The Cognitive Campus

An Architecture for Modern Campus Networks

Introduction

As the concept of campus specific networking enters its fourth decade, it has become clear that traditional approaches to campus infrastructure are struggling to cope with the proliferation of new IoT devices, connected applications and dynamic changes to users, their behavior and expectations.

The first two decades focused primarily on wired connectivity for personal workstations and printers with shared access to file servers, e-mail applications and basic intranet and internet connectivity for static content. This drove broad connectivity but required only limited capacity and security based on perimeter policies. From the early 2000s, emerging wireless and IP telephony were typically implemented as parallel solutions, carried in large stretched VLANs or isolated networks with little integration to existing infrastructure and no holistic approach to IP services.

The third decade saw substantial increases in device numbers, demand for ubiquitous wireless connectivity for portable devices and terminals, bring-your-own-device (BYOD) and a transition to latency and delay sensitive applications such as voice and video becoming broadly adopted. Cloud-based collaboration and streamed rich media hailed the transition from a best effort connectivity approach towards a true fundamental business enabler. The simple architectures and management paradigms of the early campus struggled to manage the plethora of devices and evolving security threats with limited visibility, segmentation and application awareness.

As we begin the fourth decade, the expectations of the campus network begin to shift once more as the world embraces flexible working practices, a greater reliance on Internet-based voice services and rich video conferencing, on premise and cloud based applications, convergence of legacy IT systems into cloud based IoT mandating zero down-time and unprecedented threats to data security from ever more advanced malware and malicious actors with ransomware putting organizations and businesses out of action for days if not weeks.

In order to support these new demands, a fundamental shift in campus thinking is required.

Early Evolution of the Campus

As demand for campus desktop computers and networking connectivity grew, the need to support increasing numbers of connected devices and to provide for multiple building-use configurations led enterprises to flood-wire office floor-spaces, delivering enough RJ45 jacks to support projected demands over the 10-15 refresh lifetime of a building.

This approach resulted in a large number of unused connections which, combined with relatively low throughput requirements of contemporary applications, meant that the average campus wiring closet was heavily under-utilized and therefore built to a $/port policy.

The focus on cost created a need for a class of cost-optimized products that didn't require the attributes of the dense core platforms of the day. These were realized in several ways:

- Limited function, Layer-2 only products to reduce complexity and cost
- Low cost fixed switches that could be 'stacked' to create a virtual chassis using proprietary protocols and reduce the cost of fiber
- Heavily oversubscribed modular systems with only a few uplink ports

These approaches aimed to solve the same set of problems:

- Simplified Management: Providing a single administrative management point without the need for loop-avoidance protocols
- Incremental Expansion: Providing a "pay as you grow" deployment model where additional switches/line cards may be added as more capacity is needed.
- Aggregated Connectivity: Minimizing the number of uplinks from a wiring closet in the face of fiber scarcity
- Cost Optimization: Avoiding the need for full-function devices at the network edge

Stacking switches gained popularity and quickly became the primary architecture for the campus wiring closet. The deployment of devices with limited functionality also led to de-facto design methodologies like the deployment of large Layer-2-only networks that are unsuitable for modern enterprises.

The Mature Campus

As solutions matured and enterprise demands increased, the drawbacks in deploying low-function devices with tight control plane coupling became more apparent. Underlying architectural deficiencies were band-aided with a range of measures including an explosion in the number of VLANs, Private VLAN for local user isolation and segmentation, storm and broadcast control to prevent rogue devices compromising large L2 subnets and Per-VLAN Spanning Tree (PVST and MST) implementations to decouple spanning tree topology changes and attempt to utilize uplinks more evenly.

The 'simple Layer-2' topology was no longer simple; configurations became more complex with a requirement for unique configurations per system, switches required more software updates to address software bugs and power users demanded ever more bandwidth. These new levels of complexity and demand revealed deficiencies in the physical infrastructure, fragile operational moves-add-and-changes (MAC) and maintenance needing to be implemented out of hours - for example:

- Merging the control plane of the stack members added significant complexity to the device operating system and the task of adding or removing capacity, leading to more bugs and unexpected failures and complex maintenance procedures.
- The inherent fate-sharing nature of the stack architectures meant software issues would often impact the entire stack, removing the benefit of redundancy and affecting large numbers of users. Alternative options to provide redundancy and rapid-failover brough more complexity and less reliability

- The proprietary nature of stacking architectures limited the choice of products, made intergenerational upgrades difficult or impossible and prevented multi-vendor interoperability - emerging products could not be integrated into existing stacks.
- Stacking backplanes became increasingly oversubscribed as user ports increased speed and density. The proprietary nature of a stack often resulted in an inability to determine the real available performance and difficulty in troubleshooting without vendor expert assistance.
- More deployment flexibility was required - members within a stack are distance and topology limited due to the proprietary stacking cables or the timing requirements of the centralised control plane. Servicing emerging user needs without breaking the architectural model was difficult.

