

Ruckus Enterprise Campus Network Design Guide

Supporting FastIron 08.0.90

Copyright, Trademark and Proprietary Rights Information

© 2019 ARRIS Enterprises LLC. All rights reserved.

No part of this content may be reproduced in any form or by any means or used to make any derivative work (such as translation, transformation, or adaptation) without written permission from ARRIS International plc and/or its affiliates ("ARRIS"). ARRIS reserves the right to revise or change this content from time to time without obligation on the part of ARRIS to provide notification of such revision or change.

Export Restrictions

These products and associated technical data (in print or electronic form) may be subject to export control laws of the United States of America. It is your responsibility to determine the applicable regulations and to comply with them. The following notice is applicable for all products or technology subject to export control:

These items are controlled by the U.S. Government and authorized for export only to the country of ultimate destination for use by the ultimate consignee or end-user(s) herein identified. They may not be resold, transferred, or otherwise disposed of, to any other country or to any person other than the authorized ultimate consignee or end-user(s), either in their original form or after being incorporated into other items, without first obtaining approval from the U.S. government or as otherwise authorized by U.S. law and regulations.

Disclaimer

THIS CONTENT AND ASSOCIATED PRODUCTS OR SERVICES ("MATERIALS"), ARE PROVIDED "AS IS" AND WITHOUT WARRANTIES OF ANY KIND, WHETHER EXPRESS OR IMPLIED. TO THE FULLEST EXTENT PERMISSIBLE PURSUANT TO APPLICABLE LAW, ARRIS DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, TITLE, NON-INFRINGEMENT, FREEDOM FROM COMPUTER VIRUS, AND WARRANTIES ARISING FROM COURSE OF DEALING OR COURSE OF PERFORMANCE. ARRIS does not represent or warrant that the functions described or contained in the Materials will be uninterrupted or error-free, that defects will be corrected, or are free of viruses or other harmful components. ARRIS does not make any warranties or representations regarding the use of the Materials in terms of their completeness, correctness, accuracy, adequacy, usefulness, timeliness, reliability or otherwise. As a condition of your use of the Materials, you warrant to ARRIS that you will not make use thereof for any purpose that is unlawful or prohibited by their associated terms of use.

Limitation of Liability

IN NO EVENT SHALL ARRIS, ARRIS AFFILIATES, OR THEIR OFFICERS, DIRECTORS, EMPLOYEES, AGENTS, SUPPLIERS, LICENSORS AND THIRD PARTY PARTNERS, BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL, PUNITIVE, INCIDENTAL, EXEMPLARY OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER, EVEN IF ARRIS HAS BEEN PREVIOUSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES, WHETHER IN AN ACTION UNDER CONTRACT, TORT, OR ANY OTHER THEORY ARISING FROM YOUR ACCESS TO, OR USE OF, THE MATERIALS. Because some jurisdictions do not allow limitations on how long an implied warranty lasts, or the exclusion or limitation of liability for consequential or incidental damages, some of the above limitations may not apply to you.

Trademarks

ARRIS, the ARRIS logo, Ruckus, Ruckus Wireless, Ruckus Networks, Ruckus logo, the Big Dog design, BeamFlex, ChannelFly, Edgelron, FastIron, HyperEdge, ICX, IronPoint, OPENG, SmartCell, Unleashed, Xclaim, ZoneFlex are trademarks of ARRIS International plc and/or its affiliates. Wi-Fi Alliance, Wi-Fi, the Wi-Fi logo, the Wi-Fi CERTIFIED logo, Wi-Fi Protected Access (WPA), the Wi-Fi Protected Setup logo, and WMM are registered trademarks of Wi-Fi Alliance. Wi-Fi Protected Setup™, Wi-Fi Multimedia™, and WPA2™ are trademarks of Wi-Fi Alliance. All other trademarks are the property of their respective owners.

Contents

Introduction.....	4
Overview.....	4
How to Build a Robust Campus Network.....	4
Ruckus Campus Network Solutions.....	7
Ruckus Campus Network Solutions Overview.....	7
Ruckus Wired Campus Network Designs.....	7
Conclusion.....	12

Introduction

Overview

The campus network is the portion of the enterprise network infrastructure that provides access to network communication services and resources to end users and devices that are spread over a single geographic location. The end users in a campus network may be dispersed more widely than in a single local area network (LAN) but are usually not as scattered as they would be in a wide area network (WAN). The key items that must be kept in mind while designing a campus network are the management, modularity, and resiliency in a network. All the nodes in a campus network are interconnected by optical fiber media that takes advantage of the 1/10/40/100 Gigabit Ethernet (GbE) technology. In most cases, Wi-Fi hot spots make up the user end of the network, for example, in universities, schools, and offices.

