.
- Disks spinning at 15,000 RPMs can produce about 170 IOPS.

Systems without a large cache must utilize RAID to concatenate groups of HDDs together to garner more IOPS for the system. Theoretically, building a RAID 5 from three disks would produce between 140 IOPS from SATA disks and 340 IOPS from 15,000 RPM disks.

On the other hand, a single SSD can produce about 30,000 IOPS. This is a dramatic increase and NimbleOS takes full advantage of that performance boost while keeping costs in check and data secure on HDDs.

HF-Series read operations

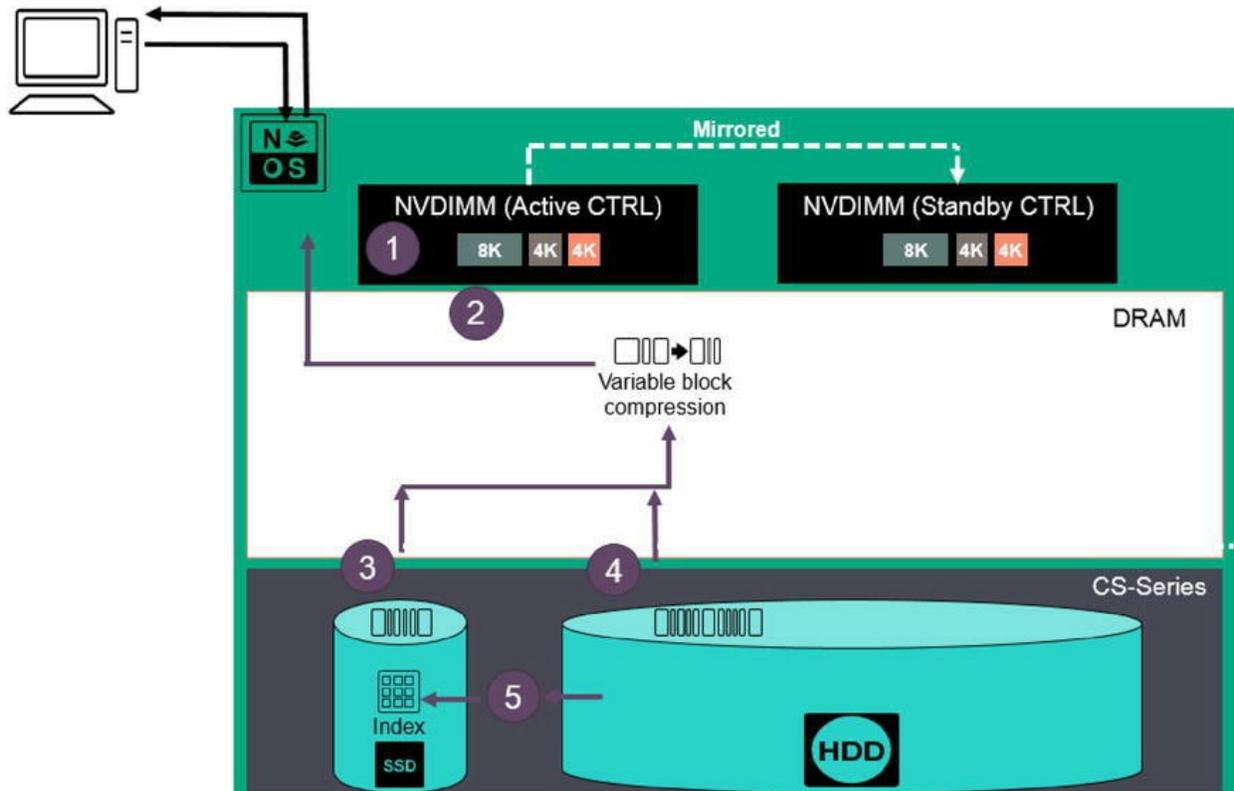


Figure 5-42 HF-Series read operation sequence

Now let us look at how reads are processed in the Nimble Array. When a read request is sent to the array, as illustrated in [Figure 5-42](#), the NimbleOS architecture will retrieve the data from the fastest area possible.

NimbleOS starts by checking in NVDIMM. If a write was just sent to the NVDIMM and is still there, it will be retrieved very quickly. If it is not in NVDIMM, NimbleOS will check DRAM, because that is the next fastest location the read can be retrieved from.

If the data is not located in the DRAM, it will check the flash cache. If found, it will validate the checksum, decompress, and return the data. If the read is still not located, NimbleOS will reference the index to locate the read on the disk. Lastly, if NimbleOS determines that the data found on the array is cache worthy, it will send a copy to the flash cache.

It is important to note that a checksum happens on both reads and writes to help ensure data integrity.

Process steps:

1. Read from NVDIMM.
2. If not found, check DRAM.
3. If not found, read from SSD. If found, validate checksum, decompress, and return data.
4. If not found, read from disk using the index on the SSD for a quick lookup. If found, validate checksum, decompress, and return data.
5. If cache-worthy, write to cache.



Note

The HF-Series Arrays have the same changed block benefit demonstrated in the AF-Series Arrays.

HF-Series triple+ parity

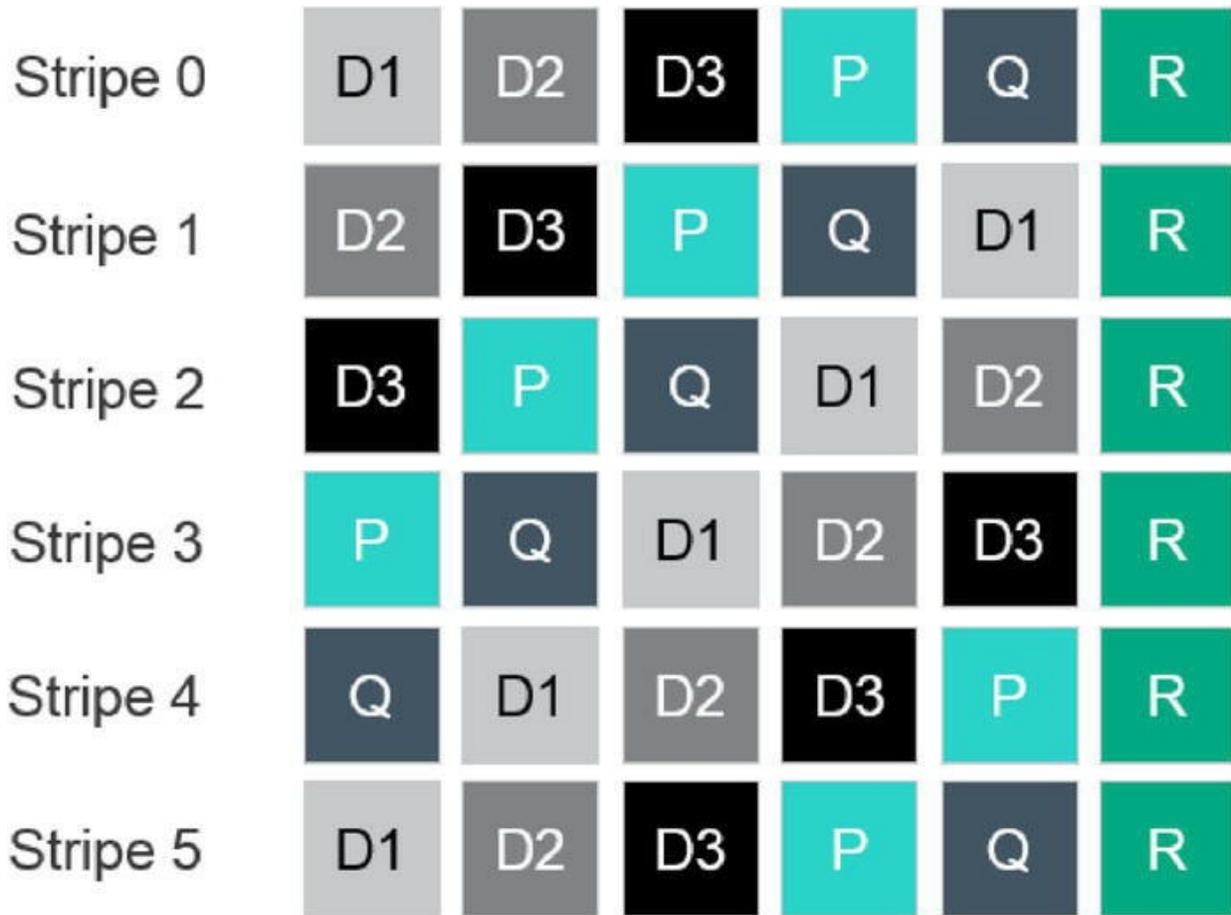


Figure 5-43 HF-Series triple+ parity RAID layout

The Adaptive Flash Platform utilizes triple+ parity RAID to concatenate all HDDs in each shelf. Nimble uses left synchronous rotation in that implementation. When writing parity, the first two parity blocks (P and Q) are rotational as illustrated in [Figure 5-43](#).

The third parity block is not rotational and is written to the disk in slot one, though it should be noted that if that disk fails, the R block may be written to another disk. To see which disks are holding the R block, obtain and run the `iostat.sh` script during a heavy read workload. The disks containing the R block should stick out with significantly less read activity.

Nimble's triple+ parity design supports the loss of three HDDs in a single shelf before it reaches zero disk redundancy. This means that three disk drives within a single shelf can fail without the system losing data.

**Note**

Failed drives are only counted if they are actively marked as failed by the controllers.

Should a single shelf reach a point where there is zero redundancy within the RAID group, the system will stop data services.

**Note**

The system will shut down if there are three disk failures prior to any one of those failed disks being rebuilt. This is a configurable option.

HPE Cloud Volumes



Figure 5-44 HPE Cloud Volumes

HPE Cloud Volumes is an enterprise-grade multi-cloud storage service. This public cloud service is run by HPE, built on HPE Nimble Storage technology, and delivered 100% as a pay-as-you-go service. HPE Cloud Volumes, as indicated in [Figure 5-44](#), provides the storage for customers who run applications in Amazon Web Services (AWS) and Microsoft Azure. Having Cloud Volumes provide the storage, while AWS or Azure provides the compute, enables customers to regain control of their data and unlock it from a single cloud provider. Customers can move data freely between Cloud Volumes and cloud-ready Nimble arrays in their data center.

The Cloud Volumes data centers are located in close proximity to AWS and Azure data centers to ensure high performance and low-latency access for your customers.

Cloud Volumes delivers proven uptime and availability and is millions of times more durable than native cloud block storage. Cloud Volumes can protect customers' data with 256-bit AES encryption, and it is certified for SoC 2 Type 1 and HIPAA compliance.



Note

Cloud Volumes may not be available in all countries, so always verify availability.

Other HPE Cloud Volumes features include:

- Amazon Web Services and Microsoft Azure Instant snapshots and clones
- 256-bit volume encryption
- Cloud visibility and reporting
- Replication between data center and cloud
- REST API
- Command-line interface (CLI)
- Multi-host access and clustering
- Real-time and predictive cost estimates
- Containers support (Docker, Kubernetes, Mesosphere)

HPE Nimble Storage ProStack objectives



Figure 5-45 HPE Nimble Storage ProStack

HPE Nimble Storage ProStack, as shown in [Figure 5-45](#), is a powerful storage array, uniting best-in-class AI for the data center with the world's bestselling server.

HPE Nimble Storage ProStack features:

- **Easy**—Intuitive setup wizard to deploy HPE Nimble Storage ProStack in less than an hour
- **Simple**—Innovative vCenter management plugin to rapidly manage and scale HPE Nimble Storage ProStack
- **AI Driven**—Improve performance, optimize resources, and simplify infrastructure planning using HPE Infosight for HPE Nimble Storage ProStack

Flexible platform and deployment options

Hardware	Virtualization	Deployment
<ul style="list-style-type: none">– HPE Nimble Storage iSCSI arrays: AF and HF– HPE ProLiant Gen9/Gen10: DL360 and DL380– HPE FlexFabric, Aruba, and Cisco	VMware vSphere: Vvol and VMFS	Greenfield: All new equipment Brownfield: Existing HPE ProLiant

Figure 5-46 Flexible platform and deployment options

Support for HPE hardware, as indicated in [Figure 5-46](#), includes:

- HPE Nimble Storage iSCSI arrays: AF and HF series
- HPE ProLiant Gen9/Gen10: DL360 and DL380
- HPE FlexFabric, Aruba, and Cisco

HPE Nimble Storage ProStack platform overview

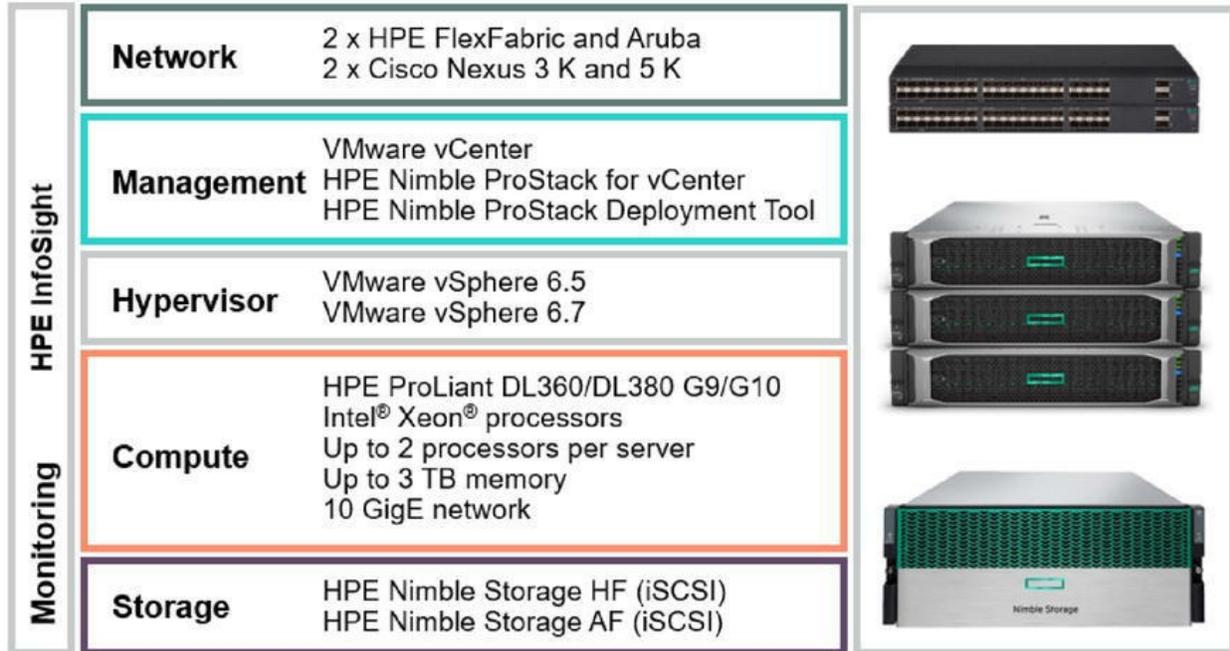


Figure 5-47 HPE Nimble Storage ProStack platform overview

HPE Nimble Storage ProStack, as shown in [Figure 5-47](#), is a flexible platform built on top of HPE hardware. Support for virtualization features and deployment options spans both greenfield installation with new equipment and brownfield with existing HPE ProLiant devices.

Activity: Configuring a storage system

1. Download the HPE Nimble Storage All-Flash Arrays QuickSpecs from: <https://h20195.www2.hpe.com/v2/getdocument.aspx?docname=a00008273enw>
2. Answer the following questions:
 - a. What are the necessary steps to configure the Nimble solution (search for Configuration Information section)?

b. What is the minimum amount of ports recommended for best performance?

HPE 3PAR StoreServ

HPE 3PAR StoreServ provides a product family to meet primary storage needs—for medium-sized enterprises and larger SMBs experiencing rapid growth, a large enterprise looking to support ITaaS, or a global service provider building a hybrid or private cloud.

HPE 3PAR StoreServ portfolio



Figure 5-48 HPE 3PAR StoreServ portfolio

HPE 3PAR StoreServ supports true convergence of block, file, and object access while offering the performance and flexibility that you need to accelerate new application deployments and support server virtualization, the cloud, ITaaS, and your future technology initiatives. It is a storage platform that allows your customers to spend less time on management, gives them technically advanced features for less money, and eliminates trade-offs that require them to sacrifice critical capabilities such as performance and scalability. With HPE 3PAR StoreServ, your customers can serve unpredictable and mixed workloads, support unstructured and structured data growth, and meet block, file, and object access storage needs from a single capacity store.

HPE 3PAR StoreServ models, as shown in [Figure 5-48](#), include:

- HPE 3PAR StoreServ 8000
- HPE 3PAR StoreServ 9000
- HPE 3PAR StoreServ 20000
- HPE 3PAR StoreServ File Controller
- HPE 3PAR File Persona

The all-flash HPE 3PAR StoreServ models—20850, 20450, 8450, and 9450—offer massive performance scalability from hundreds of thousands of IOPS to millions of IOPS all within a submillisecond response time. These all-flash systems uniquely offer performance without any compromises, be it the enterprise class Tier 1 data services (such as three data center replication) or resiliency using their scalable architecture (greater than two nodes). These systems are well suited for massive workload consolidation and should be the first choice for customers building their next-generation (all-flash) data centers.

Furthermore, HPE 3PAR StoreServ 8200 All-Flash Starter Kit is designed to offer an incredibly affordable entry point for flash performance without limiting scalability, functionality, or data mobility. The converged flash HPE 3PAR StoreServ models—20800, 8440, 8400, 8200—leverage a flash-first approach that offers flash performance with the additional flexibility to add spinning media (HDD) as a tier. In an all-flash configuration, these systems also deliver hundreds of thousands to millions of IOPS all within a submillisecond response time. The 20800 and 8440 are hyperscalable models,

supporting multi-PiB raw capacity scalability. The 8400 and 8200 models balance performance and scalability and offer a very economic starting point for customers looking to take advantage of the HPE 3PAR architecture.

It is important to remember that all HPE 3PAR StoreServ models are built on a single architecture, run the exact same HPE 3PAR operating system, and offer a common set of enterprise data services, be it the 8200, the 20800, or previous generations of HPE 3PAR StoreServ. HPE 3PAR StoreServ models can natively replicate and federate among each other without the need for any external replication or virtualization appliance.



Note

There is no upgrade path between any HPE 3PAR StoreServ arrays.

Extended SSD warranty

HPE 3PAR StoreServ 8000, 9000, and 20000 Storage systems SSD are eligible for the Extended SSD Warranty Program.

Eligible products must meet the following criteria:

- Maintained under active and continuous support coverage from HPE from time of purchase. Continuous support coverage means that the customer purchased upfront support from HPE at the time of eligible product purchase to supplement any underlying warranty coverage and has continuously renewed such support coverage with HPE without any lapses in coverage.
- Have not reached end-of-support life as determined by HPE. The warranty start date of the eligible products has not exceeded seven years.
- The SSD endurance, as solely determined by HPE, is below 5% of life left. HPE may require the customer to provide information to HPE for the purposes of determining life-left reading. If the customer does not provide the required information in a timely manner, HPE is under no obligation to provide a replacement drive pursuant to these terms.

Customers will receive a one-time only replacement SSD at no additional charge under these program terms. The request for support coverage needs to

meet the eligibility requirements for such a claim as set forth in the HPE’s standard sales terms and the Extended SSD Replacement Coverage criteria. SSDs that do not meet these program terms and have reached their maximum usage limit are not eligible for repair or replacement under HPE support coverage.

This program is subject to change or termination at any time, without notice, by HPE.

HPE 3PAR StoreServ manageability



Figure 5-49 Full set of management options

Management tools for HPE 3PAR StoreServ, as illustrated in [Figure 5-49](#), include:

- Graphical user interface (GUI) management options
- CLI
- Application programming interface (API)
- Analytics
- Management integrations with third-party products

HPE 3PAR StoreServ 8000, 9000, and 20000 software details

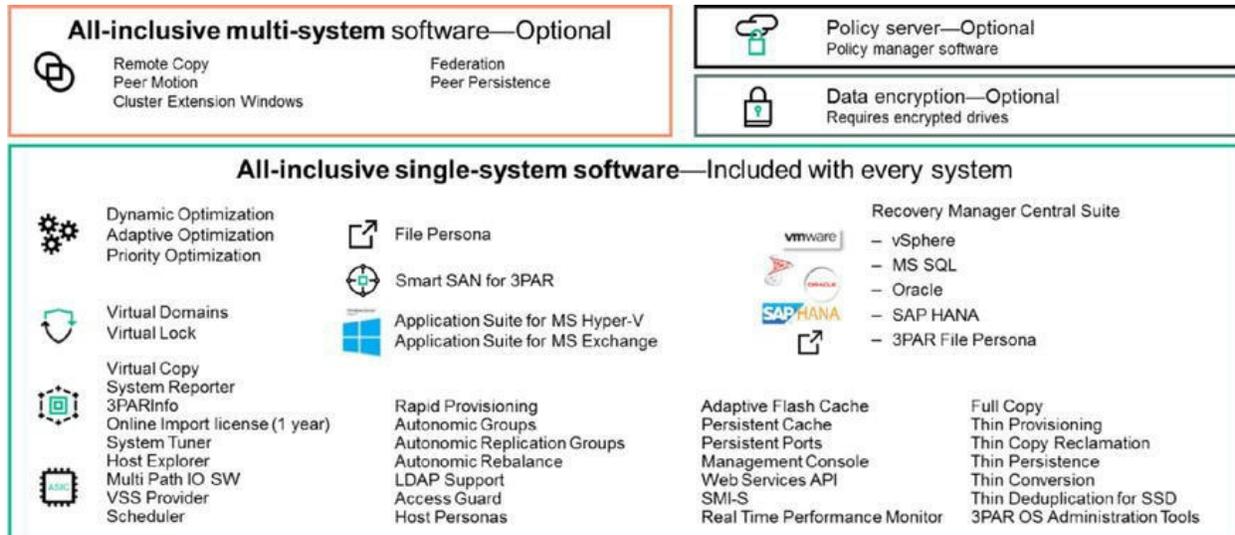


Figure 5-50 All-inclusive licensing model

The 3PAR StoreServ family includes as part of the array and drives the All-Inclusive Single-System Software, as illustrated in [Figure 5-50](#), which includes:

- OS suite
- Virtual copy
- Adaptive optimization
- Dynamic optimization
- Priority optimization
- Virtual domains
- Virtual lock
- Online import
- File persona
- Recovery Manager Central (RMC)
- Smart SAN

The all-inclusive multi-system software is an optional software suite that includes Peer Motion, Remote Copy, Peer Persistence, and Cluster Extension:

- **HPE 3PAR Peer Motion**—HPE 3PAR Peer Motion enables load balancing at will wherein movement of data and workloads between arrays is initiated without impacting applications, users, or services.
- **HPE 3PAR Remote Copy**—HPE 3PAR Remote Copy offers simple and cost-effective data protection for efficient multi-tenant disaster recovery.
- **HPE 3PAR Peer Persistence**—HPE 3PAR Peer Persistence software enables HPE 3PAR StoreServ systems located at metropolitan distances to act as peers to each other, presenting a nearly continuous storage system to VMware vSphere hosts connected to them.
- **HPE 3PAR Cluster Extension**—HPE 3PAR Cluster Extension (CLX) Software offers rapid automated protection against application downtime from fault, failure, or site disaster. CLX integrates with the Windows OS clustering software and HPE 3PAR Remote Copy to automate failover and failback between sites.

HPE 3PAR StoreServ 8000 hardware building blocks

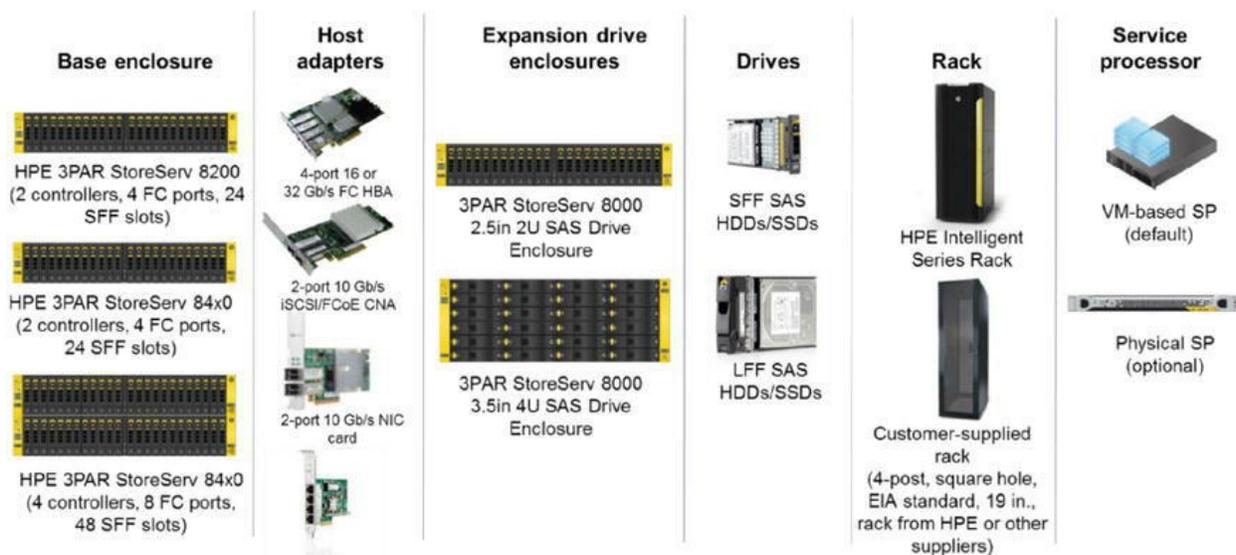


Figure 5-51 The basic building blocks of an HPE 3PAR StoreServ array

HPE 3PAR StoreServ 8000 solution, as shown in [Figure 5-51](#), typically consists of:

- Base enclosure

- Host adapters
- Expansion drive enclosures
- Drives
- Racks
- Service processor

Each HPE 3PAR StoreServ 8000 requires its own service processor. The service processor functions as the communication interface between a customer's IP network and HPE 3PAR Central by managing all service-related communications. The service processor leverages the industry-standard HTTP over Secure Sockets Layer (HTTPS) protocol to secure and encrypt data communication. The service processor can be deployed either as a virtual service processor (VSP) or a physical service processor.

A VSP is included free with the base HPE 3PAR Operating System. The VSP can be installed on a customer-provided VMware or Microsoft Hyper-V system.

The physical service processor is a dedicated storage appliance located within the storage rack providing proximity to the HPE 3PAR StoreServ 8000 Storage. The physical service processor is fully supported and maintained by HPE Services. The physical service processor has serial port connectivity that provides maintenance access for troubleshooting capabilities.

If a VMware server is not available to run the VSP, the physical service processor is the alternative for remote monitoring and remote service. The physical service processor is available in two versions: with Single Power Supply and with Redundant Power Supply.

HPE 3PAR StoreServ 8000 host adapters

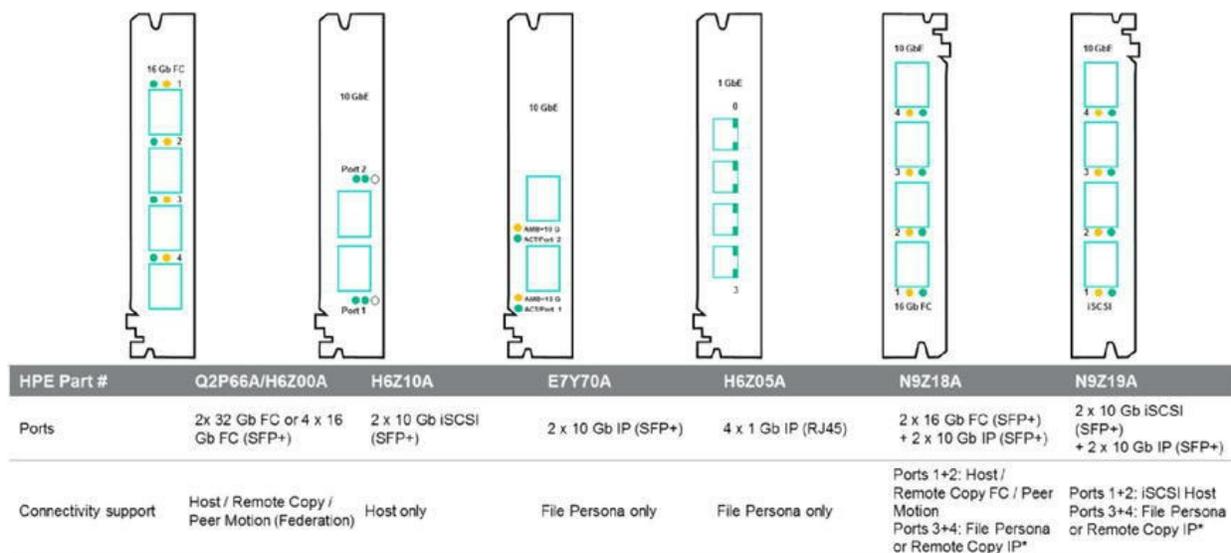


Figure 5-52 A variety of adapters to meet connectivity requirements

Host adapters, as illustrated in [Figure 5-52](#), can be ordered separately to be installed in the field, or they can be factory configured into controller nodes. Host adapter cards provide the array with additional FC ports, with 10 Gb/s iSCSI Fibre Channel over Ethernet (FCoE) ports, or with 1 GbE/s and 10 Gb/s Ethernet ports. The additional FC ports can be used for connection to hosts or used to connect to other HPE 3PAR StoreServ Storage systems in a remote copy relationship. The iSCSI/FCoE ports permit host connection in iSCSI and FCoE environments. The Ethernet ports can be used only with the HPE 3PAR File Persona Software for File services connectivity.

HPE 3PAR StoreServ 8000 HBAs include:

- HPE 3PAR StoreServ 8000 2-port 32 Gb Fibre Channel Adapter
- HPE 3PAR StoreServ 8000 4-port 16 Gb Fibre Channel Adapter
- HPE 3PAR StoreServ 8000 2-port 10 Gb iSCSI/FCoE Adapter
- HPE 3PAR StoreServ 8000 4-port 1 Gb Ethernet Adapter
- HPE 3PAR StoreServ 8000 2-port 10 Gb Ethernet Adapter
- HPE 3PAR StoreServ 8000 4-port 16 Gb Fibre Channel/10Gb NIC Combo Adapter
- HPE 3PAR StoreServ 8000 4-port 10 Gb iSCSI/10Gb NIC Combo

Adapter



Note

32 Gb FC host support requires 3PAR OS 3.3.1 Technology Release T05 or later. Check HPE Spock for latest supported 3PAR software and connectivity. The host adapter cards are optional because the Storage Base products and the Upgrade Controller Node Pair include built-in FC ports.

HPE 3PAR StoreServ 8400 node pair

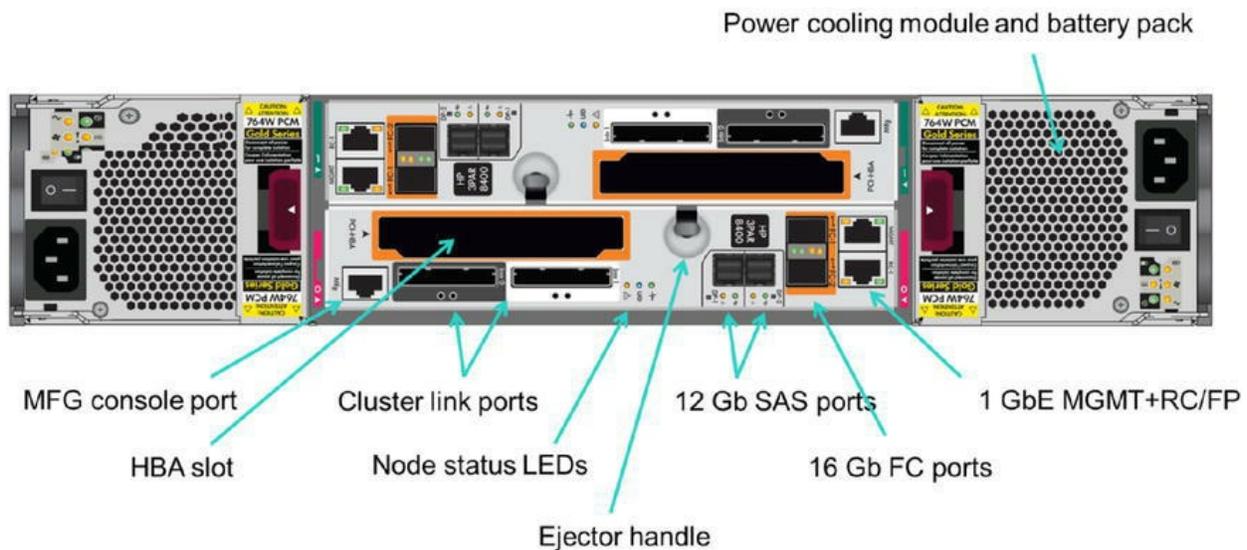


Figure 5-53 HPE 3PAR StoreServ 8400 node pair

The port numbering convention is Node:Slot:Port (N:S:P) and is used in the creation of the 3PAR World Wide Name (WWN).

The example assumes Node 0, so the port number referenced is the onboard FC port slot (slot 1) and the port referenced is port 2 (0:1:2).

The MGMT port is used to connect to the network for storage array management.

Each 8000 series controller, as shown in [Figure 5-53](#), contains:

- One built-in 1 GbE Remote Copy
 - One built-in 1 GbE Management Port
 - Two built-in 16 Gb FC ports
 - Two four-lane 12 Gb/s SAS drive chassis connections
 - Two four-node cluster interconnects
 - Optional PCI-e card slot
-



Note

Slot 3 contains the RCIP (Remote Copy over IP) data replication port (used for Remote Copy), but it can also be configured and used with the File Persona feature.



Note

Visit the “Customer Self Repair Services Media Library for more information at: <https://sml-csr.ext.hpe.com/>.

HPE 3PAR StoreServ 8000 controller nodes

Per node configuration

1 Intel® Ivy Bridge Processor	
- 8200	6-core 2.2 GHz
- 8400	6-core 2.2 GHz
- 8440	10-core 2.4 GHz
- 8450	10-core 2.4 GHz
Data cache	
- 8200	16 GB
- 8400	16 GB
- 8440	32 GB
- 8450	32 GB
Control cache	
- 8200	16 GB
- 8400	16 GB
- 8440	64 GB
- 8450	64 GB

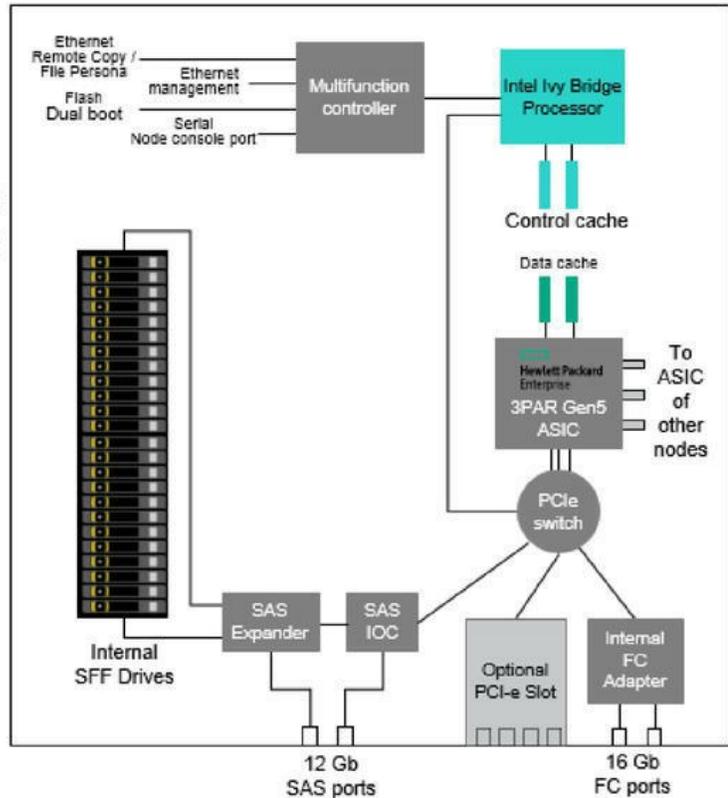


Figure 5-54 The core components of the HPE 3PAR StoreServ controllers

The HPE 3PAR StoreServ 8000 Storage requires one storage base (controller node enclosure), as illustrated in [Figure 5-54](#).

- **HPE 3PAR StoreServ 8000 Storage 2-node Storage Base**—2U controller node enclosure with 24 SFF 6.4 cm (2.5 in) drive bays at the front and two controller nodes at the rear.
- **HPE 3PAR StoreServ 8000 Storage 4-node Storage Base**—Two 2U controller node enclosures totaling 48 SFF 6.4 cm (2.5 in) drive bays at the front and four controller nodes at the rear.

HPE 3PAR StoreServ 8000 12 Gb SAS Drive Enclosures

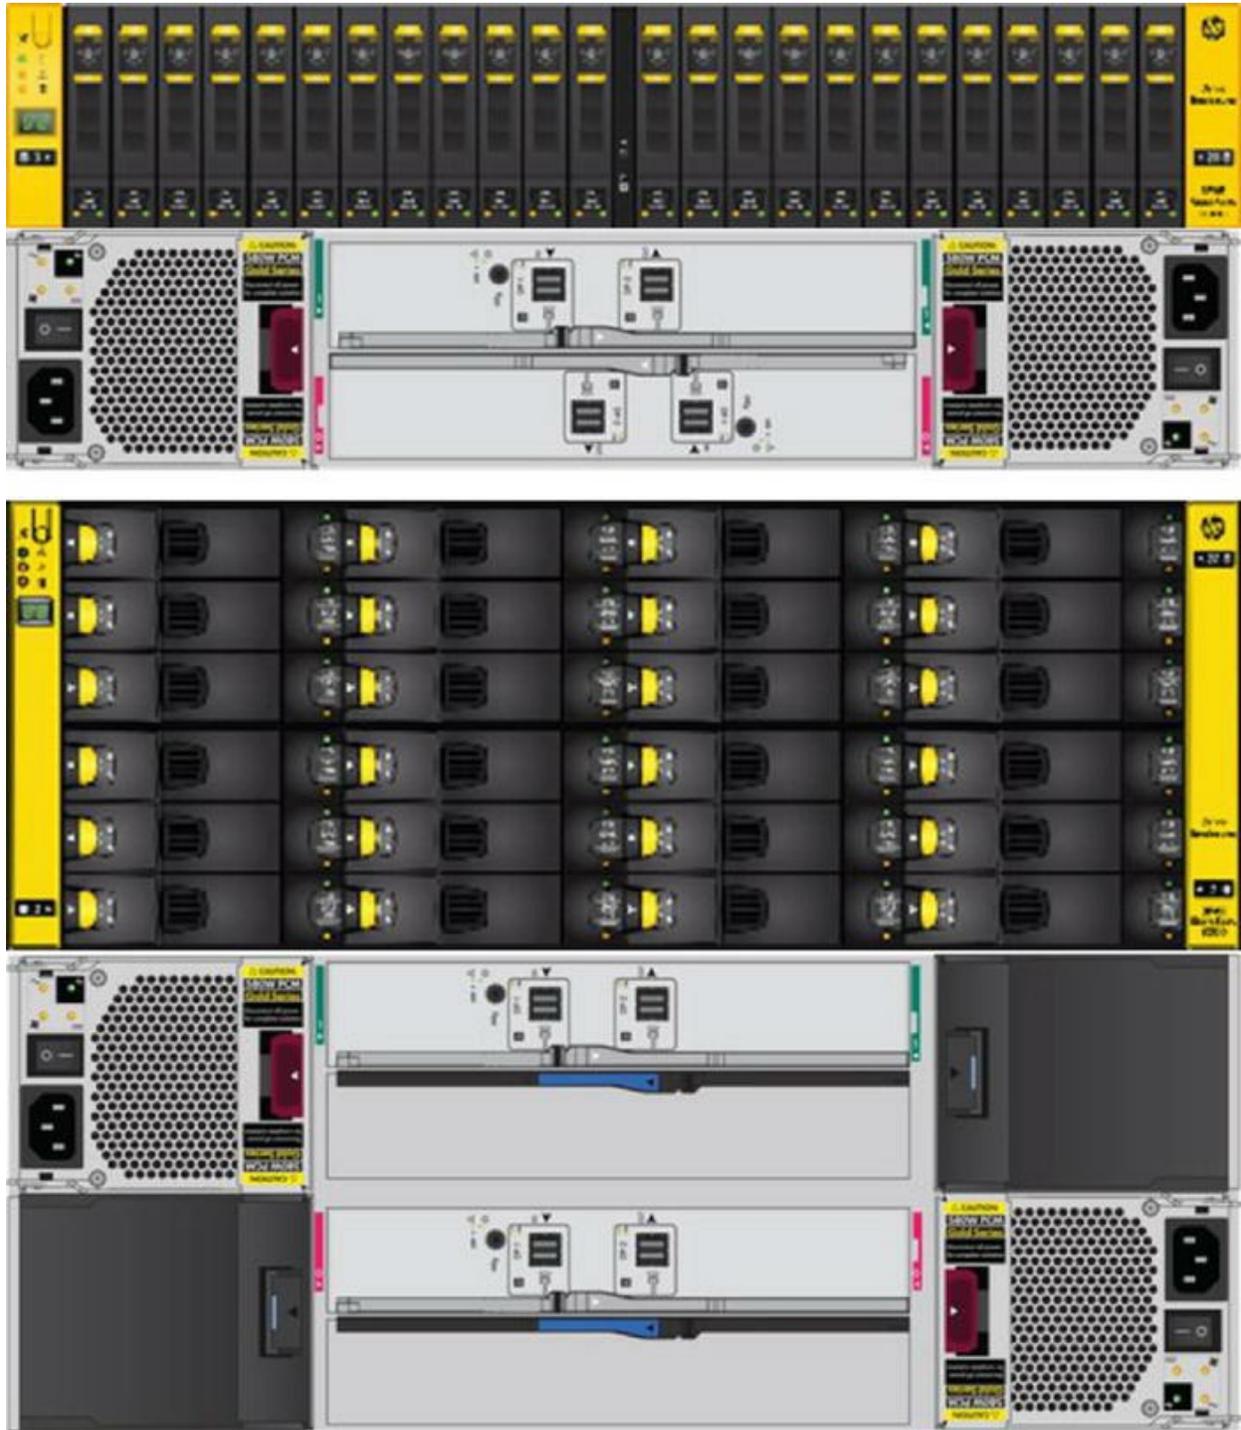


Figure 5-55 HPE 3PAR StoreServ SFF and LFF

The 8000 Series supports two types of expansion cages, as shown in [Figure 5-55](#): a 2U 24-drive 2.5” SFF cage, and a 4U 24-drive 3.5” LFF, both with a 12 Gb SAS interface.

The 8000 series expansion cages support both 6 Gb and 12 Gb SAS drives.



Note

The SFF cages (including the base enclosure) only support a 2 TB 7.2k NL disks. The LFF expansion cage supports 2, 4, and 6 TB NL disks.



Note

The LFF expansion cage does not support any 10 k or 15 k FC (fast class) disks and only supports a cMLC 480 GB SSD. No data encryption can be done on the cMLC SSD disks (including the 8450 model) in the LFF expansion cage.

The 8000 series array uses a C:M:D addressing convention where C = cage, M = magazine, and D = disk.

This is a legacy issue where the T and 10000 series used a cage with magazines. None of the current 3PARs use magazines, but the C:M:D addressing convention is still used to identify a disk. The D = Disk value will always be zero (0). For example, if a disk had failed in disk bay 12 in cage 2, the number of the drive would be 2.11.0. The reason it is 11 and not 12 is because HPE 3PAR drives are numbered from 0 to 23 in a 24-drive cage.

HPE 3PAR StoreServ 8000 hardware architecture

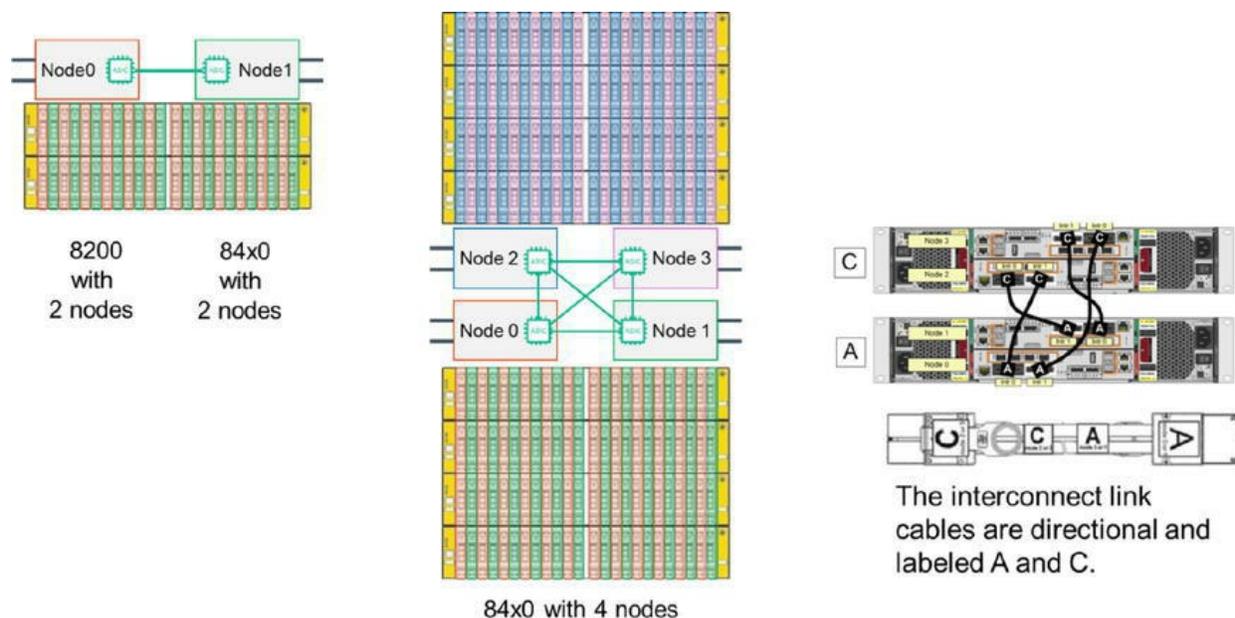


Figure 5-56 HPE 3PAR StoreServ Full-mesh backplane

Backplane interconnects within servers have evolved dramatically over the years. Most, if not all, server and storage array architectures have traditionally employed simple bus-based backplanes for high-speed processor, memory, and IO communication. Parallel to the growth of SMP-based servers, significant investments were also made to switch architectures, which have been applied to one or two enterprise storage arrays.

The move from buses to switches was intended to address latency issues across the growing number of devices on the backplane (more processors, larger memory, and IO systems). Third-generation full-mesh interconnects first appeared in the late 1990s in enterprise servers.

The HPE 3PAR StoreServ full-mesh backplane, as shown in [Figure 5-56](#), is a passive circuit board that contains slots for up to four or eight controller nodes, depending on the model. As noted earlier, each controller node slot is connected to every other controller node slot by a high-speed link (4 GB/s in each direction, or 8 GB/s total), forming a full-mesh interconnect between all controller nodes in the cluster—something that HPE refers to as a “Mesh-Active” design. These interconnects deliver low-latency, high-bandwidth communication and data movement between controller nodes through dedicated point-to-point links and a low overhead protocol that features rapid internode messaging and acknowledgment. It is important to note that, while

the value of these interconnects is high, the cost of providing them is relatively low. In addition, a completely separate full-mesh network of serial links provides a redundant low-speed channel of communication for exchanging control information between the nodes.

The HPE 3PAR StoreServ 20000 features an eight-node-capable backplane that supports two to eight controller nodes. HPE 3PAR StoreServ 8000 systems feature either a dual-node or quad-node-capable system that is essentially an equivalent of what was used in former enterprise-class arrays that offer the same high-speed links between nodes.

Active/Active versus Mesh-Active

Most traditional array architectures fall into one of two categories: monolithic or modular. In a monolithic architecture, being able to start with smaller, more affordable configurations (scaling down) presents challenges. Active processing elements not only have to be implemented redundantly, but they are also segmented and dedicated to distinct functions such as host management, caching, and RAID or drive management. For example, the smallest monolithic system might have a minimum of six processing elements (one for each of three functions, which are then doubled for redundancy of each function). In this design—with its emphasis on optimized internal interconnectivity—users gain the Active/Active processing advantages of a central global cache (for example, LUNs can be coherently exported from multiple ports). However, these architectures typically involve higher costs relative to modular architectures.

In traditional modular architectures, users can start with smaller and more cost-efficient configurations. The number of processing elements is reduced to just two, because each element is multifunction in design—handling host, cache, and drive management processes. The trade-off for this cost effectiveness is the cost or complexity of scalability. Because only two nodes are supported in most designs, scale can only be realized by replacing nodes with more powerful node versions or by purchasing and managing more arrays. Another trade-off is that dual-node modular architectures, while providing failover capabilities, typically do not offer truly Active/Active implementations where individual LUNs can be simultaneously and coherently processed by both controllers.

