



# Hewlett Packard Enterprise

---

## Hewlett Packard Enterprise Partner Ready Certification and Learning

---

This is a Hewlett Packard Enterprise Partner Ready Certification and Learning authorized course that is designed to prepare you for the associated certification exam. The course instruction provides the best source of information that can be used and studied to prepare for the associated certification exam.

HPE Partner Ready Certification and Learning provides training courses and certification exams for the most sought-after IT disciplines, including convergence, cloud computing, software-defined networking, and security. This training will help you learn how to design solutions that deliver business value, as well as gain the hands-on experience you need to progress in your career.

The Partner Ready Certification and Learning program gives you:

- A full range of learning opportunities spanning foundational skill level to master
- Personalized learning plans and resources through [mylearninghpe.com](https://mylearninghpe.com)
- Certifications that command some of the highest pay levels in the industry
- A focus on end-to-end integration, open standards, and emerging technologies
- Maximum credit for certifications you already hold
- A supportive global community of IT professionals
- A curriculum of unprecedented breadth from HPE, the world's most complete technology company

Visit [certification-learning.hpe.com](https://certification-learning.hpe.com) to learn more about HPE certifications, and find the training you need to adopt new technologies that will further enhance your IT expertise and career.





**Hewlett Packard**  
Enterprise

# HPE Building Data Center Solutions

Learner guide

Rev. 17.41

Confidential – For Training Purposes Only

Use of this material to deliver training without prior written permission from HPE is prohibited.

© Copyright 2017 Hewlett Packard Enterprise Development LP

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

This is an HPE copyrighted work that may not be reproduced without the written permission of Hewlett Packard Enterprise. You may not use these materials to deliver training to any person outside of your organization without the written permission of HPE.

Microsoft Windows Server®, Microsoft SQL Server®, Microsoft Azure®, Microsoft HyperTerminal®, and Skype for Business® are used with permission from Microsoft. VMware ESX® and ESXi® are trademarks or registered trademarks in the United States and certain other countries.

Intel® and Xeon® are trademarks of Intel Corporation in the United States and other countries.

Red Hat® and Red Hat Enterprise Linux® are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

SUSE® is a trademark of SUSE IP Development Limited or its subsidiaries or affiliates.

Mac OS® is a trademark of Apple Inc., registered in the U.S. and other countries.

Tera Term is open source and is copyrighted by the TeraTerm Project.

All other trademarks are the property of their respective owners.

Printed in USA

**HPE Building Data Center Solutions**

Learner guide

October 2017



---

# Contents

## Module 1—Recommending HPE SMB Solutions for Customer Use Cases

<b>Learning objectives</b> .....	<b>1</b>
<b>Customer scenario</b> .....	<b>2</b>
IDJT HealthCare .....	2
<b>HPE SMB server, storage, and networking solutions</b> .....	<b>3</b>
Gauge your knowledge .....	3
Discovering a customer’s business and technical requirements .....	4
Issues to consider .....	5
Assessing the customer’s applications .....	6
Customer environment requirements .....	7
Solution design considerations .....	8
Activity—Discovering a customer’s business and technical requirements .....	9
Customer scenario.....	9
Questions.....	9
<b>Mapping HPE server solutions to customer workloads based on business size</b> .....	<b>10</b>
HPE Gen10 servers .....	12
Learning check.....	14

<b>Mapping HPE storage solutions to customer requirements .....</b>	<b>15</b>
Gauge your knowledge .....	16
Nimble Storage .....	17
HPE MSA Storage .....	19
HPE StoreVirtual Storage .....	21
HPE StoreEasy Storage .....	23
StoreEasy 1000 Storage.....	24
StoreEasy 3000 Gateway Storage .....	24
HPE StoreOnce Storage.....	25
Activity—Selecting a storage system.....	27
Questions.....	28
Learning check.....	29
<b>Recommending HPE networking for SMBs.....</b>	<b>30</b>
Gauge your knowledge .....	30
SMB networking considerations .....	31
HPE networking solutions address SMB needs .....	33
HPE FlexFabric.....	35
HPE FlexManagement.....	36
HPE Switch Selector.....	37
Activity—Using the HPE Switch Selector .....	38
Customer scenario.....	38
Arista Networks.....	39
Deploying an IP storage infrastructure .....	39
CloudVision.....	40
Aruba Networks .....	41
Activity—Choosing the right Aruba networking solution .....	42
Customer scenario.....	42
Learning check.....	43
<b>Recommending HPE data center infrastructure services .....</b>	<b>44</b>
Warranty.....	44
HPE Foundation Care .....	44
Warranty compared with Foundation Care .....	45
HPE Proactive Care .....	46
HPE Proactive Care Advanced.....	46
Learning check.....	47
<b>Summary .....</b>	<b>48</b>

## Module 2—Planning and Designing HPE SMB Solutions

<b>Learning objectives</b> .....	<b>1</b>
<b>Customer scenario</b> .....	<b>2</b>
IDJT HealthCare .....	2
Gauge your knowledge.....	3
<b>Planning and designing an HPE SMB data center</b> .....	<b>4</b>
<b>Choosing a server</b> .....	<b>4</b>
HPE SMB server comparison .....	5
Intelligent System Tuning.....	6
Choosing server networking components.....	7
Gen10 server network adapters .....	8
Learning check.....	10
<b>Choosing a storage solution</b> .....	<b>11</b>
NAS.....	12
VTL instead of NAS using a StoreOnce target.....	13
SAN.....	14
HPE storage networking options .....	15
Optional storage configurations to consider .....	16
Activity—Using HPE SPOCK.....	17
Learning check.....	23
<b>Designing the network</b> .....	<b>24</b>
HPE SMB networking portfolio .....	25
Wired and wireless environments.....	26
Types of switches .....	27
VLAN connectivity.....	28
Types of VLANs.....	29
Port count and type.....	29
Hybrid switch ports .....	30
Technologies for configuring the network .....	31
Link aggregation .....	32
Link aggregation requirements .....	33
MLAG.....	34
MLAGs compared to traditional stacking.....	35
Virtual Switching Framework.....	36
Intelligent Resilient Framework.....	37
Learning check.....	38

<b>Additional infrastructure considerations .....</b>	<b>39</b>
Power and cooling .....	39
Proactive maintenance .....	40
Best practices for resiliency .....	41
Learning check.....	42
<b>Validating the design and documenting the solution .....</b>	<b>43</b>
Writing a scope of work.....	43
Developing the proposal .....	44
Learning check.....	45
<b>Summary .....</b>	<b>46</b>

## **Module 3—Installing, Configuring, and Upgrading HPE SMB Solutions and Related Components**

<b>Learning objectives .....</b>	<b>1</b>
<b>Customer scenario.....</b>	<b>2</b>
IDJT HealthCare .....	2
Gauge your knowledge .....	3
<b>Installing solutions and subsystems .....</b>	<b>4</b>
Installing data center components .....	5
Installing server options .....	6
Installing the rack infrastructure .....	7
Power requirements.....	9
Cabling requirements.....	10
Activity—Selecting a rack and power infrastructure .....	11
Customer scenario.....	11
Using the HPE Rack and Power Infrastructure Architect tool .....	11
Using the HPE Power Advisor .....	21
Learning check.....	25
<b>Configuring ProLiant servers .....</b>	<b>26</b>
BIOS and UEFI .....	27
Launching UEFI System Utilities .....	29
Using UEFI System Utilities .....	29
Selecting the boot mode .....	30
Workload profiles and performance options .....	31

HPE iLO 5 .....	33
iLO 5 licensing .....	33
Setting up an iLO management processor .....	34
iLO 5 service port on Gen10 servers .....	36
Using the iLO 5 Configuration Utility .....	37
Logging in to iLO for the first time .....	37
Intelligent Provisioning .....	38
Always On Intelligent Provisioning .....	39
Software installed with Intelligent Provisioning .....	40
Installing memory in ProLiant Gen10 servers .....	41
DIMM population rules for ProLiant Gen10 servers .....	42
Understanding balanced memory configurations .....	43
Memory configurations that are unbalanced across channels .....	43
Memory configurations that are unbalanced and balanced across processors .....	44
Installing HPE 25Gb/100Gb Ethernet adapters .....	45
Learning check.....	46
<b>Configuring HPE storage .....</b>	<b>47</b>
Designing a SAN using Fibre Channel switches .....	47
Using HPE Smart Storage Administrator .....	48
Accessing SSA .....	49
Navigating the GUI .....	51
SSA actions .....	52
Creating an array .....	53
SSA Array Details screen .....	54
SSA Configure screen .....	55
Activity—Researching SSA.....	56
Questions.....	56
Configuring HPE MSA storage solutions .....	57
Virtual storage.....	57
Disk groups.....	58
Configuring HPE StoreVirtual solutions .....	59
Creating iSCSI storage .....	60
Creating Fibre Channel storage .....	60
Configuring RAID in a StoreVirtual system.....	61

Configuring StoreOnce systems .....	62
Configuring HPE StoreEasy solutions .....	63
Heterogeneous file sharing .....	64
Learning check.....	65
<b>Configuring the network.....</b>	<b>66</b>
Configuring in-band and out-of-band system management .....	66
In-band management.....	67
Out-of-band management.....	68
Fabric designs.....	69
Single-switch fabric.....	69
Cascaded fabric.....	70
Meshed fabric .....	71
Core-edge fabric .....	72
Tiered network configurations .....	73
Leaf-spine topology .....	74
HPE Intelligent Resilient Framework .....	75
IRF and meshed stacking.....	76
Forming link aggregation connections.....	77
LACP operational modes.....	79
Link aggregation load sharing .....	79
Configuring IRF devices .....	80
IP routing.....	81
Topology that requires indirect routes .....	82
Setting up indirect IP routes .....	84
Static IP routing .....	85
Accessing and configuring HPE Comware switches .....	86
Managing configuration files.....	87
Basic CLI command levels and functions.....	88
Learning check.....	89
<b>Validating that an installation and configuration are successful .....</b>	<b>90</b>
<b>Upgrading the solution and its subsystems .....</b>	<b>91</b>
HPE SUM and iLO Federation.....	93
Nondisruptive activation for switch firmware upgrades .....	94
Learning check.....	96
<b>Summary .....</b>	<b>97</b>

# Module 4—Managing, Monitoring, Administering, and Operating HPE SMB Solutions and Related Components

<b>Learning objectives</b> .....	<b>1</b>
<b>Customer scenario</b> .....	<b>2</b>
IDJT HealthCare .....	2
<b>Gauge your knowledge</b> .....	<b>2</b>
<b>HPE management and administration tools</b> .....	<b>3</b>
Using iLO 5 to manage a data center .....	3
Agentless Management .....	4
iLO web interface .....	5
Logging in to the iLO web interface .....	5
Using iLO controls .....	6
Viewing iLO overview information .....	7
iLO server health summary.....	8
Additional iLO functions .....	9
iLO 5 Federation .....	11
iLO Federation group characteristics.....	12
iLO Federation group privileges .....	12
iLO Federation group memberships for local iLO systems .....	13
Using HPE OneView to manage a data center.....	14
HPE OneView licensing.....	15
Navigating the HPE OneView GUI .....	16
Provisioning hardware .....	18
HPE OneView Global Dashboard.....	19
Provisioning an HPE StoreVirtual volume from a server profile .....	21
Adding, removing, and editing StoreVirtual volumes online .....	22
Remote Support in HPE OneView.....	23
Learning check.....	24
<b>Managing and monitoring storage</b> .....	<b>25</b>
Nimble InfoSight.....	25
HPE Smart Storage Administrator .....	27
HPE StoreVirtual Centralized Management Console .....	28
HPE Storage Management Utility .....	28
Best practices for MSA .....	28
HPE Library and Tape Tools .....	29

**Managing and monitoring the network..... 30**

- HPE Intelligent Management Center ..... 31
  - IMC management protocols ..... 33
  - Link Layer Discovery Protocol ..... 34
- Aruba management ..... 35
  - Aruba ClearPass..... 36
- Learning check..... 37

**Summary ..... 38**

## **Module 5—Troubleshooting HPE SMB Solutions**

**Learning objectives ..... 1**

**Customer scenario..... 2**

- IDJT HealthCare ..... 2

**Gauge your knowledge..... 3**

**Introduction to troubleshooting ..... 4**

- Pretroubleshooting and replacement steps ..... 4
  - Verifying that a current backup exists..... 4
  - Graceful power-down and power-up sequences ..... 5

**HPE six-step troubleshooting methodology ..... 6**

- Step 1: Collect data..... 7
  - Collecting system data..... 8
- Step 2: Evaluate data to determine subsystems causing the issue ..... 10
- Step 3: Develop an optimized action plan ..... 10
- Step 4: Execute the action plan ..... 10
- Step 5: Determine whether the problem is solved..... 11
- Step 6: Implement preventive measures ..... 11
- Learning check..... 12

**Additional troubleshooting methods and tools ..... 13**

- Establishing a baseline ..... 13
- Resolving common problems..... 14
  - Corrupted ROM disaster recovery ..... 14
- Using published server information ..... 15
  - HPE resources for troubleshooting..... 15
- HPE Support Center ..... 17
- USB support..... 18
- Learning check..... 19

<b>Troubleshooting HPE ProLiant servers</b> .....	<b>20</b>
System configuration information .....	20
Useful server configuration information .....	20
Event logs .....	21
Beep codes .....	22
POST error messages .....	22
Checking status LEDs .....	23
ProLiant Gen9 server front panel LEDs example .....	23
Back panel LEDs example.....	24
Systems Insight Display LEDs example .....	25
Basic troubleshooting with HPE OneView .....	26
iLO event log .....	27
Viewing the iLO event log .....	27
iLO Service Port.....	28
Learning check.....	29
<b>Troubleshooting HPE storage</b> .....	<b>30</b>
HPE Smart Storage Administrator diagnostics .....	30
HPE Smart Storage Administrator Diagnostics Utility CLI .....	32
Installing the utility .....	32
Launching the utility in CLI mode .....	32
Additional tools for diagnosing array problems.....	33
Troubleshooting LUNs .....	33
Using L&TT to troubleshoot tape devices.....	34
<b>Troubleshooting the network</b> .....	<b>35</b>
Diagnosing switch problems .....	35
Diagnostic tools.....	36
Port auto-negotiation .....	36
Troubleshooting telnet issues .....	37
Troubleshooting unusual network activity.....	37
General problems .....	37
LACP-related problems .....	38
Learning check.....	39
<b>Summary</b> .....	<b>40</b>

**Appendix 1—Glossary of HPE Data Center Solutions Terms**

**Appendix 2—Scope of Work Worksheet**



---

# Recommending HPE SMB Solutions for Customer Use Cases

## Module 1

### Learning objectives

After you have completed this module, you should be able to:

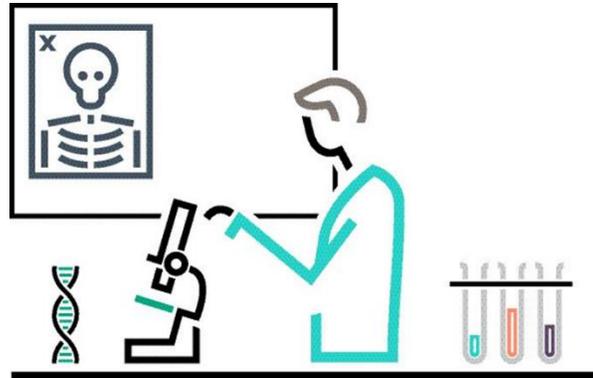
- Summarize the various Hewlett Packard Enterprise (HPE) server, storage, and networking solutions that are ideal for small to medium-sized business (SMB) customers
- Map HPE server solutions to customer workloads based on business size
- Guide customers to HPE SMB storage solutions that are appropriate for their requirements
- Recommend HPE SMB networking solutions for various customer workloads
- Recommend HPE data center infrastructure services that benefit SMBs

## Customer scenario

When working your way through this course, it is helpful to keep a customer in mind. You can imagine the customer's business and technical needs within the context of a real-world scenario and apply what you learn as the business grows and evolves through the course.

Consider the following fictitious customer as you work through this module. Use it as a model when trying to recommend and design a solution to meet the customer's business and technical objectives. Refer back to it also as you engage with activities and labs throughout this course.

### IDJT HealthCare



Healthcare organizations need to implement strong data protection to achieve and maintain compliance

For more than 80 years, IDJT HealthCare has provided renowned physicians and caregivers with the most appropriate technology across medical focus areas including cancer, cardiology, sports medicine, pain management, and sleep disorders.

IDJT HealthCare is located in the state of Pennsylvania in the United States. Newly passed legislation requires IDJT HealthCare to increase data center security in order to protect patient privacy. As a result, they are opening a remote location that will accommodate 10 employees responsible for real-time patient and lab data as well as provider information retrieval. The company plans to install physical security measures at the remote location including an access control system, pin codes, and video surveillance.

The IT manager of IDJT HealthCare, Eric Yurekha, says, "We understand that the cost of downtime can include lost revenue, damaged reputation, lost customer confidence and loyalty, and regulatory compliance exposure. We might need to wait until the best time to make any changes in our data center to minimize the impact of downtime during the move."

In addition to compliance deadlines, the company's budget requires that they repurpose existing assets. IDJT HealthCare is planning to move equipment to the remote location, including an HPE ProLiant DL360 Gen9 server running VMware ESXi and Microsoft SQL Server, and HPE StoreOnce Virtual Storage Appliance (VSA) using local disks. The server will host the company's self-programmed database application.

As a result of these changes, Yurekha says, "We plan to have better security on a dedicated system that will carry us forward as the industry evolves. Because the medical records database will be easier and faster to access and use, nurses at the hospital can spend more time caring for their patients."

## HPE SMB server, storage, and networking solutions

HPE server, storage, and networking solutions are designed to handle multiple workloads—IT infrastructure, web, business applications, collaboration, analytics, big data, and more—but each product is optimized for specific use cases. It is important to understand the basic features of the different portfolios and focus on how those details align to customer needs.

### Gauge your knowledge

Before proceeding with this section, consider the following questions to assess your existing knowledge. Pay particular attention to the answers you are less certain about. Keep in mind that there are no right or wrong answers to these questions.

1. Have you ever worked with a customer to gather their requirements for a data center solution?

---

2. What experience do you have with HPE ProLiant servers, modular smart array (MSA) storage solutions, or HPE networking switches?

---

---

3. Have you used a management tool to monitor a data center solution? If so, which one?

---

4. Do you know how to recommend an HPE service for a customer?

---

Make note of areas in which you need more information and concentrate on the material covered in those areas. Feel free to skim any material you might already know. Remember that in order to achieve the HPE ATP – Consolidated Data Center Solutions V1 certification, you must pass the exam associated with this training.

## Discovering a customer's business and technical requirements

When planning and developing an IT solution, you must first assess the business and IT needs of the customer. When all customer information has been gathered, use your experience and knowledge to recommend the best possible solution. The IT recommendation is ultimately described in the statement of work, but first 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 requirements
- Pay close attention to how much computing equipment is required
  - If over-configured, the solution does not return good value
  - If under-configured, the solution does not provide adequate service
  - Several different hardware configurations might be appropriate
- Understand the hardware components and sizing requirements
- Consider:
  - The price-to-performance ratio
  - The customer's budget
  - Total cost of ownership (TCO)/return on investment (ROI) requirements
  - The time frame compelling the purchase
- Create specifications for the configuration options
- Identify service levels, and determine how software and hardware will be used
- Consider the customer's environmental requirements

Two analyses should be performed in order to thoroughly assess each customer:

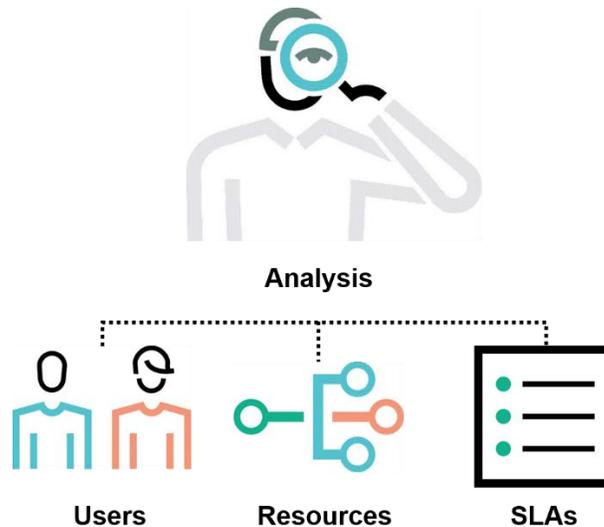
- **Applications/workload analysis**—Analyzing the types of workloads the customer is running can provide guidance in selecting a server family. It is important to find out if the customer needs support for virtualization, cloud, web infrastructure, database, app development, and so on. For example, 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 BladeSystem or HPE Synergy solutions should be considered. Specific factors to evaluate include power and cooling requirements along with interconnectivity.

---

**NOTE:** Although some server blades can map to an SMB environment, the BladeSystem C7000 enclosure is considered enterprise level and outside the scope of this training.

---

## Assessing the customer's applications

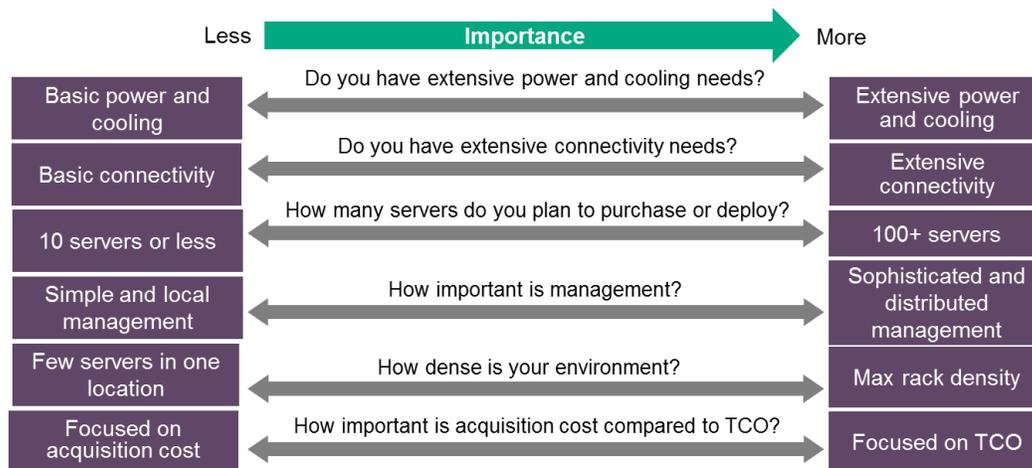


Examine the customer's existing resources, workloads, and overall operations

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. Several factors should be considered, including:

- **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.

## Customer environment requirements



Ask questions that align the customer's existing setup with standards for the server being installed

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. The HPE Blanking Panels Rack Option Kit enables you to cover open areas of the rack to better control airflow.

- **Temperature**—The maximum recommended ambient operating temperature for most server products 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 overall 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

## 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 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 he or she needs. All passwords should be changed regularly; meet the minimum character length (as defined by the operating system); and include letters, numbers, and special characters.

## Activity—Discovering a customer’s business and technical requirements

To complete this activity, read the following customer scenario. Then apply what you have learned in this section of the module to answer the questions regarding the customer’s business and technical requirements.

### Customer scenario

IDJT HealthCare has asked for your advice regarding the changes they need to make in their data center, especially in the area of security and regulation compliance. The IT manager of IDJT HealthCare, Eric Yurekha, said, “I sat down with our systems administrator, Ben Sessions, to discuss timing for the changes. He’s pretty excited to get going as soon as we can. But our finance manager, Yolonda Arena, is looking at other major expenditures we have planned for this year and is trying to hold down costs of the remote location. We worked through the numbers and we have a small budget for the new location. We are still working on our forecast for the next five years, but I hope to have that for you by the time we meet next week.”

### Questions

To prepare for the meeting, list the information that you already have that will help you make recommendations for IDJT’s data center.

Stakeholders

---

Future plans

---

---

Obstacles

---

---

Other

---

Next, make a list of questions to ask the systems administrator, Ben Sessions.

Current environment

---

---

Technical requirements

---

---

Resources

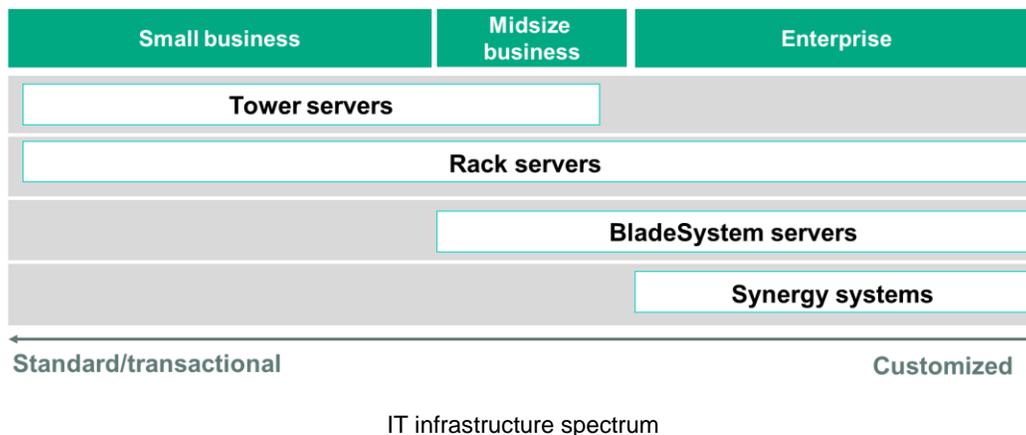
---

---

Other

---

## Mapping HPE server solutions to customer workloads based on business size



All customers have different needs, and there is no substitute for completing a full needs analysis. However, this graphic provides useful high-level guidance regarding which HPE servers might be suitable for SMBs. The enterprise column is included here for comparison.

**NOTE:** To compare HPE servers for SMBs by workload, visit the HPE Marketplace by scanning the following QR code or right-clicking it to open the hyperlink:



<https://marketplace.hpe.com/categoryLanding?catId=15351&country=US&locale=en>

Security, versatility, performance, and quality make HPE ProLiant servers ideal for small and midsize business workloads as well as for software-specific solutions such as HPE Flex Solutions for virtualization, analytics, and unified communications and collaboration solutions.

ProLiant rack and tower servers redefine compute economics by delivering more compute and storage capacity, right-sized compute with flexible choices, and lower compute energy and floor space consumption. With ProLiant rack and tower servers, customers can accelerate business results with faster compute, memory, and I/O performance, coupled with increased storage and networking performance—including lower latency. ProLiant rack and tower servers are available in these families:

- ProLiant MicroServer (tower only)
- ProLiant ML (rack and tower)
- ProLiant DL (rack only)

Additional platforms appropriate for certain SMB environments are BladeSystem solutions and Synergy solutions.

BladeSystem solutions provide multiple IT elements in one optimized package that is managed as a single platform. Blades are available as composable blocks of compute, storage, and network fabric. The building block approach makes it simple to scale infrastructure by reducing the typical complexity of configuring and separately scaling compute, storage, and networking piecemeal in traditional infrastructure. BladeSystems use optimized and preselected server profiles for optimal workload matching and performance. BladeSystems are ideal for:

- Virtual desktop infrastructure (VDI)
- Consolidation

Although ProLiant DL and ML solutions and ProLiant BL servers are viable resources in a variety of situations, the differences between them are significant. Rack-mounted servers typically are viewed as a stand-alone resource, whereas a BladeSystem has many shared components and is aggregated to function as a larger resource.

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. Synergy solutions support emerging applications in the areas of mobility, big data, and cloud native technologies. Synergy is ideal for a single infrastructure running multiple applications simultaneously:

- SQL on bare-metal compute
- Virtualized hypervisor clusters
- Client virtualization with HPE StoreVirtual VSA
- Unified communications such as Skype for Business

## HPE Gen10 servers



Gen10 servers are key to infrastructure modernization

HPE Gen10 servers are industry-standard servers that offer threat protection through a silicon root of trust, extensive standards compliance, and supply chain attack detection. Gen10 servers can recover firmware and operating systems after a denial of service attempt or detection of compromised code.

The HPE Gen10 server families for SMB environments are:

- **ProLiant DL servers**—The ProLiant DL family offers a choice of versatile, resilient servers suited to diverse workloads in data center environments of all sizes. Available in traditional 1U, 2U, and 4U form factors, they offer standard scalability, efficiency, and density suited for most applications.
  - HPE ProLiant DL360 Gen10 server—Provides high density, high performance, and energy efficiency, with flexible storage options and the ability to scale to support various workloads in a single solution. The ProLiant DL360 Gen10 server is ideal for:
    - VMware ESX hosting
    - Dynamic workloads in dense virtualized environments
    - Compute-intensive applications (web caching, data analytics)
    - Low-latency and transactional applications (warehouse and database)
  - HPE ProLiant DL380 Gen10 server—Offers increased performance with HPE Persistent Memory and greater NVMe capability. Designed with a huge storage footprint for large storage workloads and PCIe expansion with HPE FlexLOM and HPE Flexible Smart Array, the ProLiant DL380 Gen10 is ideal for customers with the following workloads or apps:
    - Highly parallel workloads such as image processing, app dev/test, and cloud computing
    - Storage-centric apps such as data warehousing and analytics
    - Customer relationship management (CRM)
    - SAP
    - Large storage capacity such as Microsoft Exchange and NoSQL databases

---

**NOTE:** For more information on HPE ProLiant DL380 Gen10, watch this video:  
<https://www.youtube.com/watch?v=B0J2pITXFzw>

---

- **ProLiant ML servers**—The ProLiant ML family offers efficient, high-performance servers that deliver expandability, manageability, and reliability.
  - The ProLiant ML350 Gen10 server is a 2P server that provides expandability and HPE integrated Lights-Out (iLO) 5 manageability that is ideal for growing SMBs and the remote/branch offices of larger businesses.
- **ProLiant MicroServer Gen10**—This server is well-suited for businesses just starting out or companies with small IT budgets. The ProLiant MicroServer Gen10 is compact, quiet, and designed for small offices and home offices. The ProLiant MicroServer Gen10 offers optimized compute for offices with up to 10 users, low acquisition costs, fast setup, and ease of use. Use cases include:
  - File sharing—Provides up to 32 GB storage, which is two times more than the previous generation.
  - Streaming—Solves media sharing and streaming problems with HPE ClearOS. The AMD Opteron X3000 series graphics capability and support for 4K display resolution is four times denser than 1080p full high definition.
  - Remote management—Offers the option to speed boot up and shutdown times using a solid-state drive (SSD) as a boot drive.
  - Storage—Allows you to customize a fully functional network-attached storage server by layering on the ClearOS open source Linux server operating system. You can easily select an app from the built-in ClearOS Marketplace.

---

**NOTE:** For more information on HPE ProLiant MicroServer Gen10, watch this video:  
[https://www.youtube.com/watch?v=y3cPqV\\_C-CU](https://www.youtube.com/watch?v=y3cPqV_C-CU)

---

- **BladeSystem**—HPE BladeSystem enables customers to transform legacy infrastructure, scale business performance, and optimize costs. BladeSystem solutions are designed to tackle the most demanding workloads including virtualization, especially VDI.
  - HPE ProLiant BL460c Gen10 server blade—Designed for a wide range of configuration and deployment options, the ProLiant BL460c Gen10 server blade provides the flexibility to optimize core IT applications with right-sized storage for the right workload for a lower TCO. Workload and applications support includes the unified API for the application ecosystem, including Docker and others.
- **Synergy**—HPE Synergy Gen10 software-defined solutions feature a flexible dual-mode controller that boosts performance for virtualization and other memory-intensive workloads. This breakthrough performance enables faster business decisions with databases and analytics workloads. Synergy is ideal for customers that need high-speed Ethernet connectivity for bandwidth-intensive workloads such as Telco NFV and increased VDI client density for up to 96 users per compute resource.
  - HPE Synergy 480 Gen10 Compute Module—This composable solution is optimized for any workload. It delivers the capacity, efficiency, and flexibility for data-intensive applications. Graphics processing unit (GPU) expansion modules and more GPU options further enable VDI workloads.

---

**NOTE:** For more information on HPE Gen10 servers, scan this QR code or right-click it to open the hyperlink.



<https://www.hpe.com/us/en/servers/gen10-servers.html>

---

## Learning check

1. Match the HPE server with the workload it is designed to handle.

<b>MicroServer</b>	High density for general purpose applications
<b>ProLiant DL</b>	File sharing and streaming with ClearOS
<b>BladeSystem</b>	GPU expansion modules and options for VDI workloads
<b>Synergy</b>	Highly scalable solution that minimizes cabling costs

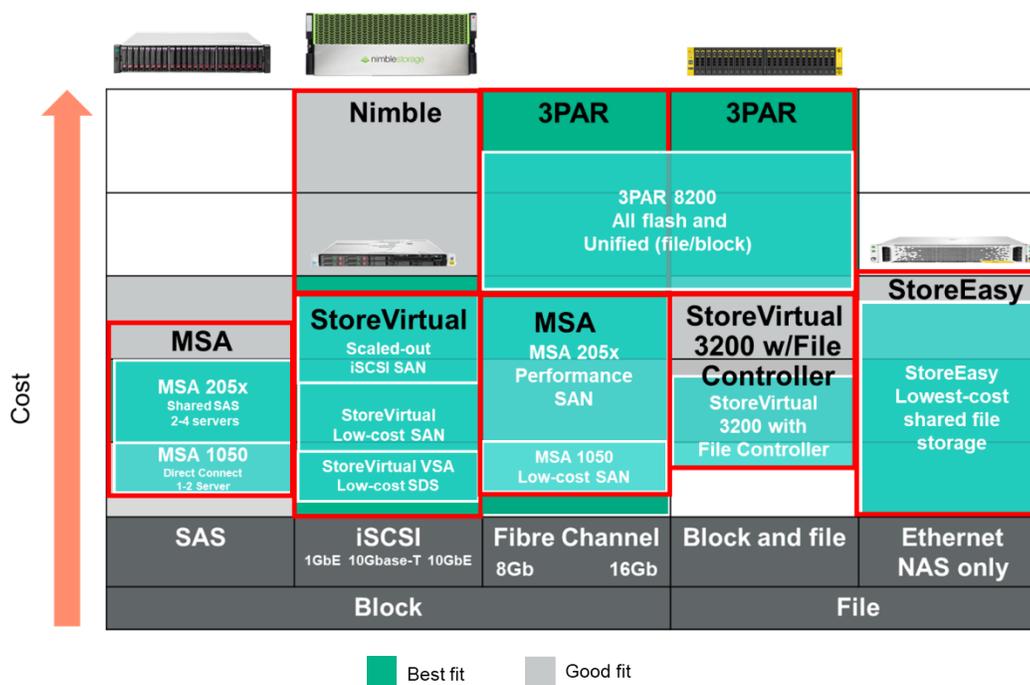
2. Fill in the blanks:

\_\_\_\_\_ and \_\_\_\_\_ are less crucial to a file and print server than they are to a database server.

3. On which considerations should potential SMB server purchases be based? (Select three.)

- Cloud readiness
- Compute density
- Business-critical application requirements
- Supported operating systems and hypervisors
- Support for flexible power options

## Mapping HPE storage solutions to customer requirements



HPE has products to satisfy SMB customers with a range of direct-attached storage (DAS), network-attached storage (NAS), and storage area network (SAN) solutions. Ideal for company-wide deployment and business-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.

The preceding graphic shows that cost is only one factor in the storage system decision-making process. HPE storage systems can be mapped to customers as follows:

- IT generalist**—Looks for the most affordable approach to boost performance in small deployments. This customer needs to consolidate a few apps with good price/performance but limited SLA requirements. This customer is a good candidate for MSA storage systems.
- Virtual machine (VM)/app-centric buyer**—Needs dedicated storage performance and cloud hooks but wants *simple* and non-traditional storage approach. This customer uses storage for VM/Microsoft apps, container databases, emerging applications, and SaaS at service providers and is a good fit for Nimble storage solutions.
- Storage savvy buyer**—Has domain experience in established data center, displacing legacy for flexibility and tier-1 capabilities. This customer uses storage systems for VMs/OLTP analytics, large private cloud computing, performance databases, Synergy Attach, and ITaaS at service providers. This buyer is an ideal HPE 3PAR customer.

## Gauge your knowledge

Before proceeding with this section, answer the following questions, which allow you to identify areas in this training that you should focus on more closely. Pay attention to the topics you are less certain about.

1. How would you explain the differences between DAS, NAS, and SAN?

---

---

---

2. Do you have any experience with storage protocols such as Serial Attached SCSI (SAS), Serial Advanced Technology Attachment (SATA), or Fibre Channel?

---

---

3. Have you ever performed a data backup? If so, what type of technology or hardware did you use?

---

---

4. What more would you like to learn about SMB storage technologies?

---

---

---

Make a note of areas in which you need more information and concentrate on the material covered in those areas. Feel free to skim any material you might already know.

## Nimble Storage



Nimble Storage All-Flash Array

Nimble Storage offers predictive and cloud ready all-flash, hybrid, and secondary flash storage. With Nimble Timeless Storage, there is no need to pay for optional software or major upgrades. All Nimble flash arrays are cloud ready, providing an easy on-ramp to the cloud. Nimble Storage offers several solutions:

- **All-flash arrays**—The Nimble Storage All-Flash Platform is engineered to meet the needs of the enterprise market space as well as the needs of the SMB. Each consecutive member of the All-Flash Platform brings with it more performance and storage capacity. Nimble all-flash arrays combine the speed of flash with the power of InfoSight Predictive Analytics—simplifying operations and delivering 33% to 66% lower TCO than other all-flash arrays. Nimble Storage all-flash arrays support iSCSI and Fibre Channel storage protocols. The all-flash models are positioned as follows:
  - The AF1000 and AF3000 are the perfect entry points for all IT organizations that require speed and economy for performance-sensitive workloads.
  - The AF5000 and AF7000 offer high performance and attractive economics for performance-sensitive workloads that need the best blend of price, performance, and scalability.
  - The AF9000 is designed for consolidating multiple large-scale performance-sensitive applications with aggressive performance and high scalability demands.
- **Adaptive flash arrays**—The Nimble Storage Adaptive Flash array combines a flash-optimized architecture with InfoSight Predictive Analytics for the fastest, most reliable access to data.
  - The CS1000/H array provides value and capacity for small to medium-sized IT organizations or remote offices, for mixed mainstream workloads.
  - The CS3000 and CS5000 are ideal for midsize IT organizations or distributed sites of larger organizations. These arrays offer the best capacity per dollar for mixed mainstream workloads and for virtual server consolidation.
  - The CS7000 offers the highest performance for larger-scale deployments or I/O-intensive mixed mainstream workloads and provides the best performance and I/O per second (IOPS) per dollar. It is designed for consolidating multiple large-scale critical applications with aggressive performance demands.

- **SF-Series Secondary Flash Array**—The Nimble Secondary Flash Array adds high-performance flash storage to a capacity-optimized storage architecture for a unique backup platform that puts backup data to work. The Nimble Secondary Flash Array is optimized for backup, disaster recovery, and secondary data storage. Instantly back up and recover data from any primary storage system. With HPE integration with leading backup software, it simplifies data life cycle management and provides a path to cloud archiving. Using flash makes data work for dev/test, QA, and analytics. The SF-Series is a line of storage arrays optimized for secondary storage tasks and designed for data deduplication to maximize effective capacity:
  - The SF100 is the initial entry-level model, targeted at midsize IT organizations or disaster recovery sites of larger organizations.
  - The SF300 is designed for larger organizations and provides twice the capacity, IOPS, and throughput of the SF100.

---

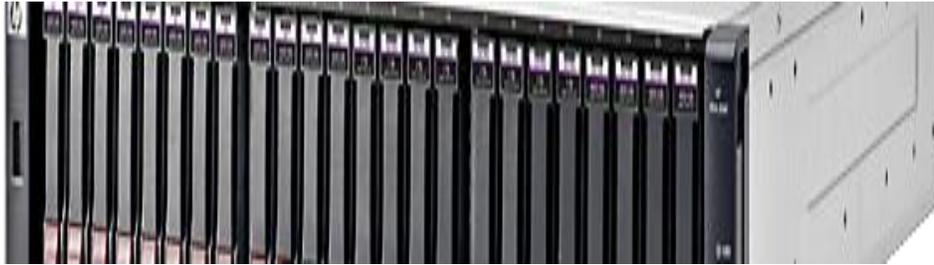
**NOTE:** For more information, scan this QR code or right-click it to open the hyperlink.



<http://www.nimblestorage.com>

---

## HPE MSA Storage



HPE MSA 2042 SFF storage

MSA storage is the ideal solution for SMB customers running Oracle, Microsoft, and SAP as well as customers deploying virtual server technologies such as VMware and Hyper-V. An MSA array enhances virtual environments, simplifies management, and reduces costs. Easy to deploy, scale, and maintain, MSA arrays ensure that crucial business data remains available. MSA storage has been the industry-leading entry storage SAN platform for the past eight years, with nearly 500,000 storage systems sold worldwide.

ProLiant and BladeSystem administrators as well as IT generalists will find storage management tasks simple and intuitive with MSA. Customers are offered operational and business benefits where they can:

- Deploy IT assets across multiple locations
- Incrementally grow storage without interruption
- Enable high availability and disaster recovery capabilities for critical applications
- Deploy a remote disaster recovery site

Products in the HPE MSA storage portfolio are positioned as follows:

- **HPE MSA 1040 Storage**—Designed for entry-level market needs, the MSA 1040 features 8 Gb Fibre Channel, 12 Gb SAS, and 1GbE and 10 GbE iSCSI at previously unavailable entry price points. The array allows users to take advantage of the latest storage technologies such as entry consolidation and virtualization initiatives in simple and efficient ways by providing a good balance between performance and budget resulting in a highly favorable ROI.
- **HPE MSA 2040 Storage**—The Energy Star certified MSA 2040 is a high-performance storage array designed for SMB customers needing 8 Gb/16 Gb Fibre Channel, 1 GbE/10 GbE iSCSI, or 12 Gb SAS connectivity with four host ports per controller. The MSA 2040 storage array provides an excellent value for customers needing performance balanced with price.
- **HPE MSA 2042 Storage**—The MSA 2042 features flash acceleration and 800 GB of flash capacity in all small form factor (SFF) and large form factor (LFF) configurations for improved IOPs performance. In addition to including two 400 GB mixed use SSDs in the base configuration, the MSA 2042 also includes a rich set of standard software features including snapshots, remote replication, and performance tiering capabilities. The price is also affordable when compared to other hybrid configurations. With the Advanced Data Services Software License (ships standard), customers can choose how to best use these SSDs to provide the optimal flash acceleration for their high-performance applications. The MSA 2042 requires little or no management overhead.

- **HPE MSA 2050 Storage**—The flash-ready MSA 2050 storage system is designed for affordable application acceleration that is ideal for small and remote office deployments. It offers a combination of simplicity, flexibility to grow now and into the future, and advanced features you might not expect in an entry-priced array. Customers can start small and scale as needed with any combination of SSDs or lower-cost midline SAS-based drives. The MSA 2050 storage system delivers performance that is double that of the previous generation at the same price. It is affordable flash-ready storage to help customers get the most performance for the lowest cost.
- **HPE MSA 2052 Storage**—The hybrid flash MSA 2052 storage system with ProLiant Gen10 branding is designed for affordable application acceleration that is ideal for small and remote office deployments. Start with 1.6 TB of flash capacity included and scale as needed with any combination of SSDs or lower-cost midline SAS-based drives.

---

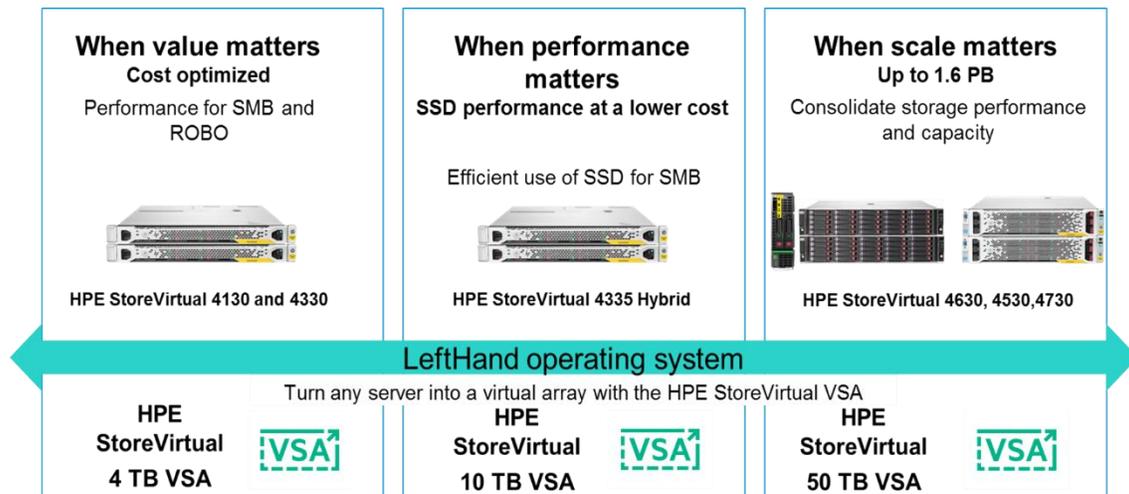
**NOTE:** For more information on the MSA product family, scan this QR code or right-click it to open the hyperlink.



<https://www.hpe.com/us/en/storage/msa-shared-storage.html>

---

## HPE StoreVirtual Storage



StoreVirtual options match various customer priorities

HPE StoreVirtual is a software-defined storage platform that scales in any direction. It provides data mobility across tiers and locations and between physical and virtual storage. It also enables linear scaling of capacity and performance. StoreVirtual systems are flexible, support continuous data growth, and include high-availability features that keep SMBs up and running. Versatile and reliable, StoreVirtual provides affordable storage for a virtualized infrastructure. Its software-defined storage VSA software with ProLiant rack-based and BladeSystem hardware models provide options to fit any infrastructure and budget. Enterprise-class storage software functionality and leading virtualization software integration are built-in. This makes StoreVirtual the ideal platform for supporting virtualization growth at all stages.

Advanced shared storage technology provides the foundation for a composable data fabric, opening up new possibilities for simplified management and scalability across the infrastructure. Software-defined storage controller software enables SMB customers to run enterprise-class storage features on the same set of hardware that also runs the application workload.

Snapshot and Remote Copy integration for virtualization platforms and Microsoft Windows applications make backup and recovery easy. With proven 99.999% availability and unique stretch cluster capabilities, customers can sustain multiple concurrent failures and still keep data online and accessible to applications. Common management and federated data services offer simplicity and flexibility in virtual data centers. With the Centralized Management Console, customers can manage hundreds of nodes in their environment with the same simplicity.

StoreVirtual solutions include:

- HPE StoreVirtual 3200**—With a multicontroller architecture built for multidimensional scalability, the StoreVirtual 3000 storage series is a cost-optimized platform built to address SMB challenges by providing app acceleration on a budget with simple management. Supercharge all workloads with a small amount of flash and HPE Adaptive Optimization software. Increase performance by adding more drives or double performance with a second system. You can also deploy in minutes with a zero learning curve. Powerful, proactive cloud-based analytics make storage optimization simple. App-integrated, app-aware snapshots simplify disaster recovery. Synchronous replication between systems and sites balances cost and resiliency. Easy to deploy, grow, and maintain, StoreVirtual 3200 ensures that critical business data remains available. It provides a unique data protection level across the entire system, reducing vulnerability without driving up costs the way traditional SANs can.

- **HPE StoreVirtual 4000**—Customers need affordable storage designed for a virtualized infrastructure that is easy to manage, supports continuous data growth, and keeps their business operational. StoreVirtual 4000 storage, based on the LeftHand OS, is a scale-out storage platform that is designed to meet the dynamic needs of virtualized environments. Intuitive, common management and storage federation provide simplicity and flexibility in today's virtual data centers. StoreVirtual enables data mobility across tiers, locations, and between physical and virtual storage.
- **HPE StoreVirtual VSA Software**—For developing converged compute and storage solutions in virtualized environments, StoreVirtual VSA software delivers high-performance, shared storage on a choice of servers and SSD or hard disk drive (HDD) media. A virtual appliance increases deployment flexibility and reduces costs through optimized use of storage resources, compute resources, rack space, and power. Built on proven data services technology, StoreVirtual VSA delivers software-defined storage by virtualizing up to 50 TB of disk capacity per server running VMware vSphere, Microsoft Hyper-V, or Linux KVM. StoreVirtual VSA eliminates the need for external shared storage required to implement advanced hypervisor features. StoreVirtual VSA uses scale-out, distributed clustering to provide a pool of storage with enterprise storage features and simple management at reduced cost. StoreVirtual VSA transforms a server's internal or direct-attached storage into a scalable, shared storage array, without dedicated storage.

---

**NOTE:** For more information on the StoreVirtual product family, scan this QR code or right-click it to open the hyperlink.



<https://www.hpe.com/us/en/storage/storevirtual.html>

---

## HPE StoreEasy Storage



StoreEasy storage use multiprotocol access to the same file share over SMB and NFS protocols

Users in many organizations store application data files or documents locally on individual laptop or desktop PCs. This practice is inefficient and has several disadvantages, including reduced mobility and the risk of data loss. HPE StoreEasy systems address these challenges by enabling organizations to store user data with more efficiency and consistency in centralized storage instead of individual user PCs.

StoreEasy storage provides a secure space for employees and teams of small and midsize businesses and distributed organizations to store and share files collaboratively. Powered by ProLiant servers and Windows Storage Server 2016 or 2012 R2, StoreEasy provides the tools and capabilities to help SMBs get the most out of their capacity, spend less time managing storage, affordably and densely scale capacity as they grow, and seamlessly back up data to the cloud using Microsoft Azure.

In heterogeneous environments, client systems running Windows or Mac OS X operating systems generally use the Server Message Block (SMB) protocol to access files on a network storage file server, whereas Network File System (NFS) protocol is typically used in networks with computers running UNIX or Linux operating systems. Sharing data between the different operating systems can be challenging in heterogeneous environments that include both UNIX and PC or Windows hosts. The administrator must consider the different methods of authenticating users, file permissions, and network protocols.

StoreEasy storage can use multiprotocol access to the same file share over both SMB and NFS protocols. Organizations with heterogeneous environments that consist of both Windows and other operating systems can take advantage of this by deploying a StoreEasy multiprotocol solution. Typically, when you deploy a Windows file server in this scenario, you want to facilitate collaboration between users on Windows and UNIX-based computers. When a file share is configured, it is shared with both the SMB and NFS protocols. In this case, Windows users access their files over the SMB protocol, and users on UNIX-based computers typically access their files over the NFS protocol.

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 based on the Advanced Encryption Standard (AES).

File Classification Infrastructure can dynamically identify files based on sensitivity and implement sophisticated access controls based on an organization's requirements using Active Directory Rights Management Services.

StoreEasy allows you to install and run endpoint protection such as antivirus on the system, reducing the cost and complexity of connecting an external endpoint protection server. You can prevent against data loss with Volume Shadow Copy Service online snapshots and simple cloud backup to Microsoft Azure.

## StoreEasy 1000 Storage

StoreEasy 14x0 and 15x0 platforms are ideal for small businesses, workgroups, or remote offices. StoreEasy 16x0 and 18x0 can accommodate medium and large IT environments. Each system ships from the factory with preintegrated hardware and preloaded software to reduce the time and complexity of installation significantly.

StoreEasy 1000 storage systems protect data while at rest with file system encryption and with Microsoft BitLocker drive encryption, while it is being transferred with SMB encryption and signing, and at the physical drive level with a lockable bezel for multi-dimensional file data security.

StoreEasy 1000 Storage increases uptime with standard features including nearly continuous health monitoring with HPE Active Health System, redundant components, file system online self-healing, online verification, and rapid repair that reduces downtime from file corruption to minutes. You can protect data from simple errors or massive failures using built-in replication, up to 100 TB, or use the optional Double-Take Availability replication software (available from a third party).

## StoreEasy 3000 Gateway Storage

With HPE StoreEasy 3850 Gateway, medium to large organizations can leverage unused block capacity for file serving and maximize their SAN investment. A StoreEasy 3850 Gateway system consists of the StoreEasy 3850 Gateway 2U chassis and one or two StoreEasy 3850 Gateway single nodes. A StoreEasy 3850 Gateway System with two nodes provides high availability by clustering the nodes together. Up to four StoreEasy 3850 Gateway systems can be configured together as an eight-node cluster.

A high-performance file serving controller enables a StoreEasy 3850 Gateway to support large numbers of users and connected devices. You can use a range of file access protocols, including iSCSI block access and Microsoft Hyper-V and SQL Server application storage over SMB. In Windows Storage Server 2016-based StoreEasy 3850 Gateways, an enhanced deduplication engine supports large files, up to 1 TB, and file systems up to 64 TB.

Built-in file sync and share enables employees, mobile workers, and distributed teams to access their work files from any Internet-connected desktop, smartphone, or tablet. Organizations can apply policies to wipe devices remotely and enforce lock screen passwords to protect data on employee devices.

Advanced data management enables you to automatically improve use of resources, comply with data retention policies, and enhance protection of sensitive files. You can reduce branch office WAN usage with Hosted BranchCache, which automatically caches frequently accessed corporate data center files locally.

---

**NOTE:** For more information, scan this QR code or right-click it to open the hyperlink.



**<https://www.hpe.com/us/en/storage/storeeasy-file-storage.html>**

---

## HPE StoreOnce Storage



Backup appliances offer affordable and flexible data protection with rapid recovery and application integration

SMB customers need efficient, high-performing, reliable backup systems that are easily integrated into their virtual environments. HPE StoreOnce provides a disk-based data protection platform that addresses data growth by using data deduplication for efficient backup data retention in virtualized and physical environments. StoreOnce systems are built on ProLiant server hardware and use embedded iLO management technology.

StoreOnce systems with StoreOnce Catalyst is the only disk-based federated deduplication solution that provides seamless backup and recovery from the smallest remote sites to the largest data centers. This solution provides a single, high-performance backup architecture that spans the entire data center. StoreOnce federated deduplication is a technology that provides deployment independence to enable movement of data across various systems without reassembling the data back into its original form. Federated deduplication is available across the StoreOnce systems portfolio, including dedicated appliances, virtual backup solutions, and on media servers and application servers with HPE Data Protector and third-party software.

Options for **midsize data centers** include:

- StoreOnce 3520 to scale from 12 to 24 TB raw (7.5 to 15.5 TB usable)
- StoreOnce 3540 to scale from 24 to 48 TB raw (15.5 to 31.5 TB usable)
- StoreOnce 5100 to scale from 48 to 288 TB raw (36 to 216 TB usable)
- StoreOnce 5500 to scale from 60 to 1120 TB raw (36 to 864 TB usable)

Options for **small and remote offices** include:

- StoreOnce VSA licensed at 4 TB, 10 TB, 32 TB, and 50 TB capacity points
- Entry-level StoreOnce 3100 appliance for 8 TB raw (5.5 TB usable)

StoreOnce VSA is a virtual appliance that delivers fast, efficient, and scalable backup. StoreOnce VSA extends the StoreOnce family to enable cost-effective data protection for virtualized environments. By deploying StoreOnce in a software-defined form factor, customers can increase flexibility and cut storage costs compared to deploying purpose-built appliances. StoreOnce VSA is ideal for centrally managed remote office/branch office (ROBO) locations with local backup and off-site backup copies. Branch offices can send backup data to a local StoreOnce VSA target with data deduplication optimized copies to a remote StoreOnce appliance located at a data center or data recovery site.

Replication, licensed by target, is available for StoreOnce systems to provide centralized data protection and disaster recovery for remote offices. StoreOnce Catalyst, licensed by the system or couplet, is one of the most efficient backup interfaces. It simplifies the management and enhances the efficiency of data movement across the enterprise.

The total number of backup targets provided by a StoreOnce backup system is split among 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 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.

In disk solutions, data is backed up from an application server over a dedicated SAN to a disk-based system and from there to a traditional tape library. This provides enhanced solutions for slow servers, single-file restores, and perishable data.

---

**NOTE:** For more information on the StoreOnce product family, scan this QR code or right-click it to open the hyperlink.



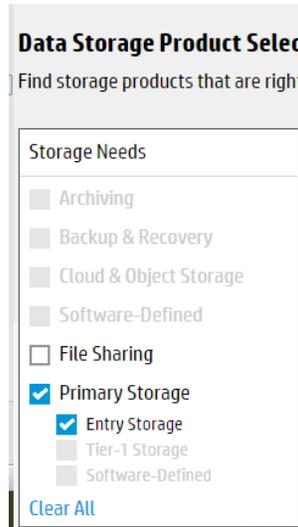
<https://www.hpe.com/us/en/storage/storeonce.html>

---

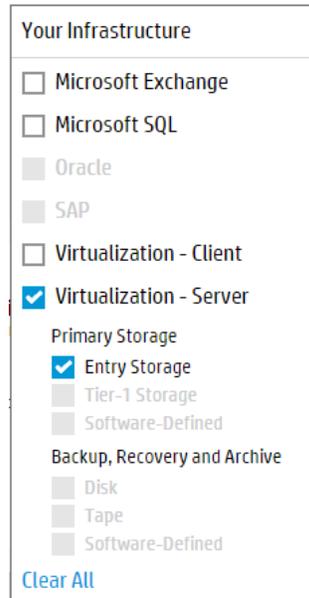
## Activity—Selecting a storage system

During this activity, you will select a storage system for IDJT HealthCare to install in their main data center. The company needs to replace the StoreOnce VSA they moved to their new location. They will use this new storage system as primary storage for frequently accessed files such as patient records. Be sure to record your selection in the worksheet provided in Appendix 2 of this learner guide.

1. To begin, open a Microsoft Internet Explorer browser window and navigate to the HPE Data Storage Product Selector: <http://h22193.www2.hpe.com/data-storage/>
2. From the Storage Needs drop-down menu, select **Primary Storage** → **Entry Storage**.



3. From the Your Infrastructure drop-down menu, select **Virtualization – Server** → **Entry Storage**.



4. Click **Learn More** from the choices you see on the screen.

## Questions

Answer the following questions regarding the storage solution you would select for IDJT HealthCare.

1. What storage solution would you select for IDJT HealthCare? Why?

---

---

---

---

---

2. Should you configure this solution as DAS, NAS, or SAN? Why?

---

---

---

---

---

3. Return to step 3 and select one of the other options from the drop-down menu. How would your selection for IDJT HealthCare change with this option?

---

---

---

---

---

---

**NOTE:** To read a solution brief on why MSA storage solutions are an affordable choice for VMware vSphere deployments, scan this QR code or right-click it to open the hyperlink.



<https://www.hpe.com/us/en/pdfViewer.html?resource=/content/hpe/country/us/en/resources/storage/Solution-brief/4AA6-6914ENW>

---

## Learning check

1. Match each HPE storage solution with the SMB customer it targets.

<b>Nimble</b>	Customers deploying virtualization and looking to simplify management and reduce costs
<b>MSA</b>	Customers requiring data mobility across tiers and between physical and virtual storage
<b>StoreVirtual</b>	Remote sites wanting a federated deduplication solution for backup and recovery
<b>StoreOnce</b>	Businesses needing all flash storage

2. Fill in the blanks:

StoreOnce VSA is a \_\_\_\_\_ that delivers fast, efficient, and scalable backup.

3. Which HPE storage solution family supports 10 GbE iSCSI at an entry-level price point?
- a. MSA
  - b. StoreOnce
  - c. StoreVirtual
  - d. Nimble
4. Which software-defined storage platform scales in any direction and enables linear scaling of capacity and performance?
- a. Nimble
  - b. MSA
  - c. StoreVirtual
  - d. StoreOnce

## Recommending HPE networking for SMBs

Small and midsize businesses are embracing mobile and cloud technologies to improve employee productivity and engage with customers. Businesses are seeking affordable, high-performance, and secure wired and wireless infrastructure that can support the growing number of mobile devices and cloud-based applications. It is helpful to understand which HPE Networking solutions meet different SMB requirements.

### Gauge your knowledge

Before proceeding with this section, consider the following questions. These initial questions help guide you in focusing on areas where you need to learn more. Take time to provide detailed answers for each question.

1. Can you explain the difference between physical (both wired and wireless) and virtual networks?

---

---

---

2. What experience do you have with network switches, routers, or access points (APs)?

---

---

---

3. What more would you like to learn about these three areas?

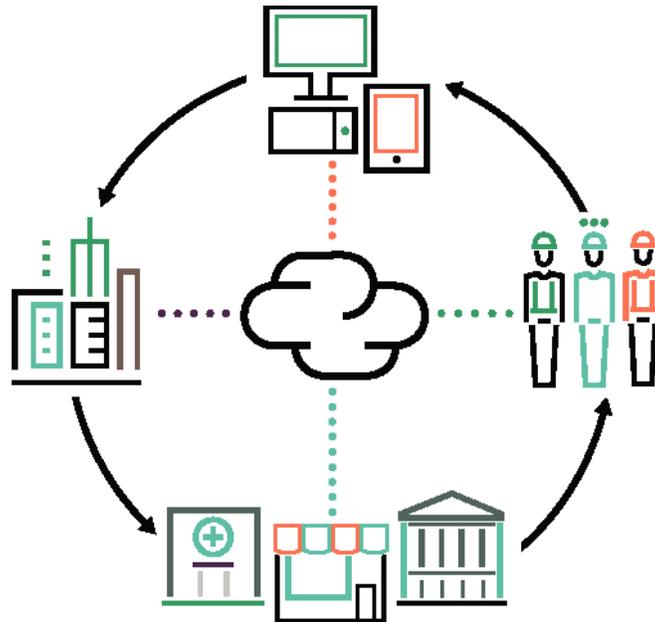
---

---

---

Make a note of areas in which you need more information and concentrate on the material covered in those areas. Feel free to skim any material you might already know.

## SMB networking considerations



Employee productivity and the Internet of Things

The combination of mobile devices and cloud-based applications is changing the way SMBs operate as well as how their employees engage with customers and interact with business processes. To support employee productivity—whether working from home, in meeting rooms, or at client offices—SMBs need a simple and reliable network infrastructure that can be supported by limited IT resources.

SMBs planning a switching refresh need to ensure that their new switches can support not only current application requirements as well as those they are planning for the future. Considerations when determining the best switch for the business include:

- **Access and aggregation**—The first consideration is whether the switch is for network access only or access and aggregation. Because switches provide the network foundation through which users, printers, and other devices connect on the network, you need to understand how many users and devices will be connecting as well as the type of applications and traffic volumes the network will need to support. Switching features vary from basic Layer 2 (local communication) to advanced Layer 3 capabilities that support advanced routing protocols that communicate with users across different buildings. Key areas to consider for selecting a switch are:
  - Scalability—Number of ports/users and Power over Ethernet (PoE)+ power requirements for access points and other IP devices is an important consideration for growth.
  - Performance—Traffic volumes, application types and data transfer speeds will determine the switch capacity needed.
  - Network reliability—Consider stacking for redundancy and redundant power.
  - Routing—Consider Layer 2 only with no routing or basic Layer 3 routing such as static or open dynamic routing or more advanced routing protocols.
  - Manageability—Determine if you need monitoring and remote management for the network via cloud management or multi-vendor network management support.

- **Performance and port speed**—Determine what port speeds the network requires. Consider primary access port speeds as well as uplink speeds. Most SMBs will be shifting to 1 Gbps on edge switch ports, although servers and uplink ports might require much faster port speeds such as 10GbE to handle increased traffic and multiple ports for redundancy.

Consider the type of applications you are deploying such as rich media collaboration applications and WLAN access points. They will need additional throughput as wireless speeds continue to rise. The latest 802.11ac Wave 2 wireless access points can process more than 1 Gbps to connected clients – which means a standard Gigabit Ethernet link may become a throughput bottleneck. To eliminate potential bottlenecks, you should consider switches that support HPE Smart Rate multi-gigabit ports which allow you to increase speeds to 2.5 Gbps, 5 Gbps, and even 10 Gbps over existing cabling.

---

**NOTE:** Wave 1 refers to the first-generation of IEEE 802.11ac products, which use 20, 40, and 80 MHz channels of bandwidth in the 5 GHz band. In some cases, this is combined in the same AP with the older 802.11n in the 2.4 GHz band.

Wave 2 is the second generation of 802.11ac products. The maximum physical (PHY) rate, which affects the data transfer throughput rate, is 2.34 Gbps. For SMBs, Wave 2 offers greater density by supporting multi-user, multiple input, multiple output, meaning the spectrum is used more efficiently for multiple connected devices. Devices can get on and off the network more easily.

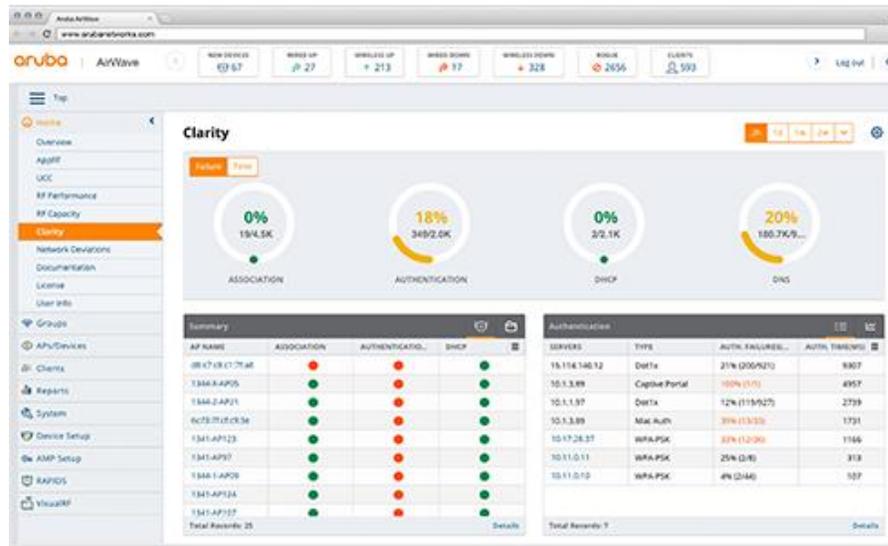
---

- **Port density**—Evaluate port requirements based on the number of users and devices and whether they are connecting via Wi-Fi or wired. VoIP and the Internet of Things (IoT) will increase port consumption.
- **Power over Ethernet (PoE)**—Many devices such as VoIP phones, WLAN access points, and IP video cameras that connect to the access switches can be powered using a PoE interface. New generations of devices such as 802.11ac APs require PoE+, which offer a higher power rating per device. Therefore, it is important to know how much power end devices will actually require so you can select a switch that provides adequate PoE+ power.
- **Stacking and redundancy**—Stackable switches are beneficial where fault tolerance and bandwidth availability are critical. Whether using virtual or physical stacking, a network can recover quickly if a single switch fails. Stacking allows multiple switches connected to each other through Ethernet connections or dedicated modules to behave as a single switch, which means you are managing just one IP device instead of many. This reduces the number of devices you need to manage while increasing network redundancy – thus better utilizing switching capacity.

Also, consider requirements for network availability. Hardware components such as redundant, hot swappable power supplies and management modules can provide additional protection for important aggregation switches.

- **Manageability**—When selecting management options, consider what models and brands are being used in the network. For intuitive Aruba wired and wireless network management, consider cloud-based with Aruba Central or for multi-vendor wired and wireless networks you can opt for on-premises with Aruba AirWave.

## HPE networking solutions address SMB needs



Aruba AirWave helps manage multi-vendor wired and wireless networks on-premises

HPE is a leading provider of networks for small and midsize businesses with:

- HPE FlexFabric**—HPE FlexFabric is designed to address more traditional data center networking architectures and requirements. It offers a full portfolio of enterprise 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 SAN convergence, Multiprotocol Label Switching (MPLS), and support for a wide variety of networking protocols. With FlexFabric switches, customers are never bound to a specific operating configuration, proprietary architecture, or network fabric. 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.

---

**NOTE:** FlexFabric is also a technology used in BladeSystem and is not part of the HPE SMB networking portfolio. Do not confuse these two terms.

---

- Aruba Networks**—For growing businesses that need reliable connectivity for today’s mobile workers, the Aruba Instant controllerless Wi-Fi solution delivers superior performance, business-grade security, and resiliency with the simplicity of zero-touch deployment. Integrated Adaptive Radio Management and patented ClientMatch technologies mitigate interference and intelligently steer clients to the strongest Wi-Fi signal to optimize the performance and reliability of the Wi-Fi network.

HPE offers data center solutions that meet SMB requirements for modern data centers:

- High performance computing
- Resiliency and high availability
- Large Layer 2 networks for virtualized environments
- Control over and visibility into virtual networks
- Network and storage convergence
- Network automation
- Comprehensive switching and routing features

All HPE data center switches provide:

- Enhanced resiliency
- Visibility and control over connections to VMs
- Converged data and storage

---

**NOTE:** For more information on HPE networking for SMBs, scan this QR code or right-click it to open the hyperlink.



<https://www.hpe.com/us/en/networking/smb-networking-overview.html>

---

## HPE FlexFabric



HPE FlexFabric switches

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.

---

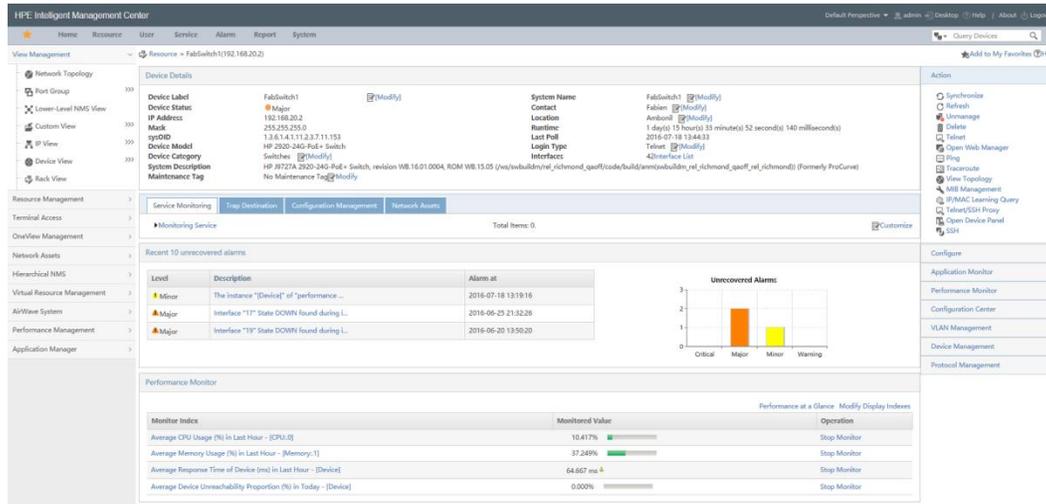
**NOTE:** A converged data center network runs both network and storage traffic over the same infrastructure. HPE converged infrastructure solutions tie virtualization, automation, and unified infrastructure management software together into pre-built, tested, and workload-optimized systems. These systems 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.