Together, these limitations created a fragile environment with high maintenance overheads and significant challenges for IT departments tasked with delivering more demanding applications such as voice, video and the secure segmentation required by the evolving business need and user behavior. In some cases, enterprises resorted to deploying data center class products at the campus edge to overcome limitations of the traditional campus.

A Cognitive Campus Architecture for 2020+

With the benefit of more than a decade of experience working with the world's most demanding customers and complex infrastructure demands, Arista's vision for the future of campus networking employs open and proven technologies deployed at enormous scale to deliver a new class of devices and a new paradigm for meeting the needs of the fourth era of campus networking.

The modern campus must deliver against new capabilities:

- Full, open automation of day 1 and day 2 management tasks
- Unprecedented reliability and uptime for the always-on enterprise
- Deep visibility into traffic flows, users and devices
- Zero-trust threat detection and mitigation, especially for IoT and BYO/unmanaged devices.
- Dynamic and secure segmentation to support flexible working and broad IoT deployment
- High performance for wired and wireless devices and rich applications
- Exceptional Quality of Experience for real-time voice and video services

The combination of Arista's industry leading EOS™ operating system and CloudVision(R) suite, with the advanced campus-focused CCS wired switching platforms, Wi-Fi and Awake security provide the foundations for enterprise campus workplaces for the next decade.

A Closer Look at the Cognitive Campus Approach

First-Class Functionality

A significant step forward in realizing an infrastructure suitable for emerging workloads is to remove the historical design compromises enforced by low function campus edge products.

Modern merchant silicon, derived from Data Center class products, provides rich functionality to support a robust and resilient Layer-2 and Layer-3 or overlay based architecture with the addition of campus focused enhancements. This inherent functionality flexibility is critical to designing and building the right foundations for a highly scalable and reliable user workspace.

Open Standards

As has been decisively proven in the Data Center, the use of open, standards based protocols in all aspects of the infrastructure is key to ensuring simple day to day operations, transparent scalability and long term flexibility for growth and interoperability.

It is no longer necessary to follow a siloed approach where campus, Data Center and WAN follow incompatible disciplines and require different skills and tools. Arista has championed the use of a modern, single image, fully modular and stateful operating system that provides fully process isolation and industry leading availability and reliability. EOS provides a rich consistent core feature set across all products to ensure that any product may be deployed in any use case if required.

Automation

In a large enterprise, automated network provisioning, maintenance and management are critical considerations when supporting distributed environments with dynamic requirements. Manually managing individual devices - whether in the campus or data center - is error prone, time consuming and operationally does not scale.

Arista products expose a common API and multiple programmable interfaces to help customers automate deployments, updates and security patches through a broad range of tools.

Additionally, Arista's CloudVision offers a unique approach to building the state-driven cognitive management plane, providing many benefits beyond single point management. CloudVision provides an end-to-end topology aware single point of management, from the data center to the campus and WAN. CloudVision makes it possible to not only automatically deploy and update the estate, but also to ensure ongoing compliance, monitor live telemetry streams and alert administrators to any software defects or security issues, providing a mitigation path.

Day 1: Commissioning & Configuration

CloudVision starts with the premise that all devices should be zero-touch provisioned into the network. CloudVision manages configuration deployment automatically using a hierarchical scheme to ensure consistent deployment throughout the infrastructure. Devices can be logically grouped into containers by function, location or other common attributes regardless of their vintage or product family without dependency on proprietary protocols. Using logical groupings with template driven configuration avoids the significant percentage of failure conditions associated with intra-stack incompatibilities.

This approach means it is possible to move from operations that apply only to stacks of switches, to floors, buildings, remote offices and entire campuses with plug and play simplicity to best fit your organization's needs.