The HPE 3PAR StoreServ architecture was designed to provide cost-effective, single-system scalability through a cache-coherent, multinode clustered implementation. This architecture begins with a multifunction node design, and like a modular array, requires just two initial controller nodes for redundancy. However, unlike traditional modular arrays, enhanced direct interconnects are provided between the controllers to facilitate Mesh-Active processing. Unlike legacy Active/Active controller architectures—where each LUN (or volume) is active on only a single controller—this Mesh-Active design allows each LUN to be active on every controller in the system, thus forming a mesh. This design delivers robust, load-balanced performance and greater headroom for cost-effective scalability, overcoming the trade-offs typically associated with modular and monolithic storage arrays.

System-wide striping

Through a Mesh-Active design and system-wide striping, the HPE 3PAR StoreServ architecture can provide the best of traditional modular and monolithic architectures in addition to massive load balancing. The HPE 3PAR StoreServ Mesh-Active design not only allows all volumes to be active on all controllers, but also promotes system-wide striping that autonomically provisions and seamlessly stripes volumes across all system resources to deliver high, predictable levels of performance. System-wide striping of data provides high and predictable levels of service for all workload types through the massively parallel and fine-grained striping of data across all internal resources (such as disks, ports, loops, cache, and processors). As a result, as the use of the system grows—or in the event of a component failure—service conditions remain high and predictable. Unlike application-centric approaches to storage, HPE 3PAR StoreServ provides autonomic rebalancing that enables the system to evenly balance and use all available physical resources.

This is particularly important with hardware upgrades because existing data should be rebalanced and stripped across new available resources. On HPE 3PAR StoreServ, this is done without service disruption or preplanning. For flash-based media, fine-grained virtualization combined with system-wide striping drives uniform IO patterns by spreading wear evenly across the entire system. Should there be a media failure, system-wide sparing also helps

guard against performance degradation by enabling a many-to-many rebuild, resulting in faster rebuilds. Because HPE 3PAR StoreServ autonomically manages this system-wide load balancing, no extra time or complexity is required to create or maintain a more efficiently configured system.

Controller node architecture

An important element of the HPE 3PAR StoreServ architecture is the controller node, and it is a powerful data movement engine that is designed for mixed workloads. As noted earlier, a single system, depending on the model, is modularly configured as a cluster of two to eight controller nodes. This modular approach provides flexibility, a cost-effective entry footprint, and affordable upgrade paths for increasing performance, capacity, connectivity, and availability as needs change. In addition, the minimum dual-controller configuration means that the system can withstand an entire controller node failure without impacting data availability. Controller nodes can be added in pairs to the cluster non-disruptively, and each node is completely hot-pluggable to enable online serviceability.

Unlike legacy architectures that process IO commands and move data using the same processor complex, the HPE 3PAR StoreServ controller node architecture separates the processing of control commands from data movement, which helps ensure that CPU bandwidth is available for control processing and is not used for bulk data transfer. This innovation eliminates the performance bottlenecks of existing platforms that use a single processing element to serve competing workloads, for example online transaction processing (OLTP) and data warehousing workloads.

HPE 3PAR StoreServ 8000 disk drives

HDD SFF	HDD LFF	SSD
HPE 3PAR 8000 300GB SAS 15 K SFF (2.5in) HDD with All-inclusive Single-system Software	HPE 3PAR 8000 4 TB+SW 7.2 K LFF HDD	SFF
HPE 3PAR 8000 600 GB+SW 15 K SFF HDD	HPE 3PAR 8000 6 TB+SW 7.2K LFF HDD	HPE 3PAR 8000 920 GB SAS SFF (2.5 in) SSD with All-inclusive Single-system Software
HPE 3PAR 8000 600 GB+SW 10 K SFF HDD	HPE 3PAR 8000 8 TB SAS 7.2K LFF (3.5 in) HDD with All-inclusive Single-system Software	HPE 3PAR 8000 1.92 TB+SW SFF SSD
HPE 3PAR 8000 1.2 TB+SW 10 K SFF HDD		HPE 3PAR 8000 3.84 TB+SW SFF SSD
HPE 3PAR 8000 1.8 TB+SW 10 K SFF HDD		HPE 3PAR 8000 7.68 TB+SW SFF SSD
		HPE 3PAR 8000 15.36 TB SAS SFF (2.5 in) SSD
HPE 3PAR 8000 2 TB+SW 7.2 K SFF HDD		

Figure 5-57 HPE 3PAR StoreServ 8000 disk drives

Customers can add drive enclosures to expand the configuration and add LFF drives to the configuration. Drive enclosures can be ordered separately for installation in the field, or they can be factory configured in a rack. Drive enclosures are optional. Because the Storage Base products and the Upgrade Node Pair include SFF drive bays the minimum configuration does not require any additional drive enclosures. For larger configurations, attach drive enclosures. Each drive enclosure includes 24 drive bays. The two-drive enclosure types can be intermixed in a single array.

- The HPE 3PAR StoreServ 8200 supports up to nine added drive enclosures.
- The HPE 3PAR StoreServ 8400 supports up to 22 added drive enclosures.
- The HPE 3PAR StoreServ 8440 supports up to 38 added drive enclosures.
- The HPE 3PAR StoreServ 8450 supports up to 18 added drive enclosures.

Drives are orderable at the time the array is purchased or can be added in the future when additional capacity is required. HPE 3PAR StoreServ 8000 drives, illustrated in [Figure 5-57](#), are sold as single drives. Note that these drives are only compatible with the HPE 3PAR StoreServ 8000 SAS Drive Enclosures.

For each drive type installed in the array, the minimum recommended initial quantity is eight drives per node pair for SSD and SAS performance HDDs, and 12 drives per node pair for nearline HDDs.



Note

Eight drives support RAID 1 and RAID 5. For RAID 6, choose 12 drives.

Requirements and recommendations include:

- The minimum number of SSDs for Adaptive Flash Cache (AFC) is two.
- Minimum upgrade quantity is four drives per node pair or two drives per enclosure, whichever is larger. Best practice is to run Autonomic

Rebalance (also known as tunesys) after adding the drives.

- RAID 6 is strongly recommended for nearline drives.
- All node enclosures must contain either zero or an even number of the same type of drives (FC, nearline [NL], SSD), with a minimum of two drive enclosures.
- A best practice is to add equal numbers of drives to all enclosures compatible with the drive type being added.
- With a four-node configuration, the best practice is to attach the same number and type of drives to each node pair.
- HPE StoreServ 3PAR Express Layout requires a minimum of six SSD drives.

HPE 3PAR StoreServ 8000 models and system features

Model	8200	8400	8440	8450 AFA
CPU	6C Ivy Bridge	6C Ivy Bridge	10C Ivy Bridge	10C Ivy Bridge
Total system memory control/data	2N: 32 GB/32 GB 1333 MHz	2N: 32 GB/32 GB 4N: 64 GB/64 GB 1333 MHz	2N: 128 GB/64 GB 4N: 256 GB/128 GB 1333 MHz	2N: 128 GB/64 GB 4N: 256 GB/128 GB 1333 MHz
Drives supported	Mixed HDD/SSD	Mixed HDD/SSD	Mixed HDD/SSD	SSD Only
MAX drives 2N/4N	HDD: 240 SSD: 120	HDD: 288/576 SSD: 120/240	HDD: 480/960 SSD: 240/480	SSD: 240/480
Raw capacity	750 TiB	1200/2400 TiB	1500/3000 TiB	922/1843 TiB
Host IO		16/32 Gb Fibre Channel 10 Gb iSCSI / FCoE 1 GbE / 10 Gb Ethernet File Persona		

Figure 5-58 HPE 3PAR StoreServ 8000 model comparison

Figure 5-58 shows a comparison between the StoreServ 8000 models.

HPE 3PAR StoreServ 8000—Key performance metrics

Platform	Workload	Perf Metric
8440/8450	8K 100% random read, RAID 1, CPVV, node local	>1M IOPS at <1ms
8440/8450	8K 100% random write, RAID 1, CPVV, distributed	>400K IOPS at <1ms
8440/8450	256K 100% seq. read, RAID 1, CPVV, node local	24 GB/s
8400	8K 100% random read, RAID 1, CPVV, node local	650K IOPS at <1ms
8400	256K 100% seq. read, RAID 1, CPVV, node local	20 GB/s

Figure 5-59 Always verify performance numbers with NinjaSTARS

Figure 5-59 shows key performance metrics for the StoreServ 8000 family.

The ability of HPE 3PAR StoreServ to maintain high and predictable performance in multi-tenant environments is made possible through architectural innovations that use all available array hardware resources at all times, thereby eliminating resource contention, supporting mixed workloads, and enhancing caching algorithms to accelerate performance and reduce latency.

Load balancing

Purpose-built for the enterprise as well as virtual and cloud data centers, the HPE 3PAR StoreServ architecture is unlike legacy controller architectures in that the Mesh-Active system design allows each volume to be active on any controller in the system using the high-speed, full-mesh interconnection that joins multiple controller nodes to form a cache-coherent Active/Active cluster. As a result, the system delivers symmetrical load balancing and utilization of all controllers with seamless performance scalability by adding more controllers and disk drives to the system.

Mixed-workload support

Unlike legacy architectures that process IO commands and move data using the same processor complex, the HPE 3PAR StoreServ controller node design separates the processing of SCSI control commands from data movement. This allows transaction-intensive and throughput-intensive workloads to run on the same storage resources without contention, thereby supporting massive consolidation and multi-tenancy. This means that, for example, the system can easily handle an OLTP application and an extremely

bandwidth-consuming data warehousing application concurrently with ease. This capability is made possible by the HPE 3PAR application-specific integrated circuit (ASIC), which offloads data processing from the control processor, where metadata is processed. By pathing and processing data and metadata separately, transaction-intensive workloads are not held up behind throughput-intensive workloads. As a result, the HPE 3PAR StoreServ platform, as compared to the architectures of traditional storage vendors—including many of today’s all-flash arrays—delivers excellent performance consistently, even in mixed-workload scenarios.



Note

Always verify performance metrics in NinjaSTARS.

HPE 3PAR StoreServ Data at Rest Encryption

Features of HPE 3PAR StoreServ Data at Rest Encryption include:

- An encrypted HPE 3PAR StoreServ array, that is, any HPE 3PAR StoreServ array that has the HPE 3PAR Data Encryption license activated or intended to be activated, must have only self-encrypted drives installed.
- Customers have the option to turn on encryption, non-disruptively, at any time, even after data has been written to the system.
- FIPS 140-2 Validated SEDs have been certified by the U.S. National Institute of Standards and Technology (NIST) and Canadian Communications Security Establishment (CSE) as meeting the Level 2 security requirements for cryptographic modules as defined in the Federal Information Processing Standards (FIPS) 140-2 Publication.
- Strengthen the Data at Rest solution with an optional FIPS 140-2 Level-2 validated external key manager, which supports KMIP 1.1 for key management communications.
- It supports HPE Enterprise Secure Key Manager 4.0 and SafeNet KeySecure k460 and k150 centralized key management.
- A data encryption license to use (LTU) is required to enable encryption

- on the array. One encryption license is required for each encrypted array.
- After encryption is enabled on the HPE 3PAR StoreServ Storage, it cannot be disabled.
 - The local key manager is included in the HPE 3PAR OS. There is not a separate orderable part number for the local key manager.

Adaptive Flash Cache

Adaptive Flash Cache Lowering read latency for HDD

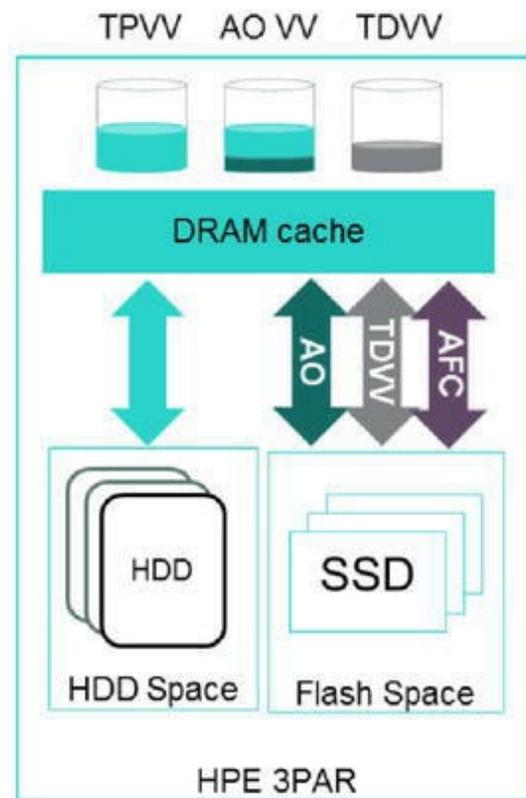


Figure 5-60 Adaptive Flash Cache

Adaptive Flash Cache (AFC), as shown in [Figure 5-60](#), provides read cache extension by dedicating a portion of SSD capacity as an augmentation of the HPE 3PAR array data cache, reducing application response time for read-

intensive IO workloads. AFC and Adaptive Optimization (AO) can coexist. If AO is configured and no available space is available on the SSD tier, a warning limit can be set on the common provisioning group (CPG) SSD tier to free up some space to allocate to AFC. Flash Cache is complementary to Adaptive Optimization—faster response times for read bursts of “cold” AO data residing on HDD.



Note

To understand how much of the existing SSD capacity should be allocated to AFC, refer to the *HPE 3PAR Adaptive Flash Cache* technical white paper at:

<https://h20195.www2.hpe.com/v2/getpdf.aspx/4AA5-5397ENW.pdf>.

Express writes

Express writes

Lowering **write** latency for SSD and HDD

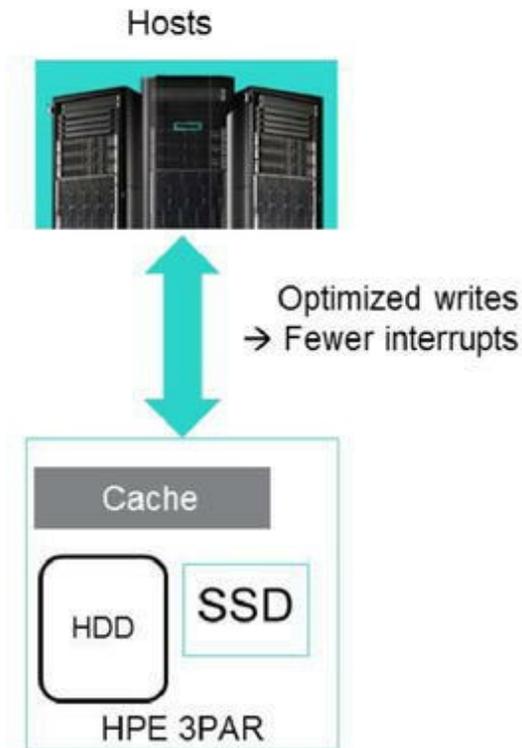


Figure 5-61 Express writes

HPE 3PAR Express Writes, as shown in [Figure 5-61](#), is a write acceleration feature built in to HPE 3PAR OS version 3.2.1 and later to optimize CPU utilization, and depending on workload, deliver greater throughput, up to 30% more IOPS, and 20% reduced latency. All currently supported HPE 3PAR StoreServ Storage HBAs and hosts also support the HPE 3PAR Express Writes feature, and its benefits extend to both spinning drives and flash-based media.

Express writes aim to deliver lower latency by delivering lower CPU interrupts per IO. This is achieved by sending the data along with the command rather than waiting for the target to request it. Previously, this was only available with FC protocol, but has now been extended to iSCSI for the 8000 and 20000 systems. This will be enabled automatically on upgrade and

can result in up to 40% improvements in latency for iSCSI writes.

HPE 3PAR StoreServ Express Layout

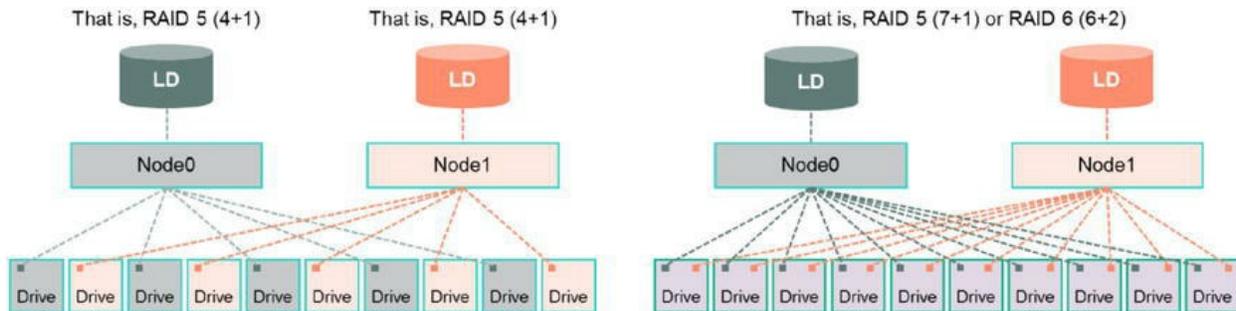


Figure 5-62 HPE 3PAR StoreServ Express Layout

This unique technology born from the HPE 3PAR three-layer virtualization technology allows 3PAR controller nodes to share access to SSDs to further drive efficiency. Replacing traditional layouts for flash, Express Layout, as illustrated in [Figure 5-62](#), allows each SSD to be actively accessed by both controllers in a node pair at the same time. This allows a node pair to use capacity from every drive to build logical capacity. For smaller configurations, like an eight-drive system, Express Layout allows the nodes to significantly reduce the overhead historically associated with parity RAID layouts and can result in more than 10% reduction in overhead which increases performance by allowing more than one controller to drive IO through the drive.

Every logical disk (LD) has an owner and a backup owner. Using the traditional layout, chunklets from any given physical disk (PD) are owned by a single node with the partner node as the backup owner; thus, every node creates LDs from the PDs it “owns.” Express Layout alters ownership for flash drives. If the set size configured requires more than 50% of the drives behind the node pair, then the LD will be created using chunklets from PDs behind the node pair, allowing each node to create LDs larger than traditionally possible. This allows smaller flash systems to create larger set sizes, reducing RAID overhead and improving capacity efficiency.

Traditional 3PAR logical disk layout:

- A physical drive is owned by one node.
- Logical disks/RAID sets are built by node with chunklets from owned drives.
- Standard layout for all configurations.
- Capacity inefficient in very small configurations.

3PAR Express Layout:

- A physical drive is owned by a node pair.
- Provides more capacity efficiency with small numbers of drives.
- LDs are built by node with chunklets from all drives*.
- Introduced for SSD with 3PAR OS 3.2.2.
- Support expanded to all drive types with 3PAR OS 3.3.1.



Note

*Caution for two-node systems only: To allow the array to create new LD even after a drive failure, keep the set size smaller than the installed number of drives, that is, with 10 drives use R5 7+1 or 6+2, with 8 drives use R5 6+1 or R6 4+2.

HPE 3PAR Adaptive Data Reduction with 3PAR OS 3.3.1



Do not compromise.

By implementing the new supported features, HPE can increase the savings and performance.

Integrated: Zero detect

Removes 16 KB pages with zeroes only

Redesigned: Deduplication

Removes duplicate 16 KB data pages inline

New: Compression

Removes redundant data inline

New: Data packing

Packs multiple compressed data pages into a single 16 KB page

Figure 5-63 ADR includes zero detection, deduplication, compression, and data packing

Adaptive Data Reduction (ADR), as shown in [Figure 5-63](#), is an add-on to the traditional 3PAR thin technologies. It is an umbrella term used to cover a set of technologies on the 3PAR designed to reduce the total capacity required to store thin volumes. The technologies included in ADR are zero detect, deduplication, compression, and data packing.

The goal of ADR is to reduce the cost of flash. With the exception of zero detect, the technologies included on the slide are only supported for flash storage on the 3PAR.

The technologies included in Adaptive Data Reduction happen in this order:

- **Zero detection**—When data for a thin provisioned volume is received from a host, it is put into the cache of the controller node, mirrored to another controller node, and the write is acknowledged back to the host. The ASIC then examines the data looking for full 16 K pages of zeros. If a 16 K page of zeros is received, no data will be written to the backend disks. If 16 K of zeros are received, where data had previously been written, space on the backend disk is reclaimed. Zero detect can be enabled or disabled for any thin provisioned virtual volume (default is enabled). Zero detect can be enabled for any thin volume from any media type: NL, FC, or SSD.
- **Deduplication**—The ASIC checks to see if the data in a 16 K page incoming write from a host matches any 16 K page of data written to SSD capacity from any volumes in the same CPG. If there is a match, the duplicate data is not written to the backend SSDs. Dedup can be enabled or disabled on a volume by volume basis, but dedup is only allowed for volumes stored on SSD.
- **Compression**—Compression is performed by the CPU on the 3PAR, and looks for opportunities to reduce the number of bits required to represent the data. Like deduplication, compression is only supported for volumes stored on SSD.
- **Data packing**—After a 16 K page of data has been compressed, it will be combined with other compressed pages into one 16 K page and written to

the backend SSDs. This reduction technology is referred to as data packing. Compression can be enabled on a volume by volume basis, and data packing will automatically be configured for all compressed volumes.

HPE 3PAR StoreServ efficiency

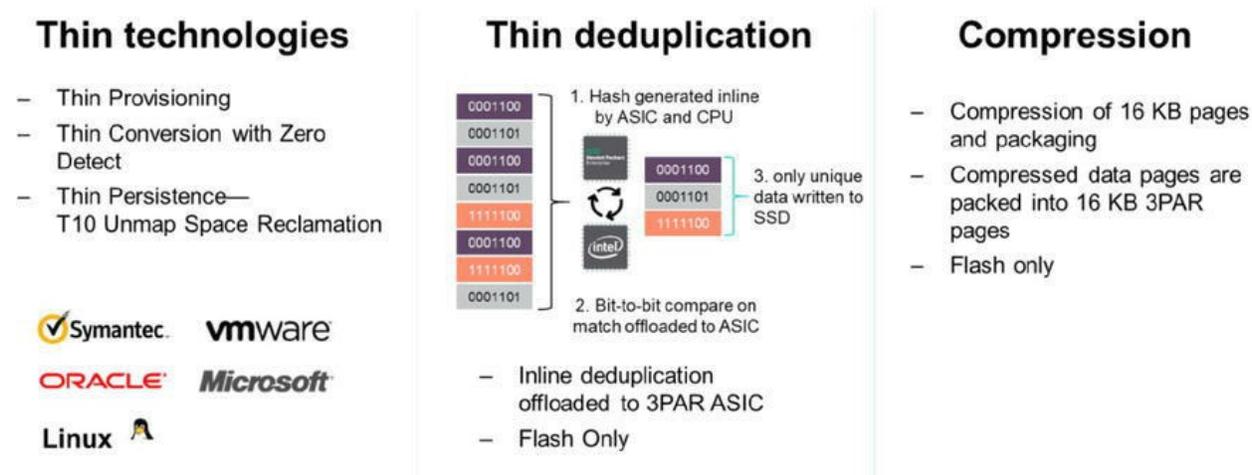


Figure 5-64 Maximizing storage efficiencies

A Thin Provisioned Virtual Volume (TPVV), as illustrated in [Figure 5-64](#), uses logical disks that belong to a logical disk pool known as a CPG. TPVVs associated with the same CPG draw user space from that pool as needed, allocating space on demand in 1 GB chunklet increments beginning with 1 GB per controller node. As the volumes that draw space from the CPG require additional storage, the system automatically creates additional logical disks and adds them to the pool until the CPG reaches the user-defined growth limit that restricts the CPG's maximum size, as indicated in [Figure 5-64](#).

Thin deduplication

For the first write, a hint is created on deduplication store (DDS) and data is saved on deduplication client (DDC). For the second write, it finds the hint and saves the data on DDS. At the same time, a background task (called conversion) is created and sent to the first write owner. The first write owner

receives a conversion request and sends data to the DDS, finds the match, modifies the exception entry to point to DDS, and frees the old DDC location. If we have more writes with the same data signature come in, the behavior will be same as TDVV.

Pros:

- Much less garbage (key benefit, should help avoid lot of issues seen in the past)
- Better performance for read since most of data is saved on local DDC
- Less metadata IO overhead

Cons:

- Extra cost on over-writes, new hint creation, and old hint retirement (page-in old data and calculate data signature)
- Extra cost for second writes of same data due to the additional overhead for hint conversion

Dynamic Optimization introduction

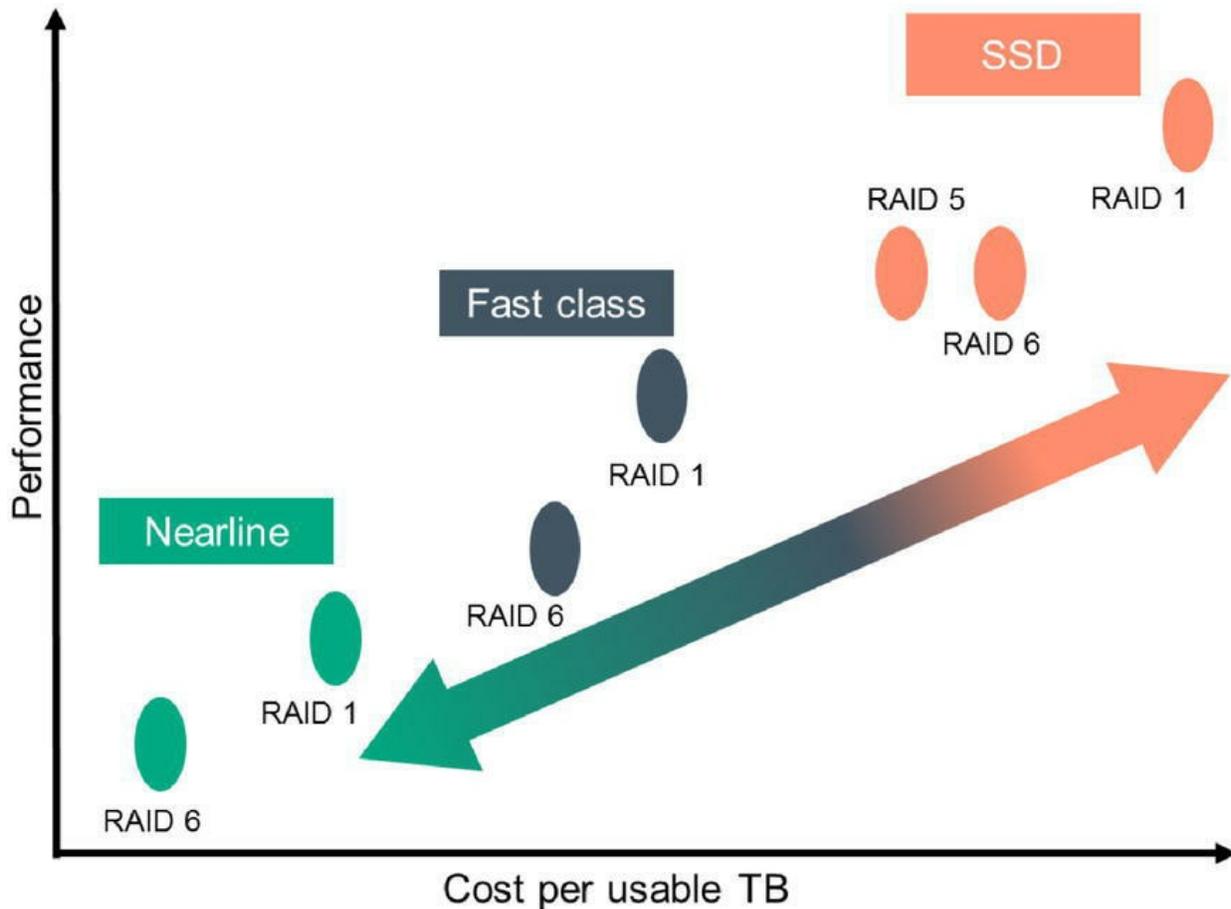


Figure 5-65 Change from disk type, RAID on the fly

HPE 3PAR Dynamic Optimization (DO), as illustrated in [Figure 5-65](#), is an HPE 3PAR OS feature that enables the administrator to dynamically tune volumes by changing volume parameters, such as RAID level, set size, and disk drive types. To perform these tasks, an HPE Dynamic Optimization license is required.

Dynamic Optimization allows for the non-disruptive restructure of virtual volumes. This functionality promotes system optimization through improved utilization of all physical resources present at a given time. In addition, DO enables altering of service levels associated with a given volume by allowing the administrator to change the volume parameters.

Using DO, the administrator can alter the service level for a specific volume after it has already been created, that is, changing from FC to NL, NL to FC, altering RAID levels, and set sizes. There is no impact to the volume while

under operation, and this can be performed during normal processing.

This feature allows the administrator to change service levels when either incorrect application performance information has been given, or history shows that the volume is either under- or overusing the resources that have been provisioned.

Often, the administrator is given certain service-level requirements where the expectation of the specific business unit has been overestimated, incorrectly suggesting the use of high-end, fast-class disks, and thus stranding the volume to a specific build type.

With DO, along with historical reporting, the user is free to change service levels to the appropriate settings after the volume has already been deployed and under use.

HPE 3PAR StoreServ Dynamic and Adaptive Optimization

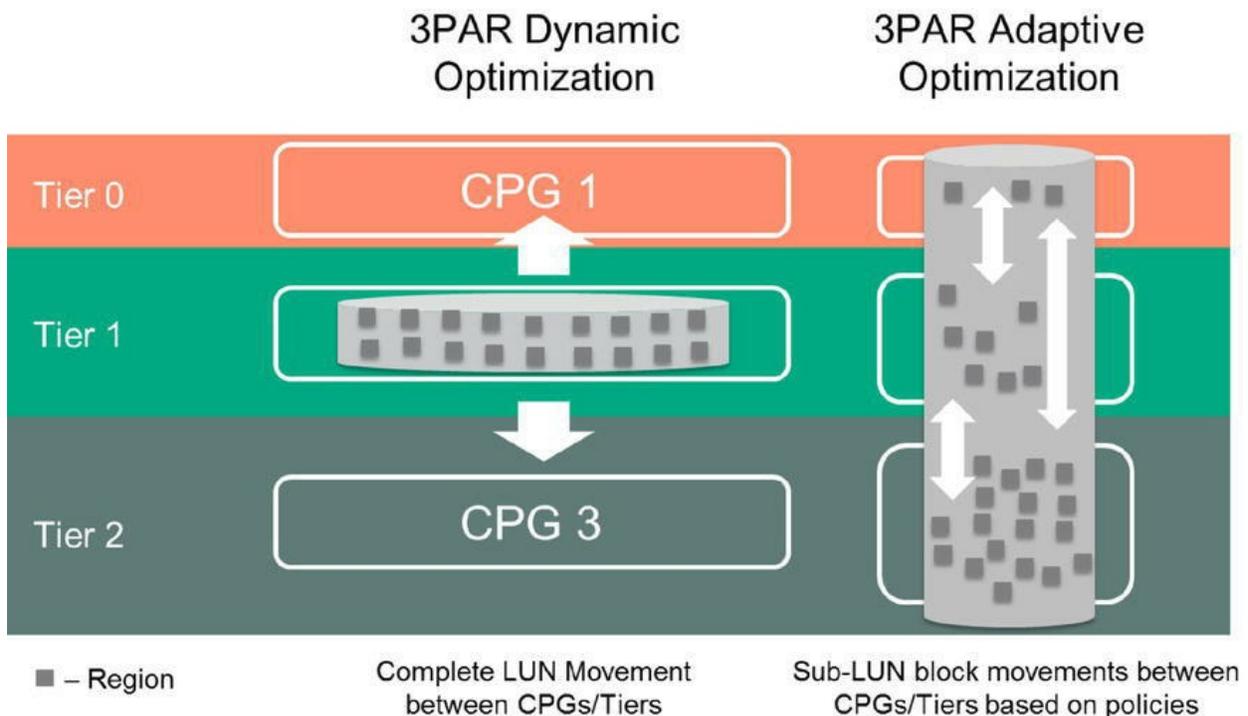


Figure 5-66 HPE 3PAR StoreServ Dynamic and Adaptive Optimization

Dynamic Optimization, as illustrated in [Figure 5-66](#), is an optional feature that allows entire VV/VLUN movement from one CPG to another for cost

reasons, performance reasons, or space utilization reasons in real time with no downtime. Adaptive Optimization is different in that data from VV/VLUN can be distributed across various media using subvolume-based algorithms.



Note

When the terms “region movement” or “region analysis” are used, this refers to the 128 MB region of user addressable space. Regions housing user data are not the only regions moved. Regions housing metadata can also become candidates for movement to a different tier. For a fully provisioned virtual volume, the region capacity is 128 MB of contiguous LD space. For a thin provisioned virtual volume, the region capacity might or might not be contiguous space.

AO analyzes performance statistics collected by System Reporter to determine whether data for a volume should move and, if so, which tier it should move to (up or down or not at all). It uses historical data to determine which data to move. AO expects that data access patterns from the past will continue for some time into the foreseeable future following the data movement.

AO operates on an element of data called a “region” and has a capacity of 128 MB of contiguous space on an LD.

Adaptive Optimization software analyzes subvolume, region-level disk access rates for a given array over a scheduled period, and then performs a data migration of regions between tiers according to a cost versus performance preference. Disk usage is optimized by moving frequently accessed data to the higher performance tier.



Note

AO and TDVV, compressed virtual volumes, and dedup/compressed (DECO) can coexist on the same HPE 3PAR array, but an SSD CPG that acts as a tier in an AO configuration cannot have any of these volumes provisioned using the CPG. The system will prevent creating an AO configuration if a CPG has these

volumes associated with it. A system can be configured with a shared pool of SSDs that may be used for sub-tiering (AO), cache augmentation (AFC), and provisioning of TDVVs, TPVVs/compressed/DECO, or fully provisioned VVs.

HPE 3PAR StoreServ Adaptive Optimization

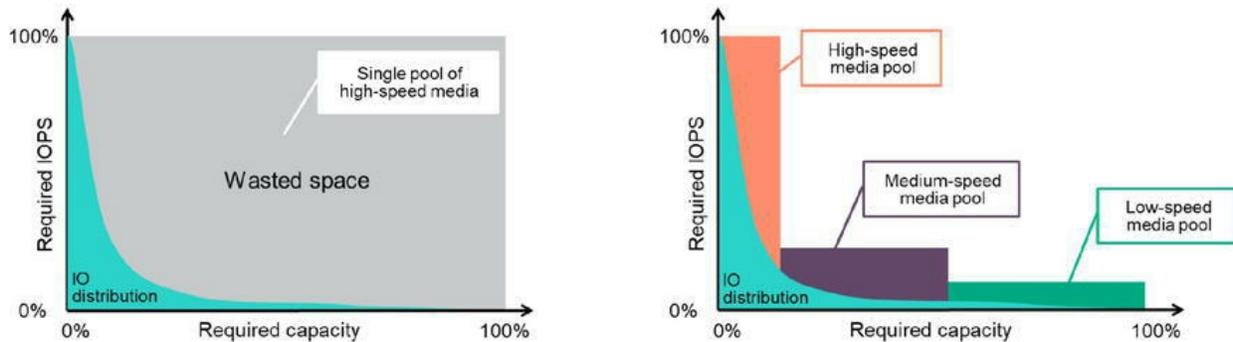


Figure 5-67 Correct sizing is important when using AO

When new data (new virtual volumes or new user space for a thin volume) is created, it will be created in the default CPG defined at volume creation time (best practice is to use the FC tier). Adaptive Optimization will not migrate regions of data to other tiers until the next time the AO configuration is executed. It is therefore important that the FC disks have enough performance and capacity to accommodate the performance or capacity requirements of new applications (or applications) that are provisioned to the system. It is a best practice to size the solutions assuming the NL tier will contribute 0% of the IOPS required from the solution. It is important to size the array correctly when setting up an HPE 3PAR array with Adaptive Optimization, as indicated in [Figure 5-67](#).

Sizing for SSD tier

It is relatively straightforward to size the FC tier based on the IOPS requirement, but the IOPS delivered by the SSD tier will be based on the number of SSD drives in the array with respect to the overall capacity as well as the application data locality. The following examples walk through how much IOPS to expect per SSD in different scenarios for an array sized

for 50 TB:

- If the SSD tier is **less than 5%** of the total capacity, then it will occupy between 1.25 TB and 2.5 TB of the total capacity. This requirement can be met with 200 GB SSD drives and each drive will get on an average 550 IOPS. Even though each SSD drive can handle a lot more IOPS, less amount of capacity used by SSD tier will translate to fewer number of IOPS per drive. This is because there are very few applications that will do very heavy IO to such a small percentage of their data space.
- If the SSD tier is **between 5% and 10%**, then it will occupy between 2.5 TB and 5 TB of the total capacity. This requirement can be met with 400 or 480 GB SSD drives, and each drive will get around 1000 to 1100 IOPS. The number of IOPS per drive increased as we added more capacity by doubling the drive size.
- If the SSD tier is **more than 10%**, then the customer will choose the 920 GB SSD drives. Each of these drives will provide on an average 2150 IOPS. This is possible because the SSD capacity has increased by doubling the size of each drive.

AO sizing for 2-tier: SSD+FC

Sizing an AO configuration with a two-tier configuration using SSD and FC will also begin by performance sizing the FC tier. The same guidelines of 150 IOPS for 10 k rpm FC drives and 200 IOPS per 15 k rpm drives are used to start. A target of 100,000 IOPS and 425 TB raw capacity, for example, can be sized by dividing 150 IOPS per 10 k FC drive by the IOPS goal of 100,000 to determine how many 10 k FC drives are required. In this example, 667 ($100000/150$) 10 k FC drives will meet the target of 100,000 IOPS.

Next, we choose a target SSD configuration and calculate how many FC drives can be removed from the configuration to remain at the target IOPS level. IOPS guidelines for SSDs can be used like the FC IOPS guidelines and traded off to maintain the target IOPS while changing the mix of drives. Use an AO configuration IOPS guideline of 1100 IOPS for the 480 GB MLC SSDs and 4000 IOPS for the 1920 GB and 3840 GB SSDs.

Choosing a configuration of 48 x 480 GB SSDs will result in an SSD tier

capable of 52,800 ($48 * 1100$) IOPS. Reduce the number of drives in the FC tier by the same 52,800 IOPS to keep the same target IOPS of 100,000. The above configuration included 10 k rpm FC drives, so we use the 150 IOPS per drive guideline to calculate the number of 10 k FC drives we can remove from the configuration. The result is 352 ($52,800 \text{ SSD IOPS} / 150 \text{ 10 k FC drive IOPS}$) 10 k FC drives can be removed from the configuration to trade the FC IOPS for the SSD IOPS. The two-tier AO configuration to provide 100,000 IOPS has now been sized with 48 x 480 GB SSDs and 315 (667-362) 10 k FC drives.

The last step is to make adjustments to meet the capacity target of 425 TB by adding FC drives. The configuration so far includes 48 x 480 GB SSDs providing 23 TB raw. Choosing 315 of the 1.2 TB 10 k FC drives will provide 377 TB of raw capacity and a total of 400 TB (23 TB SSD + 377 TB FC) raw capacities. 25 TB additional raw capacity is required to meet the target of 425 TB raw and can be achieved by adding (20) 1.2 TB 10 k FC drives. The AO configuration sized to meet 100,000 IOPS and 425 TB raw capacity using a two-tier SSD+FC configuration is 48 x 480 GB SSDs and 340 (315+25) 1.2 TB 10 k FC drives.

AO sizing for 3-tier: SSD+FC+NL

Sizing an AO configuration for a three-tier configuration using SSD, FC, and NL tiers will begin like the other examples with performance sizing of the FC tier. The sizing process for a three-tier SSD+FC+NL AO configuration follows the prior discussion of sizing a two-tier SSD+FC AO configuration with one minor modification. Sizing the FC tier and trading SSD IOPS for FC IOPS remains the same. The only difference is the last step to adjust the configuration to meet the capacity target where NL drives are added instead of FC drives.

Using a target of 125,000 IOPS and 450 TB raw as an example, start by sizing the FC tier, using the previous guidelines of 150 IOPS per 10 k rpm FC drive and 200 IOPS per 15 k rpm FC drive. Using 15 k rpm FC drives in this example, 625 drives are needed to meet the target of 125,000 IOPS. Adding 64 x 480 GB SSDs will add 70,400 IOPS, allowing the SSDs to replace the IOPS of 352 15 k rpm FC drives ($70,400 / 200 = 352$).

The AO configuration after sizing the FC and SSD tiers is made up of 64 x

480 GB SSDs holding 31 TB raw and 273 (625 – 362) 15 k rpm FC drives. If we choose 600 GB 15 k RPM FC drives, this represents 164 TB raw space and when adding the 31 TB raw from the SSD tier reaches a total of 195 TB raw. The NL tier will be sized to provide 0% of the IOPS target and capacity to fill the gap from 195 TB raw in the SSD and FC tiers to the target of 450 TB raw. The NL tier must provide 255 TB raw (450–195) capacity which can be achieved with 64 x 4 TB NL drives. In this example, a three-tier AO configuration is sized to provide a target of 125,000 IOPS and 450 TB raw with 64 x 480 GB SSDs in Tier 0, 273 x 600 GB 15 k rpm FC drives in Tier 1, and 64 x 4 TB NL drives in Tier 2.

AO sizing for 2-tier: FC+NL

Sizing an AO configuration for any set of tiers begins with performance. In all AO sizing, the NL tier will be sized to handle 0% of the IOPS making sizing for a two-tier FC+NL AO configuration straightforward. Begin by sizing the FC tier to handle 100% of the IOPS and add NL drives to meet the capacity target.

When sizing the FC AO tier, use a guideline of 150 IOPS per 10 k FC drive and 200 IOPS per 15 k drive. A target of 50,000 IOPS and 450 TB raw capacity, for example, can be sized by dividing 150 IOPS per 10 k FC drive by the IOPS goal to determine how many 10 k FC drives are required. In this case, 333 (50000/150) 10 k FC drives will meet the target of 50,000 IOPS. Choosing 600 GB 10 k FC drives will result in the FC tier of an AO configuration performing the target 50,000 IOPS and providing 200 TB of raw capacity. The addition storage needed to meet the target of 450 TB can be provided by adding (63) 4 TB NL drives.

Learning check

4. When will you recommend HPE SimpliVity 380 H over HPE SimpliVity 380 G?

5. What is the name of the MSA technology allowing an administrator to set tier for a volume?

6. Describe the read operation at HPE Nimble HF Storage.

7. What are the typical building blocks of the HPE 3PAR StoreServ 8000 solution?

HPE Primera 600 Series

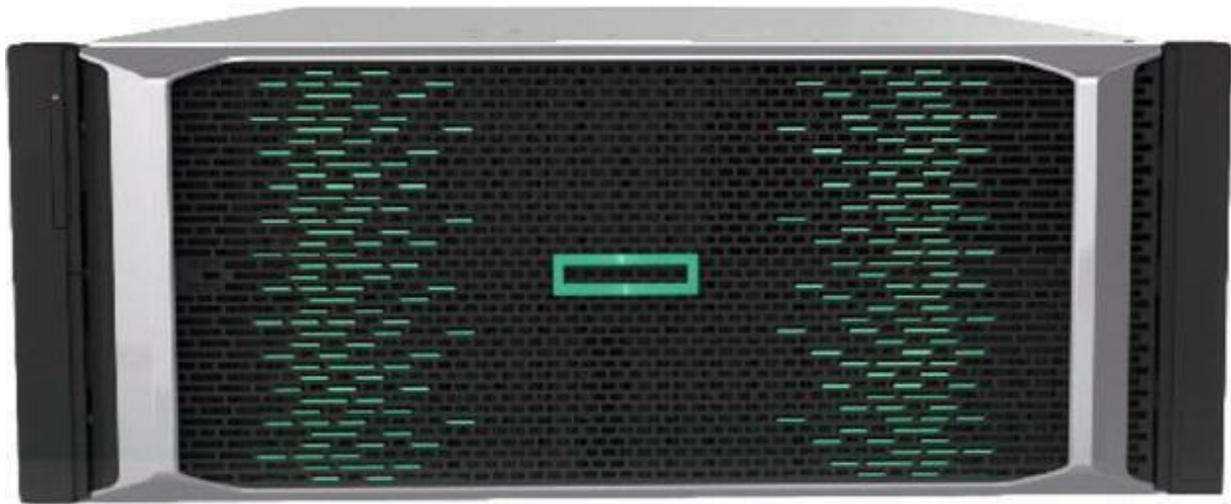


Figure 5-68 HPE Primera 600 Series

HPE announced HPE Primera, intelligent storage for mission-critical apps that delivers:

- On-demand experience for instant access to data and consumable as a service
- App-aware resiliency that sees beyond the walls of storage to predict/prevent disruptions
- Predictive acceleration to safely consolidate every mission-critical app on the same platform with extreme low-latency performance

HPE Primera, as shown in [Figure 5-68](#), is a Tier-0 enterprise storage solution that delivers extreme resiliency and performance with the agility of the cloud. Built upon proven resiliency and powered by the intelligence of HPE InfoSight, HPE Primera delivers instant access to data with storage that sets up in minutes, upgrades transparently, and is delivered as a service. Ensure always fast and always on storage for all mission-critical applications.

HPE Primera solutions are covered by HPE Proactive Care and a 100% data availability guarantee without requiring special contracts or onerous terms. Your customers can ensure no disruptions are ever felt with app-aware resiliency, guaranteed across all models of HPE Primera. If your customers experience less than 100% availability, HPE works with them to resolve the issue and provide financial credit to apply toward a future purchase of HPE Primera products.

HPE Primera is built on a series of innovations across intelligence, platform, software, and the business model:

- **Global intelligence**—HPE Primera is natively integrated with HPE InfoSight in the cloud and now in the array too. HPE has taken the most advanced AI for infrastructure and embedded it inside HPE Primera to provide end-to-end AI pipeline for self-managing storage.
- **All-active architecture**—HPE Primera has a massively parallel architecture optimized for NVMe and Storage Class Memory. Data simultaneously accesses all resources in the system and can instantly failover for extreme low-latency performance at scale and higher resiliency.
- **Services-centric OS**—A unique OS in its class that eliminates risk and simplifies management. The OS decouples the data services, so they can be independently deployed, upgraded, and restarted—eliminating update risk and enabling faster access to new innovation.
- **Timeless storage**—A new ownership experience, extended from Nimble Storage, that eliminates the pain of forklift upgrades.

HPE Primera comprises three models, as indicated in [Figure 5-69](#): HPE Primera 630, HPE Primera 650, and HPE Primera 670. Each model is available as an all-flash version (A630, A650, and A670).



Figure 5-69 Each model available as “all-flash” (A600) or “converged flash” (600)

HPE Primera technical specs

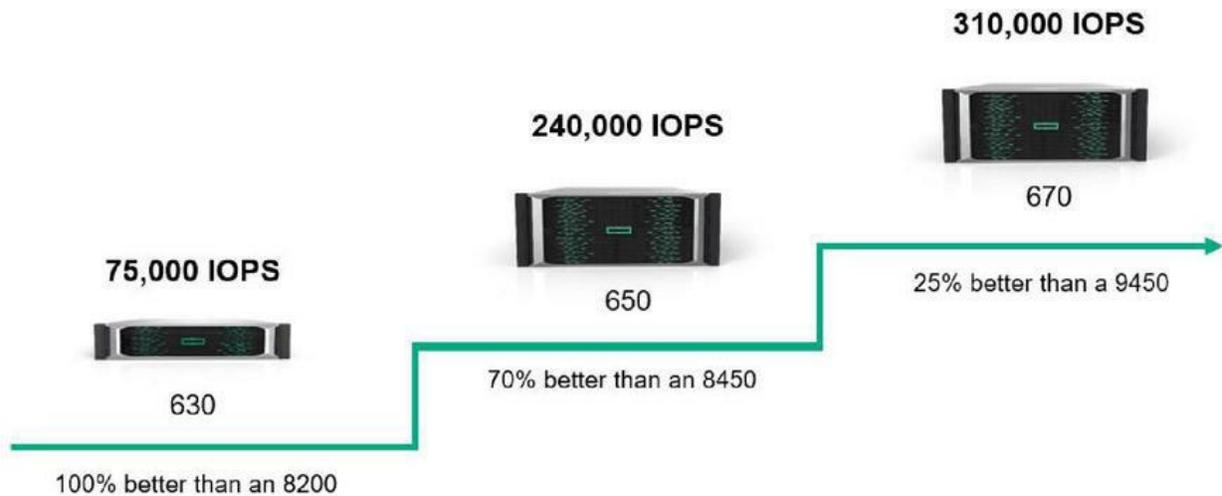
Spec	630	650	670
Base chassis	2U	4U	4U
Max number of nodes	2	4	4
Max SAS SFF / NVMe drives in base chassis	24 / 8	48 / 16	48 / 16
Maximum host ports	16	48	48
Built in 10 GbE ports (SFP+) per node	2	2	2
FC ports (32 Gb/s or 16 Gb/s) per node	8	12	12
Max number of SSDs	144	384	576
Max raw capacity (SSD only)	250 TiB	800 TiB	1600 TiB
Max number of add-on drive enclosures	5 enclosures (A630)	14 enclosures (A650)	22 enclosures (A670)



Note

Not all features are available at launch. Consult QuickSpecs.

