Architecture guide

Designing the HP Converged Campus network

Table of contents

Introduction

Campus network trends

Campus network used to be very simple. Most client's PCs were provided and controlled by the company's IT department. A large percentage were desktops connected to an Ethernet switch and some laptops were connected either to an Ethernet switch or a WLAN AP. The requirements were basically access to the servers, data bases, printers, WAN and Internet on one hand and performance and security on the other. Network management was just an overlay to enable some monitoring and make some tasks easier.

Today a new paradigm is emerging driven by several strong trends:

- Virtual meetings
- BYOD
- Cloud computing and storage
- IEEE 802.11ac WLAN

Virtual meetings

Teams are no longer in one place, with employees working from home, other cities, and even other countries and continents; meetings have become virtual and require the use of a new types of applications. The requirements include—at a minimum—voice, video, and desktop sharing.

These applications are already available and are used every day by more and more organizations. Even employees working in the same location prefer to stay at their desks because they like the agility of meeting setup and the ability of using the full power of their PCs during these meetings. The same applies to meetings with customers, vendors, and partners.

Virtual meetings also enabled a new type of collaboration between colleagues, working together for hours without leaving their desks, sharing resources, and achieving new levels of quality in their jobs while making optimal use of their time.

Collaboration applications used for this type of meetings break the client-server traffic pattern, usually called north-south, by creating a large amount of client-client or east-west traffic. They require lower latency and jitter than traditional database applications.

Bring your own device

IT departments are now required to accept mobile devices into the corporate network. These devices can be smartphones and tablets provided by the company or brought by the employees and guests.

Enterprises can only provide Internet access to mobile devices or they can create specialized apps. Corporate portals, blogs, messaging, email, and collaboration applications can be made accessible using these apps.

And hospitality can be brought to a new level of quality and interaction with apps and services delivered directly to the guest's mobile device.

The incorporation of mobile devices and apps to any campus network implies a multiplication of the number of wireless clients. Additionally, specific verticals like healthcare can create apps to access high-definition pictures and video streaming, that require high bandwidth on the wireless LAN.

Cloud computing and storage

Many users are choosing to store their documents and presentations on cloud drives to be able to access them from more than one device and share them with other users.

As more than one device per user may be accessing the cloud at the same time, bandwidth requirements may be larger.

IEEE 802.11ac WLAN

The new IEEE 802.11ac standard allows wireless clients to access the network with performance similar to their wired connections. Technologies like MIMO and Multiuser MIMO (MU-MIMO) provide a better quality and a better usage of wireless channels.

Because of this, and the proliferation of BYOD, some organizations are deploying a large number of APs per switch, meaning that the average utilization of those switch ports is several times higher than that of a typical individual user.

With many APs per access switch, comes larger amounts of traffic and the need for more bandwidth in the uplink.

Converged Campus

When all the trends described in the previous section are combined, a new reality emerges: traditional networks can neither satisfy the needs nor cope with the complexity of today and tomorrow's campus networks, and a new approach is needed.

The HP Converged Campus initiative focuses on satisfying the technical requirements of these networks without sacrificing agility, flexibility, and manageability. It unifies the campus access bringing high-speed wired and IEEE 802.11ac wireless LANs together with a single set of access control policies. Offers simplified deployment and maintenance and a set of SDN apps that allow the network to adapt to requirements in real time.

The converged campus solution includes:

- 2 or 3-tier optimized designs
- High performance, fully featured core and aggregation layer switches
- Stackable top-of-rack (ToR) switches for local server farms
- High-speed Ethernet backbone links
- WLAN controller modules and appliances
- IEEE 802.11ac dual radio access points and walljacks
- Modular and stackable access switches with high speed uplinks, PoE+, and OpenFlow/SDN support
- SDN controller and applications
- Intelligent Management Center (IMC), a modular network management solution for end-to end monitoring, maintenance, configuration and administration including WLAN management, unified access control, and BYOD
- Security features, modules, and appliances to build a secure infrastructure and provide firewalling and intrusion prevention
- Fully featured WAN routing solutions

Figure 1. Converged campus network solution outline

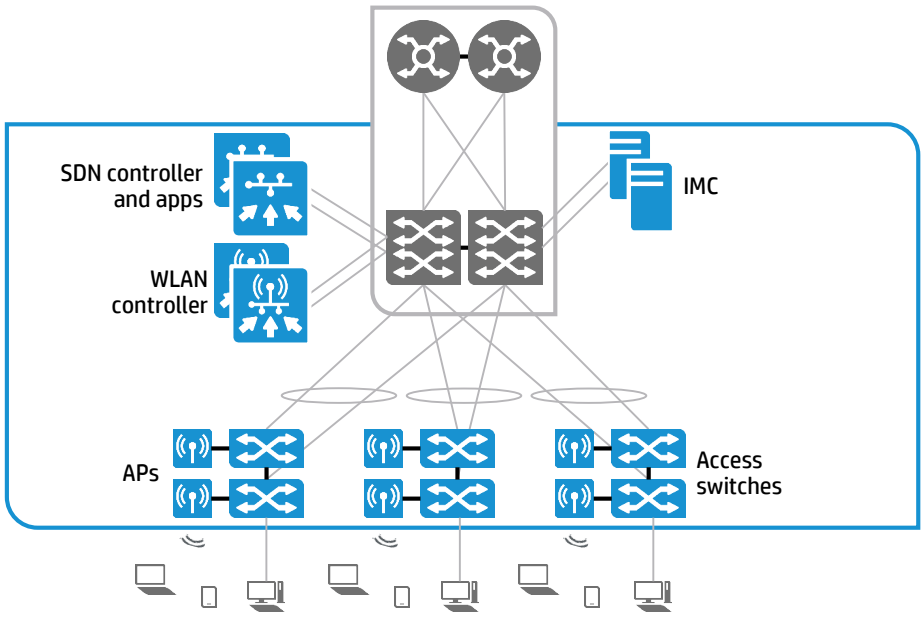

This architecture guide describes the Converged Campus edge components, and how they are integrated in a complete, converged, campus client access solution. It focuses on medium-to-large campuses.

It provides a detailed description of:

- Client access infrastructure
- Unified access control and BYOD
- Simplified deployment and configuration
- SDN applications

Designing the client access infrastructure

This section describes the converged campus client access infrastructure components, operation, and best practices.

The client access infrastructure is composed of access devices like access switches and access points; plus controllers, security, and network management applications.

Components

Depending on the particular requirements of each campus, the client access solution can be designed using different product combinations. Table 1 shows a selection of IEEE 802.11ac APs, WLAN controllers, modular and stackable campus access switches, network management platforms and modules, and SDN products.

Table 1. Converged campus client access featured products

Category	Products
802.11ac access points	• HP 560 802.11ac Dual Radio Access Point Series • HP 527 Dual Radio 802.11ac Unified Wired-WLAN Walljack • HP 525 802.11ac Dual Radio Access Point Series
WLAN controllers	• HP 870 Unified Wired-WLAN Appliance Series • HP 10500/7500 20G Unified Wired-WLAN Module • HP 850 Unified Wired-WLAN Appliance Series
Modular access switches	• HP 5400R zl2 Switch Series • HP 7500 Switch Series
Stackable access switches	• HP 5500 HI Switch Series • HP 5130 EI Switch Series • HP 3800 Switch Series • HP 2920 Switch Series
Network management platforms and modules	• HP IMC Standard Platform • HP IMC Enterprise Platform • HP IMC WLAN Services Manager • HP IMC User Access Manager (UAM) • HP IMC Endpoint Admission Defense (EAD) • HP IMC Network Traffic Analyzer (NTA) • HP IMC User Behavior Auditor (UBA) • HP IMC Branch Intelligent Management System (BIMS)
SDN	• HP VAN SDN Controller • HP Network Optimizer SDN Application for Microsoft® Lync® • HP Network Protector SDN Application

Note:
The HP campus network portfolio includes additional products: IEEE 802.11n APs and Walljacks, WLAN controllers, switches, and IMC modules. For more information, visit hp.com/networking.

Additionally, for the design of the converged campus, core switches need to be considered. HP offers advanced/high-performance core switches for large campus networks like the HP 10500 Switch Series along with core solutions for smaller campuses like the HP 5400R zl2 Switch Series and the HP 7500 Switch Series.