Day 2: Change Management, Compliance and Visibility

Post deployment, the dynamic nature of modern campuses requires a robust change control process. Using the same organizationally focused paradigm described above, CloudVision's workflows enable the planning, automation and roll back of configuration and image changes without the need for CLI intervention. This simplifies management, reduces technical prerequisites for operations personnel, and reduces deployment times from days to minutes and avoids failures introduced by human error or unexpected interdependent control plane complexities.

A deployed infrastructure must be monitored for compliance to ensure that reliability, security and other relevant regulations are upheld. This is traditionally a laborious and error prone manual process which, when neglected, leads to grey failures and unpatched security vulnerabilities.

CloudVision continually monitors for unexpected or unauthorized changes to the expected steady state environment, presenting a global view via its Compliance Dashboard. In addition to maintaining configuration consistency, CloudVision automatically provides alerts for potential software defects or newly discovered security advisories (e.g. PSIRTs) and a clear path to quickly remediate any affected devices through automated update.

Infrastructure-wide visibility is also critical to maintaining a high level of service. CloudVision acts as the engine to collect, process and visualize thousands of parameters in real time from Arista's streaming telemetry integrated into every device. Continuous monitoring of both connected devices and the health of the underlying infrastructure ensures administrators have relevant information at their fingertips to discover and investigate anomalies anywhere in the domain.

Additionally, the Awake AI-driven Security Platform provides the umbrella of real-time threat detection and modeling critical to maintaining data integrity. Awake continuously monitors network communications autonomously to discover, profile and classify devices, users and applications. The platform is capable of detecting complex adversarial threats as they emerge, building a picture of the insertion, attack surface and stages of a security breach enabling rapid response and deep forensic analysis to limit the impact of a compromise and quickly determine both remediation and the scope of any malicious activity.

Growth and Expansion

Historically, stackable architectures were perceived to provide an easy way to add capacity. In practice, adding capacity with traditional stacking architectures imposes several limitations and interdependencies. The added switch needs to not only be from the same vendor as the other members of the stack but the same model, software version, and license pack as well.

As the development of networking silicon has accelerated and the demands of new workloads and connected devices have increased, enterprises cannot afford to be trapped with a lowest-common denominator model.

Stacking solutions inherently hinder administrators from taking advantage of newer switching platforms. However an open and decoupled approach enables the introduction of up-to-date devices with important new capabilities, increased performance or simply access to better price points for an enterprise.

A stack investment is not well preserved when the only expansion option is deploying an obsolete switch, and capacity is not easily added if manual preparation and downtime needs to be scheduled.

Alternatively in the 2020 era, the Arista Extensible Operating System (EOS) provides multiple options for campus access topologies including the self-healing active-active Multi-Chassis Link Aggregation (MLAG) architecture. MLAG delivers a standards based architecture that protects against spanning tree loops while delivering high availability and traffic load balancing without the limitations of proprietary stack architectures.

Switches can be added to an Arista CMP, Cognitive Management Plane, without disruption to the other switches and independent of the switching platform. The Arista architecture is based on "always on" principles from cloud designs where hitless expansion is critical - with zero downtime. By avoiding the complexity of fate-sharing combined control planes, Arista's solution allows independent deployment of any number of devices of any product family within the wiring closet. It is possible to deploy mixed copper, fiber, POE, mGig, low cost and advanced products as suits the specific use case together without compatibility issues.

2020 Campus Connectivity

Traditional Stacking Architecture

Switches in a stacked wiring closet architecture are commonly connected with proprietary stack ports in a ring topology. The stack ports provide an inexpensive means to connect 2 or more switches, however the stack bus and its protocols are vendor proprietary and the switches often need to be in close proximity to meet strict timing requirements of the combined control plane. Below is an example of a common stack architecture that provides a theoretical 240G of bidirectional connectivity.

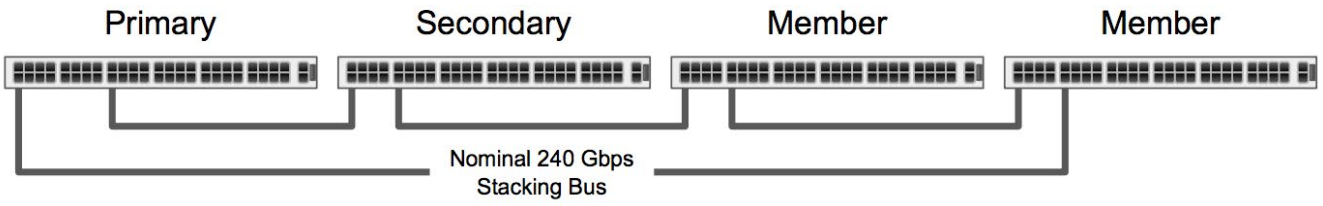

Figure 1: Common stack connectivity model

Solutions for the Wiring Closet

Arista offers a broad range of campus focused devices, ranging from low density fixed format devices to high density modular chassis. For specific requirements, customers may also choose to deploy Data Center class products such as the 7050X3 series within the campus.