A typical network is segregated into three layers: the core, aggregation, and access layers (also known as a hierarchical architecture). This document lists the different design architectures that Ruckus offers using its high-end enterprise switches to increase resiliency, as well as scalability, while reducing network complexity and network touch points.

How to Build a Robust Campus Network

Robust networks must be properly planned and cannot be constructed haphazardly by simply putting standalone components together. The network designers need to identify the network requirements, identify best solutions to meet the needs of the business, and plan for future expansion.

Network Requirements

Today's business networks must be available nearly 100 percent of the time and they should also be smart enough to protect against unexpected security threats. The networks should be designed in a manner that accommodates changing traffic loads as well as maintaining consistent application response times. The following list of requirements must be kept in mind while building an intelligent network:

- Network uptime
- Reliable application delivery, as well as response times
- Network security
- Network adaptability to growth and business changes
- Ease of network troubleshooting

Fundamental Design Principles

The following design principles must be kept in mind:

- Compiling the network requirements
- Analyzing your existing network
- Preparing for a preliminary design
- Deploying the network
- Monitoring and redesigning the network
- Maintaining design documentation

Performance and Availability

In addition to the mobility required for campus network access, a campus network must also accommodate a wide spectrum of performance and availability requirements for client application access. Many business applications are adequately supported by conventional Fast (100 Mbps) or Gigabit Ethernet (1 GbE) connectivity, although some very high-performance client applications require 10-GbE links. With current improvements to wireless technologies, access points (APs) could push up to 5 Gbps that are suitable for laptop and mobile applications. Specific applications, such as VoIP, require additional performance guarantees in the form of Quality of Service (QoS) support and traffic policing.

Redundant network devices and links for High Availability (HA) are required for all mission-critical applications and must have failover capability in the event of an individual link or interface outage. A properly designed campus network infrastructure must be sufficiently flexible to provide the required bandwidth and availability per workgroup or application as business requirements change. The variability of client connectivity requirements also impacts other layers of the network infrastructure as client traffic is funneled to the network core and data center.

Security

The large number of client devices at the campus layer poses an ongoing security challenge. A network is only as secure as its weakest link, so a large, dispersed campus network must be purposely provisioned with distributed security mechanisms to eliminate vulnerabilities. Access Control Lists (ACLs), authentication, virtual private networks (VPNs), Media Access Control Security (MACsec), Internet Protocol Security (IPsec), and other safeguards restrict network access to only authorized users and network devices and block the penetration into the campus network itself. Financial and health-related industries are now obliged to protect customer and patient information to comply with government regulations. Compared to the security mechanisms typically in place in the data center, the campus network is far more vulnerable to malicious attack. Security for the campus network must therefore be constantly reinforced and monitored to avoid exposure. The Ruckus ICX campus switches are equipped with all the necessary security features to defend against vulnerabilities in campus networks.

Management

The campus network is dispersed inherently due to the diversity of client devices. Centralized uniform management is essential for maintaining performance and availability and for enforcing corporate security policies. As new technologies, such as wireless LAN (WLAN), are introduced to facilitate user access, the campus network management framework must integrate new device and security features to ensure stable operation and provide the necessary safeguards against unauthorized intrusion. Comprehensive integrated network management tools, such as SmartZone (SZ), Ruckus Cloud, and other automation tools offered by Ruckus, can monitor traffic patterns throughout the campus network to proactively identify potential bottlenecks for network tuning on both wired and wireless network devices.

Reducing Operational Expenses

With potentially thousands of workstations, laptops, smartphones, and other end devices, and hundreds of access points, network switches, and routers, the campus network represents a substantial hardware investment. One component is the initial cost of the equipment itself, but footprint, cooling, and power consumption also contribute to the ongoing total cost of ownership (TCO) expenses. Due to the dispersed nature of the campus network infrastructure, these costs are less readily identified than comparable operational overhead in the data center. However, they should still be factored into the overall campus network design and product selection. Integrating more energy-efficient network infrastructure elements and leveraging technologies such as Power over Ethernet (PoE) dramatically reduces ongoing operational expenses (OpEx) and minimizes the impact of the network on the corporate budget. In addition, consolidation of network assets by using more efficient high-port-count switches both streamlines management and reduces energy consumption. The 1RU campus switches offered by Ruckus can essentially be a substitution for chassis because they support distributed stacking as opposed to being centralized to one location.