HPE Primera performance



Note

Performance numbers are for DECO, RAID6, 16 K 60/40 RW, and node-distributed configurations.

Learning check

8. Match the HPE Primera model with its number of controllers.



HPE backup and restore strategy

One of the most common causes of data loss is physical failure of the storage media. Normally, hard drives function for years without incident, but eventually they fail. It might happen gradually as more bad clusters accumulate, until most of the drive is unusable, or it might happen suddenly, when the hard drive just fails without warning.

What is backup?

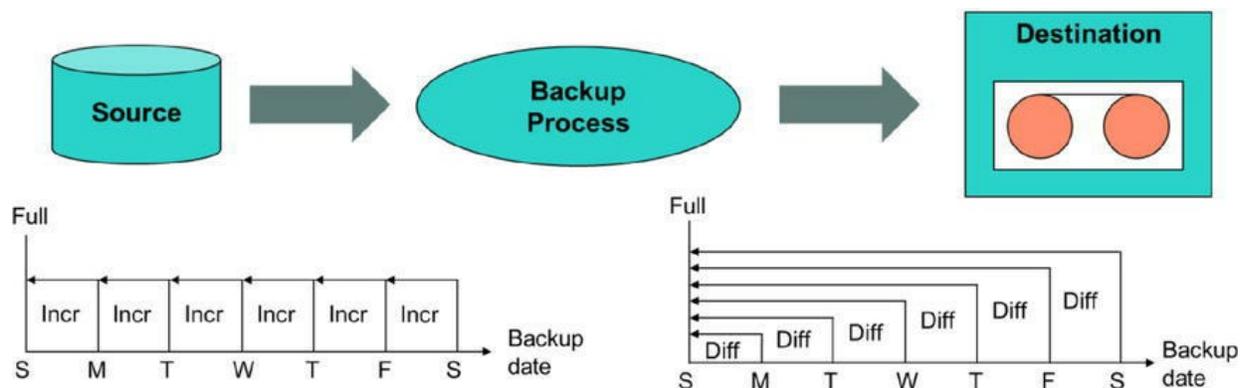


Figure 5-70 Incremental and differential backups

Backup is a process that creates a copy of data on backup media. This copy

is stored and kept for future use in case original is destroyed or corrupted.

There are two main types of backups, as illustrated in [Figure 5-70](#):

- **Incremental**—Copies all files that were changed after the last backup, regardless of what kind of backup it was. Incremental backups back up only the data that changed since the last full backup or previous incremental, and reset the archive bit.
- **Differential**—Copies all files that were changed after the last complete backup. Differential backups back up all the data since the last full backup but do not reset the archive bit.

After a full system backup has been performed, it should be verified with the existing data on the system. All backup software provides the capability to verify the backed up copy with the original data on the system. Confirm that the verification process has completed successfully before proceeding with the upgrade.

A basic approach to improve performance is to reduce the amount of data to back up. Careful planning is necessary for full and incremental backups. Full backups of all the client systems might not need to be performed at the same time.

What is restore?

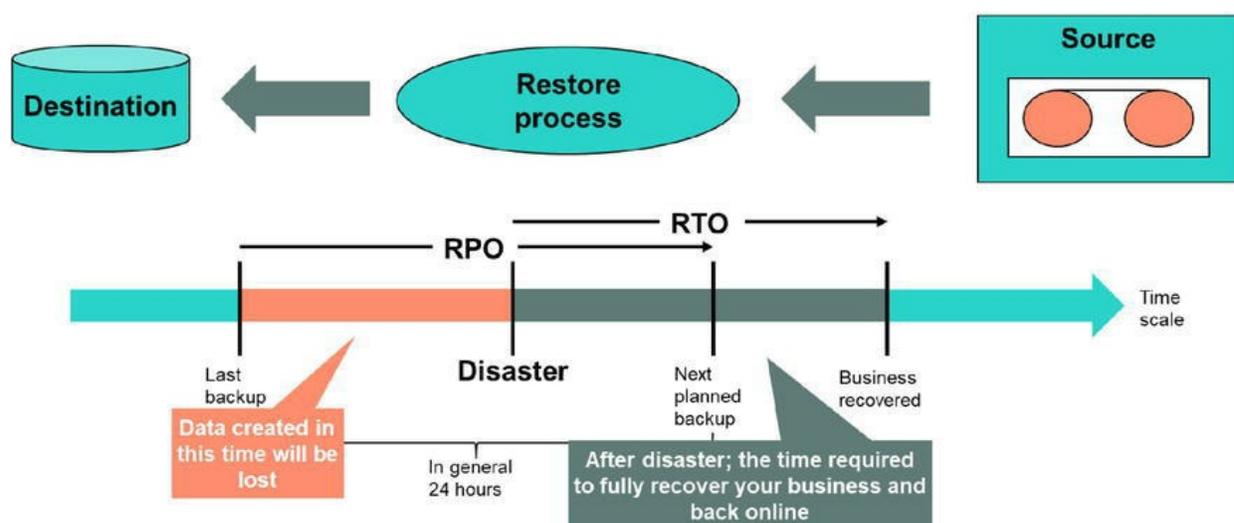


Figure 5-71 Restoring data form a backup

A restore, as illustrated in [Figure 5-71](#), is a process that recreates the original data from a backup copy. This process consists of the preparation and actual restore of data, and some post-restore actions that make that data ready for use. The source is a backup copy. A restore application is software that writes data to a destination. The destination is usually a disk to which the original data is written.

The recovery time objective (RTO) is the length of time that it is acceptable for the business to be without a specific application. This is often associated with the maximum allowable or tolerable outage. RTO is used to establish use of replication or backup to tape or disk. If downtime is not an option, RTO would be zero. In this case, a clustered or redundant infrastructure with replicated data offsite would be an ideal solution. If RTO is measured in days, tape backup might be acceptable for a specific application.

The recovery point objective (RPO) is the point (before a failure or outage) to which systems and data must be restored. The RPO dictates the allowable data loss and thus the type of data protection solution necessary.

Data deduplication can reduce the need for backup storage space by up to 20 times. Specific savings depend on the number of backups, the data retention policy, and how much the data changes.

Backing up the environment

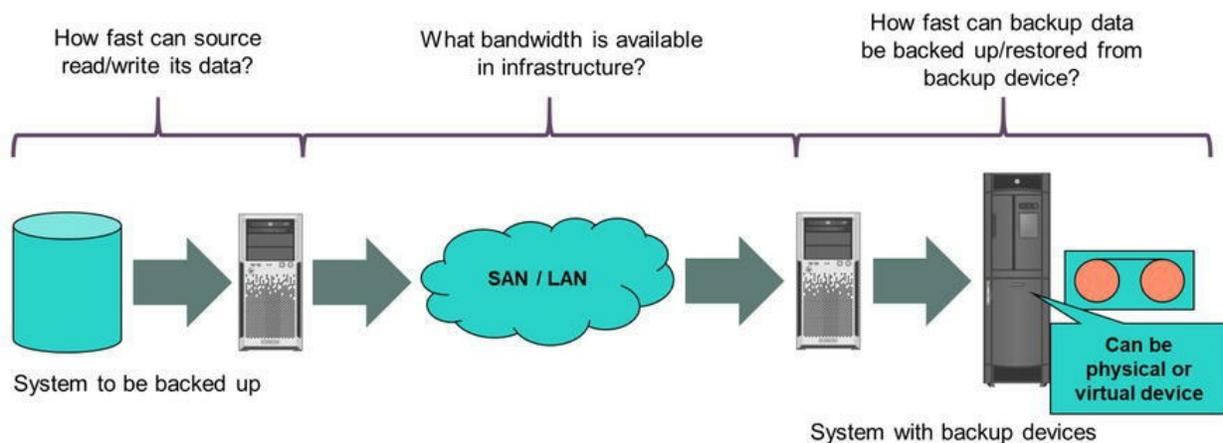


Figure 5-72 Backing up data over the SAN/LAN

During backups, as shown in [Figure 5-72](#), data is transferred over the infrastructure (SAN/LAN) from systems to be backed up to media on systems that have backup devices, where the data is stored.

Backup configuration

Backups must be performed on a disciplined, scheduled regimen, carefully designed to address certain issues, including the following:

- The frequency and schedule for backups
- Whether to use partial or full backups
- Whether partial backups should be incremental or differential

The terms full, partial, incremental, and differential describe the amount of information that is copied and backed up. All backups, whether full or partial, can be performed online or offline and are based on the status of an archive bit.

The archive bit is a file attribute that identifies whether the file has been backed up. If a file is changed after the last backup, the archive bit is cleared until the file is backed up again.

Full backup

A full backup is a complete backup of the entire server or client. A server backup includes all volumes, directories, and files. A client backup includes all drives, directories, and files.

The two types of full backups are:

- **Normal backup**—Backs up all files on the hard drive and resets the archive bit
- **Copy backup**—Similar to a normal backup but does not reset the archive bit

Backup strategies

It is important to understand that backing up is a process. As long as new

data is being created and changes are being made, backups will need to be updated. Individuals and organizations with anything from one computer to thousands (or even millions) of computer systems all have requirements for protecting data. While the scale is different, the objectives and limitations are essentially the same. Likewise, those who perform backups need to know to what extent they were successful, regardless of scale. Considerations include:

- Organizational policies
- Type of data
- Maximum downtime
- Data retention plan
- Volume of data for backup and frequency of data change
- Projected future growth of data volume
- Backup performance expectation (duration) and available infrastructure (storage and network)
- Periodicity of backup
- Process of maintaining data
- Sensitivity of data
- Disaster recovery plan

Reasons why backups fail

Do not assume that a disk-based backup eliminates and protects your customers from media-related failures. It is true that the possibility of media failures is considerably lower than tape, but it is important to add in resiliency and redundancy. For example, RAID 6, redundant power supply unit (PSU), hot spare disks, or disaster recovery systems should always be considered when creating a backup solution.

Human error is nearly always the main reason for backup failure. Clearly the best safeguard to prevent human error is to ensure best practices are followed and ensuring the backup administrators know and understand the steps for performing backup and restores. This can be improved upon by

removing the backup administrator from the loop by ensuring there is no user interaction to perform a backup operation.

Sometimes new software or new versions and patches can cause a backup failure. This is usually caused by a configuration error—back to the human problem. Modern backup software is very flexible, which provides for many variations and options, and this can create the problem that a wrong option has been chosen resulting in a backup failure or a backup that is not complete.

Backing up over SAN or network increases efficiency; this is achieved by consolidation of a backup solution, reducing the number of backup devices and sharing of this resource across different platforms. This introduces another point of failure to the backup solution HBAs, switches, and zoning. For example, if these are incorrectly configured, this could lead to a backup failing. However, as the SAN is used for other things beside a backup, any problem in this area will probably become obvious before it can create an issue with the backup.

Snapshot functionality

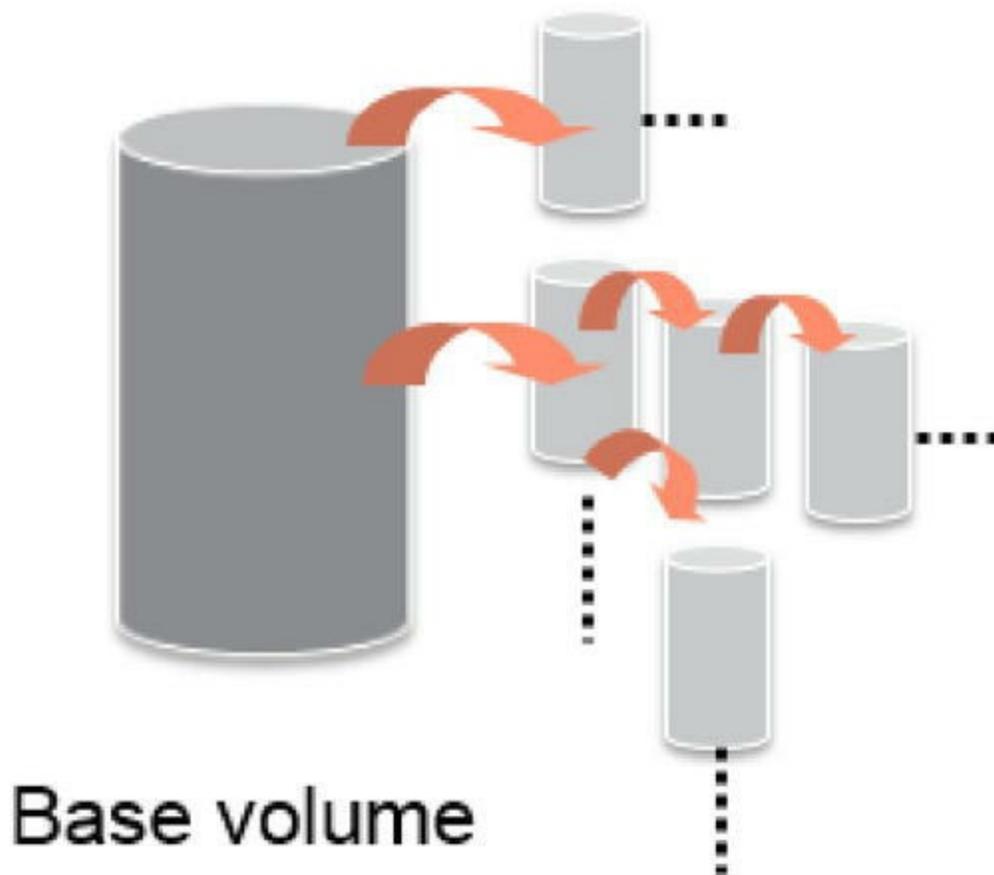


Figure 5-73 Snapshots are not a real backup

The base volume, as illustrated in [Figure 5-73](#), is the original volume that is to be copied. Unlike a clone, which is a duplicate of an entire volume, a snapshot only records changes to the base volume (copy-on-writes). This allows an earlier state of the original virtual volume to be recreated by starting with its current state and rolling back all the changes that have been made since the snapshot was created. Snapshots are ordinarily more space efficient than a clone.

Snapshots can be made of Fully Provisioned Virtual Volumes (FPVV), Thinly Provisioned Virtual Volumes (TPVVs), Thinly Deduplicated Virtual Volumes (TDVVs), clones, or another snapshot (snap-of-a-snap). Snapshots are created using copy-on-write techniques. Thousands of snapshots of each virtual volume can be created, assuming that sufficient storage space is available, and the storage array platform supports it.

A snapshot can be only removed if it has no children, but a snapshot from

different subtrees can be removed individually. Snapshot works much along the same lines as Thin Provisioning and takes advantage of space efficiencies.

Snapshots are promotable, not interdependent on each other, can be routinely created or deleted with the scheduler, and can maintain consistent snapshots over a group of volumes. Snapshots can be promoted to other base volumes besides the original volume from which the copy originated. Consistency group snapshots can be maintained for clustered environments such as VMware or databases like Oracle using Automatic Storage Management (ASM) as a disk management tool.

Snapshots are thin in nature even if the originating volume is fully provisioned. There is no need to create a separate pool of storage for snapshots, although a space must exist for the snapshot from which to gain its policy settings. It is also possible to create a snapshot that has a different quality of service (QoS) than the base volume, such as making a copy of a RAID 1 FC volume to a RAID 5 NL volume.

The snapshot product is also capable of taking multiple snapshots of existing copies and creating separate snapshot trees—again, not interdependent of the previous snapshots.

The single most important point about the snapshot technology is that it requires only a single copy-on-write to maintain the deltas from the base volume. Performance levels of the base volume only change by one additional penalty, regardless of the number of snapshots that are taken against it.

Snapshot: How it works—Create read-only snapshot

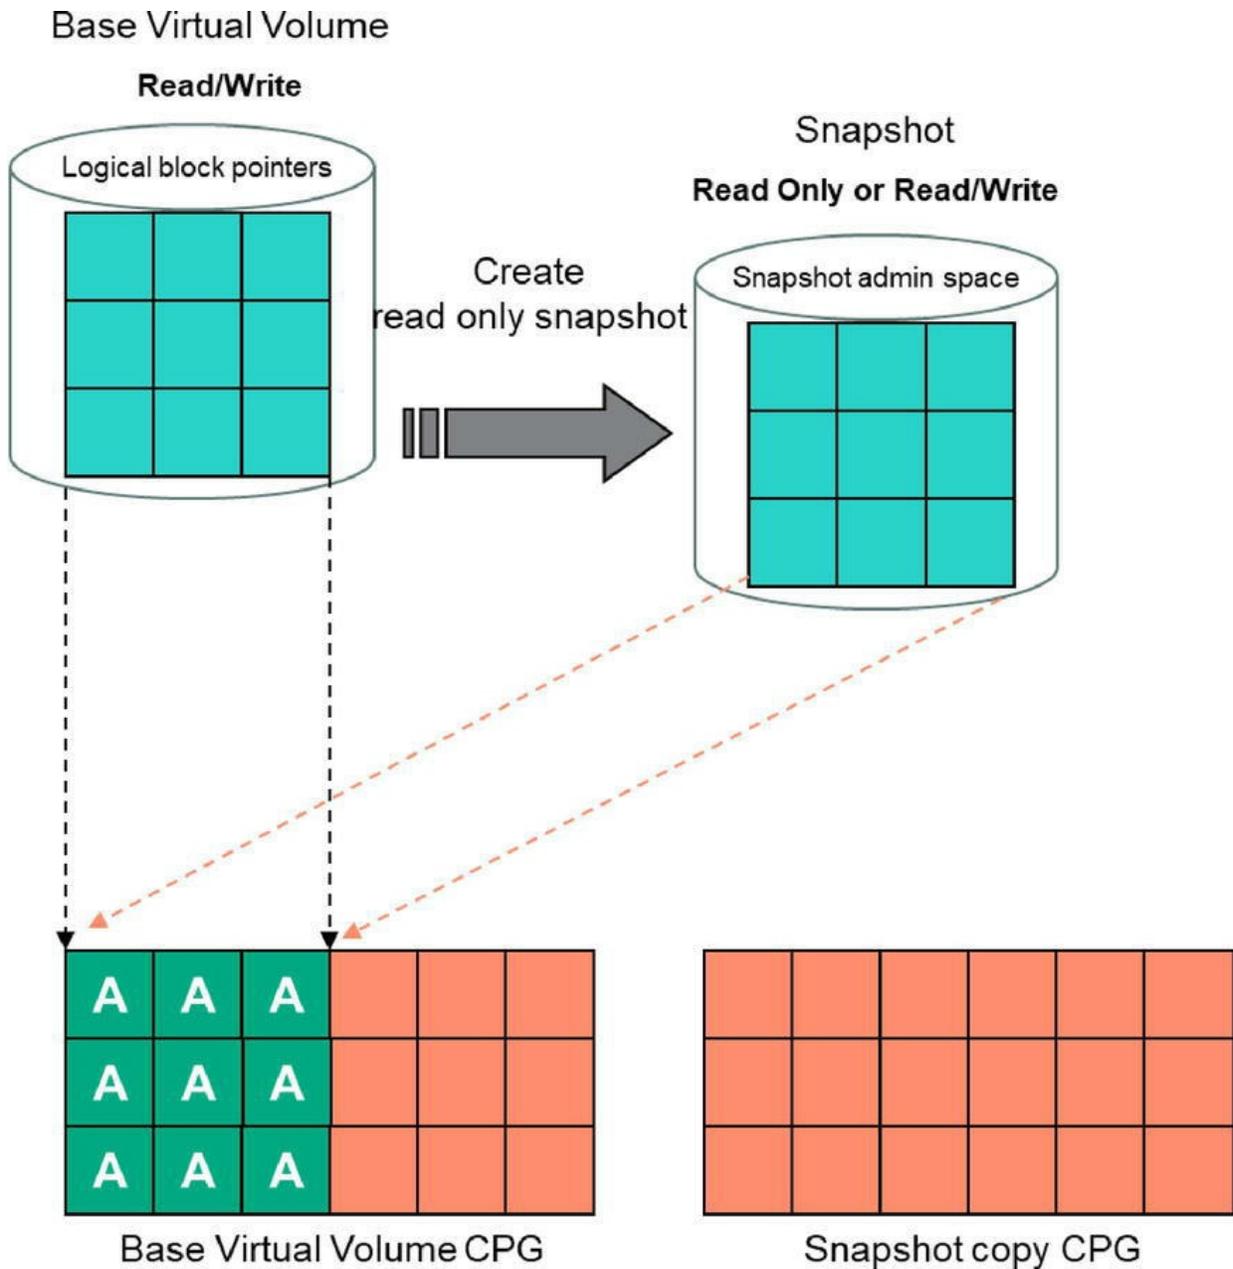


Figure 5-74 The snapshot points to all the original data

All host writes to the virtual logical unit number (VLUN), created from the virtual volume, are logged. The read-only snapshot is created by writing a mapping structure to the parent's Snapshot Administration (SA) logical disk (LD), as indicated in [Figure 5-74](#). This mapping structure points to all the data that resides in the user LD (user area of the parent source volume).

Snapshot: How it works—Update base VV

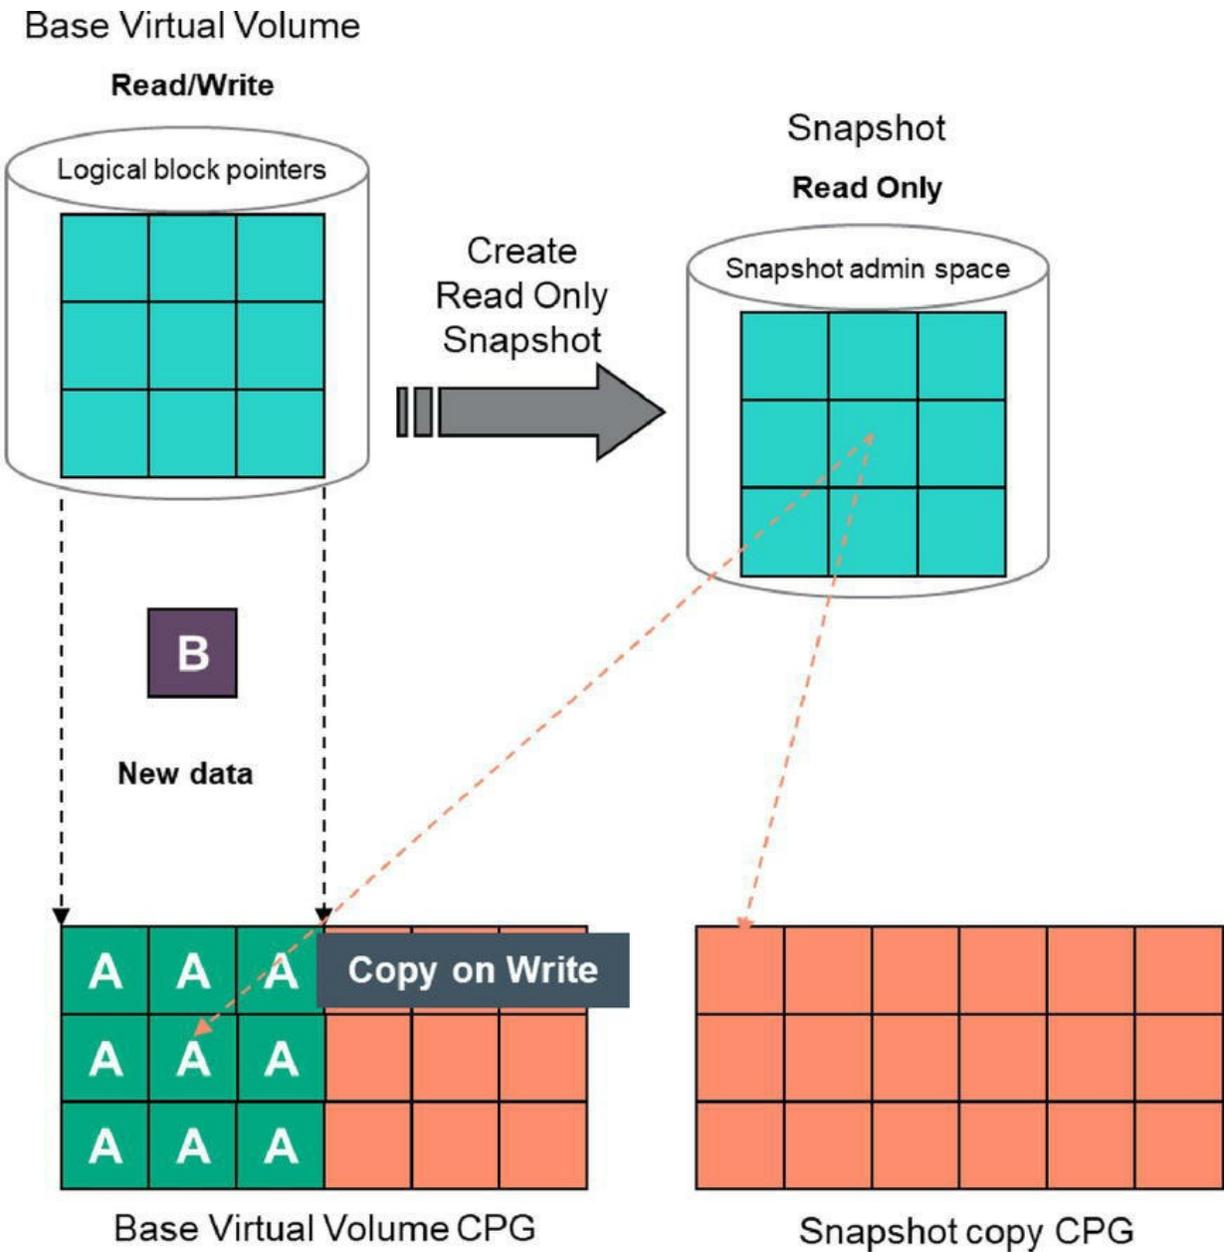


Figure 5-75 Snapshot and the copy on write feature

In the copy-on-write process, as indicated in [Figure 5-75](#), the old data is located on the parent virtual volume (or user LD). This data gets copied to the snapshot data space (SD).

The SA LD's mapping structure of the snapshot is updated with the location of the snapshot data in SD.

The new data is written to the parent virtual volume.

Clone introduction

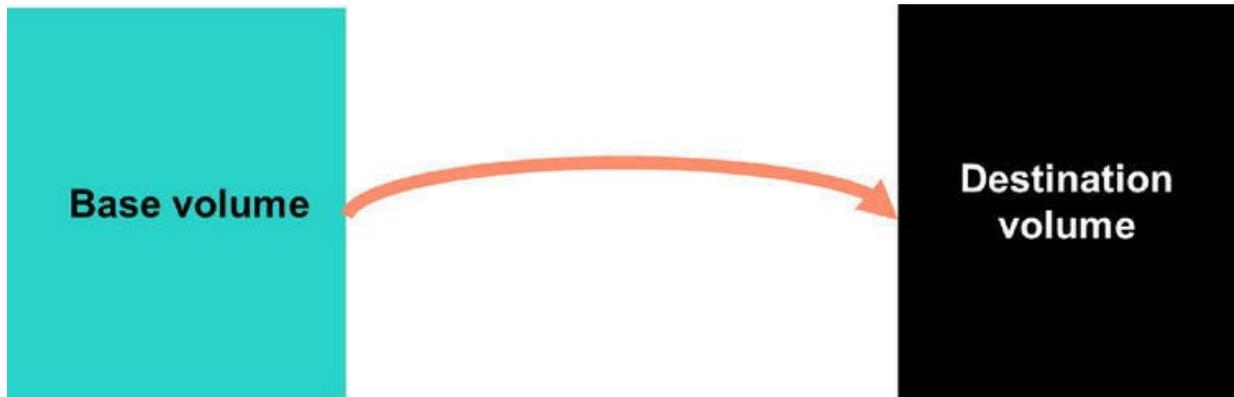


Figure 5-76 Cloning a base volume to a destination volume

A clone is a full copy of a volume. A clone duplicates all the data from one original base volume to another volume called the destination volume, as illustrated in [Figure 5-76](#). Any changes to either volume cause them to lose synchronization with each other, which is corrected by resynchronizing the two volumes. No special license is required to create a clone of a volume.

Clones can be created and managed in groups to reduce the number of management tasks. A consistency group of clones can be created from a list of virtual volumes, and they can be grouped into virtual volume sets that are managed as a single clone.

A clone can only be made from a volume with enough free space to accommodate writes to that volume during the cloning operation. In addition, the destination volume must have at least as much user space as the volume being copied.

Learning check

9. Describe the difference between an incremental and a differential backup.

HPE StoreOnce backup systems



Figure 5-77 HPE StoreOnce backup systems portfolio

The HPE StoreOnce backup system, as shown in [Figure 5-77](#), is a disk-based storage appliance for backing up servers or PCs to target devices on the backup appliance. The total number of backup targets provided by an HPE StoreOnce backup system is split between virtual tape library (VTL), NAS, or Catalyst devices. These devices can be all VTL, all NAS, or any combination of NAS, VTL, and Catalyst devices. All HPE StoreOnce devices that are configured for NAS, VTL, or Catalyst automatically make use of HPE deduplication, ensuring efficient and cost-effective use of disk space. Federated Deduplication is available across the HPE StoreOnce Systems portfolio. StoreOnce backup systems, as shown in [Figure 5-78](#), provide multi-unit or multisite replication for data redundancy.

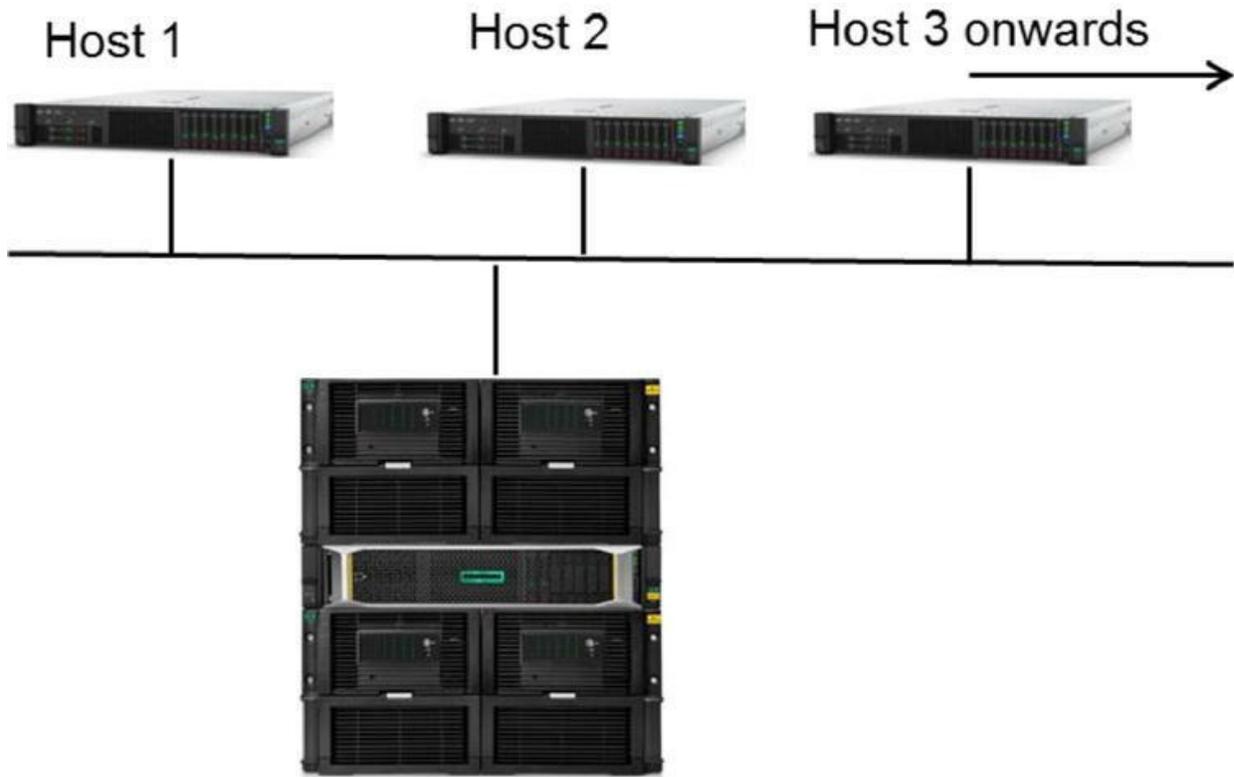


Figure 5-78 HPE StoreOnce a disk-to-disk backup solution

HPE StoreOnce Catalyst

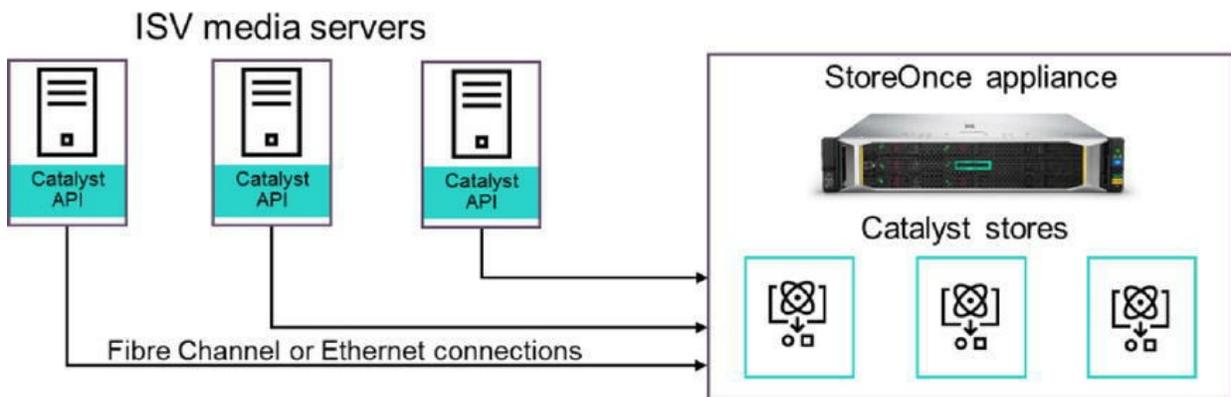


Figure 5-79 HPE StoreOnce Catalyst

A StoreOnce Catalyst, as shown in [Figure 5-79](#), is a device which independent software vendors (ISVs) can use with none of the restrictions of tape emulation and NAS shares. There are no unnecessary “geometry”

limitations such as library slots, cartridge sizes, and so on. The catalyst allows ISV software to be aware of the StoreOnce Backup system and its capabilities. The backup server is able to support node failover on multinode systems without complex restart scripts and restart from checkpoints. It also uses bandwidth-efficient methods to move (copy) data without rehydration. This enables backup over WAN or LAN, with increased performance.

Catalyst enables deduplication in different locations (Federated Deduplication) and offers tight integration with Micro Focus Data Protector, Symantec Netbackup, Backup Exec, RMAN, BridgeHead, Microsoft SQL Native, Veeam, and more. StoreOnce Catalyst enables ISV applications to control copies of backups, copy to multiple locations, and set different expiration dates. The ISV application is fully aware of all copies and jobs and does not require a StoreOnce appliance at the remote site.

If backing up using a Catalyst configuration, backup administrators can choose where the deduplication is performed. The deduplication can be performed on the StoreOnce appliance (target), in the backup server (server), or at the source of the device (client). This can minimize the volume of data being sent across a SAN or LAN and can allow a higher ingest rate at the StoreOnce appliance if the data is deduplicated at the source. Many times during backup time, the application servers have lots of memory and CPU horsepower that is underutilized during backup time. Server- and source-level deduplication can take advantage of that.

In addition, when your customer's backup application does a Catalyst Copy, it knows about both copies of the backup media that holds the same backup data. Therefore, if your customer's primary StoreOnce appliance is unavailable in their backup program, they can select to restore from the media that is on the replicated StoreOnce appliance. The backup application knows about both copies of the data.

Deduplication ratios

Data type	Backup assumptions	Typical deduplication
File systems	Weekly full backup Daily incremental backup	Up to 20:1
VMware	Daily full image backups of VMs	Up to 30:1
Exchange servers	Daily full backup	Up to 10:1
SQL databases	Daily full backup	Variable from 3:1 up to 35:1
Oracle databases	Daily full backup	Variable from 3:1 up to 20:1

Figure 5-80 Typical results by data type

Data duplication is done across all devices (VTL, NAS shares, and Catalyst stores) configured in the appliance. Data deduplication is a method of reducing storage needs by eliminating redundant data so that over time only one unique instance of the data is retained on disk. Deduplication works by examining the data stream as it arrives at the storage appliance, checking for blocks of data that are identical, and eliminating redundant copies.

If duplicate data is found, a pointer is established to the original set of data as opposed to storing the duplicate blocks—removing or “deduplicating” the redundant blocks from the volume. The key concept is that the data deduplication is being done at the block level to remove far more redundant data than deduplication done at the file level (called single instancing), where only duplicate files are removed. Data deduplication is especially powerful when it is applied to backup, since most backup data sets have a great deal of redundancy. The amount of redundancy depends on the type of data being backed up, as illustrated in [Figure 5-80](#), the backup methodology, and the length of time the data is retained.

HPE’s unique StoreOnce deduplication has been specifically designed to provide flexible, low-cost deduplication functionality for smaller and midsize data centers, allowing typically 20x more backup data to be retained on disk for longer periods of time. The actual data deduplication ratio your customer can expect will depend on many factors, including the type of data, the backup methodology used, and the length of time they retain their data. However, assuming standard business data mix and extended on-disk

retention (periods of more than 12 weeks), you could expect to see typically a 20x reduction.

HPE StoreOnce also enables data compression by looking at the incoming stream—the HPE StoreOnce will simply dedup data that is a copy, but if the data is unique (that is, a first time store), the HPE StoreOnce uses Lempel-Ziv (LZ) data compression algorithm to compress the data after the deduplication process and before storing it to disk. Typical compression ratios are between 1.5:1 and 2:1, but may be higher depending on the data type.

Adding data deduplication to disk-based backup delivers several benefits:

- It provides a cost-effective way of keeping the backup data readily available on disk for many weeks or even months.
- More efficient use of disk space effectively reduces the cost-per-gigabyte of storage and the need to purchase more disk capacity.
- By extending data retention periods on disk, the backup data is more accessible for longer periods of time before archiving to tape, making restores much faster and allowing users to return to work more quickly.

Ultimately, data deduplication makes the replication of backup data over low-bandwidth WAN links viable (providing off-site protection for backup data) as only changed data is sent across the connection to a second replicated device.

Since deduplication is applied on a per-device basis, you will not see any space-saving benefits with similar data that is stored in different libraries or shares. Customers can organize their backups so that similar data gets backed up to a single library or share. This will maximize their deduplication efficiency for the data.

This is a major best practice and a key planning step. Customers can create multiple libraries or shares, but put similar data in the same library or share for maximum deduplication within the appliance.

HPE deduplication technology is designed around compatibility and cost for users with smaller IT environments. It offers the following features and benefits:

- StoreOnce uses hash-based chunking technology, providing

deduplication technology at a lower cost. Deduplication operates independently of the backup operation, so avoiding issues with backup application support.

- Configuring deduplication is simple. There is no need to tell the deduplication software which backup application you are using or where file markers are. Deduplication works with any backup application and with any data type written to a StoreOnce configured device.
- When data is read from a VTL, a file from a NAS share, or a Catalyst store, the data is “rehydrated” into its original data stream. Data is not changed.



Note

Deduplication is enabled by the use of a check box being selected (default). Deselect the check box to create a non-deduplication library. Deduplication cannot be enabled or disabled once the library has been created.

Every customer’s data is different and will result in different deduplication ratios.

The large deduplication variation when dealing with databases is due to the different ways storage can be configured and the way transactions affect change rates. For example, multiple mount points can cause a high degree of multiplexing in backups, which reduces deduplication ratios. Also, depending on the transaction nature, large apparent database transaction changes can result in relatively minor database structural changes. This leads to higher deduplication ratios whereas minor record changes of only 1% or less can often result in changes of 5 to 10% of the total data, leading to significantly reduced deduplication ratios.

Therefore, the results shown in the preceding graphic should only be used as a guideline for discussion purposes. Actual results will be different and may be less. HPE strongly recommends completing a Proof of Concept when evaluating deduplication ratios in real user environments. The StoreOnce VSA trial can be used for a Proof of Concept.

Some customers will get more deduplication, and some customers will get

less. Some environments utilize the trial VSA license to test out their backup data to see what type of ratios they may achieve.

What is replication?



Figure 5-81 HPE StoreOnce Replication

After replication has been configured between source and target HPE StoreOnce Backup Systems, as shown in [Figure 5-81](#), it is initiated automatically whenever a cartridge is returned from a drive to a slot that has been configured for replication or when a file is created, modified, or deleted in a NAS share.

Data deduplication ensures that the amount of data to be replicated is minimized, and the impact on network traffic is negligible. However, the benefits of deduplication apply only after the first full backup. The first synchronization of the files (NAS shares) or cartridges (virtual tape devices) will require significant bandwidth. The same applies in reverse, if you need to recover data to a target from a source. Therefore, there is an advantage to being able to accelerate the replication of the first full backup. This process is known as seeding.

Replication-target libraries are not visible to the host unless you configure devices to make them visible. Use with care. If the backup application can see both source and target libraries, the application cannot distinguish between the source library and the target library because the barcodes are duplicated in the two locations.

The default values for the command and data protocol port numbers cannot be changed. If replication needs to take place through a firewall, the network administrator must open (TCP) ports 9387 (command protocol) and 9388 (data protocol).

There is no replication for Catalyst datastores since objects are copied between Catalyst datastores via an ISV backup application that supports Catalyst, so the copy can be scheduled and is known by the backup application.

Catalyst Copy is the equivalent of a virtual library and a NAS share replication. The same principles apply in that only the new data created at the source site needs to be copied (replicated) to the target site. The fundamental difference is that the copy jobs are created by the backup application and can, therefore, be tracked and monitored within the backup application catalog as well as from the StoreOnce Management GUI.

Catalyst Copy does not require a replication license where the target of replication requires a replication license for NAS and VTL replication.



Note

Replication is only for Virtual Tape Libraries and NAS Shares.

HPE StoreOnce Replication allows cartridges in slots on a library device that has been configured on one HPE StoreOnce Backup System (the source appliance) to be replicated to corresponding mapped slots on a library device on another HPE StoreOnce Backup System (the target appliance) at another location. Data is always mapped at slot level.



Note

If the replication source appliance fails and goes offline while it was performing a NAS replication job, the replication target share appliance may be left in an inconsistent state. This state may result in failure to recover from some files in the share or perform a list of files in the share because relationships between files and folders have been lost. This will automatically be repaired by the target appliance as soon as it detects the loss of the source appliance; however, it may take up to 10 minutes to complete this repair during which time the share will not be available for access.

A replication license is required for both HPE StoreOnce Backup Systems.

There are four main usage models for replication using StoreOnce VTL and NAS devices:

- **Active/Passive**—A StoreOnce system at an alternate site is dedicated solely as a target for replication from a StoreOnce system at a primary location.
- **Active/Active**—Both StoreOnce systems are backing up local data as well as receiving replicated data from each other.
- **Many-to-One**—A target StoreOnce system at a data center is receiving replicated data from many other StoreOnce systems at other locations.
- **N-Way**—A collection of StoreOnce systems on several sites are acting as replication targets for other sites.

Appliance-to-appliance connectivity via optimized StoreOnce Catalyst Copy jobs is supported over both Ethernet and Fibre Channel.

HPE StoreOnce Catalyst allows the ISV backup application to initiate and control copy of backups from one Catalyst store to another. In this way, the ISV catalog or internal database is fully aware of **all** copies of data.

The control or copy is a control function of the Catalyst API. In the case of non-HPE software such as Symantec NetBackup and Backup Exec, the OpenStorage Technology (OST) API communicates with the HPE API. The copy is called optimized duplication.

All Catalyst Copies are performed using low bandwidth. (It is not necessary to configure any “mapping” as in VTL or NAS.) A catalyst can be copied directly to another catalyst without the need for the original data to be rehydrated.

All copies are controlled by an ISV application. Data Protector will have a copy as part of the “Object Copy” function. Symantec NetBackup uses Storage Lifecycle Policy to set IP duplication.

StoreOnce Replication licenses are not needed for Catalyst devices since replication is done by the ISV backup application.

Bidirectional replication

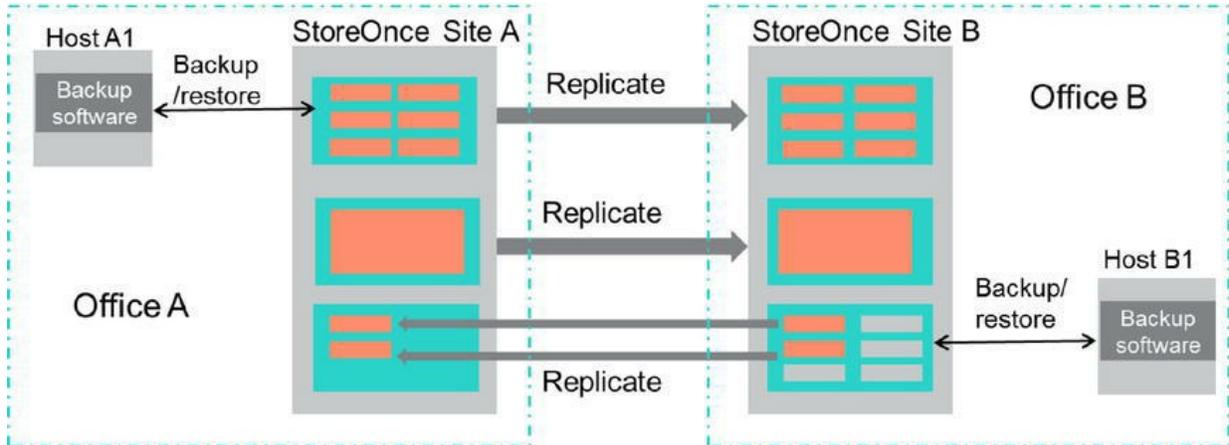


Figure 5-82 Bidirectional replication

HPE StoreOnce Replication allows data on one HPE StoreOnce Backup System (the source, local appliance) to be replicated to another HPE StoreOnce Backup System (the target, remote appliance) at another location. At the same time, a remote appliance can be configured to replicate shares or libraries to a local appliance. This provides local backup and remote replication for both sites, as illustrated in [Figure 5-82](#).



Note

If the replication source appliance fails and goes offline while it was performing a NAS replication job, the replication target share appliance may be left in an inconsistent state. This state may result in failure to recover from some files in the share or perform a list of files in the share because relationships between files and folders have been lost. This will automatically be repaired by the target appliance as soon as it detects the loss of the source appliance; however, it may take up to 10 minutes to complete this repair during which time the share will not be available for access.

A replication license is required for both HPE StoreOnce Backup Systems.

StoreOnce VSA

HPE StoreOnce VSA is a virtual appliance that delivers fast, efficient,

scalable, deduplicating backup data storage. All the features of the purpose-built StoreOnce systems are available in a software-defined backup target of up to 500 TB usable capacity.

StoreOnce VSA starts with a 4 TB license to use. This is the minimum capacity configuration and is extended in 1 TB increments using stackable 1 TB licenses. The 1 TB licenses can be added individually or in bulk. They can be added as part of the initial deployment or at any time after deployment. Up to 496 1 TB licenses can be added to give a maximum capacity of 500 TB.

For extended evaluation or non-critical production deployment there is a 1 TB freeware version from www.hpe.com/Storage/FreeBackup. It can be extended by the purchase of additional capacity up to 500 TB.

The purchased licenses include three years entitlement to HPE Pointnext technical support and software updates.

StoreOnce VSA features include:

- Protocol support: StoreOnce Catalyst, Common Internet File System (CIFS), VTL iSCSI
- Bulk deployment capability for large volume of remote offices
- Federated Deduplication and copy/replication compatible with StoreOnce hardware appliances
- Sixty-day instant on trial license

Applications supported include:

- VMware, Hyper-V, or Kernel-based Virtual Machine (KVM)
- Reporting of capacity utilization, deduplication ratio, and so on through StoreOnce Enterprise Manager
- Micro Focus Data Protector, Symantec NetBackup, BackupExec (VTL iSCSI only), Veeam

To apply an Instant On license use StoreOnce CLI commands as follows:

- # license show to show the current license status
- # license add demo to apply the 90-day license for all licensable features

StoreOnce 3640



Figure 5-83 StoreOnce 3640 Base and Expansion units

HPE StoreOnce 3640, as shown in [Figure 5-83](#), delivers scalable backup and restore for small to midsize data centers, and provides an ideal replication target device for up to 36 remote and branch offices. This scalable 2U appliance offers from 36 to 108 TB of usable local capacity using upgrade kits and up to 324 TB with Cloud Bank Storage, and it provides a solution to shrinking backup windows with speeds of up to 18 TB/hour using HPE StoreOnce Catalyst.



Note

Actual performance depends on multiple factors including configuration, data type, data deduplication, data compression, number of data streams, number of backup targets, and concurrent tasks such as housekeeping or replication. Actual capacity depends on multiple factors including data deduplication, data compression, storage formatting, log file size, metadata size, and housekeeping backlog.