This is made possible because all Arista products run the EOS operating system with its rich Layer-2 and 3 capabilities, including VXLAN overlay networking across all products - enabling customers to select the best-fit technologies for each deployment.

Figure 2: Arista's Broad Portfolio of Fixed and Modular Campus Access Platforms

Traditional deployments vary widely in terms of the number of tiers (e.g. core, distribution, IDF, access etc.) and where the boundary between Layer-2 and Layer-3 is placed. Arista's architecture is highly flexible to support both traditional deployment models, aligned to the existing physical layer, as well as optimized 'Spline' designs for higher bandwidth requirements with fewer points of over-subscription and less fiber needed.

Many organizations are migrating away from an all Layer-2 access design by moving the Layer-3 boundary further towards the network edge. Factors driving this change include improved policy deployment, automation, smaller fault domains, better segmentation for fine grained security, and easier troubleshooting with deterministic forwarding paths. Both Layer-2 and Layer-3 access modes are supported and are discussed below in Connecting the Spine.

Multi-chassis Link Aggregation (MLAG)

Layer-2 networks based on the Spanning Tree Protocol (STP) can be difficult to maintain. Negative experiences with STP are common and often stem from misconfigurations that cause large scale outages due to the broad convergence domain in large Layer-2 environments.

Arista's MLAG solution was developed to offer a flexible multi-homing mechanism that eliminated the requirement for STP in Data Center networks, where redundancy is a common requirement and loops and resultant broadcast storms are extremely disruptive.

MLAG builds on Link Aggregation Control Protocol (LACP) to allow links to be negotiated and aggregated across multiple devices without causing loops. Operating on a device to device basis, MLAG eliminates the need for a domain wide protocol such as STP and provides flow based load balancing and ensures all links are active while removing the need for network-wide topology calculations and convergence after changes.

MLAG is extremely well suited for the wiring closet as it allows administrators to build large, horizontally scalable networks with efficient traffic forwarding and minimal complexity.

Connecting Switches Within the Wiring Closet

Arista recommends the use of standard Ethernet interfaces that can be deployed either as uplinks or interconnections as appropriate for the use case. Figures 3 to 6 show various design scenarios.

In one example, an Arista wiring closet architecture constructed from ten 48 port UTP switches can provide over 200Gbps of uplink connectivity to 480 UTP clients and up to 80 10/25G clients connected via high speed optical interfaces; more than sufficient for most IDF closet PoE requirements (See figure 6).

Administrators can similarly use MLAG with higher density switches to drastically increase the scale of connectivity. For example, by using a 96 port platform, campus engineers can either double the port density, to 960 stacked UTP ports, or deploy a 480 port cluster using only 5 switches.

Similarly, MLAG can be used with modular access switches to provide load balancing of user traffic to the network core in addition to uplink resilience and redundancy for wiring closets requiring extremely high density user and device connectivity.