Flexible Design and Its Benefits

To meet the most essential design goals (including scalability, availability, security, and manageability), a network must be built for flexibility as well as growth. A hierarchical design is used to group devices into multiple networks in a layered approach. The following three basic layers make up the hierarchical design:

- Access layer
- Aggregation layer
- Core layer

Access Layer

The access layer connects end users and devices. A tiered campus network design provides the flexibility to support multiple capabilities at the access layer (or network edge). Depending on application requirements, high-performance clients can be provisioned with multiple 1- or 10-GbE interfaces for maximum throughput to the aggregation and core layers.

Due to the high availability, high-performance, and security requirements that can vary from one department to the next, the access layer switch infrastructure should provide multiple speeds, rapid failover capability, and VPN and other security protocols as required. Unified communications, such as concurrent VoIP, streaming media, and conventional data transactions, may require additional functionality for QoS delivery and PoE. In addition, applications requiring wireless connectivity need both wireless LAN (WLAN) access points as well as centralized management to ensure stable and secure connectivity.

Access layer switches are typically housed in wiring closets distributed on multiple floors of each building on the enterprise campus network. These in turn are connected to aggregation layer switches that feed traffic to other segments or to the network core. To accommodate the fan-in of multiple access layer switches to the aggregation layer, high performance uplinks are required. Currently, these are up to 100-GbE uplinks, which can be provided with ICX switches.

Aggregation Layer

The aggregation layer interconnects smaller local networks. The campus aggregation or distribution layer funnels transactions from multiple access layer switches to the network core. Because each aggregation layer switch is responsible for multiple upstream access layer switch traffic flows from hundreds of users, aggregation layer switches should have high availability architectures—including redundant power supplies, hot-swappable fans, high-performance backplanes, redundant management modules, and high-density port modules. Aggregation layer switches are typically Layer 2 or Layer 3 switches with support for robust routing protocols to service both the upstream access and downstream core layers. Ruckus ICX switches support full IPv4 and IPv6 protocols, RIPv1/v2, OSPFv2/v3, and BGP.

Core Layer

The core layer connects the aggregation layer devices. The core layer represents the heart of the data network infrastructure. Transactions from campus clients to data center servers or to external networks must pass through the core with no loss in data integrity, performance, or availability. Core switch architectures are therefore designed to support 99.999 percent ("five nines") or greater availability and high-density modules of high-performance ports. The Ruckus ICX 7850 switch can provide the services and high-bandwidth capabilities needed in the core layers.

Collapsed Core and Distribution

The three-tier model is widely used in enterprise networks that scale over a period. Certain small business networks do not necessarily grow over time; these networks are small enough to be served by a collapsed core and distribution design. Ruckus recommends a few designs to cater to these networks because it reduces the overall cost of deployment, as well as OpEx.

Ruckus Campus Network Solutions

Ruckus Campus Network Solutions Overview

Ruckus offers a full suite of Layer 2 and Layer 3 solutions that are engineered for enterprise-class applications combining high performance with resiliency and industry-leading energy efficiency. For campus LAN deployments, Ruckus offers a broad spectrum of access, wireless access, and aggregation switches to build a complete campus LAN solution for both medium and large enterprises.

Ruckus Wired Campus Network Designs

There are four different design solutions that Ruckus recommends for building campus networks, depending on the network scale (as shown in [Table 1](#)).

- Routed Access (Two/Three-Tier Model)
- Stack Architecture (Two-Tier Model)
- Two Cores with Multi-Chassis Trunking (MCT) and VRRP-E (Two/Three-Tier Model)
- Campus Fabric (Two-Tier Model)


TABLE 1 Reference Designs

Reference Designs	<500 Ports	500–3000 Ports	>3000 Ports
Routed Access	No	Yes	Yes
Stack (Core/Agg)	Yes	Yes	Yes
MCT	No	No	Yes
Campus Fabric	Yes	Yes	No

Routed Access (Two/Three-Tier Model)

The idea of placing multilayer switches in the access layer yields significant advantages simply because they can fully utilize all uplinks to the distribution layer (loops are no longer broken by the Spanning Tree Protocol (STP)). But most customer networks like to extend the Layer 3 architecture so that it is easier to implement for certain applications. In a routed access network, the Layer 3 virtual interfaces (VIs) are at the access layer. A hybrid Layer 2 and Layer 3 network can have both routed links as well as a Layer 2 access network by simply using the 802.1q (dot1q) trunking feature across these links flowing down to the access layer.