## HPE FlexManagement



IMC performs comprehensive management of the network infrastructure

Customers can manage the entire HPE networking solution—from the data center to the main campus to each branch—from the HPE Intelligent Management Center (IMC). IMC is a centralized network management platform that enables users to manage both physical and virtual networks. It uses Simple Network Management Protocol (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 bring-your-own-device (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.

## HPE Switch Selector

### Hewlett Packard Enterprise Switch Selector

Contact us
Share

---

Display  Featured  All Clear All

Switch type  Fixed port  Modular

Port count  5  8-10  16-20  24-28  30-44  48-54  More than 54

Product Line	Routing/Switching	Predominant Port Type	Uplink Speed / Media	Features
<input checked="" type="checkbox"/> Arista	<input type="radio"/> Layer 3 Advanced	<input type="radio"/> 100 GbE	<input type="checkbox"/> 100 GbE	<input type="checkbox"/> PoE <input type="checkbox"/> Replaceable Fans
<input checked="" type="checkbox"/> Altoline	<input type="radio"/> Layer 3 Dynamic	<input type="radio"/> 40 GbE	<input type="checkbox"/> 40 GbE	<input type="checkbox"/> PoE+ <input type="checkbox"/> Stacking
<input checked="" type="checkbox"/> FlexFabric	<input type="radio"/> Layer 3 Lite	<input type="radio"/> 25 GbE	<input type="checkbox"/> 10 GbE Fiber	<input type="checkbox"/> Fanless <input type="checkbox"/> Fiber Channel over Ethernet (FCoE)
<input checked="" type="checkbox"/> Aruba	<input type="radio"/> Layer 2 only	<input type="radio"/> 10 GbE	<input type="checkbox"/> 10 GbE Copper BaseT	<input type="checkbox"/> IPv6 Host / Management <input type="checkbox"/> Ultra Deep Packet Buffers
<input checked="" type="checkbox"/> FlexNetwork		<input type="radio"/> 1 GbE Gigabit Fiber	<input type="checkbox"/> 10 GbE Copper DAC	<input type="checkbox"/> IPv6 Routing
<input checked="" type="checkbox"/> OfficeConnect		<input type="radio"/> 1 GbE Gigabit Copper	<input type="checkbox"/> 1 GbE Gigabit Copper	<input type="checkbox"/> OpenFlow (SDN)
<b>Management</b>		<input type="radio"/> 100 Mb Fast Ethernet	<input type="checkbox"/> 1 GbE Gigabit Fiber	<input type="checkbox"/> Unified Wired-WLAN
<input type="radio"/> Fully managed		<input type="radio"/> All (Chassis)	<input type="checkbox"/> 100 Mb Fiber	<input type="checkbox"/> Redundant Power
<input type="radio"/> Smart managed				<input type="checkbox"/> Redundant Fabric
<input type="radio"/> Unmanaged				

Use the Switch Selector to filter HPE switches by feature

The HPE Switch Selector allows you to view the complete HPE line of switches at a glance. You can also sort the switches by port counts, speed, and more. Compare features and get complete technical specs so you can choose the right switch for your customers. Features include:

- VLAN and connectivity
- Link aggregation group (LAG) and multichassis link aggregation (MLAG)
- Port count and type
- Routing requirements
- iSCSI support
- Interoperability
- Switch virtualization
- Power and cooling requirements

## Activity—Using the HPE Switch Selector

To complete this activity, read the following customer scenario. Then navigate to the HPE Switch Selector using this link: <http://h17007.www1.hpe.com/us/en/networking/products/switches/switch-selector.aspx#.WXLAHogrJbU>. Use the customer’s information to select an HPE network switch for IDJT HealthCare.

### Customer scenario

Eric Yurekha, the IT manager at IDJT HealthCare, has gotten approval from the company’s finance manager to buy two new network switches for the main data center location. He tells you that he wants highly available top-of-rack (ToR) core switches.

“Our network is the lifeline to all the company’s patient records,” he said. “The quality of our patient care depends on it. We can’t afford for the network to go down even for an hour.” He does not outline the budget, but he mentions that Yolonda Arena, the finance manager at IDJT HealthCare, reminded him that the new switches must be compatible with the company’s medical equipment and wireless phone system in the main location. “She is still watching our total expenditures for this project and makes us count every penny,” he said. “She won’t let me hire a network admin, so a switch that helps us simplify the design and operations of our data center would be ideal.”

After talking with Ben Sessions, the systems administrator, about applications the company runs, it becomes clear that late afternoons see a surge in network traffic as nurses enter the patient records information from the day. During the same time period, doctors check their email and return phone calls and assistants record prescriptions ordered as well as medical inventory received. The switches need to handle bursts in traffic and storage easily.

### Questions

1. Based on this information, use the HPE Switch Selector to select an HPE network switch for IDJT HealthCare (<http://h17007.www1.hpe.com/us/en/networking/products/switches/switch-selector.aspx#.WXLAHogrJbU>). Which switch or switches did you select and why?

---

---

---

2. The customer tells you that they are not ready to upgrade to a 40GbE network infrastructure. In the wizard, add 10GbE as a selection criterion. How many switches are available now? Which one would you recommend?

---

---

---

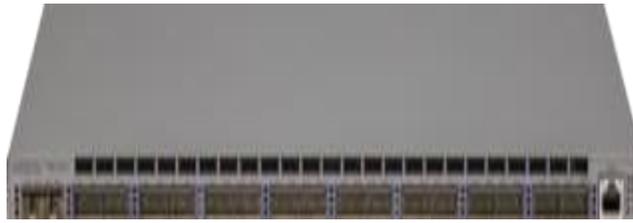
3. What other possible choices could you have made?

---

---

---

## Arista Networks



Arista 7160 Data Center Switch Series

Arista switches are the preferred solution for building high-performance, software-defined cloud networks. Arista Ethernet switches deliver efficient and performance-based 10Gb Ethernet platforms with native support for VMware virtualization, making them ideal for most demanding workloads.

Purpose-built hardware and the Arista EOS network operating system provide:

- Maximum system uptime
- Stateful fault repair
- Zero Touch Provisioning
- Latency analysis
- Fully accessible Linux shell
- Hundreds of Linux applications integrated into the hardware platforms

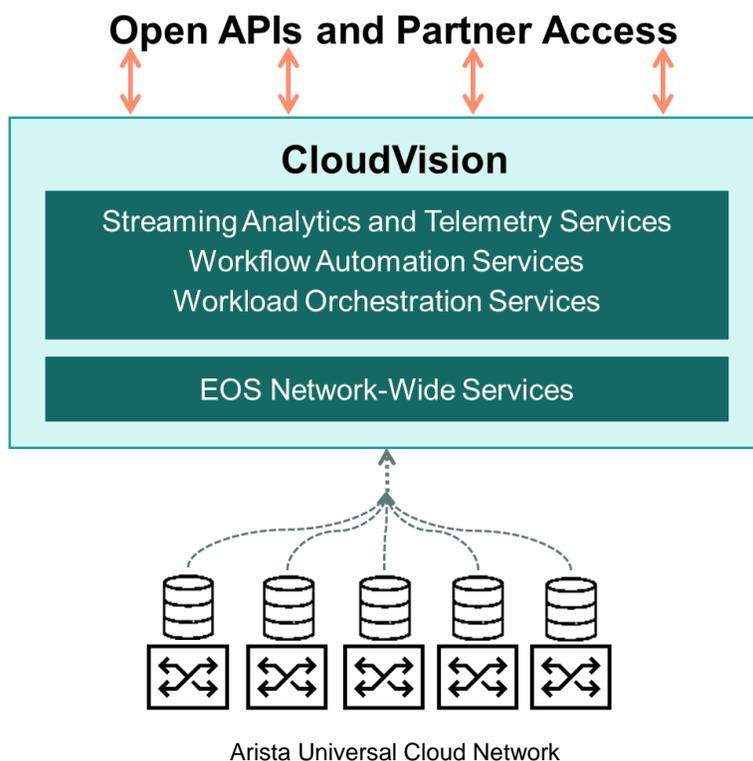
### Deploying an IP storage infrastructure

The cost of deploying and maintaining traditional storage networks is growing at an exponential rate. New requirements for compliance and new applications such as analytics mean that ever-increasing volumes of unstructured data are collected and archived. Legacy storage networks cannot meet the need to scale-out capacity and reduce capital and operational expenditures. In response to this challenge, new storage architectures based on IP/Ethernet have evolved, and are being adopted at an ever-increasing rate.

Technologies such as NAS and iSCSI are widely deployed and Fibre Channel SANs have maintained a dwindling but still strong presence in data center architectures. This paradigm is beginning to change as scale-out storage systems enabled by software-defined storage become increasingly mature. The ability to reclaim stranded DAS assets in a server infrastructure combined with the efficiencies gained by running storage traffic over the same IP/Ethernet network as standard data traffic can reduce the expenses required to deploy and run a storage infrastructure.

Traditional data center networks are not designed to reliably deploy an IP/Ethernet storage infrastructure. Arista products and solutions deliver operational and infrastructure efficiencies for a new IP/Ethernet storage fabric that were previously unavailable. Arista provides solutions that offer a superior price-to-performance ratio when designing networks to support IP/Ethernet storage. An Arista solution ensures effective use of all available bandwidth in a non-blocking mode and provides failover and resiliency when any individual chassis or port has an outage condition. MLAG and Equal Cost Multipath Routing (ECMP) provide standards-based, non-proprietary, multipath technologies at Layer 2 and Layer 3. These technologies currently scale linearly to more than 50,000 compute and storage nodes, both physical and virtual.

## CloudVision



CloudVision is a turnkey solution for network-wide workload orchestration and work flow automation. It was specifically designed to complement software-defined networking (virtualization) controller solutions that orchestrate virtual network overlays by focusing on work flow visibility, automation tasks, and initial or ongoing network provisioning across the underlying physical network.

The foundation of CloudVision is an infrastructure service, sharing, and aggregating working state of physical switches running EOS to provide network visibility and central coordination. The CloudVision components are packaged as a virtual appliance and operate as a highly available cluster with role-based privileges integrated into existing authentication tools (AAA, RADIUS, TACACS).

CloudVision can be managed with the interactive EOS CLI, the open eAPI for granular programmatic access, or a web-based portal interface. The CloudVision web-based portal combines the most common operational tasks into a dashboard view decoupled from the underlying hardware.

Workflow automation in CloudVision permits operators to execute common deployment and configuration tasks from a single visual touch point. The portal includes a turnkey solution for Arista Zero Touch Provisioning (ZTP). It also extends that from automating initial device provisioning to also include automating ongoing change controls over the operational life cycle of the device.

Using CloudVision, operators can organize devices in logical hierarchies through the use of list or configuration container views for rapid categorization of device by role, type, or other specification. Configurations can be broken down into more manageable configlets that are built and stored directly on CloudVision, ready for network-wide or group-specific provisioning. The CloudVision database also keeps historical data, including a history of network state, configuration and software versions. This state can be used for taking a network-wide snapshot for change control verification of the network, helping to simplify the change management process and reduce maintenance window times.

## Aruba Networks



Aruba 2530 Series switch

Legacy network infrastructures were designed before mobility and IoT became prevalent. Today's campus network needs to ensure that connectivity, security and smart network management complement each other so that all of these devices can securely connect. And although most users will connect via wireless, the increase in network density requirements and the need for IoT devices to connect via wired means that the wired infrastructure is just as critical.

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 the wireless network as well as with Aruba's industry-leading security and network management solutions. Aruba switches bring performance and reliability to the mobile-first campus. These industry-leading switches are scalable, secure, and feature HPE Smart Rate multi-gigabit Ethernet ports for high-speed connectivity for access points and IoT devices.

Aruba switches are designed to power and secure the intelligent edge. The Aruba 2540 switch is cost-effective and cloud-manageable, and supports role-based profiles (ACL, VLAN, policing) for the IoT. The higher-end Aruba switches add the ability to do port-based or role-based tunneling to a controller to further secure the network and extend services.

Benefits of using Aruba switches and access points include:

- **Unified wired and wireless**—Aruba's wired and wireless solutions are designed to integrate seamlessly for simple deployment, provisioning and management.
- **Layered security from edge to core**—Security to protect the network and devices. Aruba ClearPass support for Captive Portal and User Role, private VLANs, 802.1X, Web and MAC authentication, ACLs, virus throttling, and dynamic SDN Network Protector App.
- **Flexible management**—Choice of on-premise Aruba AirWave and cloud-based Aruba Central streamline support of Aruba wired and WLAN devices, whether it is a large campus or small remote branch.
- **Investment protection**—No software licensing is required. Users do not need to keep track of complex licenses or licenses for features such as stacking. The warranty includes lifetime phone support with next-business-day replacement shipments. Energy efficiency saves power and ProVision ASIC programmability adds feature enhancements with quick, no-cost operating system updates.

## Activity—Choosing the right Aruba networking solution

To complete this activity, read the following customer scenario. Then navigate to the Aruba Networking Product Wizard for Small and Midsized Businesses by using this link:

<http://www.arubanetworks.com/smb-product-wizard>

The Aruba Networking Product Wizard for Small and Midsized Businesses categorizes network switches according to the vertical markets for which they are ideally suited. Use the customer scenario information to select an Aruba switch for IDJT HealthCare. Be sure to record your selection in the worksheet provided in Appendix 2 of this learner guide.

### Customer scenario

Ben Sessions, the systems administrator at IDJT HealthCare, is happy with the increased productivity the company has experienced with the two switches they purchased for the main data center location. Now he wants to add a managed switch in the remote location to enable employees to access the Microsoft OneDrive the company has set up for prescription documentation and other files. The company also uses Microsoft SharePoint for shared policies, procedures, and other documents. He says that the budget for the switch is small, however, because they already bought the two switches for the other location.

IDJT HealthCare has 150 employees, including doctors, nurses, administrators, and other patient care professionals. In addition, they have 20 IT professionals on staff.

### Questions

1. Based on this information, use the Aruba Networking Product Wizard for Small and Midsized Businesses: (<http://www.arubanetworks.com/smb-product-wizard>) Which switch did you select and why?

---

---

---

2. What other possible choices could you have made?

---

---

---

## Learning check

1. Which HPE centralized network management platform enables you to manage physical and virtual networks from the data center to branch locations?
  - a. ClearOS
  - b. Centralized Management Console
  - c. Intelligent Management Center (IMC)
  - d. Intelligent Resilient Framework (IRF)
2. In what circumstances are stackable switches beneficial?
  - a. Where cable lengths are an issue
  - b. Where budget and space constraints are limiting
  - c. Where fault tolerance and bandwidth availability are critical
  - d. Where the distance from compute resources to client resources is short
3. Which type of devices can be powered using a PoE interface? (Select three.)
  - a. VoIP phones
  - b. IP video cameras
  - c. Stackable switches
  - d. Desktop computers
  - e. WLAN access points

## Recommending HPE data center infrastructure services

HPE offers a range of flexible and customizable services that provides customers with cost-effective support solutions tailored to their specific needs. Typical HPE services for SMBs include HPE Foundation Care, HPE Proactive Care, and HPE Proactive Care Advanced.

### Warranty

HPE warranties its products against defects in materials or workmanship under normal use. Different components might have different warranty levels and durations, and the warranty does not ensure that dedicated parts will be stocked at the local depot. Replacement depends on shipping, and the customer is responsible for installing self-repair parts, without onsite support.

### HPE Foundation Care

	HPE servers	HPE storage		HPE networking
	Entry level	Entry level	Midrange	Entry level
Optimized	Proactive Care	Proactive Care Advanced	Proactive Care Advanced	Proactive Care
Standard	Proactive Care 24x7	<b>Proactive Care</b>	<b>Proactive Care</b>	Proactive Care
Basic	<b>Foundation Care</b>	Foundation Care	Foundation Care	<b>Foundation Care</b>

Recommended service levels are highlighted in bold

With Foundation Care services, customers can save time by calling HPE first for comprehensive coverage. They can meet availability commitments with a variety of coverage levels and response times, and easily connect to HPE for fast problem resolution. Foundation Care offers comprehensive hardware and software services aimed to help increase the availability of the IT infrastructure.

HPE technical resources work with the customer's IT team to help resolve hardware and software problems with HPE and select third-party products. For hardware products covered by HPE Foundation Care, the service includes remote diagnosis and support as well as onsite hardware repair if it is required to resolve an issue. For x86 servers, HPE offers collaborative support along with hardware support that gives a single point of contact for issue diagnosis, troubleshooting, and resolution of application of known solutions, before transferring the case to the third-party software provider. This is provided even if the customer did not buy the software licenses from HPE.

Collaborative support simplifies the support experience and saves time by helping to resolve issues faster. Customers do not have to call multiple vendors or wade through support contacts for troubleshooting. Customers benefit from a single HPE point of contact for hardware and software technical support. Through a simplified process to report IT issues, customers place the first call to HPE no matter what IT issue they are reporting.

For products covered by Foundation Care, HPE offers three distinct service levels:

- HPE Foundation Care Next-Business-Day (NBD) Service
- HPE Foundation Care 24x7 Service
- HPE Foundation Care Call-to-Repair (CTR) Service

For networking products, two additional support levels are offered:

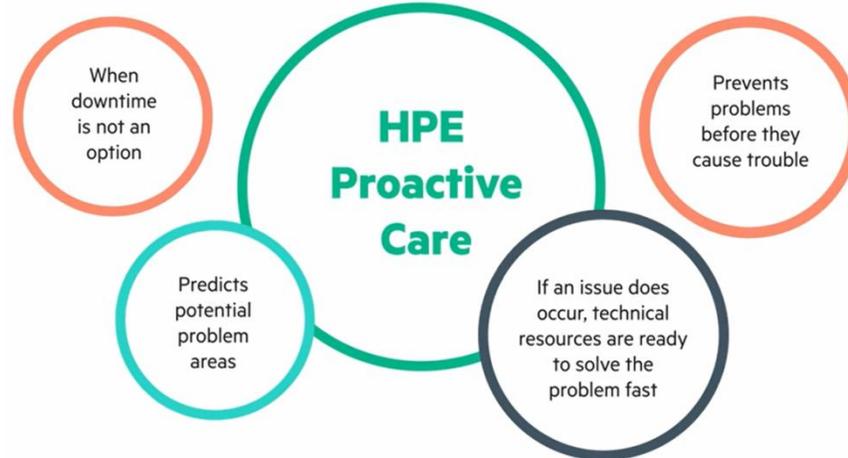
- Foundation Care 4-hour Exchange Service
- Foundation Care NBD Exchange Service

## Warranty compared with Foundation Care

	Warranty	Foundation Care NBD	Foundation Care 24x7	Foundation Care CTR
Priority	Customer is supported at best effort basis.	Customers with Foundation Care have a higher priority than customers at standard warranty.		Customer has the highest priority.
Spare parts	<ul style="list-style-type: none"> <li>The HPE stock of spare parts is not planned in relation to sales volume.</li> <li>There is no time guaranteed for the spare part arrival at destination (best effort).</li> </ul>	<ul style="list-style-type: none"> <li>The HPE stock of spare parts is planned in relation to the total sales volume, but is not customer-based or customer-specific.</li> </ul>	<ul style="list-style-type: none"> <li>The HPE stock of spare parts is planned with consideration of the numbers of customers in a certain geographic area.</li> <li>The spare parts are located in a local repository.</li> </ul>	<ul style="list-style-type: none"> <li>Customer-specific spare part planning is based on the customer's configuration and location.</li> <li>The spare part will arrive at the destination in time to ensure the promised SLA.</li> </ul>
Labor	<ul style="list-style-type: none"> <li>Most HPE equipment components have a Customer Self Repair (CSR) coding.</li> <li>The spare part will be sent to the customer's site so the customer can replace the component.</li> <li>The contingent need of HPE labor and travel will be charged and invoiced.</li> </ul>	<ul style="list-style-type: none"> <li>CSR is first choice but the customer has the right to HPE onsite labor.</li> <li>A technician will arrive according to a best-effort process, meaning customers at a higher service level will be prioritized.</li> </ul>	<ul style="list-style-type: none"> <li>CSR is first choice. Still, the customer has the right to HPE onsite labor.</li> <li>A technician will be onsite within the given time frame, given the customer has been active and of assistance during the remote support process.</li> </ul>	<ul style="list-style-type: none"> <li>There is a guaranteed time to resolution with onsite technician.</li> <li>A technician will be onsite within the SLA time frame, given the customer has been active and of assistance during the remote support process.</li> </ul>
Collaborative support	No	Yes	Yes	Yes
Duration (length)	3 months to 3 years (product dependent)	3 to 5 year support is available upfront; further support extension is available through a contractual or post-warranty renewal until end of support life.		

The preceding table compares HPE warranty with Foundation Care

## HPE Proactive Care



Some businesses are at risk of losing revenue or customers if key systems were to become unavailable. HPE Proactive Care combines proactive and reactive support to help prevent problems from occurring, and to rapidly respond to any problems that do occur. In addition to reactive hardware and software support, Proactive Care includes:

- An enhanced call experience with rapid connection to technical solution specialists (TSS) who will manage calls from start to finish
- Personalized proactive reports with analysis, recommendations, and advice in the following areas:
  - Firmware updates, software patching, and proactive scans (system health checks)
  - Trend and incident reports
  - Enhanced escalation management, if needed

The following hardware support coverage windows and onsite response times are available for both HPE Foundation Care and Proactive Care services:

- **Six-hour, call-to-repair**—This level of support includes special parts handling, call management processes, and staffing. Customers receive access to onsite service 24x7, including HPE holidays.
- **Four-hour, 24x7**—This support level offers an onsite response time of four hours. Customers receive access to onsite service 24x7, including HPE holidays.
- **Next Business Day**—This support level carries a higher priority than warranty calls. Customers get access to onsite service during HPE standard business hours on standard business days. Known as 9x5 support, this service is available nine hours a day, between 8:00 a.m. and 5:00 p.m. local time, Monday through Friday, excluding HPE holidays.

## HPE Proactive Care Advanced

Proactive Care Advanced expands on and includes everything that Proactive Care offers. Proactive Care Advanced is designed to help maximize the benefits from HPE IT investments, maintain peak performance and stability, help achieve business and IT project objectives, reduce operational costs, and allow IT staff to focus on business growth and innovation.

With Proactive Care Advanced, customers work with an assigned, local account support manager (ASM) for personalized technical and operational advice. In addition, customers have access to best practices and expertise across technologies, gleaned from many years of HPE technical experience. The ASM has access to specialist technical skills both locally and remotely when required to help the customer with specialized projects or requirements.

## Learning check

1. What is the recommended service level for entry-level HPE servers?
  - a. Foundation Care
  - b. Proactive Care
  - c. Proactive Care Advanced
2. What is the recommended service level for entry-level HPE storage systems?
  - a. Foundation Care
  - b. Proactive Care
  - c. Proactive Care Advanced
3. Which level of support, available for both Foundation Care and Proactive Care services, includes special parts handling, call management processes, and staffing?
  - a. Six-hour, call-to-repair
  - b. Four-hour, 24x7
  - c. Next Business Day

## Summary

- HPE server, storage, and networking solutions are designed to handle multiple workloads. When planning and developing an IT solution, you must first assess the business and IT needs of the customer. The customer interview can guide your assessment of business and technical requirements. Thorough planning helps to avoid potential costly mistakes and prepares an upgrade path for the future.
- HPE provides several server platforms that are appropriate for SMB environments, including ProLiant, BladeSystem, and Synergy solutions.
- HPE offers a range of DAS, NAS, and SAN solutions. HPE storage solutions for SMBs include Nimble, MSA, StoreVirtual, StoreOnce, and StoreEasy systems.
- HPE offers various networking switches for SMBs, including FlexFabric switches and Aruba Network solutions. These solutions provide an integrated wired and wireless access portfolio, and simplified network management and security solutions to minimize business disruption.
- HPE services can help SMB customers resolve hardware and software problems with HPE and select third-party products. In addition, HPE warrants its products against defects in materials or workmanship under normal use.



# Planning and Designing HPE SMB Solutions

## Module 2

### Learning objectives

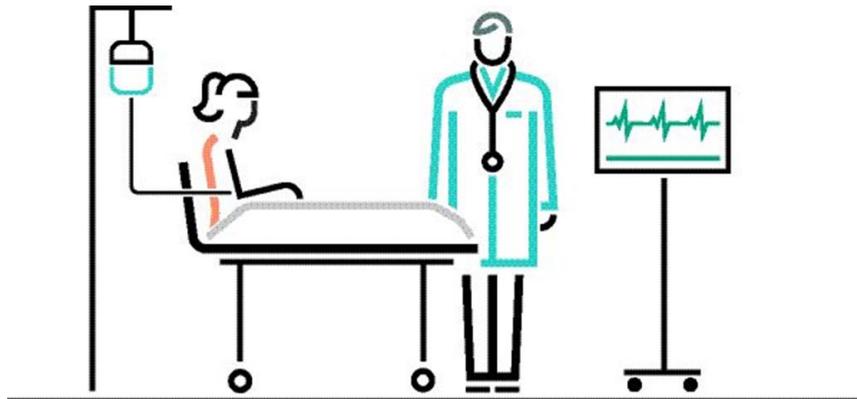
After completing this module, you should be able to:

- Plan a Hewlett Packard Enterprise (HPE) solution for a small to medium-sized business (SMB) customer
- Select a server to meet an SMB customer's needs
- Recommend a storage configuration that aligns with an SMB customer's workload and data center infrastructure
- Design a network based on an SMB customer's existing environment, resources, and workloads
- Validate the design and document the customer's solution

## Customer scenario

Consider the following fictitious customer scenario as you work through this module. Try to plan and design a solution to meet the customer's business and technical objectives. Refer back to it as you engage with activities and labs throughout this course.

### IDJT HealthCare



IDJT HealthCare serves physicians and caregivers across medical focus areas

To prepare for upcoming legislation mandating tighter security of medical records, IDJT HealthCare wants to increase security in their data center. With the planned increase in government regulations, security becomes a more important concern, as does data acceleration for data-intensive workloads. The company plans to open a remote location within the year that will accommodate 10 employees responsible for real-time patient and lab data as well as provider information retrieval.

The IT manager of IDJT HealthCare, Eric Yurekha, has told you that they are interested in upgrading the management server in their main data center location. They currently rely on one HPE ProLiant DL380p Gen8 management server and two ProLiant DL580 Gen9 database servers. They have already moved one of the ProLiant DL360 Gen9 servers to the remote location.

Their budget allows for a server upgrade this year, but they also want room for growth over the next five years, with a plan to add storage next year. Yurekha realizes the need for redundancy and backup but proposes to move forward in two steps—buy and install the server this year, and update their data protection plan next year. They are impressed by the quality of the HPE equipment as well as the support they have received from HPE.

Yurekha says, “We hope to have better security on a dedicated system that will carry us forward as the industry evolves. We hope that the medical records database will be easier and faster to access and use, which will allow the nurses at the hospital to spend more time caring for their patients.”

A site visit to the remote location showed you that the company has installed physical security measures that include an access control system, pin codes, and video surveillance. After your site visit, you tell your customer, “I think you’ll be impressed with the capabilities of the ProLiant Gen10 servers. I will address all your concerns in my proposal.” They seem receptive to hearing more, and you schedule time to present your proposal.

## Gauge your knowledge

Before proceeding with this section, take some time to consider the following questions. Answer each question as completely as possible. Pay attention to the answers you are less certain about.

1. What is the difference between a rack-mounted server and a server blade?

---

---

---

---

2. What are some of the new devices that networks now support?

---

---

3. How would you explain the difference between a switch and a router?

---

---

---

---

Make a note of areas in which you need more information and concentrate on the material covered in those areas. Feel free to skim any information you might already know.

## Planning and designing an HPE SMB data center

Taking a performance-based approach during the designing, constructing, and commissioning of the data center helps determine the resiliency of the facility with respect to levels of redundancy and reliability. Industry and operational standards as well as government regulations are in place to guide the design as well as the daily processes and procedures.

Whether you are designing a server room or a large data center, the architecture must enable the infrastructure to grow with the business. A solution-level approach designed for the company's workloads incorporates Ethernet networking, storage, computing, and security in a way that simplifies implementation.

## Choosing a server



HPE ProLiant DL360 Gen 10

ProLiant server solutions provide a foundation for business success, including security, versatility, compute performance, and quality. ProLiant servers are ideal for SMB workloads and software-specific solutions such as HPE Flex Solutions for virtualization, analytics, and unified communications and collaboration workloads.

Many SMBs have limited IT staff. This necessitates the ability to centrally manage server tasks and ensure continuity of critical data and applications. When faced with decisions regarding upgrading their IT infrastructure, IT staff needs as much help as possible in choosing which servers and components are right for their business environment.

A new generation of capabilities powered by ProLiant Gen10 solutions enables customers to accelerate applications and business insights.

ProLiant rack servers also present solutions for general-purpose workloads. SMB models include the ProLiant MicroServer Gen10, ProLiant DL380 and ProLiant DL360 Gen10 servers, and the ProLiant BL460c Gen10 server blade. All work well in space-constrained environments.

## HPE SMB server comparison

	ProLiant ML350 Gen10	ProLiant DL360 Gen10	ProLiant DL380 Gen10	ProLiant BL460c Gen10
Platform	Tower or rack	Rack	Rack	Blade
Processor	One or two Intel Xeon Skylake series	One or two Intel Xeon Scalable series	One or two Intel Xeon Scalable series	One or two Intel Xeon Scalable series
Processor cores	4 to 28	4 to 28	4 to 28	4 to 26
L3 processor cache	1.375 MB to 20 MB	8.25 MB to 38.5 MB	8.25 MB to 38.5 MB	8.25 MB to 35.75 MB
Max memory	3.0 TB Scalable Persistent Memory	3.0 TB Scalable Persistent Memory	1.0 TB Scalable Persistent Memory	2.0 TB with 128 GB DDR4
Drives	24 SFF/12 LFF HDD/SSD, eight NVMe PCIe SSDs	None ship standard	None ship standard	Two SFF SAS/SATA/SSD or two SFF NVMe (optional) SSDs

Innovations in ProLiant Gen10 two-socket servers, such as the ProLiant DL360 Gen10 and the ProLiant DL380 Gen10, allow you to run diverse workloads and applications across traditional and multi-cloud environments. For example, they feature increased in-server storage density with substantially greater NVM Express (NVMe) capacity for large write-intensive workloads and enhanced storage density with more small form factor (SFF) and large form factor (LFF) drive bays for large database workloads.

An integrated storage solution, HPE Scalable Persistent Memory runs at memory speeds with terabyte capacity using resilient technology currently deployed in your data centers. It includes a DRAM layer to accelerate applications, a dedicated flash tier for persistency, and backup power to facilitate moving data from DRAM to flash. Scalable Persistent Memory enables larger in-memory compute with persistence — up to 27x faster checkpoints operations and a 20x database restart time reduction.

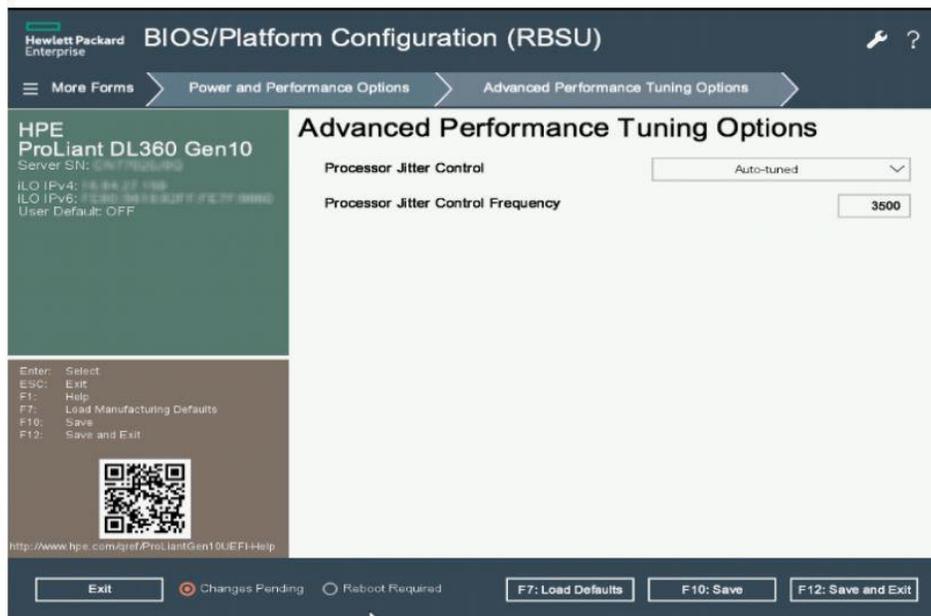
---

**NOTE:** For more information on Scalable Persistent Memory, watch this video:  
[https://youtu.be/1DgP\\_dMoBzw](https://youtu.be/1DgP_dMoBzw)

---

Although no drives ship standard in the ProLiant DL360 Gen10 and the ProLiant DL380 Gen10, they have a maximum storage capacity of 42.24 TB (ProLiant DL360 Gen10 with hot-plug SFF SATA solid state drives [SSDs] with rear drive option) and 197.68 TB (ProLiant DL380 Gen 10 with hot-plug LFF SATA with optional mid-tray and rear LFF drive cage, plus two SFF SSDs rear).

## Intelligent System Tuning



Advanced Performance Tuning Options in Intelligent System Tuning

Intelligent System Tuning (IST) is a set of server tuning technologies in ProLiant Gen10 servers. Developed in partnership with Intel, IST enables you to dynamically tune servers to match the unique needs of each workload. It levels and balances frequency fluctuation, and enables higher performance across fewer processors. IST is available exclusively on Gen10 servers with HPE iLO 5 and an iLO Advanced license or Advanced Premium Security Edition license. It includes:

- **Jitter smoothing**—Levels and balances the frequency fluctuation created when a processor runs in Turbo Boost mode. When a processor executes a frequency change, thread execution stops before the processor can run at the new chosen frequency. For a workload that depends on processor execution, a change in frequency introduces additional latency. Jitter smoothing can be used to tune for best performance in workloads that are traditionally sensitive to latency, as well as workloads that are impacted by excessive amounts of frequency shifting. It is useful for businesses that operate in ultra-high-speed markets such as high-frequency trading.

---

**NOTE:** For more information on jitter smoothing, refer to the *Configuring and tuning HPE ProLiant Servers for low-latency applications* technical white paper on the following website:  
[http://h20564.www2.hp.com/hpsc/doc/public/display?docId=emr\\_nac01804533&lang=en&cc=us](http://h20564.www2.hp.com/hpsc/doc/public/display?docId=emr_nac01804533&lang=en&cc=us)

---

- **Workload matching**—Enables you to choose from preconfigured workload profiles to tune internal server resources automatically and drive up to 9% performance improvement beyond server default settings. These preconfigured server profiles are tuned for the specific workloads most commonly used. Rather than manually tune a server profile for each workload in your environment, you simply select the workload profile you want to deploy. A Gen10 server automatically matches the internal resources to the typical requirements of that workload.
- **Core boosting**—Enables better performance across more processor cores, resulting in significant savings in core-based licensing such as Oracle databases. Core boosting is ideal for virtualized environments and big data or database workloads.

## Choosing server networking components



HPE StoreFabric CN1100R Dual Port Converged Network Adapter

HPE end-to-end server networking solutions can help customers meet the demands of next-generation data centers with features for security, virtualization, scalability, and compatibility.

- **Converged network adapter**—A CNA is a single network adapter that contains both a Fibre Channel host bus adapter (HBA) and a TCP/IP Ethernet network interface card (NIC). It connects servers to Fibre Channel-based SANs and Ethernet-based local area networks (LANs). The ability to carry both Ethernet and storage traffic on a single adapter reduces costs by requiring fewer adapters, transceivers, cables, switch ports, and PCIe slots.
- **LAN-on-motherboard**—LOM technology provides essential network connectivity without requiring an optional network card to be installed in an expansion slot. Although the close-coupled interface of a LOM design leaves standard expansion slots available for expansion functions, it also limits connectivity options. An HPE FlexibleLOM module attaches to a dedicated connector on the server blade or system board.

---

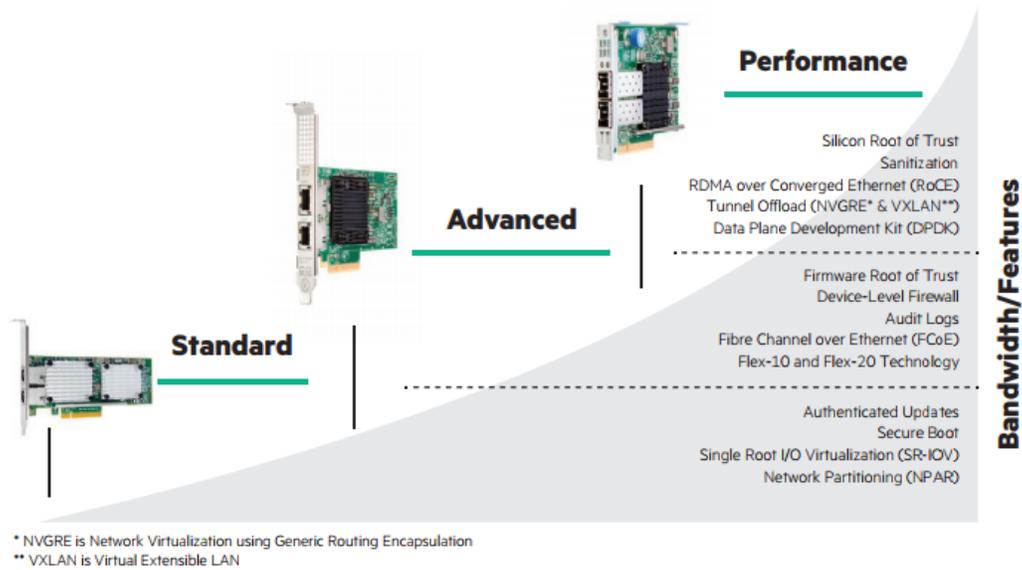
**NOTE:** For details, scan this QR code or right-click it to open the hyperlink.



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

---

## Gen10 server network adapters



HPE server network adapters

HPE server adapters, also known as 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 you 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 I/O bandwidth with lower latency**—Deliver single-root I/O 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 I/O 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**—The Standard series enables a cost-effective Ethernet solution for current server workload needs. The economic scalability of these adapters contains functionality such as single-root I/O virtualization for increased performance via direct access to hardware from a virtual environment.
  - Ethernet—All HPE adapters meet the IEEE 802.3 standard for LANs.
  - Efficiency—HPE NICs meet the price/performance goals for 1GbE and 10GbE core workloads.
  - Virtualization—SR-IOV enables basic virtualization for expanding network fabric.
  - Network partitioning—NPAR allows the user to design right-size data paths for better efficiency.
  - Authenticated updates—Updates can protect user and configuration data from unauthorized access and validate digitally signed firmware.
  - Secure Boot—Secure Boot technology ensures no rogue drivers are being executed on startup.

- **Advanced**—The Advanced series of adapters helps simplify the network and storage topology to build the new hybrid server infrastructure using CNAs. HPE FlexFabric architecture provides the configurable flexibility needed from basic Ethernet functionality, iSCSI hardware offload, and offload for Fibre Channel over Ethernet (FCoE).
  - CNA—Carries both Ethernet and storage traffic on a single adapter. This reduces costs by requiring fewer adapters, transceivers, cables, switch ports and PCIe slots.
  - FCoE— Reduces the number of NICs required to connect to disparate storage and IP networks, reducing the number of cables and switches and reducing power and cooling costs.
  - Device-level firewall—Blocks any unmanaged access to memory or storage. This ensures that on-device firmware and configuration data can only be accessed by authorized agents.
  - Audit logs—Provides traceability into authenticated firmware updates by capturing changes in standard system logs.
- **Performance**—The Performance series delivers even higher bandwidth at a lower latency, with several choices for boosting I/O bandwidth for the most demanding application workloads. These expressly fast adapters can minimize the impact of overlay networking on host performance with tunnel offload support for virtual extensible LAN (VXLAN) and Network Virtualization using Generic Routing Encapsulation (NVGRE). The 25GbE adapters are tuned to work with the latest 25/50/100 Ethernet standards and with HPE top-of-rack (ToR) switches.
  - Multispeed 10/25/50/100 Gbs—The 25GbE adapters enable optimal PCIe slots usage to reduce the total number of slots required to build a 40/50 or 100G infrastructure.
  - RoCE—Network-intensive applications such as 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.
  - Data Plane Development Kit (DPDK)—Allows software-based customization and optimization of network performance by using polling instead of traditional interrupt-driven network processing. DPDK is used in network functions virtualization (NFV) deployments for small packet acceleration.
  - Root of trust—Enables a chain of trust for authenticating updates to firmware. This blocks installation of rogue or corrupted firmware and ensures that the executing firmware is trusted.
  - Sanitization—Renders user and configuration data on the NIC irretrievable so that NICs can be safely repurposed or disposed.

In addition, HPE offers transceivers and cables to complete the data and storage fabric with HPE networking or third-party ToR switches.

Whether customers are purchasing a new ProLiant server platform or upgrading their existing infrastructure with the latest server networking configuration, the HPE ecosystem provides a seamless solution from the switch to the network adapter.

---

**NOTE:** For more information on NIC security, watch this video:  
[https://www.youtube.com/watch?time\\_continue=116&v=sBOfvLX4aG8](https://www.youtube.com/watch?time_continue=116&v=sBOfvLX4aG8)

---

## Learning check

1. Match the server networking component with its description.

**CNA**

System board module that provides network connectivity without occupying an expansion slot

**LOM**

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

2. Match the IST technology with its description.

**Jitter smoothing**

Ideal for virtualized environments and big data or database workloads

**Workload matching**

Useful for businesses that perform high-frequency trading

**Core boosting**

Tunes internal server resources automatically to boost performance

## Choosing a storage solution

The choice of an SMB storage solution depends on the customer's workload and their data center infrastructure. To find the best solution for an environment, consider the interconnect technology, controller technology and features, and drive types and configurations. An SMB storage solution is typically configured in one of three ways:

- **Direct-attached storage (DAS)**—The DAS solution you choose is based on performance, cost, and usability. The advantages of DAS are speed and affordability. Types of DAS include:
  - Storage that connects directly to the drive controller or through an expander to the drive controller
  - Storage within the server chassis or in an external chassis

A DAS solution is simple and cost-effective yet does not sacrifice flexibility or performance. Many HPE DAS solutions can be used not only as a stand-alone direct-attach enclosure, but also for expansion in data protection or big data environments. These storage enclosures can be daisy-chained together to provide support for additional drives.

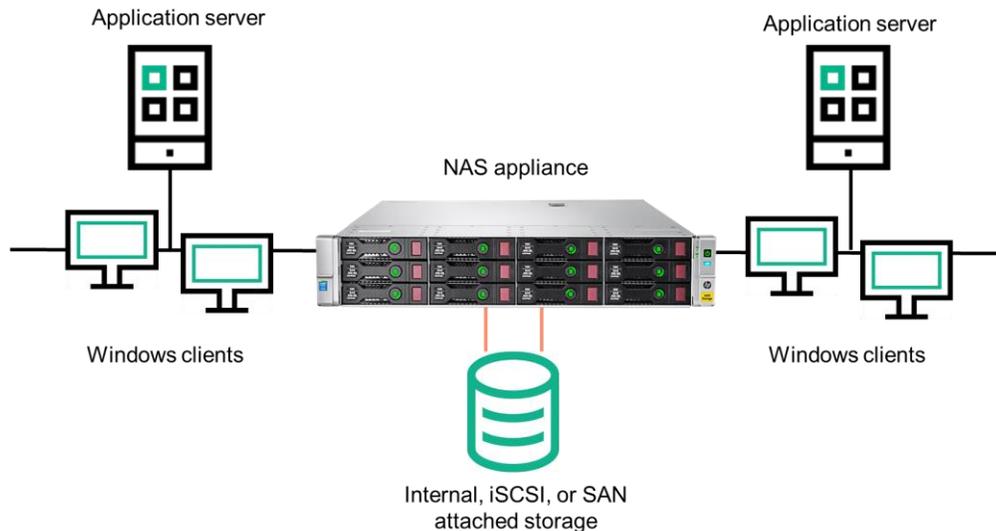
- **Network-attached storage (NAS)**—NAS consists of a specialized server or storage device that connects directly to the network. A NAS solution provides a flexible, intelligent, simple-to-manage solution for file-and-print and application-storage consolidation. NAS devices do not provide any of the typical server activities such as email, authentication, or file management. Instead, the server still handles all data processing but the NAS device delivers the data. A NAS device can be located anywhere in a LAN.

A file system is located and managed on the NAS device. Data is transferred to servers and other devices through the LAN and to clients using industry-standard file-sharing protocols such as Common Internet File System (CIFS)/SMB and Network File System (NFS). The intelligent NAS device enables data sharing among heterogeneous network clients. NAS technology simplifies manageability and improves data access to clients and applications. A NAS solution generally works with a mix of clients and servers running different operating systems. A customer with a small site or remote offices, applications, and a cost-conscious budget would find this type of storage the most cost-effective and efficient for their needs. HPE StoreEasy storage solutions are well-suited for this SMB customer.

- **Storage area network (SAN)**—A SAN 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. It is a solution, rather than a product, using shared network components. A SAN provides block-level storage for environments where high-speed transfer of large files is essential. 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. A SAN delivers better levels of performance, management, and resilience than other storage solutions.

The major difference between a SAN and NAS is that a SAN is a separate network away 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 server with several RAID storage drives that are attached directly to the network. 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, lower total cost of ownership, and relative simplicity. In a SAN architecture, administrators present storage to the server as a logical unit number (LUN) that is displayed 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.

## NAS



NAS makes storage resources more readily available

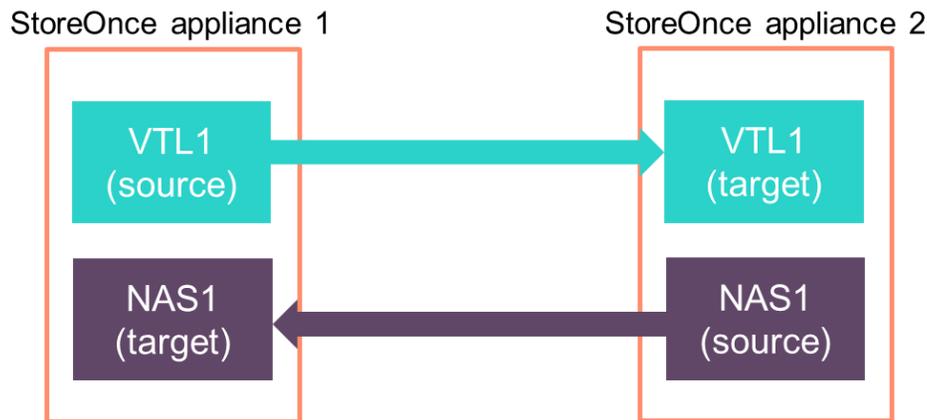
NAS devices are used primarily for unstructured data and file sharing. NAS increases storage management functionality and data protection. It is easy to install and enables data to be shared across an Ethernet network. NAS helps to alleviate the server bottlenecks commonly associated with accessing storage devices.

A dedicated NAS appliance 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. Storage devices attach to network hubs much like servers and other network devices do. All network users have equal access to the stored data and do not have to go through the server.

NAS servers provide a file-level interface to storage subsystems. Because NAS devices are server-independent, they complement and help ease the burden on overworked file servers by off-loading storage to a single, purpose-dedicated storage device. NAS devices have an operating system optimized for file sharing that does not run general server applications, eliminating a major cause of downtime.

NAS storage devices require storage cabinets that provide specialized file access, security, and network connectivity. A NIC on the server is required to access the storage. NAS provides file-to-disk block mapping and client access at the file level using network protocols.

## VTL instead of NAS using a StoreOnce target



StoreOnce in an active/active replication configuration

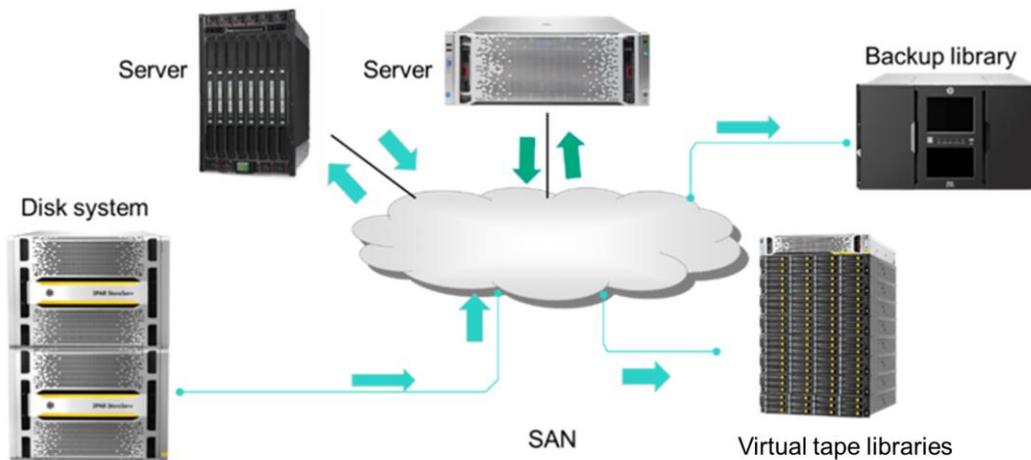
Although it was originally released as a virtual tape library (VTL), an HPE StoreOnce backup system can support both VTL and NAS targets for backup applications on a single platform. A VTL is a disk-based solution that sits between the server and the physical tape drive or library to provide faster data backup, greater data availability, and rapid restores.

StoreOnce backup systems emulate both Linear Tape-Open (LTO) autoloaders and libraries, which are then recognized by the backup application as one or more physical tape devices. Compared to other disk-based data protection solutions, the advantages of using a VTL include the ability to leverage and integrate into current tape backup environments, easy setup and management, and minimal maintenance requirements.

When using a StoreOnce backup system as a NAS target, one or more file shares can be created on each system that are then used by the backup application as CIFS targets for backup. Recognized and supported by a range of backup applications and integrated backup agents, a StoreOnce system used as a NAS target provides a simple and cost-effective method to centralize and automate the backup of multiple servers in a nontape environment.

The StoreOnce backup system is presented to the backup application as a NAS device. A StoreOnce backup system should not be used as a generic file-share device and will not perform optimally if used as such. Because it is designed and optimized for performing backups, it includes data deduplication and support for low-bandwidth replication, as shown in the preceding graphic. Both require a license on the target site, but only if replication is used.

## SAN



SAN architecture

SANs use multiple paths to connect different storage devices with associated servers and can provide backup and archival storage for multiple or remote locations. SANs use storage repositories attached to multiple host servers with centralized storage management and scalability.

General-purpose networks, such as LANs and 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 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 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.

A SAN allows block-level access and offers business continuity with redundant components, automated failover, and centralized management.

## HPE storage networking options



HPE HBAs and Smart Array controllers

HPE offers a variety of storage networking options, including HBAs. HPE HBAs are ideal for internal connectivity to hard drives and solid state drives, as well as external connectivity to tape drives and shared storage. HPE HBAs drive cost-effective and reliable scalability in data centers by providing support for SAS and SATA drives in addition to high performance for SSDs that do not require data protection. HPE HBAs can run in HBA mode or simple RAID mode. Unlike HPE Smart Array controllers, HPE HBAs do not offer acceleration or support cache modules.

HBAs available from HPE are grouped into three categories:

- **Standard: Ethernet**—Increase server optimization with cost-efficient adapters for variable networking workloads.
- **Advanced: FlexFabric**—Prepare for a hybrid server infrastructure by simplifying network and storage topologies.
- **Performance: Multi-speed**—Leverage the latest HPE adapters to boost I/O bandwidth and achieve lower latency for the most demanding applications.

In addition to HBAs, HPE Smart Array controllers can help maximize uptime and boost performance of a server-attached storage environment.

The HPE portfolio of Smart Array Gen10 controllers are designed to enhance server uptime while enabling greater flexibility for future growth. Smart Array controllers can help increase performance when connecting to internal drives on Gen10 servers. Smart Array Gen10 controllers are supported by the HPE Smart Storage Battery, which supports multiple devices and is sold separately.

HPE RAID controllers for Gen10 servers help maximize performance, data availability, and storage capacity. Customers can choose from:

- **Smart Array S-Class software RAID**—Ideal for entry-level solutions that use SATA drives in basic RAID configurations, Smart Array Software RAID delivers the efficiency needed to address evolving data storage needs. Features include RAID levels 0/1/5/10, support for 6G SATA, and access to the Unified Extensible Firmware Interface (UEFI) configuration tool. Smart Array Software RAID is supported on ProLiant rack and tower, BladeSystem, and Synergy compute modules.
- **Smart Array E-Class controllers**—Cost-effective Smart Array E-Class controllers provide simple RAID storage and enable software-defined storage with enterprise-class reliability and security. They are supported on ProLiant rack and tower servers and Synergy compute modules. Key features include RAID on Chip (ROC) and RAID levels 0/1/5/10. They offer flexibility with mixed-mode capabilities, security with HPE Smart Array SR Secure Encryption, and simplicity with the UEFI configuration tool.

---

**NOTE:** For more information on Smart Array controllers, watch this video:  
[https://www.youtube.com/watch?time\\_continue=83&v=ibXenlhEKKw](https://www.youtube.com/watch?time_continue=83&v=ibXenlhEKKw)

---

## Optional storage configurations to consider

In addition to recommending storage devices for your customers, you should also know when to recommend the following configurations:

- **Local boot from disk or flash and boot from SAN**—Traditionally, servers are configured to install the operating system on internal direct-attached storage devices. With external booting from HBAs or RAID arrays, server-based internal boot devices can be eliminated.

When HBAs are configured to boot from the disk device on the SAN, the server boots as if the SAN disk were a local disk. Booting from an external device provides high-availability features for the operating system during the boot process by configuring the HBA BIOS with redundant boot paths. After the link is established with boot LUN, no additional failover is performed.

Booting from SAN provides:

- Improved disaster recovery
- Reduced backup time when the boot path is through a SAN
- Additional SAN-managed features

This process applies to HBAs, and all supported adapter types, including CNAs, LOM modules, and mezzanine cards. For Microsoft Windows servers, only one boot path must be configured.

---

**IMPORTANT:** HPE recommends disabling HBA BIOS on HBAs that are not part of a boot-from-SAN configuration.

---

- **Backup architectures**—When you are designing a storage solution for a customer, consider which type of backup architecture the customer would benefit from.

Depending on time and money, customers can back up their data to lower-cost disk solutions, purpose-built backup appliances, or tape. Note that backing up to a local target does not protect data against a site-wide disaster; disk-based local backups must be replicated to an off-site destination, and tape must be vaulted or transported to a disaster recovery site.

In a SAN environment, I/O congestion is reduced and backup windows are improved because SANs remove backup and recovery traffic from the LAN. System performance is increased dramatically because data and communications traffic no longer compete for the limited bandwidth on a standard LAN. SANs also can be configured for serverless (active fabric) backup. Serverless backup is ideal for large databases or file servers for which there is no backup window.

Serverless backups use Network Data Management Protocol (NDMP), the standard for backing up network-attached storage, and SCSI Extended Copy to move data directly from disk to tape. These methods enable continuous, uninterrupted access to data and applications during backup and restore processes. They also eliminate the read/write processes through the application and backup hosts, resulting in up to 97% less processing power. Removing the LAN, CPU, and I/O resources from the data path eliminates network bottlenecks and maintains high application performance.

- **Deduplication**—Disk-based backup with deduplication is used for high-speed daily backup and restore. It enables you to retain data on disk longer and more cost-effectively. Several deduplication technologies are available in the marketplace. They are generally grouped into hardware-based and software-based solutions. These solutions can be further subgrouped into file-level (single-instancing) or block-level deduplication, for example.

HPE StoreOnce Catalyst provides source-side deduplication, reducing the network load during backup and copy processes. The network load reduction on the first backup of a virtual machine (VM) is lower than for subsequent backups. Starting with the second backup, the network load reduction is measureable.

## Activity—Using HPE SPOCK

In this activity, you will use the HPE Single Point of Connectivity Knowledge (SPOCK) for HPE storage products to select an HBA or CNA for the storage solution you selected for IDJT’s main data center. You will also use SPOCK to ensure interoperability of all solution components.

1. To begin, log on to SPOCK (<https://h20272.www2.hpe.com/spock>). If you do not already have an account, create one.

2. From the main screen, start by selecting the storage solution you selected for IDJT HealthCare.

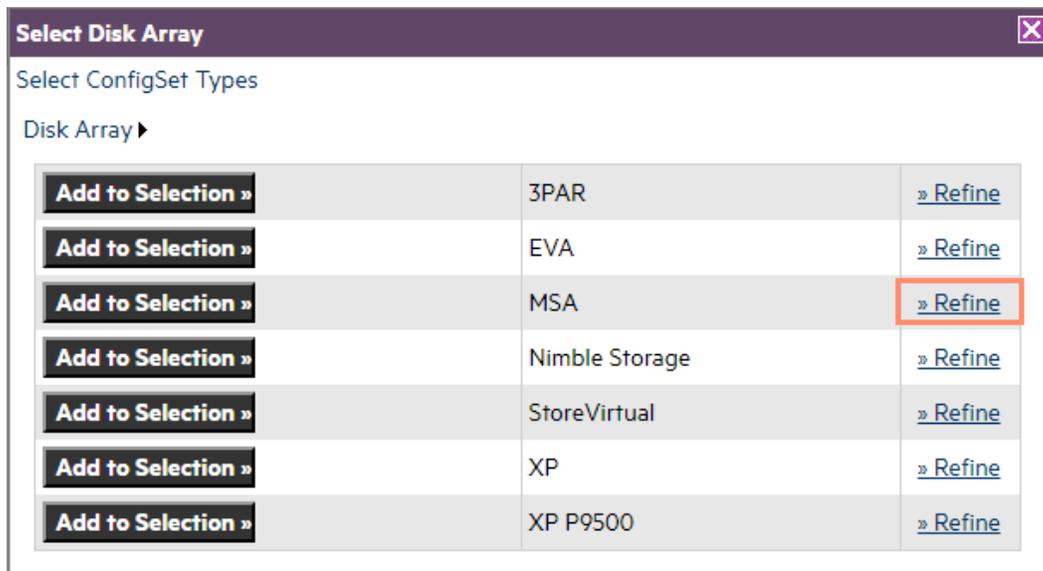
- Click the plus sign next to **Explore MSA array interoperability** and click the various links that appear in the drop-down table. Make a note of which operating systems are compatible with your hardware selections.

	Apple Mac OS X	Citrix XenServer	HP-UX	HP Virtual Machines	OpenVMS	Oracle Linux	Red Hat	Solaris	SUSE	VMware	Windows 2016	Windows 2012	Windows 2008	Windows 7	Windows 8.x
<a href="#">Explore MSA array interoperability</a>															
<a href="#">MSA 2050/2052 FC</a>	✓	✓				✓	✓		✓	✓	✓	✓			
<a href="#">MSA 2050/2052 iSCSI</a>							✓		✓	✓	✓	✓			
<a href="#">MSA 2040/2042 FC</a>	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
<a href="#">MSA 2040/2042 iSCSI</a>							✓		✓	✓	✓	✓			
<a href="#">MSA 1040 FC</a>		✓				✓	✓	✓	✓	✓	✓	✓	✓		
<a href="#">MSA 1040 1Gb iSCSI</a>						✓	✓		✓	✓	✓	✓			
<a href="#">MSA 1040 10Gb iSCSI</a>		✓				✓	✓		✓	✓	✓	✓			
<a href="#">MSA 1040 SAS</a>							✓		✓	✓	✓	✓			
<a href="#">P2000 G3 FC</a>	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		✓	✓		
<a href="#">P2000 G3 iSCSI</a>		✓				✓	✓	✓	✓	✓		✓	✓		
<a href="#">P2000 G3 10GbE iSCSI</a>						✓	✓	✓	✓	✓		✓	✓		
<a href="#">P2000 G3 FC/iSCSI</a>		✓	✓	✓	✓	✓	✓	✓	✓	✓		✓	✓		
<a href="#">P2000 G3 SAS</a>		✓	✓	✓		✓	✓	✓	✓	✓		✓	✓		
<a href="#">Explore P6000 (EVA) array interoperability</a>															
<a href="#">Explore StoreOnce Backup Systems interoperability</a>															
<a href="#">Explore StoreAll Storage Systems interoperability</a>															
<a href="#">Explore StoreEver Tape Automation interoperability</a>															

- From the left-hand pane, click **View by Array** to view the storage solution you selected previously.



- From the popup menu, click **Refine** next to the name of the storage solution you selected.



- From the list that appears, click **Refine** for more details or **Add to Selection** to add the array to your configuration set. Take some time to explore the options available before making your selection. Be sure to pay attention to the choices among SAS and SAN, iSCSI and Fibre Channel, operating systems, and drive form factors.

Select Disk Array
✕

Select ConfigSet Types

Disk Array ▶ MSA ▶

Add to Selection »	HP StorageWorks P2000 G3 10GbE iSCSI Array	<a href="#">» Refine</a>
Add to Selection »	HP StorageWorks P2000 G3 FC Array	<a href="#">» Refine</a>
Add to Selection »	HP StorageWorks P2000 G3 FC/iSCSI Combo Array	<a href="#">» Refine</a>
Add to Selection »	HP StorageWorks P2000 G3 iSCSI Array	<a href="#">» Refine</a>
Add to Selection »	HP StorageWorks P2000 G3 SAS Array	<a href="#">» Refine</a>
Add to Selection »	HPE MSA 2042 SAN Dual Controller LFF Storage	<a href="#">» Refine</a>
Add to Selection »	HPE MSA 2042 SAN Dual Controller SFF Storage	<a href="#">» Refine</a>
Add to Selection »	HPE MSA 2042 SAS Dual Controller LFF Storage	<a href="#">» Refine</a>
Add to Selection »	HPE MSA 2042 SAS Dual Controller SFF Storage	<a href="#">» Refine</a>
Add to Selection »	HPE MSA 2050/2052 SAN Storage (FC)	<a href="#">» Refine</a>
Add to Selection »	HPE MSA 2050/2052 SAN Storage (iSCSI)	<a href="#">» Refine</a>
Add to Selection »	HPE MSA 2050/2052 SAS Storage	<a href="#">» Refine</a>

- From the left-hand pane, click **View by CNA**.

SAN Compatibility

- » Compatibility Tool
  - » View by Array
  - » View by OS
  - » View by FC Switch
  - » View by FCoE Switch
  - » View by SAS Switch
  - » View by Router
  - » View by CNA
  - » View by HBA
- » My SPOCK
- » Configuration Set ZIP Files

- Review the options available and click **Refine** for more details, as needed. Select a CNA to add to your configuration set.

**Select Converged Network Adapter (CNA)**
✕

Select ConfigSet Types

Converged Network Adapter (CNA) ▶

[Show List](#)

<b>Add to Selection</b> ▶	Emulex	<a href="#">» Refine</a>
<b>Add to Selection</b> ▶	HP	<a href="#">» Refine</a>
<b>Add to Selection</b> ▶	QLogic	<a href="#">» Refine</a>

- Repeat this process to **View by HBA** and select a host bus adapter to add to your configuration set.

**Select Host Bus Adapter (HBA)**
✕

Select ConfigSet Types

Host Bus Adapter (HBA) ▶

[Show List](#)

<b>Add to Selection</b> ▶	ATTO	<a href="#">» Refine</a>
<b>Add to Selection</b> ▶	Brocade	<a href="#">» Refine</a>
<b>Add to Selection</b> ▶	Emulex	<a href="#">» Refine</a>
<b>Add to Selection</b> ▶	HP	<a href="#">» Refine</a>
<b>Add to Selection</b> ▶	HPE	<a href="#">» Refine</a>
<b>Add to Selection</b> ▶	IBM	<a href="#">» Refine</a>
<b>Add to Selection</b> ▶	Oracle	<a href="#">» Refine</a>
<b>Add to Selection</b> ▶	QLogic	<a href="#">» Refine</a>
<b>Add to Selection</b> ▶	Sun	<a href="#">» Refine</a>

- Some choices present interoperability information, including additional versions of the component you selected and configurations that include that component. You can also view the results of your choices in a PDF.

**HP H221 PCIe 3.0 SAS Host Bus Adapter**  
 Component Interoperability

The detail below lists the version of the **HP H221 PCIe 3.0 SAS Host Bus Adapter** component you selected from the supported configurations screen. You can view additional versions of the component as they exist in SPOCK, view the results in a PDF document, or view the entire configuration that contains the component version. As you move your mouse over the list of supported configurations, configuration details will be displayed.

**Selected Version:**

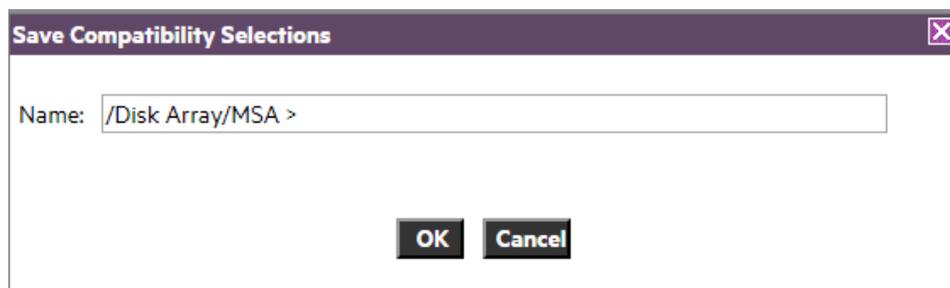
Modified Date	BIOS	Boot Image	Bus	EFI	Firmware Version	HP Part Number	Old HP Part Number	Speed	UEFI	Vendor	Vendor Part Number
4/11/2017			PCIe		15.10.10.00	729552-B21		6Gb		HP	

[» Additional versions of this component \(4\)](#)  
[» Configurations which include this version of the component \(23\)](#)  
[» View results as a PDF](#)

- Continue exploring the options available in SPOCK. Read any notes that are provided. Be sure to click through to copies of the current SAN Design Guide and other configuration guides.



- You can take several paths to arrive at a configuration set. Explore as many as you have time for. When you have finished, save your selections. You can view your final configuration set as a PDF or you can download it as an XML document in a zip file. Use the information you gained from this exercise to add components to your worksheet in the appendix of this learner guide and check their interoperability.



## Learning check

1. What is the major difference between a SAN and NAS?

---

---

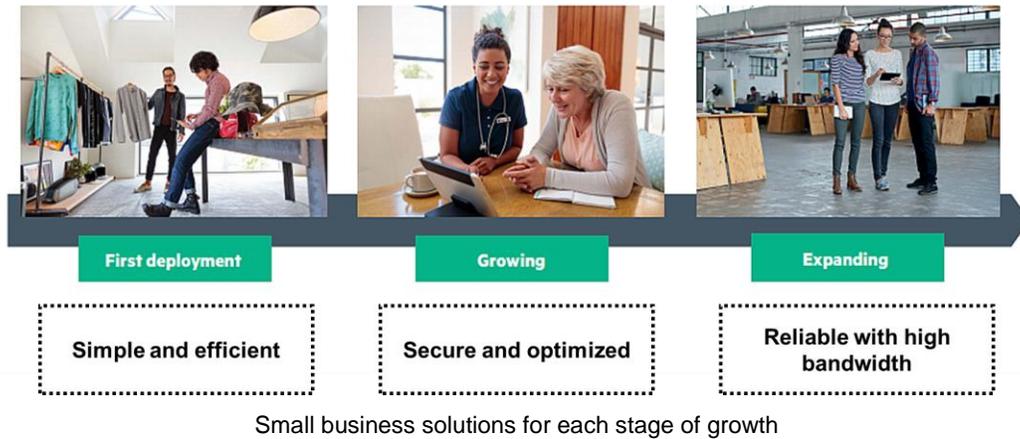
---

---

2. HPE recommends disabling HBA BIOS for HBAs that are not part of a boot from SAN configuration.

- True
- False

## Designing the network



As SMBs embrace mobile and cloud technologies to improve employee productivity and engage with customers, they need an affordable, high-performance, and secure wired and wireless infrastructure that can support the growing number of mobile devices and cloud-based applications. To support employee productivity—whether working from home, in meeting rooms, or at client offices—SMBs need a simple and reliable network infrastructure that can be supported by limited IT resources.

A common concern is having the ability to introduce new vendor hardware into an existing network. Network interoperability can be achieved by using open standards-based access layer solutions, which often involves recommending third-party vendor solutions to the customer. Solution architects must identify compatibility problems and evaluate the pros and cons of each possible solution for a given network environment. Then they are responsible for identifying appropriate solutions, whether to upgrade a current system or replace particular components.

## HPE SMB networking portfolio



Aruba 2930F Switch



HPE FlexFabric 5930 Switch Series

With wireless networks hitting gigabit speed, switches need to keep up with the performance needs of new devices and applications. A wired solution is a critical component to how applications are delivered and perceived by customers. The Aruba 2930F switch series provides a convenient and cost-effective Layer 3 (L3) access switching solution that sets up quickly with Zero Touch Deployment and built-in 1GbE or 10GbE uplinks to deliver right-size network access and performance for mobile campuses.

To avoid potential future bottlenecks, the Aruba 2930M and 3810 switch series enable SMBs to raise the data rates to 2.5Gbs, 5Gbs, and even 10Gbs. With multi-gig HPE Smart-Rate capability, these switches enable SMBs to futureproof their campus network infrastructure as new, higher-capacity wireless technologies emerge.

The Aruba Switch Series includes high-performance L2 and L3 switches that provide robust security, quality of service (QoS), stacking capabilities, energy efficiency, and ease of use for SMB campus deployments. Aruba network solutions fit SMB budgets and IT staff resources, ensuring technology staff can focus on initiatives that add business value.

Some SMBs might need converged data storage, technologies for L2 networks and multipathing, and enhanced resiliency but do not need complete L3 functionality. HPE FlexFabric switches for SMBs provide light L3 functionality and include options for 1Gbs and 10Gbs access connectivity.

The HPE FlexFabric 5920 Switch Series is ideally suited for deployments at the server access layer of large enterprise data centers. It is designed for content delivery networks, especially when they are used to reduce network congestion at the I/O that is associated with the heavy use of server virtualization, as well as bursty multimedia, storage applications, and other critical services.

The HPE FlexFabric 5930 Switch Series provides advanced features and high performance in a ToR data center switch architecture. Consisting of a 1U 32-port 40GbE QSFP+ Switch, a two-slot modular version with two 40GbE ports and a four-slot modular version, the 5930 brings high density to a small footprint.