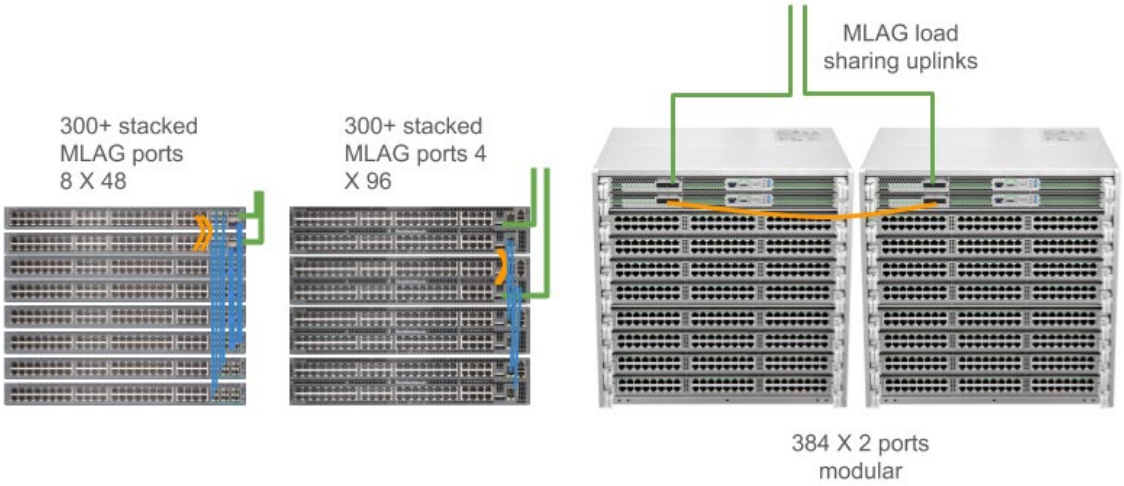

Figure 3: Scaling, Load-sharing and Resilience with MLAG

With MLAG, network administrators can easily create a standard template for any deployment configuration and scale it to meet specific performance and budgetary requirements for multiple wiring closet use cases. Below shows some examples of how switches can be connected to form an Arista wiring closet with dual load balanced uplinks traversing an IDF riser.