StoreOnce 3640 (minimum configuration 36 TB usable) features include:

- Based on HPE ProLiant Gen10 platform
- Two Intel® Xeon® 4110 Skylake CPUs (eight core)
- Twelve 4 TB LFF disks
- 2U form factor, 6U with max expansion
- Four 1 GbE embedded network adapters

- iLO 5
- Optional 10 GbaseT, 10/25 GbE, 16/32 Gb Fibre Channel
- HPE p1224 SAS RAID controller
- HPE P408i Smart Array (for boot disks)
- Two 600 GB SFF operating system disks (at rear)
- SAS expander
- Dual redundant power supplies
- Fully expanded = 108 TB of capacity

StoreOnce 3620



Figure 5-84 StoreOnce 3620 and six x 4 TB expansions disks

HPE StoreOnce 3620, as shown in [Figure 5-84](#), delivers entry-level, disk-based backup and disaster recovery, that is ideal for smaller remote or branch offices and data centers. This 2U Backup System offers 16 to 31.5 TB of usable local capacity, up to 94.5 TB with Cloud Bank Storage, and speeds of up to 14 TB/hour with StoreOnce Catalyst.

Note

Actual performance depends on multiple factors including configuration, data type, data deduplication, data compression, number of data streams, number of backup targets, and concurrent

tasks such as housekeeping or replication. Actual capacity depends on multiple factors including data deduplication, data compression, storage formatting, log file size, metadata size, and housekeeping backlog.

StoreOnce 3620 (minimum configuration 16 TB) features include:

- Based on HPE ProLiant Gen10 hardware
- Two Intel® Xeon® 4110 Skylake CPUs (8 core)
- Six 4 TB LFF disks
- 2U form factor based on ProLiant DL380
- Four 1 GbE embedded network adapters
- iLO 5
- Optional 10 GbaseT, 10/25 GbE, 16/32 Gb Fibre Channel
- HPE p1224 SAS RAID controller
- HPE P408i Smart Array (for boot disks)
- Two 600 GB SFF operating system disks (at rear)
- SAS expander
- Dual redundant power supplies
- Optional six 4 TB disk expansion kit (max. useable 31.5 TB)



Note

Expansion shelf just a bunch of disks (JBOD) is not supported on this model. No hot spare on this model.

StoreOnce 5200

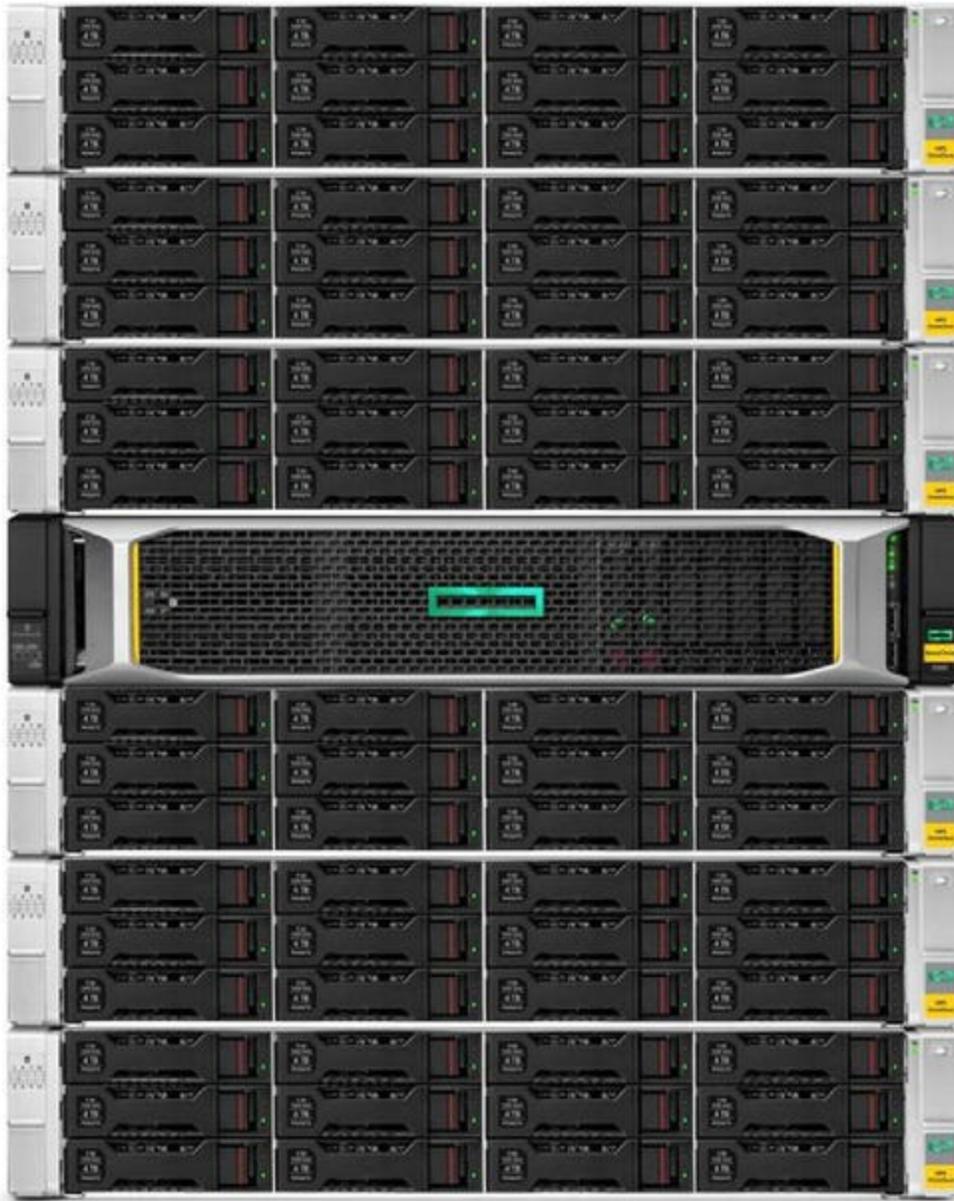


Figure 5-85 StoreOnce 5200/BB956A

HPE StoreOnce 5200, as shown in [Figure 5-85](#), delivers scalable backup and restore for small to mid-sized data centers, and provides an ideal replication target device for up to 64 remote and branch offices. This scalable 4U appliance offers from 36 to 216 TB of usable local capacity using upgrade kits and up to 648 TB with Cloud Bank Storage, and it provides a solution to shrinking backup windows with speeds of up to 33 TB/hour using HPE StoreOnce Catalyst.



Note

Actual performance depends on multiple factors including configuration, data type, data deduplication, data compression, number of data streams, number of backup targets, and concurrent tasks such as housekeeping or replication. Actual capacity depends on multiple factors including data deduplication, data compression, storage formatting, log file size, metadata size, and housekeeping backlog.

HPE StoreOnce 5200 is based on ProLiant DL380 LFF Disk Gen10 Server and includes:

- Two Intel[®] Xeon[®] 5118 Skylake CPUs (12 cores)
- Maximum write performance: 33 TB/hour (Catalyst)
- Maximum number of concurrent data streams: 512
- Maximum number of data stores: 64
- Maximum fan-in or backup targets: 32
- Redundant 800 W PSUs standard
- Dedicated, redundant operating system drives (front drive bays)
- Flexible IO configurations
 - Supports one to four IO HBAs in any combination
 - Four types available:
 - 10 base-T
 - 10 Gb or 25 Gb Ethernet
 - 16 Gb Fibre Channel (Note: 8 Gb/s Fibre Channel HBA discontinued)
 - 32 Gb Fibre Channel
- Storage
 - Minimum configuration: 36 TB

- Maximum configuration (with expansion): 216 TB
- Attain maximum storage capacity by adding up to five disk enclosures (BB964)
- Shelf data disk RAID 6 (9+2) plus hot spare
- Cloud Bank Storage (licensed) maximum capacity: 432 TB

StoreOnce 5250



Figure 5-86 StoreOnce 5250/BB958A

HPE StoreOnce 5250, as shown in [Figure 5-86](#), offers disk-based backup with deduplication for longer term on-site data retention and off-site disaster recovery with best-in-class scalability and performance for larger midsize and enterprise data centers. It also provides an ideal replication target for up to 64 remote or branch offices. This highly scalable 7U to 12U appliance delivers from 36 TB to 864 TB of usable local capacity up to 1.7 PB of Cloud Bank Storage capacity and easily tackles shrinking backup windows with backup speeds up to 41 TB/hour using HPE StoreOnce Catalyst.



Note

Actual performance depends on multiple factors including configuration, data type, data deduplication, data compression, number of data streams, number of backup targets, and concurrent tasks such as housekeeping or replication. Actual capacity depends on multiple factors including data deduplication, data compression, storage formatting, log file size, metadata size, and housekeeping backlog.

HPE StoreOnce 5200 is based on ProLiant DL380 LFF Disk Gen10 Server and includes:

- Two Intel[®] Xeon[®] 5118 Skylake CPUs (12 cores)
- Maximum write performance: 33 TB/hour (Catalyst)
- Maximum number of concurrent data streams: 512
- Maximum number of data stores: 64
- Maximum fan-in or backup targets: 32
- Redundant 800 W PSUs standard
- Dedicated, redundant operating system drives (front drive bays)
- Flexible IO configurations
 - Supports one to four IO HBAs in any combination
 - Four types available:
 - 10 base-T

- 10 Gb or 25 Gb Ethernet
- 16 Gb Fibre Channel (Note: 8 Gb/s Fibre Channel HBA discontinued)
- 32 Gb Fibre Channel
- Storage
 - Minimum configuration: 36 TB
 - Maximum configuration (with expansion): 864 TB
 - Attain maximum storage capacity by adding up to one additional (fully populated) storage enclosure (BB976A)
 - Data drive RAID 6 (9+2) plus two roaming hot spares (each drawer)
 - Cloud Bank Storage (licensed) maximum capacity: 1728 TB

What is Cloud Bank Storage?



Figure 5-87 Cloud Bank Storage

StoreOnce Cloud Bank, as shown in [Figure 5-87](#), leverages object storage for long-term retention of backup data copies. Features include:

- Enables StoreOnce to use public and private cloud services for long-term retention and disaster recovery.
- Combines the efficiency of StoreOnce deduplication with the cost benefits of cloud storage.
- Is designed to reduce cloud storage and access costs.
- Offers single management point for copy control by using the existing backup application.
- Supports public and private cloud service providers including AWS S3, Microsoft Azure, and Scality.

- Enables recovery from cloud to any StoreOnce system.
- Scales up to 2x the capacity of the associated StoreOnce system.

HPE StoreOnce Cloud Bank Storage use cases

Use Case 1—Long-term backup retention:

- Wants to retain data for a long time as cheaply as possible
- Plans to seldom read the data, if ever

Use Case 2—Offsite DR protection:

- Wants to regularly send backups offsite for the purposes of DR protection
- No secondary site to replicate to
- Only plans to read the data for DR

Use Case 3—StoreOnce VSA in the cloud backed by Object Storage:

- Using StoreOnce VSA as a backup/replication target within Cloud Compute
- Would like to reduce running costs by hosting data on cheaper object storage

Cloud Bank Storage setup

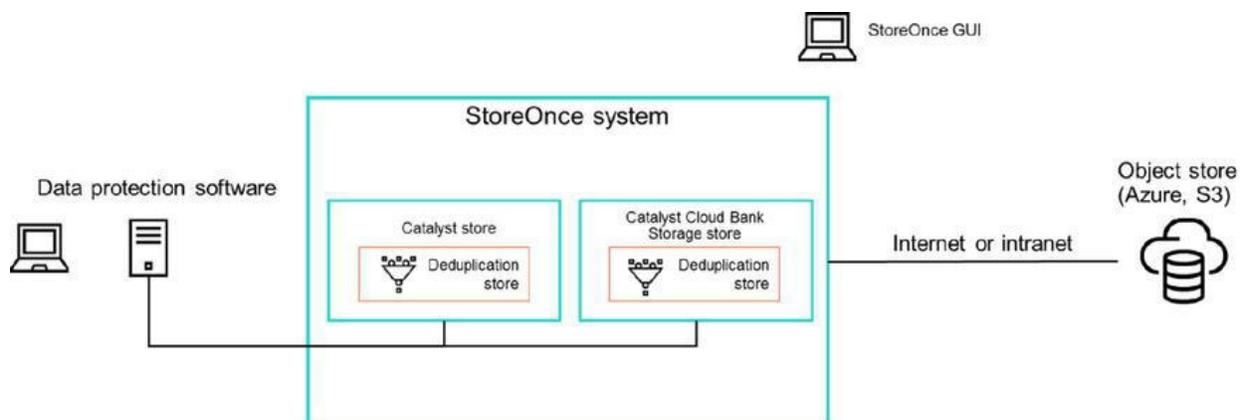


Figure 5-88 Cloud Bank Storage setup

Cloud Bank Storage setup, as illustrated in [Figure 5-88](#):

- Create a Catalyst store by using the StoreOnce GUI.
- Create a Cloud Bank Storage store, and enter credentials of the Azure or S3 object store.
- Configure the data protection software for both stores.
- Configure a backup job to the Catalyst store, and a copy job (object copy and SLP) to the Cloud Bank Storage store.

Cloud Bank Storage—Backup and copy

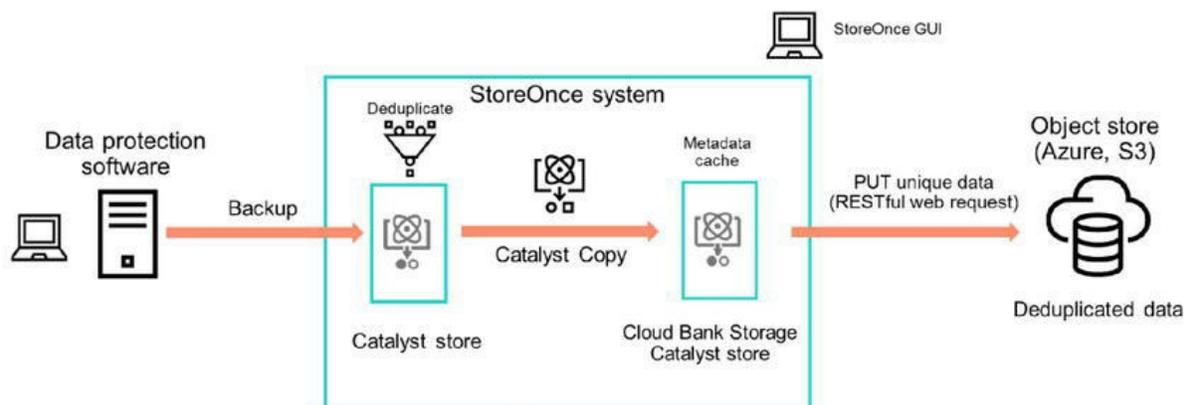


Figure 5-89 Cloud Bank Storage—Backup and copy

Cloud Bank Storage Stores, as illustrated in [Figure 5-89](#), are designed for use as Catalyst Copy targets. In this use case, the customer directs backups to a regular Catalyst Store, then for long-term backup data retention a copy job is run to copy the backup data to the Cloud Bank Storage Store.

Only unique data is sent to the cloud and data lands in the cloud already deduplicated. This can reduce storage capacity requirement by up to 20x.

The Cloud Bank Store can be on the same system as the regular Catalyst Store or on a different system. For most use cases, it makes sense to have the Cloud Bank Storage store on the same system. Where multiple StoreOnce systems are in use and only a subset of the backups need to be copied to the external object storage, these could all copy to an HPE Cloud Bank Storage on a single StoreOnce system. This would enable the setup of just one Cloud

Bank Storage store and deduplication across all the backup copies.

Cloud Bank Storage—Enhanced restore

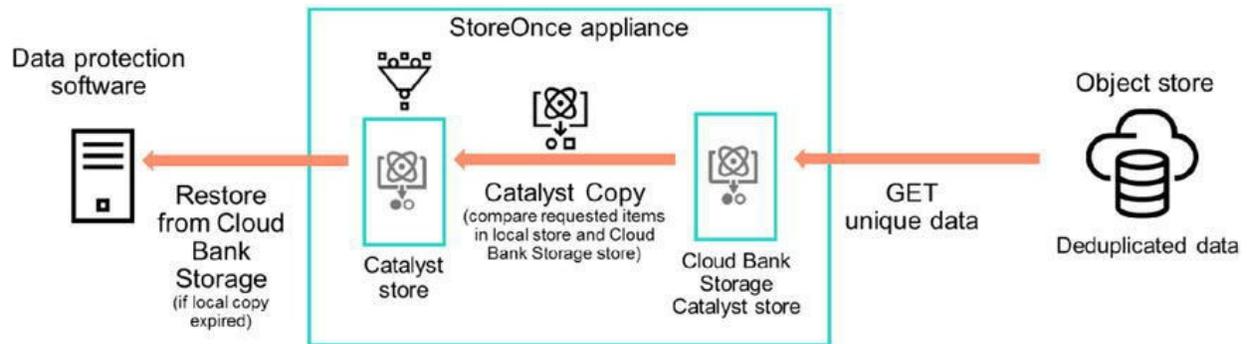


Figure 5-90 Cloud Bank Storage—Enhanced restore

The restore process with Cloud Bank Storage, as indicated in [Figure 5-90](#):

- Look for a local copy on the StoreOnce System for restoring.
- If the local copy is expired, then the user needs to restore from Cloud Bank Storage.
- Compare the requested items from the local store and the Cloud Bank Storage store.
- The final step involves restoring the unique data from the object store and completing the restoration process.

Learning check

10. Match the HPE StoreOnce model with its description.

StoreOnce VSA	6 LFF disks with optional expansion kit
StoreOnce 3640	Scalable 4U appliance with five optional disk enclosures
StoreOnce 3620	12 LFF disks
StoreOnce 5200	Virtual appliance with up to 500 TB

Configuration tools

HPE offers a variety of tools to help with the storage sizing process. HPE also 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.

HPE Storage Sizer

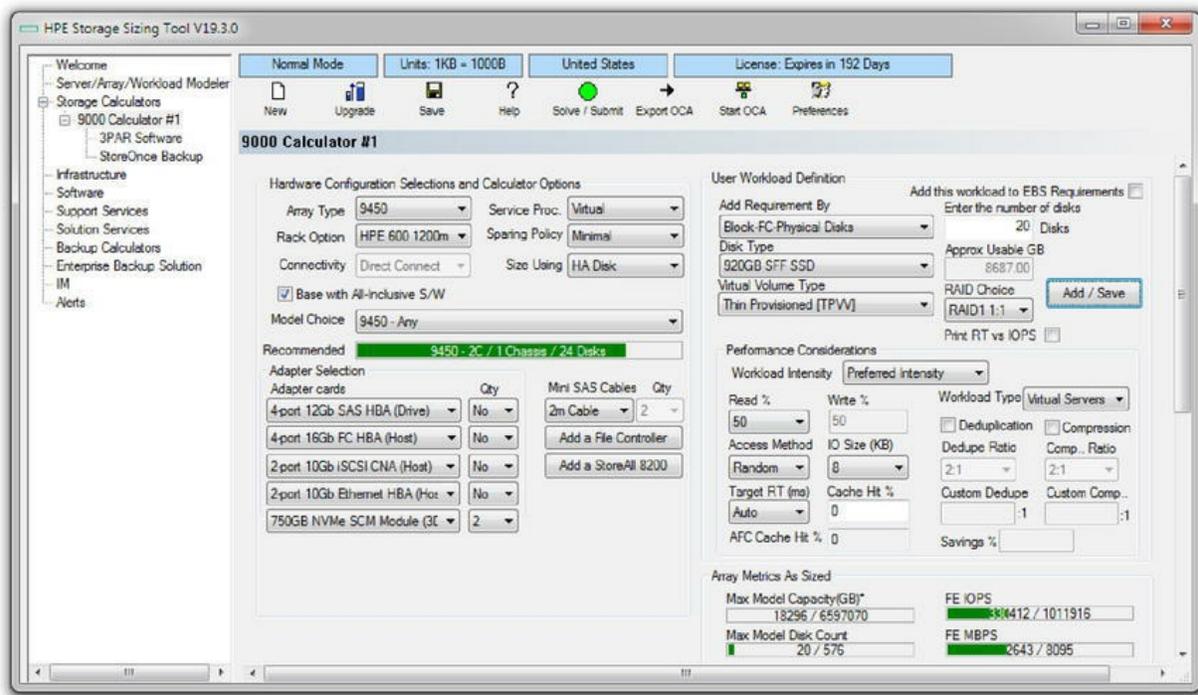


Figure 5-91 HPE Storage Sizer

The HPE Storage Sizer, as shown in [Figure 5-91](#), is a sizing tool that helps you design a storage infrastructure to meet the needs of a customer. 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, 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 One Configuration Advanced (OCA) for a quotation.
- Provides localized parts and pricing for different geographic regions.
- Includes HPE Smart Update Technology, which brings new products or functionality 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 functionalities, which were added based on user input, annual surveys, and quarterly focus groups.

HPE NinjaSTARS for Nimble/StoreOnce/3PAR/Primera

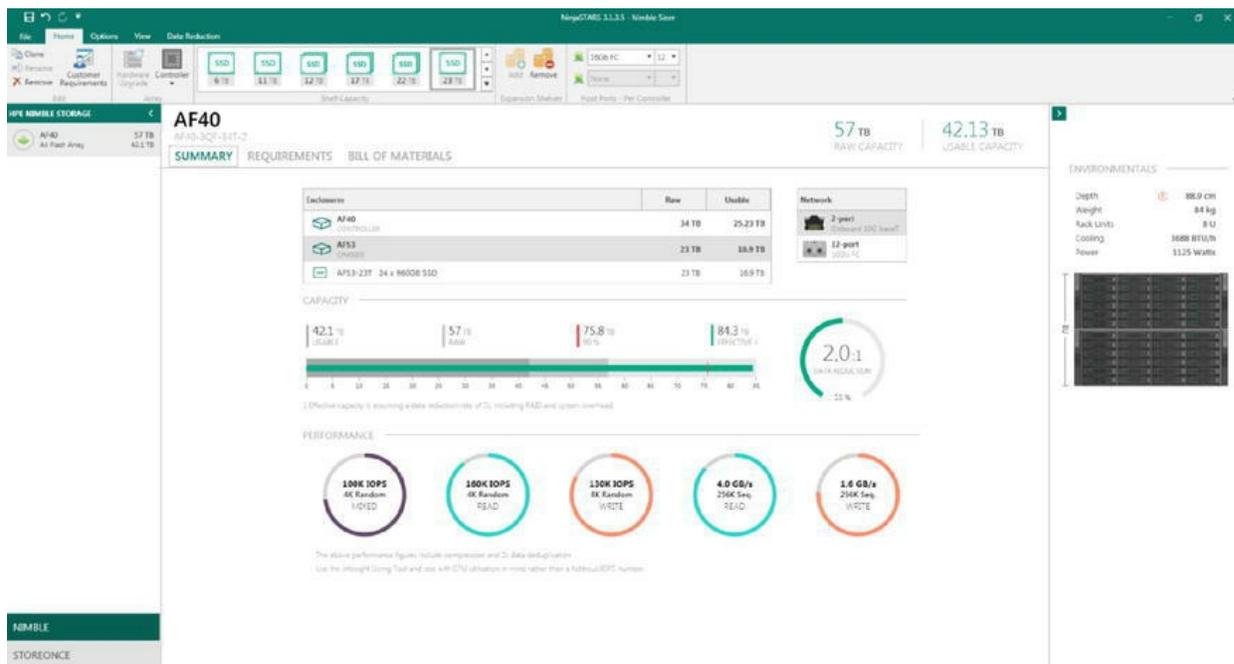


Figure 5-92 HPE NinjaSTARS for Nimble/StoreOnce/3PAR

HPE NinjaSTARS Tool, as shown in [Figure 5-92](#), is an interactive storage sizing tool for HPE 3PAR StoreServ Storage, HPE Nimble Storage arrays, HPE Primera, and HPE StoreOnce systems. It can size and configure a cost-effective and performance-optimized HPE 3PAR StoreServ or HPE Nimble Storage that can replace existing legacy arrays. Based on a set of basic capacity and performance metrics (requirements) from an existing storage environment, it converts these requirements into a ready-to-order bill-of-materials (BOM) list.

The tool has different ways of accepting the sizing requirement, including using a set of standard templates, inputting a desired usable capacity number

(with percentages per tier), or completely customizing the array. Users can interactively change the requirements to see the change in capacity and performance characteristics of the target storage system. NinjaSTARS also provides options for licensing and software components.

NinjaSTARS uses a flow-based approach in the following sequence:

- Gather configuration and sizing requirements.
- Interactively customize the target storage system.
- Select licensing and software components for final configuration.
- Generate a BOM for the selected configuration.

Installing NinjaSTARS

NinjaSTARS is an agentless, standalone tool, designed to run on any Windows machine, including a laptop, workstation, or server. It is compiled as a portable application, designed to run from a USB stick with no installation required. Scanning of the target systems does not examine any individual file, require a host agent, or leave behind any information. Network traffic generated is negligible.

Learning check

11. Which storage platforms can be configured with NinjaSTARS?

Management tools

HPE network management tools span across products and bring innovations

to complete lifecycle management. As network management becomes more complex, the risks associated with compromised data flow have also increased.

HPE 3PAR StoreServ Management Console

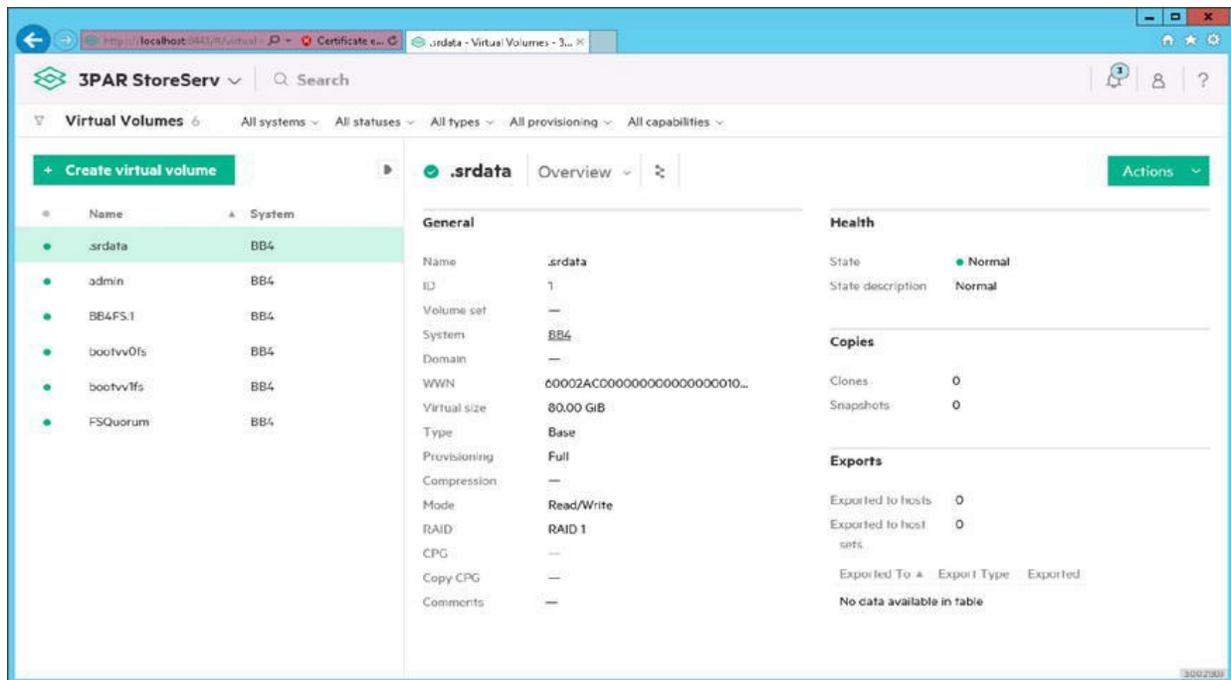


Figure 5-93 HPE 3PAR StoreServ Management Console

HPE 3PAR StoreServ Management Console (SSMC), as shown in [Figure 5-93](#), offers a modern look, a consistent feel, and a common interface and language with HPE Management tools such as HPE OneView. Designed to use the latest API and UI technologies, HPE 3PAR SSMC centralizes all HPE 3PAR StoreServ Management under a single pane of glass and offers converged management and reporting for both file and block.

Accessible through a web-based UI, you can run your operations as HPE 3PAR SSMC tasks with an improved user experience and better responsiveness. All the information you need from HPE 3PAR StoreServ is available at a glance with customizable reporting capabilities, eliminating the need for add-on software tools or diagnostics and troubleshooting.

3PAR Service Console

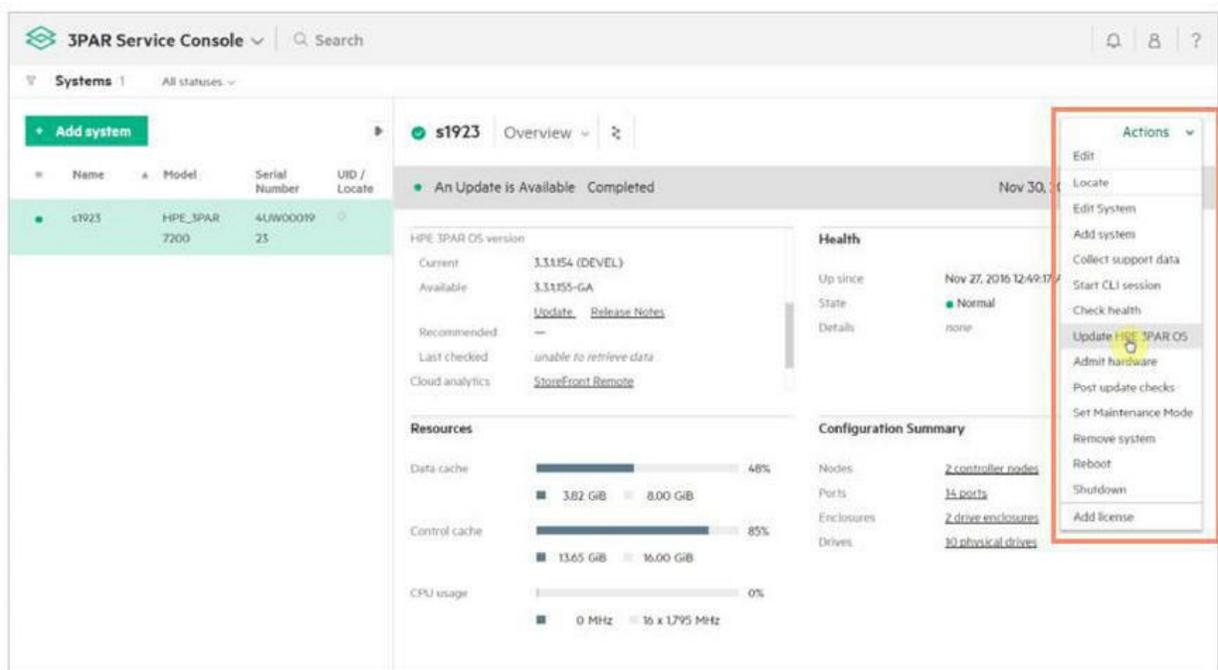


Figure 5-94 Checking for updates on the 3PAR Service Console

The 3PAR Service Console (SC), as shown in [Figure 5-94](#), also allows service functions to be performed by a company admin, HPE support, or an authorized service provider. SC replaces Service Processor Onsite Customer Care (SPOCC), the GUI for Service Processor (SP) versions 4.4 and earlier. SC functionality is similar to SPOCC.

The HPE 3PAR SC is an appliance that collects data from an attached 3PAR StoreServ Storage system in predefined intervals as well as an on-demand basis and sends the data to the call home infrastructure, if configured. The SP is the platform on which the SC runs. It sends support data back to HPE and provides a way for HPE Technical Support Engineers to remotely log in to resolve problems. The SP is also needed for HPE InfoSight support. The SC is the GUI for the SP. It provides a streamlined, more usable interface with a layout that closely resembles the HPE 3PAR SSMC.

From Actions, you can check updates for your System OS or the SP.

- **Main menu and banner**—Clicking the menu area opens the main menu of screens. Selecting a screen from the menu displays a summary of

configuration settings for the selected screen and provides Actions menus. Depending on the screen selected, system resources such as CPU usage, throughput, physical memory, and so on might be displayed.

- **Screens**—The Screens area displays tabular and graphical information and provides action dialogs for managing storage systems.
- **Detail pane**—This displays the detailed information that can be seen and allows you to select further views.
- **List pane**—When you select an item in the List pane, additional information about it is displayed in the Detail pane. When an item is selected, you can perform actions on it. Many lists can be sorted and filtered and include multiple views that you can select.
- **Actions**—The Actions menu allows you to perform actions on one or more resources that you have selected, in the list pane. If you do not have permission to perform an action, the action does not display in the menu. Also, some actions might not display because of system configurations, user roles, or properties of the selected resource.

Remote benefits of the SP include:

- **Remote diagnostics**—Key diagnostic information is maintained centrally on a historical basis. System health statistics, configuration data, performance data, and system events can be transferred frequently and maintained centrally. As a result, proactive fault detection and analysis are maximized, and manual intervention is minimized.
- **Remote serviceability**—SP provides fast, predictive response and remediation. HPE 3PAR technical support delivers rapid responses with 24x7 remote monitoring and analysis to identify issues and proactively communicate them back to the customer. As an integrated support model, HPE 3PAR technical support can remotely connect to a customer's HPE 3PAR StoreServ system through a secure IP connection to resolve issues quickly and reduce on-site visits.
- **Remote online software upgrade**—Upgrade software with no application disruption. Changes to the HPE 3PAR OS software are released for new functionalities, maintenance updates, and software patches. The ability to apply these updates can be serviced as an online upgrade, where the arrays' capability to process customer data does not

need to be disrupted during the software upgrade.

Additional features of the HPE 3PAR SP include the following:

- Collects periodic data from HPE 3PAR StoreServ (alerts, configuration, events, performance, status, and so on).
- Performs hourly health checks on the HPE 3PAR StoreServ system.
- Provides remote support capabilities over Ethernet.
- Serves as a local maintenance terminal for on-site support.
- Serves as a remote maintenance terminal for remote support.
- Contains guided maintenance scripts used for guided parts replacement activities.

NimbleOS WebUI

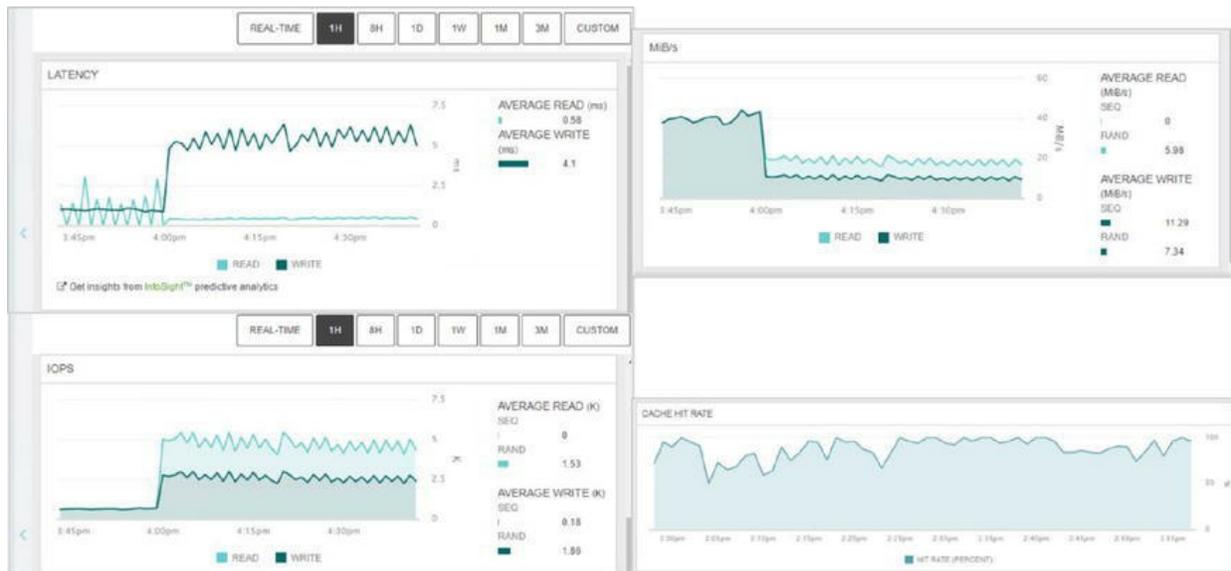


Figure 5-95 NimbleOS WebUI

The Nimble Array provides a simple-to-use, intuitive WebUI from which you manage the array. To access the WebUI, open any supported browser, and enter the array's management IP address. You are presented with the login page. Enter the password that you set during the array creation and log in.

The Home page gives you system-wide information at a glance. It lets you track performance, capacity, traffic load, and events for the entire system, as indicated in [Figure 5-95](#). Other sections include:

- **Performance**—Data throughput and IOPS are the pulse of the array. Changes in either of these may mean that system use has changed or that there is a problem somewhere in the array or from an application. The center section of the Home page gives you an at-a-glance status of data movement activity.
- **Capacity**—The Space section lets you see the overall space usage. Immediately beneath the disk space is the space savings panel, which displays the compression factor, aggregated for the entire array. By understanding the total space use, you can manage the overall array more efficiently. Individual space usage appears on the Volumes details page for each volume. The details on the efficiency of the inline compression engine lets you track actual space usage. This gives you better information when tracking usage trends and making decisions about future growth.



Note

If the volume displayed on the Volumes detail page is a replica, the WebUI shows the information relevant to the replication space.

- **Protection and Hardware**—Another key data point administrators want at a glance is the ability to quickly see that their data is protected and that the hardware is healthy. This section provides a summary of both those data points.
- **Events**—The Event Summary and Recent Events sections, displayed on the right side of the page, shows you the system health. You can then move to the Events details page to view event details.

StoreOnce Management Console

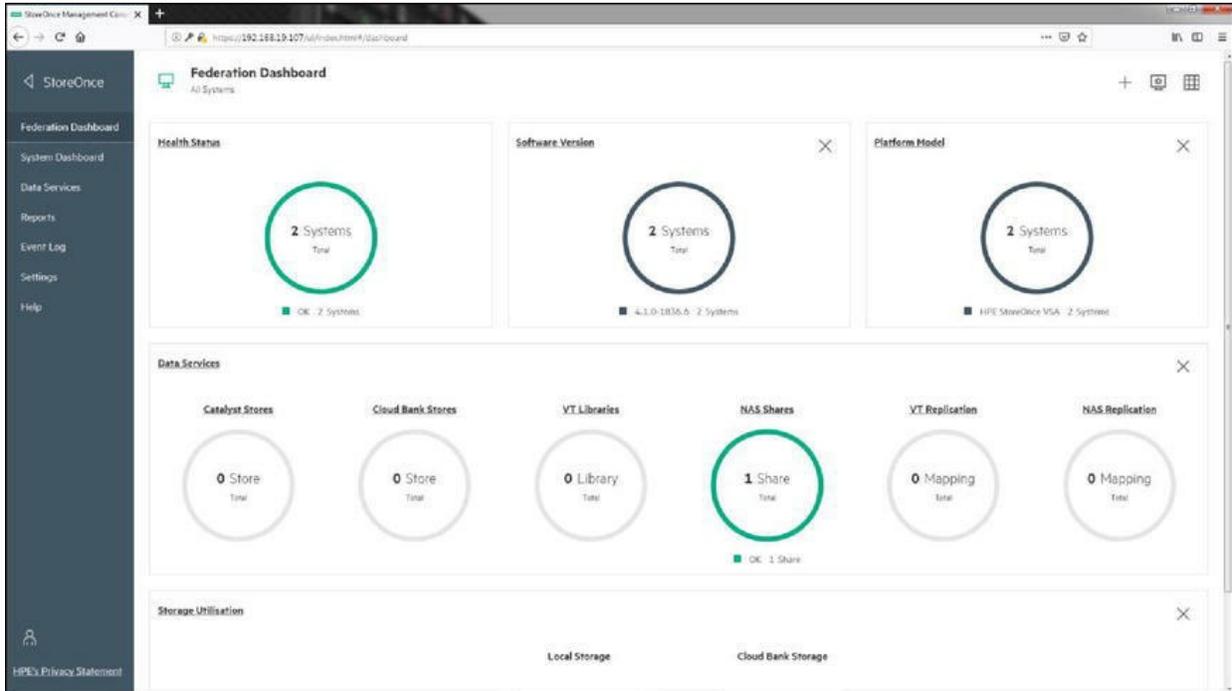


Figure 5-96 StoreOnce Management Console

The StoreOnce Management Console, as shown in [Figure 5-96](#), is the primary interface to configure and manage all backup targets and Catalyst Copy functions. To access the console, browse to the IP address or Fully Qualified Domain Name (FQDN) of the StoreOnce System.

Management federation



Figure 5-97 Federation Dashboard

Management federation allows administrators to manage multiple StoreOnce systems from a single appliance. The managing system in a federation is called the lead system, and the other systems in the federation are called member systems. When logged into a lead system, administrators can manage not only that system but also any of the member systems in the federation. For example, from the lead system you can create StoreOnce Catalyst Stores on any member systems in the federation. The Federation Dashboard screen on a lead system displays aggregated information. When logged into a member system, you can manage the system as usual; however, you cannot manage other systems in the federation from a member system.

The Federation Dashboard on a member system indicates that the system is part of a federation. Multiple federations can be configured where StoreOnce systems can be leads, members, or both within overlapping federation domains. The Federation Dashboard, as shown in [Figure 5-97](#), provides a single view for management and reporting for up to 20 VSA and physical systems (up to 100 by request).

HPE StoreOnce Recovery Manager Central

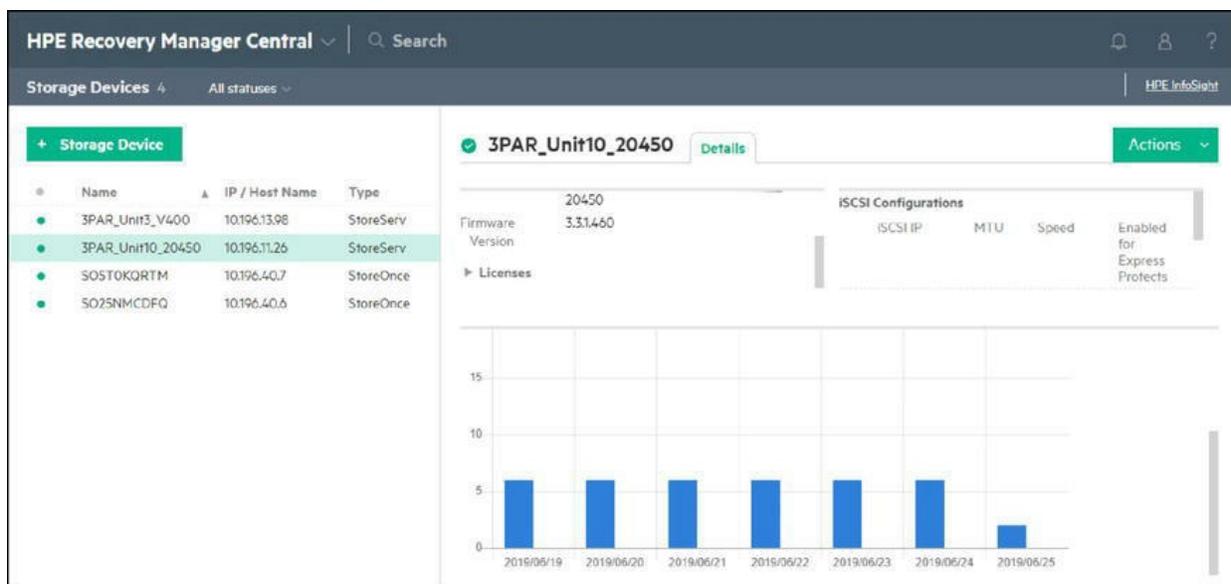


Figure 5-98 HPE StoreOnce Recovery Manager Central

The Recovery Manager Central (RMC) GUI, as shown in [Figure 5-98](#), is used to manage and monitor tasks in RMC appliance (VM deployed from the template).

Using the RMC GUI, you can perform the following:

- View all the tasks performed in RMC appliance using dashboard.
- Add storage systems and backup systems.
- Create snapshots, Express Protect backups, and Catalyst Copies.
- Add users and assign roles.
- Schedule snapshots and Express Protect backups.
- Schedule and configure Catalog Protection of the RMC appliance to the network share or Catalyst Store.
- Access all the features of RMC for Microsoft SQL (RMC-S), RMC for Oracle (RMC-O), RMC for SAP (RMC-SH), and RMC for Microsoft Exchange (RMC-E).
- Restart or shut down the appliance.
- Replicate data between source storage system and target storage system. You can also schedule data replication.
- Generate a support ticket.
- Diagnose RMC issues by checking the status of subsystems.

InfoSight

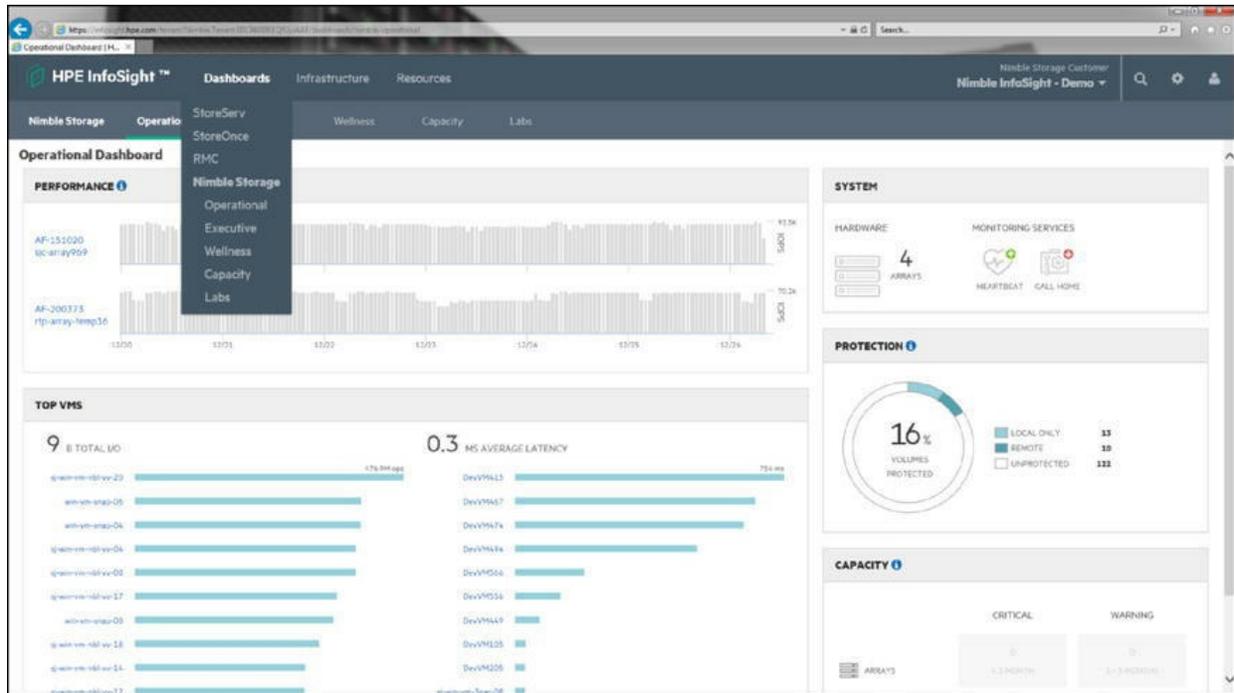


Figure 5-99 InfoSight

HPE InfoSight, as shown in [Figure 5-99](#), is built on a unique approach to data collection and analysis. It starts with having the right data. Data from the infrastructure has always been limited to log files and graphs on obvious metrics and stuck on-prem. InfoSight takes a much more comprehensive approach. Every second, InfoSight is collecting thousands of embedded sensors built into every array and pulling in data from across the infrastructure stack.

Infosight is not just analyzing individual systems, but it is also analyzing all the systems in the installed base in the cloud-based AI platform, where machine learning generates predictive analytics, continuously learns from every system, and automates recommendations that improve every environment.

InfoSight takes a comprehensive approach to data collection and analysis. Data is the key to InfoSight's global learning. Every second, millions of sensor measurements capturing the state of the systems, subsystems, and surrounding IT infrastructure in the global installed base are collected and analyzed. InfoSight continuously learns from the telemetry and develops an understanding of the ideal operating environment for every system,

workload, and application.

Advanced machine learning then drives InfoSight's predictive analytics and recommendation engines, where problematic behavior is predicted through recognition of the underlining patterns and configurations in each system.

InfoSight features:

- Predictive support automation
 - Predicts and prevents issues
 - Solves problems across infrastructure
 - Transforms the support experience
- AI-driven management
 - Makes managing effortless
 - Optimizes resource allocation
 - Sees what others cannot
- Unique product experience
 - Makes the infrastructure smarter
 - Enables the infrastructure to self-improve
 - Simulates work-load changes
- Provides access to various documents, such as HPE Nimble Storage Integration Guide

This platform drives predictive support automation (predicting and automating resolution to problems before they can impact customers), preemptive recommendations (decisions that are automatically made for IT that prevent issues and improve every environment), proactive management through advanced analytics in the cloud portal, and continuous improvement of the systems (making them smarter and more reliable).