FIGURE 1 Routed Access



Advantages

- Equal cost multiple path (ECMP) can be implemented to maximize link utilization to the core
- The design avoids Spanning Tree Protocol (STP) implementation complexity

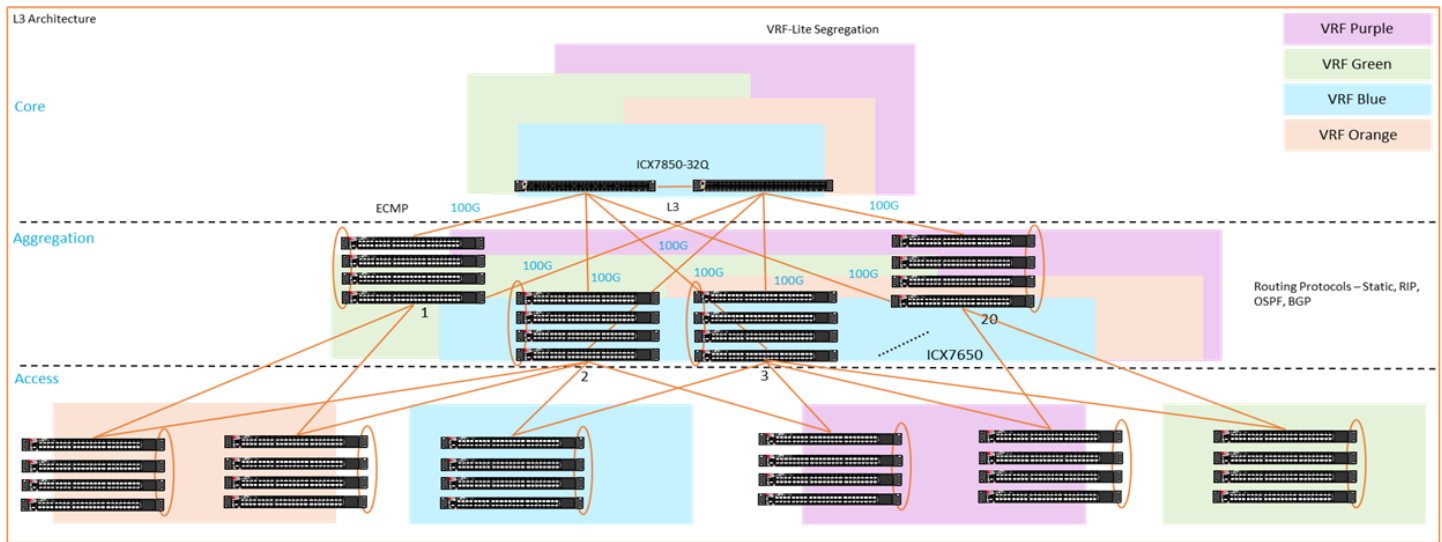
Considerations

- Campus VLANs cannot be used between the access stacks
- Complexity is increased on IP address management
- There is an added cost to the premium licensing in the switch
- VLANs cannot be re-used between stacks (Workaround would be to use VXLAN between stacks at the access layer to replicate VLANs)

VRFs in Routed Access

A variation of the routed access can be constructed by introducing VRFs at the core, distribution, and access layers to provide segmentation at the Layer 3 level. The core switches in this case can also be configured as Multi-Chassis Trunking (MCT) clusters to provide Layer 2 isolation. This hybrid design has proven to be successful in large customer enterprise networks.

FIGURE 2 Routed Access with VRF




Stack Architecture (Two-Tier Model)

A stack is a group of devices that operate and are managed as a single entity. A Ruckus stack contains from 2 to 12 units configured in either a ring or a linear topology. The units in a stack are from the same model family; that is, a stack can be any of the Ruckus ICX 7000 series switches. Ruckus stackable devices are connected through ports that can be configured for either stacking or data. The location of stacking ports and the configuration options differ by device type.

In a stack architecture design, as shown in Figure 3, the uplinks from the access switch stacks are typically aggregated to form Link Aggregation Groups (LAGs). This increases the uplink capacity, as well as providing redundancy at the access layer.