When planning and designing a data center solution for an SMB customer, it is important to consider how the network is configured in order to provide resiliency and other enhancements at any layer of the network.

---

**NOTE:** For more information on the HPE FlexFabric Network Switch Portfolio, scan this QR code or right-click it to open the hyperlink.



<https://www.hpe.com/us/en/networking/flexfabric.html#portfolio>

---

## Wired and wireless environments



Aruba 2930F Switch

For SMBs with networks that are not robust or flexible enough to keep up with the demands of a mobile environment, there is a simple, affordable option: an integrated network built around 802.11ac LAN technology, a secure and scalable switching solution, and simple cloud management. Although wireless is the primary means by which employees will access the network, the best solution for many SMBs is an end-to-end wireless and wired infrastructure that is easy to deploy and offers built-in security and tools to collect and analyze network performance data.

That solution should be built around an integrated solution of wired and 802.11ac wireless, which offers wireless access speeds three times faster than 802.11n (1.3 Gbs compared to 0.45 Gbs). An 802.11ac access point (AP) has increased signal strength and data range, with 400 Mbps at 75 feet, but 802.11n provides only 200 Mbps at the same distance. 802.11ac APs, when paired with modern switches, improve data throughput, speeds, and reliability while handling more devices simultaneously. To ensure that they operate smoothly, the wired foundation needs to be able to support these access points.

The Aruba Instant wireless APs, when implemented with Aruba switches, provide SMBs with a powerful end-to-end secure network infrastructure that grows with the business. Rogue intrusion attempts on the wireless infrastructure are automatically detected at the edge and eliminated on the wired access ports.

Considerations to keep in mind when creating the foundation for high-speed APs include the following:

- Switching must be gigabit Ethernet at the edge to keep up with increased data flow. Multi-gigabit Ethernet ports provide more capacity using existing cabling when you deploy high-speed Wave 2 802.11ac APs.
- Have 10 GbE from the access to aggregation or core switch to handle the increased traffic, eliminating any bottlenecks. At a minimum, make sure uplinks can support some type of link aggregation protocol such as Link Aggregation and Control Protocol (LACP).
- Switches supporting Power over Ethernet (PoE/PoE+) provide the power for access points. 802.11ac APs use up to 25 watts of PoE, so make sure the switching solution supports IEEE 802.3 at PoE+, which delivers up to 30 watts per port.
- Integration of deployment and management of switches together with APs can simplify the task of keeping the network up and running. A cloud-based management platform is ideal in environments limited staff.

---

**NOTE:** For more information on Aruba mobile networking, scan this QR code or right-click it to open the hyperlink.



[http://www.arubanetworks.com/assets/so/ebook\\_MobileFirst.pdf](http://www.arubanetworks.com/assets/so/ebook_MobileFirst.pdf)

---

## Types of switches

In a broad spectrum of industries—from healthcare to education, from manufacturing to transportation, from banking to government— organizations depend on their networks to deliver fast, reliable access to information. To provide this speed and reliability, networks often are constructed in (typically three) multiple layers:

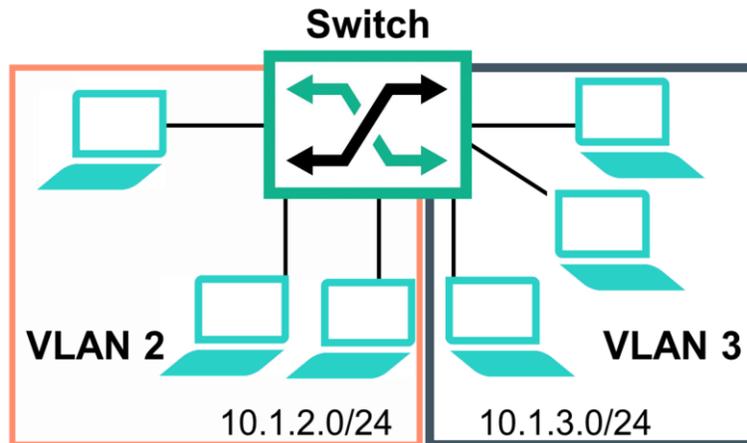
- Access (also known as network edge) layer
- Aggregation or distribution layer
- Network core layer

The access layer is usually a mesh of network switches, linked to other switches in the aggregation layer, which in turn is linked to the core. The lattice of switches provides multiple paths for network traffic, so if one link or switch goes down, traffic can continue to flow using an alternate path.

Switches can be categorized based on their level of manageability.

- **Managed**—Managed switches support Simple Network Management Protocol (SNMP) and allow you to configure each port's communication parameters and many other aspects of the switch through a command line interface (CLI). Many managed switches also provide a GUI such as a web browser interface.
- **Smart web-managed**—Smart web-managed switches, as the name suggests, can be managed through a web browser interface. This interface is designed to be intuitive, making it easy to configure and manage switch features. In addition, because these switches support SNMP, you can manage them through a centralized SNMP console.
- **Unmanaged**—Unmanaged switches provide basic L2 switching and are not configurable. These switches are commonly referred to as “plug-and-play” switches and are designed for SMBs that need basic switch functionality.

## VLAN connectivity



A collection of network nodes are logically grouped to form a separate broadcast domain

A virtual LAN (VLAN) is a group of ports designated by the switch as belonging to the same broadcast domain, or *subnet*. This means all ports carrying traffic for a particular subnet address belong to the same VLAN. Using a VLAN, you can group users by logical function instead of physical location. This helps to control bandwidth usage by allowing you to group high-bandwidth users on low-traffic segments and to organize users from different LAN segments according to their need for common resources.

A VLAN has the same general attributes as a physical LAN, but it allows all nodes for a particular VLAN to be grouped together, regardless of physical location. One advantage of using VLANs is design flexibility. VLANs allow individual users to be grouped based on business needs and protocol types. Connectivity within a VLAN is established and maintained through software configuration.

Ports on multiple devices can belong to the same VLAN, and traffic moving between ports in the same VLAN is bridged (or *switched*). To move traffic between VLANs, a routing process must occur. Basically, the router moves traffic between the VLANs (different subnets or networks). This increases security because you can easily implement access control lists (ACLs) to restrict inter-VLAN traffic flows. Also, assigning IP addresses using an intelligent IP addressing scheme based on the VLAN assignment allows you to easily manage the network.

A *static* VLAN is 802.1Q-compliant and configured with one or more ports that remain members regardless of traffic usage. Static VLANs are configured in the switch memory with a name, VLAN ID (VID) number, and port members. A *dynamic* VLAN is an 802.1Q-compliant membership that the switch temporarily creates on a port to link to another port in the same VLAN on another device.

---

**NOTE:** The IEEE 802.1q standard that applies to VLANs supports a maximum of 4,096 VLANs within a domain. A VXLAN enables you to move beyond this limitation, which is the essential to software-defined networking. For more information, visit this website:



<https://www.hpe.com/us/en/networking/sdn.html>

---

## Types of VLANs

Types of VLANs include:

- **Port-based VLANs**—Port-based VLANs group VLAN members by port. A port forwards traffic for a VLAN only after it is assigned to the VLAN. With port-based VLANs, the network administrator must manually map ports to their associated VLANs. This is the most common implementation of VLANs.
- **MAC address-based VLANs**—This VLAN feature assigns hosts to a VLAN based on their media access control (MAC) addresses. It is used together with security technologies such as 802.1X to provide secure, flexible network access for edge devices including PCs, laptops, IP phones, and so on. With a MAC address-based VLAN, the device receives and processes frames and then looks for a match using the list of MAC-to-VLAN mappings based on the source MAC address of the frame. You can define matches manually from the CLI, where you specify which MAC addresses belong to which VLAN or you can use an authentication server to associate MAC address-to-VLAN mappings. In either approach, you can specify that a range of MAC addresses, such as those for a particular vendor's voice over IP (VoIP) phone, are assigned to the same VLAN.
- **Protocol-based VLANs**—Protocol-based VLAN configuration applies to hybrid ports only. Inbound frames are assigned to different VLANs based on their Ethernet protocol types and encapsulation formats. The protocols that can be used for VLAN assignment include IP, Internetwork Packet Exchange (IPX), and AppleTalk. The encapsulation formats include Ethernet II, 802.3 raw, 802.2 LLC, and 802.2 SNAP. A protocol type and an encapsulation format comprise a protocol template. You can create multiple protocol templates for a protocol-based VLAN, and different protocol templates are assigned different protocol index values. A protocol template can be uniquely identified by a protocol-based VLAN ID and a protocol index combined. When using commands to associate protocol templates with ports, use the protocol-based VLAN ID and the protocol index to specify the protocol templates. An untagged frame reaching a port associated with protocol templates will be processed as follows:
  - If the protocol type and encapsulation format carried in the frame match a protocol template, the frame will be tagged with the VLAN tag corresponding to the protocol template.
  - If the frame does not match a protocol template, the frame will be tagged with the default VLAN ID of the port.
  - The port processes a tagged frame the same way that it processes tagged frames of a port-based VLAN.

---

**NOTE:** A tagged frame is an ID in the data packet on an Ethernet link used to identify the VLAN that the frame belongs to in a network with multiple VLANs.

---

## Port count and type

You can use the console interface on a switch to configure up to 30 port-based, IEEE 802.1Q-compliant VLANs. This enables you to use the same port for two or more VLANs and still allows interoperation with older switches that require a separate port for each VLAN. You can configure three main types of ports:

- Access ports belong to one VLAN; the port is untagged.
- Trunk ports carry multiple VLANs on a single physical link. VLANs are 802.1Q tagged. The native VLAN is untagged.
- Hybrid ports belong to multiple VLANs. They allow multiple VLANs to be untagged and tagged and are typically used for IP phone connection. Hybrid ports are also used together with protocol VLANs and IP subnet VLANs.

## Hybrid switch ports

A hybrid port is similar to a trunk port in that it can support multiple VLANs. Unlike a trunk port, a hybrid port can support multiple tagged VLANs and multiple untagged VLANs. Without a tag, a switch cannot determine the VLAN a frame belongs to if it is only using port-based VLANs. A port-based VLAN is a subset of switch ports that makes up a Layer 2 broadcast domain. Therefore, a hybrid port is typically configured with another type of VLAN. These special VLAN types include:

- **MAC-based VLAN (RADIUS-assigned VLANs for multiple endpoints)**—When a hybrid port is configured to support MAC-based VLANs, it examines the inbound MAC address of incoming traffic to determine the correct VLAN for that traffic. You can set up static MAC-to-VLAN mappings, but typically the switch learns the mappings dynamically from a RADIUS server. Any untagged traffic that is not associated with a MAC-based VLAN is assigned to the hybrid port's port VLAN ID (PVID). Outgoing traffic to these MAC addresses is sent untagged. Broadcasts for all VLANs untagged on the port are also sent untagged.

In the typical use case for MAC-based VLANs, a port connects to multiple downstream endpoints—perhaps through a less intelligent switch—and enforces authentication. The endpoints authenticate to a RADIUS server, which assigns them to VLANs based on the user identity, device type, or other criteria. The RADIUS server informs the switch of the assigned VLAN for each endpoint, and the switch port keeps track of all the source-MAC-address-to-VLAN mappings. In this way, the switch and centralized authentication server can work together to provide flexible, secure access for a variety of devices, including PCs, laptops, phones, and so on.

- **Voice VLAN (VoIP phones)**—Voice VLANs are used with VoIP phones. A hybrid port is required when both a phone and a computer connect to the same port, and the phone sends untagged traffic. The voice VLAN setting helps the port to recognize the untagged VoIP traffic and assign it to a different VLAN from the computer traffic. The switch has a vendor object identifier (OID) list so that it can recognize the source MAC address of VoIP traffic.

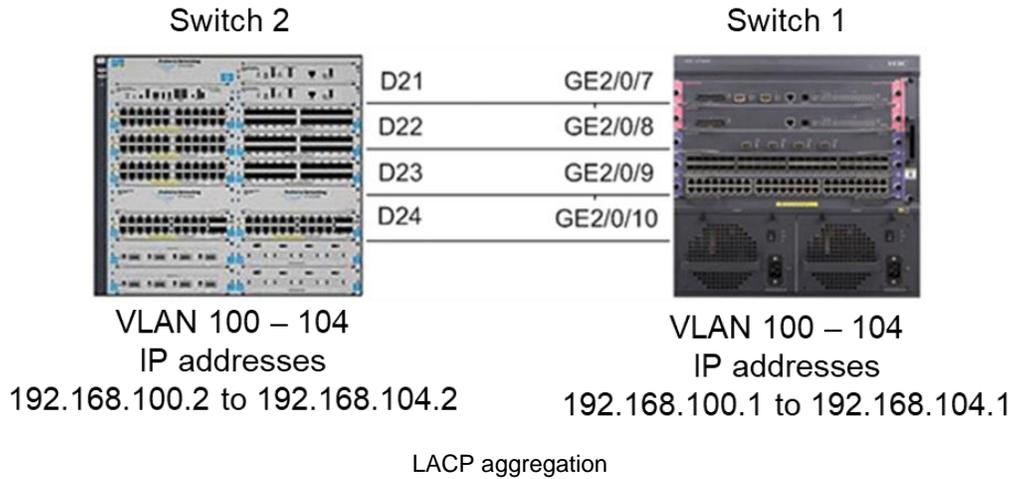
Often VoIP phones tag their traffic. Although you can use a hybrid port in this case, using a trunk port often works just as well and is simpler. The trunk port also supports voice VLAN settings, which might be required for assigning the correct priority to the VoIP traffic or for informing the phone of the correct VLAN ID.

- **IP-based VLAN**—When configured for IP subnet-based VLANs, the hybrid switch port examines incoming untagged traffic and assigns it to a VLAN based on the source IP address. Templates that map subnets to VLANs are created statically on the switch and activated on the hybrid port. The port is also configured to support the proper VLANs as untagged. Any untagged frame that does not match a template is assigned to the PVID.

This feature is not as often required in modern networks. A use case is a provider switch that connects to multiple customer switches that do not support tagging. Traffic from the customer switches arrives on the provider switch untagged, and the provider switch uses the subnets to isolate each customer switch in its own VLAN.

- **Protocol-based VLAN**—The setup for a protocol-based VLAN is similar to that for an IP-based VLAN; however, templates associate protocols with VLAN IDs. This type of VLAN is used to isolate non-IP protocols such as IPX in their own VLAN when devices connected to a port send traffic that use multiple protocols. HPE Comware switches can use a variety of network protocols carried by Ethernet to define protocol-based VLANs. Refer to the switch documentation for a complete list.

## Technologies for configuring the network



Data center networks require high levels of network bandwidth and reliability. There are several technologies you can use to configure the network for redundancy and increased bandwidth. The technology you choose depends on the type of solution you need or the type of switch you select. The choices are:

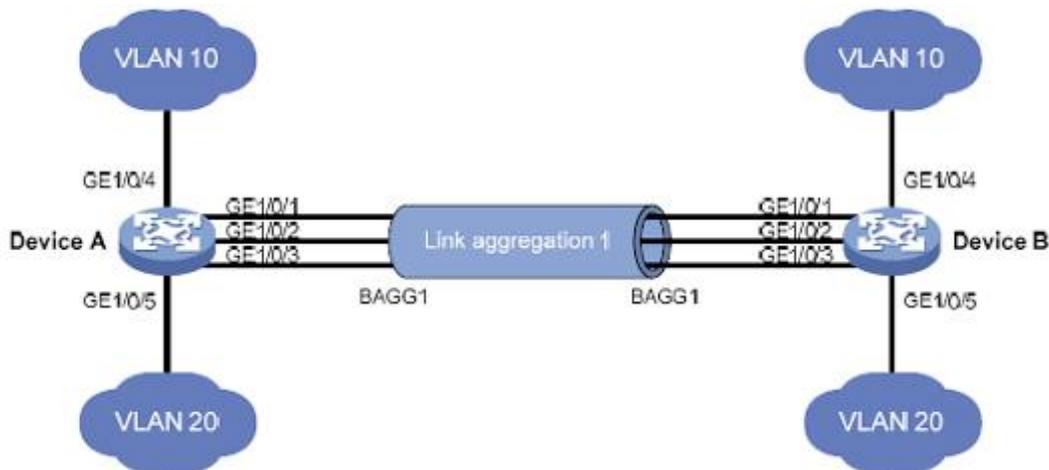
- Link Aggregation and Control Protocol (LACP)
- Multi-chassis link aggregation (MLAG)
- Virtual Switching Framework (VSF)
- Intelligent Resilient Framework (IRF)

---

**NOTE:** Networking vendors frequently use different nomenclature when referring to these technologies. For example, HPE FlexFabric switches use the term *IRF*; Arista switches use the term *MLAG*; and Cisco switches often use *virtual port channels* (vPC). In this course, MLAG is used generically to refer to the concept of aggregating ports across physically separate switches. VSF is similar to IRF on HPE Comware switches.

---

## Link aggregation



Network diagram for Layer 2 static aggregation

Certain connections in a network need more bandwidth, such as connections:

- From the access layer to the distribution layer
- Between distribution layer switches
- Between distribution and core layer switches
- Between certain servers or routers and their connected switches

Link aggregation is an L2 solution that allows you to aggregate multiple L2 Ethernet-based connections between directly connected devices. An aggregated link bundles together multiple Ethernet ports between devices, providing what appears to be a single logical interface.

Link aggregation is supported on many different types of devices, including switches, routers, and firewalls. Most server-based NICs also support this technology. This feature is commonly referred to as *NIC teaming*. Originally, NIC teaming referred to an active and a standby link between a server and the network (typically two different switches). However, most NIC teaming solutions today support link aggregation as well.

A link aggregation group (LAG) is an instance of link aggregation that uses Link Aggregation and Control Protocol (LACP) to signal the establishment of the channel between two devices. A LAG:

- Increases the available bandwidth between two devices
- Creates one logical path out of multiple physical paths

Aggregated links provide these advantages:

- **Redundancy**—If one connection in the channel fails, you can use other connections in the aggregated link.
- **More bandwidth**—Each connection can be used simultaneously to send Ethernet frames.
- **Simplified management and configuration**—Configuration is performed on the logical interface, not on each individual connection in the aggregation.

A dynamic aggregated link is automatically established and maintained by LACP between the two networking devices. A key advantage to a dynamic aggregated link 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, or any protocol, is that it is more complex to configure and typically harder to troubleshoot when two devices do not successfully establish an aggregated link.

---

**IMPORTANT:** Not all server operating systems support LACP. For example, VMware ESX/ESXi 4.0, 4.1, and 5.0 support only static link aggregation (no LACP). LACP support was added in ESXi 5.1.

---

LACP manages the dynamic link negotiation of an aggregated connection. Switches use LACP data units (LACPDUs) to exchange information and establish a dynamic aggregated link. These LACPDUs include information such as:

- Source MAC address of the device.
- A system identifier. This has two parts, the LACP system priority and the switch's MAC address.
- LACP system priority. The system priority determines which switch will select the ports that are active in the aggregated link. Smaller numbers have a higher priority.
- Port priority. This is used in the process of determining which ports are active in the aggregated link. Smaller numbers have a higher priority.

Exchanging LACPDUs allows devices to determine if the links can be aggregated. For example, the 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.

## Link aggregation requirements

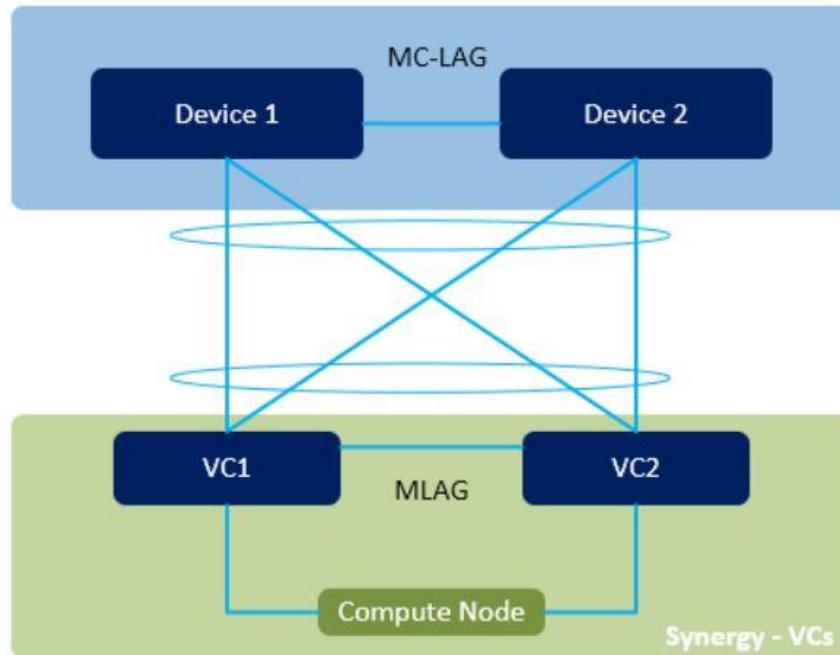
Interfaces in an aggregated link typically must be configured identically. Speed, duplexing, and VLAN settings must be the same. The protocol used to establish the connection affects what must match between the two devices.

When setting up an aggregated link, you can use up to eight interfaces bundled together, depending on the switch model:

- Up to eight active 1G Ethernet connections, providing up to 8 Gbs
- Up to four active 10G Ethernet connections, providing up to 40 Gbs

A link aggregation can only support a certain number of links per link aggregation. The number of links allowed in an aggregation depends on the switch model. For example, many switches have a maximum of eight links. Refer to the documentation for each switch for more information.

## MLAG



High-level, logical depiction of multi-chassis link aggregation

Multi-chassis link aggregation (MLAG) extends the capabilities of link aggregation by allowing a downstream switch or host to connect to two switches configured as an MLAG domain. This provides redundancy by giving the downstream switch or host two uplink paths as well as full bandwidth utilization because the MLAG domain appears to be a single switch to Spanning Tree Protocol (STP). Because the MLAG domain appears to STP as a single switch, there are no blocked ports.

MLAG enables two or more switches to function as a single switch when forming link bundles. This is useful when presenting different physical paths to hosts because it enables a host to uplink to two switches for physical diversity, but still have just a single bundle interface to manage. Also, two switches can connect to two other switches using MLAG with all links forwarding. An MLAG supports link aggregation from a server to two upstream switches, which a LAG does not support.

MLAG typically uses LACP to negotiate traffic between a host and an MLAG virtual switch or between MLAG virtual switches. MLAG spreads the benefits of link aggregation across a pair of data center switches to deliver system level redundancy as well network-level resiliency. Any networking device supporting static LAG or IEEE 802.3ad LACP can connect to an MLAG pair.

LAG and MLAG technologies are managed by a CLI with support for existing network management tools and processes.

## MLAGs compared to traditional stacking

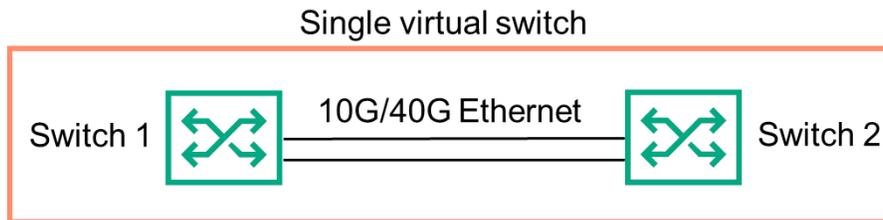
Traditionally, a network stack is a group of up to 12 same-model switches configured in a ring or linear topology. A traditional stack can contain switches that are physically connected in the same physical location. Alternatively, switches in the stack can be physically separated; the distance between them depends on the type of connector cables used. Stackable devices are connected either through dedicated stacking ports or through ports configured for either stacking or data. Even when all the switches within a stack are physically distributed, you can manage them as a single device, so you can make configuration changes easily via a single IP address. An example of a network stack is a group of ToR switches acting as a single switch to manage data center access.

Stacking involves a single control plane, which limits resilience. Although this is a benefit in MLAG, both MLAGs and stacks are independent of the control plane. Some MLAGs are the result of a unified control plane. In some Ethernet switch stacking schemes, the control plane of physically distinct switches is controlled by a master switch. The result is a virtual switch composed of several switches that function as one unified switch. IRF is an example of this.

Some MLAGs maintain separate control planes. In some MLAG schemes, two switches maintain their own control planes and are managed separately from one another. In these schemes, MLAG is a specific feature that must be enabled and configured on each switch that is intended to participate in the MLAG scheme. Design considerations for this configuration include:

- Dedicated links and VLANs between the switch members
- Rules about traffic forwarding for loop avoidance
- Changes to the overall STP topology

## Virtual Switching Framework



Virtual Switching Framework (VSF) is a stacking technology that provides scalability, resiliency, and simplicity. It enables redundancy in ToR configurations or wherever you need availability without multiple physical switches.

---

**IMPORTANT:** VSF is disabled on the switch by default.

---

VSF allows multiple switches that are connected to each other through Ethernet connections to behave like a single switch. VSF virtualizes multiple physical devices into one virtual fabric so you can manage one IP device instead of several. In addition to ease of management, VSF brings even more benefits:

- **Added resiliency**—When a VSF member switch fails, the other members take over. This process reduces recovery time, which is important for business-dependent network services.
- **More performance**—A VSF stack creates a pool of network capacity with optimized forwarding so any member of the stack can meet network demands.
- **Simplified management and configuration**—Even though VSF is made up of multiple physical devices, a VSF stack is managed as a single IP device with a single image. Working with the entire stack as a whole simplifies network design and management.
- **Convenient deployment**—VSF stack members can be located where they are needed because they are linked with the required 10GbE or 40GbE connections, module, cables, or licenses.

In addition, VSF supports VSF link aggregation. Using VSF eliminates the need for L2 redundancy protocols such as STP and L3 redundancy protocols such as Virtual Routing Redundancy Protocol (VRRP).

Configuration guidelines for VSF are:

- Support is currently limited to two members (VSF 16.x.x or later).
- Only the same model switches can join a VSF system.
- VSF links are supported on 10G and 40G Ethernet interfaces only (no 1G).
- Each switch supports only one logical VSF link.
- Logical VSF links support up to eight physical ports, and the ports can reside on different modules.

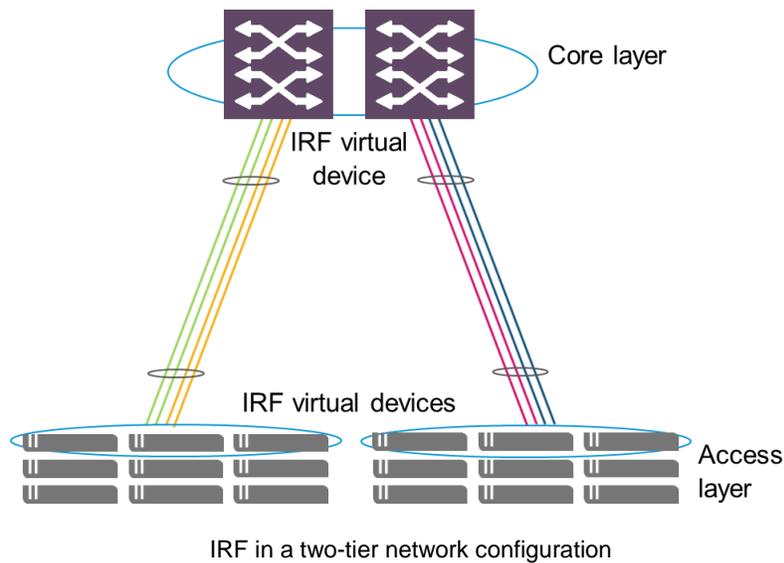



---

**CAUTION:** To avoid broadcast storms or loops in a network when configuring VSF, HPE recommends first disconnecting or disabling all ports you want to add to or remove from the VSF. After you finish configuring the VSF, enable or reconnect the ports.

---

## Intelligent Resilient Framework



HPE Intelligent Resilient Framework (IRF) helps simplify networks by allowing customers to combine multiple switches into a single, ultra-resilient virtual device. IRF virtual devices eliminate the need to run outdated protocols that slow down the network and waste valuable network bandwidth. IRF virtual devices provide the high performance and rapid failover required for today's delay-sensitive, business-critical applications. IRF offers a simpler architecture, faster failover, and greater resilience, which contributes to marked improvements in availability over architectures based on STP and other conventional network protocols.

Because it allows redundant links to be active without creating loops, IRF technology can use the full capabilities and bandwidth of each switch, ensuring greater overall efficiency and maximizing return on the investment in networking infrastructure. With IRF, businesses no longer need to connect to, configure, and manage switches individually. After the primary switch is configured, that configuration is distributed to all associated switches automatically, considerably simplifying network setup, operation, and maintenance.

By allowing a flatter, fully resilient topology, IRF can reduce the number of components in the network. This capability lowers costs for power and cooling, greatly simplifies network planning and operation, and frees staff for more productive tasks. The result is reduced overall operating and administrative expenses.

IRF implemented on FlexFabric 59xx switches forms a single IRF fabric consisting of from two to nine physical switches in one IRF domain. To other devices on the network, each IRF system is presented as one device with one MAC address and one bridge ID. Routing updates originate from this one device.

The IRF system draws on each switch's capabilities during normal operation. As a result, the IRF system provides high performance and greatly simplifies the design and operations of data center and campus networks. In addition, the IRF system provides both device-level and link-level redundancy. If a switch (or a switch component) fails or becomes unavailable, the IRF system can quickly and seamlessly fail over, preventing service interruption and guaranteeing complete continuity for business-critical applications. IRF technology virtualizes up to four chassis.

## Learning check

1. What is the difference between managed and unmanaged switches?

---

---

---

---

2. Name three types of VLANs.

---

---

---

3. What is the generic term that refers to the concept of aggregating ports across physically separate switches?

- a. LACP
- b. MLAG
- c. NIC teaming
- d. Spanning Tree

## Additional infrastructure considerations

When planning and designing a data center infrastructure for a customer, you should think about their current equipment needs as well as what they will be adding in the future. When you are choosing power, cooling, and rack components, there are many considerations to keep in mind.

---

**NOTE:** To access the HPE Rack and Power Infrastructure Architect tool, scan this QR code or right-click it to open the hyperlink.



<https://iallb.itcs.hpe.com/>

---

## Power and cooling

Power and cooling are major cost factors in data centers. When the equipment in data centers grows more diverse and more powerful, it generates more heat. Technologies supporting the deployment of IP telephony and wireless networking are additional contributors to heat generation in the data center. And because LAN switches are always on, they can be one of the biggest sources of heat generation. Server blade technologies and other high-density equipment configurations might also increase the need for effective cooling.

The power required to cool a data center can meet or exceed the power used to run the IT equipment, including the hardware running high availability power protection through an uninterruptible power supply (UPS) system. Data center capacity can be limited when cooling becomes a bigger problem than getting enough power to the equipment. Implementing cooling efficiently can save operational expenditure costs over time.

---

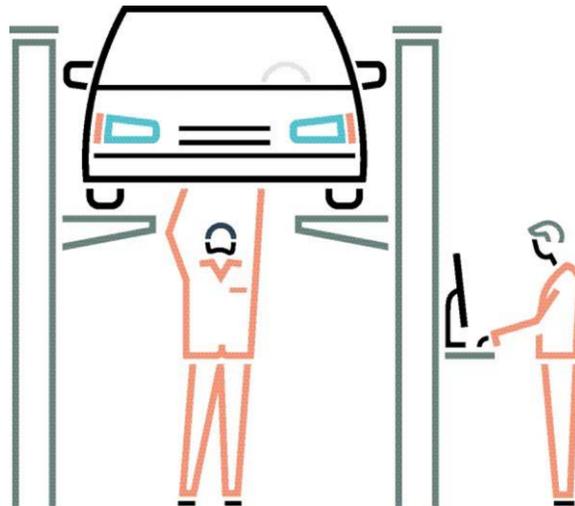
**IMPORTANT:** Rack-optimized equipment should be installed to operate at high-line voltage (200–240 V AC).

---

Configuring an appropriate data center solution involves understanding equipment power consumption and the need for redundancy. Poor planning can result in failures and downtime. When preparing a power planning strategy, consider ways to minimize power consumption and heat generation in the environment:

- Estimate power consumption and execute the heaviest possible workload under the worst-case operating conditions.
- Determine current and future space, power, and cooling requirements and plan for the addition of new equipment to the facility.
- Ensure good ventilation (bypass and recirculation) in the equipment room and keep the air inlet and outlet vents free from obstruction.
- Understand the thermal effect of changes to the environment by using what-if scenarios (changing cooling set points, implementing containment solutions, and so forth).
- Consider the placement of IT, cabling, power, and cooling equipment within the facility.
- Select the power distribution unit (PDU) to deliver the proper redundancy. The demands of the equipment should match the capabilities of the facility.

## Proactive maintenance



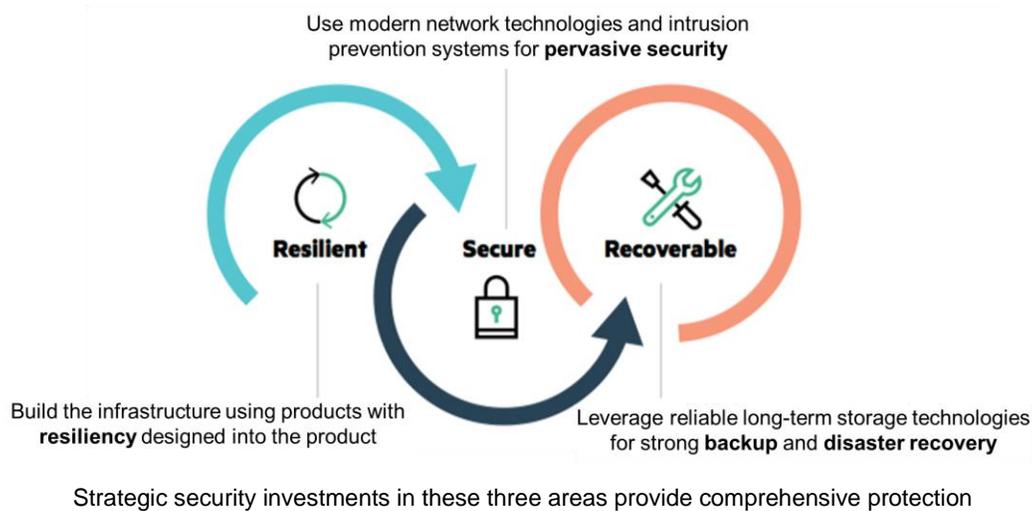
Improve troubleshooting, problem resolution, and operations by developing a maintenance plan

When planning a new solution, make sure that maintenance and troubleshooting are built into the processes. This means accounting for the potential for downtime or lost data. With foresight, these possibilities can be minimized. Make sure that your planning includes each of the following:

- **Rollback plan**—Allows you to return to the last-known good operating system, firmware, or driver version, in case of problems with newer versions
- **No single-point-of-failure design**—Reduces the risk of one device or connection leading to system failure
- **Scheduled downtime with customer**—Allows you to perform system maintenance on a regular schedule, without unexpected disruptions
- **Regular updates and alerts for status and usage**—Provides insight into the health of the customer's system, and allows you to see signs of performance degradation in advance
- **Network usage and server usage baseline**—Establishes metrics by which future performance can be measured, either in terms of system degradation or as a baseline for future upgrades
- **Regular updates for firmware and drivers**—Keeps the customer's system up to date
- **Prefailure warranty**—Ensures that when customers receive notification from system management that a critical server component might fail, the component is replaced free of charge under the warranty

You should also keep in regular contact with your customers, and review their usage and system performance measurements to help them determine the point at which they might benefit by upgrading or replacing their equipment with newer models that offer better price points or a wider set of features.

## Best practices for resiliency



Having a reliable infrastructure and a preventive security strategy is key to the survival of a business. Businesses that remain reactionary, as opposed to proactively developing a security strategy, are more susceptible to the effects of a security breach.

There are three specific areas you can fortify to safeguard a business as you deploy productivity-enhancing applications:

- Increase the resiliency of the IT infrastructure using modern products designed with security in mind.
- Protect the existing network with robust security-enhancing products.
- Plan for potential data loss by deploying backup and recovery solutions.

## Learning check

1. List three proactive maintenance strategies you can use to minimize the potential for downtime or lost data.

---

---

---

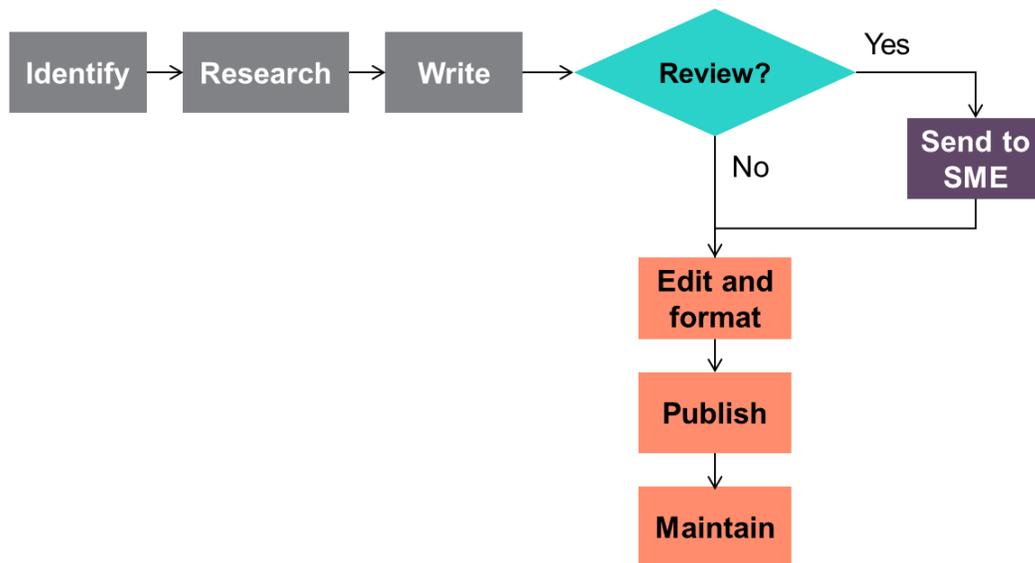
2. List three best practices for resiliency.

---

---

---

## Validating the design and documenting the solution



The process of drafting a solution proposal

After gathering the business and technical information from the customer, you can begin work on the solution proposal. The creation of timely, accurate, high-quality, proposal-ready content requires a rigorous development and maintenance process. At this point, you should focus on various aspects of the solution, including addressing architectural and transitional issues such as functional and technical design, organizational design, technology governance, and change management.

### Writing a scope of work

A *scope of work* is a project overview you prepare for the proposal. It captures the plan, time frame, required resources, and completion milestones of a project. It is a crucial document used to ensure a mutual understanding with the customer. Executive support is essential. Without leadership support, it can be difficult to implement an IT project.

---

**NOTE:** A scope of work should not be confused with a statement of work, which is a final project overview prepared for billing.

---

The scope of work should provide a summary of the plan you create for the solution, including:

- Overall time frame
- Completion milestones for each aspect of the solution
- Resources required
  - Channel partners, HPE sales representatives, and any other involved parties
  - Executive support for the project (name, position, and so on)

## Developing the proposal

When designing a solution, it is not enough to present the technical aspects. Remember to present the value of the solution in business terms that matter to the customer. Frame the technology in terms of how it addresses business drivers and initiatives, how it overcomes obstacles, and how it meets the customer's goals. Architect the solution to meet business, technical, and financial needs by:

- Describing the business value for the customer
- Developing a logical architecture that will host the solution, including:
  - Network layout
  - Server requirements
  - Application services
  - Storage capacity requirements
- Incorporating licensing options based on current HPE QuickSpecs
- Outlining how to integrate your solution into the customer's IT infrastructure

From a content perspective, proposal-ready documents typically include:

- Key benefits and differentiators
- Latest HPE marketing messages
- Customer and analysis quotes
- High-level technical information
- Proof points

---

**NOTE:** Remember to use the worksheet provided in Appendix B of this learner guide to gather information about IDJT HealthCare as you move through this course. This worksheet gives you the beginnings of an actual proposal that you could use when working with a customer.

---

The order in which you present information might be dictated by points of focus in the request for information (RFI) or request for proposal (RFP). Depending on the solution you are proposing and the resources required, you might also include support information from channel partners, program managers, special interest groups, and others. Most proposal-ready content must be reviewed before publication to ensure accuracy. Reviewers can include:

- Product managers
- Program managers
- Marketing specialists
- Special interest groups
  - HPE Solution Architect (SA) community
  - HPE Ambassador program

---

**NOTE:** Proposal-ready content that is based on nontechnical, external sources such as the HPE website does not require review by a subject matter expert (SME).

---

After review, changes must be incorporated from the reviewers to ensure that the content passes a final editing process.

## Learning check

1. Which technology enables you to aggregate up to five switches within a single interface, forming a single logical switch using direct links that run to and from each switch?
  - a. Deduplication
  - b. Meshed stacking
  - c. Intelligent Resilient Framework
  - d. Single-root I/O virtualization
2. Intelligent Resilient Framework (IRF) implemented with FlexFabric 59xx switches forms a single IRF fabric consisting of at least \_\_\_\_\_ and up to \_\_\_\_\_ physical switches in one IRF domain.
3. Proposal-ready content that is based on nontechnical, external sources such as the HPE website does not require review by a subject matter expert.  
 True  
 False

## Summary

- Whether you are designing a server room or a large data center, the architecture must enable the infrastructure to grow with the business. A solution-level approach designed for the company's workloads incorporates computing, storage, Ethernet networking, and security in a way that simplifies implementation.
- Many HPE server solutions include security, flexibility, compute performance, and other options that are ideal for SMB workloads. HPE server solutions best suited for the SMB customer include the ProLiant MicroServer Gen10, ProLiant DL360 Gen10 server, and ProLiant BL460c Gen10 server blade.
- Creating a storage configuration that meets customers' computing needs depends on the application workload and the data center infrastructure. A typical storage solution is configured in one of three ways: DAS, NAS, or SAN. Optionally, a storage configuration can be created from disk or flash and boot from SAN. Important storage considerations also include backup architectures and deduplication.
- HPE FlexFabric and Aruba network solutions for SMB customers allow for secure growth, reliability, and network interoperability. Aruba network solutions are ideal for mobile and campus networks; FlexFabric switches are a good fit in traditional SMB and larger data center environments. HPE end-to-end server networking options include HPE 25Gb Ethernet adapters for ProLiant Gen10 rack servers, server network adapters, CNAs, and LOM technology. It is important to consider various types of switches, stacking technologies, types of VLANs, hybrid switch ports, and link aggregation options for each customer's configuration.
- There are several steps that a solution architect must take to validate and document the solution design for the customer. First, ensure that the solution addresses architectural and transitional issues such as functional and technical design, organizational design, technology governance, and change management. Then write the scope of work, which is an overview of the project plan. Lastly, when developing the proposal for the customer, cite the values of the solution that go beyond the technical aspects. Frame the solution in terms that address the customer's overall business goals.



# **Installing, Configuring, and Upgrading HPE SMB Solutions and Related Components**

## **Module 3**

### **Learning objectives**

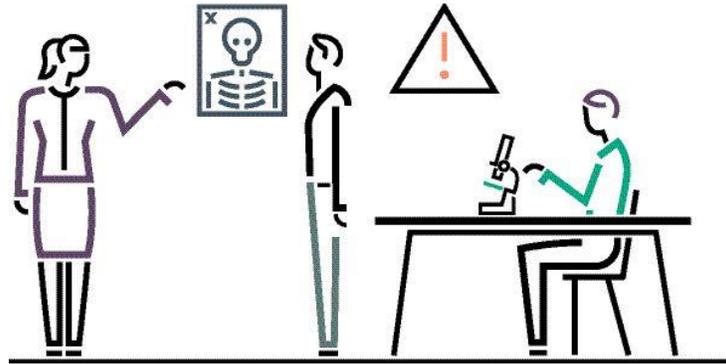
After you have completed this module, you should be able to:

- Install and configure Hewlett Packard Enterprise (HPE) data center solutions and subsystems for small to medium-sized businesses (SMBs)
  - Servers
  - Storage
  - Networking
  - Infrastructure
- Validate that an installation and configuration are successful
- Upgrade an HPE IT solution and its subsystems

## Customer scenario

Consider the following fictitious customer scenario as you work through this module. Practice planning and designing a solution to meet the customer's business and technical objectives. Refer back to this customer as you engage with activities and labs throughout this course.

### IDJT HealthCare



IDJT HealthCare serves physicians and caregivers across medical focus areas including cancer, cardiology, sports medicine, pain management, and sleep disorders

As a result of an increased number of customers coming to their sports medicine clinic, IDJT HealthCare plans to buy a new x-ray machine for their facility. They will need to transfer the captured images quickly from the machine to the servers in a way that does not create a bottleneck on the network. The IT manager of IDJT HealthCare, Eric Yurekha, is concerned that adding a new machine will cause the existing network bandwidth to be consumed quickly.

As Yurekha wrote to you in an email last week, “We are grateful for the growth of our sports medicine clinic and are ready to accept delivery of this new equipment. However, we want to be sure that our investment in this new x-ray machine does not slow us down. We understand that the cost of downtime can include lost revenue, damaged reputation, lost customer confidence and loyalty, damaged competitive position, and regulatory compliance exposure.”

The customer wants to install the x-ray machine in their existing network infrastructure so that it sends data to their Microsoft SQL Server database. You recommend upgrading a portion of the network from 1 Gb to 10 Gb for the servers that will store the Digital Imaging and Communications in Medicine (DICOM) image files. You will need to configure the network environment to support 10 Gb on a separate VLAN.

## Gauge your knowledge

Before proceeding with this training, consider the following questions to assess your existing knowledge. Pay particular attention to the answers you are less certain about. Consider your responses as you progress through the topics covered in this module.

1. What is the difference between BIOS and Unified Extensible Firmware Interface (UEFI)?

---

---

2. Which network fabric topologies are typically used in data centers?

---

---

---

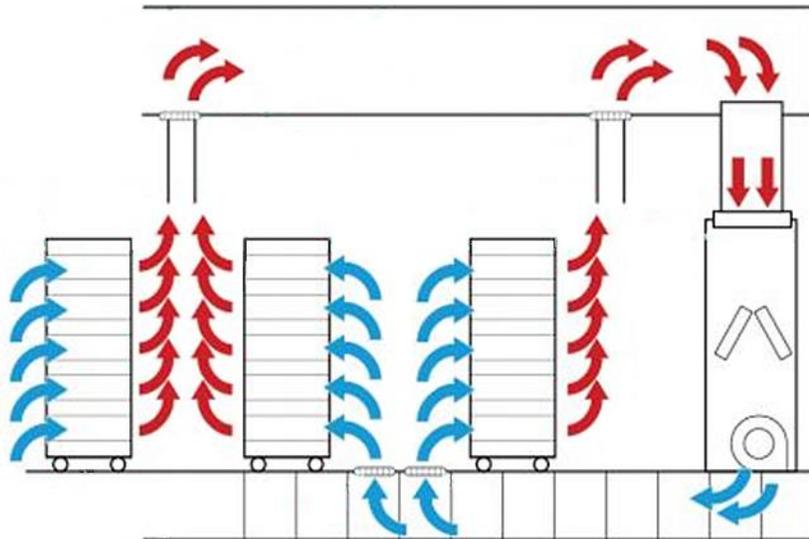
3. What experience do you have updating IT components?

---

---

Make a note of areas in which you need more information and concentrate on the material covered in those areas. Feel free to skim any material you might already know.

## Installing solutions and subsystems



Rack placement for proper airflow in a data center with a suspended ceiling and raised floor and down-flow air conditioner perpendicular to equipment rows

Before installing equipment in a data center, you should be aware of precautions and other requirements to ensure your safety and proper handling of the equipment. 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
- At least a 10 cm (4 in) clearance space at the front and back of the component for proper ventilation
- 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)

---

**NOTE:** An HPE High Performance Fan Kit is required when installing an HPE ProLiant DL360 or DL380 Gen9 or Gen10 server in an environment with an ambient temperature of 45°C. This fan kit ships standard in select ProLiant DL380 Gen9/Gen10 servers. For information about ASHRAE guidelines, refer to the HPE ProLiant Gen9 Server Extended Ambient Temperature Guidelines document by scanning this QR code or right-clicking it to open the hyperlink.



---

<http://h20566.www2.hpe.com/hpsc/doc/public/display?docId=c04513664>

---

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 of 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 or ground (PE) 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 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.

---

## Installing data center components

To view and navigate the initial server setup screens, you must first connect I/O devices to the server. To do so, connect a keyboard and mouse to the USB ports and then connect a monitor to the VGA or display port. Next, connect the network cable by connecting one end of the network cable to the network interface card (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 you 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.

## 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 installation guide for that option. If you are 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 hard drives or all solid state drives.
  - 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.

---

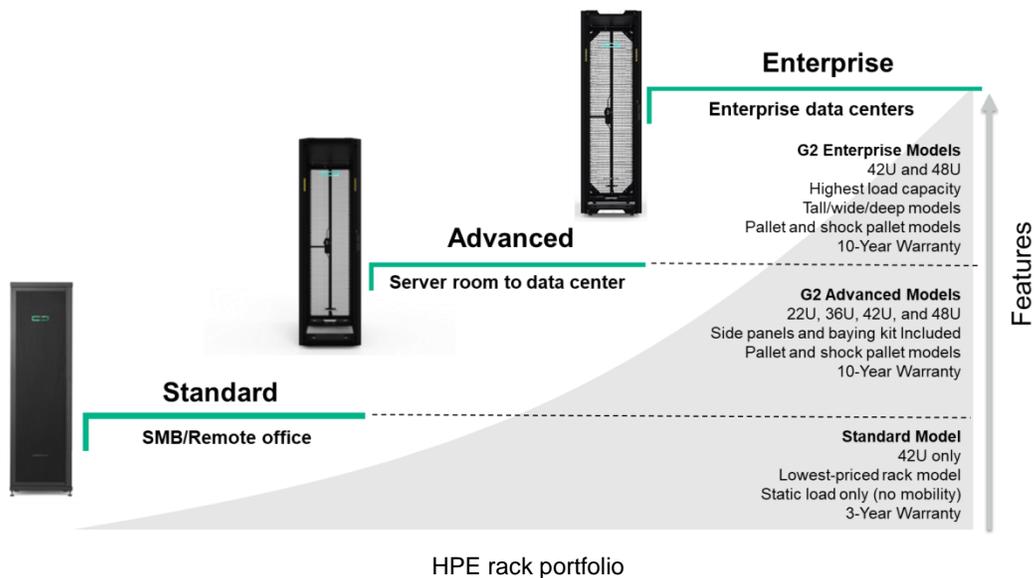
**NOTE:** For complete information on servicing ProLiant servers, visit the Hewlett Packard Enterprise Support Center and search for the server you are interested in. On the server page, click **Manuals** → **Service and maintenance information**. Then click the link to open the maintenance and service guide. You can access the Support Center by scanning this QR code or right-clicking it to open the hyperlink.



<http://h20565.www2.hp.com/portal/site/hpsc>

---

## Installing the rack infrastructure



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, power distribution units, and uninterruptible power supplies.

---

- 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 power distribution units (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 G2 Advanced Series Racks are configured in a range of sizes and with flexible features to fit right into modern data centers. They include the latest innovations to streamline installation and maintenance with a focus on simplicity. These racks have been designed and tested to work optimally with HPE servers, storage, and networking products. HPE racks carry a standard limited ten-year warranty.

HPE offers a variety of options for G2 Advanced Series 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 you to customize the rack with rack location information, branding, or other description information.

---

**NOTE:** For more information about HPE rack offerings, read the “HPE Racks for Server, Storage, and Networking” family data sheet available from:

<https://www.hpe.com/h20195/v2/getpdf.aspx/4AA6-2115ENW.pdf>

To use the HPE Rack and Power Infrastructure Architect tool, visit: <https://iallb.itcs.hpe.com/>

---

## Power requirements



Modular 0U PDUs maximize rack space

When installing more than one server, you might have to use additional PDUs 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 RETMA rails. Extension bars mount directly to the frame so you 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 ProLiant and BladeSystem servers and options around the world.

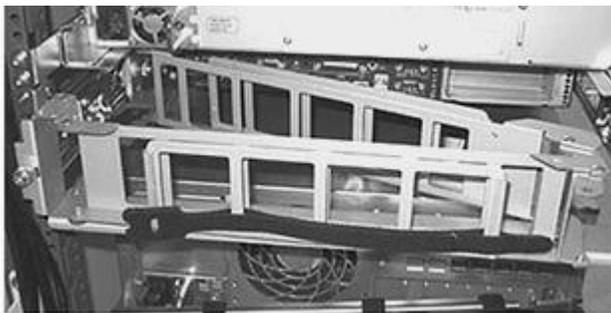
Most ProLiant servers feature auto-sensing power supplies with a power input range of 100 – 240 VAC, so they can be used globally. The appropriate power cables ship standard with all AC-powered products. HPE power cords meet the standards for the country for which they are intended. Rack-mountable HPE 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



HPE Cable Management Arm

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 host bus adapter (HBA). You should always verify that a cable/transceiver is supported with a connecting device such as a switch or NIC/iSCSI HBA.

---

**NOTE:** HPE recommends using standard cables when connecting to BladeSystem switches. High-density mini-SAS cables are not used for connecting to a disk enclosure. Direct attached copper cables are supported in 10GbE controller systems.

---

Fanout cables, sometimes known as *breakout cables*, are optical fiber cables that hold several optical fibers packaged together inside a jacket that can be terminated directly and easily. HPE recommends using fanout cables when connecting to SAS HBAs for maximum future flexibility without downtime. Note that fanout cables are not supported on all storage systems.

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 you 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 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 you to extend devices from the rack without powering them down or disconnecting cables.

## Activity—Selecting a rack and power infrastructure

In this activity, you will use HPE tools to select a rack and power infrastructure for IDJT HealthCare.

### Customer scenario

IDJT HealthCare has purchased a managed switch and an HPE ProLiant DL360 Gen10 rack-mounted server for their remote location. They also have repurposed their HPE StoreOnce VSA for the remote location. They asked you to help them choose the rack and power infrastructure that is appropriate for their remote data center.

### Using the HPE Rack and Power Infrastructure Architect tool

1. To begin this activity, navigate to: <https://iallb.itcs.hpe.com>
2. At the HPE Rack and Power Infrastructure Architect screen, click **Recommend the right infrastructure equipment**. Be sure to check the box to agree to the HPE terms and conditions.

#### What is your HPE Compute story?

Recommend the right infrastructure equipment

I know this stuff — give me the technical questions needed.

I agree to the HPE Terms & Conditions.

3. At the next screen, select **I'm new** and click **Next**.

#### Are you new to our tool or have you been here before?

Please select one of the following:

I'm new

I've been here before

**NEXT**

- At the next screen, select **IT Admin** and click **Next**.

## Tell us what your job is.

Please select one of the following:

End User

IT Admin

Purchasing

HPE Reseller

Other

- IDJT HealthCare currently has a total of three servers. Select **1 – 10** and click **Next**.

## How many servers do you have today?

Please select one of the following:

1 - 10

10 - 100

100 - 1,000

1,000+

6. IDJT HealthCare is located in the United States. From the drop-down menu, select **United States** and click **Next**.

## Where will you be installing your new IT equipment?

Please select from the dropdown menu:

**PREVIOUS** **NEXT**

7. Because IDJT HealthCare currently has fewer than five servers, select **14U** and click **Next**. If time permits, you can return to this screen and select other options.

## First, we need to determine the size (height) of the rack you want.

Please select one of the following:

- 14U** has space for 1-5 servers and additional room for other devices
- 22U** has space for up to 10 servers and additional room for other devices
- 36U** has space for up to 15 servers and additional room for other devices; also used to support 1-2 larger server systems such as HPE Apollo, BladeSystem, Moonshot, and Synergy
- 42U** has space for up to 20 servers and additional room for other devices; also used to support 1-2 larger server systems
- 47+U** has space for more than 20 servers and additional room for other devices; also used to support 3 or more larger server systems
- I'M NOT SURE** No problem, we'll show you all the options

**PREVIOUS** **NEXT**



**TIPS WORTH THINKING ABOUT**

Think about the equipment you are going to put in your rack today, plus what you'll want to add in the future.

U = standard unit of measurement for racks. 1U = 1.75 inches.

**WE RECOMMEND**

42U as this is the most popular size for a rack and offers the most options.

8. At the next screen, click **I'M NOT SURE** and then click **Next**.

## Next, let's determine what width rack you need.

Please select from one of the following:

- 600 mm** Most standard racks are 600mm wide
- 800 mm** Gives you more room on each side to manage airflow and cabling
- I'M NOT SURE**

**PREVIOUS** **NEXT**



**TIPS WORTH THINKING ABOUT**

The 800mm width is only available on 42U racks.

Most standard racks are 600mm

9. IDJT HealthCare has only ProLiant DL series servers, so click **Standard rack servers** and then click **Next**.

## Now, let's make sure your servers will fit. What type of servers will you be putting in the rack?

Please select one of the following:

- Standard rack servers** such as HPE ProLiant DL80/DL180/ DL380/DL580
- 1 or 2 large systems** such as HPE Apollo (2000/4000/6000), BladeSystem (c7000/c3000), Integrity, Moonshot, and/or Synergy
- 3 or more large systems** such as HPE Apollo (2000/4000/6000), BladeSystem (c7000/c3000), Integrity, Moonshot, and/or Synergy
- Mix of different servers**
- I'M NOT SURE** No problem, we'll show you all the options

PREVIOUS

NEXT



### TIPS WORTH THINKING ABOUT

The type of server(s) you put in the rack will determine the depth you need.

Most standard rack servers are 1075 mm deep. However, you may need additional room at the back of the rack for cabling, air flow management, or power management. If so, select "Mix of different servers" which will give you a 1200mm deep rack.

10. Because IDJT HealthCare has indicated they might move their equipment after they configure it, click **I'M NOT SURE** and then click **Next**.

## Last question on racks. Will you need to move your rack around?

Please select from one of the following:

- No – we will install and configure the rack without having to move it.**
- We may need to move the rack from room to room once we configure it.**
- Yes, we want our rack shipped to us with all the equipment pre-configured, so we'll need to move it into position.**
- Yes, we need a movable rack that we can ship to another location once we configure the equipment.**
- I'M NOT SURE** No problem, we'll show you all the options

PREVIOUS

NEXT



### TIPS WORTH THINKING ABOUT

Racks come with or without wheels and can be shipped empty or pre-configured with your equipment.

HPE offers rack configuration services through HPE Factory Express.

11. The tool displays the available racks that meet the criteria you input. In this case, there is only one rack that meets your description, as shown in the following screenshot. Note that it ships as a shock rack, which means that it can be moved after they configure it. Select the radio button for this product and then click **Next**.

## Phew — that was a lot of questions, but here are the racks that will meet your needs. Make your choice:

Please select from one of the following:



### HPE 14U 600mm x 1075mm Advanced Shock Rack

<b>Part #:</b> H6J82A	<b>Height:</b> 14U	<b>Width:</b> 600mm	<b>Depth:</b> 1075
<b>Static Load:</b> 3000 lbs 1364 kg	<b>Dynamic Load:</b> 2500 lbs 1136 kg	<b>Series:</b> Advanced	
<b>Shipped as:</b> Shock	<b>Shipped including:</b> Perforated front and rear door		

PREVIOUS

NEXT

12. The next screen displays the airflow, cable management, and stabilization options available for this rack. Make your selections based on what you have learned thus far in this course, and based on what you know about IDJT HealthCare. Click **Next**.

**Every rack likes its accessories. Our accessories will help with air flow, cable management, and stabilization. Here are the ones our customers like the most:**

Please select all the accessories you are interested in:

For Air Flow

- |                          |   |  |
|--------------------------|---|--|
| <input type="checkbox"/> | <b>Air Flow Kit</b><br>Part #: SW930A               | Add if you need to seal air gaps inside the rack, in between two bayed racks and the clearance from the floor to the rack.   |
| <input type="checkbox"/> | <b>Side Panels</b><br>Part #: Included              | Add to improve cooling efficiency and prevent hot air from recirculating through the sides of the rack.  |
| <input type="checkbox"/> | <b>1U Filler Panels (10 Pack)</b><br>Part #: SW925A | Add if you need to fill or close unused opening in your rack to prevent exhausting mixing with the cooler intake air for more efficient cooling of your equipment. |

For Cable Management

- |                          |  |  |
|--------------------------|--|--|
| <input type="checkbox"/> | <b>Cable Management Kit</b><br>Part #: AFD99A                    | Add if you need an easy way to organize and route cables within your rack. |
| <input type="checkbox"/> | <b>Quarter Turn Hook &amp; Loop (100 pack)</b><br>Part #: SW964A | Add if you need cable ties to organize cables within your rack.            |
| <input type="checkbox"/> | <b>D Rings</b><br>Part #: 168233-B21                             | Add if you need mounted rings to help stream your cables into line.        |

For Stabilization

- |                                     |   |   |
|-------------------------------------|---|---|
| <input type="checkbox"/>            | <b>Stabilizer Kit</b><br>Part #: Included | Add to prevent possible tipping when equipment is installed, removed, or accessed within the rack.    |
| <input checked="" type="checkbox"/> | <b>Baying Kit</b><br>Part #: Unavailable  | Add if you need to connect multiple racks of the same series, height, and depth to bay them together. |

**PREVIOUS** **NEXT**



**TIPS WORTH THINKING ABOUT**

These are the most popular accessories that our customers choose. We have accessories for air flow management, cable management, stabilization, rack mounting hardware, and location discovery services. Talk to your HPE partner to discuss other options.

13. IDJT HealthCare has told you that they have a limited budget for their remote location. From the PDU screen, select the **Basic** option and then click **Next**.

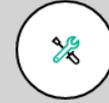
## Every rack needs a power strip. This is commonly known as a PDU (Power Distribution Unit).

Please select one of the following:

- Basic** Provides basic power distribution. Simple and reliable.
- Metered** Basic PDU PLUS the ability to remotely meter power consumption to see how much is being used.
- Metered & Switched** Basic power distribution, remote metering, AND the ability to remotely switch outlets on and off.
- Intelligent** Basic power distribution, remote metering, PLUS the ability to remotely switch individual outlets on and off using HPE's Power Discovery Solution and OneView software
- I'M NOT SURE** Show me all the options.

PREVIOUS

NEXT



### TIPS WORTH THINKING ABOUT

If you are already monitoring power centrally and simply need power outlets or are price sensitive, select a Basic PDU.

### WE RECOMMEND

a Metered PDU so you know how much power you are using and can plan for future needs while ensuring you keep within

14. You have not gotten information about the power outlet in the IDJT HealthCare remote location, so click **I'M NOT SURE** and then click **Next**.

## Okay, now we need to know what type of power you have where the rack and PDU will be located.

Please select one of the following:

<input type="radio"/> CS8265C	Reference Look 
<input type="radio"/> CS8365C	Reference Look 
<input type="radio"/> IEC 60309 30A 5-wire	Reference Look 
<input type="radio"/> IEC 60309 32A/30A 5-wire	Reference Look 
<input type="radio"/> IEC 60309 60A 4-wire	Reference Look 
<input type="radio"/> IEC 60309 60A 4-wire Watertight	
<input type="radio"/> NEMA L15-30P	Reference Look 
<input type="radio"/> NEMA L21-20P	Reference Look 
<input type="radio"/> NEMA L21-30P	Reference Look 
<input type="radio"/> NEMA L22-20P	Reference Look 
<input type="radio"/> NEMA L22-30P	Reference Look 
<input type="radio"/> NEMA L5-20P	Reference Look 
<input type="radio"/> NEMA L5-30P	Reference Look 
<input type="radio"/> NEMA L6-30P	Reference Look 
<input checked="" type="radio"/> I'M NOT SURE	



### TIPS WORTH THINKING ABOUT

You can determine the type of power by what the outlet looks like. Check the outlet where the rack will be located.

PREVIOUS

NEXT

15. At the next screen, make your best guess about the number of devices that will be plugged into the PDU at the IDJT HealthCare remote data center. You know of at least seven devices, so click **7 – 12** → **Next**. If time permits, you can return to this screen and explore other options.

## Now let's figure out how many outlets your PDU needs. How many devices (servers, storage, networking, consoles, etc.) do you plan to plug into your PDU?

Please select one of the following:

- 1 - 6
- 7 - 12
- 13 - 19
- 20+
- I'M NOT SURE

**PREVIOUS** **NEXT**



**TIPS WORTH THINKING ABOUT**

A 42U rack could potentially have 42 devices (if each device is 1U). You probably have an idea of how many servers you will have. Take that number and consider adding outlets for other devices including 1 UPS and 1 KVM console, network switch for the total number

16. IDJT HealthCare has not specified whether they want a horizontal or a vertical PDU, but your experience tells you they would be more comfortable with horizontal racks. Click **Horizontal** → **Next**.

## Now, back to school. Do you need a horizontal or vertical PDU?

Please select one of the following:

- Horizontal
- Vertical
- I'M NOT SURE

**PREVIOUS** **NEXT**



**TIPS WORTH THINKING ABOUT**

Horizontal is commonly used for smaller racks (under 36U), and some models support adding extension bars to give you more

17. A list of PDUs is displayed. A quick phone call to Eric Yurekha, the IT manager of IDJT HealthCare, confirms that this rack is compatible with the power the company has coming in to the facility. Select the **HPE G2 Basic 3.6kVA/IEC C20 Detachable 16A/100-240V Outlets (12) C13/1U Horizontal WW PDU** from the list. Make a note of the part number. Then scroll down and click **Next**.

## Way to power through it. Select from the PDUs below that fit your needs:

Please select one of the following:




HPE G2 Basic 1.9kVA/IEC C20 Detachable 16A/120V Outlets (12) 5-20R/1U Horizontal NA/JP PDU

Part #:	Category:	Type:	Phase:
P9Q31A	Basic	Horizontal	Single
<b>Input Voltage:</b>	<b>Input Connector:</b>		
100-120	IEC C20		
<b>Uheight:</b>	<b>C13 Outlets:</b>	<b>C19 Outlets:</b>	
1	0	0	




HPE G2 Basic 3.6kVA/IEC C20 Detachable 16A/100-240V Outlets (12) C13/1U Horizontal WW PDU (compatible with the optional HPE IEC Locking Power Cords for additional security)

Part #:	Category:	Type:	Phase:
P9Q37A	Basic	Horizontal	Single
<b>Input Voltage:</b>	<b>Input Connector:</b>		
100-240	IEC C20		
<b>Uheight:</b>	<b>C13 Outlets:</b>	<b>C19 Outlets:</b>	
1	12	0	



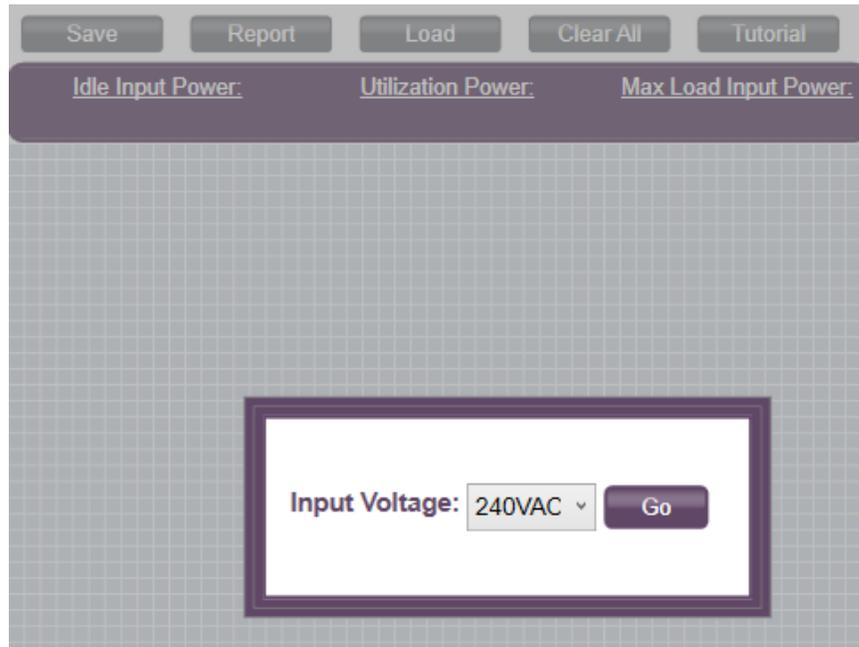
### TIPS WORTH THINKING ABOUT

One difference we didn't discuss is the amount of power your PDU will support. That's the power rating, listed as kVA — the bigger the number, the more power the PDU can supply. The amount of power a device needs can vary. This is a great conversation to have with your sales rep.

You can calculate the power you need by checking the power supply, device label, or quick spec for each device. Look for the maximum power (watts) the device can use. Get that number for each device, add them together, and that will give you the total amount of power you need for your PDU.

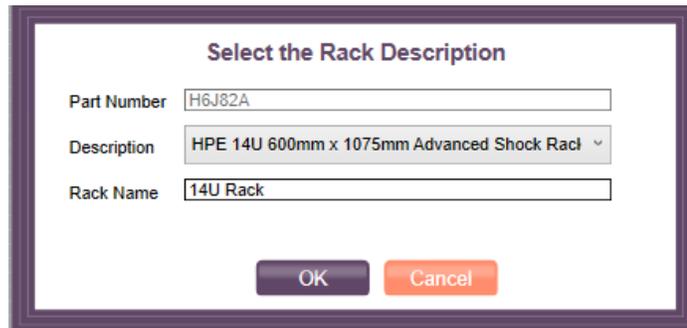
## Using the HPE Power Advisor

- Next, you need to calculate the amount of power needed in the rack you are considering for IDJT HealthCare. To do so, download the HPE Power Advisor from:  
<https://sizersllb.itcs.hpe.com/sb/installs/HPEPowerAdvisor.zip>
- After downloading and installing the tool, the first step in using HPE Power Advisor is to select an input voltage. The IDJT HealthCare data center is in the United States and runs on 240VAC. From the drop-down menu, select **240VAC**.



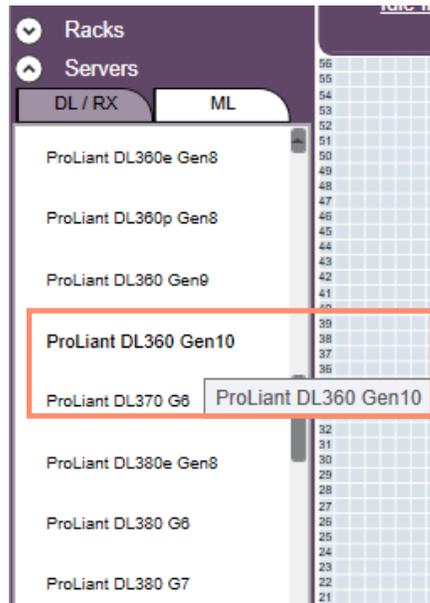
The screenshot shows the HPE Power Advisor web interface. At the top, there are five buttons: Save, Report, Load, Clear All, and Tutorial. Below these buttons are three links: [Idle Input Power:](#), [Utilization Power:](#), and [Max Load Input Power:](#). The main area of the interface is a light gray grid. In the center, there is a white box with a dark purple border containing the text "Input Voltage:" followed by a dropdown menu showing "240VAC" and a "Go" button.

