RE-ARCHITECTING ENTERPRISE NETWORKS WITH HIGH-PERFORMANCE STACKABLE SWITCHES

RUCKUS
an ARRIS company

WHITE PAPER

INTRODUCTION

In recent years the evolution of packet processor and CPU technology, the heart of an Ethernet switch, has allowed unprecedented levels of performance and functionality to be incorporated into fixed form factor (FFF) switches which were previously only available to much larger and more costly chassis-based systems. This situation was highlighted by Gartner analyst Andrew Lerner in his blog post on the topic of fixed form factor switching in the data center, which applies equally well to campus edge networks, concludes "the emergence of high-density fixed form factor switches can reduce or eliminate the need for costlier, oversized chassis-based switches in the data center. The move toward FFF switches will help network managers deliver higher-performance networks and reduce footprint, power, cooling and TCO."

Fixed form factor switches, often referred to as "pizza boxes" or "stackables", offer some compelling advantages over their chassis-based counterparts which have not been evolving at the same pace. A stackable solution offers true pay-as-you-grow economics not just in terms of acquisition costs but also ongoing support and power consumption resulting in a significantly lower Total Cost of Ownership.

MEETING THE REQUIREMENTS OF THE HIGH-PERFORMANCE ENTERPRISE

Enterprise networks have evolved from being utilities used by a limited selection of employees within an organization to become mission-critical infrastructure that supports every business function. The services that networks support are now central to the operation of every function; manufacturing data, stock control, accounting records, customer order management, or simply voice and video communication. Everything depends on the network. If the network is going to keep up with the demands of business it must be fast, flexible and secure but most of all it must be optimized for both performance and cost.

Ruckus has a long history of innovation in wired and wireless networking, driving the technologies forward to offer ever greater levels of performance and efficiency and evolving solutions to meet emerging requirements. The ICX range of switches were designed from the ground up to deliver the performance of a chassis system in a compact, cost effective and flexible package.

RE-ARCHITECTING THE NETWORK: SCALE-OUT VS. SCALE-UP

Traditional network architectures typically demand large power-hungry chassis-based devices in the aggregation and core layers of a campus network. In the past, these were necessary to provide the required resiliency and scale, but they are often underutilized within a compromised network design. This resulted in increased operational costs that could mainly be attributed to disproportionately high-power consumption and cooling expenses, often as a result of an underutilized chassis with empty slots and over specified power supplies which were sized to support a chassis full of the most power hungry cards. To address the limitations of chassis solutions, Ruckus introduced the concept of scale-out networking. Scale-out networking exploits the Distributed Chassis architecture.

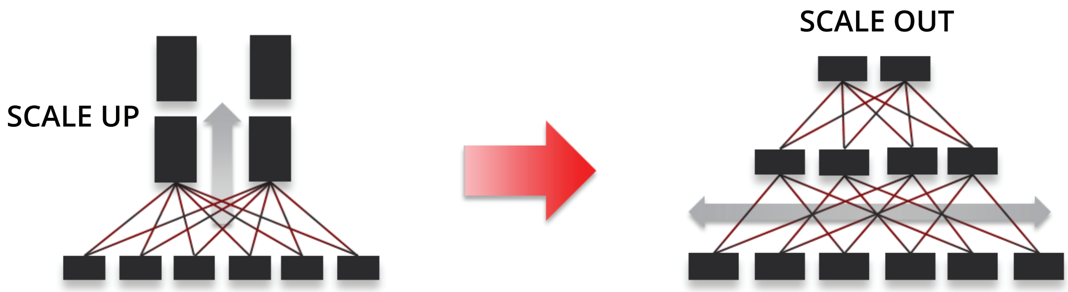

Figure 1: Evolving the network - Scale-up vs Scale-out.

Scale-out networking delivers flexible deployment options using the Ruckus ICX 7000 stackable switches that can be interconnected to create a unified system. Such a system delivers the single point of management of a chassis-based system, but with greater scalability and power efficiency matched with true pay-as-you-grow economics.