How it works

Unified Wired-WLAN controllers

One of the most important pieces of the converged campus solution is the Unified Wired-WLAN product family. This product portfolio is composed of three controller models for the campus and one model for the branch. Two of the campus controllers come in an appliance form factor while the third is a module compatible with the HP 10500 and 7500 Switches. For detailed descriptions of each one of these products visit:

- HP 870 Unified Wired-WLAN Appliance
- HP 850 Unified Wired-WLAN Appliance
- HP 10500/7500 20G Unified Wired-WLAN Module

All three Unified Wired-WLAN controllers are designed to be deployed either at the core of the campus or in a datacenter (local or remote). They communicate with the APs using a Lightweight Access Point Protocol (LWAPP) tunnel and to other controllers using an Inter-AC Tunneling Protocol (IACTP) tunnel. LWAPP tunnels can be secured using IPsec with SHA1 and AES.

The three main functions on a WLAN are AP management, client authentication, and traffic forwarding (between the WLAN and the wired LAN). AP management is performed by the controller (centralized), while client authentication and traffic forwarding can be performed either centrally, at the controller, or locally, by the APs.

In terms of traffic forwarding, there are three modes:

- **Centralized forwarding:** The controller performs traffic forwarding. Centralized forwarding comprises IEEE 802.3 centralized forwarding and 802.11 centralized forwarding. With IEEE 802.3 centralized forwarding, APs change incoming 802.11 frames to 802.3 frames and tunnel the 802.3 frames to the controller. With IEEE 802.11 centralized forwarding, APs directly tunnel incoming 802.11 frames to the controller. This type of traffic forwarding is typically used for guest traffic, or traffic requiring higher security.
- **Local forwarding:** APs directly forward data frames, while the controller may still perform authentication on clients. This forwarding mode reduces the workload of the controller, and retains the security and management advantages of the controller-based architecture. This type of traffic forwarding is typically used for traffic sensitive to latency, such as voice over WLAN (VoWLAN) traffic.

Figure 2. Centralized authentication and forwarding

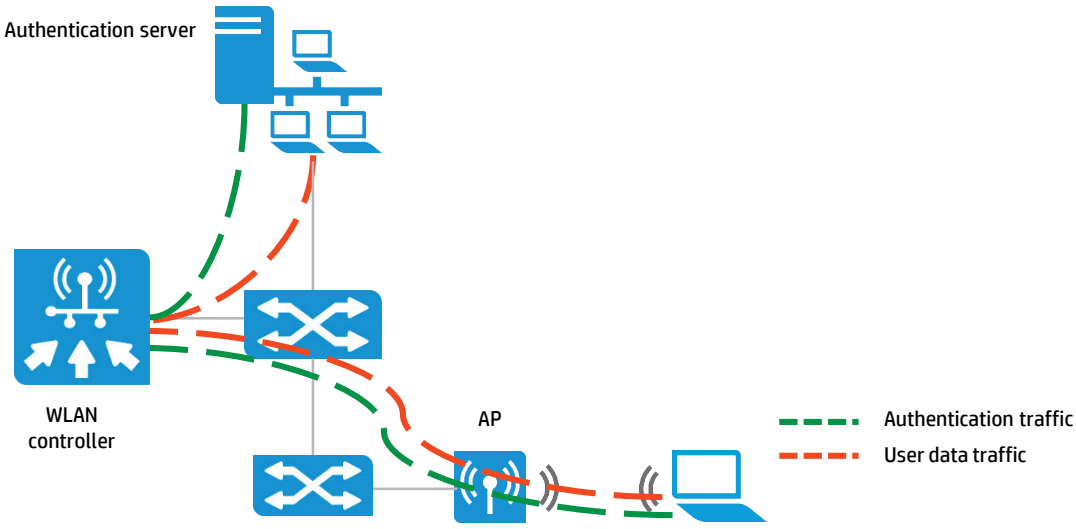

Figure 3. Centralized authentication and local forwarding

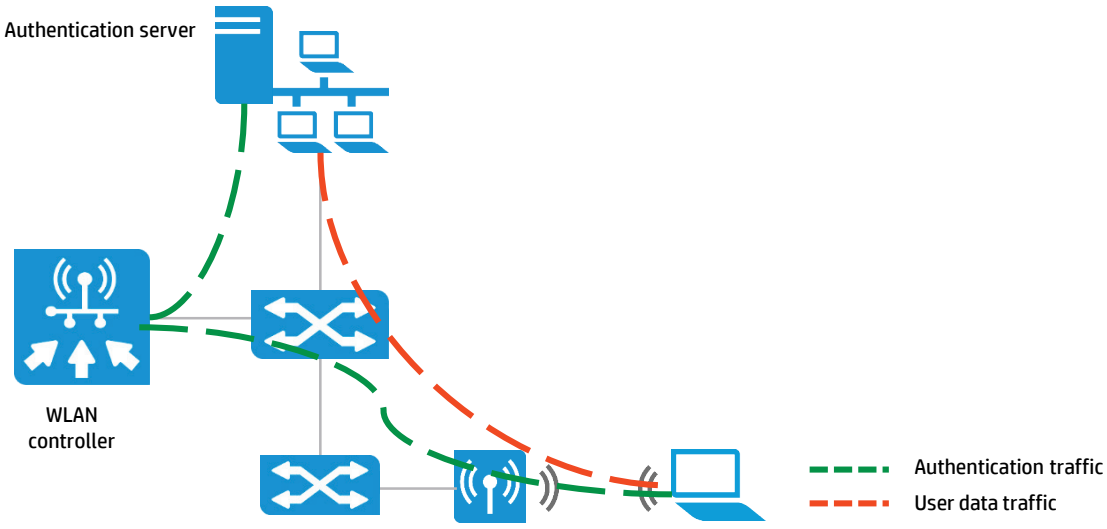

Figure 4. Local authentication and forwarding

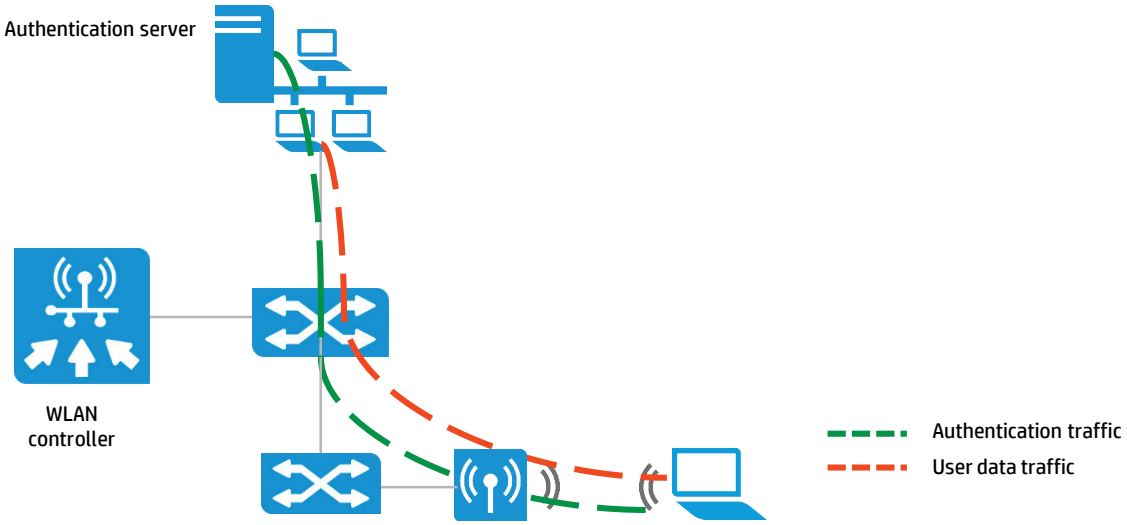

- **Policy-based forwarding:** This type of forwarding is based on a packet's destination. The AP selects local forwarding and delivers the packet directly to the adjacent AP or centralized forwarding and tunnels the packet to the controller.

Note:
Policy-based forwarding requires centralized authentication.

Table 2 summarizes the different WLAN design options available using an HP wireless solution.