- From the left-hand pane, select **14 U Rack**. At the dialog box that appears, click **OK**.

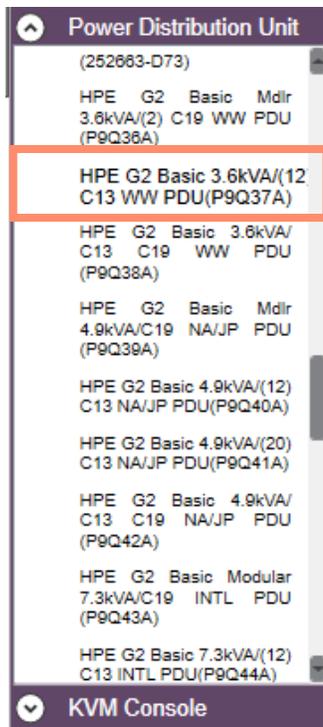


The screenshot shows a dialog box titled "Select the Rack Description". It contains three input fields: "Part Number" with the value "H6J82A", "Description" with a dropdown menu showing "HPE 14U 600mm x 1075mm Advanced Shock Racl", and "Rack Name" with the value "14U Rack". At the bottom of the dialog box are two buttons: "OK" and "Cancel".

- From the left-hand pane, expand **Servers**. From the drop-down list, select **ProLiant DL360 Gen10**, which is the new server that IDJT HealthCare will be installing in the new rack.



- From the popup window, accept the defaults and click **OK**.
- From the left-hand pane, expand **Power Distribution Unit** and then expand **Basic**. From the drop-down list, select the PDU you chose previously, which is the **HPE G2 Basic 3.6kVA/IEC C20 Detachable 16A/100-240V Outlets (12) C13/1U Horizontal WW PDU**. Double-check the part number (P9Q37A) to ensure that you are selecting the correct product.



24. From the popup window, accept the defaults and click **OK**.

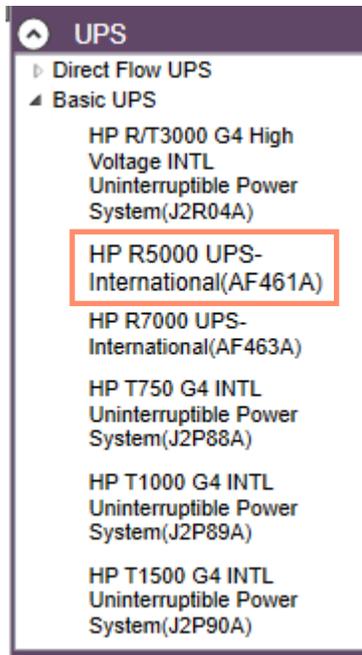


25. Next, select the KVM console. From the left-hand pane, expand **KVM Console**. From the drop-down list, select a kit that is designed to work in the US.

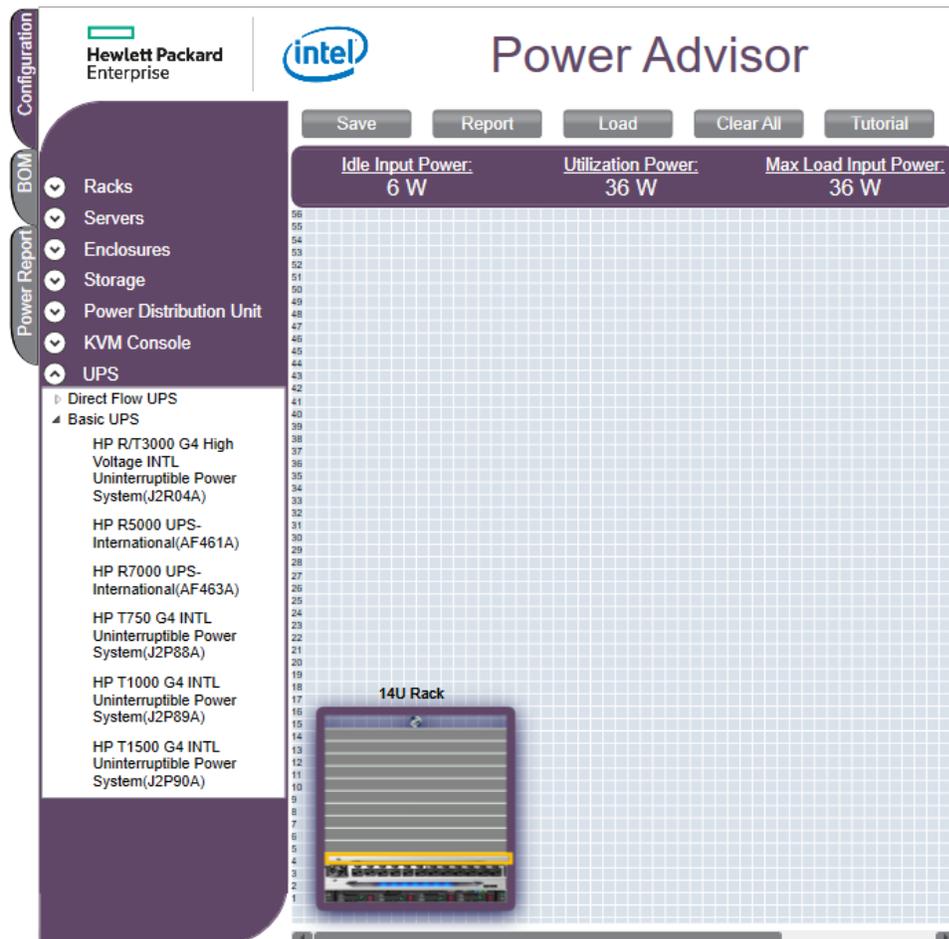
26. From the popup window, accept the defaults and click **OK**.



27. From the left-hand pane, expand **UPS** and expand **Basic UPS**. Select a basic international UPS.



28. The Power Advisor displays an illustration of the components as they would be installed in the rack, given the customer requirements and your choices. You can see that there is room for the company to grow in this rack. The required power in Watts is displayed across the top of the design field. You can choose to save the configuration or generate a power report that calculates the current data center costs or the total cost of ownership based on the components you selected. You can also generate a bill of materials, which you can access from the tabs on the left edge of the screen. To return to the configuration screen from the Power Report screen, click the **Configuration** tab.
29. Remember to add the components you selected for IDJT HealthCare to the worksheet you are using for this course. You can find the worksheet in Appendix 2 of this learner guide.
30. If time permits, return to the beginning of this activity and go through it again, this time selecting different options. Compare your results and make note of where the changes occurred based on your choices.



## Learning check

1. When you are installing equipment in a data center, what can result from a discharge of static electricity from a finger or other conductor?

---

---

---

---

2. Fill in the blanks:

Before installing and configuring IT solution components, make sure the customer has a good understanding of racking, \_\_\_\_\_, and host port \_\_\_\_\_ requirements.

3. Name up to five options available with HPE Advanced Racks.

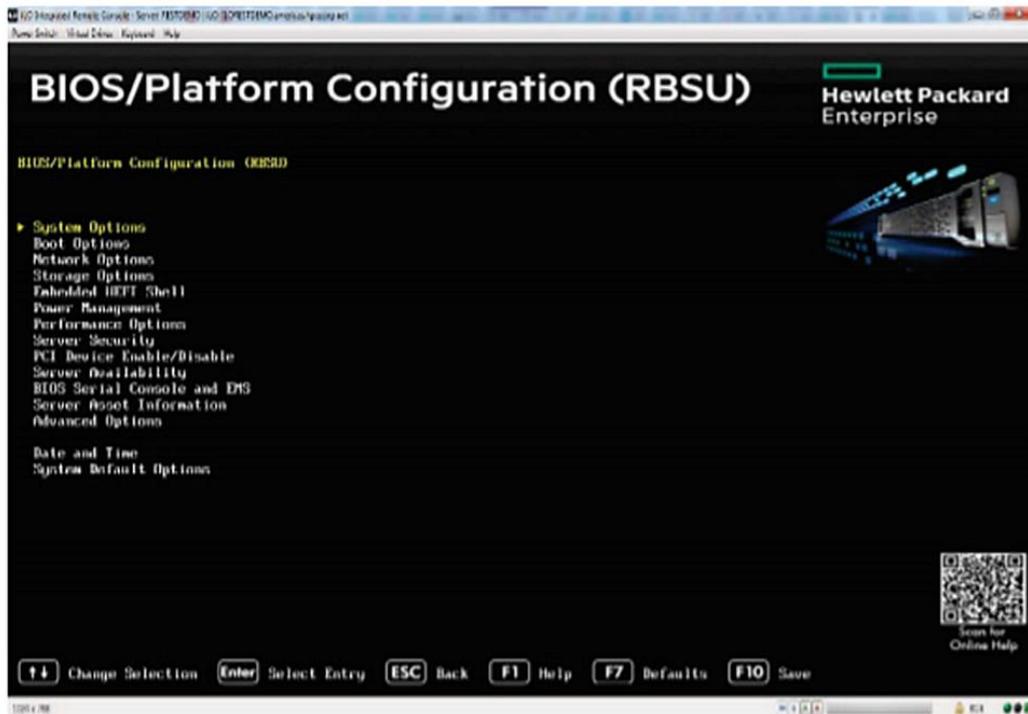
---

---

---

---

## Configuring ProLiant servers



UEFI/BIOS/Platform Configuration (RBSU) menu

There are three ways to configure a ProLiant server:

- UEFI/BIOS/Platform Configuration (RBSU)

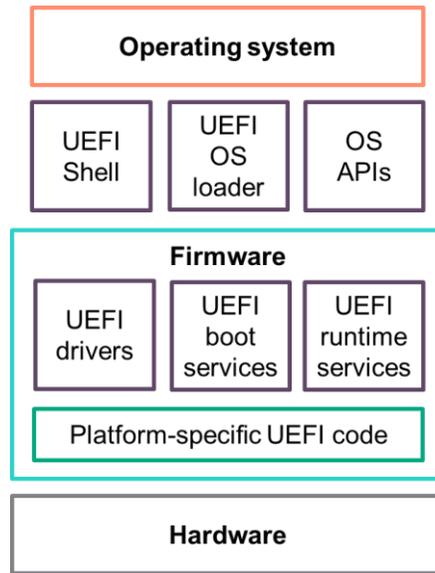
---

**NOTE:** The BIOS/Platform Configuration menu replaced the ROM-Based Setup Utility. Use this menu to access and use UEFI options.

---

- HPE iLO 5 Configuration Utility
- HPE Intelligent Provisioning

## BIOS and UEFI



UEFI modular structure of drivers and files

UEFI defines the interface between the operating system and platform firmware during the boot, or start-up process. It is embedded in the system ROM of all ProLiant Gen10 servers. UEFI employs a modular structure of drivers and files that are grouped into services that adapt well to complex and varied architectures. The UEFI framework of coding scales easily with server design trends. Agnostic to the type of processor used in a platform, UEFI enables a choice of CPUs within a server family.

Compared to BIOS, UEFI supports advanced preboot user interfaces. Its features enable you to perform a wide range of configuration activities, including:

- Configuring system devices and installed options
- Enabling and disabling system features
- Displaying system information
- Selecting the primary boot controller or partition
- Configuring memory options
- Launching other preboot environments

HPE servers with UEFI can provide:

- Support for boot partitions larger than 2.2 TB. Such configurations could previously only be used for boot drives when using RAID solutions.
- Secure Boot, which enables the system firmware, option card firmware, operating systems, and software to collaborate to enhance platform security. Secure Boot enables OEMs to implement an operating system-agnostic approach to securing systems in the preboot environment.
- An Embedded UEFI Shell that provides a preboot environment for running scripts and tools.
- Boot support for option cards that only support a UEFI option ROM.

The UEFI network stack enables implementation in a rich network-based operating system deployment environment. ProLiant Gen10 servers support a GUI-based UEFI System Utilities, which supports mouse and keyboard devices. UEFI also supports both IPv4 and IPv6 networks.

---

**NOTE:** For more information, review the *UEFI System Utilities User Guide for HPE ProLiant Gen10 Servers and HPE Synergy* by scanning this QR code or right-clicking it to open the hyperlink.



[http://h20566.www2.hpe.com/hpsc/doc/public/display?sp4ts.oid=6935826&docLocale=en\\_US&docId=emr\\_na-a00016407en\\_us](http://h20566.www2.hpe.com/hpsc/doc/public/display?sp4ts.oid=6935826&docLocale=en_US&docId=emr_na-a00016407en_us)

To access HPE UEFI System Utilities and Shell Command Mobile Help for HPE ProLiant Gen9 Servers, scan this QR code or right-click it to open the hyperlink.



[http://h17007.www1.hpe.com/docs/iss/proliant\\_uefi/UEFI\\_Gen9\\_060216/index.html](http://h17007.www1.hpe.com/docs/iss/proliant_uefi/UEFI_Gen9_060216/index.html)

---

## Launching UEFI System Utilities

1. Optional: If you access the server remotely, start an iLO remote console session.
  - a. Open a browser and enter **https://<iLO host name or IP address>** to log on to the iLO web interface.
  - b. On the login page, enter a directory or local user account name and password, and click **Log In**.
  - c. In the iLO navigation tree, select **Remote Console & Media**.
  - d. When the Launch tab is displayed, verify that your system meets the requirements for using the remote console application you want to use. Click the **Launch** button for your selected application.
 

You can also launch an iLO Remote Console session by selecting the **Integrated Remote Console** link on the **Information** → **iLO Overview** page. Click the **Console** thumbnail in the lower-left corner of the iLO web interface, and then choose the application type to launch.
2. Restart or power on the server. The server restarts and the POST screen is shown.
3. Press **F9**. The System Utilities main menu is displayed, which is the starting point for:
  - System Configuration
  - One-Time Boot Menu
  - Embedded Applications
  - System Information
  - System Health
  - Exit and resume system boot
  - Reboot the System
  - Select Language

## Using UEFI System Utilities

To use UEFI System Utilities, use the keys presented in the following table.

Action	Key
Access System Utilities	F9 during server POST
Navigate through the menus	Up and down arrows
Select items	Enter
Save selections	F10
Access Help for a highlighted configuration option (Scan the QR code on the screen to access online help for the UEFI System Utilities and UEFI Shell)	F1

Default configuration settings are applied to the server upon the first server power-up and after defaults have been restored. Default configuration settings are sufficient for typical server operations; however, you can modify configuration settings as needed. The system prompts you for access to the UEFI System Utilities each time the system is powered up.

## Selecting the boot mode

```

BIOS/Platform Configuration (RBSU)

Boot Options → UEFI Boot Order

Press the '+' key to move an entry higher in the boot list and the '-' key to move an entry lower
in the boot list. Use the arrow keys to navigate through the Boot Order list.

Generic USB Boot
Embedded Storage : Dynamic Smart Array B140i - LogVol (Lun:0 VolId:0) 232.05 GiB, RAID 0
Embedded LOM 1 Port 1 : HP Ethernet 1Gb 2-port 361i Adapter - NIC (PXE IPv4)
Embedded FlexibleLOM 1 Port 1 : HP Ethernet 1Gb 4-port 331FLR Adapter - NIC (PXE IPv4)
Embedded LOM 1 Port 1 : HP Ethernet 1Gb 2-port 361i Adapter - NIC (PXE IPv6)
Embedded FlexibleLOM 1 Port 1 : HP Ethernet 1Gb 4-port 331FLR Adapter - NIC (PXE IPv6)
Slot 2 Port 1 : Intel(R) B2574L Gigabit Network Connection (PXE IPv4)
Slot 2 Port 1 : Intel(R) B2574L Gigabit Network Connection (PXE IPv6)
▶ Embedded UEFI Shell

```

UEFI Boot Order list example

Certain boot options require that you select a specific boot mode. By default, the boot mode is set to UEFI Mode. The system must boot in UEFI Mode to use certain options, including:

- Secure Boot
- UEFI Optimized Boot
- Generic USB Boot
- IPv6 PXE Boot
- iSCSI Boot
- Boot from URL
- Fibre Channel/Fibre Channel over Ethernet (FCoE) Scan Policy

---

**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.

---

To configure the boot mode on a ProLiant server, follow these steps:

1. From the System Utilities screen, select **System Configuration** → **BIOS/Platform Configuration (RBSU)** → **Boot Options** → **UEFI Boot Order** and press **Enter**.
2. Select a setting.
  - **UEFI Mode (default)** configures the system to boot to a UEFI-compatible operating system. When booting to UEFI Mode, leave UEFI Optimized Boot enabled.
  - **Legacy BIOS Mode** configures the system to boot to a traditional operating system in Legacy BIOS compatibility mode.
3. Press F10 to save your setting.
4. Reboot the server.

## Workload profiles and performance options

Workload profiles are a configuration option to deploy BIOS settings based on which workload the customer intends to run on the server. ProLiant systems provide these workload profiles:

- **General Power Efficient Compute**—This profile is the default profile for most ProLiant servers and HPE Synergy compute modules. This profile applies the most common performance settings that benefit most application workloads and enables power management settings that have minimal impact on overall performance. The settings that are applied heavily favor a balanced approach between general application performances and power efficiency. This profile is recommended for customers that do not typically tune their BIOS for their workload.
- **General Peak Frequency Compute**—This profile is intended for workloads that benefit from processors or memory that must achieve the maximum frequency possible, for any individual core, at any time. Power management settings are applied when they ensure that any component frequency upside can be readily achieved. Processing speed is favored over any latencies that might occur. This profile is a general-purpose profile, so optimizations are done generically to increase processor core and memory speed. This profile supports workloads that benefit from faster compute time.
- **General Throughput Compute**—This profile is intended to be used for workloads where the total maximum sustained workload throughput is needed. Increased throughput does not always occur when the processor runs at the highest individual core speed. Increased throughput can occur when the processor is able to perform sustained work across all available cores during maximum utilization. Power management settings are disabled when they are known to impact the maximum achievable bandwidth. The best throughput is achieved when the workload is also non-uniform memory access (NUMA)-aware and optimized, so settings that benefit NUMA awareness are applied.
- **Virtualization - Power Efficient**—This profile is intended to be used for virtualization environments. It ensures that all available virtualization options are enabled. Certain virtualization technologies can possibly impact performance on nonvirtualized environments and can be disabled in other profiles. Power management settings might impact performance when running virtualization operating systems. This profile applies power management settings that are virtualization friendly.
- **Virtualization - Max Performance**—This profile is intended to be used for virtualization environments. It ensures that all available virtualization options are enabled. Power management settings are disabled in favor of delivering maximum performance. This profile is intended to be used by customers who want the least amount of computational latency for their workloads. Maximum speed and throughput are often sacrificed to lower overall computational latency. Power management and other management features that might introduce computational latency are also disabled. It benefits customers running real-time operating systems or transactional latency-sensitive workloads.
- **Mission Critical**—This profile is intended to be used by customers who trade off performance for server reliability beyond the basic server defaults. The profile enables advanced memory reliability, availability, and serviceability (RAS) features that are known to have more than a measurable impact to computational performance. Enabling this profile impacts maximum memory bandwidth and increases memory latency.
- **Transactional Application Processing**—This profile is intended to be used for business processing environments, such as online transaction processing (OLTP) applications that require a database back end. This includes workloads that involve a high number of user-based, transactional applications running on a single server with a cohosted database component. The profile balances the requirement of managing both peak frequency and throughput.

- **Graphic Processing**—This profile is intended for workloads that are run on server configurations that use graphics processing units (GPUs.) GPUs typically depend on maximum bandwidth between I/O and memory. Power management features that have impact on the links between I/O and memory are disabled. Peer to peer traffic is also critical and therefore virtualization is also disabled.
- **I/O Throughput**—This profile is intended to be used for configurations that depend on maximum throughput between I/O and memory. Processor utilization driven power management features that have performance impact to the links between I/O and memory are disabled.
- **Custom**—This option on the Workload Profiles menu disables workload profiles. Use this option if you want to set specific BIOS options for your deployment manually. When you select **Custom**, all the settings for the previously selected profile are carried forward. You can edit all or some of the options. Custom is not a profile and settings that you specify are not saved as a template.

Workload profile options support a variety of power and performance requirements. For most ProLiant Gen10 servers and Synergy compute modules, workload profile is set to **General Power Efficient Compute** by default. This workload profile provides common performance and power settings suitable for most application workloads. Selecting a workload profile other than the Custom profile affects other setting options. For example, selecting the **General Peak Frequency Compute** profile automatically sets Power Regulator mode to **Static High Performance**. This setting cannot be changed and is grayed out.

Multiple options are available for BIOS configuration. Not all profiles set the same options to specific settings. Each profile is designed to obtain specific performance results and sets different options to meet those results. The options that a profile sets are called dependencies. All other options are unaffected by the workload profile and are referred to as *nondependent settings*.

You apply a workload profile to have the system manage your workload according to predefined settings provided with the system. Dependent options cannot be changed and are grayed out. You can change any nondependent options in a profile.

To apply a workload profile, follow these steps:

1. From the System Utilities screen, select **System Configuration** → **BIOS/Platform Configuration (RBSU)** → **Workload Profile**.
2. Select a workload profile.
3. Optional: Change any nondependent options that you want to adjust.
4. Save and reboot to apply the workload profile.

---

**NOTE:** For more information about workload profiles, scan this QR code or right-click it to open the hyperlink.



[http://h20565.www2.hpe.com/hpsc/doc/public/display?sp4ts.oid=1009955118&docLocale=en\\_US&docId=emr\\_na-a00016408en\\_us](http://h20565.www2.hpe.com/hpsc/doc/public/display?sp4ts.oid=1009955118&docLocale=en_US&docId=emr_na-a00016408en_us)

---

## HPE iLO 5

HPE iLO 5 is a remote server management processor embedded on the system boards of ProLiant servers that provides multiple ways to configure and update servers remotely. iLO 5 makes a ProLiant server operational and boots the server.

The iLO Representational State Transfer (REST) application program interface (API) is a management interface that server management tools can use to perform server configuration, inventory, and monitoring through the iLO management processor. Using the iLO RESTful API, you can take full inventory of the server, control power and reset, configure BIOS and iLO settings, and fetch event logs, among many other functions.

The RESTful API for iLO is available on ProLiant Gen10 servers running iLO 5 1.10 or later with the iLO Standard license, although some features might not be available without an Advanced license.

### iLO 5 licensing

iLO Standard ships preconfigured on HPE servers without an additional cost or license. Features that enhance server administrator productivity and additional new security features are licensed. HPE iLO license options include:

- **HPE iLO Standard Features Support**—iLO Standard Features Support is covered under the hardware warranty for the specific server.
- **HPE iLO Support for Licensed Features**—Purchasing one- and three-year support licenses refers to the length of the support agreement for licensed iLO features. When a license is registered, HPE prepares a support agreement. After the support agreement expires, customers can continue to use licensed features. However, support for those licensed features is not available. Customers have the option to renew licensed features support.

---

**IMPORTANT:** HPE iLO license options vary by server platform and sales region. For a detailed list of options available for a specific product number, refer to the HPE iLO license product numbers. In addition, HPE iLO licenses do not expire. They are valid for the life of the server on which they are installed. For more information on iLO licensing, refer to the HPE iLO Licensing Guide: <https://www.hpe.com/us/en/servers/integrated-lights-out-ilo.html>

---

- **iLO Advanced Premium Security Edition**—Exclusively available on ProLiant Gen10 servers, the iLO Advanced Premium Security Edition license delivers all the management capabilities of the iLO Advanced license with premium security features such as Commercial National Security Algorithm (CNSA) mode, runtime firmware validation, automatic secure recovery, and secure erase of user data. CNSA mode provides the highest level of security, which is typically used for governmental confidential and top secret information. The strong cryptographic algorithms and secure protocol standards tools that CNSA mode can provide help address the need for secure, interoperable communications.

---

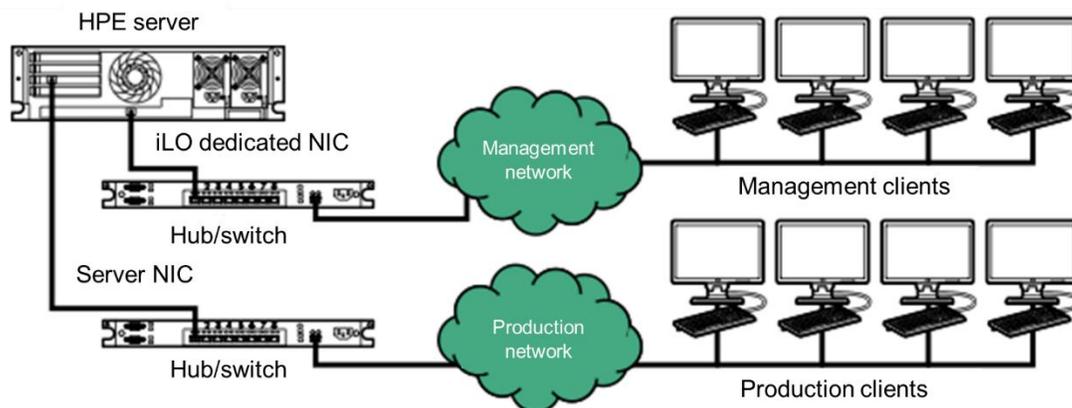
**NOTE:** To watch a video demonstrating iLO 5, scan this QR code or right-click it to open the hyperlink.



[https://www.youtube.com/watch?v=\\_rRdEekFgVA&t=24s](https://www.youtube.com/watch?v=_rRdEekFgVA&t=24s)

---

## Setting up an iLO management processor



Connecting iLO through a dedicated management network

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.
  - In a dedicated management network configuration, the iLO port is on a separate network. A separate network improves performance and security because you can physically control which workstations are connected to the network. A separate network also provides redundant access to the server when a hardware failure occurs on the production network. In this configuration, iLO cannot be accessed directly from the production network. The dedicated management network is the preferred iLO network configuration.
  - In a production network configuration, both the NIC and the iLO port are connected to the production network. In iLO, this type of connection is called the *shared network port* configuration. Certain HPE embedded NICs and add-on cards provide this capability. This connection enables access to iLO from anywhere on the network and reduces the amount of networking hardware and infrastructure required to support iLO. There are some drawbacks to using this configuration. With a shared network connection, traffic can hinder iLO performance. In addition, during the server boot process and when the operating system NIC drivers are loading and unloading, there are brief periods of time (2–8 seconds) when iLO cannot be reached from the network. After these short periods, iLO communication is restored and iLO will respond to network traffic.

- **Will NIC teaming be used with a shared network port configuration?** NIC teaming is a feature you can use to improve server NIC performance and reliability. When you select a teaming mode to use when iLO is configured to use the shared network port, iLO network communications will be blocked if the selected NIC teaming mode causes the switch that iLO is connected with to ignore traffic from the server NIC/port that iLO is configured to share. The selected NIC teaming mode sends all traffic destined for iLO to a NIC/port other than the one that iLO is configured to share. Because iLO and the server transmit and receive on the same switch port, the selected NIC teaming mode must allow the switch to tolerate traffic with two different media access control (MAC) addresses on the same switch port. Some implementations of Link Aggregation Control Protocol (LACP) (802.3ad) will not tolerate multiple MAC addresses on the same link.

If the server is configured to use HPE NIC teaming, observe the following guidelines:

- Network fault tolerance—The server transmits and receives on only one NIC (the primary adapter). The other NICs (secondary adapters) that are part of the team do not transmit server traffic and they ignore received traffic. This mode allows the iLO shared network port to function correctly. Select the NIC/port iLO uses as the preferred primary adapter.
  - Transmit load balancing—The server transmits on multiple adapters but receives only on the primary adapter. This mode allows the iLO shared network port to function correctly. Select the NIC/port iLO uses as the preferred primary adapter.
  - Switch assisted load balancing—In this mode, there is no concept of primary and secondary adapters. All adapters are considered equal for the purposes of sending and receiving data. This mode is the most problematic for iLO shared network port configurations because traffic destined for iLO can be received on only one of the server NICs/ports.
- **How will iLO acquire an IP address?** To enable iLO access after it is connected to the network, the iLO management processor must acquire an IP address and subnet mask. You can use a dynamic address or a static address.
  - **What access security is required, and what user accounts and privileges are needed?** You can use the following methods to manage access to iLO:
    - Local accounts—Up to 12 user accounts can be stored in iLO. This configuration is ideal for small environments such as labs and SMBs. Login security with local accounts is managed through iLO Access Settings and user privileges.
    - Directory services—Up to six directory groups can be configured in iLO. Use a directory to authenticate and authorize iLO access. This configuration enables an unlimited number of users and easily scales to the number of iLO devices in an enterprise. If you plan to use directory services, consider enabling at least one local administrator account for alternate access. A directory provides a central point of administration for iLO devices and users, and the directory can enforce a strong password policy.
    - Common Access Card (CAC) smartcard authentication—You can configure CAC smartcards together with local accounts and directory services to manage iLO user access.
  - **What tools will you use to configure iLO?**—iLO supports various interfaces for configuration and operation, including the iLO web interface and the ROM-based setup.

---

**NOTE:** For more information, refer to the iLO and iLO RESTful API documentation by scanning this QR code or right-clicking it to open the hyperlink.



<http://www.hpe.com/support/ilo-docs>

---

## iLO 5 service port on Gen10 servers



iLO 5 service port on a ProLiant DL380 Gen10 server

The iLO 5 service port is a USB port with the label "iLO" on the front of ProLiant Gen10 servers and Synergy Gen10 compute modules.

If you have physical access to the server, you can use the iLO service port to connect a client (such as a laptop) with a supported USB-to-Ethernet adapter to access the iLO 5 web interface, remote console, CLI, iLO RESTful API, or a scripting utility.

---

**NOTE:** HPE recommends using the HPE USB to Ethernet Adapter (part number Q7Y55A) to connect a client to the iLO service port to access the server directly.

---

The IPv4 address 169.254.1.2 is used to connect a client to any server through the iLO service port; this address cannot be changed.

To use this port, you must make sure that:

- The iLO service port and USB Ethernet adapters options are enabled.
- The client network adapter is configured to support the service port feature.

After you are connected, you can access the iLO Configuration Utility. You can also use the service port to download the Active Health System Log to a supported USB flash drive. However, you cannot use the service port to boot any device within the server, or the server itself, and you cannot access the server by connecting to the service port.

When you use the iLO service port:

- Actions are logged in the iLO Event Log.
- The server UID blinks to indicate the service port status. You can also retrieve the status by using a REST client and the iLO RESTful API.
- The connected USB flash drive is not accessible by the host operating system.
- You can connect a USB flash drive to download the Active Health System Log.

## Using the iLO 5 Configuration Utility

HPE recommends using the iLO 5 Configuration Utility to set up iLO for the first time and to configure iLO network parameters for environments that do not use Dynamic Host Configuration Protocol (DHCP), Domain Name System (DNS), or Windows Internet Name Service (WINS).

---

**NOTE:** To simplify installation, HPE recommends using DNS or DHCP with iLO 5.

---

You can use the iLO 5 Configuration Utility in the UEFI System Utilities to configure network parameters, global settings, and user accounts. The iLO default settings enable you to use most features without additional configuration. The utility has the following options:

- Network Options
- Advanced Network Options
- User Management
- Setting Options
- Set to factory defaults
- Reset iLO (active connections)

iLO server group administration automates configuring and managing large deployments of iLO processors. Using iLO extensive scripting language with the iLO 5 Configuration Utility, you can easily configure all settings for mass deployments on multiple ProLiant iLO processors.

You can access the iLO 5 Configuration Utility from the physical system console, or by using an iLO 5 remote console session. To access the iLO 5 Configuration Utility, press **F9** during POST to start the UEFI System Utilities. Click **System Configuration**, and then click **iLO 5 Configuration Utility**.

---

**NOTE:** The iLO 5 Configuration Utility is designed for the initial iLO setup, and is not intended for continued iLO administration. You can start the utility when the server is booted, and you can run it remotely with the remote console.

---

## Logging in to iLO for the first time

To log in to iLO for the first time, follow these steps:

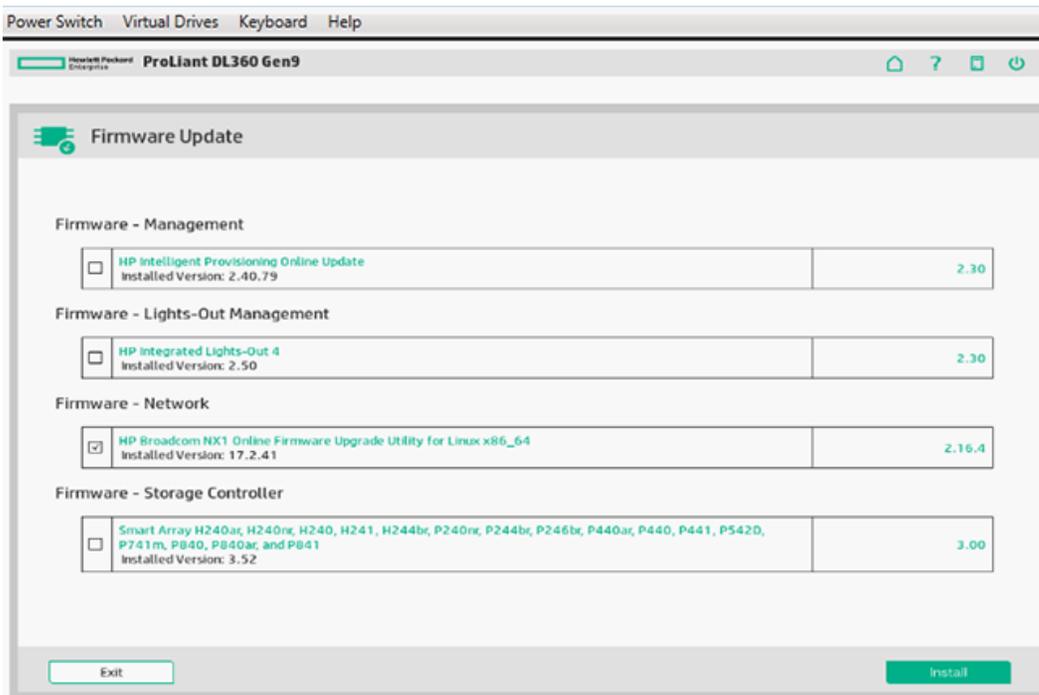
1. Enter `https://<iLO hostname or IP address>`.
2. Enter the default user credentials, and then click **Log In**. The iLO firmware is configured with a default user name, password, and DNS name. The default information is on the serial label pull tab attached to the server that contains the iLO management processor. Use these values to access iLO remotely from a network client through a web browser.
  - User name—Administrator
  - Password—A random eight-character string
  - DNS name—ILOXXXXXXXXXXXX, where the X characters represent the server serial number

---

**IMPORTANT:** HPE recommends changing the default password after you log in to iLO for the first time. If you reset iLO to the factory default settings, use the default iLO account credentials to log in after the reset.

---

## Intelligent Provisioning



Intelligent Provisioning is a single-server deployment tool embedded in ProLiant servers and Synergy compute modules. It simplifies server setup, providing a reliable and consistent way to deploy server configurations. Intelligent Provisioning enables you to perform tasks such as:

- Installing an operating system
- Updating firmware, software, and drivers
- Performing diagnostic tests

---

**NOTE:** Not all versions of an operating system are supported. For the latest information, refer to the operating system support matrix by scanning this QR code or by right-clicking it to open the hyperlink.



<http://www.hpe.com/info/ossupport>

---

## Always On Intelligent Provisioning

In addition to accessing Intelligent Provisioning by pressing **F10** from the POST screen, you can also access Intelligent Provisioning from the iLO web browser user interface using Always On. Always On Intelligent Provisioning allows you to:

- Perform functions when the server is off
- Perform tasks when running an operating system without powering off the server

The Intelligent Provisioning F10/Remote console has more options than Always On Intelligent Provisioning. F10/Remote console options include:

- Attempt Firmware Update
- Intelligent Provisioning Preferences
- Active Health System Log
- Deployment Settings
- BIOS Configuration (RBSU)
- iLO Configuration
- System Erase and Reset
- RAID Configuration

You can access Always On without having to reboot the server.

---

**NOTE:** Always On Intelligent Provisioning is limited to only using network install sources. If using Always On, only the Assisted Install option (express installation using HPE defined defaults) is listed.

---

## Software installed with Intelligent Provisioning

Intelligent Provisioning prepares the system for installing original, licensed vendor media and HPE branded versions of operating system software. Intelligent Provisioning also prepares the system to integrate optimized server support software from the HPE Service Pack for ProLiant (SPP). SPP is a comprehensive system software and firmware solution for ProLiant servers, server blades, and their enclosures as well as Synergy compute modules. These components are preloaded with a basic set of firmware and operating system components that are installed along with Intelligent Provisioning.

After the server is running, you can update the firmware to install additional components. You can also update any components that have been outdated since the server was manufactured.

When a Microsoft Windows operating system is installed by using Intelligent Provisioning with internet access, all the software applications are automatically downloaded and installed. On other operating systems or on a Windows system without internet access, the following applications are not automatically installed with Intelligent Provisioning. To install the following applications, run the SPP:

- ProLiant Agentless Management Service (AMS)
- Network Configuration Utility for Windows
- HPE Smart Storage Administrator
- Lights-Out Online Configuration Utility

The Quick Configs option in Intelligent Provisioning enables you to select a basic policy for performance, instead of power usage, without having to configure individual settings through the UEFI System Utilities menus.

The Intelligent Provisioning option lets you select the Intelligent Provisioning host override option for this boot only. It does not modify the normal boot order or boot mode settings.

---

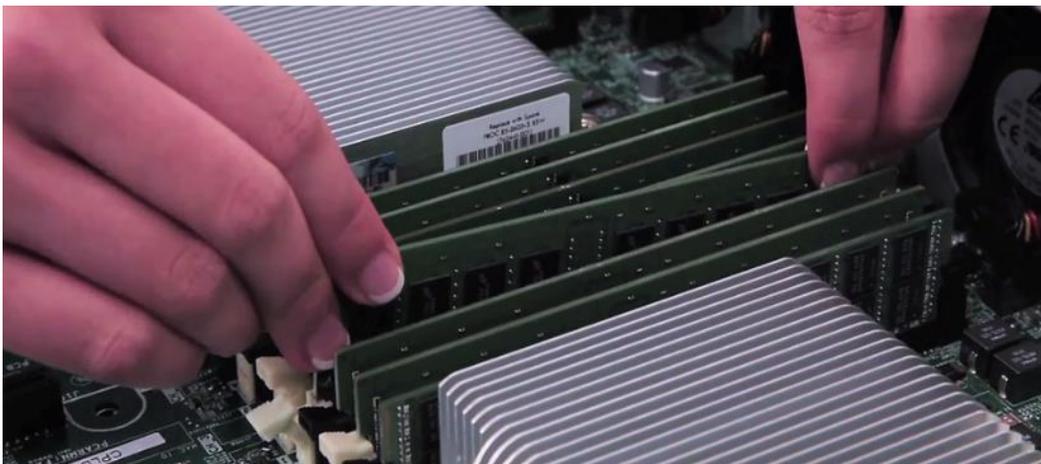
**NOTE:** For more information, refer to the Intelligent Provisioning user guides by scanning this QR code or right-clicking it to open the hyperlink.



<http://www.hpe.com/info/intelligentprovisioning/docs>

---

## Installing memory in ProLiant Gen10 servers



Internal server memory

ProLiant Gen10 servers feature a slightly different memory architecture than that introduced with ProLiant Gen9 servers. ProLiant Gen10 servers using the Intel Xeon Scalable processor family with six separate memory channels per CPU and up to 24 DIMM slots in two-socket servers and 48 DIMM slots in four-socket servers. These architectures allow large memory configurations and deliver improved memory performance. They also incorporate HPE Advanced Memory Protection technology, which improves the prediction of critical memory error conditions.

ProLiant Gen10 systems support a variety of flexible 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 ProLiant Gen10 servers with HPE DDR4 memory. Violating these rules might result in reduced memory capacity, decreased performance, or error messages during boot.

---

**NOTE:** To watch a video about HPE qualified memory, scan this QR code or right-click it to open the hyperlink.



<https://www.hpe.com/us/en/product-catalog/servers/server-memory/pip.hpe-ddr4-smartmemory.7281077.html>

For details on population rules for DDR4 memory with ProLiant Gen10 servers, scan this QR code or right-click it to open the hyperlink.



<http://www.hpe.com/docs/memory-population-rules>

---

## DIMM population rules for ProLiant Gen10 servers

The following table summarizes the overall DIMM population rules for ProLiant Gen10 servers.

Category	Population guidelines
Processors and DIMM slots	<ul style="list-style-type: none"> <li>• 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.</li> <li>• If a memory channel consists of more than one DIMM slot, the white memory slot is located farthest from the CPU. White DIMM slots denote the first slot to be populated in a channel. For one DIMM per channel, populate white memory slots only.</li> <li>• 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 memory slot. Within a given channel, populate DIMMs from the heaviest electrical load (dual rank) to the lightest load (single rank).</li> <li>• If multiple CPUs are populated, split the DIMMs evenly across the CPUs and follow the corresponding CPU rule when populating DIMMs.</li> </ul>
Performance	To maximize performance, HPE recommends balancing the total memory capacity across all installed processors and loading the channels similarly whenever possible. If the number of DIMMs does not spread evenly across the CPUs, populate as close to even as possible.
DIMM types and capacities	<ul style="list-style-type: none"> <li>• 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.</li> <li>• Do not mix RDIMMs and LRDIMMs in the same system. Do not mix 128 GB LRDIMMs with other capacity DIMMs. Unbuffered DIMMs are not supported.</li> <li>• x4 and x8 DIMMs can be mixed in the same channel.</li> </ul>
DIMM speed	<ul style="list-style-type: none"> <li>• The maximum memory speed is a function of the memory type, memory configuration, and processor model.</li> <li>• DIMMs of different speeds may be mixed in any order; however, the server selects the highest common speed among all the DIMMs per CPU.</li> <li>• HPE memory from previous generation servers is not compatible with the current generation. Certain HPE SmartMemory features such as memory authentication and enhanced performance might not be supported.</li> </ul>
Heterogeneous mix	<ul style="list-style-type: none"> <li>• There are no performance implications for mixing sets of different capacity DIMMs at the same operating speed. For example, latency and throughput will not be negatively impacted by installing an equal number of 16 GB dual-rank DDR4-2666 DIMMs (one per channel) and 32 GB dual-rank DDR4-2666 DIMMs (one per channel).</li> <li>• Take each DIMM type and create a configuration as if it were a homogeneous configuration.</li> <li>• Depending on the per-channel rules, populate the DIMMs with highest rank count in white memory slots in each channel. Populate the other DIMMs in the black memory slots in each channel.</li> </ul>

**NOTE:** For more information, refer to the *Population rules for HPE DDR4 memory with HPE ProLiant and Synergy Gen10 servers* technical white paper available from the Hewlett Packard Enterprise Information Library (<http://www.hpe.com/info/enterprise/docs>). You can also find information about DIMM population rules on the inside of the server chassis cover.

## Understanding balanced memory configurations

HPE ProLiant Gen10 12 slot per CPU DIMM population order													
1 DIMM											8		
2 DIMMs											8	10	
3 DIMMs											8	10	12
4 DIMMs			3		5						8	10	
5 DIMMs*			3		5						8	10	12
6 DIMMs	1		3		5						8	10	12
7 DIMMs*	1		3		5		7	8			10	12	
8 DIMMs			3	4	5	6	7	8	9	10			
9 DIMMs*	1		3		5		7	8	9	10	11	12	
10 DIMMs*	1		3	4	5	6	7	8	9	10			12
11 DIMMs*	1		3	4	5	6	7	8	9	10	11	12	
12 DIMMs	1	2	3	4	5	6	7	8	9	10	11	12	

\* Unbalanced, not recommended

HPE ProLiant Gen10 8 slot per CPU DIMM population order												
1 DIMM											3	
2 DIMMs			2	3								
3 DIMMs	1	2	3									
4 DIMMs		2	3							6	7	
5 DIMMs*	1	2	3							6	7	
6 DIMMs	1	2	3							6	7	8
7 DIMMs*	1	2	3	4						6	7	8
8 DIMMs*	1	2	3	4	5	6	7	8				

\* Unbalanced, not recommended

DIMM population order for ProLiant Gen10 servers with 12 slots and 8 slots per CPU

Optimal memory performance is achieved when the system is configured with a fully homogeneous and balanced memory configuration. Unbalanced memory configurations are those in which the installed memory is not distributed evenly across the memory channels or the processors. HPE discourages unbalanced configurations because they have lower performance than similar balanced configurations. The two types of unbalanced configurations have different performance implications:

- **Unbalanced across channels**—A memory configuration is unbalanced across channels if the memory installed on each populated channel is not identical.
- **Unbalanced across processors**—A memory configuration is unbalanced across processors if a different amount of memory is installed on each of the processors.

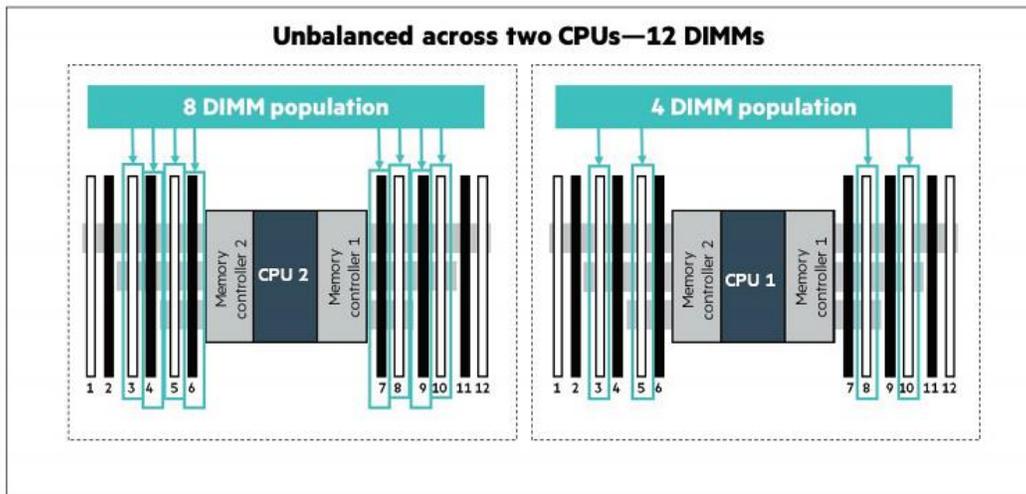
### Memory configurations that are unbalanced across channels

In unbalanced memory configurations across channels, the memory controller splits memory up into regions. In a balanced configuration, there is one region that includes all installed DIMMs. If the memory configuration is unbalanced, then it attempts to create multiple balanced regions. First, it creates the largest possible balanced region with the installed memory. The next largest region is next and so on. In this manner, the memory controller creates regions until all installed memory has been assigned to a region.

Memory should be installed as indicated in the preceding graphic based on the total number of DIMMs being installed per CPU. For example, if two DIMMs are being installed per CPU, they should be located in white memory slots numbered 8 and 10. If six DIMMs are being used per CPU, they should be installed in memory slots 1, 3, 5, 8, 10, and 12. Unbalanced configurations are noted with an asterisk in the graphic and are not recommended because memory performance will be inconsistent or degraded compared to a balanced configuration.

The primary effect of memory configurations that are unbalanced across channels is a decrease in memory throughput in those regions that span fewer memory channels. In the unbalanced example in the graphic, worst case measured memory throughput in Region 2 would be 33% or less than the throughput in the balanced example. Even in Region 1 in the unbalanced example, throughput would be limited to no more than 66% of what the single region in the balanced example could provide.

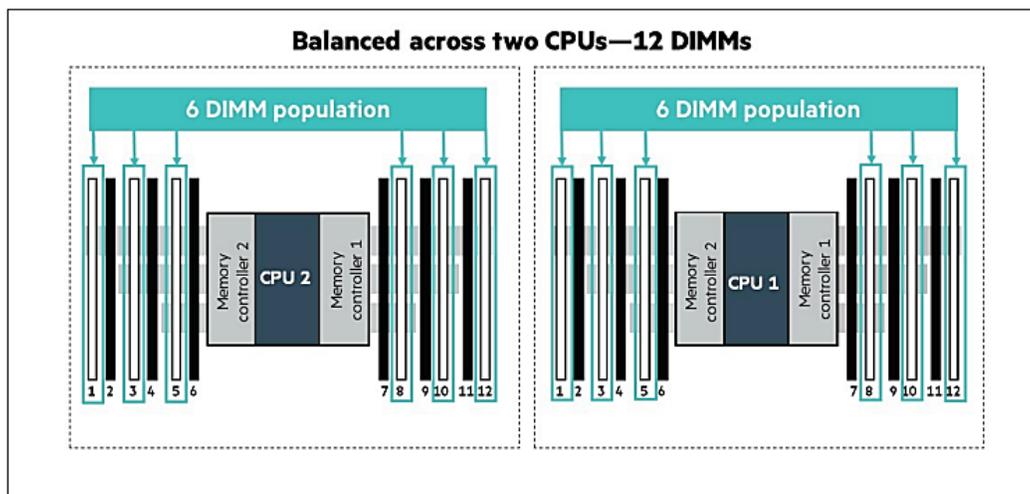
## Memory configurations that are unbalanced and balanced across processors



Example of memory that is unbalanced across two processors

The preceding graphic shows a memory configuration that is unbalanced across processors. In this example, one processor contains four DIMMs and the other CPU has eight DIMMs installed.

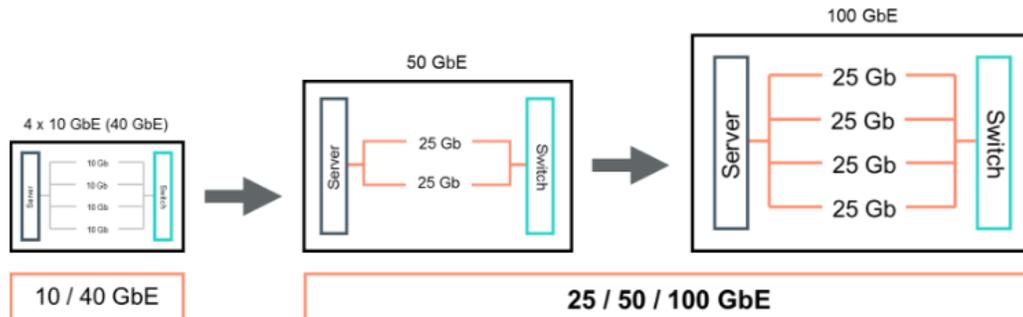
The CPU 1 threads operating on the larger memory capacity of CPU 1 might have adequate local memory with relatively low latencies and high throughput. The CPU 2 threads operating on the smaller memory capacity of CPU 2 might consume all available memory on CPU 2 and request remote memory from CPU 1. The longer latencies and limited throughput of cross-CPU communications associated with the remote memory will result in reduced performance of those threads. In practice, this might result in nonuniform performance characteristics for software program threads, depending on which processor executes them.



Example of a memory configuration that is balanced across two processors

This graphic shows an example of a configuration that is balanced across processors. In this example, both processors have six DIMMs installed.

## Installing HPE 25Gb/100Gb Ethernet adapters



HPE offers a range of Ethernet adapters

ProLiant servers with HPE 25Gb Ethernet adapters provide the largest jump in Ethernet performance at an affordable price. They also accelerate server and application performance. The standard of 25 GbE leverages technology defined for 100 GbE implemented as four 25 Gbps lanes (IEEE 802.3bj) running on four fiber or copper pairs.

These solutions are designed and tested to deliver the industry's most comprehensive 25 GbE/100 GbE networking ecosystem enabling end-to-end 25 GbE server connectivity across a 100 GbE data center network fabric.

A data center enabled with a 25 GbE/100 GbE networking infrastructure allows:

- 150% more bandwidth than 10 GbE with 50% less cabling
- 40% more streaming HD videos

Customers who would benefit from 25 GbE/100 GbE include:

- Cloud, telecommunications, and service providers experiencing 10 Gbps/40 Gbps saturation
- Financial and e-commerce sites demanding increased performance

HPE 25Gb/100Gb Ethernet adapters can be installed in a PCIe expansion slot in ProLiant DL360 and DL380 Gen9 and Gen10 servers.

---

**NOTE:** To watch a video on 25 GbE/100 GbE and enhanced NIC security, scan this QR code or right-click it to open the hyperlink.



[https://www.hpe.com/h22228/video-gallery/us/en/products/networking/networking-\(general\)/A20EEB68-3F94-4CCB-BEAC-F4D2FE330CD8/Take-the-25Gb-Express-Lane/video/](https://www.hpe.com/h22228/video-gallery/us/en/products/networking/networking-(general)/A20EEB68-3F94-4CCB-BEAC-F4D2FE330CD8/Take-the-25Gb-Express-Lane/video/)

---

## Learning check

1. Choose three functions that UEFI supports that legacy BIOS did not support. (Select three.)
  - a. IPv4 and IPv6 PXE boot
  - b. Master Boot Record (MBR)
  - c. 32-bit and 64-bit processors
  - d. Boot partitions larger than 2.2 TB
  - e. Advanced preboot user interfaces
2. You can access Intelligent Provisioning Always On without having to reboot the server.  
 True  
 False

## Configuring HPE storage

Your decision to use a direct-attached storage (DAS), network-attached storage (NAS), or storage area network (SAN) configuration in the data center determines how to connect the storage system to the devices that need to access it.

A DAS solution does not use the network because it is connected directly to the server by means of a HBA installed in a standard PCIe slot in the host. Only the host attached to the DAS solution can access the files stored there. A NAS solution is typically configured with an IP address on an Ethernet network, and is used as a central access or backup location for every host attached to it. Depending on the workloads and other requirements of the data center, a SAN solution can be configured with a dedicated gigabit Ethernet switch and a private subnet, or on a Fibre Channel network.

## Designing a SAN using Fibre Channel switches

HPE SAN switches offer the simplest and most affordable way to share the capacity of Fibre Channel storage devices among multiple servers. You can connect as few as two devices or scale up to 20 Fibre Channel devices with a single switch. Mix any combination of servers, disk storage, or backup storage for maximum flexibility and investment protection.

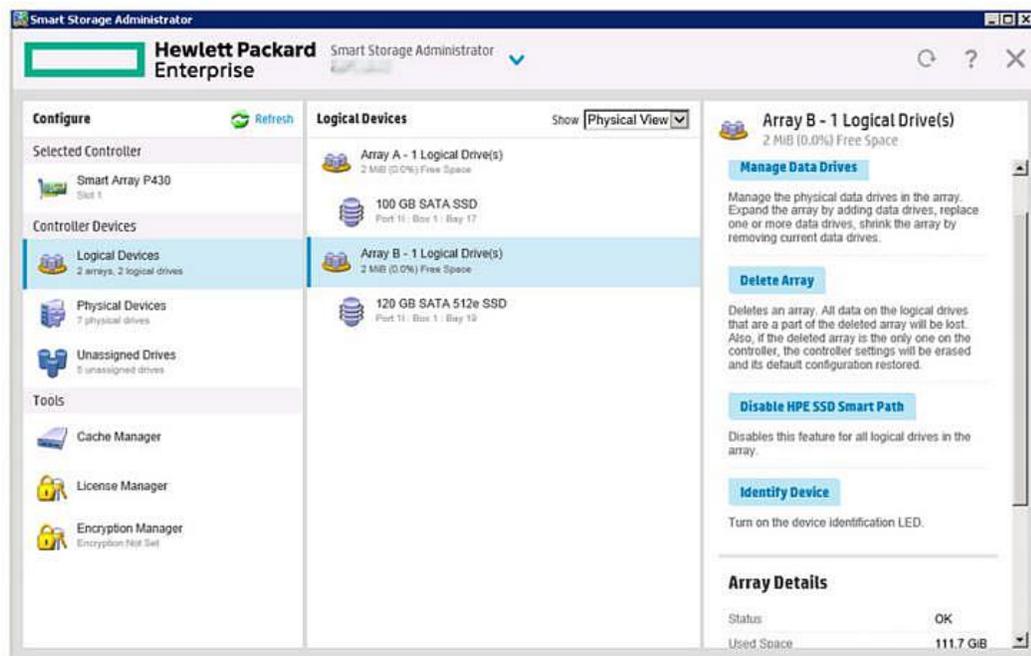
With the dependency on fewer storage devices in a consolidated SAN and the increasing number of virtualized servers on fewer server hardware platforms, the need to build redundancy into the SAN is critical.

A SAN fabric topology defines the arrangement of Fibre Channel switches in a fabric. There are three approaches to designing a SAN. You can implement:

- An HPE standard SAN fabric topology design
- A subset or variation of an HPE standard SAN fabric topology design
- A custom SAN fabric topology design

HPE offers an extensive portfolio of Fibre Channel and multipurpose/converged protocol switches, transceivers, and cables for SANs and LANs that deliver the scalability, performance, and interoperability demanded from mission-critical data and applications. Combining ease of use, flexibility, and improved security with cost-effective pricing, HPE switches are key enablers of a dynamic and adaptive storage networking environment. They provide an ideal platform for adaptation as business requirements grow and change, in terms of basic connectivity and advanced SAN or converged fabric functionality.

## Using HPE Smart Storage Administrator



Configuring logical devices with SSA

The Smart Storage Administrator (SSA) is a web-based application that helps you configure, manage, diagnose, and monitor HPE Smart Array controllers and HBAs. It is the main tool for configuring arrays on Smart Array controllers. It replaces the HPE Array Configuration Utility (ACU) with an updated design and functionality.

SSA features and functions include:

- **HPE Secure Encryption**—Is a data encryption solution for ProLiant servers that protects data at rest on any bulk storage attached to a Smart Array controller.
- **SSD Over Provisioning Optimization**—Optimizes solid-state drives (SSDs) by deallocating all used blocks before data is written to the drive. The optimization process is performed when the first logical drive in an array is created and when a failed drive is replaced with a physical drive.
- **Rapid Rebuild Priority**—Determines the urgency with which a controller treats an internal command to rebuild a failed logical drive. SSA offers four settings: low, medium, medium high, and high.
- **Auto RAID 0**—Creates a single RAID 0 volume on each physical drive specified, enabling the user to select multiple drives and configure as RAID 0 simultaneously.

SSA Scripting is a stand-alone application that is distributed with the HPE SSA CLI application. Users familiar with ACU scripting must install the HPE SSA CLI application to obtain the scripting executable. The HPE SSA scripting executable (ssascripting), formerly hpssascripting, replaces the former executable (hpacascripting) in all scripts.

Both SSA and the SSA CLI need sg driver (scsi generic) to access the Smart Array controller in Linux. Red Hat Enterprise Linux 7.1 and later does not load sg driver automatically. You should load the driver and enter `modprobe sg` before running SSA or SSA-CLI.

## Accessing SSA

SSA is accessible both offline—as a stand-alone bootable ISO image—and online:

- **Accessing SSA offline**—Using one of multiple methods, you can run SSA before launching the host operating system. In offline mode, users can configure or maintain detected and supported devices, such as optional Smart Array controllers and integrated Smart Array controllers. Some SSA features are only available in the offline environment, such as setting the boot controller and boot volume.
- **Accessing SSA online**—This method requires an administrator to download the SSA executables and install them. All formats provide support for standard configuration tasks. SSA also supports 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.

You can access and launch HPE SSA in an offline environment in three ways:

- Launch SSA with Intelligent Provisioning. During server boot, press **F10**. Select **Perform Maintenance** → **HPE Smart Storage Administrator (HPE SSA)**.
- Launch SSA during POST. During server boot, press **F10** and select a method to start Intelligent Provisioning.
- Launch SSA from an ISO image. To access the SSA CLI or SSA scripting in an offline environment, you must launch SSA from an ISO image. To prepare the image, use one of the following methods:
  - Mounting the image through iLO
  - Burning the image to a CD or DVD
  - Flashing the image to a USB memory key or SD card on an HPE UEFI bootable server
  - Installing the image on a PXE server

Booting from the ISO image on a drive, on a key, or through iLO provides the same GUI interface. The user can select to run offline SSA GUI, SSA CLI, or SSA Scripting.

To access, install, and launch SSA in the online environment, you must download the HPE SSA executables. All three formats have separate executables.

To use SSA in the online environment, obtain the executable files from the HPE website and follow the installation instructions provided with the executable.

---

**IMPORTANT:** If you are updating an existing server in an offline environment, obtain the latest version of HPE SSA through an HPE SPP before performing configuration procedures. For more information on SPPs, scan this QR code or right-click it to opening the hyperlink.



<http://www.hpe.com/info/spp>

---

**NOTE:** For more information, refer to the *HPE Smart Storage Administrator User Guide* and the *HPE Smart Array SR Gen10 Configuration Guide* by scanning the following QR codes or by right-clicking them to open the hyperlinks.



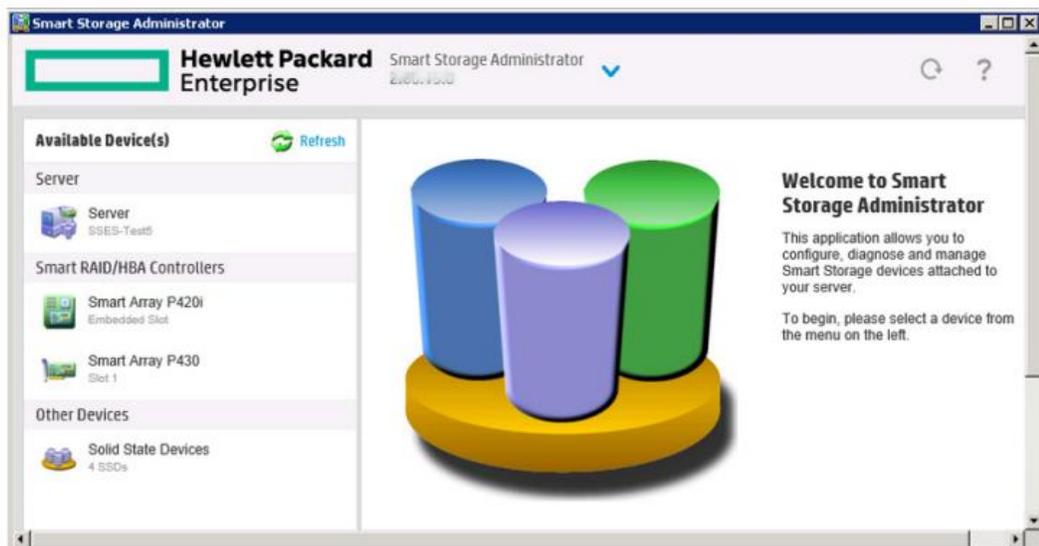
<http://h20566.www2.hpe.com/hpsc/doc/public/display?docId=c03909334>



[http://h20566.www2.hpe.com/hpsc/doc/public/display?docId=a00018944en\\_us](http://h20566.www2.hpe.com/hpsc/doc/public/display?docId=a00018944en_us)

---

## Navigating the GUI



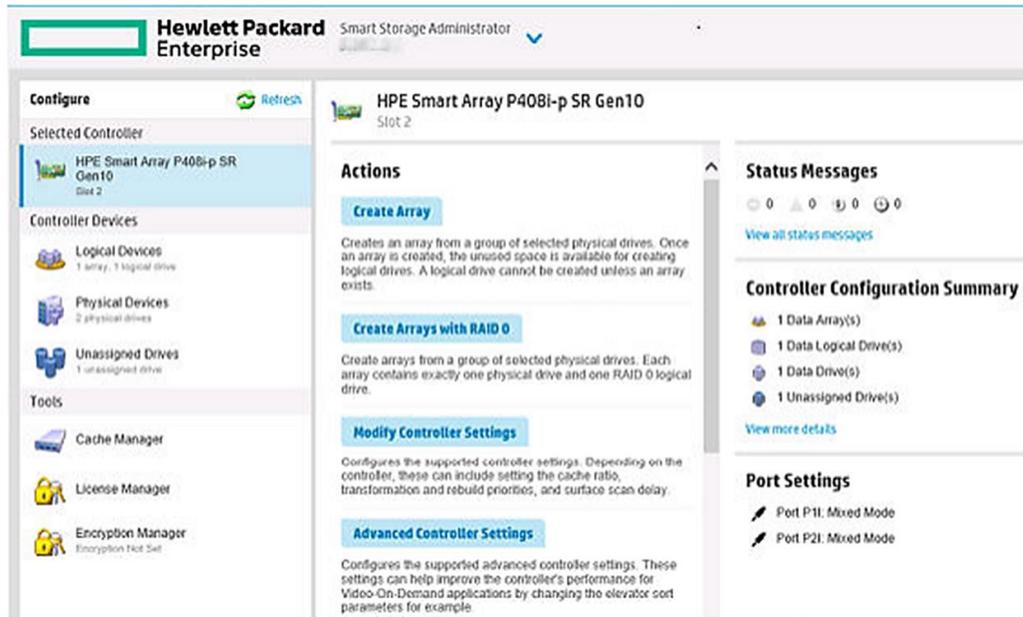
SSA Welcome screen

When you launch SSA, the application opens and SSA scans the system and detects controllers. This process can last up to two minutes. When controller detection is complete, the controllers are available on the Devices/Tools menu. After the GUI is open, tasks are distributed among categories.

When you open the SSA GUI, the Welcome screen is displayed. The following elements are visible:

- The SSA quick navigation menu is in the top-left corner of the screen. Clicking the down arrow displays the available devices, and clicking one of the available devices displays additional information and options for the device. Clicking a server or array controller reveals the available actions, alerts, and summary for that device. You can also return to a server Home screen, or you can choose **Configuration** or **Diagnostics** for a device listed.
- Available devices are listed on the left side of the screen. Clicking a server or array controller displays the available actions, alerts, and summary for that device. You can point to the status alerts to see details on an alert.
- The Rescan System icon is near the top-right of the screen. After adding or removing devices, click the icon to update the list of available devices.
- The Help button is at the top-right of the screen. To access help topics, press the **H** key or click **Help**.

## SSA actions



SSA Actions menu options

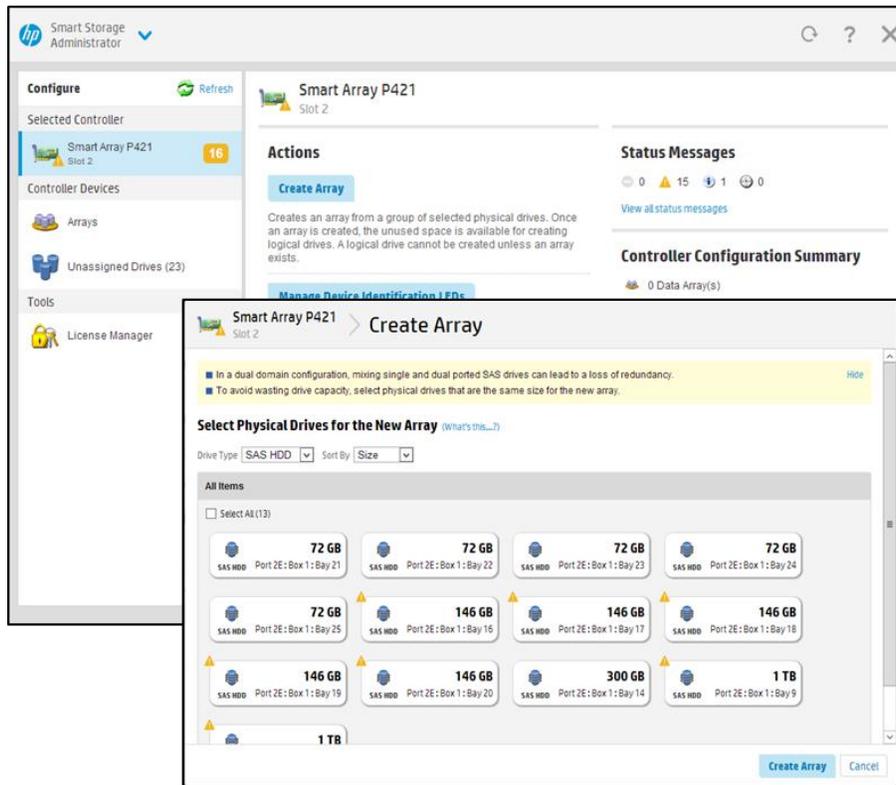
When a controller is selected from the Configure screen, the following elements are displayed:

- **Selected Controller, Controller Devices, and Tools**—This panel, on the left of the page, displays systems, controllers, arrays, physical drives, logical drives, and a cache and license manager.
- **Actions**—This panel, in the middle of the page, details:
  - Tasks that are available for the selected device based on its current status and configuration
  - Options and information pertinent to each task, after a task is selected
- **Status Messages**—This panel provides:
  - Status icons with the number of individual alerts for each category
  - A view all status messages link that displays device-specific alerts
- **Controller Configuration Summary**—This panel summarizes the data arrays, data drives and logical drives, and unassigned drives associated with a controller.

Selecting a controller from the quick navigation menu in the left pane displays the Actions page for that controller. Available actions include:

- **Configure**—Select this option to modify controller settings and advanced controller settings, modify spare activation mode, and clear the configuration, among others. To access the SSA Configure screen, select an available device from the Home screen, and then click **Configure** under the available options.
- **Diagnose**—Generate an array diagnostic report and SmartSSD Wear Gauge Report.

## Creating an array



SSA Create Array screen

When you create an array, you can select drives, specify the RAID level, and configure array settings, including stripe size and logical drive size. To create a new array (logical drive):

1. Select a controller and click **Create Array**.
2. Select the physical drives for the new logical drive.
3. Make selections for RAID level, strip size/full stripe size, sectors/track, and size.
4. Click **Create Logical Drive**.