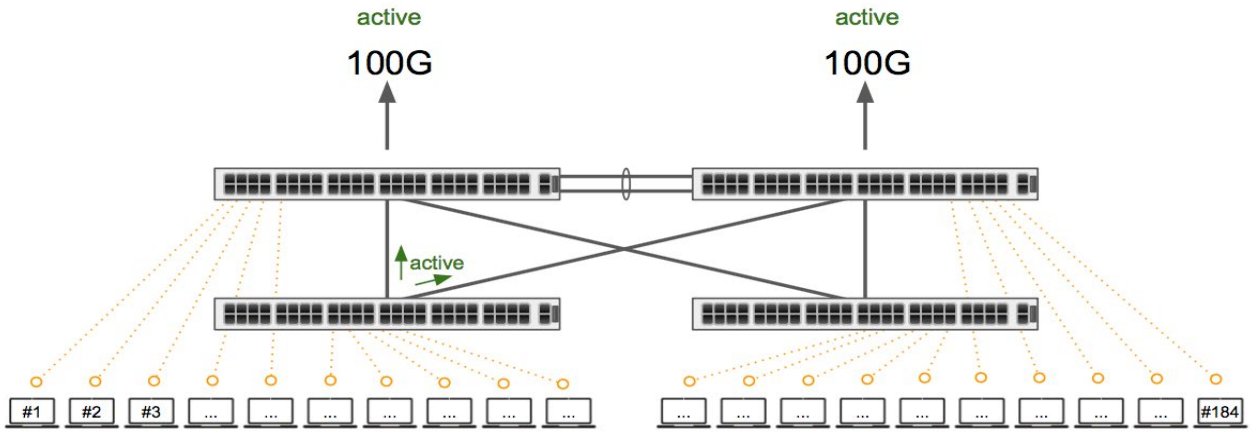

Figure 4: Example Wiring Closet for 184 client ports using 1 RU devices

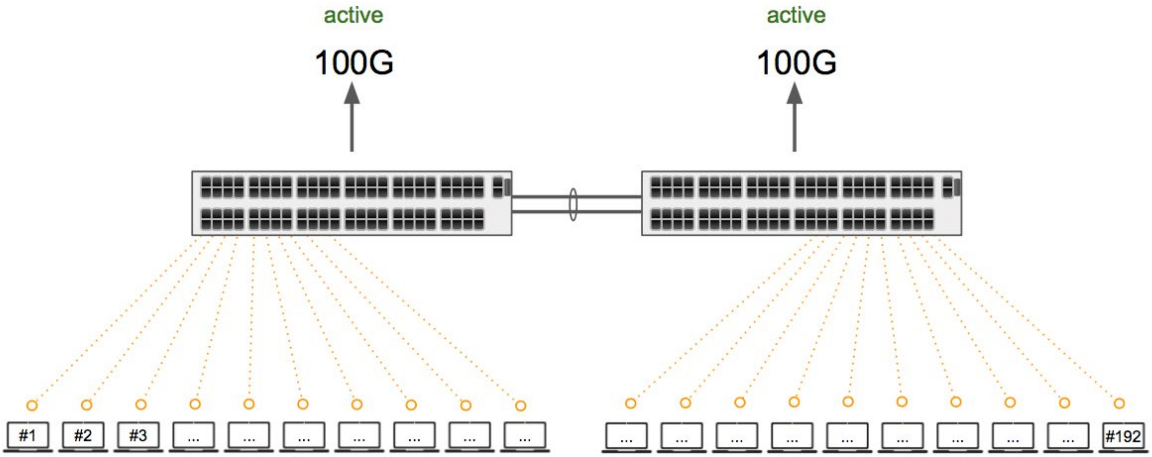

Figure 5: Example Wiring Closet for 192 client ports using 2 RU devices

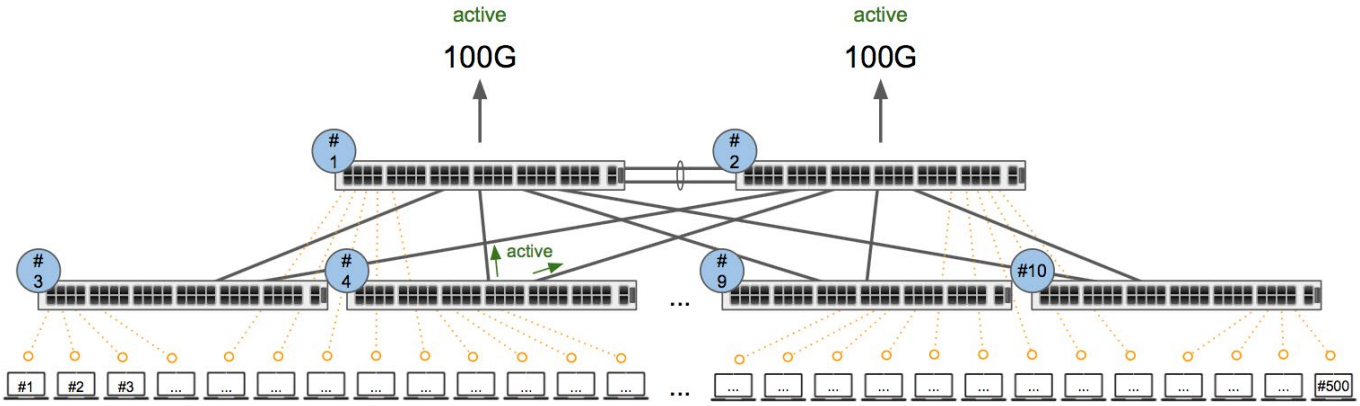

Figure 6: Example Wiring Closet for 500+ client ports using 1 RU devices

Deploying individual devices with MLAG provides many benefits. Additional switches of any model and generation can be added to service emerging needs and since the connecting links are standard Ethernet, the switches do not have to be co-located. Further, unlike a traditional stack with a ring based backplane, growth of an MLAG based closet actually provides additive bandwidth within the closet, rather than adding to the contention of the fixed size stack bus.

This flexibility allows an MLAG based wiring closet to service diverse requirements from 100Mbps to >10Gpbs client connectivity, multiple levels of Power over Ethernet (PoE) demand, unusual building layouts and use-case specific levels of aggregation/ oversubscription within a single product family.

Connecting to the Spline - Layer-2 Uplinks

MLAG may be used at multiple tiers to enable active-active load-balanced connectivity both within the wiring closet and from the closet to the Spline/Distribution layer. In an all Layer-2 access model, the Layer-2 gateways for the access VLANs would reside on the Spline/Distribution or Core layers.

Figure 7 shows connecting five 1RU, 48-port switches for a total of 240 customer facing ports. Substituting with 96 port platforms doubles the total ports to 480 ports in 10RU of rackspace. Redundant configurations consisting of both 'Square' and 'Bowtie' models are possible, depending on the availability of fiber for connectivity between the wiring closet and upper network tiers.