Table 2. Unified controller deployment options

Function	Option 1	Option 2	Option 3	Option 4
Client authentication	Centralized	Centralized	Local	Centralized
Forwarding	Centralized	Local	Local	Policy-based

The Unified Wired-WLAN Controllers manage all of the functions of the WLAN as a whole and each AP in particular. They provide radio resource management, spectrum analysis, band navigation, and WLAN configuration including QoS and security.

IEEE 802.11ac APs

Even the first generation of IEEE 802.11ac APs, called wave 1, offers a significant improvement over IEEE 802.11n.

The IEEE 802.11n (11n) amendment. High Throughput (HT) PHY specification in the IEE802.11-2012 standard) was introduced in 2009 to overcome some of the limitations of IEE 802.11a, b, and g.

Note:
The 11n amendment is now included in the IEEE 802.11-2012 standard as clause 20.

The 11n standard drastically changed the behavior of WLAN networks. The following improvements were introduced:

- **OFDM improvement:** This coding scheme, already used in 802.11a and 802.11g, was improved to bring the bandwidth of a WLAN stream from 54 Mbps to 65 Mbps. Additional efficiency techniques, usually summarized under the name Short Guard interval (SGI), brought the bandwidth up to 72.2 Mbps.
- **MIMO technology:** The multiple-input/multiple-output technology uses multiple antennas to take advantage of a signal propagation effect that was an issue in previous standards: multipath.
- **Spatial multiplexing:** This is the most popular MIMO, which splits a data stream into two or more substreams that are transmitted and received in parallel by different antennas. The result is that the bandwidth is multiplied by the number of parallel streams. The IEEE 802.11n standard supports up to four streams, however, the maximum number offered in the market is three, with many APs and clients supporting just two stream. Most vendors describe their AP capabilities with TxR:S where T is the number of antennas with transmit capability, R is the number of antennas with receive capability, and S the number of spatial streams supported; for example: 2x2:2 and 3x3:3.
- **40 MHz channels:** IEEE 802.11g and 11a radios operate in a 20 MHz channel. IEEE 802.11n technology can combine two 20 MHz contiguous channels to achieve 40 MHz, doubling the channel's capacity. This is a beneficial technology in the 5 GHz band because of the large number of channels available, but having only three nonoverlapping channels in the 2.4 GHz band makes the use of 40 MHz channels more of an issue than a solution.

In 2013 the IEEE introduced a new amendment to its WLAN standard, IEEE 802.11ac. This new technology takes the improvements of 802.11n several steps further. The first generation, Wave 1, products support:

- **5 GHz band only:** To overcome the channel limitations of the 2.4 GHz band, it only supports the 5 GHz band
- **80 MHz channels:** Up to four 20 MHz channels can be combined to obtain 80 MHz and multiply the bandwidth by 4
- **Bandwidth:** Up to 1.3 Gbps can be reached using 3 spatial streams
- **256 QAM modulation:** Both 11n and 11ac standards

Table 3. 802.11n and 802.11ac bandwidth comparison

		Number of spatial streams		
		1	**2**	**3**
20 MHz	**802.11n**	72 Mbps	144 Mbps	216 Mbps
	802.11ac	87 Mbps	173 Mbps	289 Mbps
40 MHz	**802.11n**	150 Mbps	300 Mbps	450 Mbps
	802.11ac (Wave 1)	200 Mbps	400 Mbps	600 Mbps
80 MHz	**802.11ac (Wave 1)**	433 Mbps	867 Mbps	1.3 Gbps

IEEE 802.11ac Wave 2 products will provide new features, such as:

- **Spatial streams:** Support for up to 4 streams

> **Note:**
> The standards supports up to 8 streams, however, actual products are expected to support up to 4.

- **MU-MIMO:** Multiuser MIMO allows the AP to transmit independent data streams to several wireless clients simultaneously. In other words, spatial multiplexing is used to transmit data to different receivers at the same time
- **Beamforming:** Improves the distance between transmitter and receiver by using multiple antennas. These antennas transmit the same signal slightly out of phase so that they reach the destination in phase. For that to happen, the receiver must send feedback to the transmitter. Beamforming was defined on IEEE 802.11n, but clients did not support it. It is expected that some IEEE 802.11ac wave clients will support it

Access switching component

Access switches are a key component of the Converged Campus architecture. They play two roles, the first is to provide network connectivity to Ethernet clients and the second is to connect APs to the network.

Access switches ports, in general can be classified in two categories

- **Downlinks or access ports:** These are typically 10/100/1000BASE-T with RJ45 connectors and support the PoE/PoE+ PSE (power source equipment) function. If PoE/PoE+ is supported on all ports, how much power they can supply depends on the switch model. The ports also need to support access control features like IEEE 802.11X, MAC authentication, and/or portal authentication.
- **Uplinks or backbone ports:** They connect the access switches to either the aggregation layer switches, in a 3-tier LAN, or directly to the core layer, in a 2-tier LAN. Access switches usually have a minimum of two 10GBASE-T or SFP+ uplink ports.

An important aspect of stackable switches, as their name implies, is the stacking feature. Stacking allows the interconnection of two or more switches to form a single logical unit: a stack. HP offers two stacking technologies: physical stacking and Intelligent Resilient Framework (IRF).

- **Physical stacking:** The HP 3800 Switch Series and HP 2920 Switch Series support physical stacking. This technology is similar to IRF in its externally observed behavior. A physical stack is a single entity in terms of switching, routing, and management. These switches use specialized modules and cables to build the stack. Because specialized hardware is used, the stack interconnection offers high performance and redundancy.

Figure 6. HP 3800 Switch Series physical stacking

Table 4. Physical stacking topologies

	HP 2920 Switch Series	HP 3800 Switch Series
Topology: daisy chain	Up to 4 units	Up to 10 units
Topology: ring	Up to 4 units	Up to 10 units
Topology: full mesh	N/A	Up to 5 units

- **IRF:** This technology allows the interconnection of two or more switches to form a single logical switching and routing entity called an IRF-fabric. For an external switch, IRF-fabrics behave as single switches in every aspect: single Ethernet switches, single routing peers, and single managed devices (for example: single SNMP object instances).

Figure 7. HP stackable switches IRF

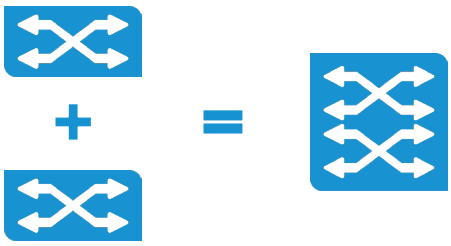

IRF is a feature of the HP Comware operating system. It uses standard 10GbE/40GbE ports to interconnect the fabric. The IRF-fabric's topology can be either a daisy-chain or a ring with each IRF-link formed by one or more physical links for high performance and availability. It is supported on all Comware-based campus switches: HP 10500 Switch Series, HP 7500 Switch Series, HP 5500 HI Switch Series, and HP 5130 EI Switch Series.

Best practices

Selecting the right site survey method

In the beginning, providing wireless connectivity was typically about coverage. Today, many networks are a combination of wired and wireless connectivity or wireless only connectivity; and users are relying on wireless for file sharing, video streaming, wireless VoIP, and more. Network administrators shifted their focus to provide users with ample bandwidth to support their usage patterns. As a result, the need for wireless site surveys has increased.
A site survey will help to determine:

- Ideal number of AP's and their optimal placement
- Critical coverage areas
- Cabling requirements, AP mounting
- Construction requirements

There are different types of site surveys:

- **Predictive site survey**

  A simulation tool is used to create a model of the RF environment. Virtual access points are placed on the imported floor plan to estimate expected coverage. The quantity and location of these virtual access points can then be adjusted. A predictive site survey does not require the onsite presence of an engineer.

- **Passive (onsite) site survey**

  An application is used to passively listen to WLAN traffic. This application detects active access points, noise level, and measures signal strength. The wireless adapter used for the survey is associated to any of the WLANs. A passive site survey requires the onsite presence of an engineer.