The benefits, as demonstrated with Nimble Storage, include:

- Predictive support enables us to deliver a radically better support experience. Eighty-six percent of problems are predicted and automatically resolved before customers even realize there is an issue,

where 54% are even beyond storage. This automation has enabled Nimble Storage to completely automate Level 1 and 2 support.

- AI-driven management has made managing infrastructure effortless and brings 79% lower storage OpEx.
- Infrastructure is no longer sitting idle and unconscious; it is self-improving. This is realized in the fact that Nimble Storage has over 99.9999% of measured, proven, and guaranteed availability across its installed base, going back to the first array that ever shipped and across all OS, models, and configurations. Historically, infrastructure gets less reliable as it ages, but Nimble has flipped that paradigm with InfoSight.

InfoSight has fundamentally transformed how infrastructure is managed and supported. Traditional infrastructures were not designed with the intelligence to effectively harness the data across a hybrid cloud world. More specifically, IT departments are constantly reacting to problems, with considerable time spent on troubleshooting and doing nonvalue-added activities. These factors dictate that a new level of intelligent resources needs to be applied to tackle the data challenge and effectively bridge from discovery to insight to action.

HPE InfoSight brings simplified, AI-driven operations to your customer's hybrid cloud world that transforms how infrastructure is managed and supported. It uses cloud-based AI to provide global insights into the status and health of infrastructure, all in one location. As it analyzes and correlates millions of sensors every minute, all customers benefit as their systems get smarter and more reliable. InfoSight watches over your customer's infrastructure 24x7, so they do not have to spend their days, nights, and weekends dealing with infrastructure issues anymore.

Analytics are only as good as the data available to analyze. Data has always been limited to one-off log files on obvious metrics or bolted on to collect data after an event—and it has always been analyzed one array at a time. We consciously architected for predictive analytics when we started. This required taking a fundamentally different approach to data collection. When we started, we began embedding sensors into every module of code in our operating system to understand all the factors that can impact the IO path.

The plan for HPE is to extend InfoSight and its AI and predictive capabilities

Summary

- Direct-attached storage (DAS) provides the most straightforward HPE storage solutions for small-to-medium businesses (SMBs).
- A storage area network (SAN) is a dedicated high-speed network that interconnects and presents shared pools of storage devices to multiple servers.
- The HPE storage array portfolio includes XP7, HPE Primera, 3PAR StoreServ, Nimble, SimpliVity, and MSA products. The HPE StoreOnce backup system is a disk-based storage appliance.
- HPE offers a variety of tools to help with the storage sizing process.
- HPE network management tools span across products and bring innovations to complete lifecycle management.

Prelearning check

Before proceeding with this section, answer the following question to assess your existing knowledge of the topics covered in this chapter. Record your answers in the space provided.

1. You are in a meeting with a customer, and they challenge you with a statement: “I heard that the management of HPE devices includes manual and repetitive tasks, with a high risk of errors.”

How should you respond?

5 Recommending HPE management and support solutions

LEARNING OBJECTIVES

- ✓ After completing this chapter, you should be able to recommend and position:
 - ✓ HPE support products
 - ✓ HPE management products
-

Recommending HPE management and support solutions for SMB customers

The wide HPE portfolio of data center products for SMB customers can be introduced using customer scenarios.

Introducing the customer scenario

The fictional customer LLP distribution company will be used as a storyline through this scenario. We will introduce the company using an interview:

- What is your primary business?
 - City-wide courier, parcel, and documents delivery
- How many employees do you currently have?
 - 25 employees + external workers

- How does your selling and delivery channel look like?
 - Contracts with dozens of local e-shops and expanding
- What does your server, storage, and network infrastructure look like?
 - ProLiant, StoreEasy, HPE Networking, Nimble
- Do you have an IT department?
 - No, external help on request
- What are your current plans?
 - 24x7 monitoring of HPE devices

Customer requirements

As a result of multiple interviews and gathering information about customer plans and customer's current infrastructure, the following requirements emerged for the new solution:

- 24x7 monitoring
- Health and inventory status
- Automated response
- Email notification
- Devices:
 - HPE ProLiant ML servers
 - HPE StoreEasy
 - HPE Nimble storage
 - Aruba 2920 and 3810M

Activity: Discovering a customer's business and technical requirements

1. Prepare a list of additional questions to ask the customer about:

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:

- Managing and monitoring single system (on system)
 - UEFI
 - iLO 5
 - RESTful API, HPE RESTful Interface Tool, and other HPE scripting tools
 - Intelligent Provisioning
 - Smart Storage Administrator
 - HPE Smart Update and Service Pack for ProLiant
 - SSH/Telnet
 - 3PAR CLI
 - StoreOnce Management Console
 - StoreEasy Management Console
 - HPE Nimble Storage graphical user interface (GUI)
- Managing and monitoring multiple systems—installed at customer site (on-premises)
 - StoreServ Management Console
 - HPE OneView
 - HPE OneView Global Dashboard
 - iLO Amplifier Pack
 - Intelligent Management Central
 - Aruba AirWave
 - Aruba NetEdit
 - iLO Federation

- StoreOnce Federation
- Managing and monitoring multiple systems—Cloud-based management
 - Aruba Central
 - Remote Support and Insight Online
 - InfoSight

Learning check

1. Name three HPE offerings for single system management.

Remote IT support

Remote IT Support services reduce unplanned downtime with automated, remote monitoring of servers, storage, and networking technologies.

HPE products can be connected to HPE to unlock all the benefits of technology investment with 24x7 monitoring by HPE. Up to 77% reduction in down time can be achieved, near 100% diagnostic accuracy, and a single consolidated view of customer environment. Remote IT Support can be used to receive prefailure alerts, automatic call logging, and parts dispatch with HPE Proactive Care service and HPE Datacenter Care.



Note

More information about Remote IT support can be located at:
<https://www.hpe.com/us/en/services/remote-it-support.html>

HPE offers options for customers to connect:

- **Through a central-connect onsite Hosting Device**—This option can be used to monitor multiple device types, including all servers, storage, and networking products. This option will limit the amount of outbound connections to a single host. Central Connection options include:
 - Insight Remote Support
 - OneView Remote Support (HPE Synergy requires OneView Remote Support)
- **Through a direct connection from each device to HPE**—Devices capable of direct connection include:
 - ProLiant Gen8/Gen9/Gen10 servers
 - BladeSystem c-Class enclosures
 - Select storage devices: 3PAR StoreServ, StoreVirtual and StoreOnce



Note

Learn about possible connection options, how to configure customer devices, and how to customize IT environment to maximize the reliability of HPE technology and enable the best possible support services at: <https://www.hpe.com/us/en/services/get-connected.html>

HPE Insight Online

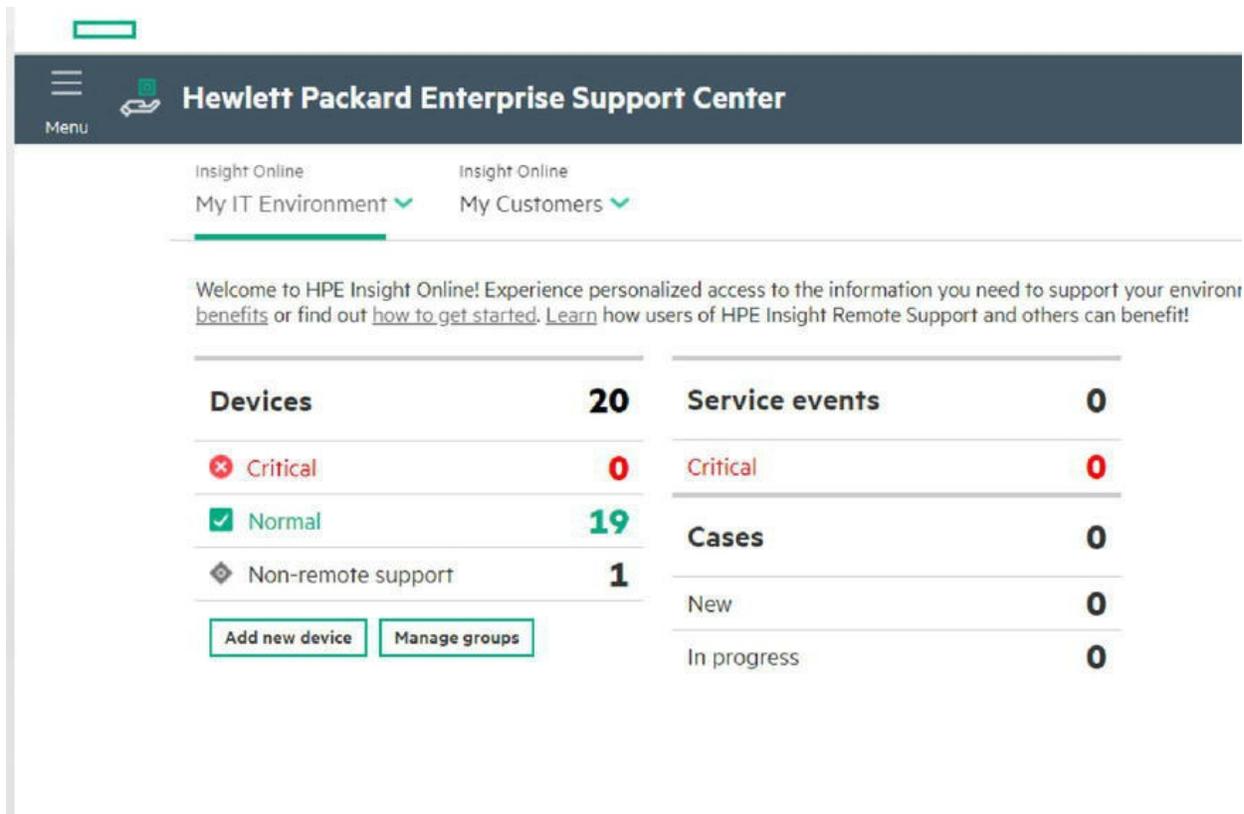


Figure 6-1 Save time and reduce downtime through increased efficiency

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, as illustrated in [Figure 6-1](#). 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. 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.

Insight Online offers:

- Easy all-in-one access, secure, personalized dashboard
- Round the clock monitoring
- Automated cases and spare parts sent to your customer's door
- Simplified warranty and contract and case management
- Support for any size, small business to enterprise
- Converged support for servers, storage, and networking
- Access to the service anywhere, anytime, from any device

Choice of HPE Remote Support tools with Insight Online

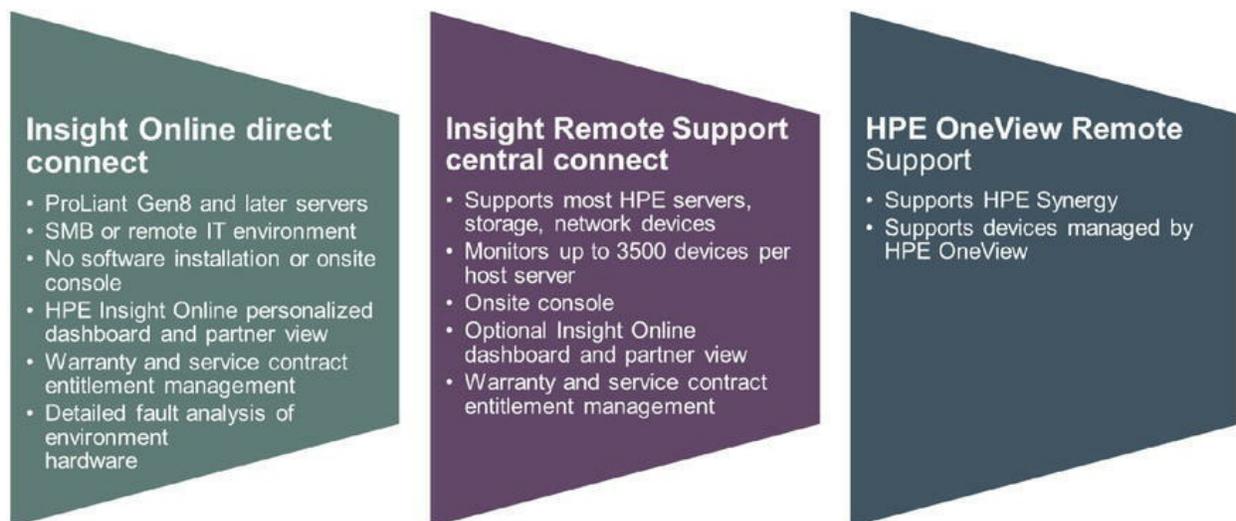


Figure 6-2 Insight Online is provided by three connection methods

To take full advantage of Insight Online, you need to install HPE remote support technologies. Remote support with Insight Online is provided by several connection methods, as shown in [Figure 6-2](#):

- **Direct connect**—Is 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, your customer uses 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**—Is available for servers, storage, and networks. Using central connect, your customers register the device to communicate with HPE through an Insight Remote Support centralized hosting device in their local environment. This method is ideal for HPE Converged Infrastructure IT environments with multiple device types.
- **HPE OneView Remote Support**—Supports devices managed by HPE OneView. Necessary for support of HPE Synergy.

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 24x7 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.

Activity: Insight Remote Support: Support Matrix

1. Open the Insight Remote Support: Support Matrix information from <http://www.hpe.com/info/insightremotesupport/docs>
2. Answer the following questions:
 - a. Is HPE ProLiant DL325 Gen10 supported?

b. What is the required protocol for monitoring and discovery from the Aruba 8320 Switch Series?

Embedded Remote Support

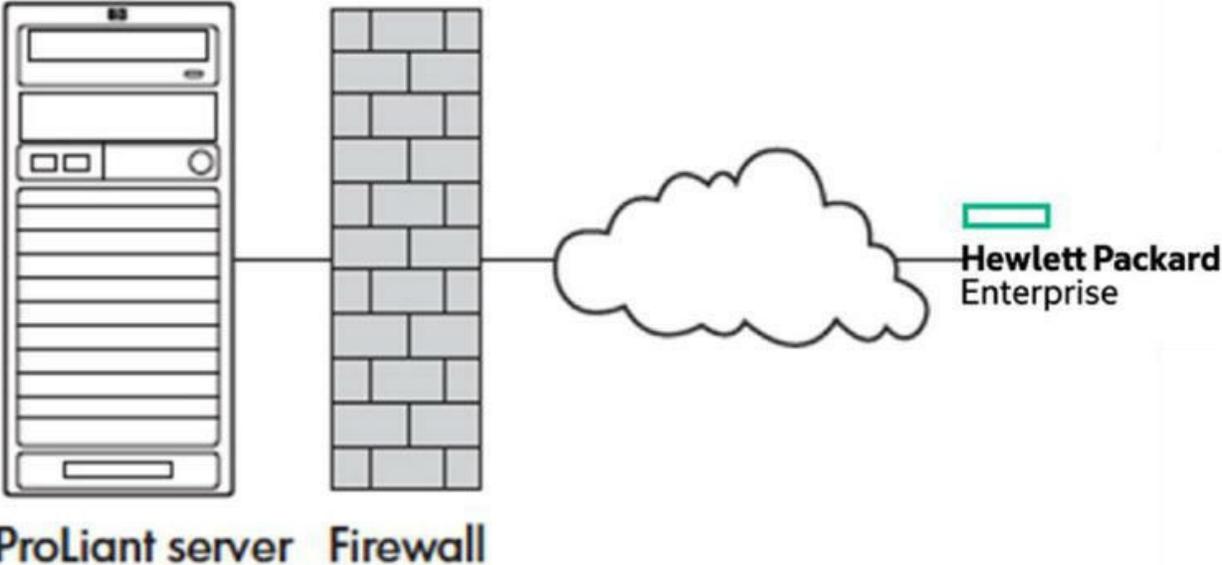


Figure 6-3 Insight Online direct connect

Embedded Remote Support enables your customers to send hardware issues to HPE Support, as part of an HPE warranty or Care Pack, as indicated in [Figure 6-3](#). System administrators can easily and securely export the Active Health file to an HPE Support professional to help resolve issues faster and

more accurately. When customers use the embedded remote support feature, they can choose from the two configuration options: Insight Online direct connect and Insight Remote Support central connect.

HPE Support can receive the necessary data automatically, so HPE Support can solve even the most elusive, intermittent issues in a minimum amount of time and with little effort by the local administrator.

Your customers can register a server to communicate directly with Insight Online without the need to set up an Insight Remote Support centralized hosting device in their local environment. Insight Online is your customer's primary interface for remote support information.

Insight Remote Support Central Connect

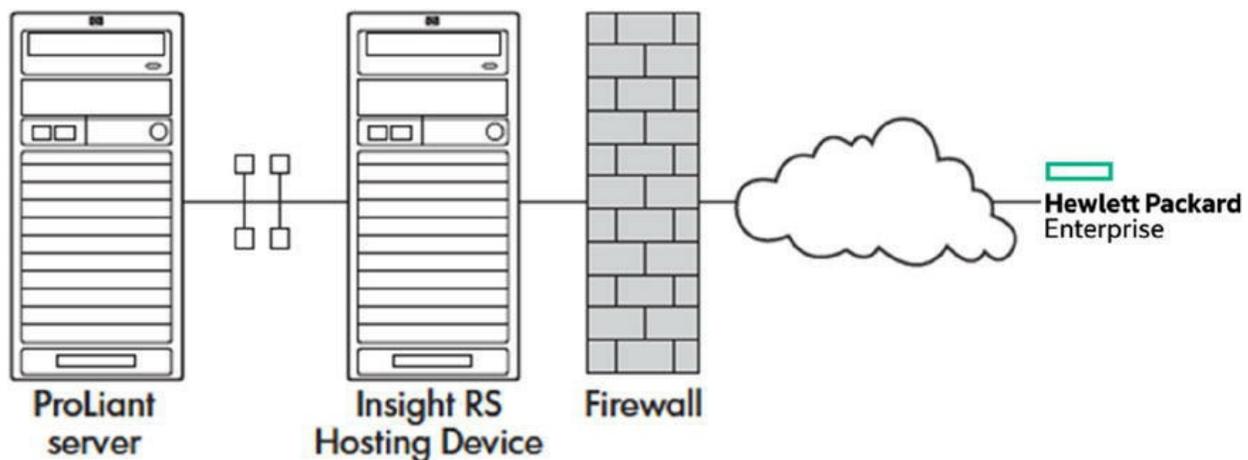


Figure 6-4 Insight Remote Support Central Connect

Your customer can register a server to communicate with HPE through an Insight Remote Support or HPE OneView, working as centralized hosting device in their environment. All configuration and service event information is routed through the hosting device, as indicated in [Figure 6-4](#). 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 6-5 HPE Insight Online—My IT Environment dashboard

HPE Insight Remote Support monitors and aggregates the server environment and alerts the customer to any issues, as shown in 6-6.

The benefits and features of using Insight Remote Support include:

- Risk mitigation
- Cost reduction
- Agentless, fast setup for ProLiant Gen10, Gen9, and Gen8 servers
- Installation wizards to discover and verify readiness of devices and to set up hosting device
- Automatic grouping by hardware type, with easy configuration
- Improved configurability of notifications
- 24x7 phone-home from HPE call centers
- Automatic problem reporting and support case submittal

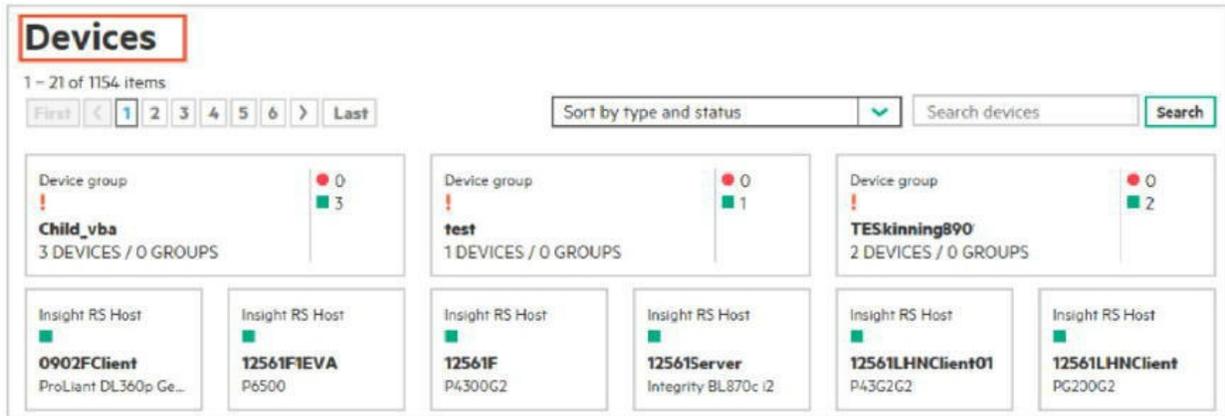


Figure 6-6 HPE Insight Online—Devices overview



Note

Additional Insight Remote Support information is available at: <http://www.hpe.com/info/insightremotesupport/docs>

Insight Online and Remote Support integration

My IT environment: For customers	My customers: For HPE authorized partners
<ul style="list-style-type: none"> – Automatically displays devices remotely monitored by HPE remote support tools – Track service events and support cases – View device configurations and proactively monitor service credit information – View and manage HPE contracts and warranties <div style="text-align: right; margin-top: 10px;">  </div>	<p>With authorization from customers, HPE partners can:</p> <ul style="list-style-type: none"> – Remotely view their customer's IT environment online – Provide service support and consultation <div style="text-align: right; margin-top: 10px;">  </div>

Figure 6-7 Insight Online and Remote Support integration for customers and partners

Insight Online provides one stop, secure web access to product and HPE support information specific to IT environments.

As indicated in [Figure 6-7](#), 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 and 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 and 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



Figure 6-8 Target markets of direct connect for Remote Support

Direct connect for remote support provides automated support, faster resolution, simplified contract, and warranty management. As indicated in [Figure 6-8](#), 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

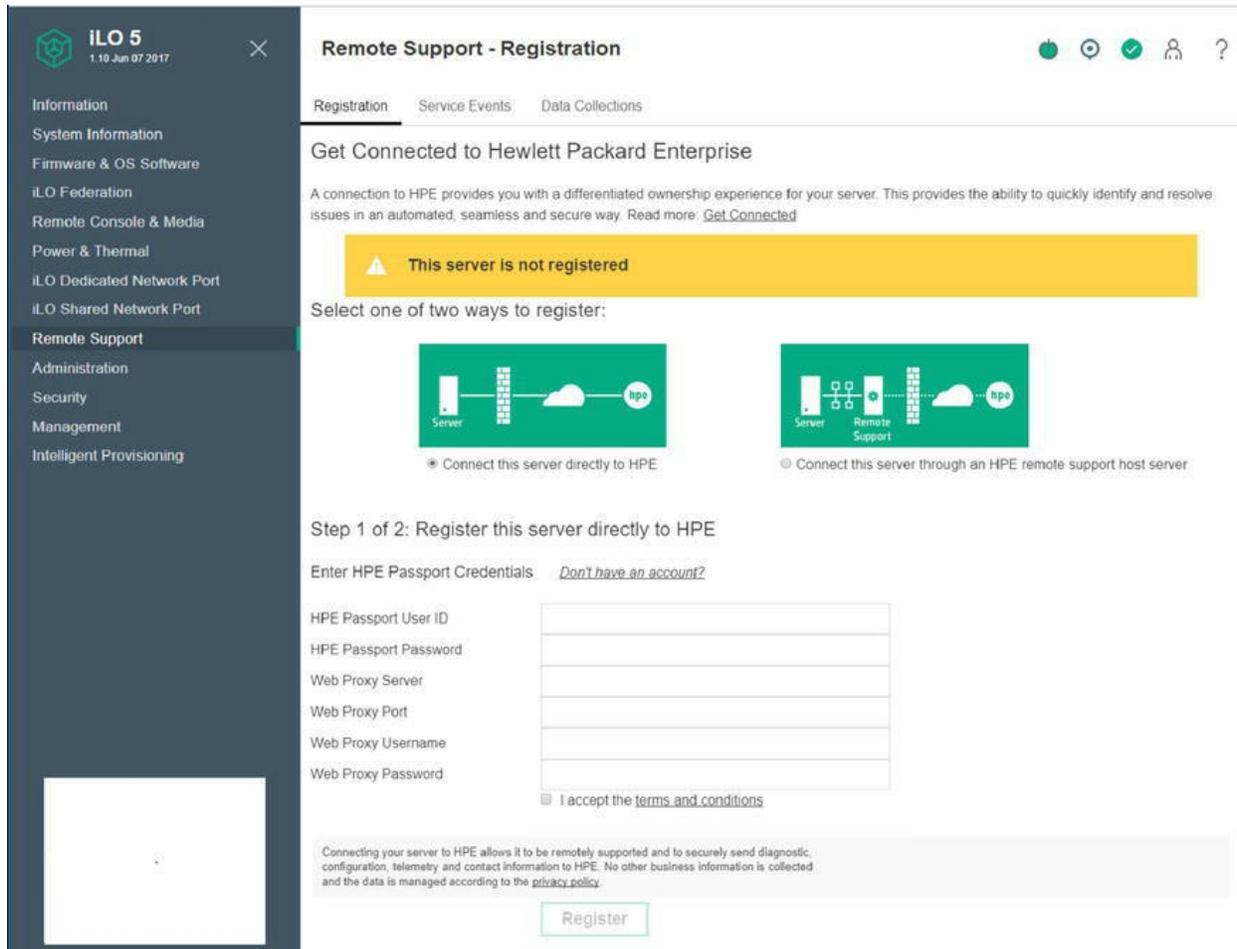


Figure 6-9 Remote Support registration

Your customer can use this procedure to register a ProLiant server for direct connect remote support. As shown in [Figure 6-9](#), by registering, they agree to send registration, service events, configuration, and Active Health System data to HPE. When your customer registers for Insight Remote Support direct connect, they must complete steps in both the Intelligent Provisioning and the Insight Online portal.

1. Verify that the server meets the prerequisites for using the Insight Remote Support solution.
2. Select **Register this server directly to HPE**. The page refreshes to show the Insight Online direct connect registration options.
3. Enter the HPE Passport credentials in the HPE Passport User ID and HPE

Passport Password boxes. In most cases, your customer's HPE Passport user ID is the same as the email address they used during the Passport registration process.

4. Optional: Enter information about the web proxy server if the ProLiant server uses a web proxy server to access the internet.
5. 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 the registration request to be fully processed.
6. Navigate to the Insight Online website at <https://support.hpe.com/hpesc/public/home>, and log in with the HPE Passport credentials.
7. Follow the instructions in Insight Online and provide site, contact, and partner information so that HPE can deliver service for the server. If your customer has multiple servers to register, they should complete Step 1 for all of the servers, and then complete Step 2 for all of the servers during one Insight Online session.
8. Return to the Insight Remote Support page in Intelligent Provisioning, and then click **Confirm**.
9. Click the **Continue** right arrow to proceed to the Intelligent Provisioning home page (Overview).

Registering for Insight Remote Support through a centralized hosting device

Your customer can use this procedure to register a ProLiant server for Insight Remote Support central connect.

1. Verify that the server meets the prerequisites for using Insight Remote Support.
2. 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.

3. Enter the Insight Remote Support hosting device host name or IP address and port number. The default port is 7906. Click **Register**.
4. 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 your customer uses the Intelligent Provisioning Recommended installation method for Windows installation.

If your customer does 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, your customer agrees 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 network interface card [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* at:
http://h20564.www2.hp.com/hpsc/doc/public/display?docId=a00026409en_us

Activity: HPE Get Connected

1. Open the Remote IT Support Getting Started Guide at:
<https://www.hpe.com/us/en/services/get-connected>
2. Answer the following questions:
 - a. What are the options to connect devices to HPE?

- b. Is there a possibility to script the configuration for remote support?

Learning check

2. Which HPE ProLiant server generations are supported for direct connection to Remote Support?

HPE OneView

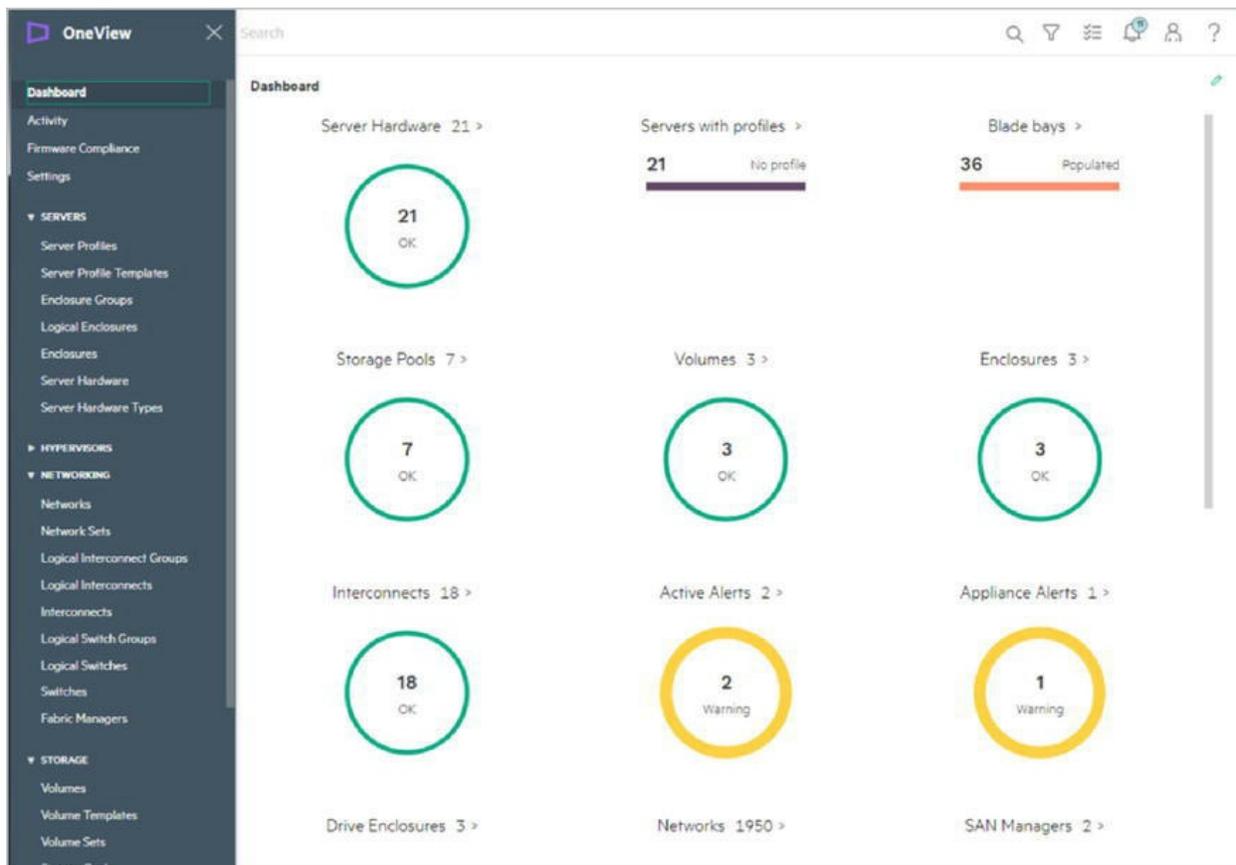


Figure 6-10 HPE OneView

HPE OneView is the foundation solution for the software-defined data center (SDDC) that helps drive the business and enable innovation. HPE OneView, as shown in [Figure 6-10](#):

- Is an API-driven infrastructure management platform.
- Is optimized for software-defined and hybrid cloud environments that experience frequent changes to infrastructure configurations and resource assignments.
- Offers a platform to uniformly define and maintain firmware and system configurations while maintaining consistent availability and control across virtualized, containerized, and bare-metal resources.

HPE OneView is a unified infrastructure management platform designed to address the operational challenges that infrastructure administrators face as they attempt to simplify operations, improve integrations across management processes and tools, and seek more unified, automated management strategies that help staff work more efficiently.

OneView is a modern, application programming interface (API)-driven infrastructure management platform optimized for software-defined and hybrid cloud environments that experience frequent changes to infrastructure configurations and resource assignments. It offers customers a platform to uniformly define and maintain firmware and system configurations while maintaining consistent availability and control across virtualized, containerized, and bare-metal resources.

OneView enables customers to streamline complex lifecycle operations and application delivery activities across network, compute, storage, and composable infrastructure, including:

- HPE Synergy (with Composer powered by OneView)
- HPE ProLiant BL, DL, and ML servers
- HPE Apollo
- HPE Superdome X systems
- HPE 3PAR StoreServ storage systems
- HPE StoreVirtual VSA

By providing a consistent management platform to configure and provision this broad range of resources, OneView can significantly improve operational productivity and overall business agility.

The key benefit of OneView is that it is one tool that uses one dataset to present one view to the administrator, combining complex and interdependent data center management capabilities in a unified interface.

Infrastructure automation made simple



Figure 6-11 OneView simplifies hybrid IT

OneView simplifies hybrid IT, transforming servers, storage, and networking into a software-defined infrastructure. OneView, as indicated in [Figure 6-11](#), makes infrastructure automation simple by increasing the speed of IT delivery for new applications and services, while still efficiently managing traditional IT.

OneView is designed to:

- Manage across compute, storage, and fabric with an easy-to-use interface for single-console infrastructure management. A single console provides a unified view of the health of servers, profiles, enclosures, storage, and networking, which enables administrators to rapidly design, provision, monitor, and update IT resources. Administrators can proactively monitor the health of the entire infrastructure and identify problems before they result in downtime.
- Leverage a modern, standards-based API and a large and growing composable partner ecosystem to increase the speed of IT delivery for

new apps and services while efficiently managing traditional IT, giving customers the confidence that they are investing in an architecture for the future with solid support.

- Reduce the risk of unauthorized user access to system management resources with enhanced authentication and encryption security features of Simple Network Management Protocol version 3 (SNMPv3), two-factor authentication, and scope-based access.
- Simplify IT operations by rapidly bringing hundreds of servers into a OneView environment with just a few clicks, dramatically shortening server setup time with the auto-discovery feature.

Simplify lifecycle operations



Figure 6-12 OneView Global Dashboard

OneView eliminates complexity, removes silos across various IT organizations, and simplifies the overall infrastructure lifecycle management. These capabilities alleviate the burden and time spent managing hardware, and free up more time to spend developing and delivering new applications and IT services.

The OneView Dashboard screen, as shown in [Figure 6-12](#), provides a

graphical representation of the general health and capacity of the resources in the data center. From the Dashboard your customers can immediately see the areas that need their attention.

Additional functionality is provided with the OneView Global Dashboard software, which provides an at-a-glance view and health status of multiple servers, profiles, and enclosures around the world and across a OneView managed environment—up to 54 enclosures and 25 appliances. With the Global Dashboard, your customers can obtain better infrastructure visibility and make faster, more informed decisions.

OneView keeps the infrastructure running at peak performance with minimal workload disruption. Infrastructure device drivers and firmware updates can be automated with the least possible impact to the production environment. Templates, used in combination with HPE Software Update Tools (SUTs), enable updates to be staged, scheduled, and installed without rebooting the system. This frictionless, non-disruptive update feature has now been extended to the Synergy composable server platform through the Synergy Composer powered by OneView. Changes to the templates are immediately reflected in profile compliance status, so your customers can quickly pinpoint systems that need attention and efficiently roll out updates.

In addition, OneView enables your customers to receive 24x7 monitoring, prefailure alerts, automatic call logging, automatic parts dispatch, and contract/warranty display through integrated remote support, which is now extended to include Synergy interconnect modules and ProLiant Gen10 firmware. HPE Proactive Care Services are also available to provide a personalized and proactive hands-on approach to maintaining an agile, healthy, and reliable infrastructure.

HPE Composable Infrastructure partner ecosystem



Figure 6-13 Growing customer adoption and support for OneView

Customer adoption and support for OneView continues to grow. Interviews with organizations that have deployed OneView to support their server, storage, and network infrastructure reveal that they are achieving significant value with OneView by making their daily IT operations more efficient and effective. These organizations, as indicated in [Figure 6-13](#), benefit from having a single console to manage their environments, as well as software-defined processes and automation.

In addition, OneView 4.2 has expanded capabilities to leverage new partners in the HPE Composable Infrastructure Partner Ecosystem program. Your customers can enable hybrid cloud management through API integrations with Composable Infrastructure Ecosystem partners, including:

- **Morpheus Data**—With Integration into OneView and 3PAR storage, Morpheus enables Unified Ops Orchestration of both DevOps pipelines and multi-cloud infrastructure including visibility, analytics, and physical state management of an HPE infrastructure running bare-metal servers, virtual machines (VMs), or containers in hybrid environments.
- **CANCOM**—OneView with CANCOM allows users to create a fully integrated digital workspace solution including hardware, software, and services to ensure seamless and automated deployment with real-time monitoring, maintenance, and daily operations capabilities.
- **InContinuum**—Using OneView with InContinuum software, organizations can perform lifecycle management and automation of bare-metal, virtual, and hybrid cloud deployments including control functions such as platform as a service, DevOps, and containers from a unified view.

Deploy infrastructure at cloud-like speed

Through software-defined intelligence, 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 administrators to develop a template once and then replicate as needed, but also helps boost productivity of administrators and software developers. In addition, change operations can be implemented by using templates so that tasks such as adding more storage to a service, modifying network connectivity, or updating firmware are implemented automatically.

Continuous, automated lifecycle operations reduce cost, save time, and increase time to value for businesses. With templates, your customers can simplify system updates and enforce compliance to ensure infrastructure stability. They can also manage deployment plans and create bootable images from capturing, cloning, or customizing golden images. Furthermore, customers can enforce compliance by using templates to quickly provision, update, or roll back images to minimize maintenance windows.

For customers with HPE BladeSystems, OneView automates the creation of Fibre Channel and Ethernet connections, reducing setup time from hours to

minutes. Integration with HPE Networking's Intelligent Management Center (IMC) and HPE Virtual Connect advances this process further. IMC listens for newly provisioned BladeSystem enclosures, and then automatically connects them to the production network.

Develop more apps faster

The OneView unified API, together with the growing ecosystem of partner integrations, enables your customers to accelerate application and service delivery; developers, IT administrators, and independent software vendors (ISVs) can automate infrastructure with a single line of code.

A substantial portion of operations work consists of routine tasks related to infrastructure lifecycle management, including designing, provisioning, monitoring, and updating. OneView is designed to automate day-to-day responsibilities by simplifying time-consuming tasks leading to increased productivity and reduced operational costs. It is an automated infrastructure provider under any environment that supports traditional, virtualized, and cloud workloads, including VMware, Microsoft, and OpenStack.

HPE OneView addresses two approaches to IT management

OneView makes daily IT operations more efficient and effective by enabling large and small enterprises to manage their environments from a single, centralized console. OneView also enables programmatic management of the infrastructure through a standards-based API. When software-defined processes are combined with automation, less IT staff time is needed to deploy, manage, and support core infrastructure, saving much-needed IT resources and simplifying management.

For example, when your customer adds a device to an environment, OneView automatically detects the hardware and prepares it for monitoring and management. If it already has a server profile, it is allocated and fully configured. If it does not yet have a server profile, the hardware is available in a global pool awaiting a new configuration. This supports dynamic hardware reconfiguration and makes it easy to provision a new server profile just like the last one. It guarantees that the server profile will successfully deploy to the allocated hardware based on knowledge of the server hardware

type and enclosure group.

In addition to facilitating daily IT operations, OneView also benefits DevOps, where the rate of IT infrastructure provisioning, configuration, and deployment activity is dramatically increasing. New code often is deployed monthly or weekly or even more frequently. The OneView unified infrastructure management platform is designed to enable customers to streamline complex lifecycle operations and application delivery activities across data center resources. OneView combines software-defined resources with template-driven provisioning and management to provide the right resource at the right time for each workload in the data center.

Two methods for an enhanced user experience

The OneView UI approach is designed to enhance the interaction among IT staff and to match work practices in the data center. It is simple, efficient, and consistent, and it is inspired by commonly used web technology. It works on desktops, tablets, and mobile devices. You can right-click the OneView UI to open in a new tab, copy and paste browser bookmarks, email links to colleagues, and print diagrams and data. Search capabilities, newsfeeds, and other functions work as expected in a web experience.

The OneView UI builds functionality around an administrator's work practices and puts resources in the menu. The combination of Java compatibility checks and security updates creates an operational burden, so OneView uses HTML5 as well as CSS3, JavaScript, and AJAX.

Your customers also have the choice of programmatic interfaces based on Representational State Transfer (REST) APIs. REST APIs involve the use of relatively basic create, read, update, and delete (CRUD) operations that are applied to resources (objects) by using standardized HTTP POST, GET, PUT, and DELETE commands. Using the REST API support available in OneView, your customers can manipulate resources in a way that is consistent with the equivalent actions they perform using the OneView UI. The REST APIs integrate with a broad ecosystem of management partners. This includes service desk, orchestration, monitoring tools, configuration management database (CMDB), and more.

The OneView UI and REST APIs are organized by resource. The online help

for each screen in the UI describes the resources and, as needed, their configuration rules.



Note

Representational state transfer (REST) is a software architectural style that defines a set of constraints to be used for creating Web services. Web services that conform to the REST architectural style, called RESTful Web services (RWS), provide interoperability between computer systems on the Internet ([Wikipedia.org](https://en.wikipedia.org), October 2019).

OneView licensing

Features	OneView Standard	OneView Advanced
Partner integrations		☑
Software-defined infrastructure (profiles, groups, sets, others)		☑
Storage provisioning and SAN zoning		☑
Virtual Connect advanced management		☑
Firmware management		☑
Power management (3D visualization)		☑
Operating system provisioning		☑
Remote management (OneView Advanced with HPE iLO Advanced)		☑
Map View	☑	☑
Smart Search, Activity View, Dashboard	☑	☑
Health monitoring	☑	☑
Inventory	☑	☑
Reporting	☑	☑
REST API access	☑	☑
Remote Support	☑	☑
TS&U	One year of 9x5 support (option)	Three years of 24x7 support (standard)

Figure 6-14 Comparison of OneView Standard and OneView Advanced licensing features

A OneView license is required for any managed server hardware. This table, as shown in [Figure 6-14](#), compares the options available with the two main types of licenses available for OneView 4.2:

- **HPE OneView Standard** is a no-fee license for the software. The standard license is licensed per physical server and enables monitoring of supported HPE hardware. With the Standard license, 9x5 support is

available for an additional fee.

- **HPE OneView Advanced** provides full-featured licenses for purchase. All OneView Advanced versions are licensed per physical server. These licenses include three years of 24x7 technical support and software updates with web-based training to build basic product proficiency. Trial versions of OneView Advanced can be used for 60 days without charge. There are two types of OneView Advanced licenses:
 - **OneView Advanced without iLO** is for server hardware that has an existing HPE iLO Advanced license or that does not require the features enabled by iLO.
 - **OneView Advanced with iLO** includes an iLO Advanced license for the server hardware, which enables advanced management features.

Managing, monitoring, and migrating server hardware

Add Server Hardware ?

iLO IP addresses or host names

ilo-d1360gen10.zeta.local
ilo-d1360gen9.zeta.local
ilo-d1380gen8.zeta.local
ilo-d1580g7.zeta.local
ilo-d1385g7.zeta.local
ilo-d1380g6.zeta.local

Add server hardware as Managed Monitored

Scope: development X
Select zero or more scopes

Credentials

User name: admin

Password: *****

Changed: User name to "admin"

Add **Add +** **Cancel**

Figure 6-15 Managing, monitoring, and migrating server hardware in OneView

Server hardware such as enclosures and HPE ProLiant rack servers, as indicated in [Figure 6-15](#), can be added to OneView in one of the following ways, depending on the enclosure type:

- **Managed**—If your customer adds a managed server to OneView, either in an enclosure or rack, they can apply configurations, deploy server profiles, monitor operation status, collect statistics, and alert users to specific conditions. Managing server hardware requires OneView Advanced licensing. For more information, see “About managed c7000 enclosures” in the OneView online help.
- **Monitored**—If your customer adds a monitored server to OneView, either in an enclosure or rack, they can monitor it for inventory and hardware status only. Monitoring server hardware uses a free OneView Standard license. For more information, see “About monitored enclosures” in the online help.

Because a monitored enclosure cannot be managed, OneView does not create various configuration objects for a monitored enclosure. These include an enclosure group, logical interconnect group, logical interconnects, and server profiles. Objects including the enclosure, server hardware, and server hardware type are created.

- **Migrated**—Enclosures from Virtual Connect Manager (VCM) and Virtual Connect Enterprise Manager (VCEM) can be migrated to OneView with the configuration information, so that the enclosure can be managed by OneView. The managed enclosure requires the OneView Advanced licensing. For more information about migrating, see “About migrating c7000 enclosures managed by other management systems” in the online help.

HPE ProLiant DL rack servers, like ProLiant BL server blades, can be monitored or managed. For a ProLiant DL server, your customer can view system health information (CPU, power, and temperature data) as they can with a monitored ProLiant BL server. They can also launch the integrated Lights-Out (iLO) console and see basic Map view information. Because your customer cannot view adapter (LAN on motherboard [LOM] and mezzanine) ports of a managed ProLiant DL server, such information is not available for a monitored ProLiant DL server.

Your customer can find server hardware instances for each server blade on the Server Hardware page. They also can find a data center rack object populated with the enclosure because this is considered a monitoring feature.

Supported devices for monitoring by HPE OneView

HPE OneView can monitor hardware health of selected HPE devices:

- ProLiant BL G6 Gen10
- ProLiant DL Gen6 Gen10
- ProLiant ML Gen8 Gen10
- ProLiant XL Gen9 Gen10
- HPE Composable Cloud for ProLiant DL
- Superdome Flex Server
- Synergy compute modules
- Synergy storage modules
- Synergy Frame 12000
- BladeSystem c7000 Enclosure



Note

Monitoring features of HPE OneView do not require HPE OneView Advanced license.

Always consult Support Matrix. Not all models and/or features are supported. Additional requirements and prerequisites exist.

Resource health monitoring operations

The primary purpose of the OneView health monitoring manager is to provide overall health monitoring of the physical resources under management of the OneView appliance. In general terms, the key objectives of the health monitoring manager are to:

- Ascertain the aggregate status of each resource.
- Present health issues to administrators.
- Minimize the need for polling by using incoming traps for status information.
- Allow SNMP trap forwarding.
- Generate emails to individuals based on alert type.

OneView offers two approaches to monitoring depending on the server:

- **Proactive**—After an HPE ProLiant Gen8 or later server has been added to OneView from the Add Server page, OneView polls for the current system status automatically.

OneView is designed to use incoming traps as a primary source for status information, but OneView also employs periodic polling of devices under management as a fail-safe approach because User Datagram Protocol (UDP)-based SNMP traps can be unreliable. Status polling is performed when a device is first added and then every 20 minutes thereafter. An event is created if the status of the device changes. In the case of server blades and rack servers, an SNMP Get request is issued for the management information base (MIB) object known as the **health status array**.

To automatically configure ProLiant Gen8 and Gen9 servers to send events to OneView, you only need to:

- Provide the Onboard Administrator credentials for server blades.
- Provide iLO credentials for rack servers.
- **Reactive**—For G7 servers and earlier, the server depends on the agent to actively send information to OneView. The appliance then accepts and organizes the information. With this method, Insight Management Agents must be installed and configured in the operating system manually. Events can only be addressed as they are received.

Using the OneView UI and the REST API, you can:

- View all alerts and tasks by description or source, and filter activities by using multiple filter criteria.

- Assign alerts to specific users.
- Annotate activities with notes from administrators, enabling the data center administrators to collaborate through the appliance instead of through outside tools such as email.
- View alerts for a specific resource from the UI page or REST API for that resource.
- Automatically forward SNMP traps from managed resources to enterprise monitoring consoles or centralized SNMP trap collectors.