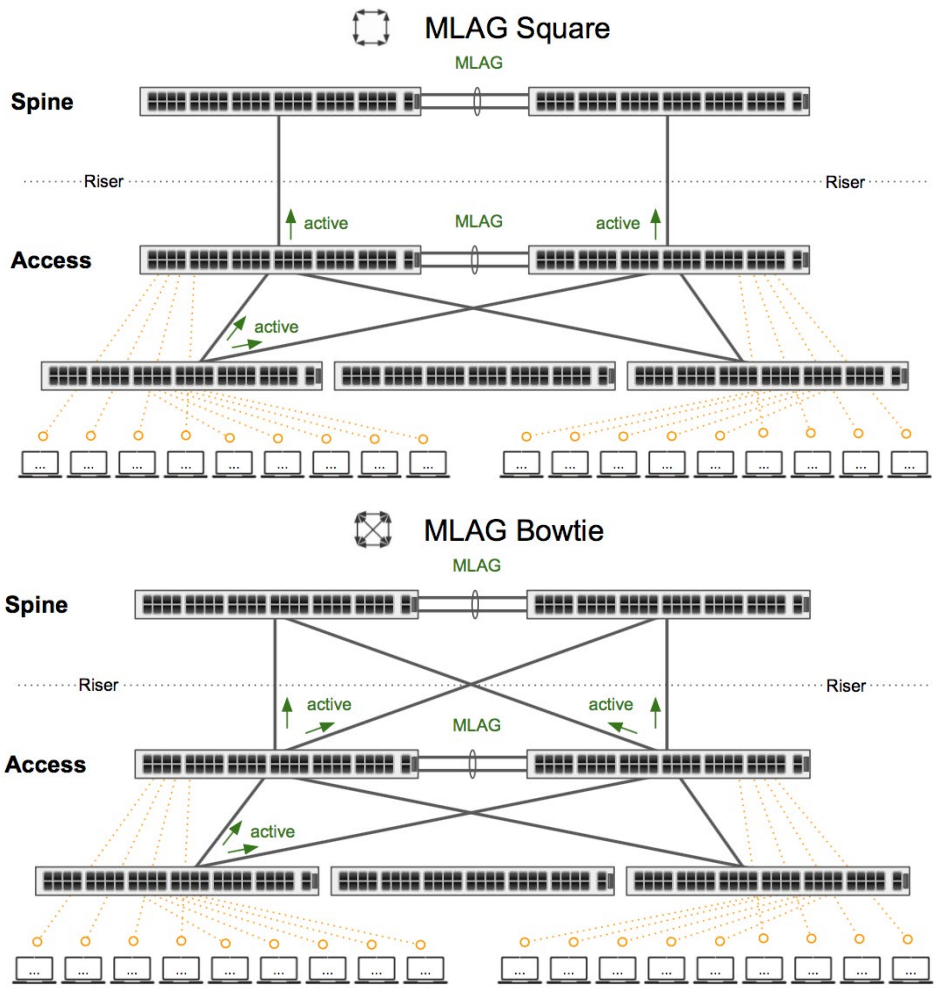

Figure 7: MLAG Square and Bowtie topologies

Connecting to the Spline Layer-3 Uplinks

Moving the Layer-3 boundary to the access layer allows for multiple new redundant topology options. If leveraging a modern leaf-spine design, uplinks benefit from Equal Cost Multi-Path (ECMP) for load balancing of user traffic and increased redundancy from multiple active paths.

Connectivity to a pair of Spline/Distribution switches is the most commonly deployed topology as it offers 1+1 redundancy and resilience to device and link failures. However some organizations, particularly those requiring higher levels of availability or more than 1+1 bandwidth resilience, or those wishing to deploy smaller form factor fixed systems in the spine rather than chassis based modular systems are deploying "n-way" designs where 'n' is commonly 3 or 4 devices, to provide higher capacity and increase scalability.

Figures 8 to 10 below illustrate typical 2-way campus Spline architectures as well as a 4-way alternative.