- **Active (onsite) site survey**

  An active site survey may use existing access points or standalone access points when existing access points are not available. The wireless adapter used for the survey is associated with one or several access points to measure round-trip time, throughput rates, packet loss, and retransmissions. An active site survey requires the onsite presence of an engineer. This type of survey is offered as a service through HP Services and many certified HP Partners.

The size and complexity of WLANs in today and tomorrow's campuses requires at least a predictive site survey. In cases where the WLAN is critical, as in most healthcare deployments and hospitality services, or in enterprises where a large percentage of clients will be wireless, an active site survey is recommended.

Choosing the right controller

All three Unified Wired-WLAN controller models have the same software features. The following table shows the differences between them and can be used to select the right model.

Table 5. Unified Wired-WLAN controller comparison

	870 UWW appliance	850 UWW appliance	10500/7500 20G UWW Module
Positioning	Large campus	Medium to large campus	Medium to large campus
Form-factor	Appliance with integrated switch	Appliance with integrated switch ports	Module for HP 10500 or 7500 Switch Series
Wireless throughput	40 Gbps	10 Gbps	20 Gbps
Maximum number of Aps	1536 (base license: 256)	512 (base license: 64)	1024 (base license: 128)
Maximum number of clients	30,000	10,000	20,000
Maximum number of 10GbE ports	4 SFP+	2 SFP+	N/A
Maximum number of GbE ports	12 RJ-45 100/1000 Mbps + 12 SFP 100/1000 Mbps	8 (10/100/1000BASE-T/SFP dual personality)	N/A

Selecting a controller redundancy option

The basic goal of redundancy is that APs, and hence wireless clients, are able to continue operating as normal, even when the AP loses contact with its primary controller. Controllers typically provide the following kinds of services:

- Configuration and firmware upgrades
- Authentication services
- Roaming services
- Captive portal services

The HP optimized architecture allows a WLAN solution to be designed such that it is optimized to handle many different kinds of failures, e.g., WAN link failure and controller failure.

HP Unified Wired-WLAN Controllers support 1+1, N+1, and N+N redundancy.

Note:
1+1 redundancy and its associated backup mechanisms are not supported on the HP 830 PoE+ Unified Wired-WLAN Switch, which is positioned for branch offices.

- **1+1 fast backup:**

  In this redundancy model, APs are connected to both an active and standby controller through control tunnels. However, only the tunnel to the active controller is up initially. The active and standby controllers exchange heartbeats for failure detection. If the active controller fails, the standby controller can immediately detect the failure and instructs the APs to switch the connections to the backup links.

  Both modules must be licensed for the number of APs that will be supported on the active controller.

- **N+1 backup:**

  If multiple HP Unified Wired-WLAN Controllers are deployed, N+1 redundancy is the best compromise between cost and reliability. In N+1 redundancy, if one of the N active controllers fail, the APs will switch their connections from the failed active controller to the standby controller. Load balancing can be achieved among the active controllers. The backup controller needs to be licensed for the largest number of APs supported by an active controller.

  The HP Unified Wired-WLAN Controllers discussed within this reference architecture can all support up to four additional controllers of the same model. If supporting different models, the number of active controllers that a backup can support is relevant to the number of APs on the network.

N+1 redundancy does not support sub second switch over, the switch over time is approximately 35s.

N+1 redundancy support switch back to original controller functionality, for example, an AP switched to backup controller, and it's original controller come back online, than the AP can automatically switch back to its original controller.

- **N+N backup and load balancing:**

  When multiple HP Unified Wired-WLAN Controllers are deployed, N+N redundancy is the most flexible redundancy mode. When an AP attempts to associate with a controller, the AP selects the optimal controller. When the controller fails, the AP selects another controller from the rest of the controllers. The selection of an optimal controller can be based on the loads on the controllers or predefined priorities. Meanwhile, load balancing can be achieved among the controllers.

  In N+N redundancy, the number of deployed APs cannot exceed the total number of APs that can be supported by all the controllers except the controller with the largest capacity. For example, if there are three controllers and the controllers can support 640, 128, and 32 APs respectively, the number of deployed APs cannot exceed 160 (128+32).

  N+N redundancy does not support sub second failover switch, the switch overtime is approximately 35s. N+N configuration does not support switch back to original controller functionality.

Note:
N+1 and N+N redundancy can be deployed over different controller models. For example, HP 850 Unified Wired-WLAN Appliances can be paired with HP 10500/7500 20G Unified Wired-WLAN modules in N+1 and N+N configurations.

Additionally, HP Unified Wired-WLAN Controllers support DHCP server, portal server, and 802.1X hot backup mechanisms that work in conjunction with the 1+1 redundancy configuration.

- **Portal server hot backup**

  Portal authentication requires high reliability, and thus is used with 1+1 redundancy. The active and standby controllers can synchronize the authentication and billing information of the clients in real time. When the active controller fails, the clients need not be re-authenticated on the standby controller and thus stay connected.

- **DHCP server hot backup**

  The HP Unified Wired-WLAN platform supports DHCP server hot backup if the built-in DHCP server is used to allocate addresses. If one active controller fails, the APs associated with the failed controller can renew their IP addresses on the standby controller without changing their IP addresses.

- **802.1X hot backup**

  802.1X hot backup is also used in conjunction with 1+1 fast backup. This feature enables controllers to synchronize 802.1X state information and the wireless clients' IEEE 802.11 information from primary to backup. When the primary controller fails, the backup controller will take over support for the APs, and the clients' 802.1X connection will not be lost.

HP recommends redundancy based on the customer's needs and requirements.

Choosing the right AP

HP offers three IEE 802.11ac APs. Table 6 provides a comparison between them and can be used to choose the right AP for each set of requirements.

Table 6. IEEE 802.11ac APs comparison

AP model	HP 560 802.11ac Dual Radio Access Point	HP 525 802.11ac Dual Radio Access Point	HP 527 802.11ac Unified Wired-WLAN Walljack
Radios	2.4 GHz 802.11b/g/n 3x3:3 MIMO 5 GHz 802.11a/n/ac 3x3:3 MIMO	2.4 GHz 802.11b/g/n 2x2:2 MIMO 5 GHz 802.11a/n/ac 2x2:2 MIMO	2.4 GHz 802.11b/g/n 2x2:2 MIMO 5 GHz 802.11a/n/ac 2x2:2 MIMO
Antennas	Internal: 6 External: N/A	Internal: 4 External: 4	Internal: 4 External: N/A
Ethernet Ports	1 x 10/100/1000BASE-T PoE PD port	1 x 10/100/1000BASE-T PoE PD port 1 x 10/100/1000BASE-T port	1 x 10/100/1000BASE-T PoE+ in uplink port 2 x 10/100/1000BASE-T downlink port 1 x 10/100/1000BASE-T PoE out downlink port 1 RJ-45 pass through port pair
PoE in requirements	Max: 14 W	Max: 12.9 W	PoE+/802.3at is required to provide power through the downlink port or USB port

Choosing the right deployment model

One important step when designing the WLAN, is deciding if authentication and forwarding are going to be centralized or distributed. This decision can be made by the WLAN service (SSID). For example:

Table 7. Deployment example

WLAN service	Authentication	Forwarding
Guest service	portal/centralized	centralized
BYOD service	portal/centralized	Local
Corporate PC service	802.1x/local	local

The following analysis uses three types of service: guest, BYOD, and corporate PC. Some of the factors to consider for this analysis are:

- Guest service

  In general guest traffic does not need to reach any destination within the LAN. Usually it will go directly to the Internet or a secure network segment (usually called a DMZ) in which guest servers are located. Centralized authentication and forwarding is recommended.

  For additional security, one of the controller ports can be directly connected to the firewall, or if that is not possible, a guest access tunnel can be created to a second controller located in the DMZ.

- Employees with corporate PC

  If the controller is located in a remote data center tunneling employee traffic back to it might not be optimal because some of that traffic has a local destination. Additionally, if the campus has direct Internet access, employees will have better performance with local forwarding.

- BYOD service (for employees)

  The BYOD service is based on portal/Web authentication. In general, the centralized portal authentication implementation is simpler, but distributed portal authentication is also supported.

  If all BYOD traffic is directed outside of the corporate network, the situation is similar to the guest service in terms of forwarding mode. If, however, personal mobile devices will have access to internal resources, like email, conferencing, or special corporate apps, and if some of these resources are located in the campus, local forwarding can be the best choice.