Distributed chassis switches offer scaling that is very often beyond what is possible with a chassis. For example, a stack of 12 x ICX 7650 switches deliver 576 downstream ports at speeds of 1/2.5/5/10GbE which can all operate at line rate concurrently. Copper ports can deliver up to 90Watts of 802.3bt compliant PoE power per port ensuring that all the latest applications and end-point configurations can be supported. In addition to the downstream ports such a system can also support multiple 10GbE, 40GbE and 100GbE connections which can be used for either stacking, computer or uplink connections. This level of scalability combined with flexible deployment options allows the creation of highly customized and optimized solutions that cannot be matched by a rigid chassis-based system.

For high-performance network core applications, the ICX 7850 provides support for 10, 25, 40 and 100GbE connections which when stacked offers unprecedented levels of performance and scalability, a single system can deliver up to 384 40/100GbE ports which can all operate at line rate concurrently in a non-blocking system.

All the Ruckus ICX 7000 platform families support up to 12 switches per stack ensuring that the benefits of scale-out networking can be exploited in all sizes of deployment and at all layers of the network.

Ruckus continues to innovate in this arena, ensuring that the network is always optimized for performance, energy efficiency and management simplicity.

CAMPUS FABRICS—DELIVERING THE NEXT GENERATION THE CAMPUS NETWORK

Built on the ICX 7000 platforms the Ruckus Campus Fabric solution collapses multiple network layers into a single logical device, combining the power of a Distributed Chassis design with the flexibility and cost-effectiveness of stackable switch building blocks. The result is a unified system that simplifies deployment and management while unifying features and services across the campus.

The traditional aggregation/core layer is replaced by a stack of high-performance ICX 7750 or ICX 7650 systems which deliver a unified network control plane that acts as the central management and traffic forwarding authority for the entire Campus Fabric domain. At the edge up to 36 high-performance ICX 7150, 7250 or 7450 switches act as virtual line-cards for the core system providing up to 1,728 user ports (with or without PoE/PoE+). Connections between any two devices in the Campus Fabric can be up to 10Km ensuring maximum deployment flexibility and scalability of the system. Together the entire Campus Fabric domain is managed as one logical device from a single point of management with traffic routing and load-balancing within the system managed automatically.

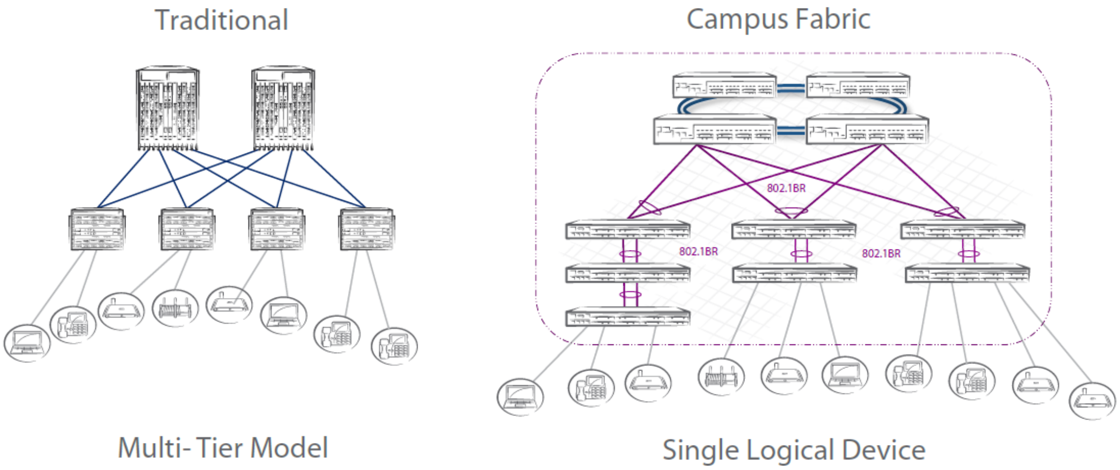

Figure 2: Ruckus Campus Fabric architecture versus a traditional multi-tier campus network.