---

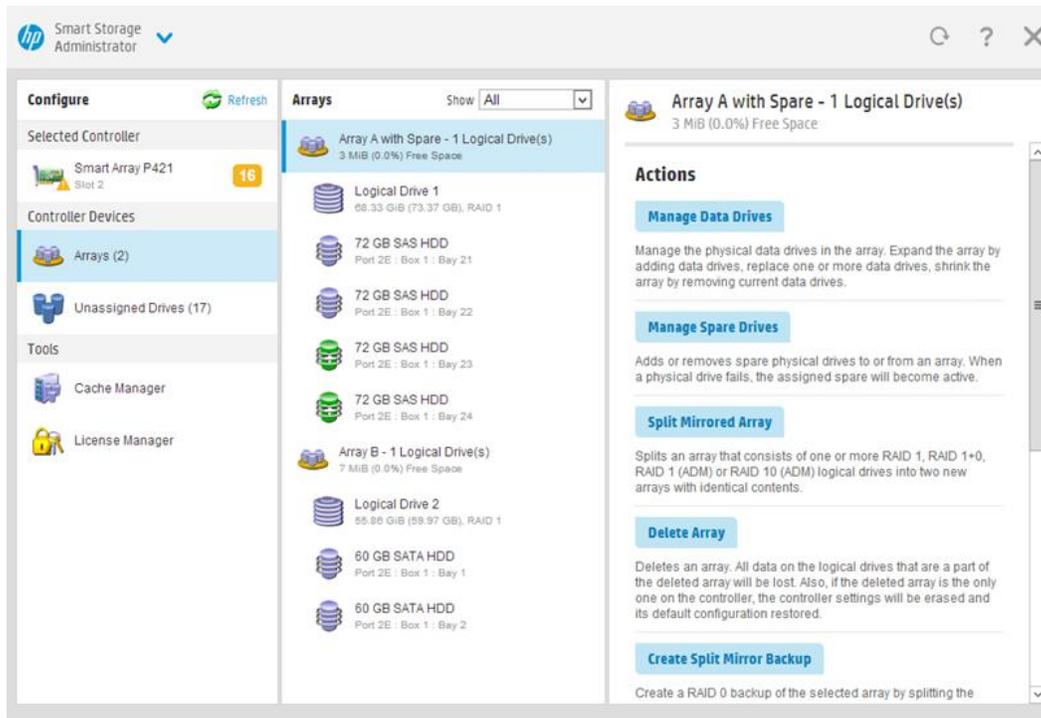
**NOTE:** If you use HPE Smart Array S100i SW RAID, SAS drives are not supported.

---

Keep in mind the following points when you build an array:

- All drives grouped in a logical drive must be of the same type (for example, either all SAS or all SATA and either all hard drives or all solid-state drives).
- For the most efficient use of drive space, all drives within an array should have approximately the same capacity. Each configuration utility treats every physical drive in an array as if it has the same capacity as the smallest drive in the array. Any excess capacity of a particular drive cannot be used in the array and is unavailable for data storage.
- The more physical drives configured in an array, the greater the probability that the array will experience a drive failure during any given period.
- To guard against the data loss that occurs when a drive fails, configure all logical drives in an array with a suitable fault-tolerance (RAID) method.

## SSA Array Details screen



Access specific details on each array

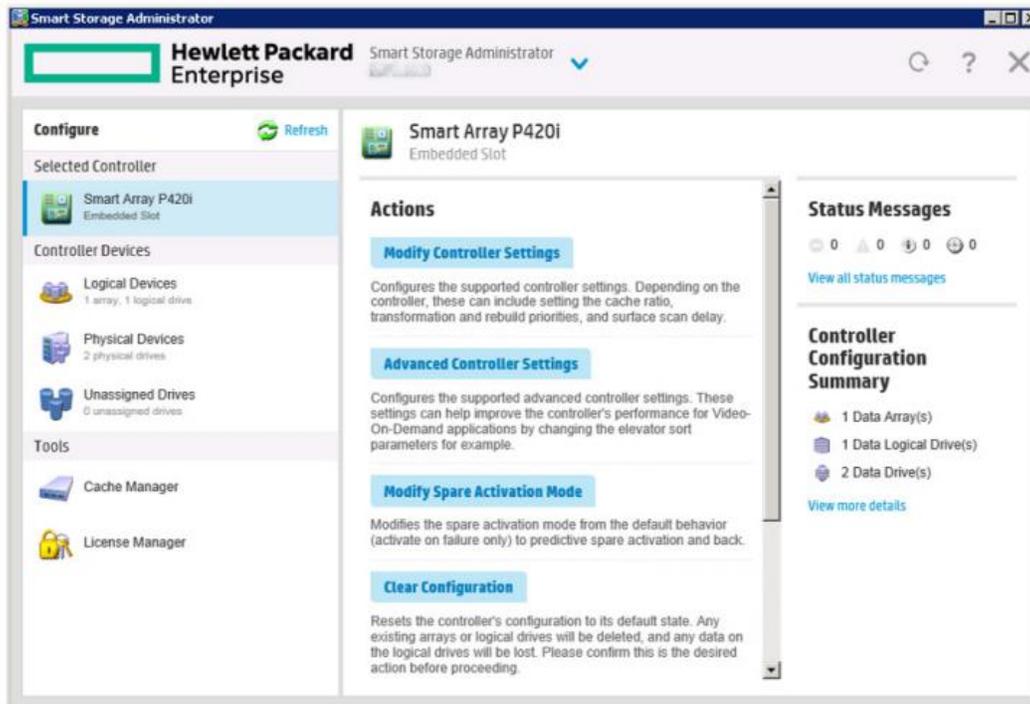
The Array Details page displays logical drives and their member physical drives. Actions include:

- **Manage data drives**—Manage the physical data drives in the array. Expand the array by adding data drives, replace one or more data drives, and shrink the array by removing current data drives.
- **Manage spare drives**—Add or remove spare physical drives to or from an array. When a physical drive fails, the assigned spare becomes active.
- **Split mirrored array**—Split an array that consists of one or more RAID 1, RAID 1+0, RAID 1 (ADM), or RAID 10 (ADM) logical drives into two new arrays with identical components.
- **Delete array**—Delete an array. All data on the logical drives that are part of the deleted array will be lost. If the deleted array is the only one on the controller, the controller settings will be erased and the default configuration restored.
- **Create split mirror backup**—Create a RAID 0 backup of the selected array by splitting an array consisting of one or more RAID 1, RAID 1+0, RAID 1 (ADM), or RAID 10 (ADM) logical drives and create two arrays: a primary array and a backup array. The primary array continues to be fully accessible to the operating system, but the backup array is hidden from the operating system.



**CAUTION:** If you select RAID 0 for any logical drives, you will experience data loss for that logical drive if one physical drive fails. Assign RAID 0 to drives that require large capacity and high speed, but pose no data safety risk.

## SSA Configure screen



SSA configuration options

To access this screen, click either a device under Configuration in the quick navigation menu or select an available device from the Home screen, and then click **Configure** under the available options.

The Configure screen displays the GUI elements from the Welcome screen and lists available actions, status messages, more detailed information, and a controller configuration summary for a selected controller. From the Configure screen, you can perform tasks related to controllers, arrays, physical drives, and logical drives.

When a controller or device is selected, the tasks shown are a subset of the total number of possible tasks for the selected item. HPE SSA lists or omits tasks based on the controller model and configuration. For example, if the selected controller has no unassigned physical drives, Create Array is not an available task.

## Activity—Researching SSA

To complete this activity, use the HPE SSA user guide to answer the questions. Keep IDJT HealthCare in mind as you consider the features of SSA.

1. Download the *HPE Smart Storage Administrator User Guide*:  
<http://h20566.www2.hp.com/hpsc/doc/public/display?docId=c03909334>

### Questions

1. What are the methods to launching SSA from an ISO image?

---

---

---

---

2. On the SSA Configure screen, if the selected controller has no unassigned physical drives, is it possible to create an array?

---

3. What is HPE SmartCache?

---

---

---

---

## Configuring HPE MSA storage solutions

HPE Modular Smart Array (MSA) storage solutions are configured via the HPE Storage Management Utility (SMU), which is a web-based application for configuring, monitoring, and managing the storage system. Each controller module in an MSA contains a web server that is accessed when you sign in to the SMU. In a dual-controller system, you can access all functions from either controller. If one controller is unavailable, you can continue to manage the storage system from the partner controller. In addition to the SMU, each controller module in the storage system has a CLI and an FTP interface, and SNMP and SMI-S interfaces. To configure and provision an MSA storage system for the first time through the SMU:

1. Sign in to the SMU. When signing in, the default user name for management is `manage` and the default password is `!manage`.
2. Verify that controller modules and expansion modules have the latest firmware.
3. Open the Configuration Wizard by hovering the cursor over the **Home** tab and selecting **Configuration Wizard**. The wizard is also accessible from the **Home** topic by selecting **Action** → **Configuration Wizard**.
4. Create virtual and linear disk groups and pools, and add dedicated spares to linear disk groups.
5. Create volumes and map them to initiators.

---

**NOTE:** An initiator is analogous to an external port on an HBA. An initiator port does not equate to a physical server, but rather a unique connection on that server. For example, a dual-port Fibre Channel HBA has two ports and therefore there are two unique initiators. The array will show two separate initiators for that HBA. MSA firmware can support 512 hosts with multiple initiators per host. An MSA system can manage 1024 initiators. The array supports the grouping of initiators under a single host and grouping hosts into a host group. The grouping of initiators and hosts allows simplification of the mapping operations.

---

6. From hosts, verify volume mappings by mounting the volumes and performing read/write tests to the volumes.

### Virtual storage

Virtual storage is a method of mapping logical storage requests to physical storage (disks). It inserts a layer of virtualization such that logical host I/O requests are mapped onto pages of storage. Each page is then mapped onto physical storage. Within each page the mapping is linear, but there is no direct relationship between adjacent logical pages and their physical storage.

A page is a range of contiguous logical block addressing (LBA) in a disk group, which is one of up to 16 RAID sets that are grouped into a pool. Thus, a virtual volume as seen by a host represents a portion of storage in a pool. Multiple virtual volumes can be created in a pool, sharing its resources. This allows for a high level of flexibility and the most efficient use of available physical resources.

Some advantages of using virtual storage are:

- It allows performance to scale as the number of disks in the pool increases.
- It virtualizes physical storage, allowing volumes to share available resources in a highly efficient way.
- It allows a volume to comprise more than 16 disks.

Virtual storage provides the foundation for data-management features such as thin provisioning, automated tiered storage, read cache, and quick rebuild.

HPE recommends using virtual storage to take advantage of the advanced virtualization features of the firmware in an MSA array.

## Disk groups

A disk group is an aggregation of disks of the same type, using a specific RAID type that is incorporated as a component of a pool, for the purpose of storing volume data. Disk groups are used in both virtual and linear storage. You can add virtual, linear, and read-cache disk groups to a pool.

All disks in a disk group must be the same type (SAS SSD, enterprise SAS, or midline SAS). A disk group can contain different disk models with different capacities and sector formats. If you mix disks with different capacities, the smallest disk determines the logical capacity of all other disks in the disk group, regardless of RAID level. For example, the capacity of a disk group with one 500 GB disk and one 750 GB disk is equivalent to a disk group with two 500 GB disks. To maximize capacity, use disks of similar size.

A **virtual disk group** requires the specification of a set of disks, RAID level, disk group type, pool target (A or B), and a name. If the virtual pool does not exist at the time of adding the disk group, the MSA system will automatically create it. Unlike linear pools, multiple disk groups (up to 16) can be added to a single virtual pool. Virtual disk groups that contain SSDs can only be created with a performance tier license. This restriction does not apply to read-cache disk groups.

---

**TIP:** For optimal performance, all virtual disk groups in the same tier within a virtual group should have the same RAID level, disk capacity, and physical number of disks.

---

When a virtual disk group is removed that contains active volume data, that volume data will drain (or be moved) to other disk group members within the pool. Disk groups should only be removed when all volume data can cleanly be drained from the disk group. Otherwise, the data will be lost. When the last disk group is removed, the pool ceases to exist and will be deleted from the system automatically.

A **linear disk group** must specify a set of disks, RAID level, disk group type, and a name. Whenever an MSA creates a linear disk group, it also creates an identically named linear pool at the same time. No more disk groups can be added to a linear pool. The legacy linear method maps logical host requests directly to physical storage. In some cases the mapping is 1-to-1, although in most cases the mapping is across groups or slices of physical storage devices. This method of mapping is highly efficient; the negative side is lack of flexibility. This makes it difficult to alter the physical layout after it is established.

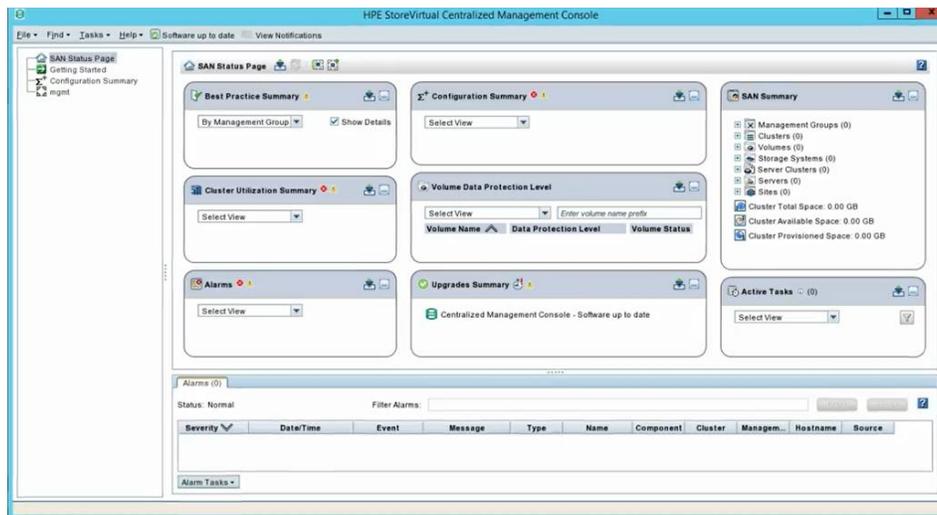
For maximum performance, all of the disks in a linear disk group must share the same classification, which is determined by disk type, size, and speed. This provides consistent performance for the data being accessed on that disk group. To dissolve a linear disk group, delete the disk group and the contained volumes are automatically deleted. The disks that compose that linear disk group are then available to be used for other purposes.

RAID types for linear disk groups created through the SMU must also be fault tolerant. The supported RAID types for linear disk groups in the interface are: RAID 1, RAID 5, RAID 6, RAID 10, and RAID 50. RAID 10 and RAID 50 are only in the interface if the system's disk configuration supports them. If RAID 10 is specified, the disk group has a minimum of two subgroups. If you select RAID 50, depending on the number of selected disks, you can create varying numbers of subgroups. Additionally, you can create fault-tolerant RAID 3 or non-fault-tolerant non-RAID or RAID 0 disk groups through the CLI.

A **read-cache disk group** is a type of a virtual disk group used to cache virtual pages to improve read performance. Read cache does not add to the overall capacity of its pool. You can add or remove it from the pool without affecting the volumes and their pool data other than to impact read-access performance.

If an MSA system uses SSDs, you can create read-cache disk groups for virtual pools if you do not have any virtual disk groups for the pool that comprise SSDs (virtual pools cannot contain both read cache and a performance tier). Only a single read-cache disk group may exist within a pool. Increasing the size of read cache within a pool requires you to remove the read-cache disk group, and then re-add a larger read-cache disk group. It is possible to have a read-cache disk group that consists of a single disk in a non-RAID configuration. The RAID type for virtual and linear disk groups must be fault tolerant.

## Configuring HPE StoreVirtual solutions



HPE StoreVirtual CMC

You can use the wizard in the StoreVirtual Centralized Management Console (CMC) to manage the storage and create a virtualized pool of storage resources.

The Configuration Interface is a CLI that uses a direct connection with the storage system to allow you to perform configuration and maintenance tasks. You might need to access the Configuration Interface if all network connections to the storage system are disabled. To connect to the Configuration Interface, use one of the following methods:

- Attach a keyboard and monitor to a USB port in the storage system (preferred).
- Attach a PC or a laptop using a null modem cable and connect using a terminal emulation program.
  - To establish a terminal emulation session on a Windows system, on the PC or laptop attached directly to the storage system, open a session with a terminal emulation program such as HyperTerminal or Procomm Plus. When the session is established, the Configuration Interface window opens. Use the following settings: 19200, 8-N-1.
  - To establish a terminal emulation session on a Linux/UNIX system, create a configuration file as root, or the root must change permissions for /dev/cua0 in order to create the config file in /etc/.

**NOTE:** For more information about using the Configuration Interface, refer to the *HPE StoreVirtual Storage User Guide*. You can access the guide by scanning the following QR code or by right-clicking it to open the hyperlink.



[http://h20628.www2.hp.com/km-ext/kmcsdirect/emr\\_na-c05049834-1.pdf](http://h20628.www2.hp.com/km-ext/kmcsdirect/emr_na-c05049834-1.pdf)

StoreVirtual supports both iSCSI and Fibre Channel storage networks. With iSCSI and optional Fibre Channel connectivity, StoreVirtual works well even in mixed-protocol environments. A single, flexible storage solution can be deployed for all locations. StoreVirtual systems enable customers to prepare for future growth, because they can migrate from one protocol to the other without having to redesign or replace storage.

## Creating iSCSI storage

To create iSCSI storage in a StoreVirtual storage system, follow these high-level steps:

1. Install storage systems on network, assign IPs, and change the host name using the Configuration Interface. Install an iSCSI initiator on the application server.
2. Create a new management group or add storage systems to an existing management group.
3. Create an iSCSI server in CMC.
4. Assign volumes to the iSCSI server.
5. Discover targets in the iSCSI Initiator.
6. Finish creating storage as necessary. For example, initialize and partition disks in Windows.

## Creating Fibre Channel storage

To create Fibre Channel storage in a StoreVirtual storage system, follow these high-level steps:

1. Install Fibre Channel-enabled storage systems on the network and Fibre Channel fabric, assign IP addresses, and change the host name using the Configuration Interface.

---

**IMPORTANT:** Do not zone storage systems until they have been put into a management group.

---

2. Ensure that the initiator HBA driver is loaded and logged in to the Fibre Channel switches.
3. Ensure that appropriate multipath I/O (MPIO) support is loaded. For example, Windows 2008 R2 or Windows 2012 uses the Microsoft MPIO feature, configured for LeftHand P4000 disks.
4. Configure storage systems by configuring 10 GbE NICs and appropriate bonding and SNMP. Change the RAID level if needed.
5. Create a new management group with two or more Fibre Channel-enabled storage systems.

---

**NOTE:** Adding two or more Fibre Channel-enabled storage systems to an existing management group creates Fibre Channel connectivity.

---

## Configuring RAID in a StoreVirtual system

Keeping multiple copies of the data ensures that data is safe and remains available in the case of disk failure. There are two ways to achieve data protection in StoreVirtual arrays:

- Configure RAID 1, RAID 10, RAID 5, RAID 5 + Spare, RAID 50, or RAID 6 within each storage system to ensure data redundancy.
- Always use Network RAID to mirror data volumes across storage systems in a cluster, regardless of RAID level, for added data protection and high availability.

Network RAID technology aggregates all resources and capacity into a single storage pool. This means that every volume in the cluster can access all available capacity. This also means all cluster resources—CPUs, RAID controllers, NICs, caches, and disk spindles—can contribute to performance.

To remove any single point of failure, the Network RAID functionality also stripes and protects multiple copies of data across a storage cluster. Redundancy is managed on a per-volume basis to improve storage utilization and match the data protection of the volume to the application's requirements.

Within each storage system, RAID 1 or RAID 10 ensures that two copies of all data exist. If one of the disks in a RAID pair goes down, data reads and writes continue on the other disk. Similarly, RAID 5, RAID 50, or RAID 6 provides redundancy by spreading parity evenly across the disks in the set.

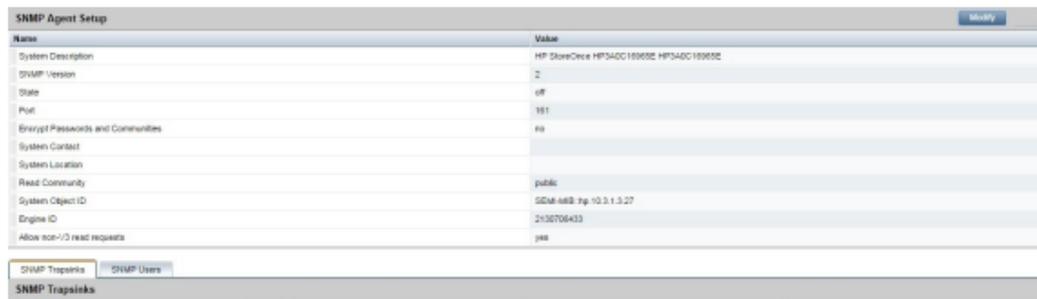
If one disk in a RAID 5 set or two disks in a RAID 6 set go down, data reads and writes continue on the remaining disks in the set. In RAID 50, up to one disk in each RAID 5 set can go down, and data reads and writes continue on the remaining disks.

---

**NOTE:** If you plan on using clusters with only a single storage system, use RAID 1 and RAID 10, RAID 5, or RAID 6 to ensure data redundancy within that storage system.

---

## Configuring StoreOnce systems



Using the StoreOnce GUI to set up the SNMP agent

All StoreOnce systems are supplied with a label that includes the iLO network name and iLO password. iLO is not required for daily management of the StoreOnce appliance but is useful in a lights-out data center situation. iLO is also useful for diagnosing hardware failures that prevent access to the appliance through the primary StoreOnce GUI or remote StoreOnce CLI interface.

To set up a basic network configuration, run the basic network script `net set address` to configure a static address for LAN port 1 (eth0). Make a note of the IP address details, which you will need to log in to the StoreOnce system from a web browser. Then use the StoreOnce CLI to configure trap sinks and users. Use the StoreOnce GUI to set up the SNMP agent and view details of traps, users, and events. To log on to the StoreOnce management GUI, use any machine connected to the same network as the appliance. The StoreOnce system uses a secure network connection.

---

**IMPORTANT:** The web browser used to communicate with the StoreOnce system requires Active Scripting or JavaScript enabled or some browser buttons will not display.

---

StoreOnce CLI commands require a Secure Shell (SSH) client application (freely available from the internet) and must be run from a SSH terminal session on a machine that is on the same network as the StoreOnce appliance: `ssh <username>@<ip_address>`

---

**NOTE:** StoreOnce CLI commands can also be run from a local console (keyboard and monitor) attached to the appliance if the network is not yet configured.

---

Three steps to setting up the network management application for use with the StoreOnce system are:

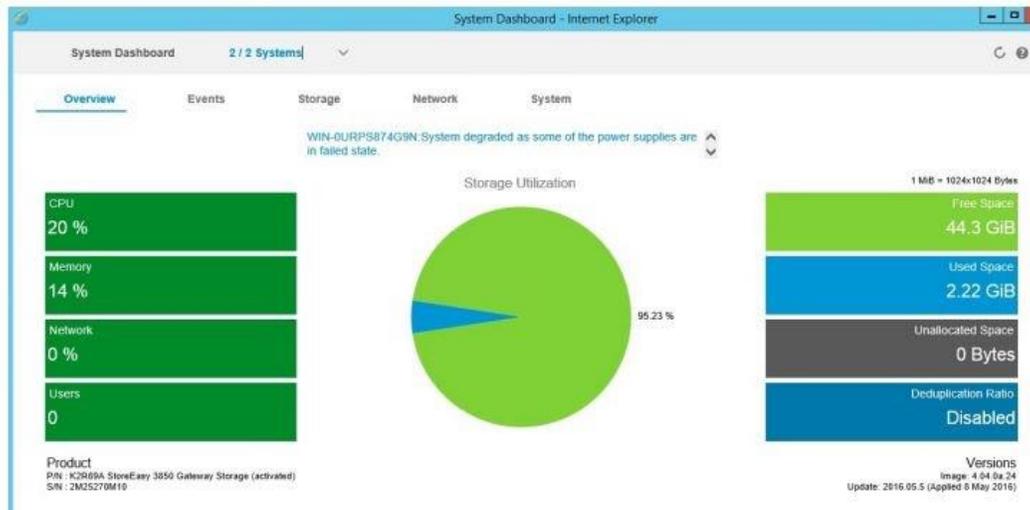
1. If this is the first StoreOnce system you are adding to the network, download the Management Information Bases (MIBs) that are used to represent the information handled by the StoreOnce system. These MIBs are SEMI-MIB.mib and TRAP-DESTINATIONS-MIB.mib. The MIBs must be registered with the network management application on the host, so information from the StoreOnce system can be transmitted to and interpreted by the network management host. To download the MIBs, go to the HPE Support website at <http://www.hpe.com/support> and search for the appropriate StoreOnce system. Select the **Software and driver downloads** option and **OS Independent** for the operating system. The MIBs are included under the MIB entry.
2. Update the existing system type database with the new system type for the StoreOnce system based on the unique system object identifier.
3. Discover the StoreOnce system under the network management application.

---

**NOTE:** When you connect subsequent StoreOnce systems to the network, you only need to update the MIBs if they are later models.

---

## Configuring HPE StoreEasy solutions



StoreEasy dashboard

StoreEasy storage ships with preconfigured hardware with the latest-generation storage operating system for purpose-built file storage. StoreEasy installation and network configuration wizards simplify installation and prevent configuration errors for faster deployment.

The StoreEasy dashboard consolidates all key status information such as performance, capacity utilization, and health in a single interface. It reduces the need to navigate multiple tools and interfaces for system management. Low-capacity warning email alerts are 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 balanced configuration based on available storage. All StoreEasy systems can act as a NAS device to provide storage for other clients. They also all support the same protocols. However, where that storage originates depends on the StoreEasy model.

StoreEasy can take direct-attached disks (internal and external) as well as storage from SANs and make it available to client systems as either file shares or iSCSI targets. You can use the same set of tools for working with the folders and LUNs. A gateway allows you to present LUNs from the SAN or array to StoreEasy, and you can then offer LUNs as Server Message Block (SMB) and Network File System (NFS) file shares to clients. This also enables other Windows-based features such as snapshots (Volume Shadow copies), data deduplication, and quotas.

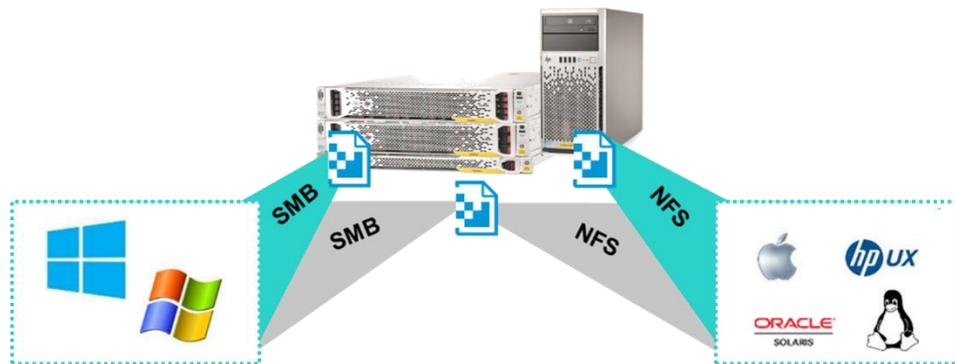
All StoreEasy models ship with Windows Storage Server 2012 R2 Standard Edition operating system preinstalled with Microsoft iSCSI Software Target included. HDD bundles make it easy for customers to buy, deploy, and provision additional storage capacity over the life of the product.

StoreEasy systems support both single domain and dual domain cabling using HPE D2000, D3000, D6020, and D6000 disk enclosures, depending on the Smart Array RAID controller installed in the system. In a single domain, external disk enclosures are connected using a single data path to the Smart Array controller in the system. In a dual domain, external disk enclosures are connected using two data paths to a single Smart Array RAID controller in the system. StoreEasy systems connect to the I/O modules on the D6000/D6020 and D2000/D3000 disk enclosures.



**CAUTION:** When cabling a dual domain configuration, do not connect multiple controllers in the system to the same disk enclosures. This is an unsupported configuration and might result in data loss. If the StoreEasy system is configured incorrectly and you attempt to use Pool Manager, Pool Manager will detect the configuration and display an error message. The controllers can be connected to different disk enclosures.

## Heterogeneous file sharing



Continuously available file sharing across multiple client platforms

StoreEasy products support heterogeneous file sharing through multiprotocol access for:

- Windows clients
  - A failover cluster running Windows Server 2012 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 must be installed on all cluster nodes.
  - Clustered file server must be configured with one or more file shares created with the continuously available property.
  - Server Message Block client computers running Windows 8, Windows Server 2012, or later. 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 3.0 provides active/active file sharing by allowing access to the same folders from multiple systems concurrently. This is achieved through Cluster Shared Volumes (CSV).
- Non-Windows clients—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.

---

**NOTE:** Refer to the following white papers at [HPE.com](http://HPE.com) for more information:

- *Deploying HPE StoreEasy Storage in a UNIX environment*
  - *Best practices for provisioning storage using SMB, NFS, and iSCSI protocols on HPE StoreEasy Storage*
  - Heterogeneous file serving on HPE StoreEasy Storage using SMB and NFS protocols
  - Heterogeneous file sharing with HPE StoreVirtual 3000 File Controller
-

## Learning check

1. HPE MSA arrays support logical drives that contain a mix of hard drives and solid-state drives.

True

False

2. A StoreEasy system can act as a NAS to provide storage for other clients.

True

False

3. What are the benefits of designing a SAN using Fibre Channel switches?

---

---

---

---

4. What must you do to access, install, and launch SSA in the online environment?

---

---

---

## Configuring the network

There are several best practices to keep in mind when configuring a data center network:

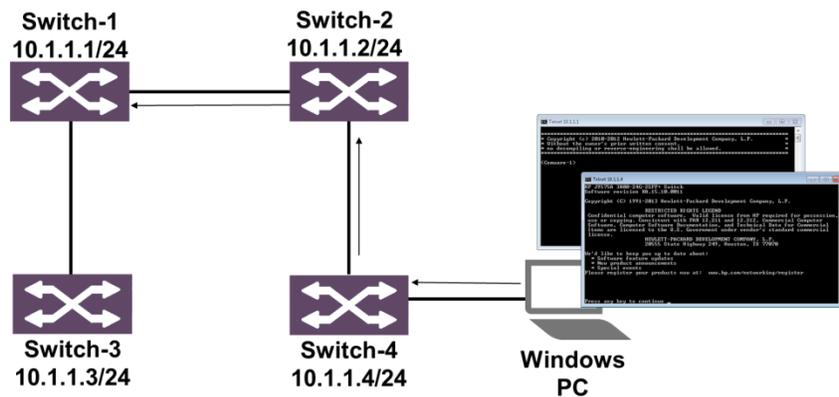
- Use static IP addresses or reserved addresses if using DHCP.
- Isolate management traffic from production traffic for better security. In this way, if the network becomes congested, you can still troubleshoot the problem because you can log in to the devices involved.
- Isolate the SAN on a separate network. If the SAN must run on a public network, use a virtual private network (VPN) to secure data.
- Create a management VLAN for activities such as device monitoring, system logging, and SNMP that is connected to all devices on the network.
- Configure all the network characteristics on a storage system before creating a management group, or before adding the storage system to a management group and cluster.
- Configure storage system settings for speed and duplex, frame size, and flow control before bonding NICs and before putting the storage system into a management group and cluster.

## Configuring in-band and out-of-band system management

A best practice for managing network devices is to use both in-band and out-of-band system management. In-band management enables you to manage devices through gigabit Ethernet ports. SNMP, the standard internet protocol, uses in-band management, which means that SNMP packets are sent over the same physical devices as data packets. However, if an in-band network is down, it becomes unusable because the network connection is lost.

Out-of-band management usually remains available even if the network is down. A common out-of-band configuration has an access server connected to the management port of each controlled device. The access server typically has a public IP address with access list controls applied to allow only certain source IP addresses. In case the network is down, you can use the access switch to manage the network remotely. For this reason, it is important to configure the access switch and port names correctly. You should also occasionally verify access to the server and make sure all links to all managed devices are operating correctly.

## In-band management



In-band access configuration

In-band access allows multiple users to access the switch through the IP network. This access uses virtual interfaces VTY0, VTY1, VTY2, and so on. At a switch's default settings, you are required to enter a password for these interfaces. To eliminate a potential security weakness, the switch does not have a default password. You must configure a unique password for your particular company.

With in-band (networked) management, management communications run over network connections. You need IP connectivity to the networking device through a direct or indirect Ethernet connection. To open an in-band management session to access the network device's CLI, you must use terminal emulation software such as PuTTY or Tera Term. This software communicates with the switch using a network protocol:

- **Telnet**—Carries the terminal session in plain text.
- **Secure Shell**—Establishes a secure tunnel for the session using authentication and encryption. SSH is the more secure option over telnet in a real-world environment.

To manage a switch through in-band access, you must configure the switch with an IP address and subnet mask that is reachable on the network.

## Out-of-band management

### Direct serial connection to console port



### Does not use or require an Ethernet connection

#### Out-of-band access configuration

When you initially configure a switch, you typically use out-of-band management. For out-of-band management, you must have physical access to the switch. You connect the management station to the switch's console port by using the serial cable that ships with the switch. This connection is dedicated to the management session. Out-of-band management does not require the switch to have network connectivity.

Out-of-band connections use an auxiliary (aux) port for console access and require no password at default settings, enabling initial access to the switch. You can leave this default authentication method (none) for out-of-band management, or you can configure the AUX0 interface to require users to log in with a password or with a user name and password.

If a user name and password are required (an authentication method called *scheme authentication*), the switch checks the credentials against a local list of users or an external authentication server, as dictated by its Authentication, Authorization, and Accounting (AAA) domain settings.

To open a management session with the switch and access the CLI, you use terminal software such as PuTTY, which is available as freeware, or Microsoft HyperTerminal, which is available free from Microsoft.

The switches use default settings for the terminal emulation software:

- 9600 bps
- Eight data bits
- No parity
- One stop bit
- No flow control

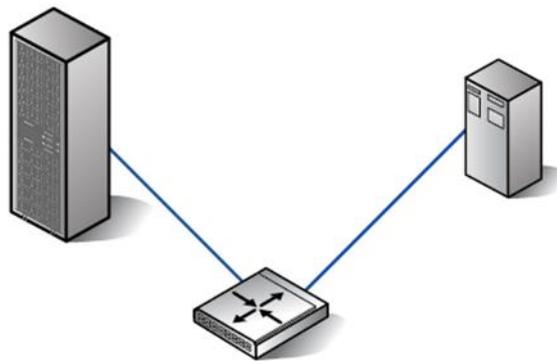
## Fabric designs

SANs provide the data communication infrastructure for advanced, cost-efficient storage systems. A typical SAN design consists of devices on the edge of the network, switches in the core of the network, and the cabling that connects it all together.

A fabric is a single switch or a set of switches connected to form a network. Fabric services manage device names and addresses, time stamps, and other functionality for the switches. A SAN fabric topology defines the arrangement of Fibre Channel switches in a fabric.

HPE recommends a tiered, core-edge topology to optimize performance, management, and scalability. This topology provides good performance without unnecessary interconnections. A tiered topology has several edge switches used for device connectivity, and a smaller number of core switches used for routing traffic between the edge switches.

### Single-switch fabric



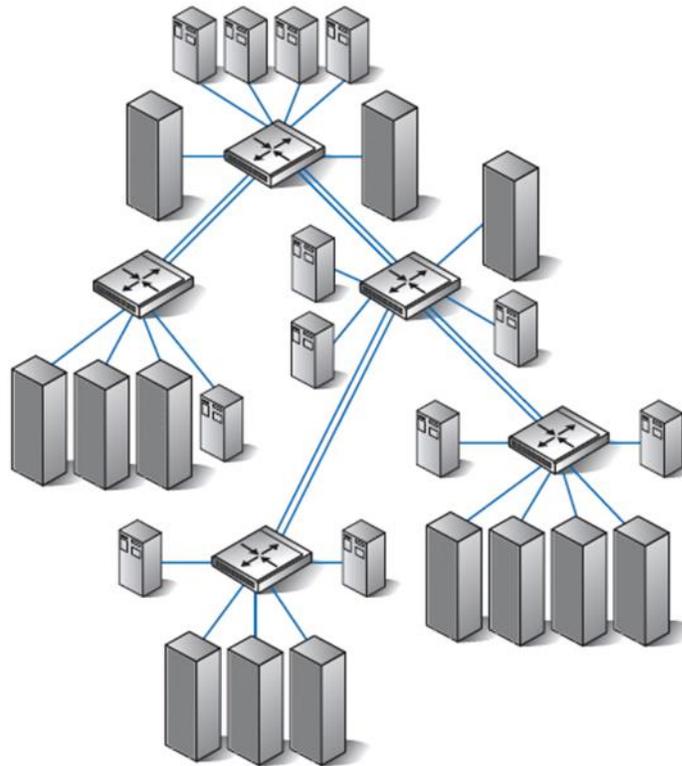
Single-switch fabric design

A single-switch fabric consists of a Fibre Channel switch, server, and storage system. This topology forms the basis for all HPE standard topologies. For example, you can connect two single-switch fabrics to create a cascaded fabric. Or you can connect three or more single-switch fabrics to create a ring fabric or a core-edge fabric.

The benefits of a single-switch fabric include:

- Easy installation and configuration of servers and storage
- Maximum fabric performance because all communicating devices connect to the same switch
- Support for local, centralized, and distributed data access needs

## Cascaded fabric



Cascaded fabric design

A cascaded fabric is a set of interconnected switches, arranged in a tree format, with one or more interswitch links (ISLs). An ISL is the link joining two Fibre Channel switches through E\_ports. You can connect one switch to one or more switches using a single ISL to each. You can also connect a pair of ISLs between two switches. Having a minimum of two ISL connections on each switch helps provide fabric path redundancy. You should consider using a cascaded fabric topology if the customer requires multiple groups of devices with localized intraswitch access.

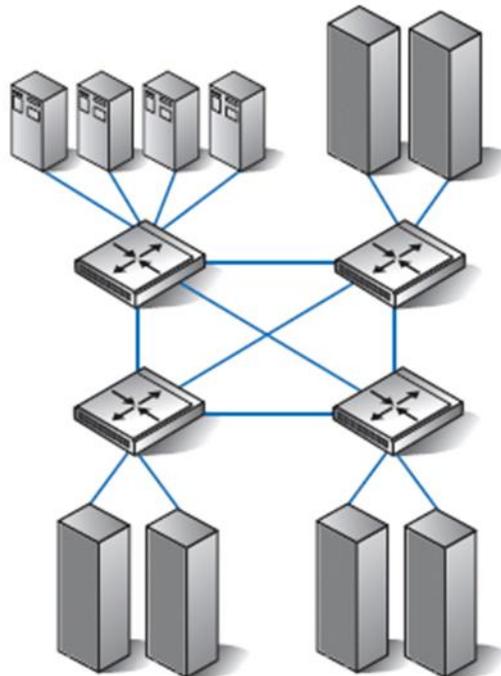
Cascading enables you to:

- Optimize I/O activity by connecting servers and storage to the same switch in the cascaded fabric
- Easily scale the fabric over time by adding cascaded switches

The benefits of a cascaded fabric include:

- Ability to connect SANs in diverse geographic locations
- Ease of scalability for increased server and storage connectivity
- Shared backup and management support
- Optimum local performance when communicating devices are connected to the same switch in the cascaded fabric
- Cost efficiency because of the large number of switch ports available
- Support for local data access and occasional centralized data access

## Meshed fabric



Meshed fabric design

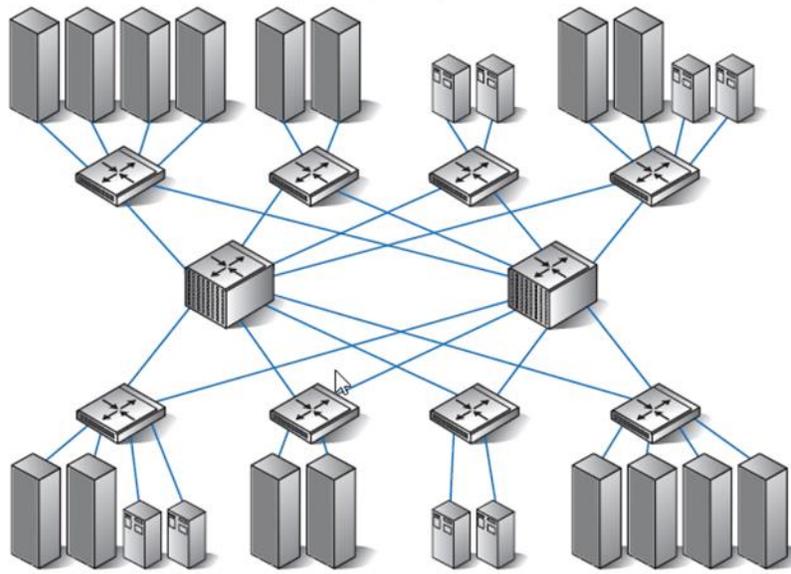
A meshed fabric is a group of interconnected switches using multiple ISLs for fabric resiliency. If one ISL fails, the switch automatically reroutes data through an alternate path in the fabric. If the alternate path includes other switches, the data must pass through those switches to reach the destination.

As you add switches, ISLs automatically connect to two or more adjacent switches to maintain mesh connectivity, ensuring path redundancy throughout the fabric. The additional ISL connectivity provides devices with more paths to communicate through the fabric. This dramatically reduces the chance that, as you add switches, you will exceed the maximum hop count.

The benefits of a meshed fabric include:

- The ability to meet multiple data access needs
- Multiple paths for internal fabric resiliency
- Ease of scalability
- Shared backup and management support
- Support for a mix of local and distributed data access
- Less impact on performance because of intraswitch traffic

## Core-edge fabric



Core-edge fabric design

HPE recommends using a core-edge fabric wherever possible. A core-edge fabric has one or more Fibre Channel switches (called *core switches*) that connect to edge switches in the fabric. The core switches provide high bandwidth and redundant connectivity to the edge switches. The edge switches provide user ports for servers and storage. You can also connect centralized storage (both disk and tape) to the core switches if centralized access is required.

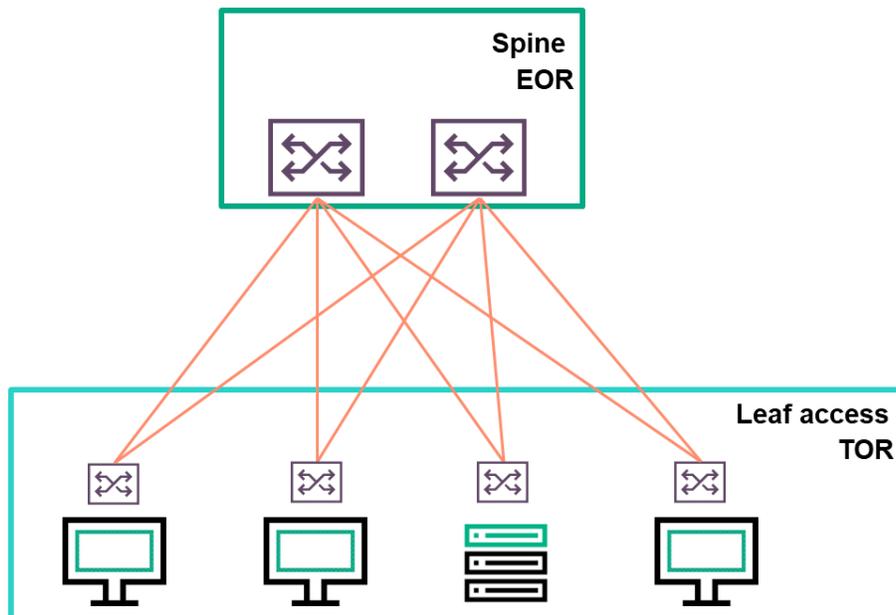
The core-edge fabric is optimal for:

- Many-to-many connectivity environments that require high performance
- Unknown or changing I/O traffic patterns
- SAN-wide storage pooling

The benefits of a core-edge fabric include:

- A maximum of two hops between switches (typically)
- Equal, centralized access to devices in the core
- Increased fabric and switch redundancy with two or more switches in the core
- Full, many-to-many connectivity with evenly distributed bandwidth
- Support for centralized and distributed data access needs
- The ability to designate an optimally located core switch as the primary management switch, with direct connections to all switches

## Tiered network configurations



A two-tier network configuration includes a leaf access layer and a spine layer

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 LAN access or server access switches and connect to the core switches.
- **LAN or server access** switches support workstations and servers.

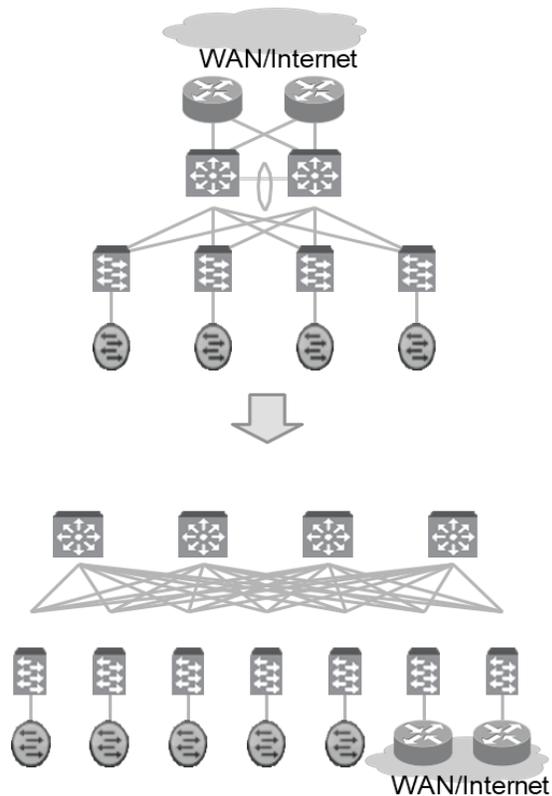
A traditional three-tier architecture includes the access layer, distribution layer, and core layer. Designed for network environments that were common 10 years ago, this architecture no longer meets the needs of many companies. Because data centers are highly virtualized, traffic patterns have changed. 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 addresses those limitations. A two-tier network consists of the leaf access top-of-rack (TOR) layer, and the spine end-of-row (EOR) 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



Layer 2 networks were widely used in the early days of networking when networks were small. As networks grew, they moved towards 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 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 *leaf-spine*. In a leaf-spine architecture, 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 TCP in-cast traffic patterns.

Leaf-spine architectures 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.

## HPE Intelligent Resilient Framework



IRF provides a simple, resilient, and versatile pay-as-you-grow business model

HPE Intelligent Resilient Framework (IRF) 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 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, 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.

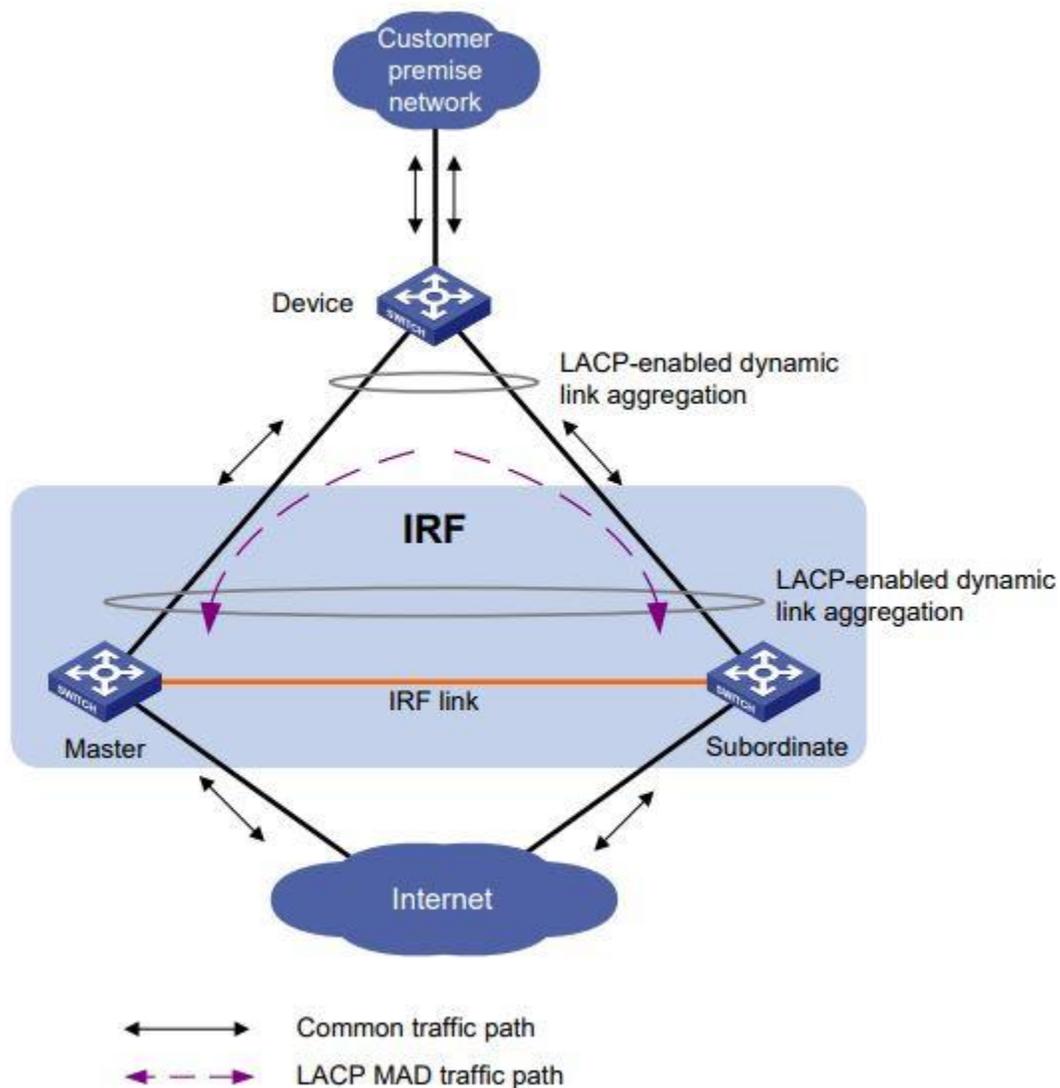
IRF delivers a network-based in-service-software-upgrade capability that allows an individual IRF-enabled switch to be taken offline for servicing or software upgrades without affecting traffic going to other switches in the IRF domain.

## IRF and meshed stacking

In addition to IRF, HPE also 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 and processing capabilities 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 both the campus LAN and data center.

## Forming link aggregation connections



LACP multi-active detection (MAD) application scenario

IRF and LACP used together 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 support only static link aggregation (no LACP). LACP support was added in 5.1.

---

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.

## LACP operational modes

LACP manages the dynamic link negotiation of an aggregated connection: Switches use LACP data units (LACPDUs) to exchange information and establish a dynamic aggregated link. These LACPDUs include information such as the source 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 LACPDUs 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.

---

**IMPORTANT:** Do not confuse static LACP with a manual link aggregation. The first uses LACP; the second does not use a protocol.

---

## Link aggregation load sharing

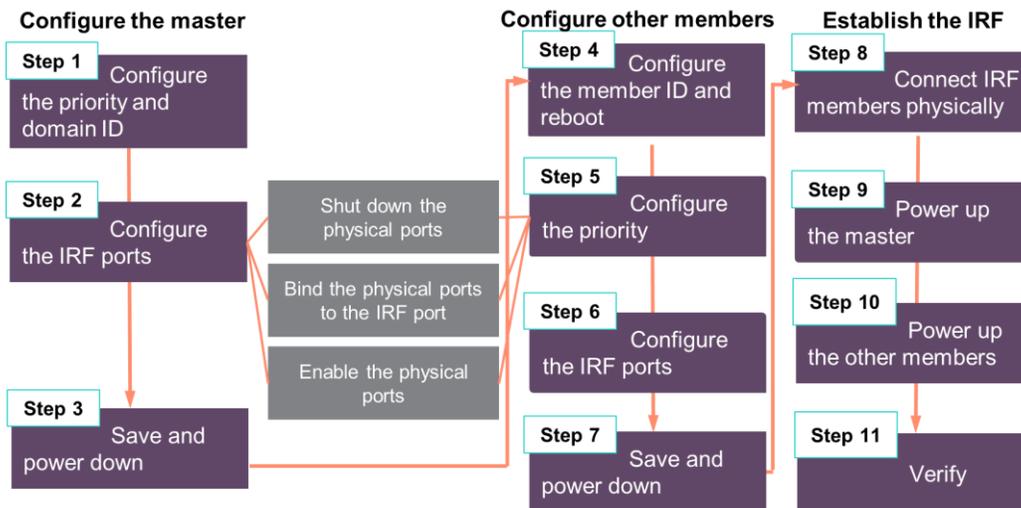
HPE FlexFabric switches provide options for load sharing, such as MAC and IP addresses and TCP or User Datagram Protocol (UDP) ports. By default, these switches use the Layer 3 option for load balancing. This option uses a hash of the source and destination IP address to select the physical interface 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 together. 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



Eleven steps are involved in configuring IRFs

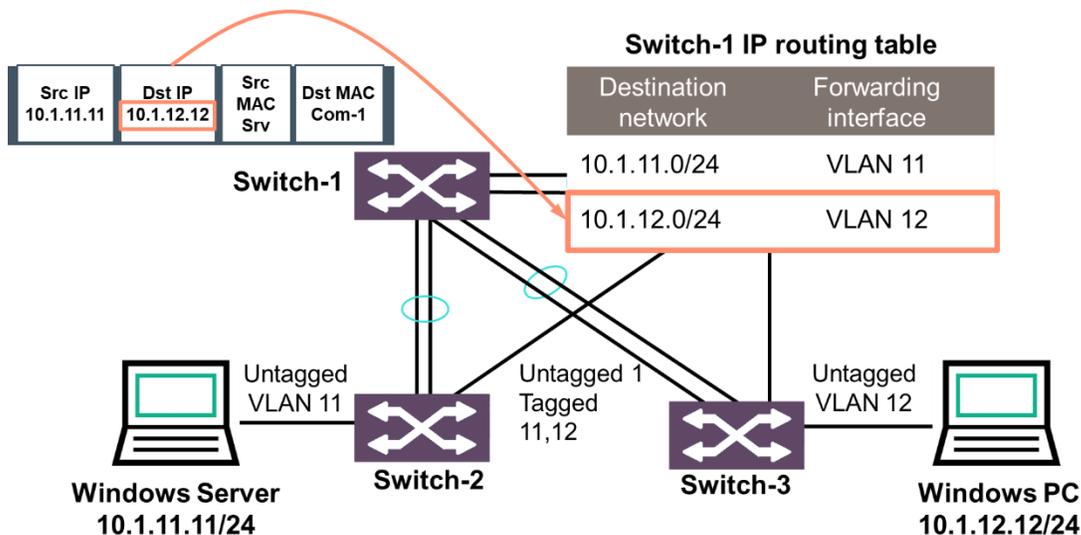
As shown in the preceding graphic, IRF virtual devices are configured in a high-level three-stage process:

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.

## IP routing



Example of overall traffic flow between three switches

Switches acting at Layer 2 make switching decisions based on destination MAC addresses. However, when a Layer 3 routing switch receives IP traffic that is destined to its own MAC address, it looks at the destination IP address inside the Ethernet frame. If the destination IP address does not belong to it, the routing switch knows that it needs to route the traffic.

The routing switch uses the destination IP address to decide how to forward the traffic. It refers to its IP routing table and finds the route with the most specific match to the packet's destination IP address—in other words, the route with a matching destination network that has the longest subnet mask.

For example, the server needs to send a packet to the client. It directs the frame to the Comware-1 MAC address because Comware-1 is its default gateway. Comware-1 receives a frame and decapsulates it. It matches the destination IP address, 10.1.12.12, to IP route 10.1.12.0/24 in its routing table.

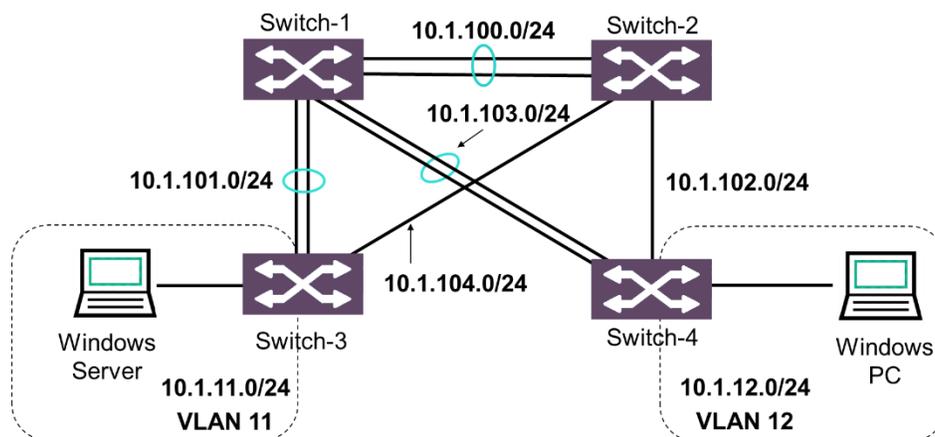
The Switch-1 routing table includes direct routes, which are sometimes called *connected routes*. Direct routes are for local networks on which the routing device has an IP address itself. They are associated with a Layer 3 forwarding interface such as a VLAN. The direct route is automatically added to the routing table when you configure the IP address on the device.

An indirect route is a route to a remote network that does not exist on the routing device. You must configure this route on the device manually (as a static route), or the device must learn it dynamically with a routing protocol.

As well as a destination network address, an indirect route also includes the next hop—the next device that will forward the packet to its final destination. This next hop is sometimes called the *gateway*. The next hop IP address should be on a network connected to the routing switch. The routing switch automatically determines the forwarding interface based on which interface connects to the next hop.

A default route is a route to 0.0.0.0/0. It is a special type of indirect route that applies to all IP traffic for which the device does not have a more specific route. For example, Switch-1 might have a default route to an external router. It uses its direct routes for all the local traffic because these routes are more specific. It only uses the default route when it must route traffic for which it does not know another route.

## Topology that requires indirect routes

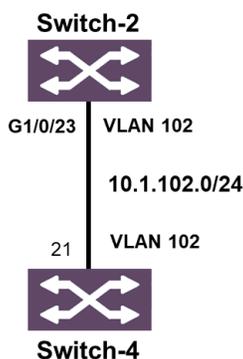


Links between switches correspond to different IP subnets

In some topologies, direct routes with a few default routes work well. All the VLANs in the network use the same default router, which routes traffic between its connected networks. There is a different logical topology that introduces a Layer 3 hop between each switch. In this topology, each switch must act as a routing switch, and the switch requires indirect routes. For example, Switch-3 is now routing traffic for VLAN 11, subnet 10.1.11.0/24. Subnet 10.1.12.0/24 is not local on this switch, so Switch-3 requires an indirect (remote) route to it.

As shown in the graphic, each switch-to-switch link is associated with a different IP subnet. You have two options for making this association:

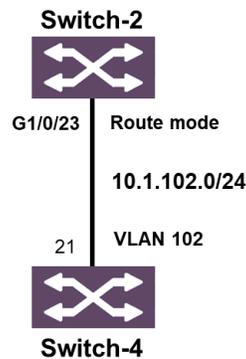
- **Dedicated VLAN interface**—You can create a VLAN that is specific to the switch-to-switch link. For example, VLAN 102 for the link between Switch-4 and Switch-2. On both sides of the link, you assign only this VLAN to the switch port. You do not assign any other ports to this VLAN. You assign an IP address to this VLAN, making it a Layer 3 interface.



Using dedicated VLANs

- **Route mode interface**—You can change the physical interface to operate in route mode. The physical interface itself becomes the Layer 3 interface, and you assign the IP address directly to it. For example, you could configure the Switch-2 G1/0/23 port in route mode and assign IP address 10.1.102.2/24 to it.

A route mode interface acts like an interface on a router. It does not switch traffic; instead, it routes traffic to other Layer 3 interfaces. It does not transmit or process any Rapid Spanning Tree Protocol (RSTP)/MSTP Bridge Protocol Data Units (BPDUs), and it terminates any spanning tree. It has a unique association with the subnet configured on it. Just as you cannot configure an IP address in the same subnet on more than one VLAN interface on a switch, you cannot configure an IP address in the same subnet on more than one route mode interface.



Using a route mode interface

No matter how you set up the network, you might choose to use a topology like this for various reasons:

- You want to eliminate RSTP/MSTP on the switch-to-switch links. In this new topology, VLANs assigned to edge ports are terminated at each switch and not carried on the switch-to-switch links. Therefore, Layer 2 loops cannot be introduced. You have seen how MSTP provides some load sharing of traffic. But using routing to determine the traffic's path instead can result in more optimal paths.

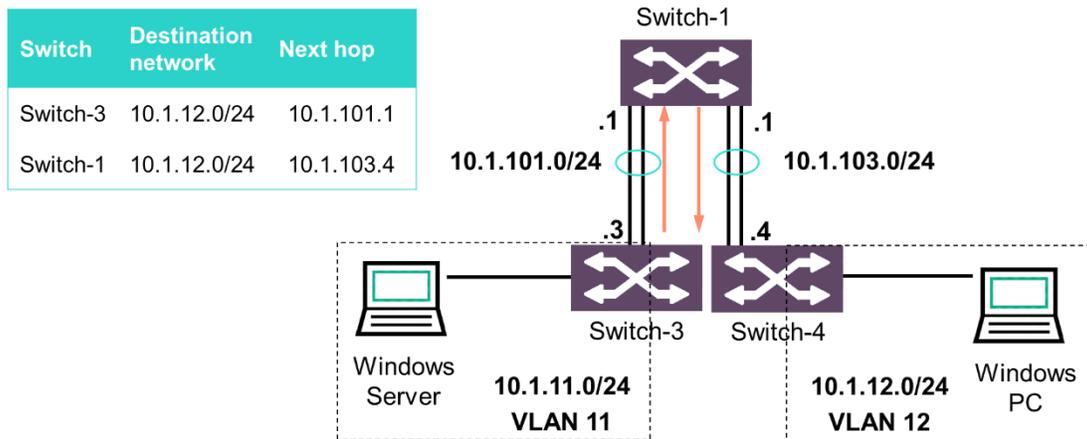
---

**NOTE:** This topology eliminates the need for spanning tree on switch-to-switch links, but you would still run spanning tree on the switches to prevent accidental loops downstream.

---

- You want to configure the default router role on different switches. In this topology, Switch-3 is the default router for devices connected to it, and Switch-4 is the default router for devices connected to it. The burden is distributed, and the negative effect of one switch failing is minimized.

## Setting up indirect IP routes



Create specific indirect routes instead of direct or default routes

The graphic shows a simplified network with just Switch-1, Switch-3, and Switch-4 (no redundant connections through Switch-2). To set up routes in this topology, you need to create specific indirect routes (as opposed to just direct and default routes). For example, you might create a route to 10.1.12.0/24 on Switch-3 so that this switch can route server traffic to clients.

In addition to the destination network address, the route specifies the next hop. It specifies *only* the next hop; as far as the local switch's routing table is concerned, the rest of the path does not exist.

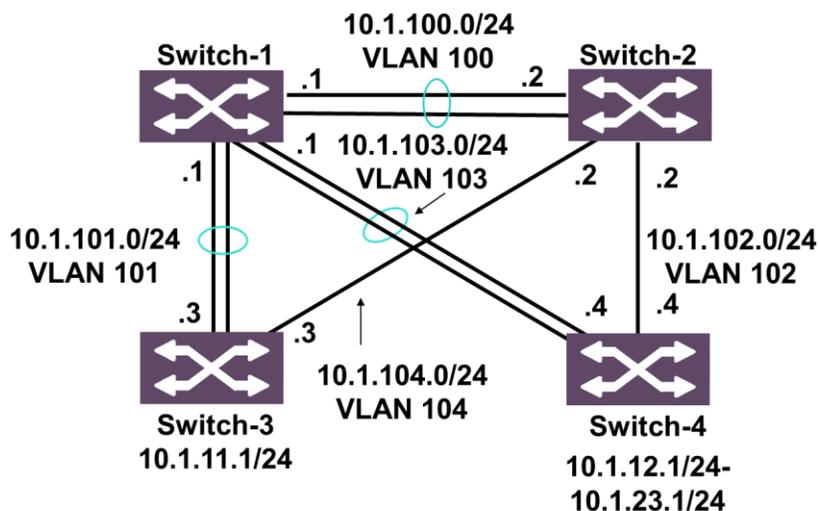
The next hop for Switch-3's route to 10.1.12.0/24 is Switch-1. Switch-1 has several IP addresses. For the next hop address, you should specify the IP address on the subnet that connects the two switches: 10.1.101.1. If you specified a different IP address on Switch-1, Switch-3 would not know how to reach that address. A route with an unreachable next hop is not added to the active IP routing table, and the switch cannot use it.

To finish the path to 10.1.12.0/24, you must also create a route to 10.1.12.0/24 on Switch-1. Again you specify the next hop in the path. This next hop is Switch-4 and, specifically, IP address 10.1.103.1 on Switch-4.

Now Switch-3 can route traffic to 10.1.12.0/24 successfully. You would need to create a return route for traffic to enable full connectivity. For example, you want to set up connectivity between servers in subnet 10.1.11.0/24 and clients in subnet 10.1.12.0/24. You have already configured a route to 10.1.12.0/24 on Switch-3 and Switch-1. You would need to configure these routes for the return traffic:

- Switch-4: A route to 10.1.11.0/24 through next hop 10.1.103.1
- Switch-1: A route to 10.1.11.0/24 through next hop 10.1.104.1

## Static IP routing



When configuring routes manually, specify next hops and create all necessary paths

Switches can learn indirect routes statically or dynamically through a routing protocol. Static routes are configured manually by the administrator. When you use static routing, you must configure the routes manually on each router or routing switch in the path. You must take care to specify the next hops correctly and create all the necessary paths.

For example, the routes shown in the graphic establish connectivity between 10.1.11.0/24 and 10.1.12.0/24. However, Switch-3 still cannot route traffic to the 10.1.102.0/24 subnet on Switch-4. You might not want to route traffic to a subnet devoted to a switch-to-switch link.

However, Switch-4 could support several other subnets with devices in them. For example, it might support 10.1.13.0/24, 10.1.14.0/24, and so on up to 10.1.23.1/24. You would need to create a route for each on each routing device in the pathway. (Or, if subnets are contiguous, you could create a route to a larger destination subnet that includes the other subnets within it.) This process can be laborious.

## Accessing and configuring HPE Comware switches

To access an interface context on Comware switches, use the following command:

```
interface <type_slot_#>/<sub_slot>/<port_#>
```

<Type> refers to the data link layer type, such as GigabitEthernet or Ten-GigabitEthernet. When you enter commands, you can use a shortcut for interface types. For example, GigabitEthernet can be represented as *g*, and Ten-GigabitEthernet can be represented as *ten*.

Replace <slot\_#> with the slot in which the module or card is installed. For both modular and fixed port switches, <slot\_#> begins at 1 on the switch and <sub\_slot> begins at 0. The ports are numbered from 1 to the maximum number of ports supported.

For example, *interface G1/0/1* indicates it is a gigabit Ethernet port, and it is the first port on module 1.

On Comware switches running the correct software version, you can create an interface range.

---

**NOTE:** The software versions that support the `interface range` command depend on the switch model. Refer to the documentation and release notes for the switches.

---

When you define the interface range, you specify the range of interfaces that belong to it, and you can also give the interface range a name. For example, you can enter:

```
[Comware] interface range name EdgePorts g1/0/1 g1/0/3 to g1/0/6
```

You can then configure settings on all of the interfaces at once. The interface range persists in the device configuration. You can access it again by specifying the range name. For example:

```
[Comware] interface range name EdgePorts
```

You can then configure additional settings on the interfaces that you previously defined in this range.

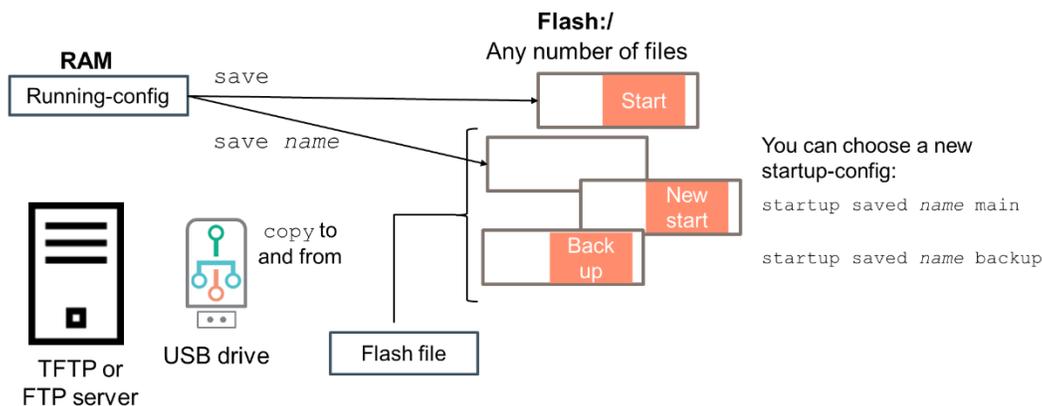
Comware switches support another way of defining settings on multiple physical interfaces at once: manual port groups. Manual port groups are available in older versions of software. You define a manual port group by name and add interfaces to the group. For example:

```
[Comware] port-group manual EdgePorts
```

```
[Comware-1-port-group-manual-EdgePorts]group-member g1/0/1 g1/0/3 to g1/0/6
```

You can then configure settings on all of the interfaces at one time. The port group does not persist in the configuration. If you want to configure settings on a range of interfaces again, you must create the port group again.

## Managing configuration files



Copying the running-config file to a startup-config file

HPE Comware devices have a concept of running-config and startup-config. However, you have more freedom in storing multiple configuration files. Comware devices have an open flash directory that stores both software images and configuration files as well as other types of files. You can store as many configuration files to the flash as the memory allows. Configuration files always require the `.cfg` extension.

The `save` command copies the running-config to a config file. When you enter the command without a filename, you are prompted to enter a filename or retain the existing startup-config filename. In either case, this file will contain your saved configurations and will become the main startup-config. The switch executes this file at bootup.

When you enter the `save` command with a filename (for example, `save myconfig.cfg`), the running-config is saved as a named file in flash. The file does not become the startup-config.

You can also copy config files to and from a number of locations:

- A TFTP server or FTP server
- A USB device
- Another file in flash

The graphic shows the options available with the `copy` command. You can alternatively use the switch's web browser interface to copy files to and from your local management station.

You can choose a new startup-config from any of these config files using the `startup saved-configuration` command. The command on its own or with the `main` option sets the specified config file as the new startup-config. Any file that used to be the main startup-config is retained in flash, but it is no longer the startup-config; a switch can only have one main startup-config.

You can optionally use the `backup` option with this command to select one file as the backup startup-config. The switch executes this file if the main startup-config cannot successfully boot for some reason.

## Basic CLI command levels and functions



Specific functions are categorized into CLI command levels

After accessing the CLI, you can enter commands. Each CLI command is associated with one of four command levels. The command level for each command is configurable, but most customers leave the commands at the default settings.