Depending on the vertical such as hospitality, sports, healthcare providers, enterprise, etc., other services may be required, and/or the relative importance of each of these services can vary. For example, in healthcare, dedicated VoIP devices may

be deployed on a separate SSID because of their special need of QoS and optimized multicast. Because latency is critical, local forwarding is the recommended solution.

Choosing the right access switch

Depending on the network requirements, different switches can be chosen. The first decision to make is which switches better satisfy the requirements, modular or stackable.

Table 8. HP campus access switches

Type	Product series	
Modular access switches	• HP 5400R zl2 Switch Series	• HP 7500 Switch Series
High-end stackable access switches	• HP 3800 Switch Series	• HP 5500 HI Switch Series
Low-end stackable access switches	• HP 2920 Switch Series	• HP 5130 EI Switch Series

For a summary of these switches' features see appendix A.

The following are some of the factors to take into account when choosing the right access switches.

- **Form factor:** Modular or stackable switches
- **Stacking technology:** Physical or IRF
- **Uplinks:** Access switches should be connected to both members of a switch pair in the next layer—aggregation or core. However, in many cases, having two uplinks per access switch may not be necessary. A stack of access switches can share the uplinks and make better use of them.

Figure 8. Access layer stacking

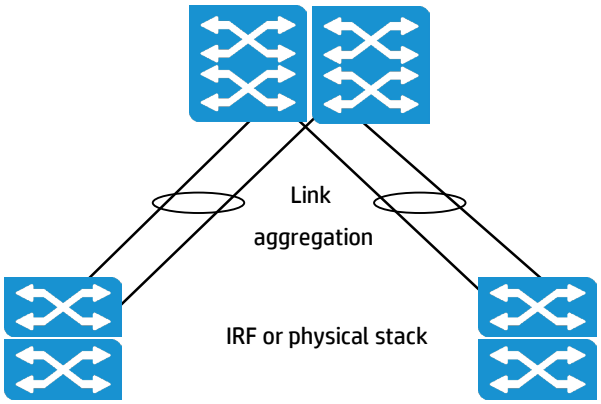

For example, a stack of six 48GbE port switches—a total of 288 ports—that requires an oversubscription rate of 10:1, will need only three 10GbE uplinks to satisfy the bandwidth needs and a fourth uplink would provide 3-1 redundancy. In this case the number of 10GbE uplinks has been reduced to a third with no effective loss of bandwidth and increased resiliency.

Reducing the number of uplinks is beneficial because it also allows, in many cases, for the consolidation of the aggregation and the core layer to form a two-tier network.

- **Comware and ProVision OS Integration:** Some large campus networks may require an IRF-optimized high-speed core and an SDN-ready, high-quality, low-cost access layer. In that case, the recommendation is a combination of Comware OS 10500 switches at the core and ProVision OS switches at the access layer (5400R zl2, 3800, or 2920).

  ProVision-based switches offer a combination of SDN support, quality, reliability, simplicity, and price that makes them an excellent choice for the access layer. The hardware-based stacking capability of the 2920 and 3800 is an example of this.

  Comware and ProVision switches integrate well because both are based on standards. For the special case of an IRF-based core, IRF-fabric members exchange information in order to detect when the fabric has been split due to a broken IRF link. The simplest way of transporting this information exchange is to use an extension of the link aggregation control protocol (LACP) called LACP-MAD, originally available only on Comware-based switches.
  The ProVision OS supports LACP-MAD today, providing an OS integration element.

Unifying access control and BYOD

An important aspect of the converged campus solution is the fact that wired and wireless clients get the same user experience, for example during the network login process.

In the past, network access was only granted to employees using PCs provided by the company or organization.

Next, the need to provide network access to guests as a hospitality service, first at hotels, schools, etc., and later at enterprises and organizations of all kinds was defined. This solution, in general, connects guests to a particular VLAN that provides access only to the Internet. Guests, in this solution are defined as "non-authenticated users".

With the proliferation of netbooks, tablets, and smartphones, employees started demanding access to the network for these devices including access to privileged company resources such as email, file servers, and databases. BYOD is the answer to this request.

In some vertical markets, BYOD is a critical part of the network access strategy. For example, in education, many students and teachers are expected to user their own devices to access learning resources, complete tasks, etc.

Some companies are starting to take advantage of their employees' mobile devices by creating or purchasing specific apps for them: collaboration tools, internal blogs and forums, podcasts, videocasts, bulletins, and marketing and HR resources.

Components

A complete access control solution requires at least three components: clients or endpoints, enforcement points, and directory and policy services that can be composed of one or more servers.

- Clients or endpoints: depending on the campus policies, these can be provided by the IT organization as in the case of corporate PCs and mobile phones, or personal devices.
- Enforcement points: HP Converged Campus Unified Wired-WLAN controllers, APs, and switches support several network login functions like Web/portal, IEEE 802.1X, and MAC-based authentication. Some of them also support Lightweight Directory Access Protocol (LDAP).
- HP IMC provides two specific IMC modules for access control, directory, and policy servers:
  - UAM provides access control functions like RADIUS and portal server, and a flexible BYOD solution. UAM can use its own user database or can access different directory services via LDAP.
  - EAD offers security policy network management and endpoint posture assessment. Through EAD, the HP BYOD solution can be extended to incorporate Mobile Device Management (MDM) solutions from vendors like Citrix® and MobileIron.
- Two additional IMC modules can be used to monitor user and traffic behavior:
  - NTA is a traffic analysis tool that receives detailed traffic statistics from the network devices through NetStream and sFlow® and stores and presents this information to allow for traffic trends and particular behaviors.
  - UBA combines information from NTA and UAM and provides a high-performance, scalable network log audit and analysis solution.

How it works

Access control and BYOD basics

The access control process can be divided into three major stages: authentication, authorization, and accounting (AAA). The main difference with traditional authentication and the BYOD approach is that in the former, only the user had to be identified while in the latter, both the user and the mobile device have to be identified and authorized. In many cases, the device is also checked for its health: antivirus state, software update state, etc.

In BYOD terms, identifying the device is called fingerprinting while the health check is called posture check.

Device fingerprinting can be performed in several steps:

1. From the MAC address OUI the device's vendor can be determined
2. Next, the DHCP process provides additional information, like device model, for example:

```
[os 907]
 description=HP 3500yl
 fingerprints=<<EOT
 1,3,4,23,67,66,43
 EOT
```

3. Finally, during the portal self-registration/authentication phase, the version of the device's browser is obtained

```
Your User Agent: Mozilla/5.0 (Windows NT® 6.1; WOW64; rv:10.0.2) Gecko/20100101
Firefox/10.0.2
```

Authentication

Authentication can be divided into three major categories: secure access, BYOD, and guest access. The main differences between these cases are how the following questions are answered:

1. User identification: Is there a user account in the corporate/organizational directory?
2. Device identification: Does the device belong to the user or is it provided by the corporation/organization?
3. Device health check: Is a device health check required for this device?

Table 9. Authentication categories

Category	Question	Answer
Secure Access	User identification	The user has an account in the corporate directory that has been created for him/her by an authorized directory administrator
	Device identification	The device has been provided by the corporation and may contain special hardware, software, and/or settings to guarantee compliance with the organization's policies
	Device health check	Corporate devices are normally checked before getting access to the network to check compliance to the company policy
BYOD	User identification	The user has an account in the corporate directory created by an authorized directory administrator. It is usually the same account used for secure access
	Device identification	The device is owned by the user. It is usually a tablet or a smartphone and in some cases a laptop or a netbook PC
	Device health check	Personal devices might be checked before getting access to the network to ensure that they are in compliance to company policy
Guest	User identification	The user does not work for the corporation and does not have an account in the directory. In some cases, a self-registration process may be established
	Device identification	The device is owned by the user. It is usually a tablet, a smartphone, or a laptop
	Device health check	Guests' personal devices do not need to be checked before getting access to the network

The three categories described above are usually applicable to enterprise organizations. With other types of organizations categories can be different. For example:

- Hotels: In a hotel, the concept of guest can be slightly different because the user may have to self-register. The registration process is usually a Web-based application tied to the hotel registration system and creates a temporary account that registers the MAC address of the device. After the registration and during the account lifetime, the device is automatically authenticated by its MAC address.
- Universities/education institutions: Here users can be categorized as employees, teachers/faculty members, students, and guests. The last two categories may not have devices provided by the organization and need to use their own devices.