The Campus Fabric solution does not require dedicated hardware (standard Ruckus ICX switches can be configured to participate in a Campus Fabric system) and is a no cost option that can be deployed simply via a configuration change to standard ICX 7000 switches.

More information is available in the Campus Fabric White Paper.

DEPLOYMENT OPTIMIZATION

In order to optimize the cabling within a campus network it is necessary to place the edge switches as close as possible to the devices which they will serve as this ensures that the length of the cable runs are minimized and the cable management headaches associated with centralized chassis-based solutions are eliminated. This reduces the cost of deployment and makes it simpler to build in resiliency by having cable runs from different wiring closets serving different areas of a building. Furthermore, shorter cable runs reduce the power loss on PoE connections thus reducing the load on the switch power supplies and saving energy, something that is becoming increasingly important as the number and power consumption of PoE devices increases. Stackable switches are perfectly optimized for distributed deployment allowing ports to be placed precisely where they are needed ensuring that all cable runs are as short as possible.

Because Ruckus switches use standard Ethernet media and transceivers for the inter-switch stack connections the deployment options are greatly increased. If standard copper stacking cables are used, then the inter-switch connections can be up to 5 meters long which is usually sufficient for locally distributed stacks such as in Top of Rack (ToR) Data Center or wiring closet applications where the switches are in close physical proximity. For broader distribution, up to 10 kilometers between switches, fiber-optic cables can be used. This allows a stack to be deployed across multiple physical locations while presenting a single management entity with a common logical data-plane. As the switches do not need to be co-located, they can be placed wherever the ports are required and then interconnected to provide a resilient system that can be managed as a single entity. Furthermore, uplinks from the distributed stack can be connected to switches in different locations but grouped together into a Link Aggregation Group (LAG) which ensures that all of the available bandwidth can be utilized while also increasing resiliency. The use of stacking technology and LAGs creates a resilient load-balanced system without the need to deploy sophisticated Layer 3 routing protocols and their associated complex configurations.

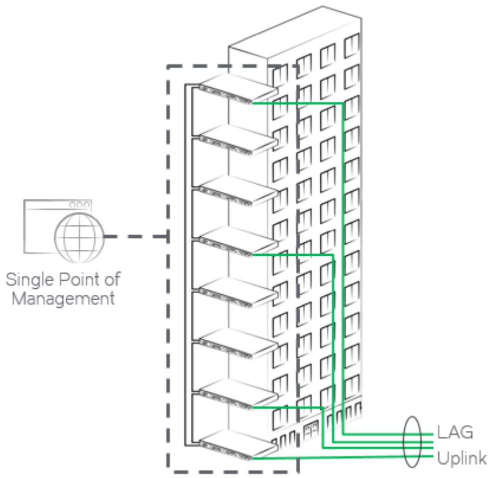

Figure 3: Distributed stacking allows switches to be placed where they are needed but still be managed as a single entity.

All of the benefits associated with stacking are also applicable to the Campus Fabric solution which also provides increased scale, up to 36 edge devices vs 12 in a stack, and the opportunity to use a variety of platforms in a unified system.

Building a comparable solution using a chassis results in either non-optimal cable routing (i.e. all cables must be routed back to a central location) or multiple chassis which will not be optimized for port density, must be managed individually and will be considerably costlier to acquire and operate. Furthermore, the stacked solution offers automated resiliency and load-balancing between the distributed switches whereas the chassis solution requires a complex protocol-based solution (e.g. using STP or L3 routing). The significant advantages of the stacking and Campus Fabric solutions can be summarized thus;

Layer 2 simplicity

Inter-switch links do not need to be considered as part of the overall network topology, so they can be used to provide resiliency without the need for Layer 3 routing to manage traffic flows.

No idle links

The inter-switch links are “internal” to the switches and as such are not seen as part of a layer 2 network, therefore all links can remain open and can all be used to carry traffic simultaneously thus maximizing throughput and ensuring very fast recovery from link failures.

Fast failover

Due to the rapid detection and recovery techniques employed on stack links, failure of a link or a switch will result in hitless failover with no impact on user services. Furthermore, complex inter-device redundancy solutions such as VRRP are not required.