FIGURE 3 Stack Architecture



Advantages

- Long distance stacking of up to 10 km is supported
- Availability of In-Service Software Upgrade (ISSU)
- A maximum of 12 ICX switches of the same model with different SKUs can be stacked together
- There is an active controller and a standby controller for redundancy
- The design provides for Layer 2 simplicity and fast failover


Two Cores with Multi-Chassis Trunking (MCT) and VRRP-E (Two/Three-Tier Model)

Multi-Chassis Trunking (MCT) can be used in a three-tier model. MCT is a technology that allows two MCT-supporting switches to cluster together and appear as a single logical device. Trunking is a technology that allows multiple links of a device to appear as one logical link. The combination of MCT and trunking allows for creating a resilient network topology that utilizes all links in the network, creating an ideal network topology for latency-sensitive applications. Dynamic LACP trunks provide link-level redundancy and increased capacity to the access layer.

VRRP-E is a Ruckus proprietary protocol that was designed to eliminate a single point of failure in a static default-route environment by dynamically assigning virtual IP routers to participating hosts. A virtual router is a collection of physical routers with interfaces that must belong to the same IP subnet. VRRP-E adds redundancy at the Layer 3 level. Each device in VRRP-E can be configured to route an upstream Layer 3 network, which essentially provides an efficient deployment.

Figure 4 illustrates an MCT design that can be deployed in enterprise networks as well as in certain data center scenarios. The MCT cluster can be formed using a pair of Ruckus ICX 7650, ICX 7750, or ICX 7850 switches. The MCT client stacks can be connected to the MCT pair subsequently. Dual links are used to provide redundancy as well as link aggregation.

FIGURE 4 MCT Cluster with VRRP-E



Advantages

- This design provides redundancy at the distribution layer with two Active-Active switches
- Subsecond failover at the distribution layer
- Flow-based load balancing
- STP-free design

Considerations

- MCT clusters along with stacking is not supported

- The maximum number of MCT clients supported is 50
- LACP on ISL is not supported
- GRE on the ISL VE interfaces is not supported
- STP is not supported on MCT VLANs
- IPv6 is not supported

Ruckus has aggressively been testing a dual-layer model that can scale out the number of MCT clients. Figure 6 shows a scaled-out design that can support up to 12 MCT client pairs, followed by 50 MCT clients to each MCT pair.