Authorization

Once the authentication process has been successfully completed and access is granted, an access policy can be assigned to the user. This policy is a set of resources the user can access with conditions.

A generic policy can be designed for each major authentication category:

- Secure access clients may be granted access to all applications relevant to the user's role and workgroup
- BYOD clients may be granted access to a subset of applications such as email, chat, other collaboration tools, corporate blogs and more, but may not be granted access to corporate databases, private and confidential information, etc.
- Guests may be granted Internet access only or a very limited set of applications

One important feature of the HP BYOD solution is that it provides a high level of flexibility in terms of policies. A different policy can be enforced based on the combination of access conditions. A certain organization may use certain policies when clients connect to a certain set of devices, switches, and APs, and a different set when they connect through any other access device.

Table 10. Access condition examples

Access condition	Description
Access area	A set of network devices (switch, AP, wireless controller)
Access IP group	A set of IPv4 or IPv6 addresses and/or subnets
SSID group	A set of SSIDs
Endpoint MAC group	A set of MAC addresses, for example a certain set of OUIs
Vendor group	A set of one or more vendors: Apple, Microsoft, HTC, Samsung, Nokia, and HP
OS group	A set of one or more operating systems: iOS, Android, Windows® 8
Endpoint Type Group	A type of endpoint device: smartphone, tablet, iPad®, PC

Best practices

The design of an access control solution for a customer depends heavily on that particular customer's needs in terms of security, resource access, etc.

Simplifying configuration

HP IMC gives customers the advantage of managing over 6000 devices, from more than 200 manufacturers, all from one single management system.

IMC manages both the physical and virtual networks, and integrates and provides both fixed and mobile access control. In addition to being a Fault, Configuration, Accounting, Performance and Security platform, IMC supports additional modules for particular network requirements, such as IPsec VPN Manager, BIMS, user behavior analysis, and EAD, to name a few.

This section describes how the HP IMC BIMS module can help customers automate their installation process for both the campus and the branch.

Components

The BIMS solution requires the deployment of the IMC Standard or Enterprise platform running on either a Windows Server® platform or Linux® platform, and the IMC BIMS module.

- HP Intelligent Management Center Standard Software Platform
- HP Intelligent Management Center Enterprise Software Platform
- HP Branch Intelligent Management System Module

The following HP campus switches support BIMS:

- HP 5120 Switch Series
- HP 5130 Switch Series
- HP 5500 EI Switch Series
- HP 5500 HI Switch Series

How it works

The BIMS module centrally manages a large number of customer premises equipment (CPEs). It solves the problems of dynamic IP management, saves network maintenance cost, and improves the network management efficiency. It provides resource, service, configuration, alarm, and system management.

Technical report-069 is the specification to use in the CPE WAN Management Protocol (CWMP). It defines an application layer protocol for remote management of end-user devices.

BIMS can be used to manage network devices in two ways:

- **Zero-touch deployment:** A TFTP server sends the boot configuration that is necessary to communicate with the BIMS server
- **Minimal touch deployment:** A USB Flash drive is used to install the autodeploy.cfg file on the device (only available on MSR and VSR routers)

Once the initial configuration is installed on the device, it communicates with the BIMS server over HTTP or HTTPS. BIMS validates the device and sends the complete configuration file. The device comes online fully configured.

Note:
On the BIMS server, various configuration files can be created with variables to add unique information to the devices, while still adding commands that represent common information.

Best practices

In a large enterprise environment, the IMC BIMS solution can sit in the campus network. It needs to have both a private and public interface connection. Switches and routers internal to the campus will connect via the private network, while branch devices will communicate over the corporate WAN or an Internet connection. As long as the external router or switch can reach the public interface of the BIMS server, communication can be achieved.

To simplify the way BIMS learns about the devices to be deployed, a bulk configuration option is available. The BIMS server can import a .csv file that can have thousands of devices listed, to make mass deployments possible.

Figure 9. BIMS bulk configuration option

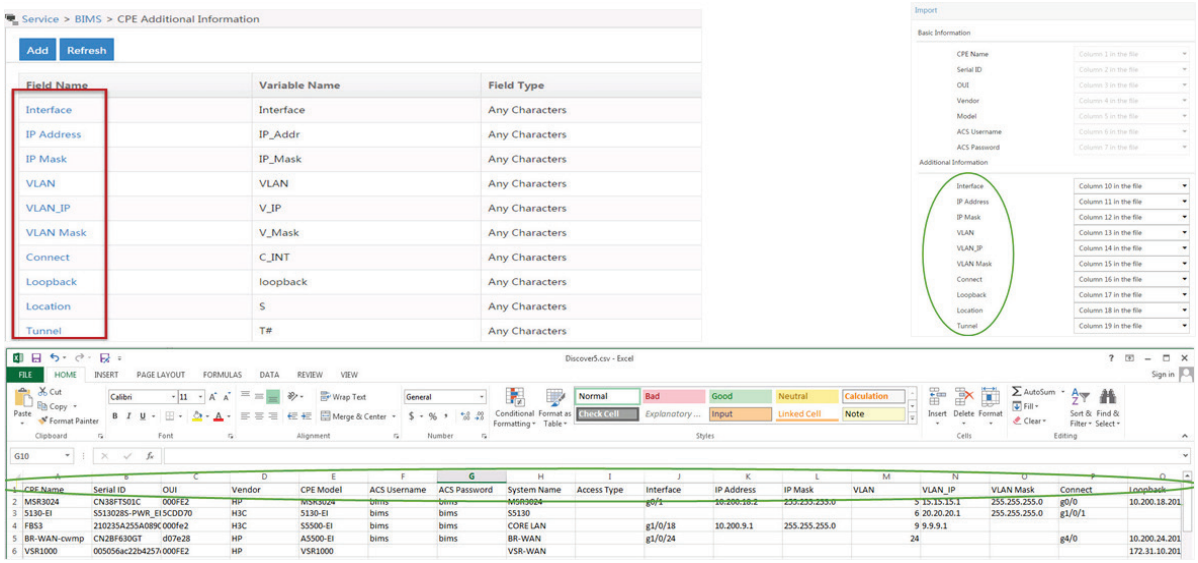

Controlling the network with SDN

HP Networking is helping to lead the way in simplifying and transforming the network to meet your organization's needs for mobility, virtualization, high-definition video, rich-media collaboration tools, and cloud computing.

By embracing a SDN, it is possible to reap the full value of your network investment. SDN, delivered through our proven solutions, will help users and organizations experience applications as never before. It will free your IT administrators from the drudgery of manual network configuration and reconfiguration because the network will be automatically tuned to application and business needs.

The IT staff can focus more on the quality of the business experience, and spend less time managing the details of the underlying networking infrastructure.

HP VAN SDN Controller

HP VAN SDN Controller Software provides a unified control point in an OpenFlow-enabled network, simplifying management, provisioning, and orchestration. This enables delivery of a new generation of application-based network services. It also provides open application program interfaces (APIs) to allow third-party developers to deliver innovative solutions to dynamically link business requirements to network infrastructure via either custom Java programs or general-purpose RESTful control interfaces.