Simplified management

Even when physically distributed all the switches can be managed as a single entity enabling one- touch configuration changes via a single IP address.

AVAILABILITY

Before the development of reliable high-speed stacking technologies deploying chassis-based systems was the only way to gain access to system-level resiliency features such as dual power supplies, fans and processors. The result was often compromised solutions as budgets were stretched to accommodate expensive chassis-based systems which were the only way to deliver the functionality that businesses demand.

Recent developments in packet processor technology matched to increasingly sophisticated software capabilities now allow the deployment of solutions based on stackable switches which offer levels of resiliency and performance that are comparable with a chassis. Furthermore, the flexibility of a stackable switch-based solution allows the design to be optimized while still meeting stringent requirements for performance and availability.

With a truly decentralized architecture stacked switches offer redundancy distributed across all devices, so there is no single point of failure. For added resiliency power supplies and fans, the traditional weak points of any system, can be duplicated to provide further protection against outages. To protect the control and forwarding planes the Ruckus solution includes a Hitless Failover function which ensure that in the event of the master switch within a stack or Campus Fabric failing the standby will take-over the control functions in under one second and without disrupting packet forwarding. If this should occur, then a new standby switch is elected so that the system is never without an active back-up for the master switch. To further enhance resiliency within the wiring closet uplinks from the distributed system can be connected to any combination of switches and combined into a Link Aggregation Group (LAG) which allows the connections to form a bonded high-speed connection which also offers high-speed recovery in the event of a link failure.

To reduce the impact of system maintenance Ruckus has developed rolling In Service Software Updates (ISSU) that allow the operating system of switches to be updated with minimal service impact. A standard software update for a stack requires a concurrent reload of all the members in the stack which can disrupt network traffic as connectivity is lost and traffic must be rerouted. Ruckus' ISSU solution allows the members of a stack to be updated and individually reloaded thus minimizing disruption as traffic is only interrupted for one switch at a time. Best practice dictates that connections from a Distributed Chassis to other nodes in the network should be duplicated and connected to different members of the stack therefore connectivity is maintained for the duration of the update process and the end users do not experience any interruption of service.

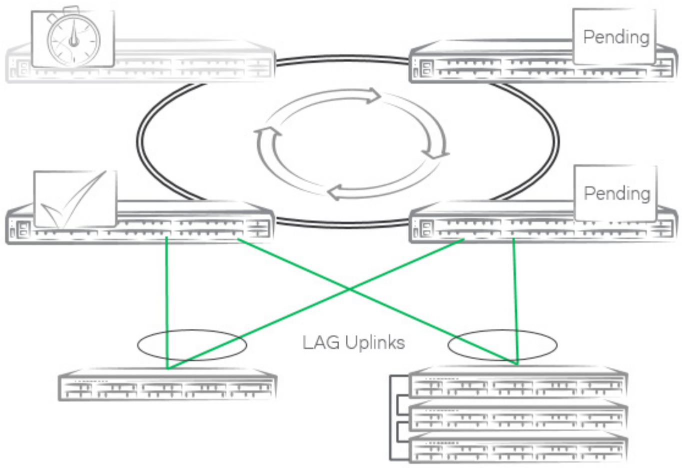

Figure 4: Rolling ISSU ensures that dual-homed devices never lose connectivity.

PAY-AS-YOU-GROW—CHASSIS CAPABILITIES, STACKABLE PRICES

Ruckus redefines the economics of enterprise networking by providing chassis-like levels of performance and availability matched with the flexibility of fixed form factor stackable platforms.

One of the biggest challenges that chassis-based systems pose is the up-front costs associated with the system infrastructure (chassis, power supplies, switching matrix, etc.) that must be borne even when deploying a small number of ports. When building a network consideration must always be given to likely future growth both in terms of the number of connected devices and also the traffic volumes that it will be required to carry. Therefore, the sizing of the chassis is critical if highly disruptive "fork-lift" upgrades are to be avoided as demand grows and the result is often that systems are over-specified creating a very high initial cost and the deployment of underutilized and unnecessarily power-hungry systems.