SNMP trap receiver processing

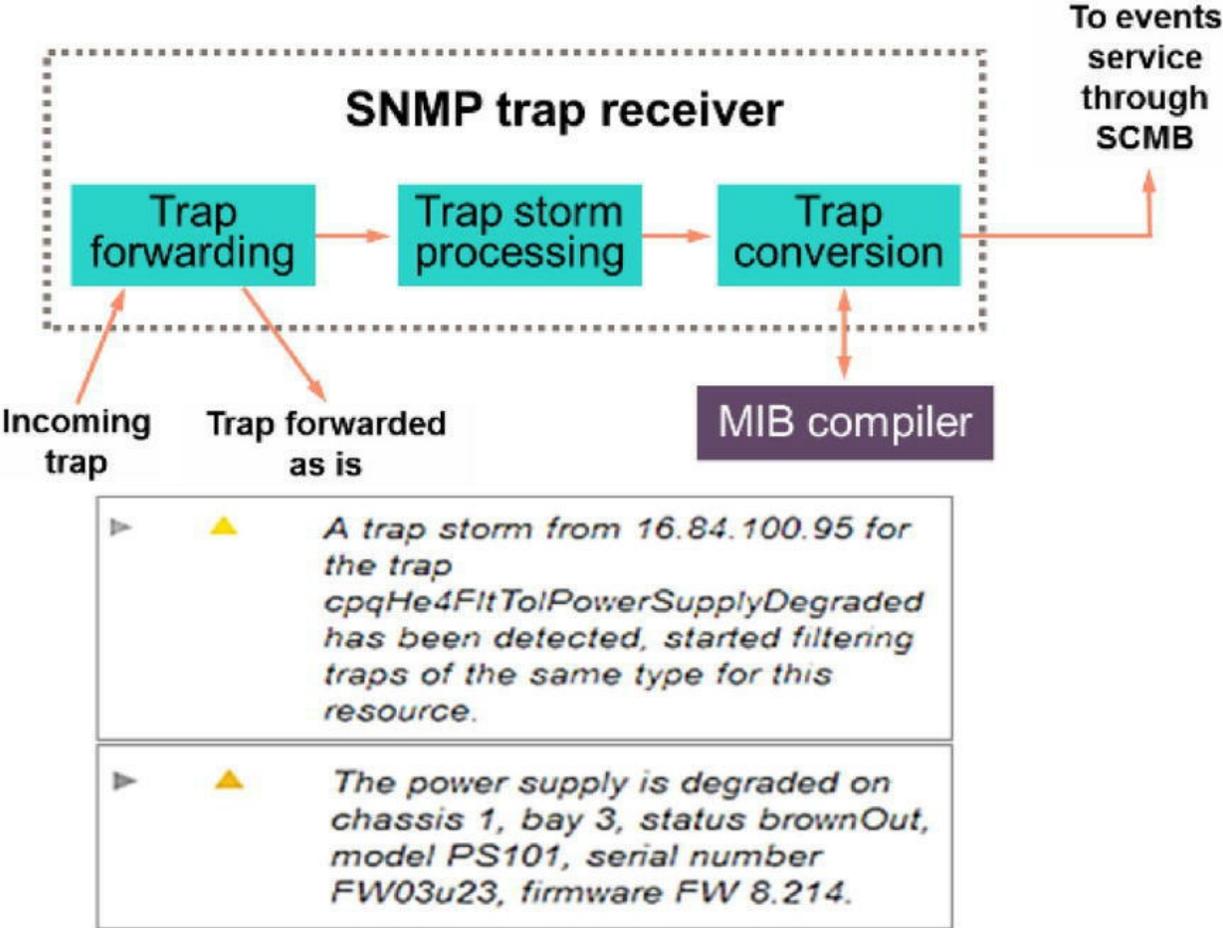


Figure 6-16 The primary functions of the SNMP trap receiver

Figure 6-16 provides a closer look at how the SNMP trap receiver component

operates. The SNMP trap receiver has three primary functions:

- Processing incoming SNMP traps
- Monitoring for SNMP trap storms
- Converting SNMP traps into descriptive events

An SNMP **trap storm** is defined as four or more identical traps occurring within a 30-second period. To be considered **identical**, an SNMP trap must meet three criteria:

- **Is from the same device**—A device that has the same IPv4 address
- **Has the same trap identifier**—The same MIB trap Object ID (OID)
- **Has the same number of variable string values**—In the same order and with the same values

If an SNMP trap storm does occur, then the first three traps are converted into alerts. After the third trap, OneView creates a trap storm event. Intermediate trap storm events will also be created to continue informing administrators about the issue.

It takes up to one minute after the last trap arrives to determine when a trap storm has ceased. At that time, a trap storm ended event is generated.

Alert Aggregation Service processing

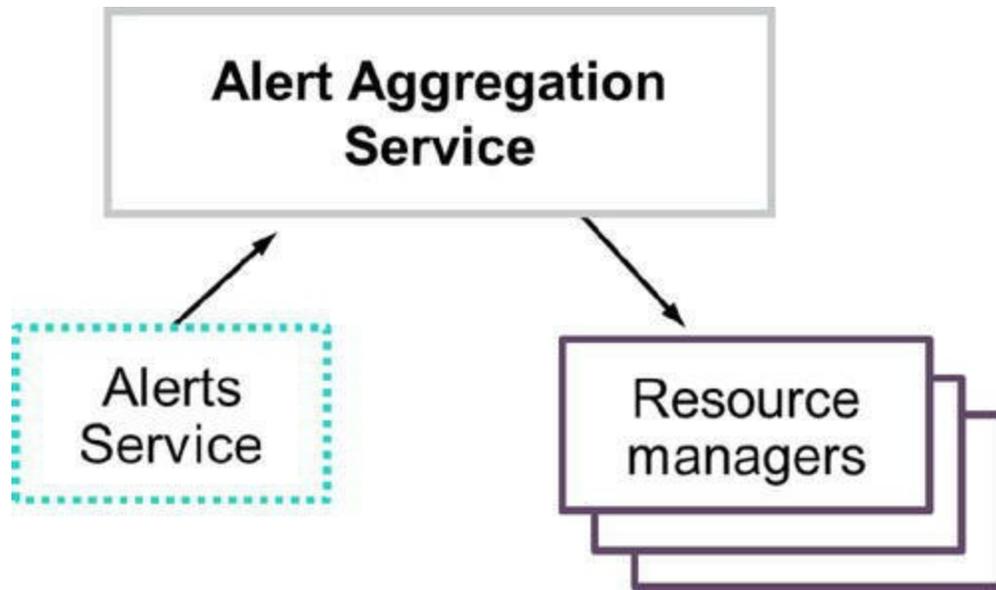


Figure 6-17 The Alert Aggregation Service computes a summary view of resource status

An important subcomponent of the resource health monitoring system is the Alert Aggregation Service, as illustrated in [Figure 6-17](#). Its primary task is to take the worst-case status among all the health categories for a resource and generate a summary view.

The aggregate status is a consolidated view based on the observed operational status from resource managers combined with the recent alert status. The worst-case status has the highest criticality based on relatively standardized definitions such as Critical, Warning, and OK (normal).

System health monitoring

Efficient data views and effective control enable your customers to respond to issues when managing the health of ProLiant servers. When managed resources are added to the appliance, they are automatically set up for monitoring, including the automatic registration of SNMP traps and scheduling of health data collection. ProLiant Gen8 and Gen9 servers are monitored immediately without requiring your customers to invoke additional configuration or discovery steps.

All monitoring and management of data center devices is agentless and out-

of-band for increased security and reliability. Operating system software is not required, open SNMP ports on the host operating system are not required (for ProLiant Gen8 and Gen9 servers), and zero downtime updates can be performed for these embedded agents. ProLiant Gen8 and later servers support agentless monitoring by iLO. OneView uses SNMP in read-only mode to the iLO only, not to the host operating system. ProLiant G6 and G7 servers require host operating system SNMP agents.



Note

Read-only mode means SNMP uses gets and traps, but not sets.

OneView provides proactive alert notifications by email (instead of using SNMP trap forwarding) and automated alert forwarding. Your customers can view, filter, and search their alerts using Smart Search. Alerts can be assigned to specific users and annotated by administrators. Notifications or traps can be automatically forwarded to enterprise monitoring consoles or centralized SNMP trap collectors.

The customized dashboard capability allows your customers to select and display important inventory, health, or configuration information and to define custom queries for new dashboard displays. The single user interface provides additional summary views of firmware revisions and of the hardware inventory for servers, storage, and networks. Other data and inventory elements are visible through the user interface and REST API, and they can be found using Smart Search.

SAN health and diagnostics

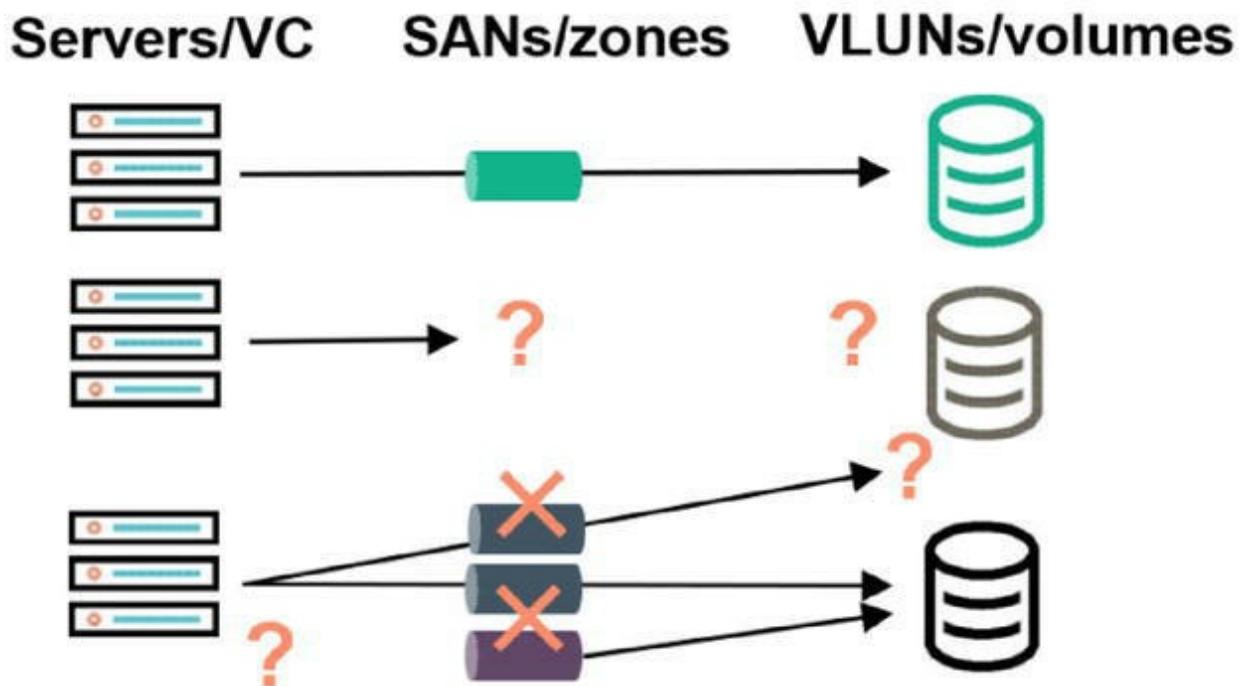
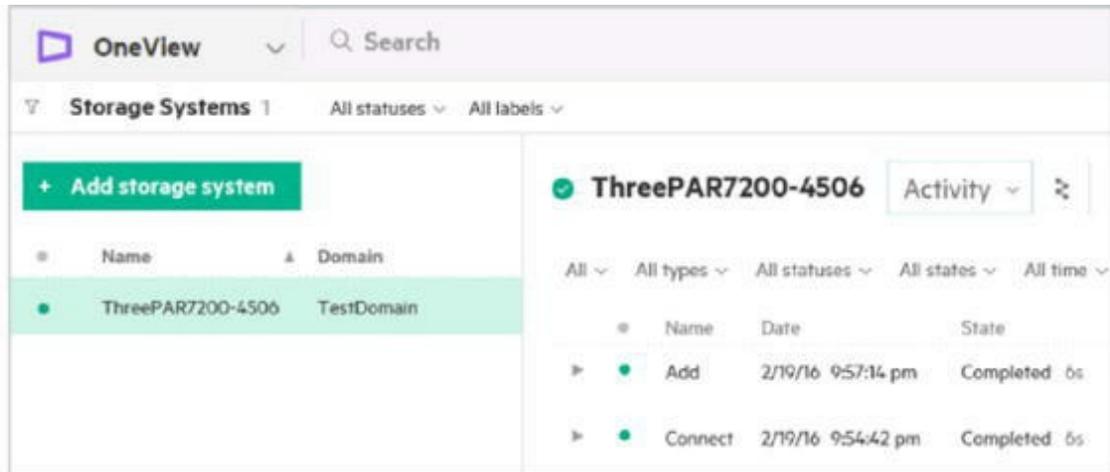


Figure 6-18 SAN resources are immediately exposed in the topology map

Using HPE 3PAR StoreServ storage within OneView, as indicated in [Figure 6-18](#), is as simple as selecting a storage template and a server profile. OneView automation carves out the storage volume, zones the Fibre Channel storage area network (SAN), and attaches the storage to the server profile.

After they are rolled out, the SAN resources are immediately exposed in the topology map. This includes multi-hop Fibre Channel and Fibre Channel over Ethernet (FCoE) architectures. In OneView, proactive alerts are provided

when the expected and actual connectivity and states differ or when SAN health issues are immediately visible in the topology map. OneView provides SAN configuration reports, which include guidance for SAN efficiency and help in resolving potential SAN issues before there is a business impact.

Learning check

3. Match the way of adding hardware to OneView with its description.

Managed	Import devices managed by VCM
Monitored	Deploy server profiles, monitor status, collect statistics, and alert users
Migrated	Monitor for inventory and hardware status

HPE OneView partner integrations

Typically, shifting from one management tool to another, each with a partial view of available data, is both time-consuming and complex. OneView includes integrations that reduce the time needed to make important administrative changes. These integrations provide additional support for partner management platforms.

OneView integrations deliver comprehensive system health and alerting, driver and firmware updates, operating system deployment, detailed inventory, and HPE Virtual Connect fabric visualization. System administrators can gain insight and control of virtualized environments while reducing the time it takes to make important changes, increase capacity, or manage planned and unplanned downtime. You can gain greater control of environments by integrating the management features of:

- HPE Synergy
- HPE ProLiant servers

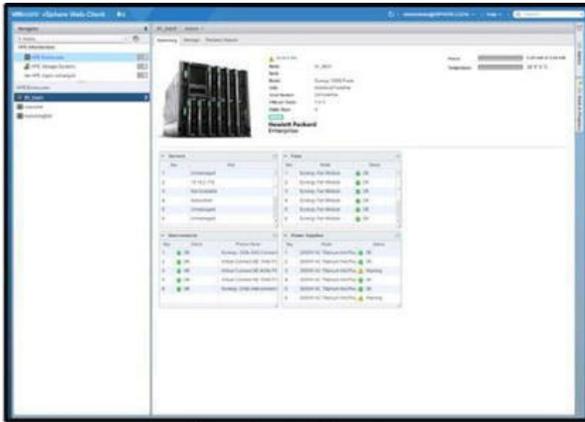
- HPE BladeSystem
- HPE Storage

When used with the automation power of OneView, data center best practices can be defined once and reused many times to provision an entire cluster with compute and storage fully configured. This integration provides consistency for software deployment and updates and enables a faster response in the event of server or storage failure, reducing the risk of downtime.

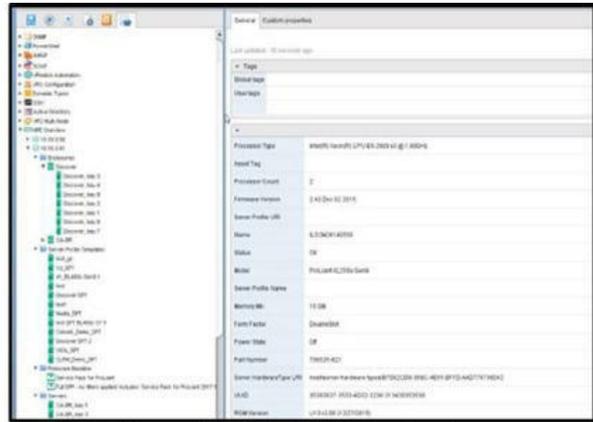
OneView offers the following plug-ins for virtualization management products:

- **HPE OneView for VMware vCenter**—Delivers useful HPE hardware management capabilities to virtualization administrators, enabling comprehensive deployment, provisioning, monitoring, remote control, and power optimization directly from the vCenter console. Integration includes an end-to-end connectivity diagram available in vCenter.
- **HPE OneView for Microsoft System Center**—Supports health monitoring and alerting with server profile provisioning to create or grow a Hyper-V cluster.

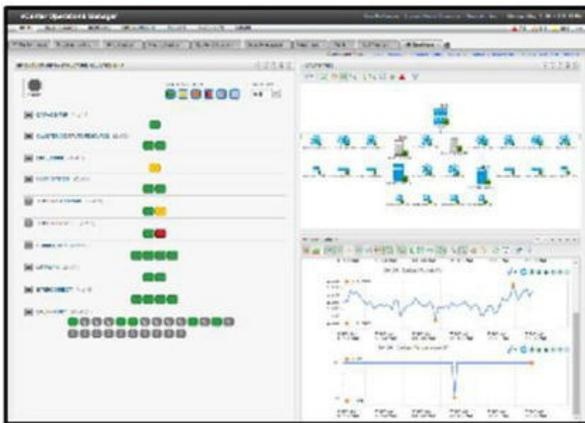
HPE plugins for VMware



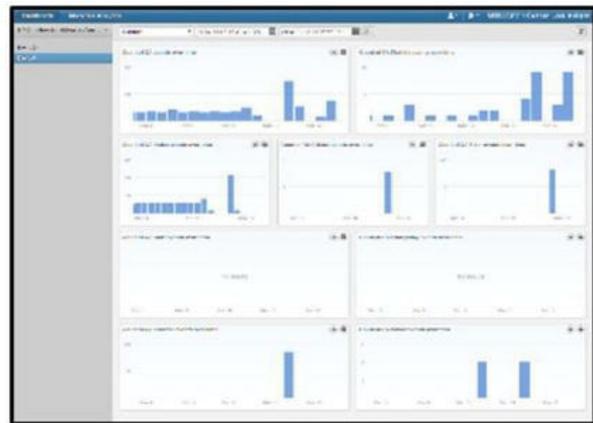
vCenter Server



vRealize Orchestrator



vRealize Operations Manager



vRealize Log Insight

Figure 6-19 Delivering HPE hardware integration directly within VMware native interfaces

OneView for VMware vCenter brings the native manageability of the HPE infrastructure to VMware environments, as indicated in [Figure 6-19](#).



Note

OneView for VMware vCenter supports Synergy Gen10 compute modules, Synergy Composer, HPE Storage products (3PAR StoreServ, MSA, StoreVirtual), and Synergy Image Streamer.

OneView for VMware vCenter 9.0 seamlessly integrates the manageability features of ProLiant, BladeSystem, Virtual Connect, and storage with VMware solutions. It reduces the time needed to make changes, increase

capacity, or manage planned and unplanned downtime.



Note

OneView for VMware vCenter 9.0 is compatible with BladeSystem c7000 and Synergy enclosures with OneView 3.10 and 4.x.

By leveraging OneView software-defined templates and the OneView REST API, OneView for VMware vCenter simplifies the process of deploying a complete vSphere cluster. Users can deploy ESX/ESXi hosts directly from vCenter by right-clicking a cluster or deploying a VMware hypervisor on a bare-metal server. When the process is complete, the new hypervisor is added to the appropriate cluster.

Other functions allow users to:

- Simplify administration with VMware console access to HPE infrastructure management for health, inventory, and configuration monitoring.
- Reduce planned and unplanned downtime by automating responses to hardware events and access to detailed resolution information for health alerts.
- Proactively manage changes with detailed relationship dashboards that provide insight into the relationship between the physical and virtual infrastructure, from the VM to the network edge.
- Maintain stability and reliability with online firmware inventory and deployment.
- Consistently and repeatedly deploy bare-metal servers.
- Take control by launching trusted HPE management tools.
- Provision on-demand server and storage capacity.
- Provides deep analytics, automation, and troubleshooting via integration with vRealize Operations, Orchestrator, and Log insight.

HPE OneView for VMware vCenter: Brings the power of HPE

OneView to VMware environments



Figure 6-20 HPE OneView for VMware vCenter brings the power of HPE OneView to VMware environments

HPE OneView for VMware vCenter, as shown in [Figure 6-20](#), helps system administrators with:

- Deploying the infrastructure faster
 - Template-based provisioning that leverages the HPE OneView automation engine
 - Create or expand a VMware cluster in just five simple steps
- Simplifying lifecycle operations
 - Non-disruptive, comprehensive health monitoring and alerting
 - Firmware and driver updates
 - HPE fabric visualization
- Increasing productivity
 - Consistently integrates directly into VMware consoles
 - Use familiar VMware tools for HPE management tasks
 - Context-sensitive launch of HPE tools directly from vCenter console

The OneView for VMware portfolio provides comprehensive lifecycle management of the HPE converged infrastructure directly from the VMware vCenter consoles. These plug-ins:

- Are ideal for customers who have standardized on VMware tools as their management platform.
- Integrate directly into VMware consoles.

- Leverage the OneView infrastructure automation engine.
- Reduce complexity by using the same VMware tools for HPE management tasks.
- Simplify the admin learning curve by using the VMware tools with which they are already familiar.

The following VMware extensions are available as part of each OneView Advanced license:

- **HPE OneView for VMware vRealize Operations** is a development and process automation tool. It provides a predefined set of libraries and workflows that can be used to access and control third-party applications, technologies, and infrastructure. This plug-in provides health, utilization, and performance metrics in the context of the HPE hardware hierarchy so administrators can monitor critical trend changes. Its dashboards facilitate the identification of root cause problems and impacted resources across the converged infrastructure.
- **HPE OneView for VMware vCenter Log Insight** allows deep troubleshooting of an environment by analyzing unstructured data contained in iLO and Onboard Administrator (OA) logs. Information is displayed in the dashboards of VMware vCenter and vRealize Log Insight, allowing counts of critical events to be quickly identified and investigated for optimal resource utilization and rapid problem resolution.
- **HPE OneView for VMware Operations Manager** reveals critical trend changes. It includes dashboards that facilitate the identification of root cause problems and impacted resources across the data center. The OneView for VMware vRealize Operations Manager provides OneView server and infrastructure topology, status, and alerts to the vRealize Operations Manager environment.

OneView for VMware supports enhanced integration with VMware vCenter Operations Manager and vRealize Log Insight with HPE unique entitlement to use vCenter Operations Manager integration with the standard edition.

HPE OneView for VMware vCenter: Overview

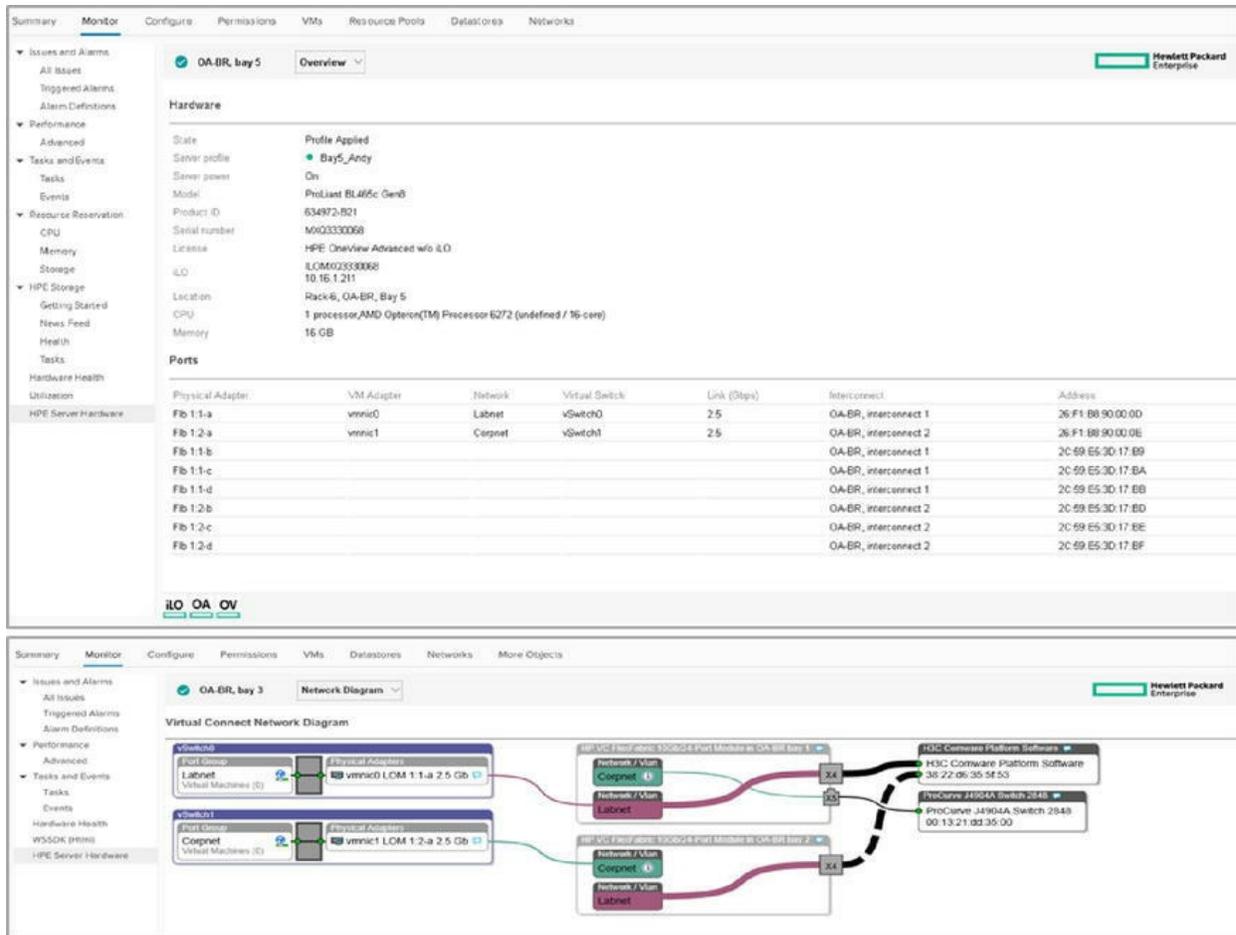


Figure 6-21 HPE OneView for VMware vCenter

User interface of HPE OneView for VMware vCenter, as shown in [Figure 6-21](#), offers the following HPE-related view and configuration sections:

- Hardware overview
- Firmware inventory
- Network ports
- Fabric diagram
- Enclosures
- Alerts
- Proactive HA
- Launch iLO, OA, HPE OneView

- Grow/edit/shrink/import cluster
- Cluster aware firmware updates
- Host and cluster consistency check and remediation

HPE server hardware: Overview

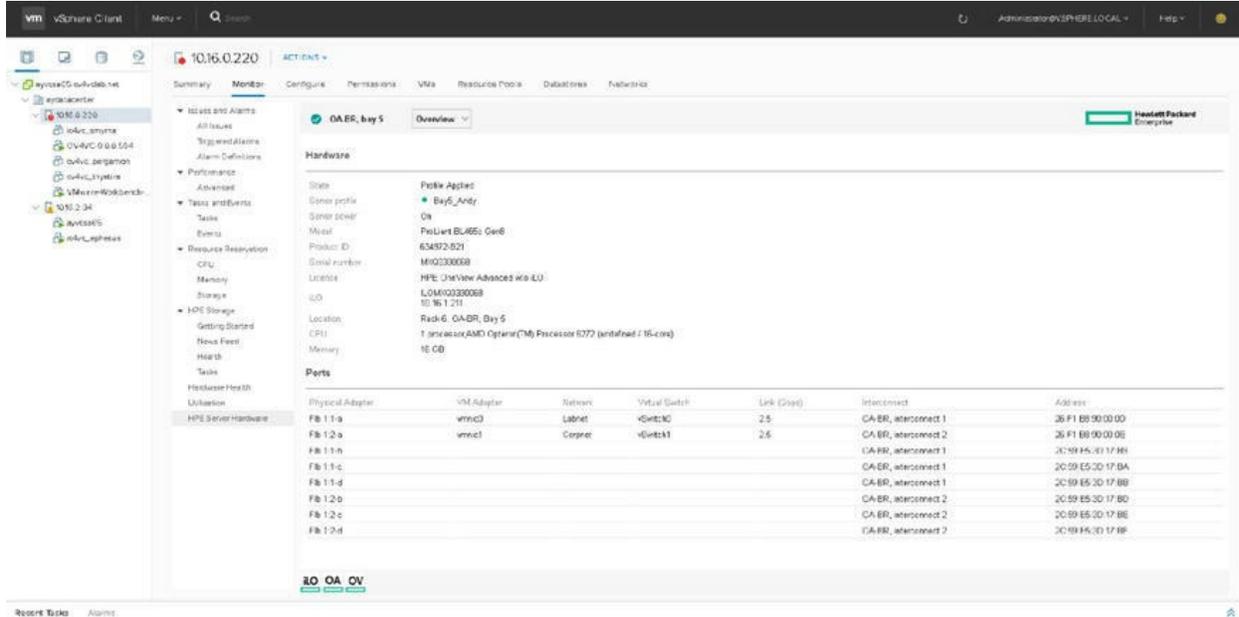


Figure 6-22 HPE Server hardware: Overview

The HPE Server Hardware menu for the Monitor tab, as shown in [Figure 6-22](#), provides information on the configured vCenters under following categories:

- **Overview**—Provides the details of the selected host and ports to which it connects. The Overview page displays the hardware and ports information for the selected host server and provides links to launch the following HPE tools:
 - Integrated Lights-Out
 - Onboard Administrator
 - Virtual Connect Manager
- **Hardware**—Provides detailed information about the selected host server

such as the health state of that server, server profile, model type, location of the server, iLO IP address, memory, serial number, product ID.

- **Firmware**—Provides the details of the firmware component installed for the selected host.
- **Network Ports**—Provides the details of the network ports IP address and virtual switch.
- **Activity**—Provides information about any active alerts in HPE OneView managed servers.
- **Network Diagram**—Displays a graphical representation of the network topology of the server. This feature requires HPE BladeSystem enclosures or Synergy frames with HPE interconnects.
- **Remote Support**—Provides support or contract information for specific vCenter host.
- **HPE OneView**—This link is displayed if the selected host is managed by HPE OneView. If a selected cluster contains hosts managed by HPE OneView, this link will be displayed at the cluster level also. If the cluster includes hosts managed by separate instances of HPE OneView, mouse-over the HPE OneView icon to display a popup menu with links to each HPE OneView instance.

When configured for vCenter administrators, auto-login is enabled for direct launch into the HPE tools listed in this section with the exception of HPE OneView, which does not support auto-login. Users with read-only access cannot access the HPE tools. Other non-administrative vCenter Server users are prompted to enter a user name and password to access the HPE tools. For DL servers, the Onboard Administrator launch link is unavailable. Links that appear to be transparent indicate that the tool is inaccessible or that no data is available.

Integration with HPE OneView Remote Support

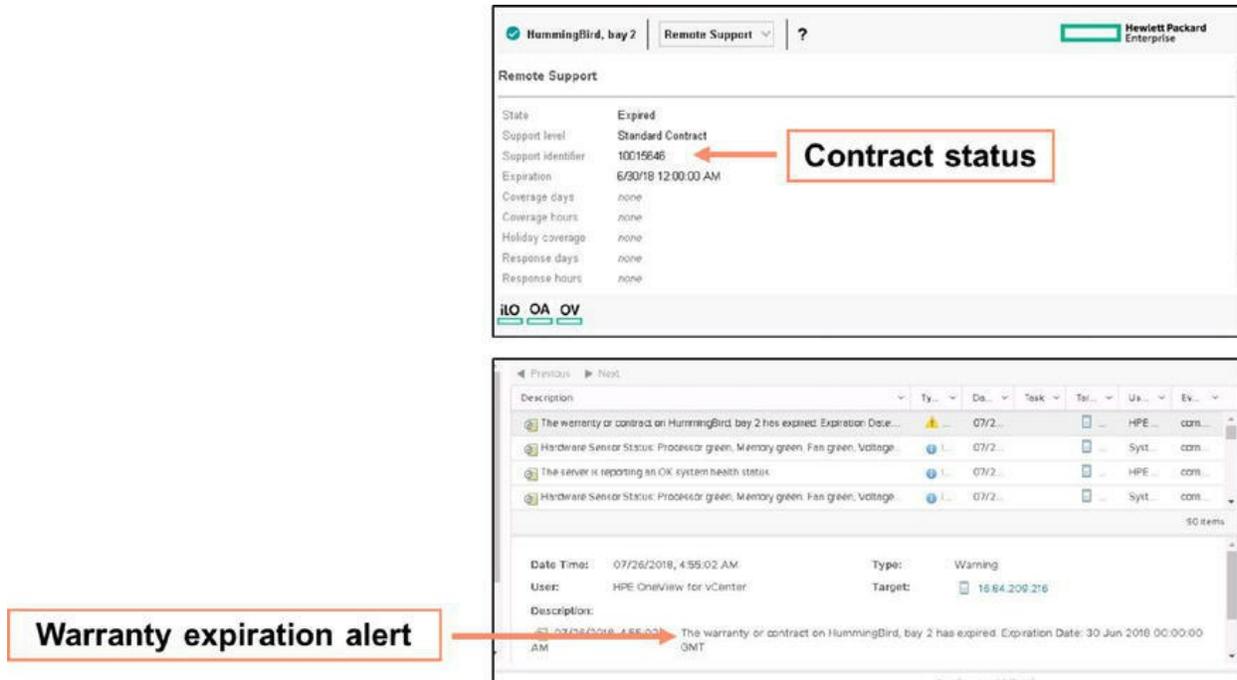


Figure 6-23 Integration with HPE OneView Remote Support

HPE OneView for VMware vCenter adds Remote Support information for the warranty expiration date for server hardware and enclosures. Two types of events are supported, as indicated in Figure 6-23:

- Support/contract about to expire
- Support/contract already expired

Current state of HPE OneView Remote Support information for enclosure and server hardware is visible in the VMware client.

HPE OneView Integrations for Microsoft System Center

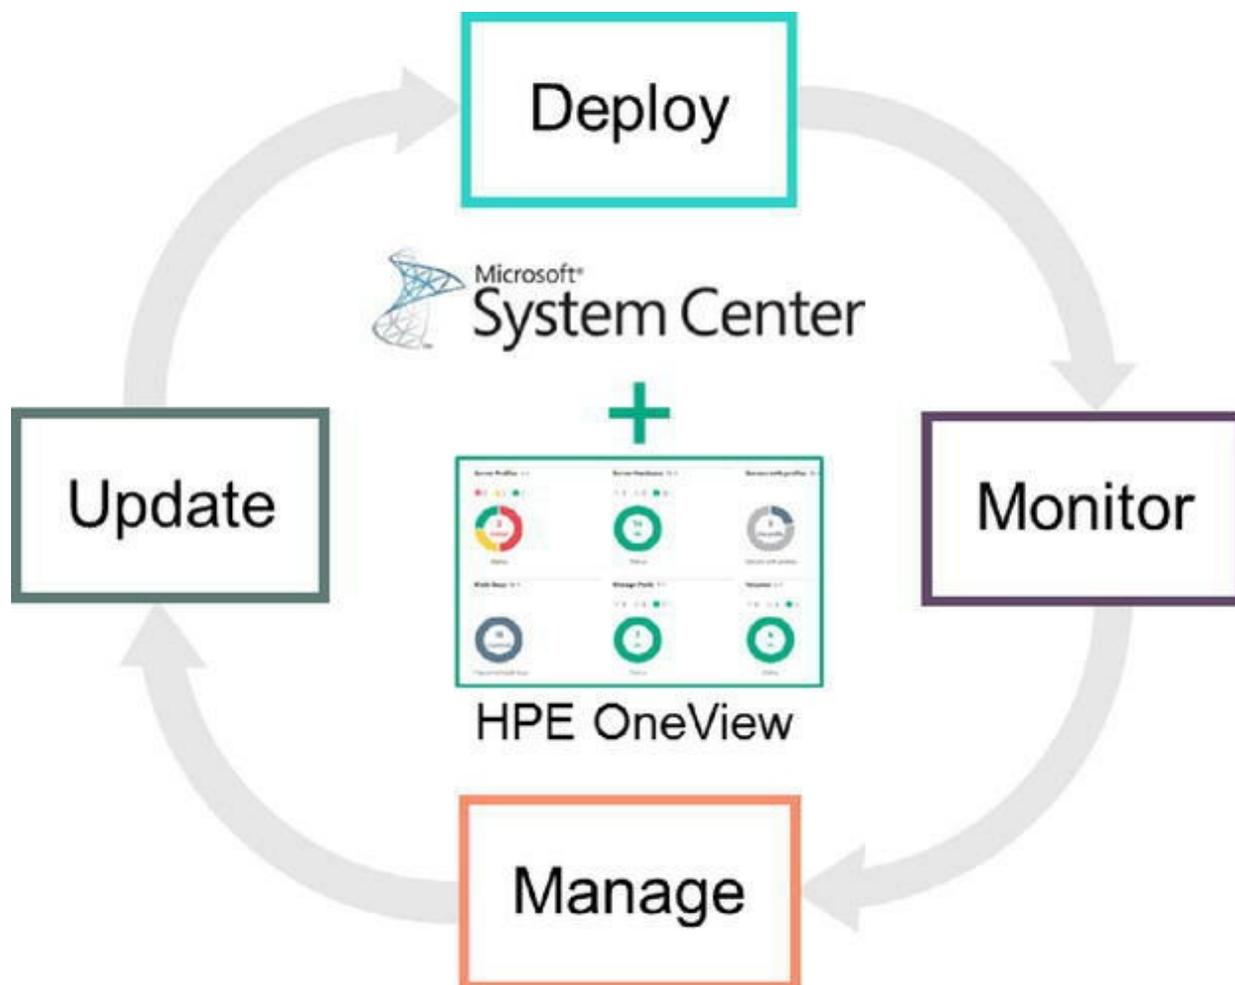


Figure 6-24 HPE OneView Integrations for Microsoft System Center

HPE OneView Integrations for Microsoft System Center, as shown in [Figure 6-24](#), brings the native manageability of HPE hardware to System Center environments. This integration is ideal for customers who have standardized on Microsoft System Center as their management platform:

- Integrates directly into System Center consoles
- Leverages the HPE OneView and HPE Synergy infrastructure automation engine
- Reduces complexity by using the same System Center tools for HPE management tasks
- Simplifies admin learning curve by using the System Center tools with which customers are already familiar

OneView for Microsoft System Center integrates ProLiant, BladeSystem, and Synergy manageability features into Microsoft System Center. This integration provides comprehensive system health and alerting, driver and firmware updates, and operating system deployment. Using this plug-in for Microsoft System Center, your customers can provision hosts using OneView server profiles to create or grow a Hyper-V cluster consistently and reliably. In addition, the health monitoring and alerting information from a OneView infrastructure perspective shows relationships clearly.

OneView for Microsoft System Center integrates with Synergy Gen10 compute nodes and Synergy Composer to provide:

- Comprehensive health monitoring and alerting for Synergy frames, compute nodes, storage modules, and interconnects via the HPE OneView Management Pack for System Center Operations Manager (SCOM).
- Fabric visualization and automated Hyper-V deployment and updates (using Synergy Composer server profile templates) occur via the HPE Fabric Management Add-in for System Center Virtual Machine Manager (SCVMM).
- Compute module firmware and Windows driver updates occur via the HPE ProLiant Updates Catalog for System Center Configuration Manager (SCCM) and SCVMM.

This integration delivers server profile-based deployment and automated HPE StoreVirtual or StoreOnce deployment for secure backup and recovery. In addition, System Center displays the network mapping, storage connections, and hardware status. OneView for Microsoft System Center includes the following components to provide a single repository of Microsoft System Center suite of products:

- **HPE OneView Storage System Management Pack**—Part of the HPE OneView SCOM Integration Kit, it integrates an HPE 3PAR Storage Management Pack and BladeSystems/Virtual Connect Management Pack. It provides a unified view to alerts/events and topological view of HPE hardware managed under OneView, enabling quick response to hardware events on HPE storage and servers running Windows and Linux, as well as BladeSystem enclosures and Virtual Connect.

- HPE Storage Management Pack for System Center**—Also part of the SCOM Integration Kit, it provides seamless integration with Microsoft System Center Operations Manager (SCOM) to enable predefined discovery and monitoring policies, event processing rules, and topology views for HPE storage.

Seamless integration into Microsoft System Center

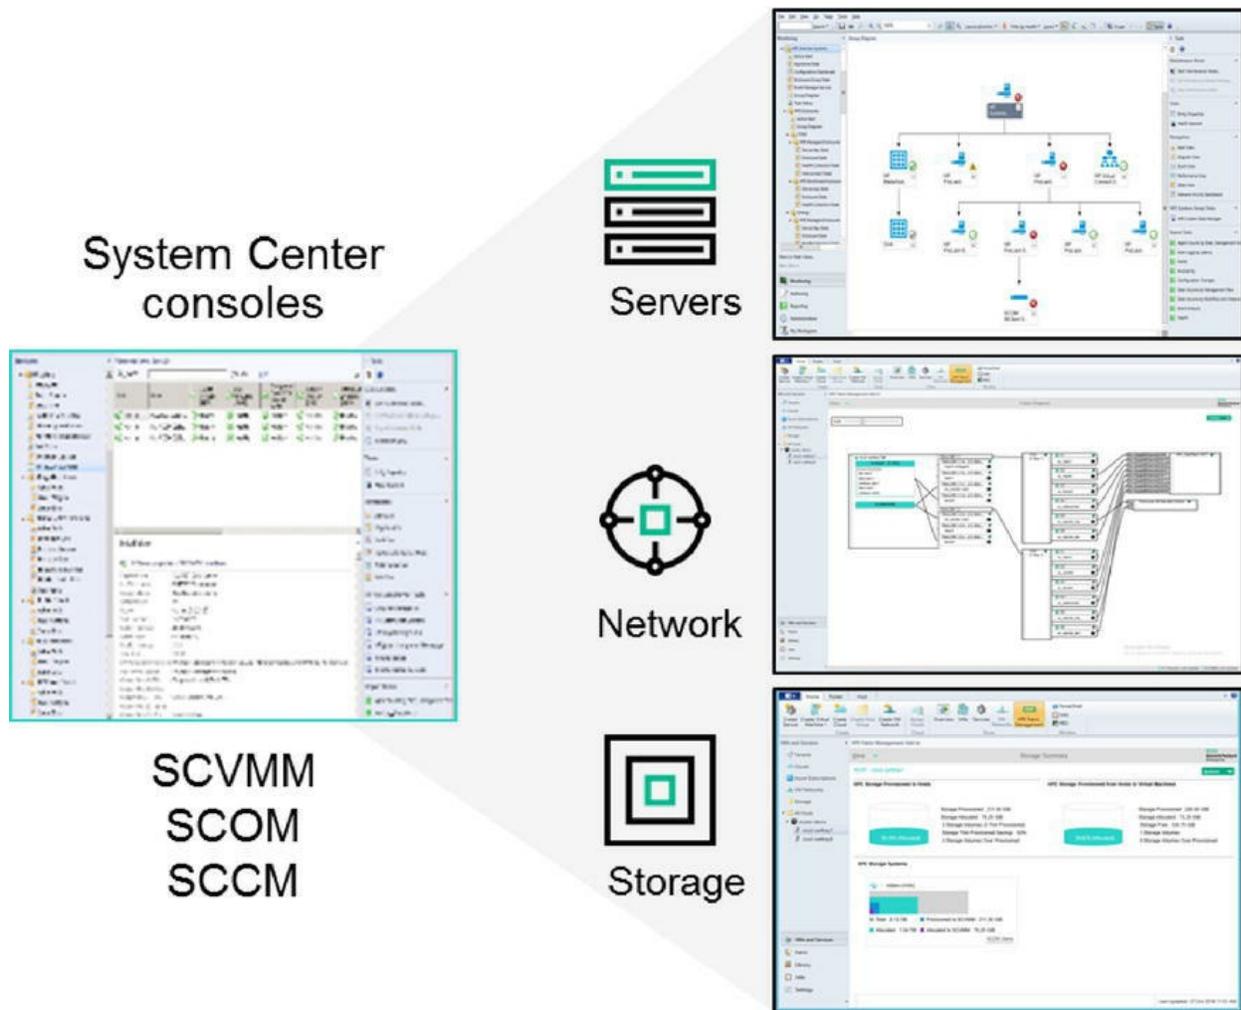


Figure 6-25 Manage the virtual and physical environment from Microsoft System Center

Other functionalities of HPE OneView for Microsoft System Center, as illustrated in [Figure 6-25](#):

- Simplify administration with single-console access to health, inventory, and configuration monitoring.
- Reduce planned and unplanned downtime with detailed resolution information for health alerts.
- Take control by launching trusted HPE management tools in context.
- Proactively manage changes with detailed insight into the relationship between the physical and virtual infrastructure, from the VM to the network edge.
- Provide consistent and repeatable bare-metal server deployment.
- Maintain stability and reliability of the environment with simplified driver and firmware updates.



Note

Health monitoring (SCOM) and driver/firmware updates (SCCM/SCVMM) are supported with HPE OneView Standard or Advanced licenses.

System Center consoles

SCVMM	OpsMgr	ConfigMgr
<ul style="list-style-type: none"> • Automates HPE storage management • Provides an integrated view of virtual machines and storage • Provides enhanced provisioning that uses profiles to grow clusters • Offers HPE fabric visualization 	<ul style="list-style-type: none"> • Manages hardware health on servers running Windows and Linux, BladeSystem enclosures, and VC • Monitors the health of servers that do not have an operating system loaded 	<ul style="list-style-type: none"> • Uses ICsp for deployment to bare metal HPE servers • Provides component-level inventory of every managed server

Figure 6-26 Microsoft System Center consoles supporting integration with HPE management features

Microsoft System Center, as illustrated in [Figure 6-26](#), is a collection of extensions that expose HPE management features within the context of the System Center consoles:

- **System Center Virtual Machine Manager (SCVMM)**
 - Automates HPE storage management and provides an integrated view of VMs and associated storage resources.
 - Provides enhanced provisioning that uses OneView profiles to create or grow Microsoft Hyper-V clusters.
 - Offers HPE fabric visualization using the HPE ProLiant Updates Catalog.
- **System Center Operations Manager (OpsMgr)**
 - Manages hardware health on servers running Windows and Linux, as well as BladeSystem enclosures and Virtual Connect.
 - Monitors the health of servers that do not have an operating system loaded, as well as ProLiant Gen8 and Gen9 servers running any operating system that has a supported Agentless Monitoring Service

(such as ESXi).

- **System Center Configuration Manager (ConfigMgr)**—Provides component-level inventory of every managed server using the HPE ProLiant Inventory Tool.

HPE OneView for Microsoft System Center



Figure 6-27 HPE OneView for Microsoft System Center brings the power of HPE OneView to System Center environments

HPE OneView for Microsoft System Center, as illustrated in [Figure 6-27](#), helps system administrators with:

- Deploying the infrastructure faster
 - Template-based provisioning that leverages the HPE OneView automation engine
 - Create or expand a Hyper-V cluster in just few simple steps
- Simplifying lifecycle operations
 - Non-disruptive, comprehensive health monitoring and alerting
 - Firmware and driver updates
 - HPE fabric visualization
- Increasing productivity
 - Consistently integrates directly into System Center
 - Use familiar System Center tools for HPE management tasks
 - Context-sensitive launch of HPE tools directly from System Center console

System Center Virtual Machine Manager integrations

Core server integrations with SCVMM include:

- OS deployment (base)
- Server firmware and driver updates via automated, rotating workflow
- No licenses required

Enhanced integration for HPE OneView include:

- HPE OneView profile-based deployment for blade servers
- End-to-end Virtual Connect networking view
- Cluster configuration view
- Expand an existing cluster
- HPE Synergy support
- Licensed as part of HPE OneView Advanced and HPE Synergy

HPE Storage integrations include the following:

- Visualize the relationship between the Hyper-V virtual machines, hosts servers, and the HPE 3PAR Storage
- Active management such as create/expand/delete volumes for HPE 3PAR Storage
- No licenses required

Learning check

4. Which four management solutions are supported by OneView integrations?

HPE OneSphere

HPE OneSphere is a software-as-a-service (SaaS)-based multi-cloud management solution. Through role-based access to HPE OneSphere, internal stakeholders (IT operations, developers, business executives) can manage hybrid clouds capable of supporting both virtualized and cloud-native applications.

HPE OneSphere is not sold as a separate product; it is a part of HPE GreenLake services.

The simplified deployment and management features of HPE OneSphere provide a cloud-like experience with on-premises infrastructure, as indicated in [Figure 6-29](#). It further allows organizations to integrate Amazon Web Services (AWS) and Microsoft Azure public cloud as well as onboard existing VMware vSphere, Kubernetes containers, and Kernel-based Virtual Machine (KVM) as private cloud resource providers. HPE OneSphere also provides core functions such as self-service access to users across lines of business, development, and IT as well as detailed usage, cost, and budgeting metrics that deliver insights across public and on-premises private clouds.

Building on our momentum to simplify hybrid cloud

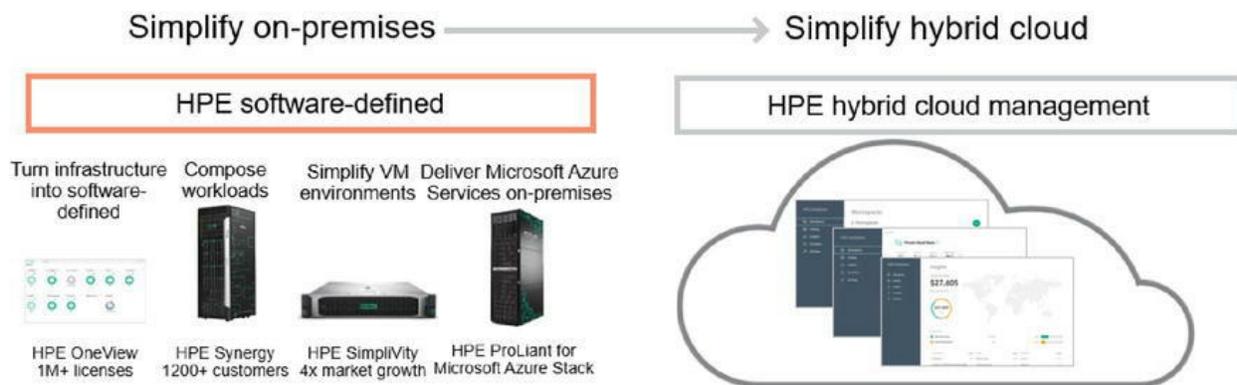


Figure 6-28 Building on our momentum to simplify hybrid cloud

Hybrid infrastructure use is pervasive—and it is resulting in complexity. The orchestration tools are maturing, yet we are still challenged without consolidated management. Moreover, apps are being developed differently—

they are cloud-native or being optimized for the cloud and span multiple clouds and sites.