Figure 10. HP SDN architecture

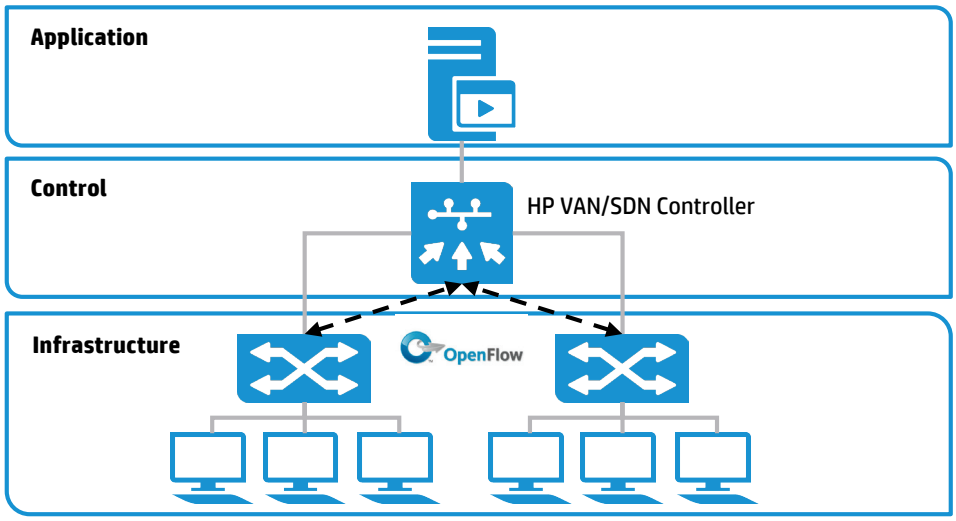

The VAN SDN Controller is designed to operate in campus, data center, and service provider environments. It offers:

- Enterprise-class platform for the delivery of a broad range of network innovations
- Compliant with OpenFlow 1.0 and 1.3 protocols
- Support for over 50 OpenFlow-enabled HP switch models
- Open APIs to enable third-party SDN application development
- Extensible, scalable, and resilient controller architecture

HP VAN SDN solutions for the campus

Two of the SDN solutions for the campus LAN from HP include the HP Network Protector SDN Application and HP Network Optimizer SDN Application for Microsoft Lync®.

HP Network Protector SDN Application utilizes SDN and the OpenFlow protocol to push security to the edge of the network where clients connect. Instead of malicious traffic traversing a network to the core and getting blocked by an IPS, it can now be blocked at the edge without having to add a dedicated security appliance. HP Network Protector can be used to turn a traditional access layer switch into a security appliance.

HP Network Optimizer SDN Application for Microsoft Lync utilizes SDN and the OpenFlow protocol to dynamically provision QoS for Microsoft Lync voice, video, and application sharing calls which is not possible with traditional QoS.

SDN hybrid deployment

With HP, it is possible to get the benefits of SDN without re-architecting an existing network.

This is accomplished with a hybrid deployment where SDN is used to provide enhancements while allowing traditional technologies and protocols to forward traffic throughout the network. The benefit of this is that it is only necessary to deploy OpenFlow enabled switches at the edge or access of a network.

HP Network Protector SDN Application

HP Network Protector SDN Application was developed out of the continued need to enhance network security.

The HP Network Protector SDN Application stops threats at the network access layer before they can cause damage. HP Network Protector can be used in any network environment where security is a concern, including BYOD, data center, and cloud computing environments. HP envisions a network where HP Network Protector can be implemented on any network device anywhere in the network for unprecedented network visibility, event correlation accuracy, and security control.

HP Network Protector features include:

- Runs on HP Virtual Application Networks (VAN) SDN Controller
- Consumes real-time reputation security intelligence from the HP TippingPoint DVLabs cloud service
- Protects from over 1,000,000 botnets, malware, and spyware malicious sites
- Provides native integration for improved visibility and accuracy with HP ArcSight solutions
- Uses OpenFlow-enabled switches to detect malware and botnets
- Has the ability to implement a custom whitelist and blacklist
- Has dynamic switch learning with HP OpenFlow-enabled switches, which distributes detection into the switch infrastructure
- Provides security enforcement decision feedback with HP IMC

The value of HP Network Protector SDN Application is in pushing security to the access layer of a network without requiring additional hardware. It is desirable to block malicious traffic as close to the source as possible; but it is not economically feasible to deploy an IPS between every access layer edge port and every network host. HP Network Protector solves these issues.

How it works

HP Network Protector utilizes OpenFlow on access layer switches. When a switch boots and connects to the HP VAN SDN Controller with Network Protector, a default flow is pushed to the device. This is in addition to the flows installed on the device by the base controller to support a hybrid environment.

Figure 11. HP Network Protector dashboard

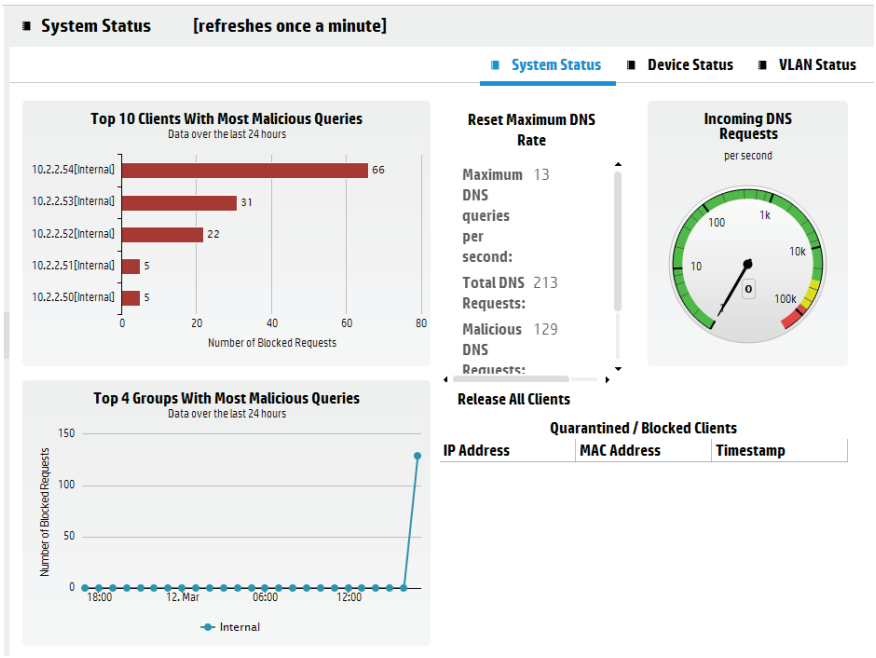

When a switch receives DNS traffic, it is forwarded to the controller. On the controller, the request is compared to the HP TippingPoint Reputation Security Monitor (RepSM). If there is a match, meaning the traffic is malicious in nature, HP Network Protector creates a DNS response and sends it to the switch to be forwarded to the client. There are two response options. HP Network Protector can respond with a host not found, providing an immediate failure to the requesting client so that it doesn't have to wait for a timeout or attempt further resolution. Or, the second option is to redirect a client to a server of the administrator's choice. In this case, the client could be provided feedback that the request was blocked due to company security policy. With either option, the malicious traffic will be blocked.

If the DNS request doesn't match the RepSM, meaning the traffic should be permitted, HP Network Protector instructs the switch to forward the traffic normally.

HP Network Optimizer SDN Application for Microsoft Lync

Deploying trusted and granular QoS can be extremely complex and require implementing tedious and time-intensive manual configurations on a device-by-device basis. In fact, it is nearly impossible to implement traffic policy using deep packet inspection (DPI) for soft clients with legacy networks because of SIP TLS encryption and dynamic application ports used by UC&C applications. This results in poor application traffic visibility.

The HP Network Optimizer SDN Application will automate policy deployment dynamically on a per-connection basis for voice, video, and application sharing to deliver a better user experience and reduce operational costs. When a desktop sharing, voice, or video connection is initiated using the Microsoft Lync client in the campus or branch office, the Lync Server in the data center provides the HP Network Protector SDN Application with call details such as source and destination IP address, protocol type, application ports, and bandwidth requirements at the start and end of every call. HP Network Optimizer then uses these per-connection application details to dynamically provision the end-to-end network path and QoS policy via the HP VAN SDN Controller using OpenFlow.

Once the QoS policies and path are programmed via OpenFlow, the call is connected to the destination client. The HP Network Optimizer SDN Application uses the intelligence from Lync Server and the Lync SDN API, along with the robust capabilities of the HP VAN SDN Controller, to implement consistent QoS across the network. All of this is done dynamically through a central point of control; eliminating the need for manual, device-by-device configuration via the CLI, and greatly simplifying policy deployment and reducing the likelihood of human errors.

Figure 12. HP Network Optimizer dashboard

How it works

The HP Network Optimizer SDN Application utilizes OpenFlow to dynamically prioritize traffic at the edge of a network. Traditionally, there are four ways that traffic can be identified so that it can be prioritized.