The command levels and basic functions associated with each are as follows:

- **Level 0–Visitor** includes commands for network diagnosis and commands for accessing an external device. Commands at this level include `ping`, `tracert`, `telnet`, `ssh2`, and certain display commands. These are read-only type commands, which cannot be saved and do not survive a device restart. Upon device restart, the commands at this level will be restored to the default settings.
- **Level 1–Monitor** includes commands for system maintenance and service fault diagnosis. Commands at this level include `debugging`, `terminal`, `refresh`, `reset`, `send`, and certain display commands. Again, command configurations cannot be saved and do not survive a device restart.
- **Level 2–System** includes service configuration commands, including interface, VLAN, and routing configuration commands, as well as commands for configuring services at different network levels. In fact, this level permits all configuration commands except those at the manage level.
- **Level 3–Manage** includes commands that influence the basic operation of the system and configure system support modules. These commands include ones for managing the file system; configuring FTP, TFTP, and Xmodem download; managing users and privileges; and configuring system-local parameter settings (settings that are not defined by any protocols or RFCs).

A user at a certain privilege level can enter any commands at that level or a level below. For example, a user with privilege Level 3 can enter all of the `manage`, `system`, `monitor`, and `visitor` commands.

You can change which commands are associated with various privilege levels. The command for making those changes is a Level 3 command by default.

## Learning check

1. What is a best practice in system management when managing network devices?
  - a. Use in-band management only.
  - b. Use out-of-band management only.
  - c. Use both in-band and out-of-band management.
  - d. Use in-band management when the out-of-band network is down.
2. Why is IRF an important technology? (Select two.)
  - a. It extends network control over multiple active switches.
  - b. It configures a fully homogeneous and balanced memory configuration.
  - c. It automatically provides an array diagnostic report and SmartSSD Wear Gauge Report.
  - d. It configures all the network characteristics on a storage system before creating a management group.
  - e. It simplifies networks by allowing the user to combine multiple switches into a single, ultra-resilient virtual device.

## Validating that an installation and configuration are successful

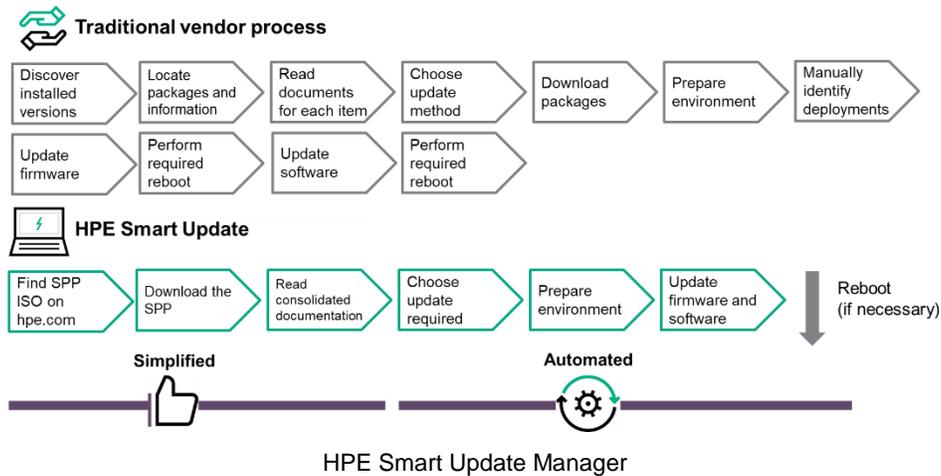
After you have designed and installed a solution for your customer, the next important step is validating that both the installation and the configuration you installed were successfully completed. Check the following components to confirm that the installation is complete:

- **Redundant power**—HPE racks use redundant PDUs and power supplies to provide power redundancy throughout the system. Before powering on the system, verify that it has a valid redundant power configuration connected to independent electrical power providers. Each of the PDUs must connect to a power socket with power being separately supplied by an independent electrical power source.
- **Verifying power connections**—In order to complete the installation, you must verify the power connection and final positioning and power on the system. Verify the following components before powering on:
  - Power cord connections
  - AC cord connections
  - Main power cords
- **Power cord connections**—Verify that the power cord is plugged into a grounded (earthed) electrical outlet that is easily accessible at all times. Do not route the power cord where it can be walked on or pinched by items placed against it. Pay particular attention to the plug, electrical outlet, and the point where the cord extends from the system.

Each enclosure in an MSA storage system has two power cooling modules (PCMs), one on the left side and one on the right side of the enclosure. When routing PCM power cords, ensure power redundancy is maintained by connecting each PCM within a shelf to a different PDU. When observed from the rear, the left side cabling (ID #0) is black and the right side cabling (ID #1) is gray. Each PCM on an enclosure must be connected to a different PDU for power redundancy. Use the black power cable on the left PCM and the gray power cable on the right PCM to help differentiate the PDU power source each is attached to. Before powering on the storage system, verify that:

- The AC cords are correctly configured and properly connected to the component
- The AC cord locks and cable ties on the PDUs and PCMs are properly connected
- If there is a physical service processor, it is properly connected to the PDU
- **Positioning the server and storage systems**—Install the doors of the server or storage system and filler panels that were not completed during installation, if needed. If the user has not positioned the system in a permanent operating location, place it in its final position before powering on.
- **Switches**—After you complete a switch installation, verify that:
  - There is enough space for heat dissipation around the switch, and the rack or workbench is stable.
  - The grounding cable is securely connected.
  - The correct power source is used.
  - The power cords are properly connected.
- **Interface cables**—Verify that all the interface cables are cabled indoors. If any cable is routed outdoors, verify that the socket strip with lightning protection and lightning arresters for network ports have been properly connected.

## Upgrading the solution and its subsystems



HPE servers and their installed hardware options are preloaded with the latest firmware. However, updated firmware might be available. HPE offers two products to keep firmware updated:

- **HPE Smart Update Manager (SUM)** is a product used to install and update firmware, drivers, and systems software on ProLiant servers, while taking into account all interdependencies. SUM is preloaded with all relevant interdependency information from extensive HPE testing before each SPP release. It does not require an agent for remote installations because it copies a replica of itself to each of the target servers only for the duration of the installation.

SUM has an integrated hardware and software discovery engine that finds the installed hardware and current versions of firmware and software on target servers. It identifies associated targets that should be updated together to avoid interdependency issues. SUM installs updates in the correct order and ensures that all dependencies are met before deploying an update, including updates for HPE Onboard Administrator and HPE Virtual Connect. It prevents an installation if there are version-based dependencies that it cannot resolve.

SUM provides a GUI and a command-line scriptable interface for deployment of systems software for single or multiple ProLiant servers and network-based targets, such as iLOs and HPE Fibre Channel modules. SUM supports deploying up to 50 nodes at a time in all SUM modes. Deployment times depend on system and server environment variables.

A deployment screen provides details on components that need updates, including estimated deployment time. SUM can perform updates immediately or on a schedule and can reboot immediately after update, drive the operating system to delay the reboot by up to an hour, or wait for a reboot that is part of a regularly scheduled maintenance protocol.

---

**NOTE:** Scan this QR code or right-click it to open the hyperlink for more information and documentation on SUM.



<https://www.hpe.com/us/en/product-catalog/detail/pip.5182020.html>

---

- **HPE Service Pack for ProLiant and firmware bundles**—These are collections of comprehensive systems software (drivers and firmware) delivered as a single package with major server releases. This solution uses SUM as the deployment tool and is tested on all supported ProLiant servers. SPP can be used in online mode on a Windows or Linux hosted operating system or in offline mode, where the server is booted to an operating system included on the ISO file. In offline mode, the server can be updated automatically with no user interaction or updated in interactive mode.

Each downloadable set is heavily tested for coherency and interdependencies in the HPE labs. These convenient bundles are released with most major ProLiant server releases and twice per year (April and October) as a Maintenance Supplement Bundle (MSB). An MSB is a bundle that is associated with the most recently released SPP. It includes hot fixes, security bulletin-related components, and operating system supplement components. It also includes components needed to provide delayed application or device support and new device support that has been released since the most recently released SPP. An MSB leverages the SUM version found in the most recently released SPP.

The results of the interdependency testing are coded into SUM. Each firmware or driver update is a self-contained executable that takes care of updating the existing firmware or driver with a newer release, and double checks that it is indeed executing against the right hardware. They perform the update when the operating system (Windows, Linux, or VMware) is up and running.

The SPP is currently delivered as a full ISO. Users can also create a custom SPP of those components and operating systems used in their environment and save it as a zip, ISO, or bootable ISO. The SPP custom download also allows users to easily add supplements, hotfixes, or MSBs to applicable SPPs.

---

**IMPORTANT:** Although a ProLiant server is preloaded with firmware and drivers, you should update the firmware upon initial setup to ensure you have the latest versions. Also, downloading and updating the latest version of Intelligent Provisioning ensures the latest supported features are available. For ProLiant servers, firmware is updated using the Intelligent Provisioning Firmware Update utility. Do not update firmware if the version you are currently running is required for compatibility.

---

---

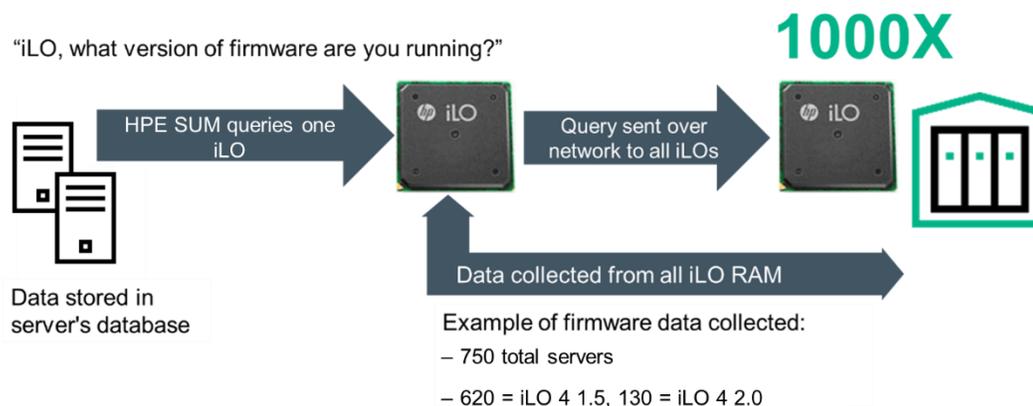
**NOTE:** For more information on SPPs, scan this QR code or right-click it to open the hyperlink.



<http://www.hpe.com/servers/spp>

---

## HPE SUM and iLO Federation



Process of collecting firmware data

When you log in to SUM, it automatically searches for iLO Federation groups on connected networks. SUM searches each group and displays the nodes that respond to the SUM search. The Edit scalable update group screen is used 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 displays the Power Management Controller (PMC), complex programmable logic device (CPLD), system ROM, and iLO firmware. 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 launching SUM. Having 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.

Other systems ask iLOs for data on a periodic basis, and store that data in a database. They might request server data such as temperatures, profiles, or firmware versions from iLO once an hour and store that in the server's database or HPE OneView appliance. This is less scalable than the iLO Federation approach, which includes virtual real-time updates.

iLO Federation management provides scalability enhancements when used with SUM:

- Automatically discover iLO Federation groups on the management network.
- Update the iLO and ROM firmware online on ProLiant servers in the iLO Federation Group through the iLO.
- Update all applicable firmware on ProLiant servers in the iLO Federation Group through the iLO using offline firmware deployment.
- Apply updates to all members of an iLO Federation group (Advanced iLO license required).

## Nondisruptive activation for switch firmware upgrades

Load and activate firmware upgrades on an operating switch without disrupting data traffic or having to reinitialize attached devices. If the nondisruptive activation fails, you will usually be prompted to try again later. Otherwise, the switch will perform a disruptive activation. A disruptive activation interrupts Fibre Channel data traffic on the switch, but a nondisruptive activation does not. For information about nondisruptive firmware versions, see the firmware release notes.

When installing firmware on more than one switch in the fabric, wait until the activation is complete on the first switch before installing firmware on a second switch. If you attempt to activate firmware on a second switch before activation is complete on the first, a message advises you to wait and perform a hot reset later on the second switch to complete the installation. Ports that change states during the nondisruptive activation are reset.

When the nondisruptive activation is complete, SAN Connection Manager sessions, QuickTools sessions, and Enterprise Fabric Management Suite sessions are automatically reconnected. However, telnet sessions must be restarted manually.

A custom firmware installation downloads the firmware image file from a remote host to the switch, unpacks the image file, and resets the switch in separate steps. This allows you to choose the type of switch reset and whether the activation will be disruptive (`reset switch` command) or nondisruptive (`hotreset` command). The following example illustrates a custom firmware installation with a nondisruptive activation. To log in to the switch through telnet:

1. Open a command line window on a workstation and enter the telnet command followed by the switch IP address:

```
# telnet ip_address
```

The IP address can be one of the following:

- 4-byte IPv4 address
- 16-byte IPv6 address
- DNS host name (requires a DNS server)

The telnet window opens, prompting you to log in.

2. Enter an account name and password. The default account name is `admin`, and the password is `password`.

```
Switch login: adminpassword: xxxxxxxx
```

The following warning is displayed each login until the user changes the default password:

**Warning:** User account password has not been changed. HPE strongly recommends that it be done before proceeding.

3. Download the firmware image file from the workstation to the switch.

- If the workstation has an FTP server, enter the `image fetch` command:

```
SN6000 FC Switch #> admin start
SN6000 FC Switch (admin) #> image fetch account_name ip_address
filename
```

- If the workstation has a TFTP server, enter the `image tftp` command to download the firmware image file:

```
SN6000 FC Switch (admin) #> image tftp ip_address filename
```

- If the workstation has neither an FTP nor a TFTP server, open an FTP session and download the firmware image file by entering the following FTP commands:

```
>ftp ip_address or switchname
user:images
password: images
ftp>bin
ftp>put filename
ftp>quit
```

4. Display the list of firmware image files on the switch to confirm that the file was loaded.

```
SN6000 FC Switch #> admin start
SN6000 FC Switch (admin) $> image list
```

5. Unpack the firmware image file to install the new firmware in flash memory.

```
SN6000 FC Switch (admin)
$> image unpack filename
```

6. Wait for the unpack action to be completed.

```
Image unpack command result: Passed
```

7. A message will prompt you to reset the switch to activate the firmware. Use the `hotreset` command to attempt a nondisruptive activation. If the nondisruptive activation fails, you will usually be prompted to try again later. Otherwise, the switch will perform a disruptive activation.

```
SN6000 FC Switch (admin)
$> hotreset
```

## Learning check

1. Which components need to be checked in order to validate that both the installation and the configuration you installed were successfully completed?

---

---

---

---

2. Which HPE product prevents an installation if there are unresolved version-based dependencies?
  - a. HPE Smart Update Manager
  - b. Maintenance Supplement Bundle
  - c. HPE Service Pack for ProLiant
  - d. HPE Intelligent Resilient Framework

## Summary

- New equipment for an SMB is typically installed in a single-server environment and often in a remote office. 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. Before installing and configuring IT solution components, it is important to make sure that the customer understands the physical requirements for the installation.
- There are three ways to configure a ProLiant server:
  - UEFI/BIOS/Platform Configuration (RBSU)
  - HPE iLO 5 Configuration Utility
  - HPE Intelligent Provisioning
- Your decision to use a DAS, NAS, or SAN configuration in the data center determines how to connect the storage system to the devices that need to access it.
  - The Smart Storage Administrator is a web-based application that helps you configure, manage, diagnose, and monitor HPE Smart Array controllers and HBAs.
  - HPE MSA storage solutions are configured by using the Storage Management Utility, which is a web-based application for configuring, monitoring, and managing the storage system.
  - StoreVirtual storage enables you to create a virtualized pool of storage resources and manage a SAN.
  - HPE StoreEasy systems enable organizations to store user data with greater efficiency and consistency in centralized storage instead of individual user PCs.
- There are several best practices to keep in mind when configuring a data center network. HPE recommends using both in-band and out-of-band system management. HPE also recommends a tiered, core-edge topology to optimize performance, management, and scalability. This topology provides good performance without unnecessary interconnections.
- After you have designed and installed a solution for your customer, the next important step is validating that both the installation and the configuration you installed were successfully completed.
- HPE servers and their installed hardware options are preloaded with the latest firmware. HPE offers two products to keep firmware updated: HPE Smart Update Manager and HPE Service Pack for ProLiant.





**Hewlett Packard  
Enterprise**

---

# **Managing, Monitoring, Administering, and Operating HPE SMB Solutions and Related Components**

## **Module 4**

### **Learning objectives**

After completing this module, you should be able to:

- Explain how small to medium-sized business (SMBs) can use Hewlett Packard Enterprise (HPE) tools to manage and administer data centers
- Manage and monitor storage in data centers
- Manage and monitor data center networks

## Customer scenario

When working your way through this course, it is helpful to keep a customer in mind. In this way, you can imagine the customer's business and technical needs in a real-world scenario, and apply what you learn to this scenario as the business grows and evolves. Consider the following fictitious customer scenario. Refer back to it as you engage with labs and activities throughout this course.

### IDJT HealthCare



IDJT HealthCare serves physicians and caregivers across medical focus areas

IDJT HealthCare has decided to process its own medical insurance transactions. To support the new applications, they need a virtualization environment that runs on VMware ESXi. IDJT plans to manage and maintain its own servers to support the new application. The application requires virtual machines (VMs) that run Microsoft Windows Server 2016. They also need management software that will allow them to monitor and manage their datastores across their data center.

At a customer meeting, Eric Yurekha, the IDJT HealthCare IT manager, said, "We have struggled with various management solutions in the past because we had different software for our server, storage, and networking devices. We need a program that will help us manage the data center easily with no unplanned downtime during firmware updates."

Keep this statement in mind as you work through this module. Try to form a response to this comment by the end of the module.

## Gauge your knowledge

Before proceeding with this section, answer the following questions. Pay particular attention to the answers you are less certain about. Keep in mind that there are no right or wrong answers to these questions.

1. What considerations do you think are important when managing a data center?

---

---

2. How do you think managing a server is different from managing a storage array or a networking device?

---

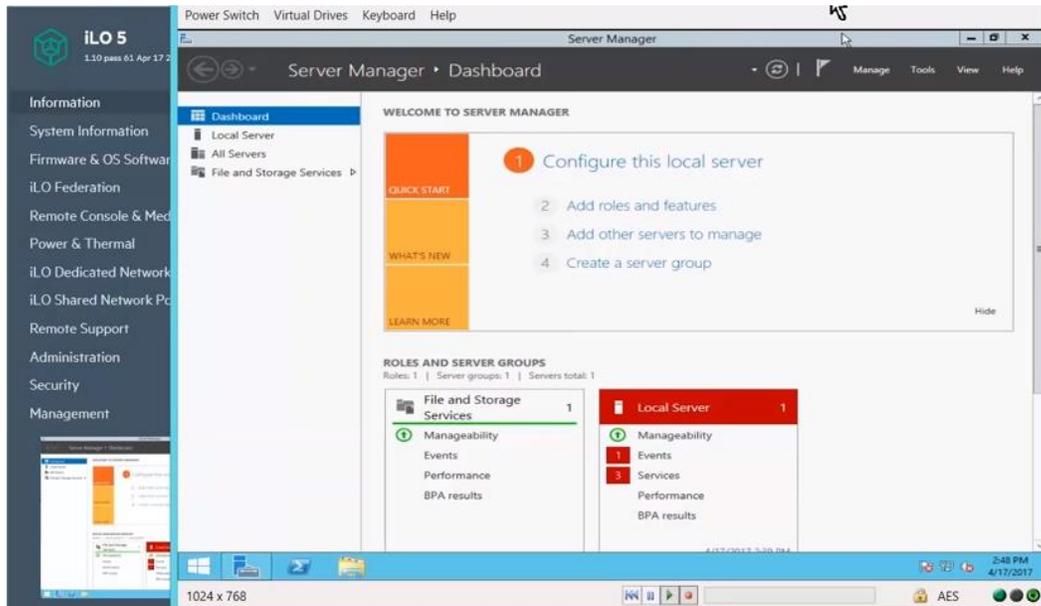
---

---

## HPE management and administration tools

HPE provides a comprehensive set of management programs to help SMB companies meet their administration requirements at every stage of the server life cycle. Two offerings allow IT administrators to manage and monitor data systems: HPE integrated Lights-Out (iLO) 5 for HPE ProLiant servers and HPE OneView.

### Using iLO 5 to manage a data center



iLO 5 enables you to remotely access the Server Manager dashboard

iLO management enables you to monitor and repair servers from remote locations. Its architecture is composed of an independent microprocessor running an embedded operating system that ensures that most iLO functionality is available, regardless of the host operating system. The installed operating system and system configuration determine the installation requirements.

Access iLO from a remote network client by using a supported browser and providing the default Domain Name System (DNS) name, user name, and password. iLO 5 supports the latest versions of the following browsers:

- Microsoft Edge and Microsoft Internet Explorer
- Mozilla Firefox
- Google Chrome mobile and desktop

The browser must meet the following requirements:

- The iLO web interface uses client-side JavaScript; this is not enabled by default in Internet Explorer.
- Cookies and pop-up windows must be enabled for certain features to function correctly. Verify that pop-up blockers are disabled.
- To access the iLO web interface, Transport Layer Security (TLS) 1.0 or later must be enabled.
- The iLO driver enables HPONCFG and the HPE Agentless Management Service to communicate with the iLO processor. With Agentless Management, hardware management capabilities run on the iLO hardware, independent of the host operating system and processor. Health monitoring and alerting begin working the moment you supply power to the server.

## Agentless Management

Configuring iLO Agentless Management settings for SNMP

Agentless Management allows ProLiant servers to collect and deliver hardware and other server management information to management consoles such as HPE OneView and HPE Insight Control without requiring the installation of traditional agents or providers on the host operating system.

With Agentless Management, the Simple Network Management Protocol (SNMP) management software operates within the iLO firmware instead of the host operating system. It monitors and generates SNMP traps and additional operating system data independently of the operating system and processor. This frees memory and processor resources on the server for use by the operating system and server applications.

Another advantage is that in addition to monitoring all key internal subsystems (such as thermal, power, and memory), iLO sends SNMP management notifications directly to management consoles, even with no host operating system installed. Agentless Management not only simplifies agent management regardless of the host operating system, but also provides an iLO-dedicated management network isolated from the regular data network.

## iLO web interface

You can use the iLO web interface to manage iLO. 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.

You can also use a remote console XML configuration and control scripts, Systems Management Architecture for Server Hardware (SMASH) Command-Line Protocol (CLP), or the iLO Representational State Transfer (RESTful) application programming interface (API). The iLO RESTful API is a management interface that server management tools use to perform configuration, inventory, and monitoring tasks by sending basic HTTPS operations (GET, PUT, POST, DELETE, and PATCH) to the iLO web server.

---

**NOTE:** To learn more about the iLO RESTful API, scan this QR code or right-click it to open the hyperlink.



<http://www.hpe.com/info/restfulinterface/docs>

For information about using the iLO RESTful API to automate tasks, visit the following site to access libraries and sample code. Scan this QR code or right-click it to open the hyperlink.



<http://www.hpe.com/info/redfish>

---

## Logging in to the iLO web interface

1. Enter: `https://<iLO host name or IP address>`

When you access the iLO web interface, you must use HTTPS (HTTP exchanged over an SSL encrypted session).

The iLO login page opens. If a login security banner is configured, the banner text is displayed in the NOTICE section.

2. Do one of the following:
  - On the login page, enter a directory or local user account name and password, and then click **Log In**.
  - Click the **Zero Sign In** button. If iLO is configured for Kerberos network authentication, the Zero Sign In button is displayed below the Log In button. You can use the Zero Sign In button to log in without entering a user name and password.
  - Connect a smartcard, and then click the **Log in with Smartcard** button. If iLO is configured for Common Access Card (CAC) smartcard authentication, the Log in with Smartcard button is displayed below the Log In button. Do not enter a login name and password when you use CAC smartcard authentication.

## Using iLO controls



Click each icon to complete various tasks

When you log in to the iLO web interface, the iLO controls are available from any iLO page. From left to right, these icons are:

- **Power icon**—Click this icon to access the Virtual Power Button features. The Virtual Power Button section on the Server Power page displays the current power state of the server, as well as options for remotely controlling server power. System Power indicates the state of the server power when the page is first opened. The server power state can be ON, OFF, or Reset. Use the browser refresh feature to view the current server power state. The server is rarely in the Reset state.
- **UID icon**—Click this icon to turn the UID LED on and off. The UID LED helps you identify and locate a server, especially in high-density rack environments. The possible states are UID ON, UID OFF, and UID BLINK. If the iLO Service Port is in use, UID BLINK status includes the Service Port status. The possible values are UID BLINK (Service Port Busy), UID BLINK (Service Port Error), and UID BLINK (Service Port Finished). When the UID is blinking, and then it stops blinking, the status reverts to the previous value (UID ON or UID OFF). If a new state is selected while the UID LED is blinking, that state takes effect when the UID LED stops blinking.



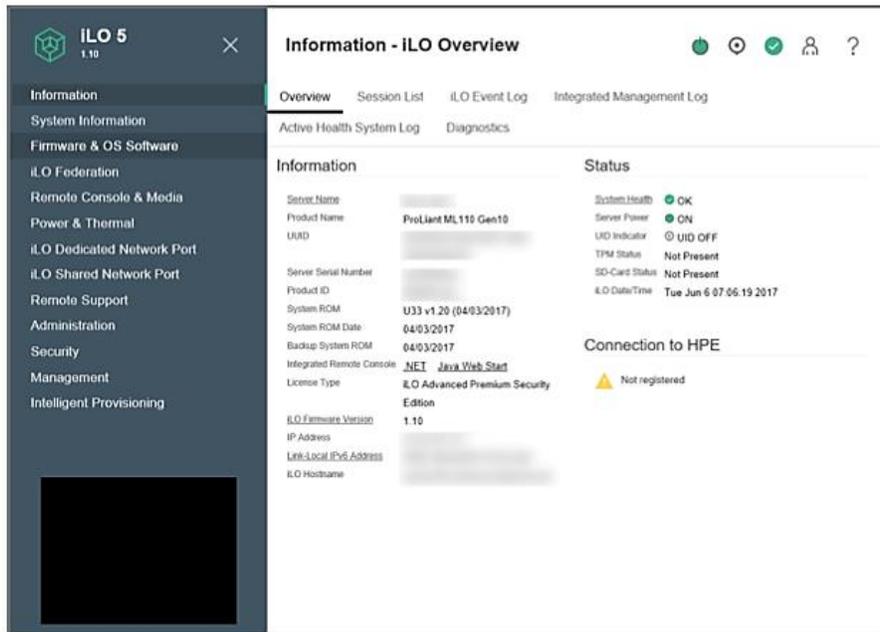

---

**CAUTION:** The UID LED blinks automatically to indicate that a critical operation is underway on the host, such as Remote Console access or a firmware update. Do not remove power from a server when the UID LED is blinking.

---

- **Language**—Click this icon to select a language for the current iLO web interface session. To view or modify the language settings, click the **Language** icon, and then select **Settings**.
- **Health icon**—Click this icon to view the overall health status for the server fans, temperature sensors, and other monitored subsystems. For all components except the HPE Agentless Management Service (AMS), click a component to view additional details.
- **User icon**—Click this icon to do the following:
  - To log out of the current iLO web interface session, click the **User** icon, and then select **Logout**.
  - To view the active iLO sessions, click the **User** icon, and then select **Sessions**.
  - To view or modify iLO user accounts, click the **User** icon, and then select **Settings**.
- **Help icon**—Click this icon to view online help for the current iLO web interface page.
- **Ellipsis icon**—This icon is displayed on the Firmware & OS Software page when the browser window is too small to show the full page. To access the Update Firmware and Upload to iLO Repository options, click this icon. These options are available on all Firmware & OS Software tabs.

## Viewing iLO overview information



Information – iLO Overview screen

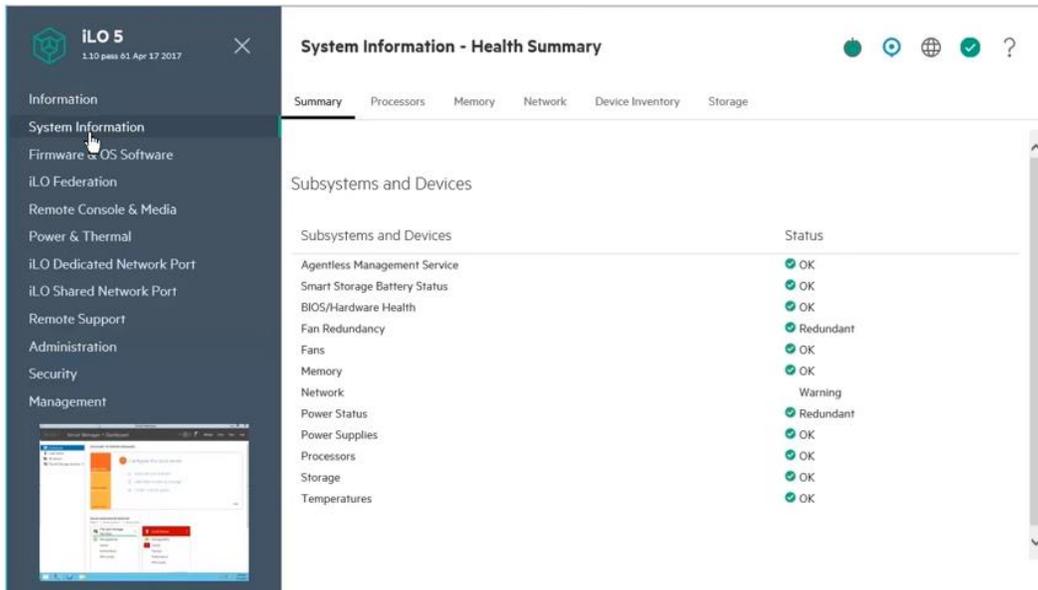
To view iLO overview information, navigate to the Information – iLO Overview page. The Information – iLO Overview page displays high-level details about the server and the iLO subsystem, as well as links to commonly used features.

**NOTE:** For more information on configuring iLO and monitoring systems, refer to the *iLO 5 User Guide*. You can access this document by scanning this QR code or by right-clicking it to open the hyperlink.



<http://www.hpe.com/support/ilo5-ug-en>

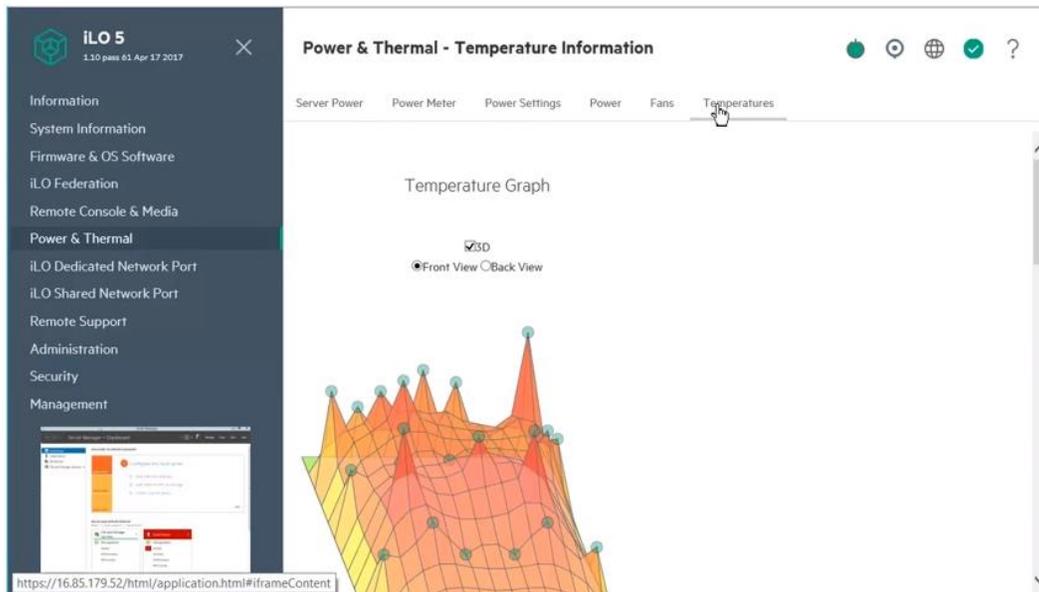
## iLO server health summary



System Information – Health Summary screen

With ProLiant Gen9 servers, users can access iLO server health summary by pressing the UID button when the server is powered off. In Gen10, this summary can be accessed even if the server is running, and not just on auxiliary power. The screen displays asset information and critical IML entries. A thumbnail of the live screen is shown on the left side of the screen.

## Additional iLO functions



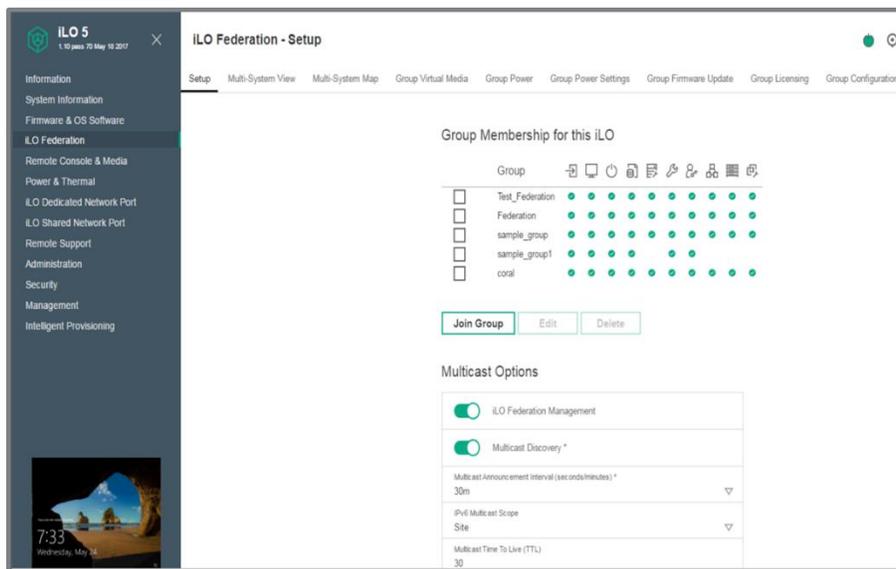
iLO 5 3D thermal monitoring

The following functions, both standard and licensed, are key benefits of iLO 5:

- **Provision**—Rapid discovery and remote access features enable you to inventory and deploy servers using Virtual Media and iLO Federation remotely with the iLO web interface, remote console, command-line interface (CLI), or mobile application.
  - Discovery Services automatically reports server locations to HPE SIM and Insight Control. Power Discovery Services is an enhancement to the iPDU technology. It automatically reports iPDU power status.
- **Monitor**—System health and performance protection with advanced power and thermal control enables you to achieve maximum power efficiency.
  - Advanced Power Management provides access to power-related data from any of the three iLO interfaces (browser, script, or command line). Available information includes time spent in Power Regulator Dynamic Savings mode; average, peak, and minimum power consumption over 24-hour intervals; and iLO Power Meter (when supported by the host).
  - Embedded system health monitors fans, temperature sensors, and power supply sensors without loading the System Management driver. You can access the status of these components from all iLO for ProLiant user interfaces (browser, Systems Management Architecture for Server Hardware (SMASH) command line protocol (CLP), and script), independently of the host operating system. The management processor also reports sensor status to the operating system through an Intelligent Platform Management Interface (IPMI) specified interface. The intelligence of iLO manages the Sea of Sensors thermal control, directs the Dynamic Power Capping technology, and monitors the health of server components.

- **Optimize**—iLO management optimizes core lifecycle management functions for instant availability.
  - Integrated Remote Console (IRC) enables you to access, control, and configure system properties. IRC turns a supported browser into a virtual desktop, giving you full control over the display, keyboard, and mouse of the host server. You can use IRC to access the remote file system and network drives, observe POST boot messages as the remote host server restarts, and initiate ROM-based setup routines to configure the remote host server hardware. When you are installing operating systems remotely, IRC enables you to view and control the host server monitor throughout the installation process.
  - Kerberos authentication integrates iLO authentication into a customer security system. If the client workstation is logged in to the domain, and the user is a member of a directory group for which iLO is configured, Kerberos enables a user to log in to iLO without supplying a user name and password.
- **Support**—Core instrumentation that operates whether the operating system is up or down enables you to view the Integrated Management Log and Active Health System logs.
  - iLO provides a free Active Health System viewer that you can use to quickly view and self-diagnose any server issues. You can also send logs to HPE support for additional assistance.
  - Active Health System Diagnostics provides a 24/7 control center for servers. It enables you to continuously monitor more than 1600 system parameters and receive consolidated health and service alerts.

## iLO 5 Federation



iLO Federation – Setup screen

With an HPE iLO Advanced license, you can enable the full implementation of iLO Federation management. iLO Federation enables you to manage multiple servers from one system by using the iLO web interface.

When configured for iLO Federation, iLO uses multicast discovery and peer-to-peer communication to enable communication between the systems in an iLO Federation group. When an iLO Federation page loads, a data request is sent from the iLO system running the web interface to its peers, and from those peers to other peers until all data for the selected iLO Federation group is retrieved.

iLO 5 supports the following features:

- **Group health status**—View server health and model information.
- **Group virtual media**—Connect scripted media for access by the servers in an iLO Federation group. It enables you to replicate same script media within an iLO Federation group.
- **Group power control**—Manage the power status of the servers in an iLO Federation group. It lets you power several servers on and off with a single click.
- **Group power capping**—Set dynamic power caps for the servers in an iLO Federation group. It allows servers to share power consumption among a group of servers.
- **Group firmware update**—Update the firmware of the servers in an iLO Federation group.

---

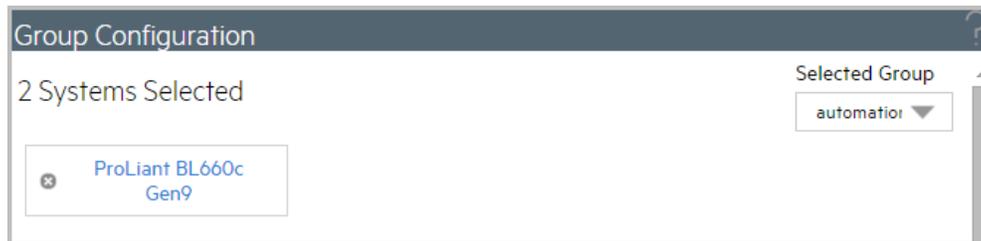
**NOTE:** For information about configuring iLO Federation, refer to the *HPE iLO Federation User Guide for iLO 5*. You can access this guide by scanning this QR code or right-clicking it and opening the hyperlink.



[http://h20566.www2.hp.com/hpsc/doc/public/display?docId=a00018325en\\_us](http://h20566.www2.hp.com/hpsc/doc/public/display?docId=a00018325en_us)

---

## iLO Federation group characteristics



The iLO Federation→Group Configuration page

iLO Federation groups allow iLO systems to encrypt and sign messages to other iLO systems in the same group. iLO Federation groups can overlap, span racks and data centers, and be used to create management domains.

An iLO system can be a member of up to 10 iLO Federation groups. There is no limit on the number of iLO systems that can be in a group. You must have the Configure iLO Settings privilege to configure group memberships. You can use RIBCL XML scripts to view and configure group memberships.

HPE recommends installing the same version of the iLO firmware on iLO systems that are in the same iLO Federation group.

## iLO Federation group privileges

All iLO systems are automatically added to the DEFAULT group, which is granted the Login privilege for each group member. You can edit or delete the DEFAULT group membership. Additionally, a group can be granted the following privileges:

- **Remote Console**—Group members can remotely access the managed server Remote Console, including video, keyboard, and mouse control.
- **Virtual Power and Reset**—Group members can power-cycle or reset the managed server.
- **Virtual Media**—Group members can use scripted virtual media with the managed server. Group Virtual Media enables you to connect scripted media for access by the servers in an iLO Federation group.
- **Host BIOS**—Group members can configure the host BIOS settings by using the UEFI System Utilities.
- **Configure iLO Settings**—Group members can configure most iLO settings, including security settings and group memberships, and can remotely update firmware.
- **Administer User Accounts**—Group members can add, edit, and delete iLO user accounts.
- **Host NIC**—Group members can configure the host NIC settings.
- **Host Storage**—Group members can configure the host storage settings.
- **Recovery Set**—Group members can manage the recovery install set. This privilege is not available if you start a session when the system maintenance switch is set to disable iLO security.

## iLO Federation group memberships for local iLO systems



Configuring group memberships for a local iLO system specifies members' privileges

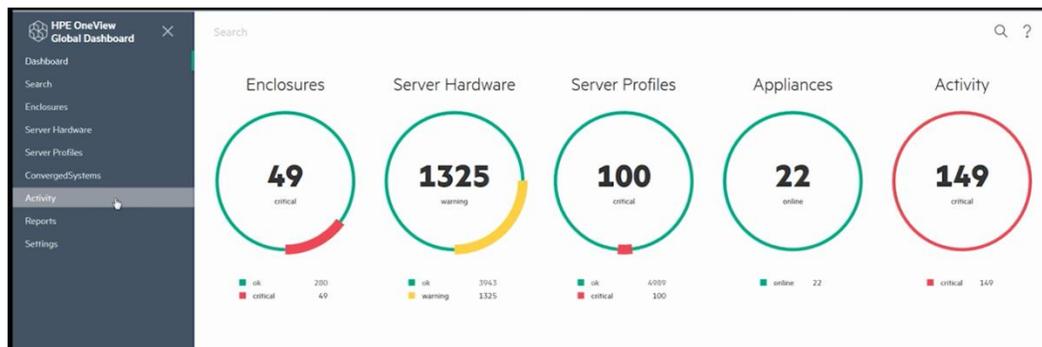
You can use the iLO web interface to configure group memberships for a local iLO system or a group of iLO systems. You also can use the iLO RESTful API to configure group memberships or RIBCL XML scripts to view and configure group memberships. You must have the Configure iLO Settings privilege to configure group memberships.

When you configure group memberships for a local iLO system, you specify the privileges that members of a group have for configuring the local managed server. For example, if you add the local iLO system to group1 and assign the Virtual Power and Reset privilege, the users of other iLO systems in group1 can change the power state of the managed server. If the local iLO system does not grant the Virtual Power and Reset privilege to group1, the users of other iLO systems in group1 cannot use the group power control features to change the power state of the managed server. If the system maintenance switch is set to disable iLO security on the local iLO system, the users of other iLO systems in group1 can change the state of the managed server, regardless of the assigned group privileges.

Group memberships for the local iLO system are configured from the **Setup** tab on the **Administration** → **iLO Federation** page. You can perform the following tasks for a local iLO system:

- View group memberships
- Add and edit group memberships
- Remove group memberships

## Using HPE OneView to manage a data center



HPE OneView Global Dashboard Activity

HPE OneView is a single integrated platform, packaged as an appliance, which is a preconfigured virtual machine ready to be deployed on a hypervisor host. HPE OneView implements a software-defined approach to managing a physical infrastructure through its entire life cycle. It is designed for the physical infrastructure needed to support virtualization, cloud computing, big data, and mixed computing environments.

HPE OneView automates tasks and streamlines processes delivering the entire life cycle management of the infrastructure across compute, storage, and fabric resources. IT staff can control resources programmatically through a unified API, enabling easy integration with workflow and management tools. HPE OneView supports key scenarios such as deploying bare-metal servers, deploying hypervisor clusters from bare metal, performing ongoing hardware maintenance, and responding to alerts and outages. HPE OneView makes it possible to easily monitor, configure, and manage physical and logical server, network, and storage resources through either a GUI or by using RESTful APIs.

Management software is integrated with HPE OneView for seamless operation. Environmental management such as power, cooling, and space planning requires that you consider all the equipment in the entire data center, including equipment not managed by HPE OneView. HPE OneView consolidates data center power and cooling information into one interface view.

HPE OneView combines complex and interdependent data center provisioning and management into one simplified and unified interface. You can:

- **Provision the data center**—After you install the HPE OneView appliance and perform the initial configuration tasks, you can quickly bring existing hardware under management, and prepare for and deploy hardware to the data center.
- **Manage and maintain firmware and configuration changes**—When you add a resource to the appliance to be managed to ensure compatibility and seamless operation, the appliance automatically updates the resource firmware to the minimum version required to be managed by the appliance.
- **Monitor the data center and respond to issues**—You use the same interface for monitoring that you use to provision resources. There are no additional tools or interfaces to learn.

HPE OneView also offers predefined reports to help you manage your appliance and its environment. You can view the reports in the UI or generate them by using REST API. You can also save the reports as a Microsoft Excel workbook (\*.xlsx) or CSV MS-DOS (\*.csv).

## HPE OneView licensing

HPE OneView manages servers and enclosure networking resources, supports connections from enclosures to storage, and provides information to help you manage data center power and cooling. Servers are represented and managed through server profiles and server profile templates. Servers can be added to HPE OneView in one of the following ways:

- **Managed**—If you add a managed server to HPE OneView, you can apply configurations, deploy server profiles, monitor operation status, collect statistics, and alert users to specific conditions. Managing server hardware requires HPE OneView Advanced licensing.
- **Monitored**—If you add a monitored server to HPE OneView, you can monitor it for inventory and hardware status only. Monitoring server hardware uses the free HPE OneView Standard license.
- **Migrated**—HPE BladeSystem enclosures can be migrated to HPE OneView with their configuration information, so that the enclosure can be managed by HPE OneView. The managed enclosure requires HPE OneView Advanced licensing.

The HPE OneView management appliance controls licenses. The same management appliance can be used for both HPE OneView Advanced licenses and for HPE OneView Standard. You can make this choice when you initially add a system to the HPE OneView management appliance.

- HPE OneView Advanced provides full-featured licenses, which are licensed per physical server. These licenses include three years of 24x7 Technical Support and Updates (TS&U) with web-based training (WBT) to build basic product proficiency. Full HPE OneView Advanced licenses also provide integrated right-to-use HPE Insight Control at no additional charge. Trial versions of HPE OneView Advanced can be used for 60 days without charge. HPE OneView Advanced management software licenses can be used for 60 days without charge.
- HPE OneView Standard can be used for inventory, health monitoring, alerting, and reporting without additional fees. The user interface is similar to the HPE OneView Advanced version, but the software-defined functionality is not available. An annual 9x5 support offering is available for an additional fee. After the license is deployed, both storage and servers are monitored in HPE OneView, and the storage topology is accessible in Map View. HPE OneView Standard management software can be used without charge.

The HPE OneView management appliance controls the licenses, and it can be obtained in two ways:

- Software download (OVA) from the HPE Software Depot
- Purchase of the HPE OneView Media Kit (contains a USB flash drive)

## Navigating the HPE OneView GUI



Components of the HPE OneView GUI

The preceding graphic illustrates the following components of the HPE OneView GUI:

- Use the top-level menu for navigating the various sections of the UI. Each section is categorized based on function and role (Servers, Networking, Storage, and Facilities). Click the downward caret icon or anywhere in the area to expand the menu.
- The view selector enables you to control the information displayed about a resource so that you can focus only on what you are interested in.
- The Map View icon provides a graphical representation of the relationships between the selected resource and other resources. To see these relationships, select the network icon or select **Map view** in the view selector.
- The Activity control icon expands (or hides) the Activity sidebar for a recent appliance, resource, or user activity (from the current login session and browser window). The Activity sidebar window shows recent alerts and task activity for the current resource. The Activity window displays the current activity of the administrator performing various actions within their session, which can be expanded or collapsed to increase the viewing dimensions of the details pane. Individual activity items can be selected and then directly navigated to in the case of Create and Update actions.
- Your user account enables you to edit some account information, depending on your credentials.
- Online Help control expands (or hides) a sidebar that provides access to UI and REST API help, the EULA and Written Offer, and the HPE OneView online user forum.

- The Details pane provides all information known about a selected resource instance. To see details about a particular resource instance, click its name in the resource selection/filter pane. The Activity Details section within the Details pane displays the most recent activity, including actions performed by an administrator and automated alerts. It can be expanded to view further details, and the administrator can then navigate to the activity item to clear, assign, or provide notes regarding the event.
- The resource selection/filter pane lists all resource instances that have been configured on the appliance. In some cases, a status icon indicates general health of the resource.
- The Actions drop-down menu lists the actions that are available to run on the current resource. Actions include but are not limited to adding, creating, deleting, removing, and editing a resource instance. If you do not have the appropriate permissions to perform an action, the action is not shown on the Actions menu.

In addition to the screen components shown in the graphic, every UI screen has a notifications area that alerts you when an event or activity requires your attention. Some screens also have a filters sidebar that enables you to control the type of information displayed in the master pane.

## Provisioning hardware

**Create Server Profile Template**
General ▾

**General**

---

Name

Description

---

**Server Profile**

Server profile description

Server hardware type

Enclosure group

Affinity

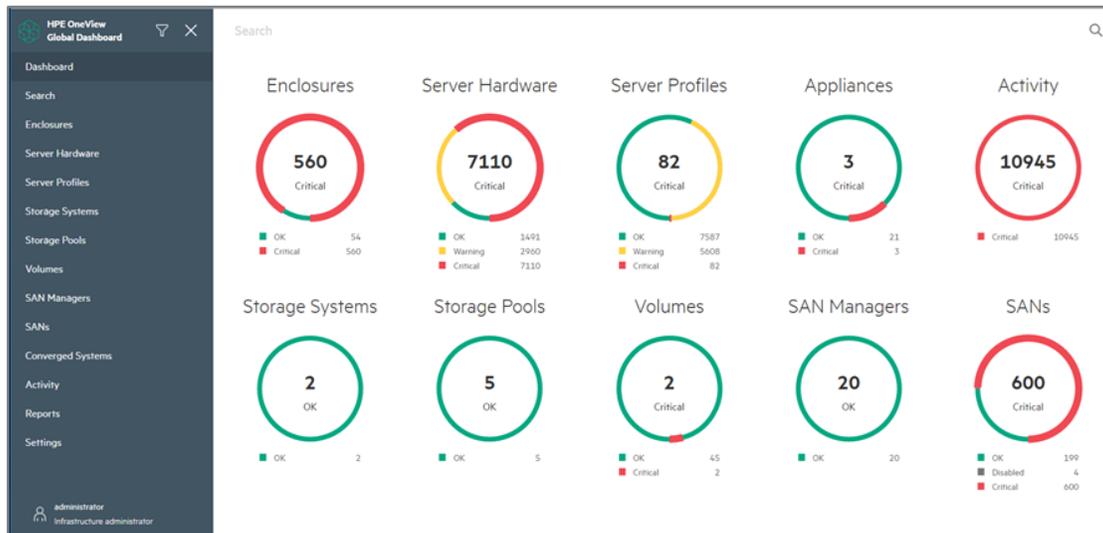
Creating a server profile template in HPE OneView

After you install the HPE OneView appliance and perform the initial configuration tasks, you can quickly bring existing hardware under management and prepare for and deploy hardware to the data center.

Features for provisioning hardware and bringing resources under management include:

- **Resource templates, groups, and sets**—Groups and templates enable you to define configurations that are specific to the environment you want to build, such as virtual hosts, Microsoft Exchange environments, external or internal web servers, or corporate database servers.
- **Server profiles and server profile templates**—Server profiles and server profile templates enable you to provision hardware and resources quickly and consistently according to your best practices. Store your best practice configuration in a server profile template and then use the server profile template to create and deploy server profiles.
- **Streamlined process for bringing hardware under management**—HPE OneView simplifies the process of bringing server hardware and BladeSystem enclosures and interconnects under management.
- **Operating system deployment**—Server profiles and enclosure groups make it easier to prepare a bare-metal server for operating system deployment.
- **Storage provisioning and management**—HPE OneView provides automated, policy-driven provisioning of supported storage resources. It is fully integrated with server profiles so that you can manage new or existing storage infrastructure. With HPE OneView you can view and manage storage system and storage pools. You can add existing volumes, create new volumes, and create volume templates to provision multiple volumes with the same configuration.
  - Switched fabric, direct attach, and virtual SAN (vSAN) topologies are supported.
  - Storage system and storage pools are added to the appliance followed by volumes, which are associated with networks. The volumes can then be attached to server profiles. You can also add SAN managers to make their managed SANs available to the appliance. Managed SANs can be associated with Fibre Channel or Fibre Channel over Ethernet (FCoE) networks on the appliance to enable automated zoning and automatic detection of connectivity.

## HPE OneView Global Dashboard



Assess the health of data center infrastructure across multiple appliances and resources

Each time you log in to the appliance, the HPE OneView Global Dashboard is the first screen you see. Select **Dashboard** from the main menu any time you want to see the Dashboard charts.

---

**IMPORTANT:** The dashboard is blank the first time you log in to the appliance because you have not yet configured any resources. If this is the first time you are logging in to the appliance, refer to the appropriate HPE OneView user guide to define the data center environment and bring the infrastructure under appliance management.

---

The dashboard charts provide a visual representation of the general health and status of the appliance and managed resources. From the dashboard, you can immediately see resources that need attention. For direct access to resources that need attention, click the resource name.

You can customize the dashboard display by adding, deleting, and moving resource panels. The dashboard shows the status of the most relevant resources that are associated with assigned user roles. If you are assigned multiple roles, such as network and storage roles, the default dashboard displays the combination of resources that each role would see on the dashboard.

Hover your cursor over a chart slice to view the resource instances represented by that slice. If you hover over a different slice in the same chart, the text and count displayed in the center of the chart changes. Click a slice to access the resource page filtered by the status or value associated with the slice.

The dashboard updates in real-time, displaying resource changes as they happen on the appliance. A full refresh of all resources occurs every 12 hours by default. The dashboard offers quick access to resources by providing links to the relevant page on the appliance. User roles ensure that only approved users can access monitored appliances.

The HPE OneView Global Dashboard is delivered as an OVA file for installation on a VMware vSphere hypervisor host.

---

**NOTE:** To download the Global Dashboard OVA file from the HPE Software Depot, scan this QR code or right-click it to opening the hyperlink.



<https://h20392.www2.hpe.com/portal/swdepot/index.do>

---

## Provisioning an HPE StoreVirtual volume from a server profile

**Add Volume**

LUN  Auto  Manual

**Volume properties**

Filter pools by defined connections

Storage pool: Gary13VSACluster

Capacity: 350.00 GiB

Provisioning: Thin

Data protection level: Network RAID-0 (None)

Sharing: Private

Permanent

**iSCSI Authentication**

Mutual CHAP will be used to authenticate this attachment. Names and secrets will be autogenerated.

**Storage paths**

Connection ID	Network	Storage Targets	Enabled
7	iscsi-san Ethernet Untagged	pending assignment	<input checked="" type="checkbox"/>

Changed: Capacity to \*350.00\*

**Add** **Add +** **Cancel**

Add Volume screen

To create a StoreVirtual volume and attach it to the server, select **Add Volume** in the profile and then follow these steps:

1. Select a StoreVirtual storage pool in which to create the volume.
2. When you create a volume from a StoreVirtual storage pool, you can configure StoreVirtual volume settings (Thin and thick provisioning).
3. iSCSI Mutual CHAP authentication is automatically generated and configured (boot initiator and target).
4. Create storage paths on iSCSI networks connecting the server and the storage system.

## Adding, removing, and editing StoreVirtual volumes online

**Edit myserver-datavolume** General ?

**General**

Name: myserver-datavolume

Description:

Volume template: None

Storage pool: Gary13VSACluster

Storage system: Gary13VSACluster

**Volume Properties**

Capacity: 500.00 GiB **1**

Sharing:  Private  Shared

**Advanced**

Provisioning: Thin

Data protection level: Network RAID-10 (2-Way Mirror) **2**

Permit Adaptive Optimization: Yes **3**

OK Cancel

Edit server screen

When you add resources to StoreVirtual appliance, they are automatically configured for monitoring health, activity, alerts, and utilization. You can monitor resources immediately without performing additional configuration or discovery steps.

A server reboot is not required for adding private or shared volumes, deleting storage volumes, or editing storage volumes. When you edit a storage volume, you can:

1. Grow volume capacity
2. Change the volume data protection level
3. Change the Adaptive Optimization setting for the volume

## Remote Support in HPE OneView



Steps to enable Remote Support in HPE OneView

By registering for Remote Support in HPE OneView, you enable HPE Proactive Care and automatic case creation for hardware failures on ProLiant Gen8 and later servers and enclosures. Remote Support enables Proactive Care services including Proactive Scan reports and firmware/software analysis reports with recommendations that are based on collected configuration data.

Setup is a one-time process and no license is required; Remote Support works with monitored hardware. After it is enabled, all eligible devices added in the future will be automatically enabled for Remote Support. HPE will contact the customer to ship a replacement part or send an engineer for devices that are under warranty or support contract.

Remote Support is secure. No business data is collected, only device-specific configuration and fault data. All communications are outbound only and use industry-standard TLS encryption to ensure the confidentiality and integrity of the information.

The graphic shows the steps for enabling Remote Support in HPE OneView:

1. On the Settings page, select **Remote Support**.
2. Click the **Enable Remote Support** radio button.
3. Enter contact information and site location, and click **Register**.

Basic collection sends configuration information to HPE for analysis and proactive services in accordance with warranty and service agreements. This data is transmitted every 30 days.

Information sent to HPE includes data about the server's health, configuration, and run-time telemetry. This information is used to troubleshoot issues and perform closed-loop quality analysis. This data is transmitted every seven days.

To disable Remote Support, simply uncheck the **Enable remote support** box. When Remote Support is disabled:

- Service events are not sent to HPE to trigger support cases
- Data collections are not run as scheduled
- You do not need to re-enter any registration details because these are still saved in HPE OneView and at HPE

---

**NOTE:** Enabling Remote Support gives you access to configuration and contract warranty reports in Insight Online in the HPE Support Center. Enabling Insight Online is optional and is typically done in enterprise (not SMB) data centers for quick and efficient cloud-based management and to share information with HPE authorized partners for support and proactive planning.

---

## Learning check

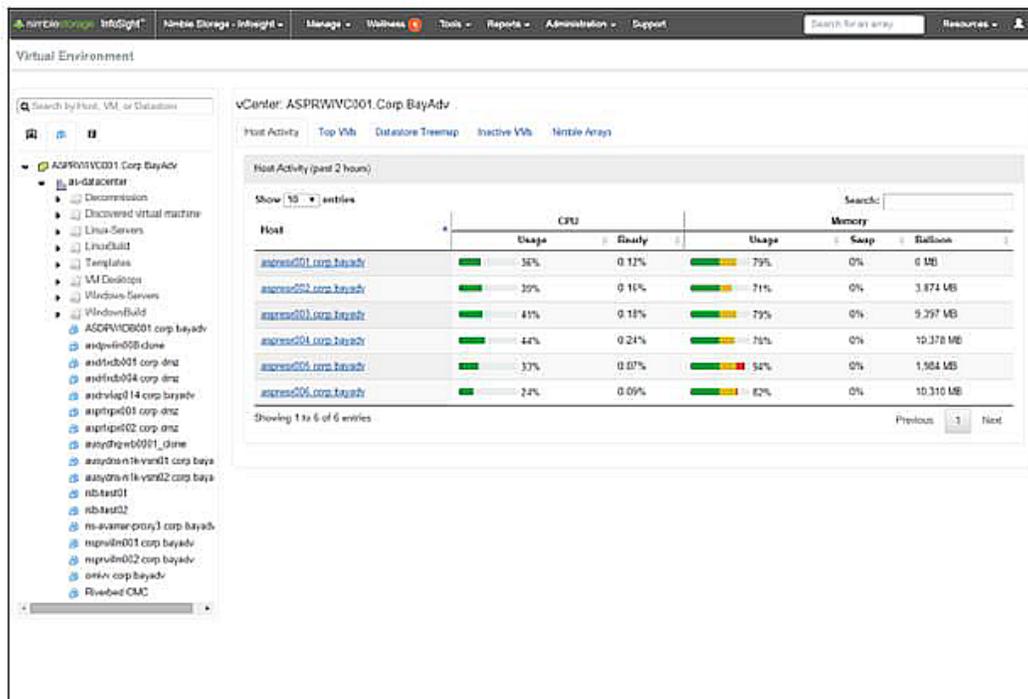
1. Which management interface do server management tools use to perform configuration, inventory, and monitoring tasks by sending basic HTTPS operations (GET, PUT, POST, DELETE, and PATCH) to the iLO web server?
  - a. SMASH CLP
  - b. iLO RESTful API
  - c. RIBCL XML scripts
  - d. A remote console XML configuration
2. How many iLO Federation groups can an iLO system be a member of?
  - a. Only one
  - b. Up to 10
  - c. Up to 100
  - d. There is no limit
3. Which HPE OneView GUI element shows recent alerts and task activity for the currently selected resource?
  - a. Details pane
  - b. Actions menu
  - c. Activity sidebar window
  - d. Resource selection/filter pane
4. HPE OneView Remote Support transmits information about the server's health, configuration, and run-time telemetry to HPE every seven days.
  - True
  - False

## Managing and monitoring storage

HPE provides a variety of tools for managing and administering storage in SMB data centers, including:

- **InfoSight**—For managing HPE Nimble storage arrays
- **Smart Storage Administrator**—For managing storage controllers in a ProLiant server
- **StoreVirtual Centralized Management Console**—For managing StoreVirtual arrays
- **Storage Management Utility**—For managing HPE MSA storage
- **Library and Tape Tools**—For managing HPE tape storage and magneto-optical storage products

## Nimble InfoSight



InfoSight dashboards pinpoint VM-related issues

Nimble Storage has developed a unique and innovative approach to the entire storage life cycle that leverages the power of sophisticated data sciences. The result is Nimble Storage InfoSight, a cloud-based management and support system that integrates, automates, and substantially simplifies storage administrative tasks—ensuring the optimal health of all Nimble Storage arrays.

Unlike most conventional, reactive storage support frameworks that involve disparate tools and exhaustive manual processes, InfoSight simplifies management and with intelligent, analytics-driven automation, can predict and pre-empt issues. This frees IT resources for more strategic activities.

InfoSight automatically opens 90% of all support cases and generates resolutions for more than 80% of them, dramatically reducing the amount of time and effort an IT team spends on support processes. As a key foundational component of the Nimble Adaptive Flash platform, InfoSight provides intelligent, actionable recommendations on how best to scale storage resources in alignment with changing business requirements. With InfoSight, all of this information is easily accessible from a single cloud-connected interface. InfoSight also plays an instrumental role in the services provided by the Nimble Storage support organization and team of concierge managers, enabling them to be highly effective and efficient in managing customers' accounts.

InfoSight enables you to:

- Easily and quickly meet changing service-level agreements with expert guidance on storage resource planning
- Streamline storage administration through intelligent automated support and resolution
- Obtain detailed planning data for more informed decision making
- Maintain peak storage health with accurate, real-time performance analytics
- Simplify storage management with a single cloud-connected portal
- Free IT staff to focus on projects that deliver greater business value

Nimble InfoSight delivers the industry's first proactive, efficient, cloud-connected storage management and support experience, driven by powerful data sciences.

---

**NOTE:** For more information, scan this QR code or right-click it to opening the hyperlink.



**<https://www.nimblestorage.com/technology-products/infosight/>**

---

## HPE Smart Storage Administrator

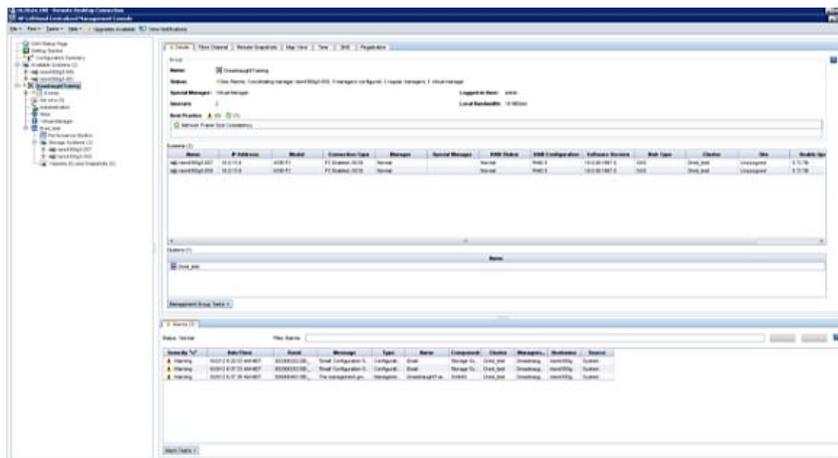


SSA Diagnostics screen

HPE Smart Storage Administrator (SSA) offers the following features and functionalities:

- SSD Over Provisioning Optimization optimizes solid-state drives (SSDs) by deallocating all used blocks before any data is written to the drive.
- Rapid Rebuild Priority determines the urgency with which the controller treats an internal command to rebuild a failed logical drive. There are four settings: low, medium, medium high, and high.
- Auto RAID 0 creates a single RAID 0 volume on each physical drive specified, enabling you to select multiple drives and configure as RAID 0 simultaneously.
- SSD Smart Path improves SSD read performance by bypassing the Smart Array firmware for the optimal performance path to the SSD.
- A scripting capability enables customizable configurations that can be deployed to multiple servers.
- Predictive failure activation mode enables rebuilding to begin before the drive fails, reducing the likelihood of data loss that could occur if an additional drive fails.
- Support for HPE SmartSSD Wear Gauge reporting benefits users by monitoring their usage of SSDs. SSA notifies users of the estimated remaining life of the drive based on the current workload.

## HPE StoreVirtual Centralized Management Console



Common management and federated data services provide simplicity and flexibility in virtual data centers

The HPE StoreVirtual Centralized Management Console (CMC) is the intuitive, GUI-based administrative interface for StoreVirtual devices running LeftHand OS. Use the CMC to configure and manage storage volumes spanning clustered storage nodes and view a single graphical layout of the storage environment. Tab view provides details of the selected object.

## HPE Storage Management Utility

HPE Storage Management Utility V3 (SMU) is the management GUI for modular smart array (MSA) storage. It is available with GL200 firmware or later. Existing MSA customers can choose to use SMU V3 or continue using SMU V2 if newer virtualization features are not required.

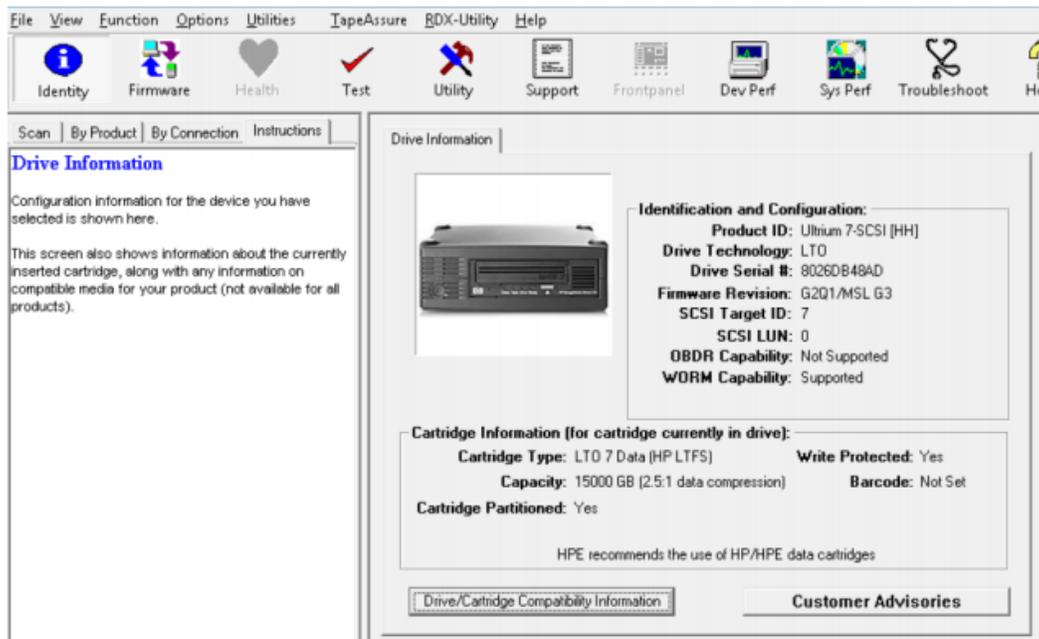
Thin Provisioning allows storage allocation of physical storage resources only after they are consumed by an application. Thin Provisioning also allows overprovisioning of physical storage pool resources, allowing ease of growth for volumes without predicting storage capacity upfront. Thin Provisioning is available with GL200 firmware or later.

### Best practices for MSA

A best practice for monitoring MSA array health is to set up the array to send notifications, which is important for troubleshooting and log retention. SMU is the recommended method for setting up email and SNMP notifications. Setting up these services is easily accomplished by using a web browser. To connect, enter the IP address of the management port of the HPE MSA.

HPE recommends enabling SNMP traps. Version 1 SNMP traps can be sent to up to three host trap addresses. To send version 3 SNMP traps, create an SNMPv3 user with the trap target account type. Use SNMPv3 traps rather than SNMPv1 traps for greater security. SNMP traps can be useful in troubleshooting issues with HPE MSA arrays.

## HPE Library and Tape Tools



L&TT main screen

HPE Library and Tape Tools (L&TT) is a free, downloadable diagnostic tool for all HPE tape storage and magneto-optical storage products. It is ideal for a wide range of customers, including users who want to:

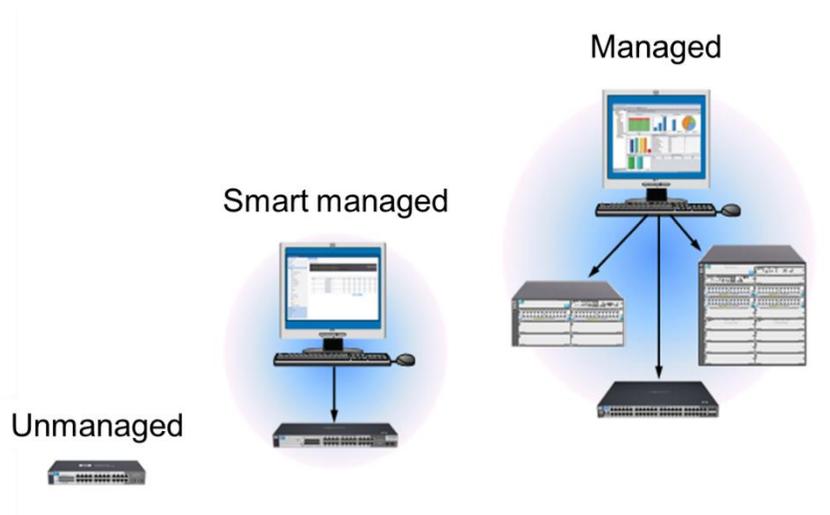
- Verify their installation
- Ensure product reliability
- Perform their own diagnostics
- Achieve faster resolution of tape device issues

L&TT performs firmware upgrades, verification of device operation, failure analysis, and a range of utility functions. Performance tools assist in troubleshooting bottlenecks, and system configuration checks warn of common host issues. It also integrates with HPE support by generating and emailing test results and support tickets. HPE Support requires the use of L&TT to troubleshoot most device issues, so HPE recommends that a support ticket is pulled and the device assessment test is run before calling.