Customers have complex, highly distributed environments, each operating as a silo with no way to manage it all in aggregate. IT Ops want to spend less time managing the infrastructure and more time devoted to enabling application development, and they want to eliminate operational friction. Developers want fast app deployment. The typical wait time for the handshaking between developers and IT just does not cut it. Developers want access to what they want, where and when they want it—and IT wants to provide them that agility and the business insights in real time, not at the end of a consumption or billing cycle. We need intelligence and insights to optimize workload deployments and placement, rather than getting notified at the end of the month of their investment.

HPE OneView helps our customers transform their infrastructure into a software-defined infrastructure, as illustrated in [Figure 6-28](#). It takes a software-defined, programmatic approach to managing an infrastructure with efficient workflow automation, a modern dashboard, and a comprehensive partner ecosystem. With HPE OneView, the infrastructure can be configured, monitored, updated, and repurposed with a single line of code, allowing IT teams to more effectively meet changing application needs.

HPE Synergy is a single software-defined infrastructure that combines compute, storage, and fabric so all resources are available to run any application. HPE Synergy composer leverages integrated software-defined intelligence to accelerate operations through a single interface. This allows IT to precisely compose and recompose logical infrastructures very rapidly. HPE Synergy offers a dynamic, flexible, cost-effective infrastructure that introduces efficient operations and staff productivity.

HPE SimpliVity is our hyperconverged infrastructure solution that organizations use to simplify virtualized environments. This all-flash, pre-integrated, hyperconverged building block dramatically simplifies IT by combining all hyperconverged infrastructures and advanced data services for virtualized workloads—including VM-centric management and mobility, data protection, and guaranteed data efficiency—delivered on HPE ProLiant DL380 Servers. HPE SimpliVity hyperconvergence consolidates discrete software and hardware devices and functionality in a single building block.

This reduces the cost and complexity of IT environments.

We introduced our ProLiant-based Microsoft Azure Stack appliance to enable Azure cloud on-premises. The API-compatible services running on-premises allow workloads to be deployed in either an Azure Stack private cloud or an Azure Stack public cloud without modification.

HPE ProLiant for Microsoft Azure Stack:

- Brings an integrated solution for delivering on-premises Azure-consistent services.
- Offers 50% greater VM computing capacity and more configuration flexibility than competing solutions.
- Is optimized for Azure Stack and built to scale.

HPE also recently introduced HPE OneSphere, our offering to simplify management of customers' hybrid cloud environments.

HPE OneSphere simplifies hybrid cloud environments

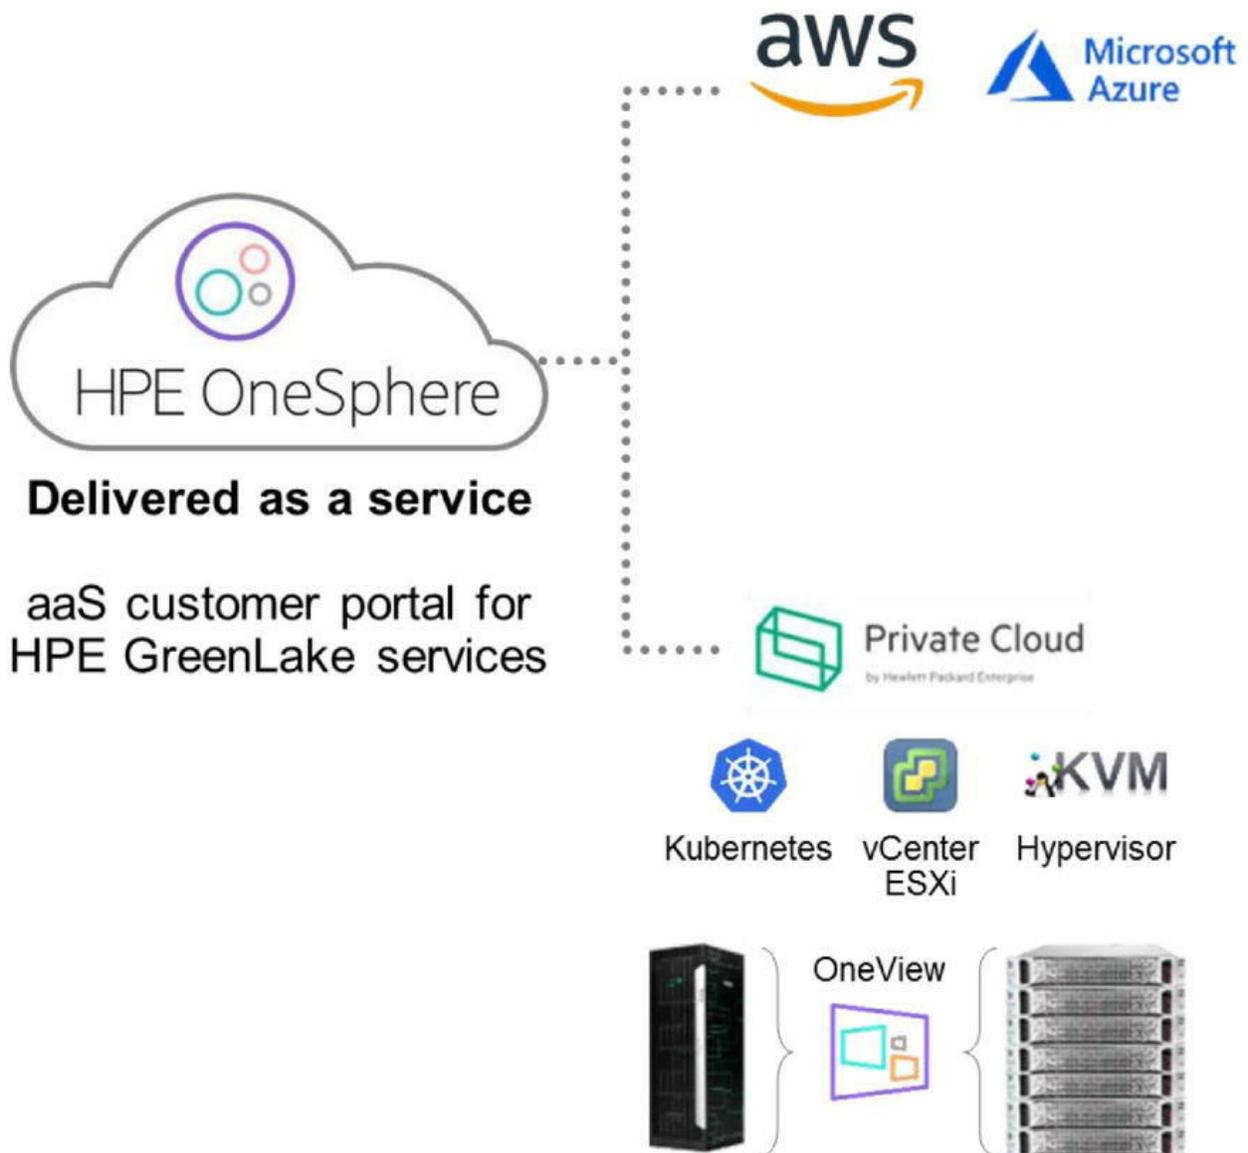


Figure 6-29 HPE OneSphere simplifies hybrid cloud environments

HPE OneSphere, as illustrated in [Figure 6-29](#), simplifies public cloud:

- Start up instantly via SaaS management and access to public cloud services.
- Build on-premises clouds and connect to public clouds in a few clicks.
- Get real-time insights on cost and utilization across clouds.
- Optimize workload deployment.
- Streamline DevOps workflow.

- Provide guardrails for the data center with governance and compliance.

HPE OneSphere simplifies **private cloud**:

- Speed resource vending.
- Automate infrastructure lifecycle management.
- Automate everything through APIs.
- Optimize on-premises utilization.
- Choose from consumption models to optimize costs.

HPE OneSphere is designed for IT, business execs, and developers to work together faster. Developers want to work unencumbered. They want to be self-sufficient and work with the tools and clouds with which they are familiar and comfortable. They also want access to services, tools, and templates that will save them time. They do not want to trip over IT when it is time to deploy. They want to consume everything as a service, so they can move faster. HPE OneSphere eliminates the need to submit tickets and wait for infrastructure. HPE OneSphere enables one or more developers to access project workspaces that have a quota of resources and a hybrid catalog, so everyone can self-service.

IT Ops want to be more strategic to the business. They want to seamlessly enable IT services to their organization—be an IT service broker. No other solution makes it this simple to transform virtual clusters, containers, and bare-metal servers into a true private cloud with multi-tenancy, a self-service portal, and a catalog.

We not only eliminate silos across on-premises and public clouds, but we also enable IT to move faster, from weeks or months to minutes. Business executives, typically the force behind digital transformation, want to know what is happening with their investments right away—and not the end of a billing cycle.

Getting accurate costs, usage, and utilization across on-premises and public clouds is complex and often a manual process—and it is even harder to accurately allocate costs to workgroups or projects.

OneSphere shows a near-real-time view of consumption insights across public and on-premises clouds, which enables business leaders to drive

resource utilization up and costs down to make everything more efficient.



Note

Verify the support for HPE OneSphere in your country.

Key features

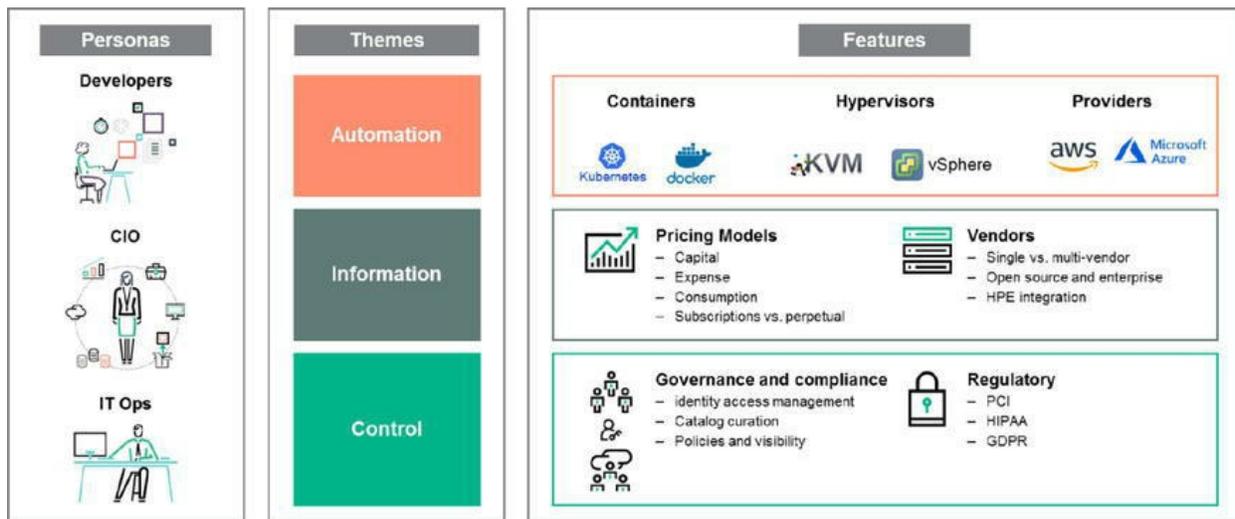


Figure 6-30 HPE OneSphere features

HPE OneSphere provides features for all the members of IT team, as indicated in [Figure 6-30](#):

- Developers
- Management/CIO
- IT Ops

Automation features of HPE OneSphere include the following:

- Is offered as SaaS Solution.
- Provides automated vCenter and KVM provisioning of on-premises virtual resources.
- Imports services and applications into a curated catalog.

- Deploys workloads from the self-service catalog and APIs.
- Provides KVM provisioning on private preconfigured infrastructure.
- Uses HPE OneView appliance and reports metrics for server resources.
- Discovers and deploys HPE OneView Server Profiles.
- Provisions production-grade Kubernetes clusters on AWS and on-premises VMware environments with a few clicks.
- Enables customers to import AWS/Azure public or private template definitions into the OneSphere catalog.
- Deploys consistent and approved templates on a per project basis for AWS CloudFormation and Azure Resource Manager (ARM).

In addition, the OneSphere admin can import public and private Docker images and Helm Charts into the Catalog.

HPE OneSphere is working with various pricing models: capital, expense, consumption, subscriptions vs perpetual. Environments with multiple vendors are also supported.

Learning check

5. Describe the HPE OneSphere offering.

HPE InfoSight: AI for hybrid cloud

HPE InfoSight is built on a unique approach to data collection and analysis. It starts first with having the right data. Data from infrastructure has always been limited to log files and graphs on obvious metrics and stuck on-

premises. InfoSight takes a much more comprehensive approach. Every second, InfoSight is collecting thousands of embedded sensors built into every array and pulling in data from across the infrastructure stack.

HPE is not just analyzing individual systems, but is analyzing all the systems in our installed base in our cloud-based AI platform, where we use machine learning to generate predictive analytics, continuously learn from every system, and automate recommendations that improve every environment.

HPE InfoSight: Improving efficiency through AI

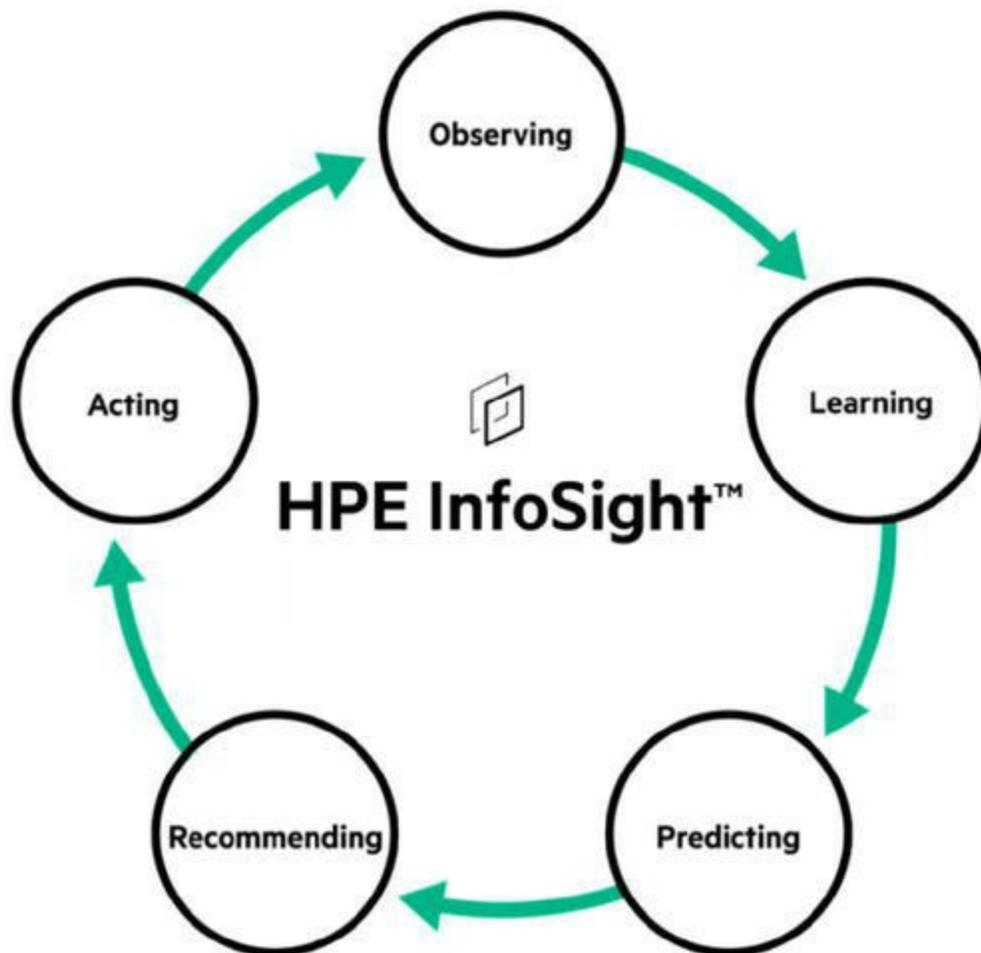


Figure 6-31 HPE InfoSight key steps

There are five key steps in the machine learning and predictive analytics process of InfoSight that are built on top of each other, as shown in [Figure 6-](#)

31:

1. **Observing:** The thousands of data points and sensors built into the HPE server and storage products
2. **Learning:** Applying advanced pattern recognition to the sensor data collected across all devices globally
3. **Predicting:** Anticipating problems based on the observations and learnings
4. **Recommending:** Intelligent decisions that prevent issues, improve performance, and optimize resources
5. **Acting:** Automation resulting in game-changing benefits and outcomes

By performing the five keys steps iteratively over time, InfoSight gets smarter every second and sees beyond the limits of other tools because of its unique approach to data collection and analysis.

HPE InfoSight—AI for hybrid cloud world

InfoSight takes a comprehensive approach to data collection and analysis. Data is the key to InfoSight's global learning. Every second, millions of sensor measurements capturing the state of the systems, subsystems, and surrounding IT infrastructure in our global installed base are collected and analyzed. InfoSight continuously learns from the telemetry and develops an understanding of the ideal operating environment for every system, workload, and application.

Analytics are only as good as the data available to analyze. Data has always been limited to one-off log files on obvious metrics or bolted on to collect data after an event—and it has always been analyzed one array at a time. We consciously architected for predictive analytics when we started. This required taking a fundamentally different approach to data collection. When we started, we began embedding sensors into every module of code in our operating system to understand all the factors that can impact the IO path.

Advanced machine learning then drives InfoSight's predictive analytics and

recommendation engines, where problematic behavior is predicted through recognition of the underlining patterns and configurations in each system.

- Predictive support automation
 - Predicts and prevents issues
 - Solves problems across infrastructure
 - Transforms the support experience
- AI-driven management
 - Makes managing effortless
 - Optimizes resource allocation
 - Sees what others cannot
- Unique product experience
 - Makes infrastructure smarter
 - Enables infrastructure to self-improve
 - Simulates work-load changes

This platform drives predictive support automation (predicting and automating resolution to problems before they can impact customers), preemptive recommendations (decisions that are automatically made for IT that prevent issues and improve every environment), proactive management (through advanced analytics in our cloud portal), and continuous improvement of our systems (making them smarter and more reliable).

The benefits have been game changing. As demonstrated with Nimble Storage:

- Predictive support enables us to deliver a radically better support experience, with 85% less time spent managing storage issues. Eighty-six percent of problems are predicted and automatically resolved before customers even realize there is an issue, where 54% are even beyond storage. This automation has enabled Nimble Storage to completely automate Level 1 and 2 support.
- AI-driven management has made managing infrastructure effortless with 79% lower storage OpEx.

- Infrastructure is no longer sitting idle and unconscious. It is self-improving. This is realized in the fact Nimble Storage has over 99.9999% of measured, proven, and guaranteed availability across its installed base, going back to the first array that ever shipped and across all OS, models, and configurations.

InfoSight has fundamentally transformed how infrastructure is managed and supported.

Traditional infrastructure was not designed with the intelligence to effectively harness the data across a hybrid cloud world. More specifically, IT departments are constantly reacting to problems, with considerable time spent on troubleshooting/doing non-value-added activities. These factors dictate that a new level of intelligent resources needs to be applied to tackle the data challenge.

HPE InfoSight brings simplified, AI-driven operations to your customer's hybrid cloud world that transforms how infrastructure is managed and supported. It uses cloud-based artificial intelligence to provide global insights into the status and health of infrastructure, all in one location.

The plan for HPE is to extend InfoSight and its AI and predictive capabilities across the portfolio. InfoSight supports many platforms, such as Nimble, 3PAR StoreServ, and ProLiant, Synergy, and Apollo servers with iLO5 and iLO4.

See once, and prevent for all



Figure 6-32 Learning from the installed base

InfoSight was founded on the belief that vendor support must evolve from a painful break-fix model because support has always been reactive.

“See once, and prevent for all” means that if we see a problem in any customer environment, we make sure no other customer will see the same problem. For any problem experienced in the installed base, InfoSight learns to predict the issue and determine if any other customer in our installed base would be susceptible to the issue using pattern matching algorithms—and these are problems that go across the stack.

The self-healing, self-managing, and self-optimizing data center is all about taking advantage of all the rich telemetry that is available across the installed base and doing something with that telemetry to make intelligent decisions, as shown in [Figure 6-32](#).

Only InfoSight goes beyond the obvious

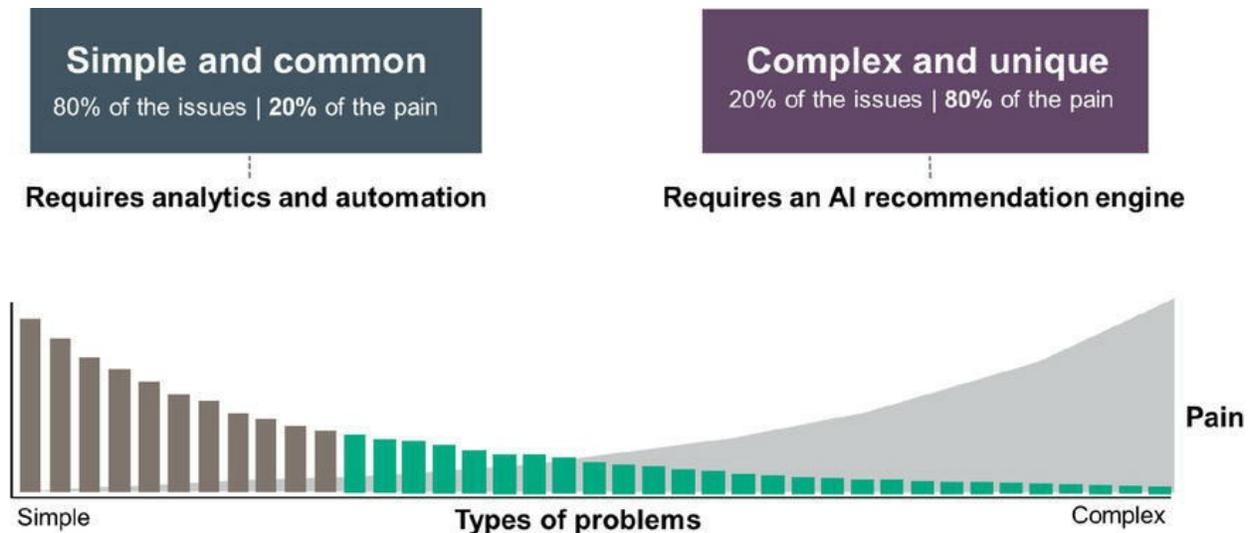


Figure 6-33 80% of issues are simple issues requiring 20% of the time when solving

Traditional troubleshooting techniques can be efficient for simple and common problems. AI can help with complex, unique problems, that are not some common, but are responsible for most of the pain connected to

troubleshooting, as illustrated in [Figure 6-33](#).

InfoSight for servers—Putting it all together

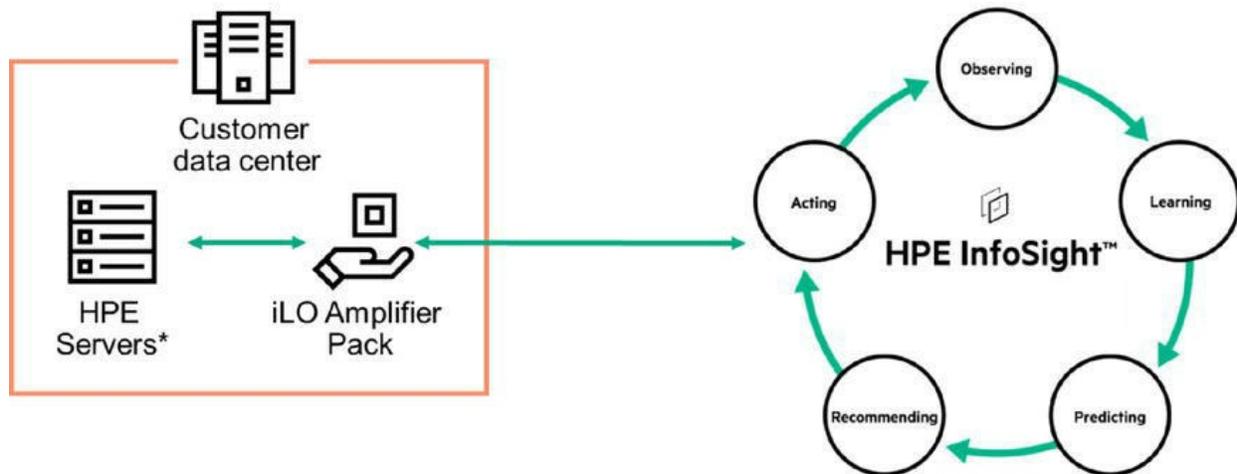


Figure 6-34 InfoSight for servers—Putting

InfoSight for Servers is the extension of InfoSight for ProLiant, Synergy, and Apollo Gen10, Gen9, and Gen8 servers with iLO 5 and iLO 4, as shown in [Figure 6-34](#).

Bringing InfoSight to servers will be a multi-step journey that starts with a basic integration with InfoSight and will become more extensive over time. The initial release of InfoSight for Servers will include:

- Predictive analytics to predict and prevent problems
 - Data analytics for server security
 - Predictive data analytics for parts failure
- Global learning that provides wellness and performance dashboards for global inventory of servers
- A recommendation engine to eliminate performance bottlenecks on servers

Monitoring servers from InfoSight is based on three major components:

- **InfoSight for Servers**—With InfoSight for Servers, HPE is combining the cloud-based machine learning of InfoSight with the health and

performance monitoring of Active Health System (AHS) and iLO to optimize performance and predict and prevent problems on Gen10, Gen9, and Gen8 ProLiant, Synergy, and Apollo servers.

- AHS
 - AHS is like a “flight recorder” for servers that provides continuous, proactive health monitoring and recording thousands of system parameters and diagnostic telemetry data 24x7 on the server.
 - AHS is available on Gen10, Gen9, and Gen8 servers with iLO 5 and iLO 4.
- iLO
 - iLO is HPE’s management processor that is embedded on the ProLiant, Apollo, and Synergy servers.
 - iLO and AHS are tightly integrated and work together to gather and record the thousands of system parameters and diagnostic telemetry data.
 - Access to AHS is through iLO (customers access the AHS data through iLO).

It is important to note that although InfoSight for Servers is new, the capturing and recording of health, configuration, and performance sensor and telemetry data at the individual server level has been around since AHS was introduced in Gen8 with iLO 4. What is new with InfoSight for Servers is the automatic collection of this data into InfoSight.

InfoSight Cross Stack Analytics for VMware environments

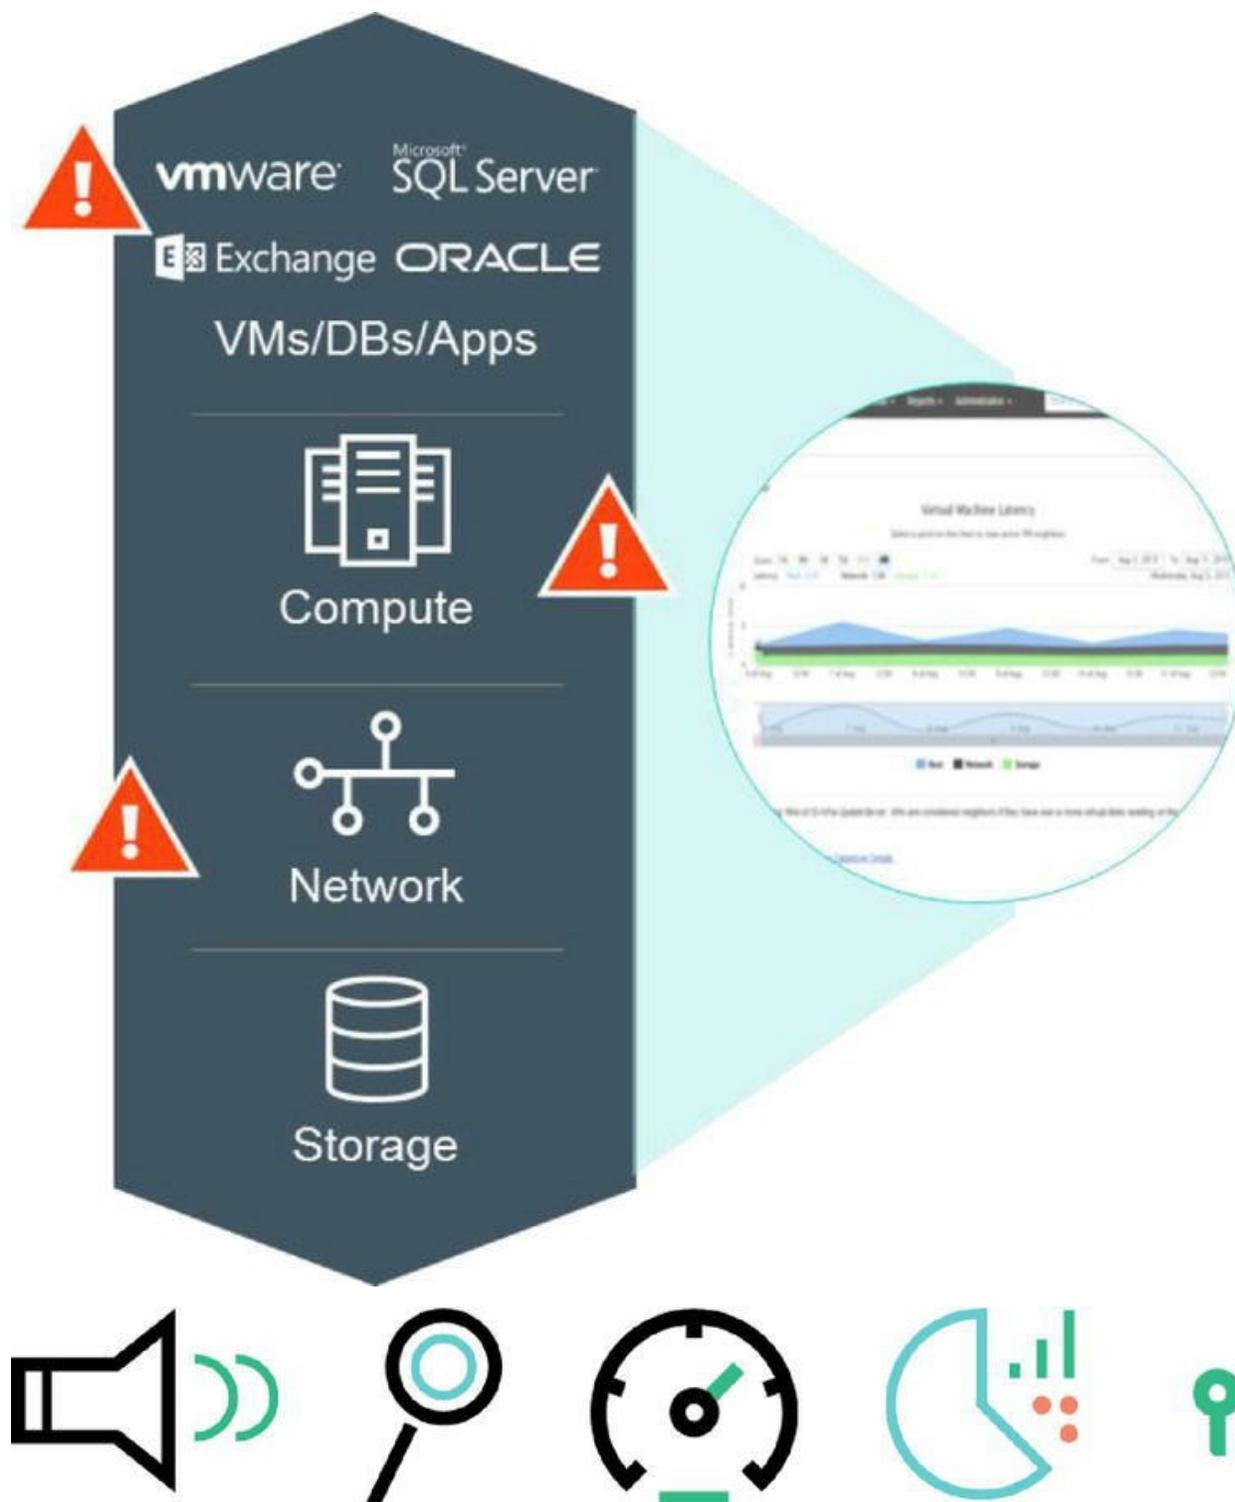


Figure 6-35 Cross Stack Analytics for VMware environments

A difficult and common problem for customers is identifying and resolving

performance bottlenecks in virtualized environment. InfoSight Cross-stack Analytics save customers' tremendous time. The solution correlates storage with VMware to quickly diagnose the root cause to performance problems, and identify if it is coming from the host, VMs, network, or array.

Cross Stack Analytics for VMware environments, as illustrated in [Figure 6-35](#), includes the following features:

- **Noisy neighbor**—Determine if VMs are hogging resources from another VM
- **Host and memory analytics**—Visibility into host CPU and memory metrics
- **Latency attribution**—Identify root cause across host, storage, or SAN
- **Inactive VMs**—Visibility into inactive VMs to repurpose/reclaim resources
- **Top performing VMs**—Visibility into Top 10 VMs by input/output operations per second (IOPS) and latency

Learning check

6. Which tool is responsible for forwarding data from HPE ProLiant servers connected to HPE InfoSight?
 - A. VMware vCenter server
 - B. iLO Amplifier Pack
 - C. HPE OneSphere
 - D. Smart Update Tools

HPE RESTful API and PowerShell cmdlets

While a GUI is available for the vast majority of Hybrid IT components in the data center, these tools are usually not usable for integration and automation.

REST communication

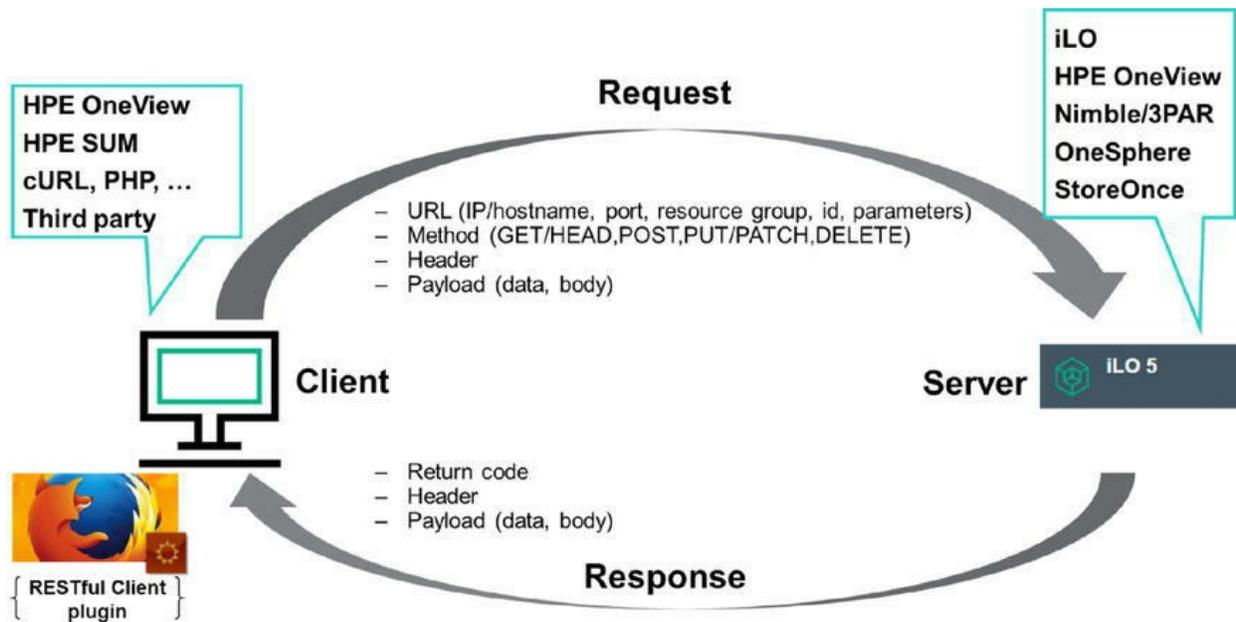


Figure 6-36 REST communication

The HPE RESTful API, as illustrated in [Figure 6-36](#), is a management interface that server management tools can use to configure, inventory, and monitor many components of the data server. It is an architectural style consisting of a coordinated set of architectural constraints applied to components, connectors, and data elements within a distributed hypermedia system.

The open, industry-standard HPE RESTful API provides a programmable interface and lightweight data model specification that is simple, remote, secure, and extensible. REST has become a popular communication protocol on service-oriented architecture styles. It enables IT staff to quickly and securely customize configurations and provisioning, and at the same time, provide a common interface for integration to many ecosystems.

REST is a web service that uses basic CRUD operations performed on resources using HTTP Post, Get, Put, Delete, and Patch. REST is an alternative to more complex programming mechanisms such as SOAP, CORBA, and RPC. Simply put, a REST call is an HTTP request to the server.

A REST client sends HTTPS operations to the server (such as iLO, HPE OneView, Nimble array, and so forth) to GET and PATCH JSON-formatted data, and to configure supported settings, such as Unified Extensible Firmware Interface (UEFI) BIOS settings.

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 RESTful API enables you to do programmatically. For example, you can use a scripting language such as Microsoft PowerShell to perform tasks by using RESTful API calls that you might otherwise complete through the web-based UI.

Today, many available tools for server management via scripting bring with them limitations around automation, orchestration, and management. Because scripting interfaces are not common across HPE management tools, HPE is using the HPE RESTful API as a standardized scripting solution to address key challenges around:

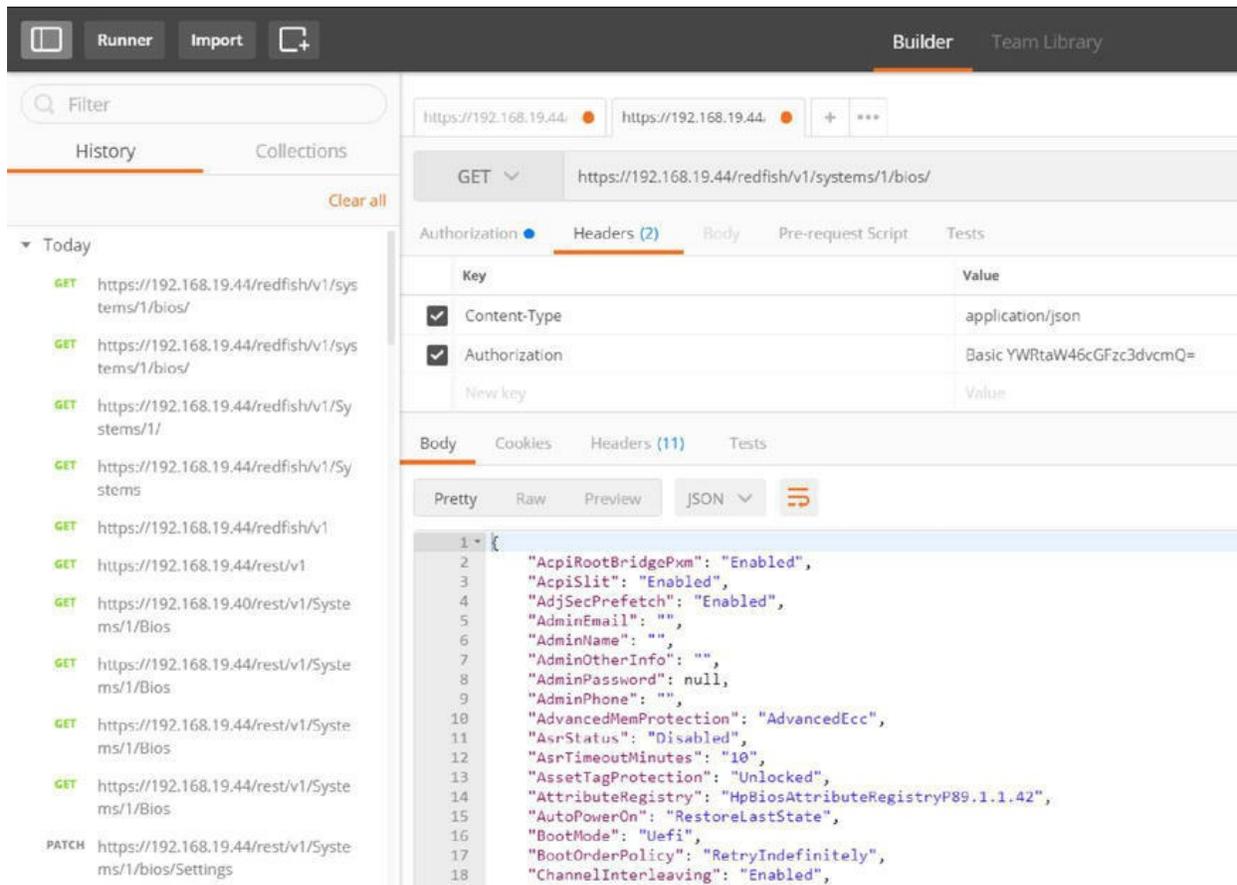
- **Unsecure remote capabilities**—Remote scripting is often not secure, triggering the need for another mechanism to transport scripts to target nodes.
- **Learning and deployment**—This can be time-consuming because a single command utility does not work across server components with existing scripting tools. The learning curve increases because administrators are required to learn different types of interfaces across the data center.
- **Scripting efficiency**—Using different tools creates complexity. Running the server through PXE for updates also delays scripting. Running scripts on too many servers is not readily scalable.

Designed for ProLiant Gen9 and Gen10 servers, the HPE RESTful API directly addresses scripting challenges in a way that is:

- **Simple**—Easier access to information eliminating multiple tools to run scripts and provision server
- **Remote and secure**—Capabilities leveraging an industry-proven HTTPS protocol
- **Extensible**—Ability to script and expose new functionalities with few or

no firmware upgrade dependencies

Example of HPE RESTful API implementation—iLO



The screenshot displays a REST client interface with a list of requests on the left and a detailed view of a GET request on the right. The GET request is to the URL `https://192.168.19.44/redfish/v1/systems/1/bios/`. The headers section shows `Content-Type: application/json` and `Authorization: Basic YWFRtaW46cGFzc3dvcmQ=`. The response body is a JSON object with the following structure:

```
1 {
2   "AcpiRootBridgePxm": "Enabled",
3   "AcpiSlit": "Enabled",
4   "AdjSecPrefetch": "Enabled",
5   "AdminEmail": "",
6   "AdminName": "",
7   "AdminOtherInfo": "",
8   "AdminPassword": null,
9   "AdminPhone": "",
10  "AdvancedMemProtection": "AdvancedEcc",
11  "AsrStatus": "Disabled",
12  "AsrTimeoutMinutes": "10",
13  "AssetTagProtection": "Unlocked",
14  "AttributeRegistry": "HpBiosAttributeRegistryP89.1.1.42",
15  "AutoPowerOn": "RestoreLastState",
16  "BootMode": "Uefi",
17  "BootOrderPolicy": "RetryIndefinitely",
18  "ChannelInterleaving": "Enabled",
```

Figure 6-37 HPE RESTful API implementation example

The HPE RESTful API for iLO is designed using the Hypermedia as the Engine of Application State (HATEOAS) REST architecture. This architecture allows the client to interact with iLO through a simple fixed URL—`rest/v1`. This has the advantage of the client not needing to know a set of fixed URLs. When creating a script to automate tasks using the HPE RESTful API for iLO, your customer only needs to hardcode this simple URL, and design the script to discover the RESTful API URLs that are needed to complete a task.

**Note**

For more information about HATEOAS, visit:
<http://en.wikipedia.org/wiki/HATEOAS>

The HPE RESTful API for iLO, as shown in [Figure 6-37](#), is the main management API for iLO-based HPE servers. Using this API, it is possible to take a full inventory of the server, control power and reset, configure BIOS and iLO settings, and fetch event logs, in addition to performing many other functions.

This API follows the internet trend in moving to a common pattern for new software interfaces. Many web services in a variety of industries use RESTful APIs because they are easy to implement and easy to consume, and they offer scalability advantages over previous technologies. HPE OneView, OpenStack, and many other server management APIs are now RESTful APIs. Most HPE management software offerings, as well as the entire SDDC architecture, are built upon RESTful APIs.

The HPE RESTful API for iLO has the additional advantage of consistency across all present and projected server architectures. The same data model works for traditional rack-mount servers and blades as well as newer types of systems such as HPE Moonshot. The HPE RESTful API for iLO provides this advantage because the data model is designed to self-describe the service's capabilities to the client and has room for flexibility designed in from the start.

**Note**

For more information on the HPE RESTful API, visit:
<https://hewlettpackard.github.io/ilo-rest-api-docs/ilo5/>

HPE RESTful Interface Tool

```
C:\Program Files\Hewlett Packard Enterprise\RESTful Interface Tool>ilorest
ilorest : RESTful Interface Tool version 2.5
Copyright (c) 2014, 2019 Hewlett Packard Enterprise Development LP
-----
ilorest > login 192.168.19.46 -u admin
Password:
Discovering data...Done
ilorest > bootorder

Current Persistent Boot Order:
1. HD.EmbrAID.1.3 (Embedded RAID 1 : HPE Smart Array P408i-a SR Gen10 - 279.3 GiB, RAID1 Logical Drive 2(Target:0, Lun:1
))
2. HD.EmbrAID.3.1 (Embedded RAID 1 : HPE Smart Array P408i-a SR Gen10 - 139.7 GiB, RAID0 Logical Drive 1(Target:0, Lun:0
))
3. NIC.LOM.1.1.Httpv4 (Embedded LOM 1 Port 1 : HPE Ethernet 1Gb 4-port 331i Adapter - NIC (HTTP(S) IPv4))
4. NIC.LOM.1.1.IPv4 (Embedded LOM 1 Port 1 : HPE Ethernet 1Gb 4-port 331i Adapter - NIC (PXE IPv4))
5. NIC.LOM.1.1.Httpv6 (Embedded LOM 1 Port 1 : HPE Ethernet 1Gb 4-port 331i Adapter - NIC (HTTP(S) IPv6))
6. NIC.LOM.1.1.IPv6 (Embedded LOM 1 Port 1 : HPE Ethernet 1Gb 4-port 331i Adapter - NIC (PXE IPv6))
7. Generic.USB.1.1 (Generic USB Boot)
8. HD.SD.1.2 (Internal SD Card 1 : Generic USB3.0-CRW)
9. HD.EmbrAID.1.2 (ESXi EFI DISK 1)

Continuous and one time boot options:
1. None
2. Cd
3. Hdd
4. Usb
5. SDCard
6. Utilities
7. Diags
8. BiosSetup
9. Pxe
10. UefiShell
11. UefiHttp
12. UefiTarget

Continuous and one time boot uefi options:
```

Figure 6-38 HPE RESTful Interface Tool version 2.0

The HPE RESTful Interface Tool simplifies server configuration by using industry-recognized RESTful APIs, enabling your customers to script provisioning on ProLiant Gen9 and Gen10 servers. The RESTful Interface Tool, as shown in [Figure 6-38](#), offers a single command-line interface (CLI) to configure various server components, plus document-allowed server configurations and dependencies. The HPE RESTful API for UEFI can be used to configure settings through this CLI.

The RESTful Interface Tool is key to enabling software-defined computing. Benefits include:

- **Easy customization**—A single command line to simplify customizing workflows and scripts by standardizing a set of commands that interacts with all server components
- **Reduced travel costs**—Capability to remotely manage servers
- **Reduced deployment complexity**—Enablement of any of the three modes—interactive, scriptable, or file-based—to program and execute scripts easily
- **Simplified scripting**—Self-descriptive tool to reduce the learning curve

adoption

Using PowerShell to manage HPE devices

Windows PowerShell provides a powerful set of utilities that you can use to perform various configuration tasks on Hybrid IT data center components. Windows PowerShell cmdlets for managing HPE devices are designed for customers familiar with Windows PowerShell. It is the Microsoft task automation framework, consisting of a command-line shell and associated scripting language built on Microsoft .NET Framework. It follows the standard PowerShell syntax and scripting model, making it easy for customers to incorporate these functions into their administrative scripts.

The HPE Scripting Tools for Windows PowerShell uses lightweight commands (cmdlets) that better enable integration with the current IT ecosystem, allowing retrieval of firmware versions from multiple enclosures and servers. PowerShell Onboard Administrator cmdlets enable retrieval of firmware versions from multiple BladeSystem enclosures and servers, and pipe information to cmdlets that update enclosures, blade, and enclosure options.