FIGURE 5 Dual-Layer MCT Design

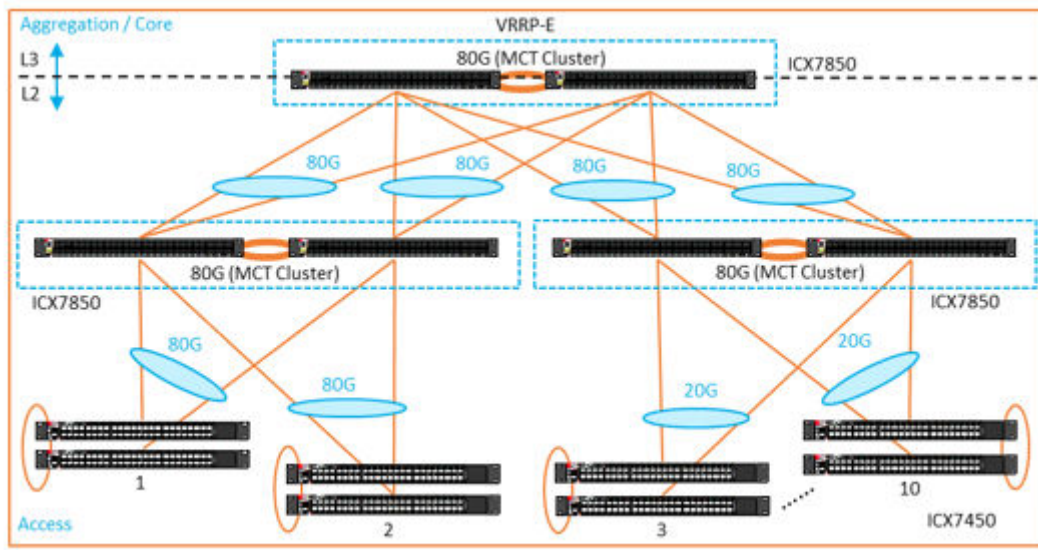
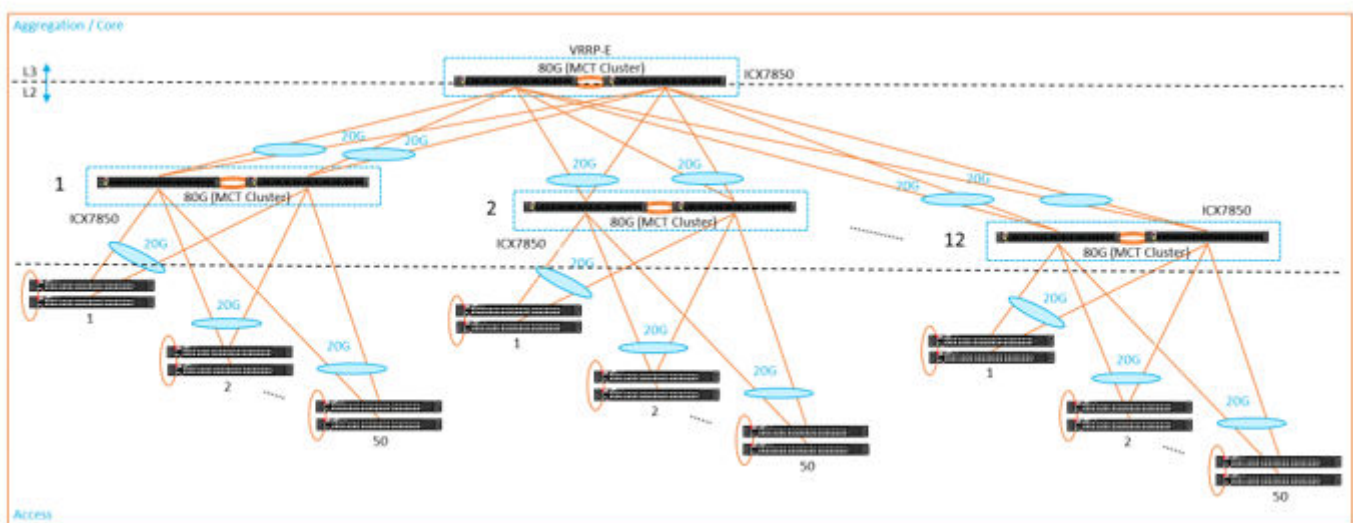



FIGURE 6 Scaled-Out Dual-Layer MCT Design



Campus Fabric (Two-Tier Model)

Campus Fabric creates a more scalable architecture based on the IEEE 802.1BR standards. The Ruckus Campus Fabric architecture shown in Figure 7 can support Ruckus ICX 7750 or ICX 7650 switches as stack units that can be configured as Control Bridges (CBs) and the Ruckus ICX 7150, ICX 7250, and ICX 7450 switches as Port Extender (PE) units.

FIGURE 7 Campus Fabric



The Campus Fabric design can be used primarily as a two-tier model that is suitable for medium-sized campus networks where the core and distribution layers can be collapsed into a single layer. The Campus Fabric domain can contain from 1 to 4 CB units. A maximum of 36 PE units can be supported in a domain. Assuming there are 48 ports in each device, the domain can support up to 1800 ports.

Advantages

- Campus Fabric supports a distributed architecture as opposed to a bulky chassis architecture
- Seamless mobility with Layer 3 boundaries between physical locations
- Redundancy at the Control Bridge, aggregation, and core levels is available

Considerations

- The Control Bridge can support up to four Ruckus ICX 7650 or ICX 7750 switches, limiting the number of 10-GbE ports to 192
- IP addressing is not supported on the PE units
- Reload is required with the use of STP

Conclusion

Campus networks have unique design criteria that provide greater flexibility to support a wide range of client devices and applications. Wired access, wireless, and PoE devices coexist at the network edge. Mission-critical client applications require high bandwidth, high availability, and security. A tiered campus network architecture with fan-out to client devices at the access layer, consolidation of access layer traffic through the aggregation layer, and centralized routing through the network core provides a scalable model for growing the campus network over time and accommodating higher traffic volumes and multiple protocols, as required. The full suite of Ruckus intelligent campus network IP infrastructure solutions and comprehensive network management tools enable customers to build and expand robust, cost-effective, and business-optimized campus networks that meet both current and future corporate requirements.



© 2019 ARRIS Enterprises LLC. All rights reserved.
Ruckus Wireless, Inc., a wholly owned subsidiary of ARRIS International plc.
350 West Java Dr., Sunnyvale, CA 94089 USA
www.ruckuswireless.com