Features and functions of L&TT include:

- **Device identification**—L&TT clearly identifies the storage products connected to the system, along with key information on product configuration and status.
- **Troubleshooting tests and utilities**—L&TT provides tests to verify product functionality and isolate product issues. Tests include device self-tests, read/write tests on drives, and specific device utilities. L&TT includes utilities for working with tape devices, such as erasing and initializing tapes and configuring tape drives.
- **Support ticket generation**—If you experience a problem with a storage product, L&TT can generate a support ticket that includes essential information for troubleshooting the problem. You can email the support ticket to a support center for assistance.
- **Firmware upgrades**—L&TT provides a convenient way of updating product firmware, enabling users with an internet connection to take advantage of ongoing enhancements.

## Managing and monitoring the network



Switches can be managed, smart managed, and unmanaged

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 graphical user interface, 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 Intelligent Management Center



IMC converts meaningless network data to actionable information

HPE Intelligent Management Center (IMC) 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.

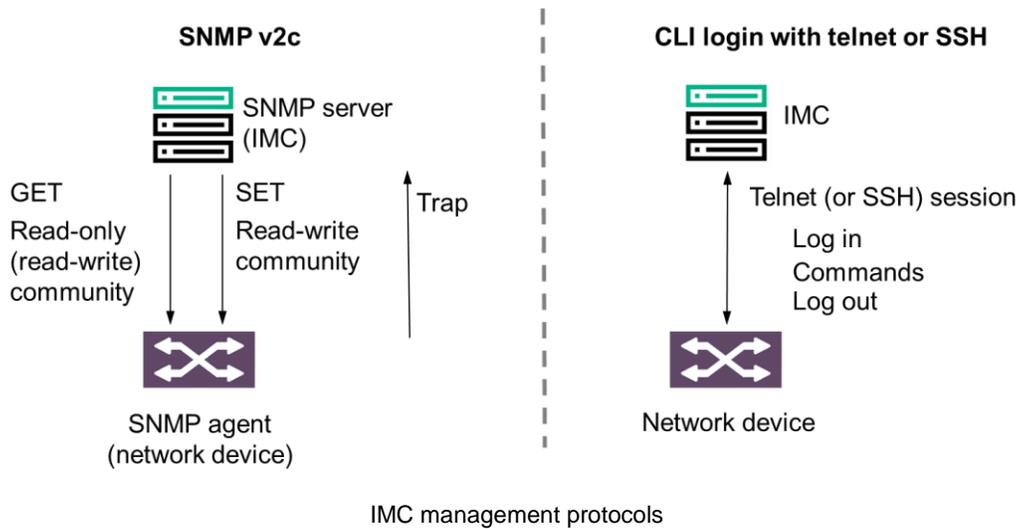
With the IMC auto-discovery feature, you can locate all devices on the network. IMC categorizes discovered devices into types of network devices (such as switches, routers, servers, access points, and desktops) and maps them on a network topology. You can view the network devices based on IP address 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 functionality. Examples of the management capabilities you can add through IMC modules include:

- 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 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.
- User Access Manager (UAM) allows you to translate business policies for access controls into network configurations. You can create service policies and then apply these service policies to user accounts. These policies can be enforced no matter where and how users connect, whether through local Ethernet connections, wirelessly through access points, or remotely through a VPN.
- Endpoint Admission Defense (EAD) helps protect networks from internal threats. Companies create a "posture" which defines the minimum requirements that endpoints must meet before accessing the network. For example, companies can dictate that devices and applications must be patched and running current antivirus software before accessing the network, so that hackers cannot exploit these vulnerabilities.
- 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.
- If you need more visibility into applications, servers, or VMs, you can add the Application Health Manager, Service Health Manager, and Virtualization Monitor (vMON).

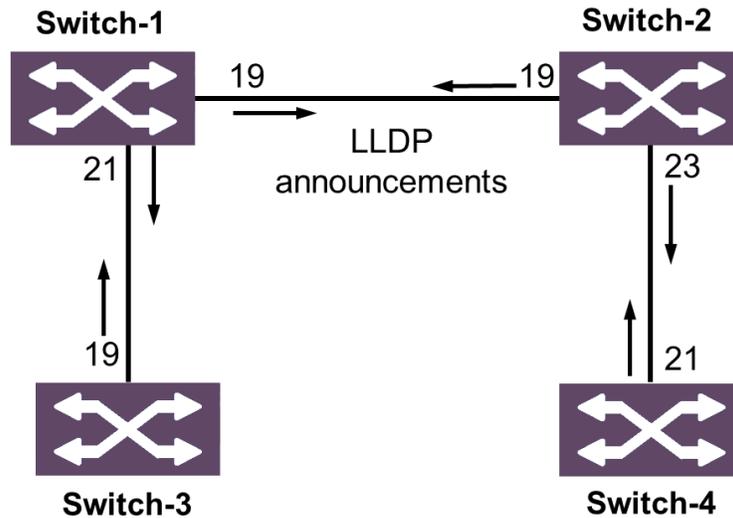
## IMC management protocols



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 read-only 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 is not secure because the community is included in messages in plain text. Therefore, the community string cannot function as a true password.

## Link Layer Discovery Protocol



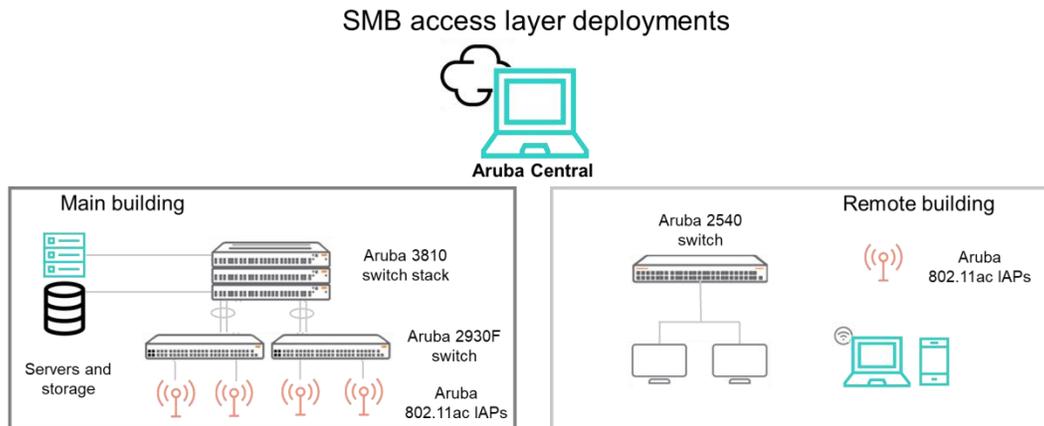
Comware switch configuration example

In a heterogeneous network, devices from different vendors need to be able to discover one another and exchange configuration information. The Internet Engineering Task Force (IETF) defined Link Layer Discovery Protocol (LLDP) in IEEE 802.1AB enables this exchange of information. The protocol operates at the data link layer, enabling directly connected devices to exchange information.

With LLDP, a device exchanges local device information such as its major functions, management IP address, device ID, and port ID. Each device sends this information as type, length, and value (TLV) in LLDP data units (LLDPDUs) to directly connected devices. At the same time, the device receives LLDPDUs from neighbors that support LLDP. The local device saves the information it receives in a standard management information base (MIB). SNMP programs such as IMC can use the LLDP information stored in MIBs to quickly detect Layer 2 network topology changes and identify each change.

All current HPE networking devices support LLDP.

## Aruba management



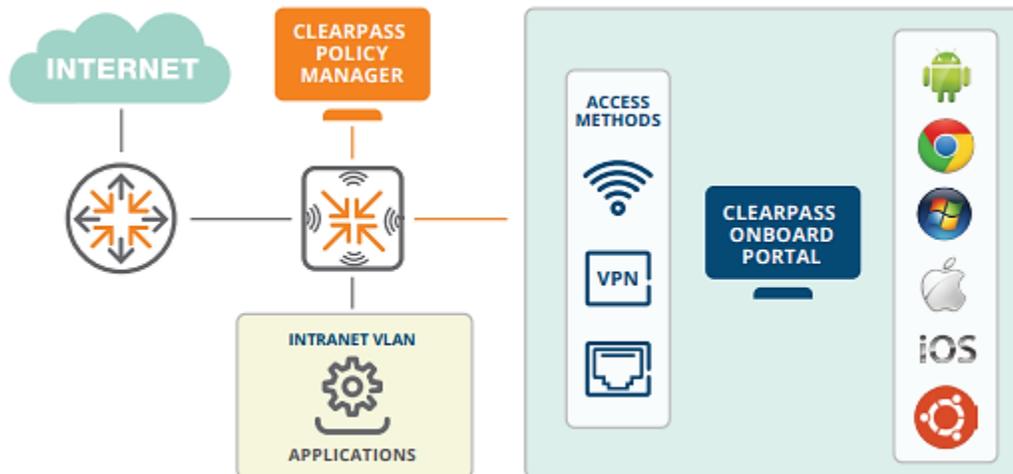
Aruba Central simplifies management of switches, wireless APs, and branch controllers

Instant access points include a free local management interface that provides visibility into the network and all capabilities to manage the Instant network. Aruba Central is a cloud-based software-as-a-service subscription that simplifies the management of Aruba Instant, switches, and mobility controllers. It offers remote monitoring and troubleshooting, centralized configuration and firmware management, compliance reporting, and Zero Touch Provisioning. The Aruba Central cloud management platform offers IT administrators a simple, secure, and cost-effective way to manage their Aruba Instant wireless APs, switches, and branch controllers. Visitors, contractors, and suppliers can access guest Wi-Fi through a customizable enterprise-grade portal.

Without additional network management hardware and software to install and maintain, network administrators can take advantage of an OPEX-friendly cloud subscription that includes automatic firmware updates, full technical support, and PCI compliance reporting. Aruba Instant APs, switches, and branch controllers can be shipped directly to remote sites where anyone without technical expertise can simply power them up and connect to the network. Configuration is automatically pushed from Central so the network is up and running in minutes.

Aruba Central mobile application for iOS ensures that monitoring the entire network, receiving notifications, and provisioning an AP, switch, or branch controller with easy barcode scanning is easy. Aruba Central also allows you to tap into smart network and presence analytics to better understand network usage, improve operations, and generate new revenues by enhancing customer engagement and making intelligent merchandising, layout, and staffing decisions.

## Aruba ClearPass



Automating device provisioning for secure BYOD with ClearPass Onboard

Aruba ClearPass has a low overhead model in terms of equipment deployment. The all-in-one ClearPass appliance is sold as either a hardware appliance or as a virtual machine. As a form factor, a hardware appliance ships in an all-in-one box, with all the software features on the same appliance.

The number of endpoints determines the number of appliances a business will want to purchase, although Aruba advises the necessary number of appliances, plus-one appliance, to serve as a backup and provide redundancy.

Managing the onboarding of personal devices for bring-your-own devices (BYOD) deployments can put a strain on IT and help desk resources, and can create security concerns. ClearPass Onboard lets users safely configure devices for use on secure networks all on their own. Device-specific certificates even eliminate the need for users to repeatedly enter login credentials throughout the day. Using ClearPass Onboard, the IT team defines who can onboard devices, the type of devices they can onboard, and how many devices per person. A built-in certificate authority lets IT support personal devices more quickly as an internal public key infrastructure (PKI), and subsequent IT resources are not required.

With secure network access control (NAC), customers can achieve an assured security posture without additional resources. Aruba ClearPass empowers IT teams with:

- Visibility into all connecting and connected devices, wired and wirelessly
- Control of Internet of Things (IoT) devices, BYOD, and corporate devices, across multiple network vendors
- Response with seamless integration of security tools for automated threat detection, escalation, and unified policy enforcement

The Aruba NAC solution, ClearPass, provides a single RADIUS-based security and verification point for all wired and wireless networks, applications, IoT devices, employee, contractor, and guest devices. ClearPass allows organizations to create, define and enforce a consistent access policy of what can connect to which elements of the network, based on the type of device, who is using it, where and when it is being used, the type of connection and its health status.

## Learning check

1. What do you use to administer StoreVirtual devices running LeftHand OS?
  - a. InfoSight
  - b. HPE OneView
  - c. Smart Storage Administrator
  - d. Centralized Management Console
2. What do SNMP programs use to quickly detect and identify Layer 2 network topology changes?
  - a. SNMP v2c
  - b. Client-side JavaScript
  - c. LLDP information stored in MIBs
  - d. Endpoint Admission Defense (EAD)
3. All current HPE network devices support LLDP.
  - True
  - False

## Summary

- HPE offers management tools designed to manage the entire system:
  - iLO 5 and iLO Federation
  - HPE OneView
- HPE tools designed to manage and monitor storage systems are specific to the product line:
  - InfoSight is designed for managing HPE Nimble storage arrays.
  - SSA is designed for managing the storage controllers in a ProLiant server.
  - Centralized Management Console is designed for managing StoreVirtual arrays.
  - HPE Storage Management Utility is the management GUI for HPE MSA storage.
  - Library and Tape Tools is designed for managing HPE tape storage and magneto-optical storage products.
- 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.
  - Smart-managed switches provide basic capabilities and can be managed through a web browser interface.
  - Unmanaged switches provide basic Layer 2 switching and are not configurable.
- IMC is a centralized network management platform that allows you to manage physical networks (both wired and wireless) and virtual networks.



---

# Troubleshooting HPE SMB Solutions

## Module 5

### Learning objectives

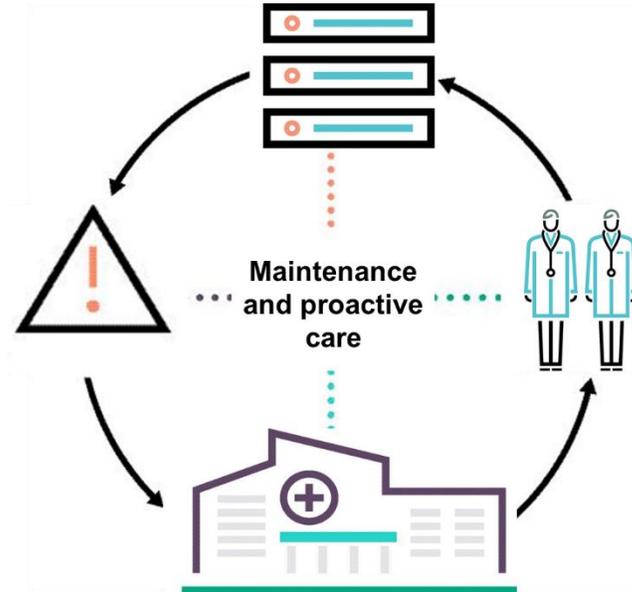
After completing this module, you should be able to:

- Summarize the basic approach to troubleshooting
- Recall the steps in the Hewlett Packard Enterprise (HPE) six-step troubleshooting methodology
- List additional methods and tools used in troubleshooting efforts
- Describe recommended solutions to common problems involving:
  - HPE servers
  - HPE storage
  - HPE networks

## Customer scenario

Consider the following fictitious customer scenario as you work through this module. Use it as a model when considering how to plan and design solutions to meet customers' business and technical objectives. Refer back to it as you engage with activities and labs throughout this course.

### IDJT HealthCare



IDJT HealthCare serves physicians and caregivers across medical focus areas

IDJT HealthCare understands the importance of maintenance and proactive care to avoid downtime. They would like to gather data proactively from the servers that run their critical business applications. In order to be prepared for a worst-case scenario, they want to simulate a situation in which a server is not functioning as intended, and an application is unavailable. They want to ensure that they know how to use the available tools to troubleshoot and collect important health and diagnostic data from the data center components.

You need to demonstrate to the customer how to use these tools to manage and maintain their equipment.

## Gauge your knowledge

Before proceeding with this section, think about your existing knowledge of the topics covered in this module by answering the following questions. Record your answers in the space provided and be prepared to discuss your responses with the group.

1. Which troubleshooting methodologies have you used in the past? How successful have they been?

---

---

---

---

2. Which tools have you used to diagnose system issues? What have been the challenges and rewards?

---

---

---

---

3. What experience do you have sending trouble tickets to HPE? What insights can you share about the process?

---

---

---

---

## Introduction to troubleshooting

Troubleshooting is a process that involves knowledge, experience, and intuition. An important part of this process is the use of proactive measures such as record keeping. The goal of this module is to promote a structured approach to troubleshooting that highlights common system administration trouble spots.

Some people address problems by just "jumping in" and trying things. However, unless you have experienced or solved the problem before, it is best to gather information before acting. This can be accomplished through testing to see if the problem can be replicated, looking at log files, and talking to others who were present when the problem occurred. A simple problem might not require documentation of your troubleshooting steps. However, as problems become more and more complex, documenting your actions becomes crucial.

Remember that what is working can be just as important as what is not working. "Has it *ever* worked?" is a key question to ask. If not, there will likely be a different set of steps to follow to resolve the issue.

## Pretroubleshooting and replacement steps

Before beginning any maintenance, upgrade, or troubleshooting procedures, it is important to follow all product safety and operating instructions. Always refer to the documentation (printed or electronic) supplied with the product. Observe all warnings on the product and in the operating instructions to reduce the risk of bodily injury, electric shock, fire, and damage to the equipment.

---

**NOTE:** For important safety, environmental, and regulatory information, refer to *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products* by scanning the following QR code or right-clicking it to open the hyperlink.



[http://h20566.www2.hpe.com/hpsc/doc/public/display?sp4ts.oid=5316427&docLocale=en\\_US&docId=emr\\_na-c03471072](http://h20566.www2.hpe.com/hpsc/doc/public/display?sp4ts.oid=5316427&docLocale=en_US&docId=emr_na-c03471072)

---

HPE field subassemblies for specific computer product line equipment are designed to be repaired in the field when possible, when failed parts can be isolated. Many servers and storage devices are designed with field-replaceable units (FRUs) that can be replaced in the field by qualified personnel trained by HPE. These FRUs are not designed to be replaced by customers. For details, refer to the maintenance and service guide and the user guide for the system that needs replacing.

## Verifying that a current backup exists

Before starting any procedures, you should verify that a current backup of the system data exists. After a full system backup has been performed, it should be verified with the existing data on the system. All backup software enables you to verify the backed up copy with the original data on the system. Confirm that the verification process has completed successfully before proceeding.

A basic approach to improve backup 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.

## Graceful power-down and power-up sequences

During unplanned power outages, HPE highly recommends following the proper power-up sequence to resume normal operations. Improper power down, power up, and restart sequences of SAN storage and attached hosts will affect the normal functionality.

The general sequence for scheduled power-down, power-up, or restart operations for maintenance and upgrades is as follows:

- Powering down sequence:
  - a. Host servers
  - b. Storage devices (libraries, storage, and so forth)
  - c. SAN switches
- Powering up sequence:
  - a. SAN switches
  - b. Libraries (take a long time to initialize)
  - c. Other storage devices (disk arrays, NAS, and so forth)
  - d. Host servers (only when all other devices are up)

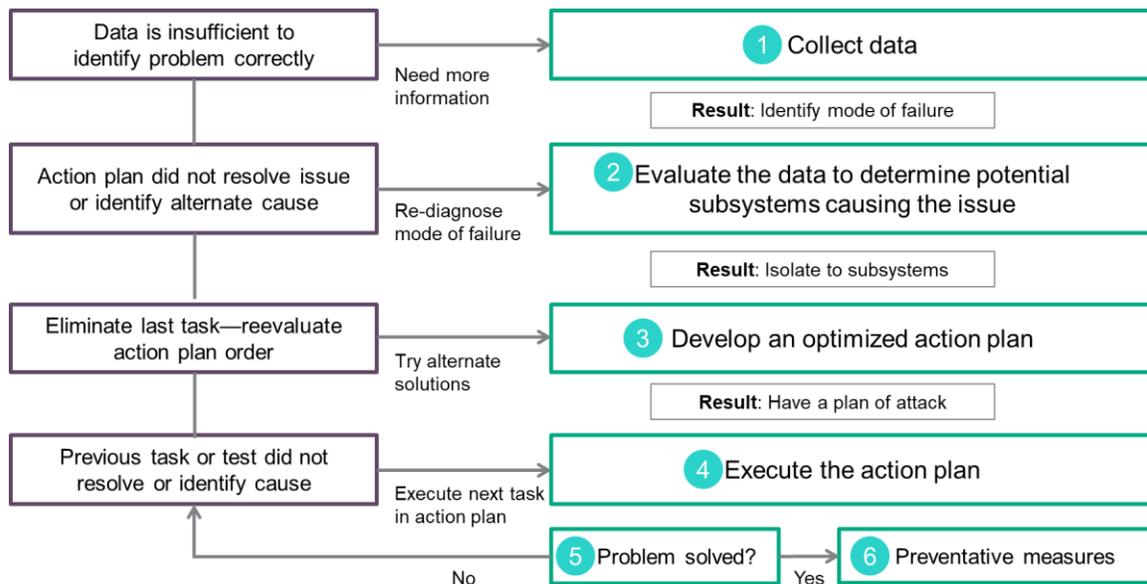
These are the steps to completely shut down a data center:

1. Shut down all application servers such as development and staging servers, Microsoft SharePoint portal server, enterprise resource planning (ERP) software related servers, and so forth.
2. Shut down all servers such as database servers, Exchange server, all virtual machines, and infrastructure management server.
3. Shut down all VMware hosts, then domain controllers and Dynamic Host Configuration Protocol (DHCP) servers.
4. Shut down all storage solutions such as libraries, disk arrays, and NAS; then shut down all related appliances.
5. Shut down the network, routers, and firewalls including SAN switches, gateways, and so forth.
6. Lastly, shut down the uninterruptible power supply (UPS).

These are the steps to power up a data center:

1. Power on the UPS.
2. Power on the network, routers, and firewall including SAN switches, gateways, and so forth.
3. Power on all storage solutions such as libraries, disk arrays, or NAS solutions. Wait for a few minutes for the disks to spin up and the libraries to initialize completely.
4. Power on the domain controller and DHCP servers. Test the domain and DHCP service before powering on other servers.
5. Power on all host servers such as database servers and then application servers.

## HPE six-step troubleshooting methodology



HPE troubleshooting flowchart

This graphic shows the basic six steps of the troubleshooting process recommended by HPE. The high degree of interaction between the system, hardware options, operating system, and application software can make it difficult to isolate the root cause of a problem. Intermittent problems and issues caused by multiple subsystem malfunctions can be especially difficult to troubleshoot.

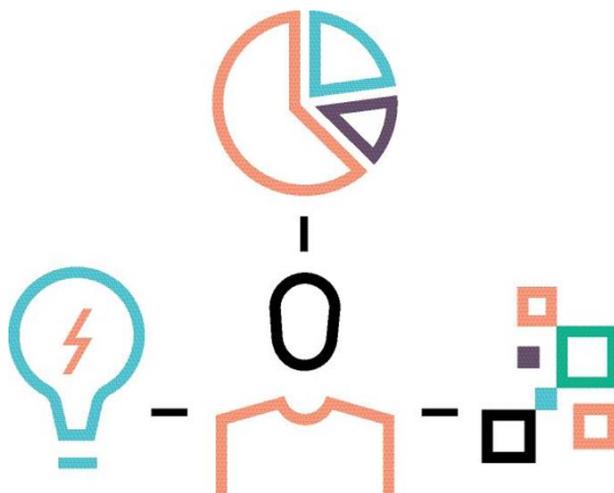
Using a methodical approach to troubleshooting helps to systematically identify the core of a problem, resolve it, and take steps to prevent it from reoccurring. The six steps recommended by HPE are:

1. Collect data.
2. Evaluate the data to determine potential subsystems causing the issue.
3. Develop an optimized action plan.
4. Execute the action plan.
5. Determine whether the problem is solved.
6. Implement preventive measures.

As part of any change management process, altering only one variable at a time can show the impact of that change. If manipulating one variable does not successfully restore the system to normal operation, revisiting the data to determine the next variable to manipulate is the next step. By following a troubleshooting methodology, you can use a standard approach to eliminate possible causes until a solution is found.

This methodology provides a logical framework that can be used to troubleshoot system problems and reach problem resolution. A logical framework provides a consistent and solid foundation for other technicians and system engineers to work from when escalation is necessary.

## Step 1: Collect data



Gather as much information as possible

Troubleshooting a problem involves spending time and effort gathering helpful information. Collecting data about the reported problem requires asking the right questions to understand what failed and in what context. Continue to ask questions to learn as much detail as possible.

To arrive at an accurate problem description, you need to:

- Determine and use the most appropriate tools for each situation
- Understand how the system reacts in a failure scenario

Collecting data includes:

- Identifying the hardware and software components in the system
- Asking questions to understand what failed and in what context
- Continuing to ask questions to learn as much detail as possible
- Gathering information about the failure such as:
  - Failure conditions from LED indicators
  - Power-on self-test (POST) messages
  - Diagnostic, test, and log reports
  - Critical error log messages
  - Remote log messages
  - Stop, abnormal end of task, and trap messages
  - Status and error reports from systems management tools
  - Status information reported by management agents
- Organizing the collected data

## Collecting system data

To identify the source of the problem and suggest a solution, you must gather data from hardware subsystems and the operating system. OEMs provide many tools to help you collect data.

The vendor-specific tools listed here are examples of the ways to collect system data.

- **HPE**—HPE provides several methods for viewing system data. Some of these utilities are integrated into the server, and all complement each other.
  - **HPE OneView**—HPE OneView supports inventory management, health monitoring, alerting, and reporting capabilities on ProLiant servers. Predictive analytics can be added to let you troubleshoot rapidly by using HPE Operations Analytics for HPE OneView.
  - **HPE Insight Remote Support**—HPE strongly recommends that customers register their devices for remote support to enable enhanced delivery of an HPE warranty, support services, or contractual support agreement. Insight Remote Support supplements on-premise continuous monitoring to ensure maximum system availability. It provides intelligent event diagnosis and automatic, secure submission of hardware event notifications to HPE, which initiates a fast and accurate resolution based on the product's service level. Notifications can be sent to an authorized HPE channel partner for onsite service, if configured and available in that country. Insight Remote Support is available as part of an HPE warranty, support services, or contractual support agreement.
  - **HPE Active Health System**—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. As an essential component of the HPE integrated Lights-Out (iLO) Management Engine portfolio, Active Health System provides:
    - Continuous health monitoring of more than 1600 system parameters
    - Logging of all configuration changes
    - Consolidated health and service alerts with precise time stamps
    - Agentless monitoring that does not affect application performance
  - **HPE System Management Homepage**—This web-based interface provides a consolidated view of all system health and configuration information and simplifies access to HPE web-enabled management software running on ProLiant servers.
  - **HPE Smart Storage Administrator (SSA)**—SSA is an advanced utility that provides high-availability configuration, management, and diagnostic capabilities for all HPE Smart Array products.
  - **Integrated Management Log (IML)**—This is a record of significant events that occur during system operation. iLO for ProLiant captures and stores the server's IML for access through a browser or command line, even when the server is not operational. The IML time stamps each event with one-minute granularity. This capability can be helpful when troubleshooting remote host server problems.
  - **Insight Diagnostics (online and offline editions)**—This proactive server management tool gathers critical hardware and software information and provides comprehensive server configuration information.

---

**NOTE:** For more information or to download the Insight Diagnostics utility, visit the HPE website: <http://www.hpe.com/info/InsightDiagnostics>. The Insight Diagnostics Online Edition is available from the HPE Service Pack for ProLiant (SPP) site: <http://h17007.www1.hpe.com/docs/enterprise/servers/gen9/tsg/239280.htm>.

---

- **Linux**—Red Hat Enterprise Linux (RHEL) and SUSE Linux Enterprise Server offer monitoring tools for obtaining information about the running kernel:
  - **iostat**—Measures I/O and processor statistics for devices and partitions
  - **vmstat**—Reports virtual memory information
  - **netstat**—Monitors network information
  - **top**—Measures top processor processes
  - **sar**—Measures system activity information and compiles cumulative statistics over a specified period of time

RHEL log files are located in the `var/log` directory. An especially good resource for hardware troubleshooting is the kernel startup file, which is named `boot.log`. This plaintext file can be opened with any text editor. It lists the devices on the system and their status when the system booted. Another useful file is `/var/log/messages`, which includes boot messages and current system log information.

- **Microsoft**—Windows Server 2008 and 2012 feature tools for event management, monitoring and alerting, reporting, and trend analysis:
  - **System Monitor**—This tool is used to monitor local or remote system performance and to set counters as needed. You can collect and view extensive data about the usage of hardware resources and the activity of system services on computers you administer. It provides a general, customizable mechanism to view various counters and other metrics within a running application.
  - **Performance Monitor**—This visualization tool allows you to view performance data, both in real time and from log files. You can also examine performance data in a graph, histogram, or report.
  - **Task Manager**—This tool provides information about programs and processes running on a computer.

## Step 2: Evaluate data to determine subsystems causing the issue

After you collect data and identify the symptoms, evaluate the facts in order to:

- Determine which components might have caused what happened
- Isolate faults to a hardware or software subsystem
- Understand the mode of failure

## Step 3: Develop an optimized action plan

After collecting the facts and isolating the specific mode of failure, perform these actions to develop an optimized action plan:

- Identify specific root causes for the specified mode of failure.
- Identify possible solutions for each possible root cause.
- Prioritize the solutions by balancing the time and cost needed to implement each solution against the likelihood that each solution will fix the issue. Alternatively, you can prioritize the solutions based on the potential value of the information gained if each solution is inadequate.
- Identify the steps necessary to implement each solution.
- Compile all the steps into an optimized action plan by eliminating redundancy and ensuring that only one variable is being manipulated at a time.
- Inform and update the customer about the action plan.
- Incorporate an escalation strategy into the action plan and be prepared to incorporate technical assistance.
- List the order of people or groups to contact and the information needed from each.

## Step 4: Execute the action plan

Implement the written optimized action plan by following these steps:

1. Execute each step completely, as agreed on by the customer.
2. Apply only one solution or change only one variable at a time.
3. Carefully observe and record the results of each step including any error messages or changes in functionality.

If your action plan did not solve the problem, it might:

- Result in enough information to solve the issue
- Provide enough pertinent information to go back to Step 2

With a newly created action plan, you might be able to find the root cause of the problem and reach a solution.

---

**IMPORTANT:** To keep your customers involved, you must ensure that they are informed of the steps taken and the actions planned.

---

## **Step 5: Determine whether the problem is solved**

Observe and evaluate the results of each step in each solution until the problem has been isolated and resolved. If the problem is not resolved:

1. Collect more data.
2. Evaluate the information gathered from the implementation of the action plan.
3. Develop another optimized action plan.
4. Implement the new optimized action plan.

Repeat these steps until you resolve the problem. As additional information is gathered, new action plans can be optimized, executed, and evaluated.

## **Step 6: Implement preventive measures**

As soon as the problem is resolved, look for opportunities to implement preventive measures to stop the problem from occurring again. Find other ways to improve or increase availability.

To implement preventive measures:

1. Identify the root cause of the problem.
2. Determine proactive steps that can prevent the problem from reoccurring.
3. Devise a system test to verify changes and procedures before implementing them.
4. Implement a new set of procedures, software processes, and administrative maintenance guidelines to attain a higher level of availability.
5. Perform preventive maintenance, including reseating boards and securing proper airflow.
6. Add fault-tolerant elements to critical subsystems, where applicable.

A carefully planned system software maintenance strategy maximizes server stability and availability. By developing well-regulated system software baselines for business servers, you can reduce the time required to update or troubleshoot existing servers and ensure that new servers are set up with tested and stable software configurations. Also, an appropriate service agreement for hardware and software can minimize downtime.

## Learning check

1. What is the proper sequence for powering down a data center? Put the steps in the correct order.

Order	Step
	Shut down the UPS.
	Shut down the network, routers, and firewalls.
	Shut down all database servers and virtual machines.
	Shut down all VMware hosts, and then domain controllers and DHCP servers.
	Shut down all storage solutions and related appliances.
	Shut down all application servers.

2. Match the HPE tool for collecting system data with its function.

Tool	Function
HPE OneView	Monitors and records changes in the server hardware and system configuration
Insight Diagnostics	Provides high-availability configuration, management, and diagnostic capabilities for Smart Array products
Smart Storage Administrator	Gathers critical hardware and software information and provides comprehensive server configuration information
Integrated Management Log	Supports inventory management, health monitoring, alerting, and reporting capabilities on ProLiant servers
Active Health System	Records significant events that occur during system operation

## Additional troubleshooting methods and tools

In addition to the six steps included in the HPE approach to troubleshooting, there are several other methods that can be leveraged regardless of the complexity of the issue. Effective troubleshooting methods include establishing a baseline, using published server information, collecting problem and system data, using online and offline diagnostic tools, and checking status LEDs. These strategies can be extremely effective in reaching a resolution.

Diagnostic tools are typically used during troubleshooting because they simplify the process of effectively identifying, diagnosing, and isolating hardware issues. Test features often let you test the functionality of all the major hardware components in the system. Depending on the tool, tests can be interactive or unattended and vary in duration and completeness.

Diagnostic tools fall into two main categories:

- **Online**—Executed under the control of a running operating system
- **Offline**—Executed after shutting down the operating system

Offline diagnostic tools are normally more exhaustive because they can test all system components without avoiding components that would otherwise be in use by the running operating system.

### Establishing a baseline

Before you can begin to identify a system failure, you must understand how the system should operate under normal circumstances. Recognizing when a system is performing normally and understanding what a system requires to operate properly will help you to identify a malfunctioning system.

Creating a performance baseline is essential in determining normal operating parameters for a server. A baseline is a set of critical data used as a control. It allows you to capture the current state of the server or network performance and use this information as the starting point for comparing future performance levels.

---

**IMPORTANT:** This baseline should be captured after initial installation of the equipment, during normal daily operation. The baseline should be updated after any hardware or software additions to the system.

---

Establishing the baseline will help you recognize the warning signs of a failing component. These signs or symptoms might include consistent or intermittent error codes, loss of functionality, or a change in the time required to perform a task.

Among other things, you should understand:

- When, why, and which LEDs illuminate
- In what order system components power up, especially when external storage enclosures are configured. The power cycle order should normally be:
  - Power up—External storage enclosures first, then server
  - Power down—Server first, then external storage enclosures
- The boot load order of files

Failing to understand how a system or subsystem operates can lead to unnecessary parts replacement, software upgrades that might not address the problem being reported, wasted time and effort, unnecessary downtime, and reduced customer satisfaction.

## Resolving common problems

Complications range from the most basic connector issues to complex software configuration problems. Despite the level of complexity, there are several common procedures and solutions that apply to a variety of troubleshooting levels for ProLiant servers.

Tasks that might address common issues include:

- **Checking for loose connections**—To confirm that this is not the source of the problem, there are a few steps you can take, including:
  - Be sure all power cords are securely connected.
  - Be sure all cables are properly aligned and securely connected for all external and internal components.
  - Remove all data and power cables, and check for damage. Be sure no cables have bent pins or damaged connectors.
  - If a fixed cable tray is available for the server, be sure the cords and cables connected to the server are correctly routed through the tray.
  - Be sure each device is properly seated.
  - If a device has latches, be sure they are completely closed and locked.
  - Check any interlock or interconnect LEDs that could indicate a component is not connected properly.
  - If problems continue to occur, remove and reinstall each device, checking the connectors and sockets for bent pins or other damage.
- **Referring to service notifications**—Locate the latest service notification on the manufacturer's website or find another process recommended by the manufacturer.

## Corrupted ROM disaster recovery

Another frequently experienced issue is a corrupted ROM. If the system ROM of a ProLiant ML, DL, or BL server is corrupted, the system automatically switches to the redundant ROM in most cases. If the system does not automatically switch to the redundant ROM, perform the following steps:

1. Power down the server.
2. Remove the server from the rack, if necessary.
3. Remove the access panel.
4. Change positions 1, 5, and 6 of the system maintenance switch to **on**.
5. Install the access panel.
6. Install the server into the rack or enclosure.
7. Power up the server.
8. After the system beeps, repeat steps 1 through 3.
9. Change positions 1, 5, and 6 of the system maintenance switch to **off**.
10. Repeat steps 5 and 6.
11. If both the current and backup versions of the ROM are corrupt, return the system board for a service replacement.

## Using published server information

There is an extensive amount of information, resources, and documents you can use to expand your troubleshooting efforts. This should not be considered a comprehensive guide that can address all system administrator problems or provide links to all the necessary tools. However, this information can provide an effective starting point.

Depending on your situation, an online or a printed reference might be useful for finding a part number, upgrade information, or other details. The following types of information from online and printed publications can be helpful to you:

- **Illustrated parts catalogs**—Provide an illustrated reference for spare parts and a comprehensive list of spare part numbers specific to that machine
- **POSTs**—Describe the internal system diagnostic programs that are executed automatically when you power on the system
- **Error messages and codes**—List POST and diagnostics error codes, the required course of action to resolve the problem described by each error code, and LEDs and their meanings
- **LED indicators**—Contain graphics and tables showing LED locations and functions for the system interlock status; hot-pluggable I/O fans, power supplies, and SCSI hard drives; PCI hot plug components; and system I/O boards
- **Specifications**—Provide operating and performance specifications for the specific HPE computer for which a particular guide is developed

## HPE resources for troubleshooting

Troubleshooting resources are available for HPE products in the following documents:

- The *Troubleshooting Guide for HPE ProLiant Gen10 servers* provides procedures for resolving common problems and comprehensive courses of action for fault isolation and identification, issue resolution, and software maintenance.

---

**NOTE:** To access the troubleshooting guide, refer to the Hewlett Packard Enterprise Information Library by scanning this QR code or right-clicking it to open the hyperlink.



<http://www.hpe.com/info/gen10-troubleshooting>

---

- The *HPE iLO 5 User Guide* provides more information on IML, including the steps to view the IML and other details.

---

**NOTE:** To access the iLO 5 guide, scan this QR code or right-click it to open the hyperlink.



[http://h20565.www2.hpe.com/hpsc/doc/public/display?docId=a00018324en\\_us](http://h20565.www2.hpe.com/hpsc/doc/public/display?docId=a00018324en_us)

---

You can access additional documentation from the following locations:

- **Hewlett Packard Enterprise Information Library**—Select the product, subcategory, or information type to narrow the list to fewer documents.
- **Product-specific information library**—Refer to the Start Here document that ships with the product.
- **HPE QuickSpecs**—For more information about product features, specifications, options, configurations, and compatibility, refer to the QuickSpecs document for each product.

---

**NOTE:** To access QuickSpecs, scan this QR code or right-click it to open the hyperlink.



<http://h41370.www4.hp.com/quickspecs/overview.html>

---

- **White papers**—White papers are electronic documents that provide information on complex technical topics. Some white papers contain in-depth details and procedures. Topics include HPE products, HPE technology, operating systems, networking products, and performance.

---

**NOTE:** White papers are available from the Hewlett Packard Enterprise Information Library and other areas of the HPE website.

---

- **Email subscription services**—Receive support alerts, product communications, driver updates, software releases, firmware updates, and customer replaceable component information by signing up for relevant emails through the HPE website.

---

**NOTE:** To see the subscriptions options, scan this QR code or right-click it to open the hyperlink.



<https://h41360.www4.hpe.com/>

---

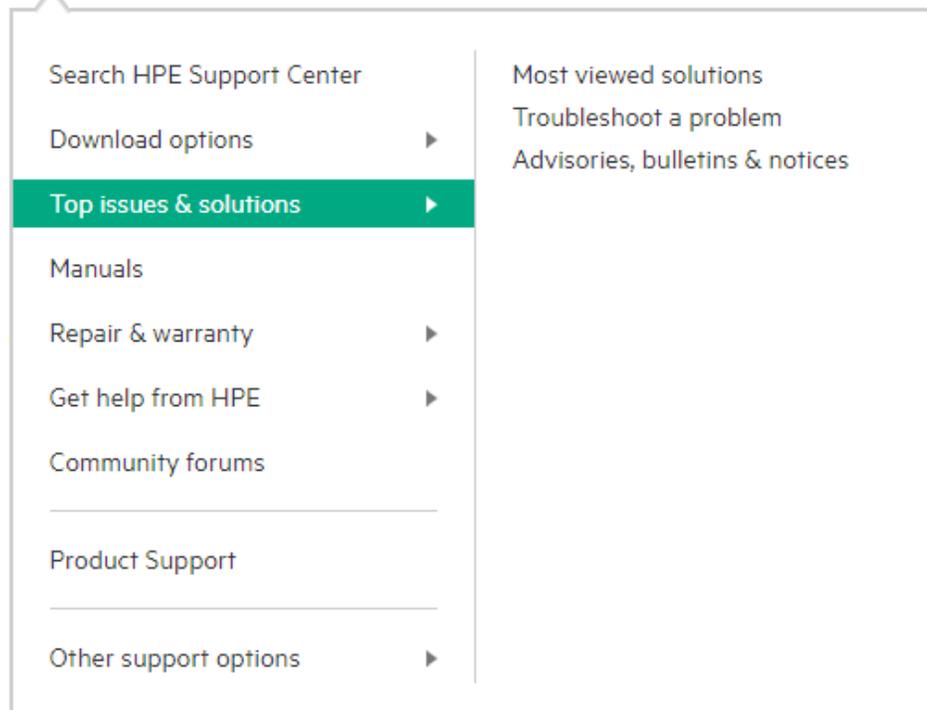
## HPE Support Center

### Hewlett Packard Enterprise Support Center

Insight Online

Product Support ▼

My IT Environment ▼



HPE Support Center Product Support menu

The HPE Support Center website is a repository for troubleshooting tools and information, as well as the latest drivers and flash ROM images. This site also provides access to product manuals and warranty information.

From the Product Support drop-down menu, you can choose to see the solutions most often viewed; find information about troubleshooting an issue with the product you have selected; and view advisories, bulletins, service notices, and compatibility tables.

**NOTE:** To access the HPE Support Center website, scan this QR code or right-click it to open the hyperlink.



[www.hpe.com/support/hpesc](http://www.hpe.com/support/hpesc)

## USB support

HPE provides external USB support to enable local connection of USB devices for server administration, configuration, and diagnostic procedures. For additional security, external USB functionality can be disabled through USB options in Unified Extensible Firmware Interface (UEFI) System Utilities.

HPE servers support USB 2.0 ports and USB 3.0 ports. Both port types support installing all types of USB devices (USB 1.0, USB 2.0, and USB 3.0), but might run at lower speeds in specific situations:

- USB 3.0 capable devices operate at USB 2.0 speeds when installed in a USB 2.0 port.
- In UEFI Boot Mode, HPE provides legacy USB support in the preboot environment before the operating system loading for USB 1.0, USB 2.0, and USB 3.0 speeds.
- In Legacy BIOS Boot Mode, HPE provides legacy USB support in the preboot environment before the operating system loading for USB 1.0 and USB 2.0 speeds. USB 3.0 ports can be used with all devices in Legacy BIOS Boot Mode but are not available at USB 3.0 speeds in the preboot environment. Standard USB support (USB support from within the operating system) is provided by the operating system through the appropriate USB device drivers. Support for USB 3.0 varies by operating system.

For maximum compatibility of USB 3.0 devices with all operating systems, HPE provides a configuration setting for USB 3.0 Mode. Auto is the default setting. This setting impacts USB 3.0 devices when connected to USB 3.0 ports in the following manner:

- **Auto (default)**—If configured in Auto Mode, USB 3.0 capable devices operate at USB 2.0 speeds in the preboot environment and during boot. When a USB 3.0 capable operating system USB driver loads, USB 3.0 devices transition to USB 3.0 speeds. This mode is compatible with operating systems that do not support USB 3.0 but allows USB 3.0 devices to operate at USB 3.0 speeds with state-of-the-art operating systems.
- **Enabled**—If enabled, USB 3.0 capable devices operate at USB 3.0 speeds at all times (including the preboot environment) when in UEFI Boot Mode. Do not use this mode with operating systems that do not support USB 3.0. If operating in Legacy Boot BIOS Mode, the USB 3.0 ports cannot function in the preboot environment and are not bootable.
- **Disabled**—If configured for Disabled, USB 3.0 capable devices function at USB 2.0 speeds at all times.

The preoperating system behavior and default operation of USB ports is configurable in UEFI System Utilities.

---

**NOTE:** For more information, review the UEFI System Utilities User Guide for HPE ProLiant Gen10 Servers and HPE Synergy by scanning this QR code or right-clicking it to open the hyperlink.



[http://h20566.www2.hpe.com/hpsc/doc/public/display?sp4ts.oid=6935826&docLocale=en\\_US&docId=emr\\_na-a00016407en\\_us](http://h20566.www2.hpe.com/hpsc/doc/public/display?sp4ts.oid=6935826&docLocale=en_US&docId=emr_na-a00016407en_us)

---

## Learning check

1. When should a performance baseline be captured?

---

---

---

---

2. If the system ROM of a ProLiant ML, DL, or BL server is corrupted, and the system does not automatically switch to the redundant ROM, what should you do first?
  - a. Remove the access panel.
  - b. Change positions 1, 5, and 6 of the system maintenance switch to **on**.
  - c. Remove the server from the rack.
  - d. Power down the server.

## Troubleshooting HPE ProLiant servers

HPE provides procedures for resolving common problems and comprehensive courses of action for fault isolation and identification, error message interpretation, issue resolution, and software maintenance on ProLiant servers and server blade models. In addition, problem-specific flowcharts can help you navigate complex troubleshooting processes.

In addition to troubleshooting issues with hardware, some applications provide device management capabilities that consolidate and integrate management data from server and third-party devices. Some management applications offer consolidated views of system health indicators, logs, tools, and diagnostics from multiple sources, which can simplify the gathering of data for troubleshooting.

Web-based management applications allow system administrators to perform normal administrative tasks from any remote location using a web browser. These tools include monitoring capabilities and logs that can be used in troubleshooting both local and remote servers.

### System configuration information

System configuration information is an important element of troubleshooting. You can obtain this information in several ways. Your choice could be determined by the state of the machine, such as whether the operating system is available.

Two options are the manufacturer's ROM-based setup and configuration utilities. These tools usually provide the ability to view the system configuration. Management applications and diagnostic tools from the server manufacturer and operating system vendor are frequently used and likely to offer more features.

When troubleshooting an issue, check the firmware revisions and hardware configuration of the server. In general, all firmware in the server should be at the latest revision. Use tools to look for devices that have firmware revisions with known issues, which can then be upgraded. The hardware configuration can be checked to ensure that all hardware is configured properly for its intended use.

### Useful server configuration information

Basic server configuration information that might be useful when troubleshooting includes:

- Explanation of the issue, the first occurrence, and frequency
- Any changes in hardware or software configuration before the issue surfaced
- Third-party hardware information
  - Product name, model, and version
  - Company name
- Specific hardware configuration
  - Product name, model, and serial number
  - Number of processors and speed
  - Number of DIMMs and their size and speed
  - List of controllers and NICs
  - List of connected peripheral devices
  - List of any other optional HPE hardware
  - Network configuration

- Specific software information
  - Operating system information
  - List of third-party and server manufacturer's software installed
  - Verification of latest drivers installed
  - Verification of latest ROM/BIOS/UEFI version
  - Verification of latest firmware on array controllers and drives
- Results from attempts to clear nonvolatile RAM (NVRAM)

## Event logs

Logs are important tools that can be used to help isolate a problem. Events are unusual system occurrences. Event logs are designed to record all server-specific events by tracking event details recorded by the health drivers, such as operating system information, ROM POST codes, and errors during normal operations.

Logs are usually classified by three types: application logs, system logs, and security logs. Each log records events and errors for a postdiagnosis review, helping to identify server failures promptly. The information displayed for each log entry usually includes a description of the event, the number of times the event has occurred, the initial time and date the event occurred, and the time and date the event was last updated. The types of event and severity levels are important information to view. It is important to quickly identify a critical, failed, or degraded component. Event types include:

- Information
- Warning
- Error
- Success or OK
- Failure

Logs can vary in the level and format of detail provided. Events can sometimes be filtered, for example, by severity, class, or count. Aside from including detailed error information, some logs recommend repair actions to assist troubleshooting. The ability to mark repaired items and enter maintenance notes is sometimes an option depending on the log. Other logs integrate logs from different sources to simplify troubleshooting. Some tools can use logs to generate reports or alerts based on a list of critical events.

Logs can be reviewed easily through an event viewer. An example of a simple logged event follows:

```
**001 of 010**  
---caution---  
03/19/2016  
12:54 PM  
FAN INSERTED  
Main System  
Location:  
System Board  
Fan ID: 03  
**END OF EVENT**
```

## Beep codes

Many servers provide a series of audio beeps to alert you to a problem. The number, duration, and sequence of the beeps indicate the problem. For example, after installing a new processor, the system might emit one long and one short beep, indicating an unsupported processor.

---

**NOTE:** For error message details, check the *HPE ProLiant Gen9 Troubleshooting Guide: Volume II—Error Messages*, scan this QR code or right-click it to open the hyperlink.



<http://h20566.www2.hpe.com/hpsc/doc/public/display?docId=c04443553>

---

**NOTE:** Server blades do not have speakers and thus do not support audio output.

---

## POST error messages

Diagnosing problems on servers can often be accomplished—if the system boots and the video is working—by interpreting POST messages. These diagnostics, which reside in the BIOS ROM, isolate server-related logic failures and indicate the component that needs to be replaced by displaying an error message. Most server hardware failures are accurately isolated during POST. The number of tests displayed depends on the configuration of the server.

Some messages are only informational and do not indicate any error. When POST detects a system failure, it displays a POST error message or emits a series of beep codes. All POST error messages and beep codes occur before an operating system has loaded. Therefore, a server generates only the codes that are applicable to its BIOS configuration and hardware options that are installed.

To read the POST error message on screen, the video subsystem of the server must be working. If the video is not working, use the beep codes to troubleshoot the cause of the failure during the server POST. A failed processor, memory, video controller, or system board can all cause a video failure during POST.

Whenever a nonfatal error occurs during POST, an error message describing the problem is presented on-screen. These text messages are displayed in normal video (white text on black background). It shows the details of the error. The following is an example of a POST error message:

```
Error message 1 of 1: Error code 0103
Keyboard not detected - Keyboard error
```

In some cases, an error message might include recommendations for troubleshooting or require that you press the **Enter** key to display recommendations. Follow the instructions on the screen.

---

**NOTE:** For more information about beep codes and POST error messages for ProLiant servers, scan this QR code or right-click it to open the hyperlink.



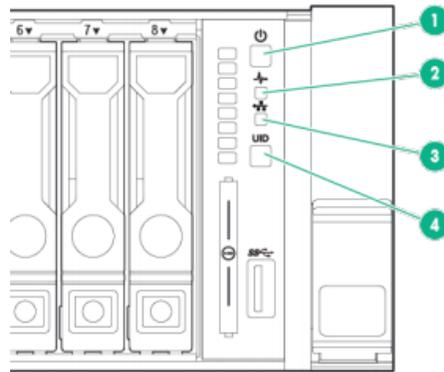
[http://h20565.www2.hpe.com/hpsc/doc/public/display?sp4ts.oid=502439&docId=emr\\_na-c01702138&docLocale=en\\_US](http://h20565.www2.hpe.com/hpsc/doc/public/display?sp4ts.oid=502439&docId=emr_na-c01702138&docLocale=en_US)

---

## Checking status LEDs

ProLiant servers provide status indicators and LEDs that report the system status. Viewing system health indicators should be a first step in troubleshooting. LEDs can assist you in monitoring the activity and the health of internal and external system components, and assist you in identifying and resolving problems. The type, number, and conditions of LEDs vary with the type of server.

### ProLiant Gen9 server front panel LEDs example



Front panel LEDs on a ProLiant DL380 Gen9 server

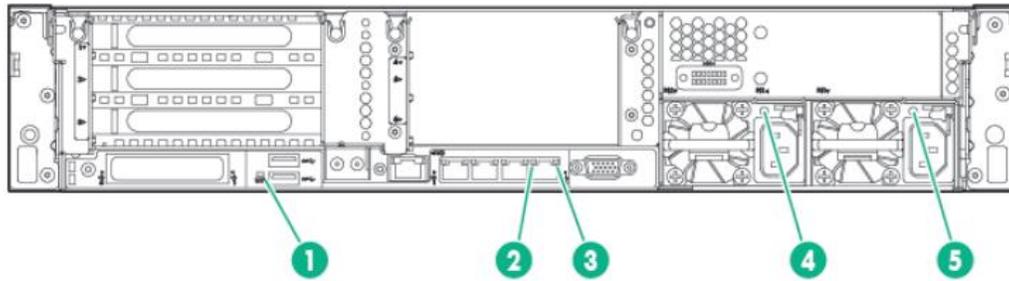
The graphic shows the location of the front panel LEDs on a ProLiant DL380 Gen9 server. The table describes each item.

Item	Description	Status
1	Power on/standby button and system power LED <sup>1</sup>	<ul style="list-style-type: none"> <li>• Solid green = System on</li> <li>• Flashing green (1 Hz/cycle per sec) = Performing power on sequence</li> <li>• Solid amber = System in standby</li> <li>• Off = No power present<sup>2</sup></li> </ul>
2	Health LED <sup>1</sup>	<ul style="list-style-type: none"> <li>• Solid green = Normal</li> <li>• Flashing green (1 Hz/cycle per sec) = iLO is rebooting</li> <li>• Flashing amber = System degraded</li> <li>• Flashing red (1 Hz/cycle per sec) = System critical<sup>3</sup></li> </ul>
3	NIC status LED <sup>1</sup>	<ul style="list-style-type: none"> <li>• Solid green = Link to network</li> <li>• Flashing green (1 Hz/cycle per sec) = Network active</li> <li>• Off = No network activity</li> </ul>
4	UID button/LED <sup>1</sup>	<ul style="list-style-type: none"> <li>• Solid blue = Activated</li> <li>• Flashing blue:                             <ul style="list-style-type: none"> <li>– 1 Hz/cycle per sec = Remote management or firmware upgrade in progress</li> <li>– 4 Hz/cycle per sec = iLO manual reboot sequence initiated</li> <li>– 8 Hz/cycle per sec = iLO manual reboot sequence in progress</li> </ul> </li> <li>• Off = Deactivated</li> </ul>

#### NOTES:

1. When all four LEDs flash simultaneously, a power fault has occurred.
2. Facility power is not present, power cord is not attached, no power supplies are installed, power supply failure has occurred, or power button cable is disconnected.
3. If the health LED indicates a degraded or critical state, review the system IML or use iLO to review the system health status.

## Back panel LEDs example



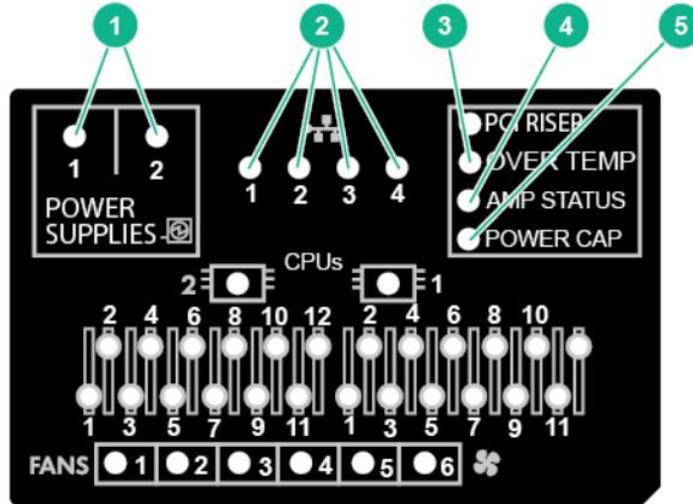
Rear panel LEDs on a ProLiant DL380 Gen9 server

The graphic shows the location of the rear panel LEDs on a ProLiant DL380 Gen9 server and the table describes each item.

Item	Description	Status
1	UID LED	<ul style="list-style-type: none"> <li>• Off = Deactivated</li> <li>• Solid blue = Activated</li> <li>• Flashing blue = System being managed remotely</li> </ul>
2	NIC link LED	<ul style="list-style-type: none"> <li>• Off = No network link</li> <li>• Green = Network link</li> </ul>
3	NIC activity LED	<ul style="list-style-type: none"> <li>• Off = No network activity</li> <li>• Solid green = Link to network</li> <li>• Flashing green = Network activity</li> </ul>
4	Power supply 2 LED	<ul style="list-style-type: none"> <li>• Off = System is off or power supply has failed</li> <li>• Solid green = Normal</li> </ul>
5	Power supply 1 LED	<ul style="list-style-type: none"> <li>• Off = System is off or power supply has failed</li> <li>• Solid green = Normal</li> </ul>

**NOTE:** For more information, refer to the maintenance and service guide for the specific server model.

## Systems Insight Display LEDs example



LEDs on a ProLiant DL380 Gen9 server

The Systems Insight Display LEDs represent the system board layout. The display enables diagnosis without having to remove the access panel.

The graphic shows the Systems Insight Display LEDs on a ProLiant DL380 Gen9 server. The following table describes each item numbered in the graphic.

Item	Description	Status
1	Power supplies	<ul style="list-style-type: none"> <li>• Off = Normal</li> <li>• Solid amber = Power subsystem degraded, power supply failure, or input power lost</li> </ul>
2	NIC link/activity	<ul style="list-style-type: none"> <li>• Off = No link to network; if the power is off, view the rear panel RJ-45 LEDs for status</li> <li>• Flashing green = Network link and activity</li> <li>• Solid green = Network link</li> </ul>
3	Over temp	<ul style="list-style-type: none"> <li>• Off = Normal</li> <li>• Solid amber = High system temperature detected</li> </ul>
4	AMP status	<ul style="list-style-type: none"> <li>• Off = AMP modes disabled</li> <li>• Solid green = AMP mode enabled</li> <li>• Solid amber = Failover</li> <li>• Flashing amber = Invalid configuration</li> </ul>
5	Power cap	<ul style="list-style-type: none"> <li>• Off = System is in standby or no cap is set</li> <li>• Solid green = Power cap applied</li> </ul>
N/A	All other LEDs	<ul style="list-style-type: none"> <li>• Off = Normal</li> <li>• Amber = Failure</li> </ul>

## Basic troubleshooting with HPE OneView

HPE OneView has a variety of troubleshooting tools you can use to resolve issues. By following a combined approach of examining screens and logs, you can obtain a history of activity and the errors encountered. Start troubleshooting in the following areas:

- The Activity screen displays a log of all changes made on the appliance, whether user-initiated or appliance-initiated. It is similar to an audit log, but with finer detail and it is easier to access from the UI. The Activity screen also provides a log of health alerts and status notifications.
- The audit log contains a record of actions performed on the appliance, which you can use for individual accountability. Monitor the audit logs because they are rolled over periodically to prevent them from getting too large. Download the audit logs periodically to maintain a long-term audit history. Each user has a unique logging ID per session, enabling you to follow a user's trail in the audit log. Some actions are performed by the appliance and might not have a logging ID.

---

**NOTE:** You must have infrastructure administrator privileges to download the audit log.

---

- Create a support dump file to gather logs and other information required for debugging into an encrypted, compressed file that you can send to an authorized support representative for analysis.

---

**IMPORTANT:** If the appliance is in an error state, a special Oops screen is displayed. This screen indicates that the appliance encountered a serious error and could not recover from it. Restarting the appliance might resolve the error. The error message will advise you to create a support dump file and contact an authorized support representative. Anyone can create an encrypted support dump file from that screen without the need for logging in or other authentication.

---

- Review reports for interconnect, server, and enclosure status. Reports can also provide inventory information and help you see the types of server models and processors in the data center. They can also show you what firmware needs to be updated.

The HPE OneView User Guide has more information on troubleshooting, including:

- The appliance and the appliance network setup
- Firmware bundles
- Networks
- Storage
- Server hardware and server profiles

---

**NOTE:** To access the HPE OneView 3.0 User Guide, scan this QR code or right-click it to open the hyperlink.



[http://h20565.www2.hpe.com/hpsc/doc/public/display?sp4ts.oid=5410258&docLocale=en\\_US&docId=emr\\_na-c05098916](http://h20565.www2.hpe.com/hpsc/doc/public/display?sp4ts.oid=5410258&docLocale=en_US&docId=emr_na-c05098916)

---

## iLO event log

The iLO event log provides a record of significant events recorded by the iLO firmware. Logged events include major server events such as a server power outage or a server reset, and iLO events such as unauthorized login attempts. Other logged events include successful or unsuccessful browser and Remote Console logins; virtual power and power-cycle events; clearing the log; and some configuration changes, such as creating or deleting a user.

iLO provides secure password encryption, tracking all login attempts and maintaining a record of all login failures. The Authentication Failure Logging setting allows you to configure logging criteria for failed authentications. The event log captures the client name for each logged entry to improve auditing capabilities in DHCP environments, and records the account name, computer name, and IP address.

For a list of errors that might be displayed in the event log, see the error messages guide for the server.

### Viewing the iLO event log

1. In the navigation tree, click **Information** and then click the **iLO Event Log** tab.
2. Optional: Use the event log sort, search, and filter features to customize the log view.
3. Optional: To view the event details pane, click an event.

Event log details include:

- **ID**—The event ID number. Events are numbered in the order in which they are generated. By default, the event log is sorted by the ID, with the most recent event at the top.
- **Severity**—The importance of the detected event.
- **Description**—The description identifies the component and detailed characteristics of the recorded event. If the iLO firmware is rolled back to an earlier version, the description `UNKNOWN EVENT TYPE` might be displayed for events recorded by the newer firmware. You can resolve this issue by updating the firmware to the latest supported version, or by clearing the event log.
- **Last Update**—The date and time when the latest event of this type occurred. This value is based on the date and time stored by the iLO firmware. If the iLO firmware did not recognize the date and time when an event was updated, `[NOT SET]` is displayed.
- **Count**—The number of times this event has occurred (if supported). In general, important events generate an event log entry each time they occur. They are not consolidated into one event log entry.

When less important events are repeated, they are consolidated into one event log entry, and the Count and Last Update values are updated. Each event type has a specific time interval that determines whether repeated events are consolidated or a new event is logged.

- **Category**—The event category, such as Hardware Failure, Firmware Failure, Power, or Administration.

---

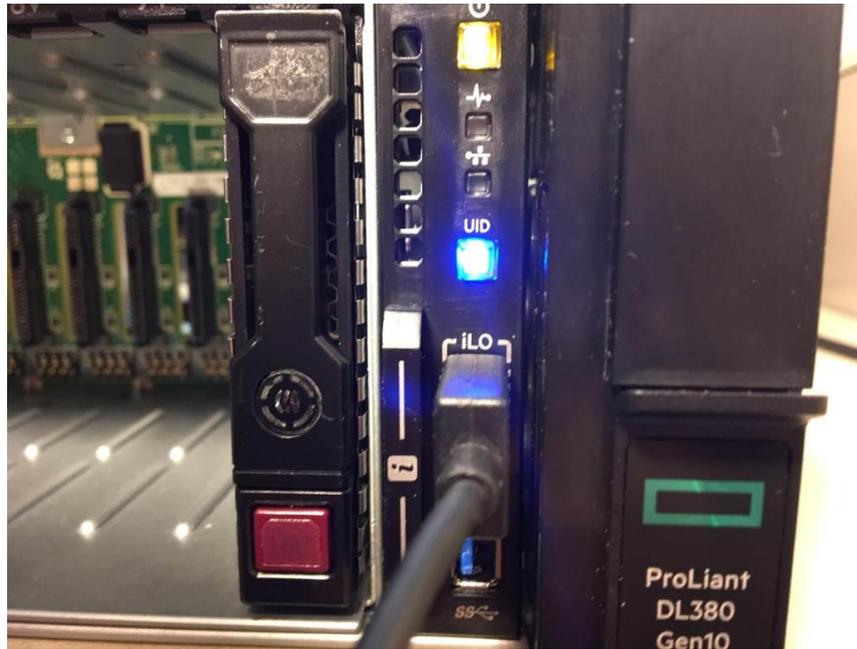
**NOTE:** To access the *HPE Error Message Guide for HPE ProLiant Gen10 servers and HPE Synergy*, scan this QR code or right-click it to open the hyperlink.



[https://h20565.www2.hp.com/hpsc/doc/public/display?sp4ts.oid=7252836&docLocale=en\\_US&docId=emr\\_na-c05362118](https://h20565.www2.hp.com/hpsc/doc/public/display?sp4ts.oid=7252836&docLocale=en_US&docId=emr_na-c05362118)

---

## iLO Service Port



iLO Service Port on a ProLiant DL380 Gen10 server

The dedicated USB port for iLO provides:

- Plug-in Ethernet for all iLO functionality, direct to a laptop
  - Troubleshooting
  - Health information
  - Logs
  - Remote console and virtual media
- Plug-in USB key with a script for easy Active Health System download
- Full customer control over devices and security

To ensure that a secondary remote console client can share video with the primary remote console client, connect to iLO from a network that can communicate directly with the network used by the primary remote console client. The remote console supports sharing by forwarding video directly from the primary remote console client to the secondary remote console client, which is why both clients must have direct communication with each other.

## Learning check

1. What are the three types of log classification?

---

---

---

2. If the iLO firmware is rolled back to an earlier version, the description UNKNOWN EVENT TYPE might be displayed for events recorded by the newer firmware. How can you resolve this issue?

---

---

---

---

## Troubleshooting HPE storage

HPE provides a variety of utilities that will help you diagnose array and other issues with HPE storage systems.

### HPE Smart Storage Administrator diagnostics



SSA Diagnostics screen

From the Diagnostics screen, SSA generates the following reports and logs:

- **Array diagnostic report**—This report contains information about all devices, such as array controllers, storage enclosures, drive cages, as well as logical, physical, and tape drives. For supported solid-state drives (SSDs), this report also contains SmartSSD Wear Gauge information. Array diagnostics is available with SSA 9.0 and later.
- **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.

Each 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

For each controller, or for all of them, you can view or save a diagnostic report and a SmartSSD Wear Gauge report.

For the view tasks, SSA generates and displays the report or log. For the save tasks, SSA generates a report without the graphical display. For either type of task, the report can be saved. In online and offline environments, SSA saves the diagnostic report to a compressed folder, which contains an XML report, a plain text report, and a viewer file so the report can be displayed and navigated using a web browser.

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

SSA can detect and correct performance issues caused by nonoptimal logical drive alignment for 512e physical drives. The following scenarios indicate drive support is needed:

- Multiple logical drives exist in a single array.
- An array consists of one or more 512e physical drives.
- At least one of the logical drives in the array is not aligned on a native block boundary. For current 512e drives, the native block boundary is 4K.

---

**NOTE:** 512e drives feature an advanced format in which the physical sector size is 4,096 bytes, but the logical sector size emulates 512 bytes in sector size. The 4K drives feature an advanced format in which the physical sectors and logical sectors are both 4,096 bytes in size.

---

As a response, HPE SSA will display a warning indicating the logical drive is not optimally aligned and that performance of the logical drive will not be optimal. Additionally, the array will present a Re-align Logical Drive button if the following scenario is met:

- There is enough free space in the array to move the logical drive to be aligned to the native 4K boundary.
- The controller is capable of performing the transformation (requires a cache module with a fully charged battery or capacitor connected).
- The controller does **not** have HPE SmartCache enabled.

## HPE Smart Storage Administrator Diagnostics Utility CLI

The HPE Smart Storage Administrator Diagnostics Utility CLI is a stand-alone diagnostic utility that provides configuration and error information about array controllers, storage enclosures, drive cages, logical drives, physical drives, and tape drives. For any supported SSDs, the utility provides current usage level and remaining expected lifetime.

The utility collects all possible information about storage devices in the system, detects all problems, and provides a detailed configuration report in .zip format. After downloading and installing the software, you can run the utility as a CLI in an online environment. The functionality in this utility is mirrored in the diagnostics features of SSA, which can be run in an offline environment.

### Installing the utility

To install the Smart Storage Administrator Diagnostics Utility CLI, follow these steps:

1. Browse to the HPE Smart Storage Administrator website: <http://www.hpe.com/servers/ssa>.
2. Click **Download software**.
3. Select an operating system.
4. Identify the preferred software and version, and then click **Download**.
5. Save and then run the executable file. By default, the software installs at C:\Program Files\HP System Tools\.

### Launching the utility in CLI mode

1. Click **Start** → **All Programs** → **Windows System** → **Smart Storage Administrator Diagnostics Utility** → **Read Me**.
2. Open a command prompt.
3. Change directory (cd) to the location where ssaducli.exe is installed. This directory is commonly C:\Program Files\Smart Storage Administrator\ssaducli\bin.
4. Do one of the following:
  - Generate a diagnostic report with the following command:  

```
ssaducli -f adu-report.zip
```
  - Generate a SmartSSD Wear Gauge report with the following command:  

```
ssaducli -ssd -f ssd-report.zip
```
  - For more options, use the following command:  

```
ssaducli -help
```

## Additional tools for diagnosing array problems

To troubleshoot array problems and generate feedback about arrays, use the following diagnostic tools:

- **Event Notification Service**—This utility reports array events to the Windows system event log and IML. Access the utility from the Support Center website: <http://www.hpe.com/support/hpesc>.
- **HPE Insight Diagnostics**—Insight Diagnostics is a tool that displays information about the system hardware configuration and performs tests on the system and its components, including drives if they are connected to Smart Array controllers. This utility is available from the Support Center website: <http://www.hpe.com/info/InsightDiagnostics>.
- **POST messages**—Smart Array controllers produce diagnostic error messages (POST messages) at reboot. Many POST messages suggest corrective actions. For more information about POST messages, refer to *HPE ProLiant Gen9 Troubleshooting Guide: Volume II—Error Messages*.

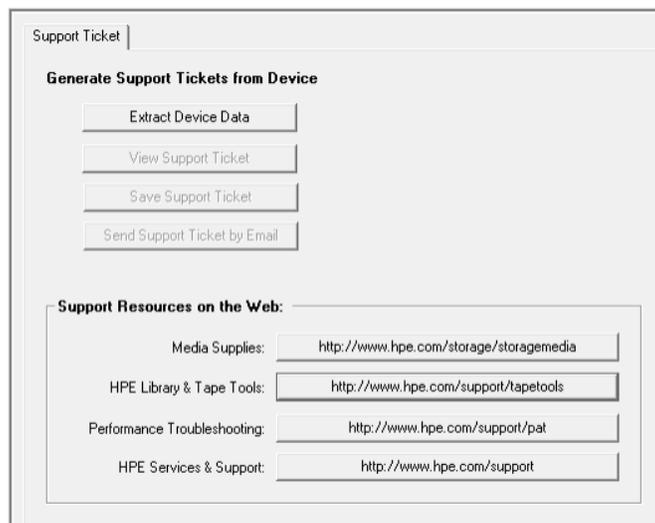
## Troubleshooting LUNs

Logical unit number (LUN) issues are difficult to diagnose because they usually arise from specific features of a specific system's hardware and software. The most common issue is that one or more LUNs are not visible, but this can have several reasons.