Example of PowerShell cmdlets—HPE iLO with PowerShell

```
PS D:\> Get-Command -Module HPEiLOcmdlets
-----
CommandType      Name
-----
Cmdlet           Add-HPEiLODirectoryGroup
Cmdlet           Add-HPEiLOFederationGroup
Cmdlet           Add-HPEiLOInstallationQueueWaitTask
Cmdlet           Add-HPEiLOInstallSet
Cmdlet           Add-HPEiLOLanguagePack
Cmdlet           Add-HPEiLOMaintenanceWindow
Cmdlet           Add-HPEiLORepositoryComponent
Cmdlet           Add-HPEiLOSMPAlertDestination
Cmdlet           Add-HPEiLOSMPv3User
Cmdlet           Add-HPEiLOSSORRecord
Cmdlet           Add-HPEiLOUser
Cmdlet           Backup-HPEiLOSetting
Cmdlet           Clear-HPEiLOAHSDATA
Cmdlet           Clear-HPEiLOESKMLLog
Cmdlet           Clear-HPEiLOEventLog
Cmdlet           Clear-HPEiLOHotKeyConfig
Cmdlet           Clear-HPEiLOIML
Cmdlet           Clear-HPEiLOInstallationTaskQueue
Cmdlet           Clear-HPEiLOInstallSet
Cmdlet           Clear-HPEiLOLog
Cmdlet           Clear-HPEiLOMaintenanceWindow
Cmdlet           Clear-HPEiLOPowerOnTime

PS D:\> Find-HPEiLO 192.168.19.40-49

IP                : 192.168.19.40
Hostname          : ilo-d1380pgen8.zeta.local
SPN               : ProLiant DL380p Gen8
FWRI              : 2.7
PN                : Integrated Lights-Out 4 (iLO 4)
SerialNumber      : C2221805MX
CUUID             : 31323436-3032-5A43-3232-313830354058

IP                : 192.168.19.44
Hostname          : ilo-d1360gen9.zeta.local
SPN               : ProLiant DL360 Gen9
FWRI              : 2.7
PN                : Integrated Lights-Out 4 (iLO 4)
SerialNumber      : CZJ44308RD
CUUID             : 34343737-3733-5A43-4A34-343330385244

IP                : 192.168.19.45
Hostname          : d1580g7
SPN               : ProLiant DL580 G7
FWRI              : 1.91
PN                : Integrated Lights-Out 3 (iLO 3)
SerialNumber      : CZJ2491SKK
CUUID             : 30333436-3638-5A43-4A32-343931534848

IP                : 192.168.19.46
Hostname          : ilo-d1360gen10.zeta.local
SPN               : ProLiant DL360 Gen10
FWRI              : 1.4
PN                : Integrated Lights-Out 5 (iLO 5)
SerialNumber      : CZJ7510GV3
CUUID             : 31363738-3030-5A43-4A37-353130475633
```

Figure 6-39 Example of PowerShell cmdlets—HPE iLO with PowerShell

PowerShell is the standard environment in Windows and uses simple syntax: VERB-NOUN or ACTION-ResourceName. It is also available for Linux and as a container. PowerShell allows your customers to extend the language by importing custom modules and adding more CmdLets, as shown in [Figure 6-39](#).



Note

Many modules are available at the Microsoft PowerShell Gallery: <https://www.powershellgallery.com/>. HPE is shipping several modules for PowerShell, including HPE OneView, iLO, Storage, OneSphere, BIOS, and Smart Array.

Learning check

7. What are the typical requirements for a REST API request to iLO?

Customer scenario

The wide HPE portfolio of data center products for SMB customers can be introduced using customer scenarios.

Introducing the customer scenario

The fictional customer LLP distribution company will be used as a storyline through this scenario. We will introduce the company using an interview:

- What is your primary business?
 - City-wide courier, parcel, and documents delivery
- How many employees do you currently have?
 - 25 employees + external workers
- How does your selling and delivery channel look like?
 - Contracts with dozens of local e-shops and expanding
- What does your server, storage, and network infrastructure look like?
 - ProLiant, StoreEasy, HPE Networking, Nimble
- Do you have an IT department?
 - No, external help on request
- What are your current plans?
 - 24x7 monitoring of HPE devices

Customer requirements

As a result of multiple interviews and gathering information about customer plans and customer's current infrastructure, the following requirements emerged for the new solution:

- 24x7 monitoring
- Health and inventory status
- Automated response
- Email notification
- Devices:
 - HPE ProLiant ML servers
 - HPE StoreEasy
 - HPE Nimble storage
 - Aruba 2920 and 3810M

Summary

- Remote IT Support services reduce unplanned downtime with automated, remote monitoring of servers, storage, and networking technologies.
- HPE OneView is the foundation solution for the software-defined data center (SDDC) that helps drive the business and enable innovation. OneView offers plugins for several virtualization management products.
- HPE OneSphere is a software-as-a-service (SaaS)-based multi-cloud management solution.
- InfoSight takes a comprehensive approach to data collection and analysis using AI.
- The HPE RESTful API is a management interface that server management tools can use to configure, inventory, and monitor many components of the data server.

Learning check answers

Chapter 1

1. Which HPE server platform is used in SMB offers?
 - **HPE ProLiant Servers**
2. You are in a meeting with a customer, and they challenge you with a question: “What are the most important barriers of transforming IT environment for an edge-centric, cloud-enabled, data-driven world?”
How should you respond?
 - **Technology**
 - **People**
 - **Economics**
3. How does HPE make Hybrid IT simple?
 - **Define the Hybrid IT strategy**
 - **Power Hybrid IT platforms**
 - **Optimize Hybrid IT consumption and delivery**

Chapter 2

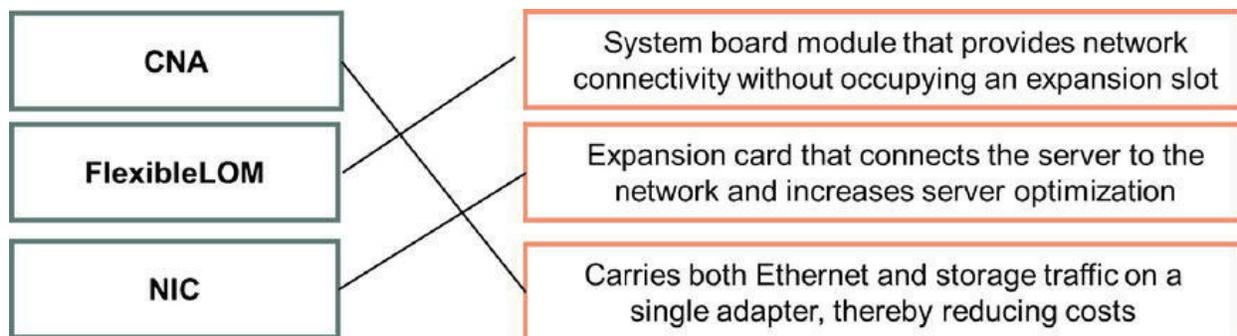
1. What are typical storage considerations when assessing customer requirements?
 - **RAID or non-RAID configurations**
 - **Storage availability**

- **Capacity**
 - **Performance**
 - **Fault tolerance**
2. A customer requires a tower server for environment without a server room with redundant, hot-plug fans. Which server supports these requirements?
- A. **ML350**
 - B. ML110
 - C. ML30
 - D. DL380
3. Which processor will support a 4-socket system?
- A. Silver 4208 Processor
 - B. **Gold 5215L Processor**
 - C. **Gold 6240Y Processor**
 - D. **Platinum 8280 Processor**
 - E. Pentium Gold G5400
4. What are the requirements for using HPE Persistent Memory featuring Intel[®] Optane[™]?
- **Supported HPE ProLiant server**
 - **Supported Intel[®] Gen2 processor**
 - **Can only be mixed with either RDIMMs or LRDIMMs**
 - **Maximum of 12 HPE Persistent Memory DIMMs**
 - **Configurations exceeding 1 TB/socket, the “M” series (2TB/socket) or “L” series processors (4.5 TB/socket) are required**
5. Name at least 3 advantages of hardware-based RAID over software-based

RAID.

- **Faster, more reliable performance**
- **Decreased processor load**
- **User-friendly configuration utilities**
- **No operating system interface for rebuild**
- **Safe write caching**

6. Match the server networking component with its description.



7. What is the maximum inlet ambient temperature for systems compliant with ASHRAE Class A4?
- A. 30°C
 - B. 35°C
 - C. 40°C
 - D. **45°C**
8. Recommend a single socket server providing up to 32 cores and 4 TB in a 1U rack form.
- **DL325 Gen10**
9. What are the typical components of an HPE ProLiant for Microsoft Azure Stack solution?
- **TOR and management switches (typically FlexFabric)**

- **Management and compute nodes (DL360 Gen10, DL380 Gen10)**
 - **Racking and power options (for example, PDU, KVM switch, rackmount console)**
 - **Azure Stack support**
 - **Deployment service**
10. Name at least 3 criteria that the SSD Selector Tool can use to filter down to the correct SSD drive your customer.
- **Type of the workload**
 - **SSD type (read intensive, mixed use, write intensive)**
 - **Type of the server and the server model**
 - **Capacity**
 - **SSD Interface (SAS, SATA, NVMe)**
 - **Form factor (SFF, LFF, M.2, add-in card, ...)**
 - **Category (recommended vs. extended)**
11. You are in a meeting with a customer, and they challenge you with a statement: “I heard that the configuration process required to prepare a server configuration quote is time consuming, with a high of risk of errors and risk of ordering incompatible components.”
- How should you respond?
- **HPE provides a range of tools and resources for simplifying the configuration and quoting process. These include Reference Architectures, HPE Switch Selector, HPE Proposal Web, HPE OneConfig Advanced, HPE Power Advisor, HPE Synergy Planning Tool, HPE Server Memory Configurator, SSD Selector Tool, HPE iQuote Universal, and a range of HPE solution sizers.**
12. Name the available response levels with HPE Foundation Care services.
- **HPE Foundation Care Call to Repair (CTR) Service**
 - **HPE Foundation Care 24x7 Service**

- **HPE Foundation Care Next Business Day (NBD) Service**
- **HPE Foundation Care 4 Hour Exchange Service**
- **HPE Foundation Care Next Business Day Exchange Service**

13. Name 3 available tools to configure local storage in an HPE ProLiant server.

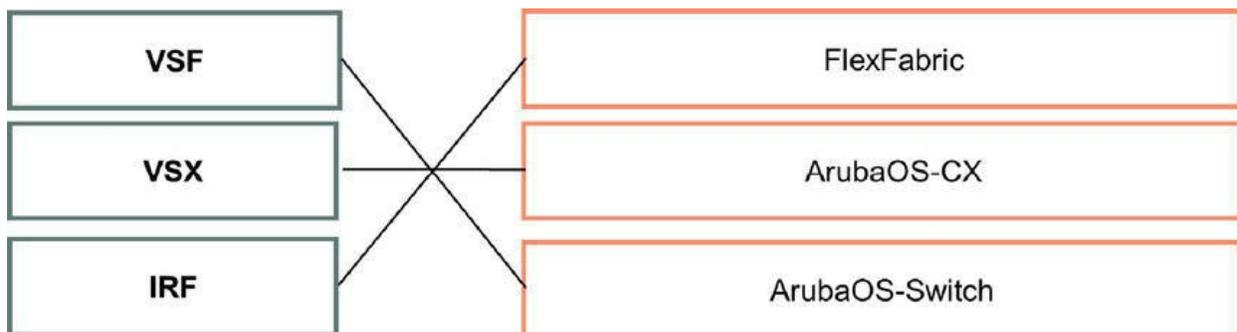
- **Smart Storage Administrator GUI**
- **Smart Storage Administrator CLI**
- **UEFI**
- **REST API**
- **PowerShell cmdlets for Array Controller**
- **“Always On” IP started from iLO 5 (Intelligent Storage Configuration)**

Chapter 3

1. Recommend an HPE networking platform for an out-of-band management network.

- **ArubaOS-Switch**
- **FlexNetwork**

2. Match the virtualization and HA solution with the networking platform.



3. Which ArubaOS-Switch series supports IPv4 BGP?
 - A. Aruba 2530
 - B. Aruba 2540
 - C. Aruba 2930F
 - D. **Aruba 5400R zl2**

4. Recommend a ToR switch with 25 GbE downlinks.
 - A. **Aruba 8325 Switch Series**
 - B. Aruba 8320 32p
 - C. Aruba 8400
 - D. Aruba 8320 48p

5. Which transceiver type supports 100 Gb?
 - A. **QSFP28**
 - B. QSFP+
 - C. SFP28
 - D. SFP

6. How many network devices can be managed by the HPE Intelligent Management Center Basic Software Platform?
 - **50**

Chapter 4

1. Recommend a NAS platform for a customer currently using an HPE 3PAR StoreServ Array.
 - **HPE Storage File Controller**

2. Name 3 typical use cases for HPE StoreEasy platform.
 - **User, department, and corporate file shares**

- **Home directory consolidation**
 - **File sync and share for desktop and mobile**
 - **Connected ROBO**
 - **Bulk data preservation of unstructured file data**
 - **Compliant archives using iTernity Compliant Archive Software**
 - **Dedicated backup system for Microsoft apps, Hyper-V VMs, and Windows PCs**
 - **Simple data store for SQL Server and Hyper-V over SMB or iSCSI**
3. Which technology allows efficient replication of folders in an HPE StoreEasy solution?
- A. DFS-N
 - B. NTFS
 - C. **DFS-R**
 - D. ReFS
4. How can security be integrated at the file level in HPE StoreEasy?
- **A lockable bezel**
 - **Drive encryption**
 - **File system encryption, file classification, and dynamic access**
 - **Endpoint assurance and encryption**
 - **Automatically enabled anti-malware**
5. You are in a meeting with a customer, and they challenge you with a statement: “I heard that HPE StoreEasy products are only supported in a Windows environment.”
- How should you respond?
- **HPE StoreEasy platform offers extensive protocol support, including SMB, NFS, iSCSI, WebDAV, HTTP/HTTPS,**

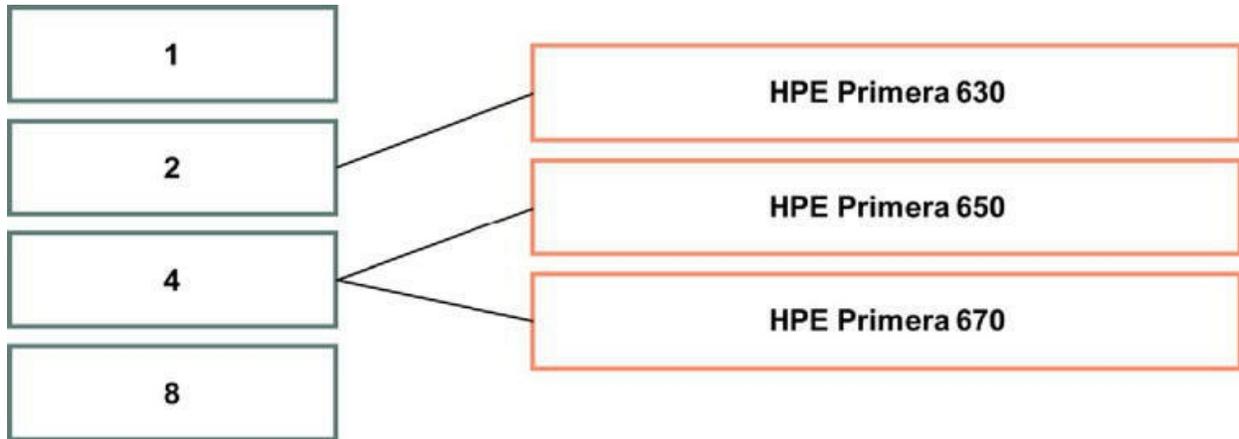
FTP/FTPS, that can easily be used in non-Windows environment.

6. Describe the SMB 3 multichannel features.
 - **Enables SMB clients to automatically identify all the network interfaces on a target server and to simultaneously use all the network interfaces on the client that can connect to those on the server.**

Chapter 5

1. What is the difference between D3610 and D3710?
 - **D3710 support SFF drives**
 - **D3610 support LFF drives**
2. Name at least 3 typical use cases for SAN.
 - **High availability**
 - **Better utilization of backup and restore solutions**
 - **Business continuance**
 - **Server and storage consolidation**
3. How many controllers does AF20Q have?
 - A. 0 (external controllers are needed)
 - B. 1
 - C. 2
 - D. 4
4. When will you recommend HPE SimpliVity 380 H over HPE SimpliVity 380 G?
 - **HPE SimpliVity 380 H is optimized as a node for backup and recovery.**

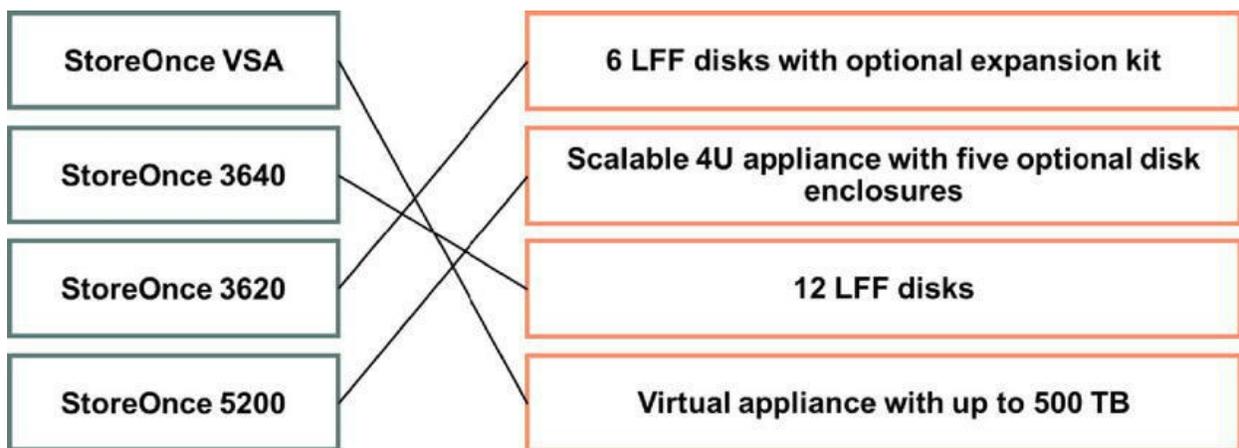
- What is the name of the MSA technology allowing an administrator to set tier for a volume?
5. tier for a volume?
 - **Volume tier affinity**
 6. Describe the read operation at HPE Nimble HF Storage.
 - **Read from NVDIMM.**
 - **If not found, check DRAM.**
 - **If not found, read from SSD.**
 - **If found, validate checksum, decompress, and return data.**
 - **If not found, read from disk using the index on the SSD for a quick lookup.**
 - **If found, validate checksum, decompress, and return data.**
 - **If cache worthy, write to cache.**
 7. What are the typical building blocks of the HPE 3PAR StoreServ 8000 solution?
 - **Base enclosure**
 - **Host adapters**
 - **Expansion drive enclosures**
 - **Drives**
 - **Rack**
 - **Service processor**
 8. Match the HPE Primera model with its number of controllers.



9. Describe the difference between an incremental and a differential backup.

- **Incremental backup**—Backup that selects only files that have changed since a previous full or incremental backup
- **Differential backup**—Incremental backup that backs up changes made since the last full backup

10. Match the HPE StoreOnce model with its description.



11. Which storage platforms can be configured with NinjaSTARS?

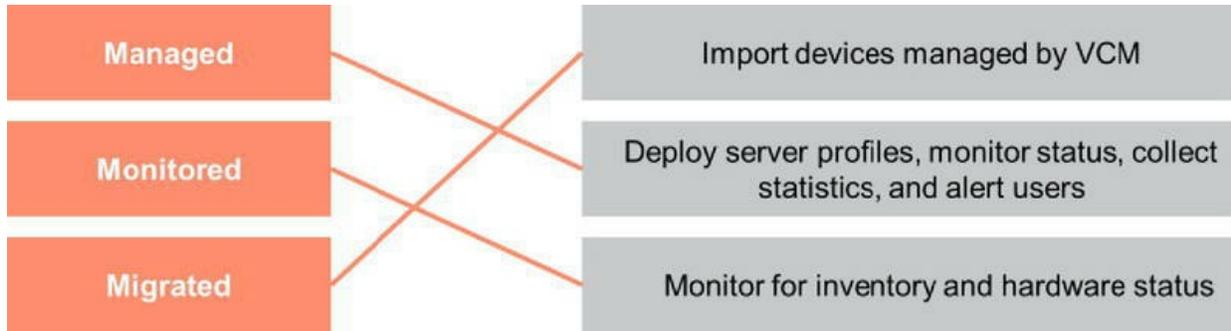
- **NinjaSTARS allows configuring and sizing of 3PAR, Nimble, Primera, and StoreOnce systems.**

12. Which tool can be used to create Catalyst copies?

- A. HPE 3PAR StoreServ Management Console
- B. HPE StoreOnce Recovery Manager Central**
- C. HPE StoreOnce Management Console
- D. HPE InfoSight

Chapter 6

1. Name 3 HPE offerings for single system management.
 - **UEFI**
 - **iLO 5**
 - **RESTful API, HPE RESTful Interface Tool, and other HPE scripting tools**
 - **Intelligent Provisioning**
 - **Smart Storage Administrator**
 - **HPE Smart Update and Service Pack for ProLiant**
 - **SSH/Telnet**
 - **3PAR CLI**
 - **StoreOnce Management Console**
 - **StoreEasy Management Console**
 - **HPE Nimble Storage GUI**
2. Which HPE ProLiant server generations are supported for direct connection to Remote Support?
 - **Gen9**
 - **Gen10**
3. Match the way of adding hardware to OneView with its description.



4. Which 4 management solutions are supported by OneView integrations?
 - **VMware vCenter server**
 - **VMware vRealize Log Insight**
 - **Microsoft System Center Virtual Machine Manager**
 - **Microsoft System Center Operations Manager**

5. Describe the HPE OneSphere offering.
 - **HPE OneSphere is a SaaS-based multicloud management solution. Through role-based access to HPE OneSphere, internal stakeholders (IT operations, developers, business executives) can manage hybrid clouds capable of supporting both virtualized and cloud-native applications.**

6. Which tool is responsible for forwarding data from HPE ProLiant servers connected to HPE InfoSight?
 - A. VMware vCenter server
 - B. iLO Amplifier Pack**
 - C. HPE OneSphere
 - D. Smart Update Tools

7. What are the typical requirements for a REST API request to iLO?
 - **URL/URI**
 - **HTTP method**

- **HTTP header**
- **Optional payload/body/data for some types of requests**

Practice Test

INTRODUCTION

The HPE ATP–Hybrid IT Solutions V2 certification validates a successful candidate’s ATP level skills to plan, design, recommend, and demonstrate HPE Hybrid IT solutions and deliver a proof of concept for a solution.

Ideal candidate for this exam

A successful candidate will have the competence to:

- Describe, differentiate, and apply industry-standard, foundational SMB architectures and technologies.
- Differentiate the functions, features, and capabilities of HPE product and solution-specific components and offerings.
- Recommend and position HPE SMB products, solutions, and appropriate services for customer use cases (positioning/use cases).
- Evaluate customer environment, and plan and design solutions using the HPE SMB portfolio to meet customer business requirements.
- Validate, install, configure, and upgrade HPE SMB solutions and their components.
- Troubleshoot, diagnose, and replace HPE SMB solution components.
- Manage, monitor, administer, and operate HPE SMB solution components.

Exam details

The following are details about the exam:

- Exam ID: HPE0-V14
- Exam Type: Proctored
- Number of items: 60
- Item types: Input text, input numbers, matching, multiple choice (single-response), multiple choice (multiple-response), and point and click
- Exam time: 1 hour 30 minutes
- Passing score: 68%
- Reference material: No online or hard copy reference material will be allowed at the testing site.

HPE0-V14 testing objectives

- 9%—Describe, differentiate, and apply industry-standard architectures and technologies.
- 17%—Recommend and position HPE Hybrid IT products, solutions, tools, and appropriate services for customer use cases.
- 26%—Evaluate customer environment, and plan and design solutions using the HPE SMB portfolio to meet customer business requirements.
- 25%—Validate, install, configure, and upgrade HPE SMB solutions and their components.
- 12%—Troubleshoot and repair HPE SMB solution components.
- 11%—Manage, monitor, administer, and operate HPE SMB solutions and their components.

Test preparation questions and answers

The following questions will help you measure your understanding of the material presented in this book. Read all the choices carefully, as there might be more than one correct answer. Choose all correct answers for each question.

Questions

1. Which RAID configuration must contain at least four drives and can survive two drive failures?
 - a. RAID 0
 - b. RAID 1
 - c. RAID 5
 - d. RAID 6

2. What is used to identify a server in an iSCSI SAN fabric?
 - a. IQN
 - b. WWNN
 - c. WWPN
 - d. SCSI ID

3. Match the principal way of implementing storage to its description:

Storage principle	Description
DAS	Dedicated network of storage systems providing block-based storage
NAS	Directly attached internal or external storage
SAN	Specialized storage device connected to Ethernet network providing file-based storage

4. Which command can be used to display LLDP information from a neighborhood device on a FlexFabric device?
 - a. display lldp neighbor-information
 - b. display connected-switches
 - c. show connected-switches
 - d. show lldp neighborhood

5. Match the virtualization and HA solution with the networking platform.

Networking platform	Virtualization and HA solution
---------------------	--------------------------------

FlexFabric	VSX
ArubaOS-CX	VSF
ArubaOS-Switch	IRF

6. A customer needs to move to a disk-based backup solution using HPE StoreOnce. They only have a single data center but want to send backup data offsite for the purposes of DR protection. The off-site data will only be read for DR purposes.

They have a business requirement to keep backup data totaling 500TB.

Identify the HPE StoreOnce system that will help the customer meet their goals in the most cost-effective manner.

- a. HPE StoreOnce 3620
 - b. HPE StoreOnce 3640
 - c. HPE StoreOnce 5200
 - d. HPE StoreOnce 5250
7. A system administrator needs to display packet statistics for port 1/0//10 facing the VMware ESX server, which two commands can be used on HPE FlexFabric device?
- a. display counters inbound interface Ten-GigabitEthernet 1/0/10
 - b. display counters outbound interface Ten-GigabitEthernet 1/0/10
 - c. ping interface Ten-GigabitEthernet 1/0/10
 - d. ping 192.1.0.10
 - e. ping 192.1.0.10 -p Ten-GigabitEthernet 1/0/10
8. Which HPE storage platform is using Triple-Parity RAID?
- a. HPE StoreOnce
 - b. HPE Nimble
 - c. HPE MSA

- d. HPE 3PAR StoreServ
9. A customer needs to update a server's firmware offline using the latest Support Pack for ProLiant (SPP). What is the correct way to boot from SPP?
- a. Prepare a USB drive and boot from the drive
 - b. Boot the server in emergency mode with F8
 - d. Use the iLO floppy functionality
 - e. Press F10 during POST and download SPP from HPE Support Center
10. Refer to the exhibit.



- Which RAID level can be selected at the next form? (Select two.)
- a. RAID 0
 - b. RAID 1
 - c. RAID ADM
 - d. RAID 5
 - e. RAID 6
 - f. Triple-parity
11. Which HPE storage platform can be connected to a server using SAS protocol?
- a. HPE StoreEasy
 - b. HPE StoreOnce

- c. HPE MSA
- d. HPE Primera

12. A local municipal institution is implementing a new virtualization solution to support their applications. They have the following requirements:

- Three VMware ESXi Hosts, each with two AMD processors, 1TB of RAM, 16Gbps FC Dual Port HBA
- HPE Nimble HF20 with 34TB useable capacity
- Integration into their existing Aruba network
- The solution must fit in 6U of rack space

If the customer needs more capacity in the future, they will use scale-out growth. Which server model should you recommend to this client?

- a. HPE ProLiant DL20
- b. HPE ProLiant DL325
- c. HPE ProLiant DL385
- d. HPE ProLiant DL580

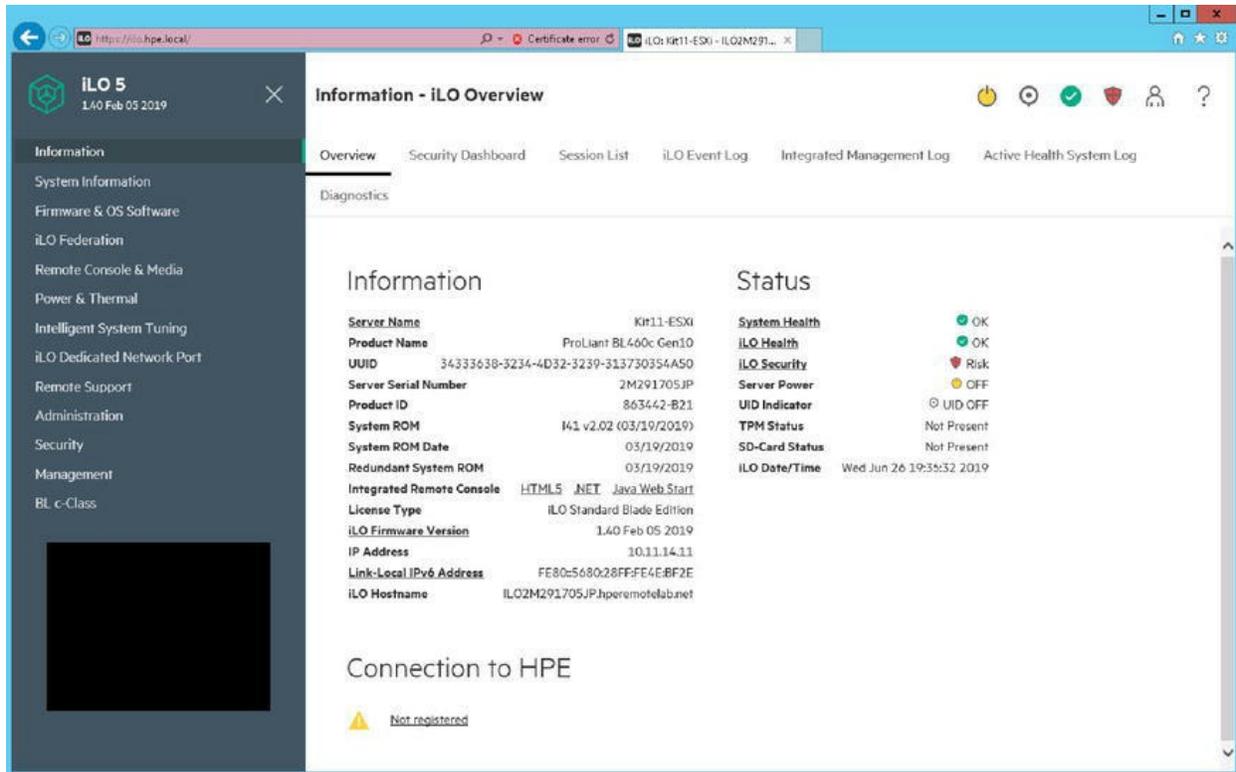
13. What is the possible SAS speed used by HPE D3710 Disk Enclosure? (Select two.)

- a. 6 Gb/s
- b. 8 Gb/s
- c. 10 Gb/s
- d. 12 Gb/s
- e. 16 Gb/s

14. Which of the following are typical components of a direct-attached storage (DAS) solution? (Select three.)

- a. Server
- b. Fibre Channel switch

- c. Array controller
 - d. Network-attached storage
 - e. Cables
 - f. Fibre Channel hub
15. Which of the following is a valid Fibre Channel switch differentiator?
- a. Automatic traffic rerouting
 - b. Maximum amount of ports
 - c. Maximum capacity of shared storage
 - d. Maximum amount of concurrent users
16. Which command can be used to check details about DIMM module using REST API?
- a. `curl -k https://ilo.hpe.local/details/show?DIMM`
 - b. `curl -k ftp://ilo.hpe.local/rest/v1/Systems/1/Memory/proc1dimm8-uadmin:password`
 - c. `curl -k https://ilo.hpe.local/redfish/v1/Systems/1/Memory/proc1dimm8-uadmin:password`
 - d. `curl -k https://admin:password@ilo.hpe.local/hw/Memory/-efault`
17. Refer to the exhibit.



Which link can be used to display information related to login/logout events to management processor?

- a. Security Dashboard
- b. iLO Event Log
- c. Integrated Management Log
- d. Management

18. Match the type of capacity with the description.

Type of storage capacity	Description
Raw	Capacity corrected of RAID overhead
Usable	Capacity corrected of deduplication and compression
Effective	Sum of drive capacities

19. Which functionality of iLO provides continuous, proactive health monitoring and recording of thousands of system parameters and diagnostic telemetry?
- a. Active Health System Log
 - b. Integrated Management Log
 - c. OneView
 - d. Insight Remote Support
20. Which HPE devices can be monitored and analyzed by HPE InfoSight for Servers? (Select two.)
- a. Synergy Gen9
 - b. Apollo G7
 - c. ProLiant with iLO3 or newer
 - d. Moonshot 1500
 - e. ProLiant G7
 - f. HPE ProLiant Gen8 and newer
21. A customer needs the HPE networking solution you are proposing to have advanced Layer 3 features and capabilities.

Which HPE switches meet this requirement? (Select two.) a. Aruba 2540

- a. Aruba 2930
 - b. Aruba 8320
 - c. FlexFabric 5700
 - d. FlexFabric 5710
22. Which transceiver type does support 25GbE?
- a. QSFP+
 - b. SFP28
 - c. MiniSAS HD

- d. SFP+
23. Which HPE platforms are supported for monitoring and analysis by HPE InfoSight? (Select two.)
- a. ProLiant servers with iLO3
 - b. Synergy Gen8
 - c. Apollo with iLO4
 - d. Moonshot 1500
 - e. ProLiant Gen10 servers
24. A system administrator is using an HPE OneView for monitoring DL servers and requires a CSV file with firmware versions of their server infrastructure. Which procedure can be used to obtain such report?
- a. Export CSV firmware report from the Dashboard.
 - b. Export Current firmware from Logical enclosures screen.
 - c. Open Server hardware page, select all servers, and click Actions menu.
 - d. Export to CSV from Server firmware inventory report.
25. A system administrator is configuring zoning at B-series FC switch. Which command can be used to add a server WWN to the switch?
- a. `alicreate "ESX1"," 20000000FD27E5BD"`
 - b. `addzone "ESX1"," 20000000FD27E5BD"`
 - c. `create server ESX1 20000000FD27E5BD`
 - d. `create server ESX1; modify server ESX1 -wwnn 20000000FD27E5BD"`
26. Which HPE networking device should you recommend to a customer requiring an affordable, data center switch supporting 10 Gb downlink ports and 40 Gb uplinks for a small-scale environment?
- a. FlexFabric 5710

- b. FlexFabric 5940
- c. Aruba 8325-32C
- d. HPE Altoline 9960

27. Match the backup-related factor to its description:

Term	Description
RTO (Recovery Time Objective)	How much data is allowed to go unprotected, and how far back in time the data must be recovered.
Retention	How long the customer is willing to wait for the data to be recovered and the maximum allowable downtime.
RPO (Recovery Point Objective)	How long the data needs to be kept available. Can range from seconds to decades, depending on company policies and government regulations.

28. Which HPE management tool provides analysis and cloud monitoring for StoreServ devices in a converged infrastructure?

- a. HPE 3PAR SSMC
- b. HPE OneView for storage
- c. HPE OneView
- d. HPE InfoSight for HPE 3PAR

29. Match the HPE server to the feature:

HPE Server	Description
ML350	Depth below 50cm / 20in
Depth below 50cm / 20in ML30	Hot-plug fans and two processor sockets
DL325	Supports AMD EPYC processor

30. You have configured a port on an HPE FlexFabric 5900 series switch that is connected to a VMware ESXi server to support multiple VLANs.

description server ESXi01-eth1

```
port link-type trunk
port trunk permit vlan 50
port trunk permit vlan 2
port trunk pvid vlan 50
```

Given the configuration above, what should happen to the traffic tagged with VLAN 50 coming from the VMware ESXi host?

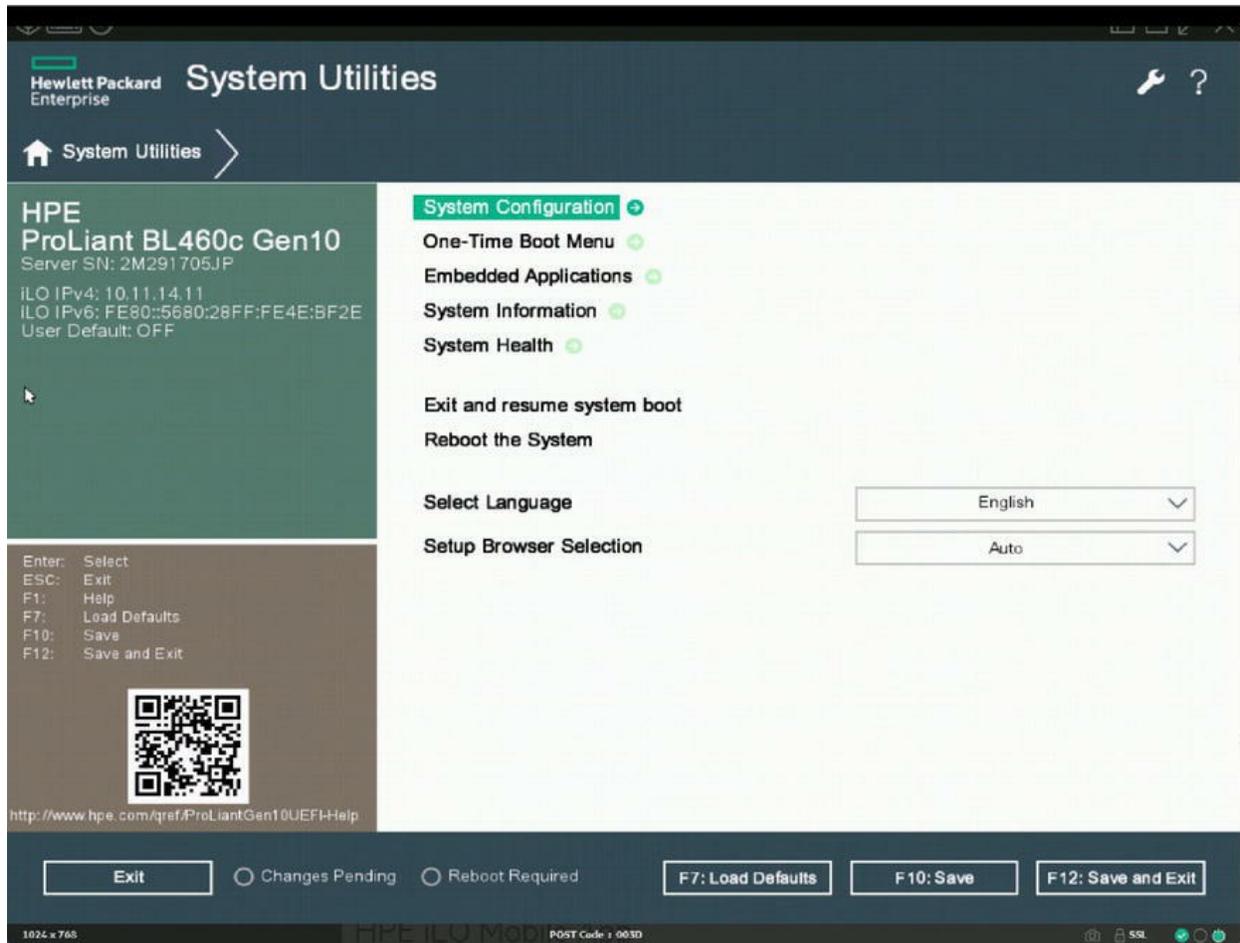
- a. Traffic will be dropped
- b. Traffic will be retransmitted as untagged
- c. Traffic will be transmitted in vlan 1
- d. Traffic will be transmitted in vlan 50

31. You are reviewing a storage design using MSA2050 and 30TB of flash storage. The customer requires protection against a single drive failure at the lowest price point.

Which of the disk drive designs is adequate for this site?

- a. Twelve HPE MSA 3.2TB drives, configured with RAID6
- b. Twenty HPE MSA 3.2TB drives, configured with RAID1+0
- c. Ten HPE MSA 3.2TB drives, configured with RAID0
- d. Eleven HPE MSA 3.2TB drives, configured with RAID5

32. Refer to the exhibit.



A system administrator needs to run UEFI Server Hardware Diagnostics to verify a hardware health issue. What is the correct menu option to reach this functionality?

- a. System Configuration
- b. Embedded Applications
- c. System Information
- d. System Health

Answers

1. **D** is correct. RAID 6 can survive two drive failures and requires a minimum of four drives.
 A, B, and C are incorrect. RAID 0 does not offer protection, RAID 1 requires a minimum of two drives, and RAID 5 requires three drives as a minimum and protects against a single drive

failure.

For more information, see [Chapter 2](#).

2. **A** is correct. IQN is used to identify a server in an iSCSI SAN.
- B, C, and D** are incorrect. WWNN and WWNP are used in FC SAN; and SCSI ID is typically used between SCSI controller and drives.

For more information, see [Chapter 5](#).

3. Correct matching is shown in the table below:

Storage principle	Description
SAN	Dedicated network of storage systems providing block-based storage
DAS	Directly attached internal or external storage
NAS	Specialized storage device connected to Ethernet network providing file-based storage

For more information, see [Chapter 4](#).

4. **A** is correct. display lldp neighbor-information command will display the necessary information.
- B, C, and D** are incorrect. These commands will fail on a FlexFabric device.

For more information, see [Chapter 3](#).

5. Correct matching is shown in the table below:

Networking platform	Virtualization and HA solution
FlexFabric	IRF
ArubaOS-CX	VSX
ArubaOS-Switch	ArubaOS-Switch

For more information, see [Chapter 4](#).

6. **D** is correct. Only HPE StoreOnce 5250 provides support for more than 500TB of storage.
- A, B, and C** are incorrect. HPE StoreOnce 3620, HPE StoreOnce 3640, and HPE StoreOnce 5200 do not provide support for 500TB.

For more information, see [Chapter 6](#).

7. **A** and **B** are correct. The display counters commands will display packet statistics per port.
 C, **D**, and **E** are incorrect. While the **D** is a correct command, it does not solve the problem. **C** and **E** will fail on an HPE FlexFabric device.
For more information, see [Chapter 3](#).
8. **B** is correct. Triple-Parity RAID is a feature of HPE Nimble Storage.
 A, **C**, and **D** are incorrect. None of StoreOnce, MSA, or 3PAR StoreServ use Triple-Parity RAID.
For more information, see [Chapter 5](#).
9. **A** is correct. SPP provides a utility to create USB bootable drive for offline firmware updates.
 B, **C**, and **D** are incorrect. SPP will not fit to floppy image, booting operating system is not an offline maintenance, and SPP cannot be downloaded from Intelligent Provisioning accessible via F10.
For more information, see [Chapter 2](#).
10. **A** and **B** are correct. RAID 0 and RAID 1 can be used with two drives.
 C, **D**, and **E** are incorrect. RAID ADM and RAID 5 require three drives; RAID 6 requires three drives. Triple-parity RAID cannot be configured using the UEFI utility at the server level.
For more information, see [Chapter 2](#).
11. **C** is correct. HPE MSA Storage offers controllers with SAS connectivity.
 A, **B**, and **D** are incorrect. StoreEasy is a NAS system connected to a network and HPE StoreOnce provides stores over TCP/IP or FC SPP. HPE Primera is using SAS for drive enclosures, but not for servers.
For more information, see [Chapter 5](#).
12. **C** is correct. Only the DL385 can be equipped with two AMD processors.
 A, **B**, and **D** are incorrect. DL20 and DL580 use Intel processors; DL325 uses AMD, but is a single-socket system.
For more information, see [Chapter 2](#).
13. **A** and **D** are correct. HPE D3710 supports both 6 Gb/s and 12 Gb/s SAS connections.
 B, **C**, and **E** are incorrect. 8 Gb/s, 10Gb/s, and 16 Gb/s are not speeds used for SAS connections.
For more information, see [Chapter 5](#).
14. **A**, **C**, and **E** are correct. Typical DAS solution consists of server, HBA or array controller, cables, and storage enclosure.
 B, **D**, and **F** are incorrect. Fibre Channel components are used in SAN and Network-attached storage is not a DAS component.
For more information, see [Chapter 5](#).

15. **B** is correct. Amount of ports is typical FC switch differentiator.
 A, C, and D are incorrect. Automatic traffic rerouting is a feature of all FC switches. Maximum amount of concurrent users and shared storage capacity is not limited by FC switch.
For more information, see [Chapter 2](#).
16. **C** is correct. This is the only command that contains the correct URL including redfish and login credentials.
 A, B, and D are incorrect. These commands are missing the string redfish in the URL, use incorrect authentication, and/or non-existing parts of the URL.
For more information, see [Chapter 2](#).
17. **B** is correct. Login, logout, and power-related events are part of iLO Event Log.
 A, C, and D are incorrect. Security Dashboard does not contain any log/events, IML is oriented more to hardware health events, and Management is used for configuration.
For more information, see [Chapter 2](#).
18. Correct matching is shown in the table below:

Type of storage capacity	Description
Raw	Sum of drive capacities
Usable	Capacity corrected of RAID overhead
Effective	Capacity corrected of deduplication and compression

19. **A** is correct. AHS log provides continuous, proactive health monitoring and recording of thousands of system parameters and diagnostic telemetry.
 B, C, and D are incorrect. IML contains hardware-related events, but no recording of system parameters. OneView and Insight Remote Support do not contain diagnostic telemetry.
For more information, see [Chapter 6](#).
20. **A and F** are correct. InfoSight for Servers is the extension of InfoSight for ProLiant, Synergy, and Apollo Gen10, Gen9, and Gen8 servers with iLO 5 and iLO 4.
 B, C, D, and E are incorrect. G7 servers with iLO3 or older are not supported, nor are Moonshot systems.
For more information, see [Chapter 6](#).
21. **C and D** are correct. Aruba 8320 and FlexFabric 5710 support advanced Layer 3 capabilities.
 A, B, and E are incorrect. Aruba 2540 and FlexFabric 5700 provide lite Layer 3 features and Aruba 2930 basic features. Advanced Layer 3 features, such as BGP or BFD, are missing.
For more information, see [Chapter 3](#).

22. **B** is correct. SFP28 supports 25GbE.
 A, C, and D are incorrect. QSFP+ is used for 40GbE, MiniSAS HD supports 12Gb and does not support Ethernet. SFP+ supports 1 or 10GbE.
For more information, see [Chapter 3](#).
23. **C and D** are correct. InfoSight for Servers is the extension of InfoSight for ProLiant, Synergy, and Apollo Gen10, Gen9, and Gen8 servers with iLO 5 and iLO 4.
 A, B, and E are incorrect. G7 systems with iLO3 and older are not supported. The Moonshot platform is not supported by InfoSight.
For more information, see [Chapter 6](#).
24. **D** is correct. Server firmware inventory report supports CSV export.
 A, B, and C are incorrect. The Dashboard, Logical enclosures, and Server hardware screens do not have the possibility to export firmware in CSV format.
For more information, see [Chapter 6](#).
25. **A** is correct. alicreate is the command used for creating aliases and adding WWN to a B-Series switch.
 B, C, and D are incorrect. These commands do not work on a B-Series switch.
For more information, see [Chapter 5](#) and [B-Series command reference](#).
26. **A** is correct. FlexFabric 5710 provides the required functionality.
 B, C, and D are incorrect. FlexFabric 5710 is the best option to provide the required functionality at the lowest price point.
For more information, see [Chapter 3](#).
27. Correct matching is shown in the table below:

Term	Description
RTO (Recovery Time Objective)	How long the customer is willing to wait for the data to be recovered and the maximum allowable downtime.
Retention	How long the data needs to be kept available. Can range from seconds to decades, depending on company policies and government regulations.
RPO (Recovery Point Objective)	How much data is allowed to go unprotected and how far back in time the data must be recovered.

For more information, see [Chapter 5](#).

28. **D** is correct. HPE InfoSight provides the required functionality.
 A, B, and C are incorrect. HPE OneView and 3PAR SSMC are not used as management tools in a cloud.

For more information, see [Chapter 5](#).

29. Correct matching is shown in the table below:

HPE Server	Description
ML350	Hot-plug fans and two processor sockets
ML30	Depth below 50cm/20in
DL325	Supports AMD EPYC processor

For more information, see [Chapter 5](#).

30. **D** is correct. The port trunk permit vlan 50 command allows receiving tagged traffic with VLAN 50.

A, B, and **C** are incorrect. Tagged traffic with VLAN 50 is not allowed with these commands.

For more information, see [Chapter 4](#).

31. **C** is correct. Eleven drives with RAID 5 will provide the required capacity and protection against a single drive failure.

A, B, and **D** are incorrect. RAID6 or RAID 1+0 will exceed the functionality, but not at the lowest price point. RAID 0 does not protect against a drive failure.

For more information, see [Chapter 2](#).

32. **B** is correct. The Embedded Applications menu contains a link to UEFI Server Hardware Diagnostics.

A, C, and **D** are incorrect. The System Configuration, System Information, and System Health menu options cannot be used to reach UEFI Server Hardware Diagnostics.

For more information, see [Chapter 2](#).