- First, it is possible to prioritize all traffic from a device. This method is used with traditional VoIP phones by placing the phone in a voice VLAN and prioritizing all traffic in that VLAN. Typically, an ACL would also be used to stop all traffic not destined for the VoIP server. This solution is not possible with Microsoft Lync because the client is a soft client running on PC.
- Second, if a solution uses a predefined TCP or UDP port number, traffic matching that port can be prioritized. This is not an option with Lync because it dynamically assigns ports from a wide range so that it can support multiple calls between many parties simultaneously.
- Third, it is possible to copy network traffic to an analyzer to determine its nature. Even when this solution is possible, it requires a significant amount of network bandwidth and processing power on the analyzer, which is a waste of expensive and precious resources. However, in the case of Lync this is not possible because all traffic is encrypted, making analysis impossible.
- Lastly, it is possible for the client to tag traffic as important and configure the network to trust the tags. While this will work and Lync does support it, this solution requires a level of trust from network clients that is not reasonable. As soon as the network trusts a client, there will be users who abuse the trust and artificially prioritize all of their traffic. In other words, a user could use a company's network to watch Netflix movies in full HD.

This left HP and Microsoft to determine a new method to prioritize traffic. It was realized that Lync Server had complete knowledge of all calls happening in an environment. Microsoft, in collaboration with HP, developed an API that installs on Lync Server and can make RESTful API calls to the HP Network Optimizer SDN Application with all of the call details, including users, type of call, and bandwidth requirements. HP Network Optimizer can then dynamically prioritize traffic on the network for the duration of the call.

There are two ways to prioritize the interesting traffic based on the capabilities of the network. When HP first demonstrated this solution, it was assumed that the entire network was OpenFlow enabled. This is a great solution because it doesn't require any QoS configuration on the network. According to some enterprises, it can take more than six months to deploy a QoS solution on a network.

However, it became apparent that the assumption of a 100 percent OpenFlow-enabled network was not reasonable. Therefore, it was decided to approach the first release of the product with the assumption of a hybrid network where only the edge, or access devices, were OpenFlow enabled. In this case, the Lync SDN solution does DSCP remarking at the edge of the network and the rest of the network is configured to trust the markings supplied by the access layer device. It was previously described that trusting QoS settings was a bad idea. But in this case, the access layer devices are doing the marking and not the end user clients. When HP Network Optimizer boots, a default flow is pushed to all access devices that remarks all traffic to normal priority in the specified VLANs. Then it is possible for HP Network Optimizer to dynamically prioritize the Lync traffic to an administratively assigned priority.

Out of the box, this solution will work without additional configuration for clients that are attached to OpenFlow-enabled devices. In the case where one client is not directly connected to an OpenFlow-enabled device, it is possible for a network administrator to configure a gateway for a known group of devices. This will enable prioritization to be dynamically assigned for the network under an administrator's control.

Creating a converged campus with HP

The HP Converged Campus solution provides you with a single, optimized, and scalable unified network for secure and consistent access to business critical applications. It is designed to meet emerging trends such as:

- Virtual meetings
- BYOD
- Cloud computing and storage
- IEEE 802.11ac WLAN

HP offers a comprehensive portfolio of campus access technologies, allowing businesses to deliver high-performance, reliable network services to the ever-growing number of mobile users, devices, and applications—and meet users' expectations for wireless connectivity, BYOD, unified communications, and security.

Appendix A. Basic facts for HP campus access switches

Table 11. HP 5400R zl2 Switch Series basic facts

Positioning	Small-to-medium campus core (main option) Small-to-medium campus aggregation Client access with ProVision OS, SDN support, and modular form-factor
Form-factor	Modular
Operating system	ProVision OS
Virtualization technology	Distributed Trunking (DT)
Maximum number of 10GbE ports	96
Maximum number of GbE ports	288
Maximum number of PoE+ ports	288 (at 30 W)
Unicast routing protocols	RIPv2, OSPFv2, BGP
Multicast routing protocols	PIM-DM and PIM SM
SDN	Ready—Applications available
Service modules	HP Advanced Services v2 SSD Module HP Advanced Services v2 HDD Module HP MSM775 zl Premium Controller Module

Table 12. HP 7500 Switch Series basic facts

Positioning	Small-to-medium campus core and aggregation with complex routing features requirements Client access with HP Comware OS and modular form factor Client access with high bandwidth uplinks (10GbE or 40GbE)
Form-factor	Modular
Operating system	HP Comware v5
Virtualization technologies	IRF—up to 4 units
Maximum number of 40GbE ports	Standalone switch: 40—IRF fabric: 152
Maximum number of 10GbE ports	Standalone switch: 80—IRF fabric: 312
Maximum number of GbE ports	Standalone switch: 480—IRF fabric: 4312
Maximum number of PoE+ ports	240 (at 30 W)
Unicast routing protocols	RIP, OSPF, BGP, ISIS (all: IPv4 and IPv6)
Multicast routing protocols	PIM DM, PIM SM, PIM SSM, BIDIR PIM, MSDP, MBGP (all: IPv4 and IPv6)
MPLS	MCE, MPLS, MPLS L3VPN, MPLS L2VPN, VPLS, Multicast VPN
Service modules	HP 10500/7500 20G Unified Wired-WLAN Module HP 10500/11900/7500 20 Gbps VPN Firewall Module

Table 13. HP 3800 Switch Series basic facts

Positioning	Client access—advanced feature requirement
Form-factor	Stackable
Operating system	ProVision OS
Virtualization technology	Distributed trunking (DT) Physical stacking—up to 10 units (meshes stack up to 5 units)
Maximum number of 10GbE ports	Standalone switch: 4—stack: 40 (10GBASE-T or SFP+)
Maximum number of GbE ports	Standalone switch: 48—stack: 480
Maximum number of PoE+ ports	36 (at 30 W)
Unicast routing protocols	RIPv2, OSPFv2, BGP
Multicast routing protocols	PIM-DM and PIM SM
SDN	Ready

Table 14. HP 5500 HI Switch Series basic facts

Positioning	Client access—advanced routing/MPLS VPN requirements
Form-factor	Stackable
Operating system	HP Comware v5
Virtualization technologies	IRF—up to 9 units
Maximum number of 10GbE ports	Standalone switch: 6—IRF stack: 36 Note: 2 fixed 10GbE ports + 2 expansion slots x 2-port modules
Maximum number of GbE ports	Standalone switch: 48—IRF fabric: 432
Maximum number of PoE+ ports	48 (at 30 W)
Unicast routing protocols	RIP, OSPF, BGP, ISIS (all: IPv4 and IPv6)
Multicast routing protocols	PIM DM, PIM SM, PIM SSM, BIDIR PIM, MSDP, MBGP (all: IPv4 and IPv6)
MPLS	MCE, MPLS, MPLS L3VPN, MPLS L2VPN, VPLS, Multicast VPN

Table 15. HP 2920 Switch Series basic facts

Positioning	Client access—basic feature requirement
Form-factor	Stackable
Operating system	ProVision OS
Virtualization technology	Physical stacking—up to 4 units (meshes stack up to 5 units)
Maximum number of 10GbE ports	Standalone switch: 4—stack 16 Note: require additional modules for stacking
Maximum number of GbE ports	Standalone switch: 48—stack: 192
Maximum number of PoE+ ports	36 (at 30 W)
L3 protocol set	Static and RIPv2
SDN	Ready—Applications available

Table 16. HP 5130 EI Switch Series basic facts

Positioning	Client access—basic routing requirements
Form-factor	Stackable
Operating system	HP Comware v7
Virtualization technologies	IRF—up to 4 units
Maximum number of 10GbE ports	Standalone switch: 4—IRF stack: 4 Note: 2 expansion slots x 2-port modules
Maximum number of GbE ports	Standalone switch: 48—IRF fabric: 432
Maximum number of PoE+ ports	24 (at 30 W)
Unicast routing protocols	Static and RIPv2

Learn more at
hp.com/networking

Sign up for updates
hp.com/go/getupdated

Share with colleagues

Rate this document

4AA5-7049ENW, February 2015