Because LUNs sit at the crossroads of physical and virtual storage, they can be affected by both hardware and software. Everything from host bus adapters (HBAs) to operating system configuration to storage management software can cause issues with LUNs. If the LUNs do not work properly, follow these steps to diagnose the issue:

1. Start with the documentation. It should indicate what each device or software service expects from the LUNs. Almost invariably, these are configuration issues, rather than hard limits. To fix them, understand what is expected. For example, does the hardware auto-configure? Most hardware does, but there is some hardware that does not. Another possible issue with a Windows boot-from-SAN installation occurs when the storage controller does not follow the SCSI-3 standard of assigning LUN 0 to the controller, rather than to a disk. To start from SAN, Windows requires that LUN enumeration adhere to the SCSI-3 standard.
2. Get the right drivers. Make sure the latest drivers are used in the HBAs and other parts of the system.
3. Check the LUN masking policies. One of the most common issues is an inappropriate LUN masking policy, especially if LUN 0 is masked from hardware or software that needs it. Although LUN masking is basic to storage security, check policies and make sure that everything that needs to see LUN 0 can find it.
4. Make sure that the system is configured for the right number of LUNs. Many kinds of hardware and software assume that there are eight LUNs or less. If there are more LUNs than that, configuration parameters might have to be changed.
5. Determine whether some LUNs are visible. If some LUNs are visible but not all, it can be assumed that the issue is not fabric-specific.
6. If the situation gets worse, determine what is known and not known, and then follow the LUN creation process from the beginning to see where it fails. Make sure the system finds the physical disks, and then work through setting up the LUNs. Try setting up LUN 0 separately, and then adding LUNs.
7. When Microsoft Multipath I/O (MPIO) and Microsoft Device Specific Module (DSM) are configured on Windows server operating systems, a storage device can be connected to the host by using redundant physical paths. In this scenario, when multiple paths are lost concurrently, the LUN for the storage device is lost on the host. There is a hotfix for Windows Server 2012 R2, Windows Server 2012, or Windows Server 2008 R2.

## Using L&TT to troubleshoot tape devices



Support ticket generation screen

A major feature of HPE Library and Tape Tools (L&TT) is the ability to generate reports and support tickets. For a hardware problem, a report or support ticket can provide vital information to help diagnose and resolve the problem. Reports and support tickets can be generated on all operating systems supported by L&TT.

To extract and send the report by email:

1. Extract the support ticket.
2. Save the support ticket.
3. View the support ticket.
4. Click **Send Support Ticket by Email** to send the report or support ticket immediately by email to an HPE support center or other destination.

If L&TT determines that a tape drive is not functioning correctly and needs to be exchanged, the accompanying failure information provides a faster response from HPE support. With L&TT, you can perform periodic health checks by running the L&TT Drive Assessment test and the L&TT Media Analysis test on tapes holding key data.

You can also run L&TT in Installcheck mode, which performs several checks to determine if the hardware is connected properly. If a tape drive is not performing to specification, L&TT has different tests to measure performance:

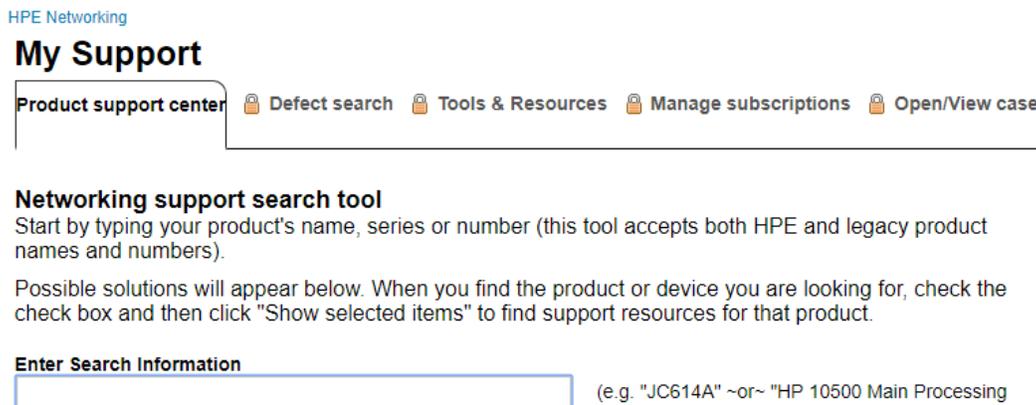
- The Drive Performance test measures the transfer rate of the drive, HBA, and cabling.
- The System Performance test consists of two separate tests:
  - The Backup Performance test measures the data source transfer rate from the disks or network for backup.
  - The Restore Performance test measures the data restore transfer rate to the disks or across the network for restores.

These tests measure the actual transfer rates, enabling you to determine if a performance problem exists, and if so, whether the problem is with the drive or the system. In most cases, system performance is the limiting factor.

## Troubleshooting the network

Performance-related network problems can be caused by topology, switch configuration, and the effects of other devices or their configurations on switch operation. For switch-specific information on hardware problems indicated by LED behavior, cabling requirements, and other potential hardware-related problems, refer to the installation guide that shipped with the switch.

### Diagnosing switch problems



HPE Networking Lookup/Support search tool main screen

Use these approaches to diagnose switch problems:

- Use the HPE Networking Lookup/Support search tool to search for and download software and firmware for switch/router/hub/wireless devices and to find product-specific manuals and FAQs: <https://h10145.www1.hpe.com/support/SupportLookup.aspx>.
- Check the switch LEDs for indications of proper switch operation:
  - Each switch port has a Link LED that should illuminate whenever an active network device is connected to the port.
  - Problems with the switch hardware and software are indicated by flashing the Fault and other switch LEDs.

For a description of LED behavior and information on using LEDs for troubleshooting, refer to the installation guide that shipped with the switch.

- Check the network topology and installation. For topology information, refer to the installation guide that shipped with the switch.
- Check cables for damage, the correct type, and proper connections. You should also use a cable tester to check the cables for compliance to the relevant IEEE 802.3 specification. For correct cable types and connector pin-outs, refer to the installation guide that shipped with the switch.
- Use the Port Utilization Graph and Alert Log in the WebAgent included in the switch to help isolate problems. These tools are available through the WebAgent:
  - Port Utilization Graph
  - Alert log
  - Port Status and Port Counters screens
  - Diagnostic tools (link test, ping test, configuration file browser)

- For help isolating problems, use the easy-to-access switch console built into the switch or telnet to the switch console. These tools are available through the switch console:
  - Status and Counters screens
  - Event Log
  - Diagnostic tools (link test, ping test, traceroute, and port auto-negotiation)

## Diagnostic tools

The ping test and the link test are point-to-point tests between the switch and another IEEE 802.3-compliant device on the network. These tests can tell you whether the switch is communicating properly with another device.

- **Ping test**—Tests the path between the switch and another device on the same or another IP network that can respond to IP packets (Internet Control Message Protocol [ICMP] echo requests). The syntax of the command is:

```
ping <destination IP> source
```

- The source IP you want to use must be one of the switch's IP addresses.
- **Link test**—Tests the connection between the switch and a designated network device on the same LAN (or VLAN, if configured). During the link test, IEEE 802.2 test packets are sent to the designated network device in the same VLAN or broadcast domain. The remote device must be able to respond with an 802.2 Test Response Packet.
- **Traceroute**—Traces the route from the switch to a host address. This command outputs information for each (router) hop between the switch and the destination address. Note that every time you execute traceroute, it uses the same default settings unless you specify otherwise for that instance of the command.

## Port auto-negotiation

When a link LED does not light (indicating loss of link between two devices), the most common reason is a failure of port auto-negotiation between the connecting ports. If a link LED fails to light when you connect the switch to a port on another device, do the following:

1. Ensure that the switch port and the port on the attached end-node are both set to **Auto** mode. To check the mode setting for a port on the switch, use either the Port Status screen in the menu interface or execute the appropriate command in the CLI.
2. If the attached end-node does not have an Auto mode setting, you must manually configure the switch port to the same setting as the end-node port.

## Troubleshooting telnet issues

Off-subnet management stations can lose telnet access if you enable routing without first configuring a static (default) route. That is, the switch uses the IP default gateway only while operating as a Layer 2 device. If routing is enabled on the switch, the IP default gateway is not used. You can avoid this problem by using the `ip route` command to configure a static (default) route before enabling routing.

Telnet access might be disabled by the `Inbound Telnet Enabled` parameter in the System Information screen of the menu interface:

### 2. Switch Configuration

#### 1. System Information

The switch might not have the correct IP address, subnet mask, or gateway. Verify by connecting a console to the switch's console port and selecting:

### 2. Switch Configuration

#### 5. IP Configuration

If you are using DHCP to acquire the IP address for the switch, the IP address "lease time" might have expired so that the IP address has changed. For more information on how to reserve an IP address, refer to the documentation for the DHCP application that you are using.

If one or more IP-authorized managers are configured, the switch allows inbound telnet access only to a device having an authorized IP address.

## Troubleshooting unusual network activity

Network activity that fails to meet accepted norms can indicate a hardware problem with one or more of the network components, possibly including the switch. Such problems can also be caused by a network loop or simply too much traffic for the network as it is currently designed and implemented. Unusual network activity is usually indicated by the LEDs on the front of the switch or measured with the switch console interface or with a network management tool.

A topology loop can also cause excessive network activity. The Event Log FFI messages can be indicative of this type of problem.

## General problems

General issues to prepare for include:

- **The network runs slow; processes fail; users cannot access servers or other devices**— Broadcast storms might be occurring in the network. These might be caused by redundant links between nodes.
  - If you are configuring a port trunk, finish configuring the ports in the trunk before connecting the related cables. Otherwise you might inadvertently create a number of redundant links (that is, topology loops) that will cause broadcast storms.
  - Check for FFI messages in the Event Log.
- **Duplicate IP addresses**—This issue is indicated by this Event Log message:
 

```
ip: Invalid ARP source: IP address on IP address
```

  - In this case, both instances of `IP address` are the same address, indicating that the switch's IP address has been duplicated somewhere on the network.

- **Duplicate IP addresses in a DHCP network**—If you use a DHCP server to assign IP addresses in the network, and you find a device with a valid IP address that does not seem to communicate properly with the server or other devices, a duplicate IP address might have been issued by the server. This can occur if a client has not released a DHCP-assigned IP address after the intended expiration time and the server "leases" the address to another device. This can also happen, for example, if the server is first configured to issue IP addresses with an unlimited duration, and then is subsequently configured to issue IP addresses that will expire after a limited duration. One solution is to configure "reservations" in the DHCP server for specific IP addresses to be assigned to devices having specific media access control (MAC) addresses.

One indication of a duplicate IP address in a DHCP network is this Event Log message:

```
ip: Invalid ARP source:
<
```

```
IP-address> on <
```

```
IP-address>
```

- In this case, both instances of `IP-address` are the same address, indicating that the IP address has been duplicated somewhere on the network.

- **The switch has been configured for DHCP/Bootp operation, but has not received a DHCP or Bootp reply**—When the switch is first configured for DHCP/Bootp operation, or if it is rebooted with this configuration, it immediately begins sending request packets on the network. If the switch does not receive a reply to its DHCP/Bootp requests, it continues to periodically send request packets, but with decreasing frequency. Thus, if a DHCP or Bootp server is not available or accessible to the switch when DHCP/Bootp is first configured, the switch might not immediately receive the desired configuration.

After verifying that the server has become accessible to the switch, reboot the switch to restart the process.

## LACP-related problems

If you cannot enable Link Aggregation Control Protocol (LACP) on a port with the `interface <port-number> lacp` command, the switch displays the message:

```
Operation is not allowed for a trunked port
```

You cannot enable LACP on a port when it is configured as a static trunk port. To enable LACP on a static trunk port:

1. Use the `no trunk <port-number>` command to disable the static trunk assignment.
2. Execute: `interface <port-number> lacp`




---

**CAUTION:** Removing a port from a trunk without first disabling the port can create a traffic loop that can slow down or halt the network. Before removing a port from a trunk, HPE recommends that you either disable the port or disconnect it from the LAN.

---

## Learning check

1. What is a common issue with a LUN that is not operating properly?

---

---

---

---

2. Match the network diagnostic tool with its purpose.

<b>Tool</b>	<b>Purpose</b>
ping	Trace the path from the switch to a host address
link	Test the path between the switch and another device on the same or another IP network that can respond to IP packets
tracert	Test the connection between the switch and a designated network device on the same LAN (or VLAN, if configured)

3. Before removing a port from a LAG, you should either disable the port or disconnect it from the LAN.

- True
- False

## Summary

- Troubleshooting is a process that involves knowledge, experience, and intuition. An important part of this process is the use of proactive measures such as record keeping. Before beginning any maintenance, upgrade, or troubleshooting procedures, it is important to follow all product safety and operating instructions. Always refer to the documentation supplied with the product.
- The HPE six-step troubleshooting process helps to systematically identify the core of a problem, resolve it, and take steps to prevent it from reoccurring. The six steps are:
  1. Collect data.
  2. Evaluate the data to determine potential subsystems causing the issue.
  3. Develop an optimized action plan.
  4. Execute the action plan.
  5. Determine whether the problem is solved.
  6. Implement preventive measures.
- Effective troubleshooting methods include establishing a baseline, using published server information, collecting problem and system data, using online and offline diagnostic tools, and checking status LEDs. These strategies can be extremely effective in reaching a resolution.
- HPE provides procedures for resolving common problems and comprehensive courses of action for fault isolation and identification, error message interpretation, issue resolution, and software maintenance on ProLiant servers and server blade models.
- HPE provides a variety of tools and utilities that will help you diagnose array and other issues with HPE storage systems, including:
  - HPE Smart Storage Administrator and SSA Diagnostics Utility CLI
  - Event Notification Service
  - HPE Insight Diagnostics
  - POST messages
  - L&TT
- Performance-related network problems can be caused by topology, switch configuration, and the effects of other devices or their configurations on switch operation.



---

# Glossary of HPE Data Center Solutions

## Terms

### Appendix 1

---

#### A

---

AAA	See <i>Authentication, authorization, and accounting</i> .
Access control list (ACL)	A list of permissions associated with an object. An ACL specifies the users or system processes that are given access to objects, as well as which operations are allowed on given objects.
Access layer	Is responsible for providing end user devices with a connection to network resources. It is also known as the <i>network edge</i> or <i>desktop layer</i> because it connects client nodes to the network.
Access point (AP)	A networking hardware device that creates a WLAN allowing users to connect to the network. The AP often connects to a router via a wired network as a stand-alone device, but it can also be a component of the router itself.
ACL	See <i>Access control list</i> .
Advanced Encryption Standard (AES)	A specification for the encryption of electronic data established by the U.S. National Institute of Standards and Technology. It is a U.S. government standard for secure and classified data encryption and decryption.
AES	See <i>Advanced Encryption Standard</i> .
Agentless data collection	A collection of inventory and performance data from managed systems without requiring installation or configuration of agents on managed systems.
Aggregation layer	Contains multilayer switches that perform an aggregation function connecting server farms that span multiple access switches. Its primary purpose of the distribution layer is cable reduction and network management, taking the many uplinks from edge switches and aggregating them into higher speed links. The aggregation layer also contains devices that provide services such as firewalls, load balancers, and SSL terminations. It is also known as the <i>distribution layer</i> .

---

---

Alert	Represents one or more events or possible problems with a resource; usually has a severity, status, and description associated with it.
American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE)	A global society advancing human well-being through sustainable technology for the building environment. The society produces guidelines such as A3 and A4, which are used to classify equipment temperature tolerances. Equipment that meets the A3 or A4 guideline allows data centers more flexibility in their power and cooling strategies.
AP	See <i>Access point</i> .
API	See <i>Application programming interface</i> .
Application lifecycle management	The process of planning, provisioning, and customizing applications and their resources, then patching and returning them to the pool of free resources.
Application programming interface (API)	A set of programming language libraries used to communicate with a particular service, application, or server.
Array	A physical storage device or an aggregate set of devices containing one or more storage pools that serves storage to clients and provides an API or a management console to allow an administrator to manage the storage provided by the device.
ASHRAE	See <i>American Society of Heating, Refrigerating and Air-Conditioning Engineers</i> .
Availability	The ability of a user to access information or resources in a specified location and in the correct format.
Authentication, authorization, and accounting (AAA)	A framework for intelligently controlling access to computer resources, enforcing policies, auditing usage, and providing the data necessary to bill for services.
<b>B</b>	
Bare metal server	A server that is not booted with a production operating system (could be a new server with no operating system); a server not yet known to the management software.
Basic input/output system (BIOS)	Non-volatile firmware used to perform hardware initialization during the booting process. It also provides runtime services for operating systems.
Backup, recovery, and archiving (BURA)	An approach to data loss prevention that uses backups (copies of data made in case the primary copy is lost or damaged), recovery (restoration of data files from backup files), and archives (migration of infrequently accessed data from primary storage to a lower-cost long-term tier).
BfS	See <i>Boot from SAN</i> .

---

---

Bill of materials (BOM)	A list of parts needed to create a unit; a list of components or sub-components, such as part numbers, of an equipment or a hardware configuration.
BIOS	See <i>Basic input/output system</i> .
Block storage	Storage solution based on presenting a disk or a virtual disk to the operating system as if it were a local disk drive. Typically the operating system accesses the device directly, or builds a local file system on top of that device. The system typically assumes exclusive access to the presented device. Block storage is deployed in a SAN and other large-scale storage solutions.
BOM	See <i>Bill of materials</i> .
Boot from SAN (BfS)	Bootstrapping of server operating system from a boot device (logical drive) located on a storage area network (SAN) as opposed to an internal storage device.
BPDU	See <i>Bridge protocol data units</i> .
Bridge protocol data units (BPDU)	Data messages that are exchanged across the switches within an extended LAN that uses a Spanning Tree Protocol (STP) topology. BPDU packets contain information on ports, addresses, and priorities and ensure that the data is transmitted to the appropriate destination.
Bring your own device (BYOD)	Policies and technologies that allow employees to bring personally owned mobile devices (laptops, tablets, and smart phones) to their workplace and use those devices to access privileged company information and applications.
British Thermal Unit (BTU)	A measure of the amount of energy required to raise the temperature of one pound of water by one degree Fahrenheit.
BTU	See <i>British Thermal Unit</i> .
BURA	See <i>Backup, recovery, and archiving</i> .
Business continuity	Planning and activities performed by an organization to ensure that critical business functions continue to be available on a daily basis, and not only during a time of disaster.
BYOD	See <i>Bring your own device</i> .

---

## C

---

Cache	Memory built into the CPU chip. Cache memory pools store information that the processor is most likely to need next.
Central processing unit (CPU)	The computer component responsible for interpreting and executing most of the commands from the computer's other hardware and software.
CIFS	See <i>Common Internet File System</i> .
CLI	See <i>Command line interface</i> .

---

---

Cluster	A parallel or distributed computing system comprising many discrete systems that form a single, unified system. Clusters can provide high availability, load balancing, and parallel processing.
CNA	See <i>Converged network adapter</i> .
Command line interface (CLI)	A means of interacting with a computer device, a program, or an operating system where the user issues text-based line commands.
Common access card	A single corporate card or token that can be used for user authentication for PC, network, and application login and building access.
Common Internet File System (CIFS)	A protocol that defines a standard for sharing files across corporate intranets and the internet. It provides the foundation for file sharing in Windows.
Composer	A component of HPE Synergy that provides infrastructure management for assembling and re-assembling fluid pools of compute, storage, and fabric resources to meet any workload.
Converged infrastructure	An IT environment where server, storage, and networking resources are standardized, virtualized, organized into, and managed as resource pools.
Converged systems	Factory-integrated platforms that deploy virtual machines quickly, simplify IT operations, and reduce costs. They also provide integrated management and lifecycle support.
Converged network adapter (CNA)	A network adapter that combines the functionality of a host bus adapter (HBA) with a network interface controller (NIC).
Core	The data-processing engine within a processor. A single processor might have multiple cores. A core might support multiple execution threads.
Core boosting	Technology that enables better performance across more processor cores, resulting in significant savings in core-based licensing.
CPU	See <i>Central processing unit</i> .
CPU socket	The connector on a system board that houses a CPU. It creates the electrical interface and contact between the CPU and the system board.
Create, read, update, and delete (CRUD)	Basic HTTP methods in the REST API. They also represent the four basic functions of persistent storage.
CRM	See <i>Customer relationship management</i> .
CRUD	See <i>Create, read, update, and delete</i> .
Customer relationship management	A term that refers to practices, strategies and technologies that companies use to manage and evaluate customer interactions and data throughout the customer lifecycle. The goal of CRM is to improve business relationships with customers, facilitate customer retention, and increase sales. CRM systems are designed to compile information on customers across different channels or points of contact between the customer and the company.

---

---

**D**

---

DAS	See <i>Direct-attached storage</i> .
Data center bridging	A set of IEEE standards that create a lossless fabric on top of Ethernet. RDMA works best when the underlying wires implement a lossless fabric. Data center bridging represent a set of enhancements to the Ethernet local area network communication protocol for data center environments.
Data Plane Development Kit (DPDK)	A set of libraries and drivers designed to run on any processor to enable fast packet processing. DPDK is an open source BSD licensed project that allows software-based customization and optimization of network performance by using polling instead of traditional interrupt-driven network processing. It is used in network functions virtualization (NFV) deployments.
Deduplication	Technology similar to audio and video compression that compresses data files so that a smaller storage volume can hold the data. Redundant data blocks are removed and replaced with pointers to the unique data copy. It reduces storage and processing costs.
Desktop virtualization	A VM running a desktop operating system that users access from multiple locations using a variety of devices. Also called <i>client virtualization</i> .
DHCP	See <i>Dynamic Host Configuration Protocol</i> .
DIMM	See <i>Dual in-line memory module</i> .
Direct-attached storage	A method to provide data access by directly connecting the digital storage device to the computer accessing it. Examples include hard drives, solid state drives, and optical disk drives. Benefits include speed and affordability.
Disaster recovery	A set of processes, policies, and procedures necessary to prepare for recovery or continuation of technology infrastructure critical to an organization after a natural or human-caused disaster.
Discovery	A feature within a management application that finds and identifies network objects.
Disk group	An aggregation of disks of the same type, using a specific RAID type that is incorporated as a component of a pool. The purpose of disk groups is to store volume data.
Domain	A collection of computers and devices on a network that are administered as a unit with common rules and procedures. Within the internet, domains are defined by the IP address.
DRAM	See <i>Dynamic random access memory</i> .
DPDK	See <i>Data Plane Development Kit</i> .
Dual in-line memory module (DIMM)	A module that contains one or several random access memory (RAM) chips on a small circuit board with pins that connect it to the computer motherboard.

---

Dynamic Host Configuration Protocol (DHCP)	A network protocol that enables a server to automatically assign an IP address to a computer from a defined range of numbers.
Dynamic random access memory (DRAM)	High-capacity, low-cost volatile memory. It is fast and reliable but cannot retain data when power is lost.

---

## E

---

East/west traffic	The network traffic that goes between servers in a given data center (in the context of the data center LAN), or traffic that goes between servers in different data centers (in the context of the WAN). Can also refer to client-to-client traffic in collaboration tools.
ECC memory	See <i>Error correcting code memory</i> .
EoR	See <i>End of row</i> .
End of row (EoR)	A network design in which each server in a rack is connected directly to an aggregation switch rather than to individual switches in each rack. Although this design reduces the number of network devices and improves the port utilization of the network, the horizontal cabling required is more complex.
Error correcting code (ECC) memory	A protocol that uses a special algorithm to encode information in a block of bits. This block contains enough detail to permit a single bit error to be recovered in the protected data. ECC memory detects both single-bit and multi-bit errors, and corrects single-bit errors transparently.
Electromagnetic interference (EMI)	A disturbance generated by an external source that affects an electrical circuit. Actions should be taken to prevent EMI.
EMI	See <i>Electromagnetic interference</i> .
Enclosure	A chassis that contains multiple server blades and interconnect devices.

---

## F

---

Fabric	Fibre Channel network composed of one or more interconnected Fibre Channel switches.
Failover	The operation that takes place when a primary service (network, storage, or CPU) fails and the application continues operation on a secondary unit.
FCoE	See <i>Fibre Channel over Ethernet</i> .
Fibre Channel	A network technology primarily used for storage networks. It is an industry-standard interconnection protocol featuring a serial architecture. Currently, Fibre Channel supports high-speed, full-duplex connections in multi-topology environments.

---

---

Fibre Channel over Ethernet (FCoE)	An encapsulation of Fibre Channel frames over Ethernet networks. This allows Fibre Channel to use 10 Gbps or higher-speed Ethernet networks while preserving the Fibre Channel protocol.
Field replaceable unit (FRU)	A device part or component that can be replaced in the field by qualified personnel trained by HPE.
File storage	File storage enables users to access files and folders through a hierarchal formatted file system, such as NTFS. Files are stored based on name and file extension. File storage is mostly used on hard-drives and NAS solutions.
File Transfer Protocol (FTP)	A standard network protocol used for the transfer of computer files between a client and server on a computer network.
Flash drives	High-performance drives designed for customers requiring boot-from-flash for integrated hypervisors and first tier operating systems.
Flash memory	Nonvolatile memory that can be erased and written to more than once.
FlexFabric	Solution that provides the network infrastructure for the data center; simplifies data center infrastructure with converged network, computer, and storage resources across both virtual and physical environments to accommodate cloud computing models.
FRU	See <i>Field replaceable unit</i> .
FTP	See <i>File Transfer Protocol</i> .

---

**G**

---

GPU	See <i>Graphics processing unit</i> .
Graphics processing unit (GPU)	A single-chip processor that creates lighting effects and transforms objects when a 3D scene is redrawn on a computer. These are mathematically-intensive tasks that otherwise would put a significant strain on the CPU.

---

**H**

---

Hard disk drive (HDD)	A data storage device that uses magnetic storage to store and retrieve digital information using one or more rigid, rapidly rotating disks coated with magnetic material.
HBA	See <i>Host bus adapter</i> .
HDD	See <i>Hard disk drive</i> .
High availability	System design approach and associated service implementation that ensures that a pre-arranged level of operational performance will be met during a contractual measurement period.
Host	A system or partition that is running an instance of an operating system.

---

---

Host bus adapter	A circuit board or integrated circuit adapter that provides input/output processing and physical connectivity between a server and a storage device.
Hosting	The process of making services available over a network.
HPE Active Health System	A solution that monitors and records changes in the server hardware and system configuration. It assists in diagnosing problems and delivering rapid resolution when server failures occur.
HPE CISSS	See <i>HPE Converged Infrastructure Solution Sizer Suite</i> .
HPE CMC	See <i>HPE StoreVirtual Centralized Management Console</i> .
HPE ConvergedSystems	HPE turnkey systems that integrate hardware, software, and services. Tested, integrated, optimized, and scalable for virtualization, cloud, and next-generation applications.
HPE Event Notification Service	A utility that reports array events to the Windows system event log and Integrated Management Log.
HPE FlexFabric	An HPE solution that provides the network infrastructure for the data center; simplifies data center infrastructure with converged network, computer, and storage resources across both virtual and physical environments to accommodate cloud computing models.
HPE FlexibleLOM	HPE technology that uses a dedicated system board edge connector that accepts an optional FlexibleLOM NIC. A FlexibleLOM architecture maintains the close-coupled interface of a LOM but allows users to select the connectivity they need now and adapt to network changes in the future without using a standard PCIe slot.
HPE iLO	See <i>HPE integrated Lights Out</i> .
HPE IMC	See <i>HPE Intelligent Management Center</i> .
HPE InfoSight	A cloud-based management and support system for Nimble storage arrays that integrates, automates, and simplifies storage administrative tasks.
HPE Insight Control	System management software that manages a wide variety of systems, including HPE servers, clusters, desktops, workstations, and portables.
HPE Insight Online	A web-based, personalized IT and support dashboard; provides the information users need to monitor HPE devices in their IT environment from anywhere, at any time, and at no additional cost.
HPE Insight Remote Support	A comprehensive solution that enables enhanced delivery of an HPE warranty, support services, or contractual support agreement. It provides intelligent event diagnosis and automatic, secure submission of hardware event notifications to HPE, which initiates a fast and accurate resolution based on the product's service level.

---

---

HPE integrated Lights-Out (iLO)	A management processor that virtualizes system controls, enables remote management, simplifies server setup, and engages in health monitoring and power and thermal control of HPE servers. It is embedded on the system boards of ProLiant servers and provides multiple ways to configure and update servers remotely.
HPE Intelligent Management Center (IMC)	A centralized network management platform from HPE that allows you to manage both physical and virtual networks; supports both HPE and third-party network devices.
HPE Intelligent Provisioning	HPE technology that configures HPE ProLiant Gen8 and later servers and Synergy compute modules. It deploys "off-the-shelf" and HPE branded versions of leading server operating systems; replaces HPE SmartStart.
HPE Intelligent Resilient Framework (IRF)	Extends network control over multiple active switches. IRF helps simplify networks by allowing the combination of multiple switches into a single, ultra-resilient virtual device. The two-tier topology is connected with aggregated links—increasing bandwidth and providing redundancy.
HPE iQuote Universal	An HPE cloud service for channel sales configuration and quoting. It generates quotes for products across the HPE portfolio, including servers, storage, networking, commercial desktops, laptops, and workstations. It includes real-time information on promotional pricing and stock, and notifies users of technical errors.
HPE Library and Tape Tools (LT&T)	A free, downloadable diagnostic tool for all HPE tape storage and magneto-optical storage products. It performs firmware upgrades, verification of device operation, failure analysis, and a range of utility functions.
HPE LL&T	See <i>HPE Library and Tape Tools</i> .
HPE Networking Online Configurator	A web-based tool that enables product quotes according to precise specifications. It contains the most current HPE Networking product and pricing information.
HPE OCA	See <i>HPE One Config Advanced</i> .
HPE OCS	See <i>HPE One Config Simple</i> .
HPE One Config Advanced (OCA)	A web-based tool that configures large and/or complex solutions. It replaces HPE SalesBuilder for Windows and Watson.
HPE One Config Simple (OCS)	A web based tool that helps sales and non-technical people provide customers with initial configurations quickly and easily. The configuration can be sent to HPE for special pricing or configuration help.
HPE OneView	A converged infrastructure, software-defined management platform providing a single integrated view of the IT infrastructure, a software-based approach to lifecycle management, and an open development environment.
HPE Power Advisor	An easy-to-use tool that estimates data center power requirements for server and storage configurations. Power Advisor is available in an online Windows application as well as a downloadable version for working offline.

---

HPE Products and Solutions Now (PSNow)	A portal that provides access to a repository of current product and solution information. It provides real-time access to selling information across HPE on any device.
HPE PSNow	See <i>HPE Products and Solutions Now</i> .
HPE Rack and Power Infrastructure Architect Tool	A web-based tool used to select racks and the power options for a solution. It also provides options for accessories such as air flow kits and panels, and cable management or stabilization products.
HPE reference architectures	Designs for complete, open solutions optimized for specific workloads. Each presized and tested configuration follows a proven deployment methodology to help reduce provisioning time, cost, risk and errors, while also speeding application access.
HPE SAN Design Reference Guide	A single source reference for SAN configuration support, with extensive information about interoperability, HPE SAN architecture, including Fibre Channel, iSCSI, Fibre Channel over Ethernet, SAN extension, and hardware interoperability. It also provides end-to-end storage networking architectural information.
HPE Server Memory Configurator	A web-based tool specifically designed to assist with populating DDR Memory in ProLiant servers. The tool provides optimal configurations based on memory population guidelines. Non-optimal configurations are also shown for customers who require a specific memory configuration.
HPE Single Point of Connectivity Knowledge (SPOCK)	An HPE portal for obtaining detailed information about supported HPE storage product configurations. SPOCK is the primary place to find information about HPE storage connectivity—interoperable product and solution configurations that are certified, qualified, and supported by HPE.
HPE Sizer for Server Virtualization	An automated, downloadable tool that provides quick and helpful sizing guidance for server and storage configurations running in VMware vSphere or Hyper-V environments.
HPE Smart Selling Tool	A web-based tool to access relevant assets and collateral. It is an internal tool available to HPE employees and authorized resellers.
HPE Smart Storage Administrator (SSA)	A web-based application that helps configure, manage, diagnose, and monitor HPE Smart Array controllers and HBAs. It is the main tool for configuring arrays on Smart Array controllers and replaces the HPE Array Configuration Utility (ACU).
HPE Smart Storage Administrator Diagnostics Utility CLI	A stand-alone diagnostic utility that provides configuration and error information about array controllers, storage enclosures, drive cages, logical drives, physical drives, and tape drives.
HPE Smart Update Manager (SUM)	A technology included in many HPE systems software maintenance and management products for installing and updating firmware and system software components on HPE ProLiant and HPE Integrity servers, enclosures, and options.

HPE SMU	See <i>HPE Storage Management Utility</i> .
HPE Solution Demo Portal	A portal from HPE that provides a central location for demonstrations, webinars, and supporting collateral that showcase how HPE technologies lead, innovate, and transform enterprise business; includes live and pre-recorded demonstrations.
HPE SPOCK	See <i>HPE Single Point of Connectivity Knowledge</i> .
HPE SSA	See <i>HPE Smart Storage Administrator</i> .
HPE SSD Selector Tool	A web-based tool to help select solid state drives for a customer solution. It allows choice of the SSD type, server type, drive capacity, interface type, and form factor. It also guides SSD selection based on customer workloads.
HPE Storage Management Utility (SMU)	The management GUI for modular smart array (MSA) storage. It facilitates easy setup of email and SNMP alerts and notifications.
HPE Storage Sizing Tool	A downloadable tool that simplifies the process of designing a storage solution by applying storage design rules, as well as licensing and services rules. The output is a valid, supported configuration with localized parts and pricing.
HPE StoreOnce	A family of modern storage systems from HPE that provide a single, high-performance backup architecture that spans the entire enterprise. StoreOnce storage systems greatly reduce the amount of backup data companies need to store.
HPE StoreServ	A family of modern storage systems from HPE that are autonomic, efficient, and federated.
HPE StoreVirtual	A family of modern storage systems from HPE that provide affordable storage for a virtualized infrastructure. Advanced shared storage technology provides the foundation for a composable data fabric, opening up new possibilities for simplified management and scalability across the customer's infrastructure.
HPE StoreVirtual Centralized Management Console	An intuitive, GUI-based administrative interface for StoreVirtual devices running LeftHand OS. It can be used to configure and manage storage volumes spanning clustered storage nodes and view a single graphical layout of the storage environment.
HPE SUM	See <i>HPE Smart Update Manager</i> .
HPE Switch Selector	A web-based tool to use to select the right network switch for a customer's application.
HPE Synergy Planning Tool	A web-based, easy-to-use tool that helps in planning a Synergy solution based on HPE Synergy product rules and best practices.
HPE System Management Homepage	A web-based interface that provides a consolidated view of all system health and configuration information and simplifies access to HPE web-enabled management software running on ProLiant servers.
HPE Virtual Connect	HPE hardware virtualization product, primarily for server blades.

HTTP	See <i>Hypertext Transfer Protocol</i> .
Hybrid cloud	The cloud infrastructure contains two or more clouds (private, community, or public) that remain separate but are bound together by technology that enables data and application portability.
Hyper-threading	Intel technology that enables certain processors to create a second virtual core that allows additional efficiencies of processing. This is not a true multi-core processor, but it adds performance benefits.
Hypertext Transfer Protocol (HTTP)	An application protocol for distributed, collaborative, and hypermedia information systems. It is the foundation of data communication for the internet.
Hypervisor	A virtualization software layer that distributes the physical server resources among multiple virtual machines.
<b>I</b>	
IEEE	See <i>Institute of Electrical and Electronics Engineers</i> .
IGMP	See <i>Internet Group Management Protocol</i> .
Image Streamer	A component of HPE synergy that acts as a repository of bootable images, which can be streamed across multiple compute modules. Image Streamer Is a physical appliance that plugs into the appliance bay at the side of the Synergy frame.
IML	See <i>Intelligent Management Log</i> .
Integrated Management Log (IML)	A record of significant events that occur during system operation. iLO captures and stores the server's IML for access through a browser or command line, even when the server is not operational.
Institute of Electrical and Electronics Engineers (IEEE)	A professional association that sets and maintains the standards that drive the functionality, capabilities, and interoperability of networking products and services.
Intelligent Platform Management Interface (IPMI)	Defines a set of computer interface specifications to a system and provides out of band management and monitoring capabilities independently of the host system's processors, BIOS or UEFI, and operating system. It allows a system administrator to manage a processor that is powered off or otherwise experiencing power issues.
Intelligent power distribution unit (iPDU)	A power distribution unit that can be automatically discovered and controlled.
Intelligent System Tuning	A set of server tuning technologies in ProLiant Gen10 servers, developed in partnership with Intel, that enables administrators to dynamically tune servers to match the unique needs of each workload.
Interconnect module	A hardware module that enables communication between server hardware in a BladeSystem enclosure and data center networks.

---

Interleaving	A memory technology that maximizes performance by spreading memory addresses evenly across memory banks. The result is greater memory throughput because the time spent waiting for memory banks to become available is reduced. It can compensate for slow data transfer in DRAM or core memory by increasing bandwidth.
Internet Group Management Protocol (IGMP)	Communications protocol used on IP networks to establish multicast group memberships.
Internet of Things (IoT)	A system of interrelated computing devices, mechanical and digital machines or objects that have unique identifiers and the ability to transfer data over a network without requiring human-to-human or human-to-computer interaction.
Internet Small Computer System Interface (iSCSI)	The SCSI protocol mapped to TCP/IP and run over standard Ethernet technologies. It enables universal access to storage devices and SANs over standard Ethernet-based TCP/IP networks.
Interswitch link (ISL)	The link joining two Fibre Channel switches through expansion ports.
IP address	A unique string of numbers separated by periods that identifies each computer using the Internet Protocol to communicate over a network.
IoT	See <i>Internet of Things</i> .
iPDU	See <i>Intelligent power distribution unit</i> .
IPMI	See <i>Intelligent Platform Management Interface</i> .
IRF	See <i>HPE Intelligent Resilient Framework</i> .
iSCSI	See <i>Internet Small Computer System Interface</i> .
ISL	See <i>Interswitch link</i> .

---

## J

---

Jitter smoothing	Patent-pending technology that levels and balances the frequency fluctuation created when a processor runs in Turbo Boost mode.
Jumbo frame	Permits up to a 9K byte MTU when running Ethernet I/O traffic. Enables networks to achieve greater throughput performance and improve CPU utilization, and are particularly useful for database transfer and tape backup operations.

---

## K

---

Kerberos	A network authentication protocol that uses secret key cryptography to provide authentication between clients and servers. It is used in Active Directory implementations to enable single sign-on (SSO).
----------	---

---

## L

LACP	See <i>Link Aggregation and Control Protocol</i> .
LAG	See <i>Link Aggregation Group</i> .
LAN on motherboard (LOM)	Typically, a chip or chipset embedded on a PC system board to handle network connections; less expensive than the traditional network interface cards; also frees up a PCI slot.
Large form factor (LFF)	A size characteristic of hard disk drives. An LFF drive is 3.5 inches in diameter.
Layer 2 (L2) switch	Forwards traffic based on the frame's Data Link Layer information, specifically the hardware address, which is called the Media Access Control, or MAC address.
Layer 3 (L3) switch	Routes traffic based on Network Layer information. To route traffic, Layer 3 switches must have the appropriate IP route. Layer 3 switches support static routes and routes learned through routing protocols.
LFF	See <i>Large form factor</i> .
Linear Tape Open (LTO)	Tape drives use Linear Tape Open, or LTO technology. For additional security, LTO tape drives have built-in hardware encryption, for high security without any loss of performance or capacity.
Link Aggregation Group (LAG)	An instance of link aggregation that uses Link Aggregation and Control Protocol (LACP) to signal the establishment of the channel between two devices. A LAG increases the available bandwidth between two devices and creates one logical path out of multiple physical paths.
Link Aggregation and Control Protocol (LACP)	A method to control the bundling of several physical Ethernet ports together to form a single logical channel; allows a network device to negotiate an automatic bundling of links by sending LACP packets to the peer. 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.
Link Layer Discovery Protocol (LLDP)	A protocol that enables devices from different vendors to be able to discover one another and exchange configuration information.
LLDP	See <i>Link Layer Discovery Protocol</i> .
Logical disk	A partition that may contain a file system or database or be used by a volume manager or hypervisor to present higher-order volumes. Volumes are backed up by capacity that has been created from lower-level capacity.
Logical interconnect	An HPE OneView construct that represents available networks, uplink sets, and stacking links for a set of physical interconnects in a single enclosure.
Logical interconnect group	Within HPE OneView, acts as a recipe for creating a logical interconnect representing available networks, uplink sets, stacking links, and interconnect settings for a set of physical interconnects in a single enclosure.

Logical unit	The entity within a SCSI target that executes I/O commands. Each logical unit exported by an array controller corresponds to a volume.
Logical unit number (LUN)	A number used to identify a logical disk or volume on a storage system. It is used in the SCSI protocol to access an array within a target. A LUN acts as a logical reference that is used to refer to a volume or logical disk on the SAN, so that it can be presented to and accessed by a target device.
Logical volume	The storage medium that is associated with a logical disk. It typically resides on one or more hard drives.
Load reduced dual in line memory module	A type of DIMM that replaces the hardware register from an RDIMM with a chip that isolates the electrical loads. This results in greater density of DIMMs in the server.
LOM	See <i>LAN on motherboard</i> .
LRDIMM	See <i>Load reduced dual in line memory module</i> .
LTO	See <i>Linear Tape Open</i> .
LUN	See <i>Logical unit number</i> .
<b>M</b>	
MAC address	See <i>Media Access Control address</i> .
Management LAN	A LAN dedicated to the communications necessary for managing systems. Typically of moderate bandwidth (10/100 Mbps) and secured.
Management protocol	A set of protocols, such as WBEM, HTTP, or SNMP, used to establish communication with discovered systems.
Maximum transmission unit (MTU)	The maximum size IP frame that a switch can receive for Layer 2 frames inbound on a port. A Layer 2 switch drops any inbound frames larger than the MTU allowed on the port.
Media Access Control (MAC) address	A unique identifier assigned to network interfaces for communications on the physical network segment.
Meshed stacking	Allows aggregation of up to five switches to form a fully meshed stack for resiliency and management in a single interface.
Microsoft Hyper-V	Microsoft type-1 hypervisor for deploying and serving virtual computers. It can create virtual machines on x86-64 systems running Windows.
Microsoft SQL Server	A relational database management system developed by Microsoft.

---

Mirroring	Refers to the real-time operation of copying data, as an exact copy, from one location to a local or remote storage medium. A mirror is an exact copy of a dataset. Data mirroring is most frequently used when multiple exact copies of data are required in multiple locations.
MLAG	See <i>Multi-chassis link aggregation</i> .
Monitoring	The process of collecting server, storage, network, and application availability and performance data across the physical and virtual infrastructure.
MPIO	See <i>Multipath I/O</i> .
MPLS	See <i>Multiprotocol Label Switching</i> .
MTU	See <i>Maximum transmission unit</i> .
Multi-chassis link aggregation (MLAG)	Extends the capabilities of link aggregation by allowing a downstream switch or host to connect to two switches configured as an MLAG domain.
Multipathing	A routing technique of using multiple alternative paths through a network. This can yield a variety of benefits such as increased bandwidth, improved security, or fault tolerance.
Multipath I/O	A fault-tolerance and performance-enhancement technique that defines more than one physical path between the CPU in a computer system and its storage devices through the buses, controllers, switches, and bridge devices that connect them.
Multiprotocol Label Switching (MPLS)	A type of data-carrying technique for high-performance networks. MPLS directs data from one network node to the next based on short path labels rather than long network addresses. This avoids complex lookups in a routing table.
Multi-threading	The ability of an application and operating system to allow parallel computing by dividing processing between multiple processors or cores.

---

**N**

---

NAND flash	Type of storage used in SSDs and other storage devices that retains data without power. It is more cost-effective than DRAM but has lower performance and endurance.
NAS	See <i>Network-attached storage</i> .
National Institute of Standards and Technology (NIST)	An organization dedicated to promoting U.S. innovation and industrial competitiveness by advancing measurement science, standards, and technology. HPE applies National Institute of Standards and Technology security controls to a solution stack of storage, networking, servers, and software. Federal agencies follow NIST guidelines for cybersecurity protection and the private sector is beginning to do the same.

---

---

NDMP	See <i>Network Data Management Protocol</i> .
Nearline drives	Drives that are typically used for secondary storage because the data stored on them does not need to be highly available. Users do not access data stored on nearline drives very often, but they do need to access it on demand.
Network	Medium for interconnecting a specific set of server links; typically relates to a single L2 broadcast domain; may map to a specific Ethernet VLAN, a multi-tenant network, or a specific IP subnet.
Network-attached storage (NAS)	Uses industry-standard protocols to connect to host devices over a network. Consists of a specialized server or storage device that connects directly to the network. A NAS solution provides a flexible, intelligent, simple-to-manage solution for file-and-print and application-storage consolidation.
Network core layer	Provides high-speed, highly redundant forwarding services to move packets between distribution-layer devices in different regions of the network. Core switches and routers are usually the most powerful, in terms of raw forwarding power. Core network devices manage the highest-speed connections.
Network Data Management Protocol (NDMP)	The protocol used to transport data between network attached storage (NAS) devices and backup devices.
Network File System (NFS)	A protocol used by client systems running Unix or Linux operating systems to access files on a network storage file server.
Network functions virtualization (NFV)	A way to use software to virtualize network services instead of proprietary, dedicated hardware.
Network Time Protocol	Network protocol for time synchronization between devices.
Network virtualization	Takes physical network devices and adds an abstraction layer that brings these devices into a single network resource pool from which virtual networks can be provisioned.
Network interface card (NIC)	An expansion card that enables a computer to connect to a network; such as a home network or the internet using an Ethernet cable. It is also referred to as an Ethernet card or a network adapter.
Network overlay	A virtual network that runs independently on top of another one. Overlay networks shift the network complexity from the physical network to the overlay network with software, and provide network resources from a single management point without changing the physical network.
Network virtualization using generic routing encapsulation (NVGRE)	A network overlay for Microsoft environments. In NVGRE, a virtual Layer 2 overlay network, or tunnel, is automatically created on top of a Layer 3 network. Virtual machine-to-virtual machine traffic crosses this virtual network, which means that a virtual machine can be freely migrated across the data center over an overlay network without reconfiguration. Inserting the NVGRE header on an Ethernet frame, as well as calculating the new checksum value, creates a tremendous burden on throughput, host CPU utilization, and power consumption. This limits the number of VMs per physical server platform.

---

---

NFS	See <i>Network File System</i> .
NFV	See <i>Network functions virtualization</i> .
NIC	See <i>Network interface card</i> .
NinjaCrawler	A host-based deduplication and compression estimation tool. It supports Windows, Linux, and VMware platforms, and scans customer volumes to identify data reduction with 3PAR dedupe and compression.
NinjaSTARS	A capacity and performance sizing tool that helps design cost-effective storage solutions that meet customer's service level agreements. It was developed for worldwide channel partners to assess legacy storage environments for customers interested in HPE storage arrays. "STARS" stands for Storage Assessment, Recommendation, and Sizing.
NIST	See <i>National Institute of Standards and Technology</i> .
Non-uniform memory access (NUMA)	A computer memory design choice for processors. Because the memory access time depends on the location of the memory relative to the processor, with NUMA, accessing local memory is faster than accessing memory associated with other NUMA nodes. The number of processors within a NUMA node depends on the hardware vendor.
Nonvolatile dual in line memory module (NVDIMM)	A nonvolatile persistent memory solution that combines NAND flash, DRAM, and an optional power source into a single memory subsystem. NVDIMMs provide persistent memory and are installed in standard DIMM slots.
Non-Volatile Memory Express (NVMe)	A host controller interface and storage protocol created to accelerate the transfer of data between enterprise and client systems and solid-state drives over a computer's high-speed PCIe bus.
Nonvolatile random access memory (NVRAM)	A type of RAM that retains data after the power is turned off.
North/south traffic	The client server network traffic that goes between users in a branch office and the data center that hosts the application that they are accessing.
NTP	See <i>Network Time Protocol</i> .
NUMA	See <i>Non-uniform memory access</i> .
NVDIMM	See <i>Nonvolatile dual in line memory module</i> .
NVGRE	See <i>Network virtualization using generic routing encapsulation</i> .
NVRAM	See <i>Nonvolatile random access memory</i> .
NVMe	See <i>Non-Volatile Memory Express</i> .

---

---

**O**

---

OLTP	See <i>Online transaction processing</i> .
Online transaction processing (OLTP)	A description of applications that require a database back end. This includes workloads that involve a high number of user-based, transactions running on a single server with a cohosted database component.
Operating system	The software that supports a computer's basic functions, such as scheduling tasks, executing applications, and controlling peripherals. It is the most important program that runs on a computer.
Open Systems Interconnection (OSI) Model	A conceptual model that was a result of the OSI project at the International Organization for Standardization. The seven layers of the model are interconnected, in that each layer serves the layer above and below it. A key reason for using the OSI model is that it provides a standard for communication between devices, regardless of their infrastructure or environment.
OSI	See <i>Open Systems Interconnection Model</i> .
Oversubscription	A condition in which a router must queue packets in order to prevent dropping them. The amount of queuing available on an interface determines the amount of momentary oversubscription that the router can tolerate on that interface before dropping packets and causing performance degradation.

---

**P**

---

Packet buffer	Memory space set aside for storing packets awaiting transmission over a network or that have been received over a network. Packet buffering creates a reserve of packets that can be used during packet transmission delays or retransmission requests.
Parity checking	Used for memory error detection and correction. It uses parity bits to check accuracy of data transmission.
PCIe	See <i>Peripheral Component Interconnect Express</i> .
Peripheral Component Interconnect Express (PCIe)	A standard type of connection for internal devices in a computer. Generally, PCI Express refers to the actual expansion slots on the motherboard that accept PCIe-based expansion cards and to the types of expansion cards themselves. Although computers may contain different types of expansion slots, PCI Express is considered the standard internal interface.
Persistent memory	Nonvolatile memory that delivers the performance of memory with the persistence of traditional storage. Persistence is achieved by combining the speed and endurance of DRAM with the nonvolatility of NAND flash.
PDU	See <i>Power distribution unit</i> .

---

Peer Motion	StoreServ software that balances workloads across multiple arrays in the same location or across geographies, and to shift workloads between systems dynamically without affecting application performance or availability.
Physical capacity	The maximum amount of storage available on an array or pool, as constrained by the physical structure of the storage system.
PoE	See <i>Power over Ethernet</i> .
PoE+	See <i>Power over Ethernet Plus</i> .
POST error messages	Diagnostic messages that isolate server-related logic failures and indicate the component that needs to be replaced. Most server hardware failures are accurately identified during POST.
Power distribution rack	A data center power distribution rack that improves power management by moving power distribution to the row level. It decentralizes power, improves cable management, decreases diagnostic time for problems, and saves installation costs.
Power distribution unit (PDU)	Rack device that distributes conditioned AC or DC power within a rack.
Power over Ethernet (PoE)	A technology for wired Ethernet LANs that allows the electrical current necessary for the operation of each device to be carried by the data cables rather than by power cords.
Power over Ethernet Plus (PoE+)	A technology that is similar to PoE, except that the PoE+ standard allows for a maximum of 25.5 watts to be carried through the Ethernet cables, while the PoE standard allows for a maximum of 15.4 watts.
Private cloud	A cloud infrastructure that is operated solely for a single organization. Can be owned or operated by that organization or by a third party. Can be on or off premises.
Processor	An integrated chip on a computer system board that handles requests from the programs that run the computer. A processor is also known as a <i>CPU</i> .
Product Bulletin	A convenient central resource for technical overviews and specifications for hardware and software in an HPE solution. The downloadable Product Bulletin application is loaded with features to aid with the purchase, sale, and support of HPE products. It contains updated QuickSpecs and a Quick Quote resource that is integrated with the Internet Price List feature. This allows users to create, save, and print quotes with up-to-date prices.
Provisioning	The process of deploying an operating system, application, or service from a template.
Public cloud	A cloud infrastructure that is made available to the general public or a large industry and is owned and operated by the organization selling cloud services.

---

---

**Q**


---

QoS	See <i>Quality of Service</i> .
Quality of Service (QoS)	A combination of qualitative and quantitative factors such as uptime, response time, and available bandwidth, that collectively describe how well a system performs.
QuickSpecs	A convenient central HPE resource that provides technical overviews and specifications for HPE hardware and software.

---

**R**


---

Rack	A set of components cabled together to communicate between themselves. A rack is a container for an enclosure.
RAID	See <i>Redundant array of independent disks</i> .
Raw storage space	A measure of a drive's total storage capacity. The raw storage space in an array is the sum of each drive's raw capacity.
Registered dual in line memory module	A type of DIMM with a memory technology that uses a hardware registry for better reliability.
RDIMM	See <i>Registered dual in line memory module</i> .
RDMA	See <i>Remote direct memory access</i> .
RDMA over converged Ethernet (RoCE)	An accelerated I/O delivery mechanism that allows data to be transferred directly from the memory of the source server to the memory of the destination server, bypassing the operating system kernel. RoCE has better I/O performance and CPU efficiency than TCP/IP.
Redundancy	A system design in which a component is duplicated so if it fails there is a backup.
Redundant array of independent disks (RAID)	A RAID volume consists of more than one drive, but appears to the operating system to be a single logical disk. RAID improves performance by disk striping, which involves partitioning each drive's storage space into units and placing those on multiple disks.
Remote office/branch office (ROBO)	A term used to refer to an off-site office with an external connection to the organization's WLAN or LAN.
Remote direct memory access (RDMA)	A method of moving buffers between two applications across a network efficiently. RDMA works over TCP/IP with RDMA-enabled Ethernet NICs. It provides increased throughput, lower latency, and lower CPU utilization than traditional TCP communications.

---

---

Representational State Transfer (REST)	An architectural style consisting of a coordinated set of constraints applied to components, connectors, and data elements within a distributed hypermedia system. REST describes a way to create, read, update, or delete information on a server by using simple HTTP calls such as POST, GET, PUT, and DELETE.
Replication	The process of copying data from one location to another. It helps an organization possess up-to-date copies of its data in the event of a disaster.
Request for information (RFI)	A standard business process used to collect written information about the capabilities of various suppliers.
Request for proposal (RFP)	A document that an organization posts to elicit bids from potential vendors for a desired solution.
Resiliency	Ability to provide and maintain an acceptable level of service in the face of faults and challenges to normal operations.
REST	See <i>Representational State Transfer</i> .
RESTful	Applications conforming to REST constraints are known as RESTful.
Return on investment (ROI)	A measure of how much a company earns for each dollar (or monetary unit) of investment it makes. It is used to evaluate the efficiency of an investment or to compare the efficiency of a number of different investments.
RFI	See <i>Request for information</i> .
RFP	See <i>Request for proposal</i> .
ROBO	See <i>Remote office/branch office</i> .
RoCE	See <i>RDMA over converged Ethernet</i> .
ROI	See <i>Return on investment</i> .

---

**S**

---

SAFanalyze	See <i>Storage Assessment Foundry analyze tool</i> .
SAFcollect	See <i>Storage Assessment Foundry collect tool</i> .
SAN	See <i>Storage area network</i> .
SAS	See <i>Serial-Attached SCSI</i> .
SATA	See <i>Serial Advanced Technology Attachment</i> .
Scalability	Ability to expand a computing resource to match the current workload.
SCMB	See <i>State change message bus</i> .

---

---

Scope of work	A project description generally prepared as part of a proposal. It typically includes the overall time frame of the project, completion milestones for each aspect of the project, and the required resources to complete the work.
SDN	See <i>Software-defined networking</i> .
Secure Boot	A security standard to help ensure that all PCs boot using only software that is trusted by the PC manufacturer.
Secure HTTP	An extension to the HTTP protocol that supports sending data securely over the web.
Secure Shell	A network protocol that allows users to log into another system over a network and execute commands on that system. It also enables users to move files from one system to another, and provides authentication and secure communication over insecure channels.
Secure Sockets Layer (SSL)	A standard protocol layer that lies between HTTP and TCP and provides privacy and message integrity between a client and server. A common usage is to provide authentication of the server, so clients can be assured of communication with that server.
Serial Advanced Technology Attachment (SATA)	A drive technology that is built for reliability and larger capacity for non mission-critical server applications and storage environments. These high-capacity drives provide the lowest cost per GB, and the best price advantage for non-mission critical applications with low workloads.
Serial Attached SCSI (SAS)	A drive technology that satisfies the data center requirements of scalability, performance, reliability, and manageability. SAS midline drives provide the lowest dollars per gigabyte and economical reliability and performance.
Server	A computer designed to process requests and deliver data to other computers over a LAN or the internet.
Server blade	A dense server system containing processor, memory, and network connections that can be inserted into an enclosure to share power supplies, fans, switches, and other components with other blades.
Server Message Block (SMB)	A protocol used by client systems running Windows or Mac OS X operating systems to access files on a network storage file server.
Server Message Block (SMB) Direct	A technology that supports the use of network adapters that have remote direct memory access, or RDMA, capability. RDMA network adapters can function at full speed with very low latency and use very little processor power.
Server profile	A description of a server including its identity (serial number and UUID), hardware requirements (CPU, memory, and so on), configuration settings, and connections to data center networks and fabrics in a manner independent of the underlying hardware.
Server virtualization	A single physical server hosts a number of VMs, each with an operating system, and includes virtualized devices such as hard drives, NICs, optical drives, and printers.

---

---

Service level agreement (SLA)	Contract between a consumer of computing resources and the provider of those resources. It includes the minimum level of service, particularly in terms of availability.
Service Pack for ProLiant (SPP)	A comprehensive collection of firmware and system software components, all tested together as a single solution stack that includes drivers, agents, utilities, and firmware packages for HPE ProLiant servers, controllers, storage, server blades, and enclosures.
SFF	See <i>Small form factor</i> .
Simple Network Management Protocol (SNMP)	An internet-standard protocol which monitors network and device health using agents.
Simple Object Access Protocol (SOAP)	A lightweight protocol for exchange of information in a decentralized, distributed environment.
Single point of failure	A device or connection in a system that can cause the entire system to fail.
Single sign-on (SSO)	Ability to log in once with a single set of credentials and gain access to multiple systems without having to log in to each individually.
SLA	See <i>Service level agreement</i> .
Small form factor (SFF)	A size characteristic of hard disk drives. An SFF drive is 2.5 inches in diameter.
Small to medium-sized business (SMB)	A business which, due to its size, has different IT requirements and often faces different IT challenges than large enterprises, and whose IT resources, including budget and staff, are often highly constrained.
SmartSSD Wear Gauge report	Contains information about the current usage level and remaining expected lifetime of SSDs attached to the system.
SMB	See <i>Small to medium-sized business</i> or <i>Server Message Block</i> .
Snapshots	Enable fast recovery of data at a known point in time. With a snapshot, customers can quickly make a point-in-time copy of the primary version of the data and save it for future use. Because of the speed of this operation and the low disruption to productivity, customers can take frequent snapshots and thus create many recovery points.
SNMP	See <i>Simple Network Management Protocol</i> .
SNMP trap	An asynchronous event generated by an SNMP agent that the system uses to communicate a fault.
Soft errors	Randomly occurring events that account for most errors in memory. Some causes are faulty system boards, physical memory errors, static electricity, and incorrect memory system timings. Also known as <i>transient errors</i> .
Spanning Tree Protocol (STP)	A networking protocol that creates a single path over a network. STP prevents loops from occurring, even if there are multiple paths to the same destination.

---

---

SOAP	See <i>Simple Object Access Protocol</i> .
Software-defined networking (SDN)	An approach to networking that enables network administrators to manage network services through abstraction of lower-level functionality. SDN abstracts the control plane and presents a single, simple interface for interacting with the network dynamically.
Solid state drive (SSD)	A drive that reduces power consumption for applications that require high random read and write IOPs performance. SSDs are categorized as read-intensive, mixed-use, and write-intensive so users can choose the SSD that meets the demands of their workload.
SPP	See <i>Service Pack for ProLiant</i> .
SRAM	See <i>Static random access memory</i> .
SSD	See <i>Solid state drive</i> .
SSH	See <i>Secure Shell</i> .
SSL	See <i>Secure Sockets Layer</i> .
SSO	See <i>Single Sign-On</i> .
State change message bus (SCMB)	An interface that uses asynchronous messaging to notify subscribers of changes to managed resources, both logical and physical.
Static random access memory (SRAM)	Nonvolatile memory used for high-speed cache. SRAM is able to retain data by using an alternative source of power, such as a battery.
Scope of work	A project overview prepared for a proposal that captures the plan, time frame, required resources, and completion milestones of a project.
Storage area network (SAN)	A network or subnetwork that connects data storage devices with associated data servers. A SAN is typically part of an overall network of computing resources. It features a dedicated private network to carry storage traffic using one of a variety of protocols, such as Fibre Channel or iSCSI.
Storage Assessment Foundry analyze tool	A tool that provides a rapid summary of key metrics that are required as input into the solution design process and NinjaSTARS. It takes less than five minutes to analyze and five minutes to get a summary report.
Storage Assessment Foundry collect tool	A tool that collects information from unfamiliar devices to better understand a customer's current environment.
Storage federation	Technology that enables online, non-disruptive movement of storage volumes between arrays.
Storage virtualization	Adds an abstraction layer that brings physical storage devices into a single storage resource pool from which virtual disks are provisioned.
STP	See <i>Spanning Tree Protocol</i> .

---

---

**T**


---

TCO	See <i>Total cost of ownership</i> .
TCP/IP	See <i>Transmission Control Protocol/Internet Protocol</i> .
Telnet	<i>Terminal</i> or <i>teletype protocol</i> . A network protocol used to provide a CLI for communicating with a device. Enables an administrator to use a computer to log on remotely to another computer that is part of the same network
TFTP	See <i>Trivial File Transfer Protocol</i> .
Tiered storage	A storage networking method in which data is stored on various types of media based on usage, performance, availability and recovery requirements.
Top of rack (ToR)	A networking design in which servers in the same or adjacent rack connect to a network switch installed inside the rack. Although the actual physical location of the switch does not necessarily need to be at the top of the rack, it is the most common location because it is easier to access and simplifies cable management. ToR switches provide a unified, converged fabric over 10-Gb Ethernet for LAN and SAN traffic. This unification enables network consolidation, reducing the number of adapters and cables required and eliminating redundant switches.
ToR	See <i>Top of rack</i> .
Total cost of ownership (TCO)	A financial estimate intended to help buyers and owners determine the direct and indirect costs of a product or system.
Trivial File Transfer Protocol	A simple protocol used to copy files across a network. Because it does not provide authentication or security mechanisms, it is not frequently used to send files over the internet. However, it can be useful during the boot process because it requires a very small amount of memory.
Transmission Control Protocol/Internet Protocol	A set of communication protocols used to interconnect network devices on the internet. TCP/IP can also be used as a communications protocol in a private network.
<hr/> <b>U</b> <hr/>	
UEFI	See <i>Unified Extensible Firmware Interface</i> .
UI	See <i>User interface</i> .
Unified Extensible Firmware Interface (UEFI)	Specification that defines the software interface between an operating system and the platform firmware. Unified EFI (UEFI) supersedes the original EFI specification and replaces the BIOS firmware interface.
Uninterruptible power supply (UPS)	A device that allows a computer to keep running when the primary power source is lost. It also may provide protection from power surges.

---

UPS	See <i>Uninterruptible power supply</i> .
-----	---

---

Usable storage space	The amount of storage remaining after RAID and other overhead have been subtracted from the raw storage space.
----------------------	--

---

User interface (UI)	A system through which users interact with a computer.
---------------------	--

---

**V**

---

VDI	See <i>Virtual desktop infrastructure</i> .
-----	---

---

Virtual desktop infrastructure (VDI)	An data center design that places user desktops in the data center and replaces user PCs with thin clients; reduces the need for desktop support resources and lowers desktop computer capital expenditures.
--------------------------------------	--

---

Virtual device	An emulation of a physical device. This emulation, used as a device by an virtual machine, effectively maps a virtual device to an entity on the VM host.
----------------	---

---

Virtual disk	A virtual logical disk or volume in storage virtualization applications.
--------------	--

---

Virtual extensible local area network (VXLAN)	A network overlay for VMware and Linux environments. A virtual Layer 2 overlay network, or tunnel, is automatically created on top of a Layer 3 network. Virtual machine-to-virtual machine traffic crosses this virtual network, which means that a virtual machine can be freely migrated across the data center over an overlay network without reconfiguration. Inserting the VXLAN header on an Ethernet frame and calculating the new checksum value creates a burden on throughput, host CPU utilization, and power consumption.
---	---

---

Virtualization	An approach to IT that partitions servers, storage, and networking into several virtual machines.
----------------	---

---

Virtual local area network (VLAN)	A standard that enables network administrators to group end users by logical function rather than by physical location. Created on switches to segment networks into smaller broadcast domains, enhance network security, and simplify network management.
-----------------------------------	--

---

Virtual machine (VM)	A software simulation of a fully operational computer that can have its own operating system, storage, and applications.
----------------------	--

---

Virtual private network (VPN)	Extends a private network across a public network, and enables users to send and receive data across shared or public networks as if their computing devices were directly connected to the private network. Applications running across the VPN may therefore benefit from the functionality, security, and management of the private network.
-------------------------------	---

---

Virtual SAN (vSAN)	A software-defined partition in a storage area network that enables you to isolate traffic within specific portions of that storage area network. It is used primarily in virtualization environments and in cloud computing.
--------------------	---

---

Virtual storage	A method of mapping logical storage requests to physical storage (disks).
-----------------	---

---

---

Virtual storage appliance (VSA)	A storage controller that runs on a virtual machine to create shared storage without the cost of additional hardware.
Virtual switching framework	A technology that virtualizes two physical devices in the same layer into one Virtual Fabric. This provides high availability and scalability.
Virtual tape library (VTL)	A data storage virtualization technology. It presents a storage component such as hard disk drives as tape libraries or tape drives to be used with existing backup software.
VisioCafe	A web-based tool that provides graphics to create diagrams and schemas for proposed solutions. HPE provides storage subsystem graphics in this tool as well as graphics of servers or HPE networking devices. All graphics are free to download from the VisioCafe website.
VLAN	See <i>Virtual local area network</i> .
VLAN tagging	A method used to identify the VLAN that the frame belongs to in a network with multiple VLANs.
VM	See <i>Virtual machine</i> .
VM host	A server running HPE Integrity Virtual Machines, VMware ESX, VMware ESXi, or Microsoft Hyper-V that provides multiple virtual machines, each running its own instance of an operating system.
VMware vCenter	VMware enterprise-level virtualization management product.
VMware ESXi	VMware enterprise-class, type-1 hypervisor for deploying and serving virtual computers.
VPN	See <i>Virtual private network</i> .
Volume	A logical drive provided by a storage pool that can be presented to a host system.
VSA	See <i>Virtual storage appliance</i> .
vSAN	See <i>Virtual SAN</i> .
VSF	See <i>Virtual switching framework</i> .
VTL	See <i>Virtual tape library</i> .
VXLAN	See <i>Virtual extensible local area network</i> .

---

---

**W**

---

WiFi	A technology allowing computers, smartphones, or other devices to connect to the internet or communicate with one another wirelessly within a particular geographic area.
WINS	See <i>Windows Internet Name Service</i> .
Windows Internet Name Service (WINS)	The Microsoft computer name registration and resolution service that maps computer NetBIOS names to IP addresses.
Wireless local area network (WLAN)	A wireless distribution method for two or more devices that uses high-frequency radio waves and often includes an access point to the internet. A WLAN allows users to move around the coverage area, often a home or small office, and maintain a network connection.
Wizard	A sequential series of pages that transforms a complex task into simple steps and guides a user through them. The wizard makes sure that the user provides all required information and does not omit any steps.
WLAN	See <i>Wireless local area network</i> .
Workload	A collection of processes in a stand-alone server, nPartition compartment, virtual partition compartment, or virtual machine compartment.
Workload matching	Time-saving technology that uses preconfigured workload profiles to tune internal server resources automatically.
Workload-optimized server	A strategy based on the fact that certain workloads perform better on specific compute architectures. Organizations with stringent performance requirements need servers designed for specific tasks. By matching the workload to the right compute platform, organizations can gain web-scale performance, reliability and scalability while reducing total cost of ownership.
World Wide ID (WWID) or World Wide Name (WWN)	A unique identifier assigned to a Fibre Channel device.
WWID	See <i>World Wide ID</i> or <i>World Wide Name</i> .
WWN	See <i>World Wide ID</i> or <i>World Wide Name</i> .

---

**X**

---

---

**Y**

---

---

**Z**

---





---

# Planning and Designing HPE SMB Solutions Worksheet

## Appendix 2

### Introduction

Use this worksheet as you go through the activities in this training to gather the details about the solution you would propose for IDJT HealthCare. Be sure to refer back to the customer scenarios throughout the course to learn about the business and IT requirements of this customer.

### Customer overview

IDJT HealthCare is located in the state of Pennsylvania in the United States. For more than 80 years, the company has provided renowned physicians and caregivers with the most appropriate technology across medical focus areas including cancer, cardiology, sports medicine, pain management, and sleep disorders.

IDJT HealthCare is opening a remote location that will accommodate 10 employees responsible for real-time patient and lab data as well as provider information retrieval. The company plans to install physical security measures at the remote location including an access control system, pin codes, and video surveillance.

IDJT HealthCare has 150 employees, including doctors, nurses, administrators, and other patient care professionals. In addition, they have 20 IT professionals on staff. The IT infrastructure includes one HPE ProLiant DL380p Gen8 management server and two HPE ProLiant DL580 Gen9 database servers in their main data center.

## Worksheet

Use this worksheet to record the selections you make during the activities in this course. You will use this information to produce a statement of work for IDJT HealthCare.

### Data center solution components

Component	Technical details	Key benefits	Business value
<b>Network</b>			
	Main location		
	Remote location		
<hr/>			
<b>Servers</b>			
	Main location		
	Remote location		
<hr/>			
<b>Storage</b>			
	Main location		
	Remote location		
<hr/>			
<b>Rack</b>			
	Main location		
	Remote location		
<hr/>			

<b>Component</b>	<b>Technical details</b>	<b>Key benefits</b>	<b>Business value</b>
<b>Power</b>			
	Main location		
	Remote location		
<b>Licensing</b>			
	Main location		
	Remote location		
<b>Management</b>			
	Main location		
	Remote location		
<b>Other components</b>			
	Main location		
	Remote location		

## Notes