By contrast stackable switches offer true pay-as-you-grow economics with a lower cost-per-port whether you are deploying large numbers of 10/100/1000 Mbps edge ports, Multigigabit (2.5/5/10 Gbps) ports or high density 10/40/100 Gbps core systems. With each switch that is added to the system the forwarding capacity is increased in line with the increase in port count thus ensuring that performance is never compromised. Furthermore, as each device is added to the system the number of Inter Switch Links (ISL) is increased as the connections between the devices are an extensible backplane that grows as the system expands ensuring line-rate and non-blocking performance at all times.

NO COMPROMISE PERFORMANCE

Ruckus' Distributed Chassis solution delivers flexible deployment options that allow ports to be located where they are needed while providing a fully extensible high-speed backplane that is sized to match the capacity of the base platform. Ruckus' ICX 7000 series platforms offer the following Distributed Chassis bandwidths:

PLATFORM	MAXIMUM BANDWIDTH PER STACK TRUNK	MAXIMUM NUMBER OF SWITCHES IN A STACK	MAXIMUM STACK BANDWIDTH
ICX 7150	20Gbps	12	480Gbps
ICX 7250	20Gbps	12	480Gbps
ICX 7450	40Gbps	12	960Gbps
ICX 7650	100Gbps	12	2.40Tbps
ICX 7750	240Gbps	12	5.76Tbps
ICX 7850	400Gbps	12	9.60Tbps

Furthermore, each time a switch is added to the system it increases the total forwarding capacity as each switch performs local switching for packets that enter and leave on the same device, packets destined for a port on another switch in the stack are forwarded directly to the target device. This contrasts with most chassis-based platforms which have a fixed forwarding capacity that is shared between the system slots regardless of the number or type of line cards installed. Some higher capacity chassis systems do benefit from distributed forwarding, but these come with a very high cost penalty.

The performance advantage and extensible scalability of the Distributed Chassis is provided without compromising operational efficiency as the resulting cluster of devices is managed via a single IP address and viewed as a single entity even when switches are deployed in different locations. And if a resilient topology is used (e.g. ring) then the system can facilitate fast recovery in the event of a link failure thus ensuring uninterrupted service to any connected devices.

All of the ICX 7000 platforms can be deployed in homogeneous stacks using local or long-distance links of up to 10Km for maximum flexibility. For more details please go to www.ruckusnetworks.com/icx.

ENERGY EFFICIENCY

As there are no empty slots in a Distributed Chassis solution its power consumption is directly related to the number of ports deployed, with chassis-based solutions power and cooling are sized assuming that the chassis is fully populated with the most power-hungry cards so over specified and therefore not optimized. The net result is that chassis-based systems typically suffer from unnecessarily high-power consumption, which in turn drives up the cooling requirements further increasing the Total Cost of Ownership (TCO) of the complete system.

When the deployment of a network is optimized to ensure that it is precisely sized for the application without compromising scalability or performance, the impact on the overall power consumption of the system is significant. Ruckus continues to innovate in this arena, ensuring that the network is always optimized for both performance and energy efficiency.

More information on this topic is available in the Ruckus white paper Energy Efficiency in Campus Networks.

CONCLUSION

In the past only chassis-based systems could offer the performance, scale and resilience that organizations demand in critical areas of their networks but advances in processor and software technology mean that today stackable fixed form factor switches can comfortably meet or exceed the most demanding requirements. As a result, the latest fixed form factor switches offer the most cost-effective solution to acquire and operate without the need to compromise on functionality or performance.

Designed to meet the most demanding technical and commercial requirements Ruckus' ICX range of edge switches allow the creation of highly scalable and reliable networks that ensure that every enterprise is equipped to support the most demanding applications. And when combined with innovative solutions such as Campus Fabrics enterprises can build highly converged networks while minimizing operational overheads.

RUCKUS™ an ARRIS company

350 West Java Dr., Sunnyvale, CA 94089 USA

www.ruckusnetworks.com

19-02-D