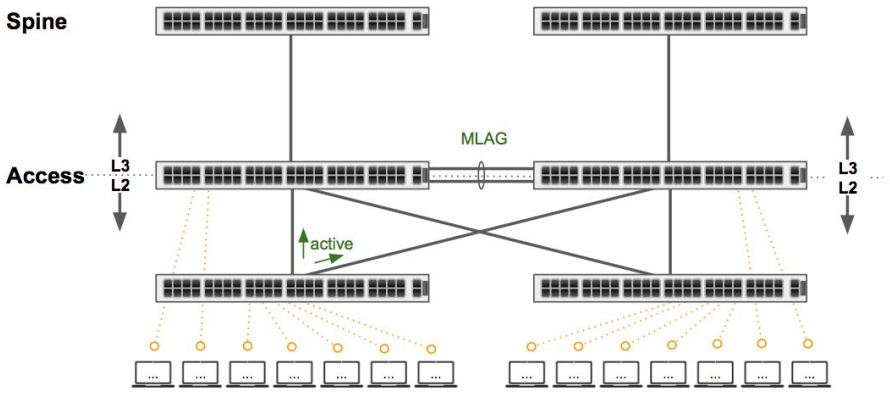

Figure 8: Three-Tier Layer-3 Access with 2-Way Spine

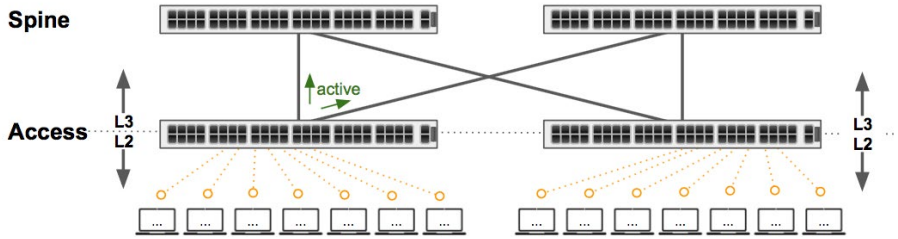

Figure 9: Two-Tier Layer-3 Access with 2-Way Spine

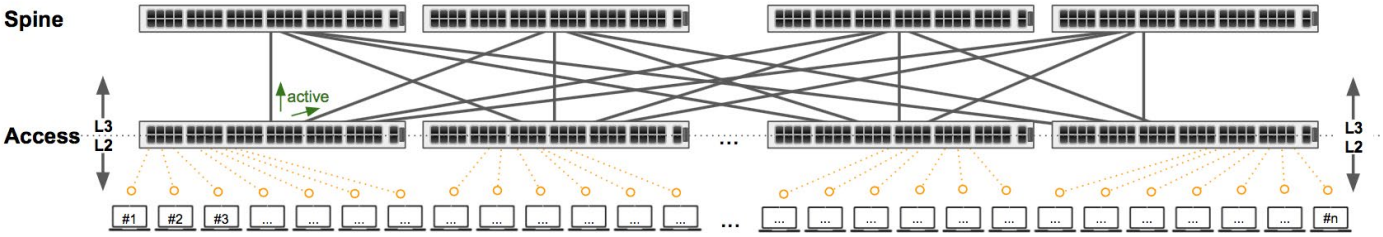

Figure 10: Two-Tier Layer-3 Access with 4-Way Spine

VXLAN for Segmentation and Extending L2 Adjacency over L3 Network

VXLAN is a network overlay technology that provides dynamic provisioning of network segments in the form of Software Defined tunnels overlaid on a normal Layer-2 / 3 topology. VXLAN provides for both Layer-2 adjacency of clients that are separated by Layer-3 infrastructure as well as the possibility of deploying departmental Layer-3 VPNs without the need for MPLS or the complexity of a VRF-Lite model.

In cases where legacy applications require Layer-2 adjacency between multiple closets, moving the Layer-3 boundary to the wiring closet would normally prevent the application from reaching its peers. VXLAN overlays provide a means to maintain Layer-2 adjacency for such legacy applications while allowing modernization of the remainder of the environment.

Furthermore, as the number of connected endpoints continues to grow, there is an increased demand for segmentation as well as abstraction of user networks from infrastructure. Separating user traffic from the protocols that define the underlying topology removes the need to modify the network to accommodate user adds, moves and changes. This simplifies the design and provides for the additional security by not exposing the network topology, network protocols or devices to malicious users.

All Arista platforms for Wide Area, data center and campus networks support VXLAN and EVPN. As a true multi-vendor open standard, VXLAN/EVPN ensures organizations can have a single, continuous, end-to-end approach to segmentation and Layer-2 extension.

Maintainability and Troubleshooting Considerations

Maintainability and troubleshooting are important considerations in a campus. Due to the shared control plane of a traditional stacking solution, malfunctions, bugs or configuration mistakes that lead the stack to split into two can cause significant outages that are difficult to troubleshoot or recover from without manual intervention.

This 'split-brain' condition arises because each divided portion of the stack holds the same configuration and if both are simultaneously online but operating independently of each other, they will introduce duplicate IP addresses and protocol state into the network.

Arista's recommended architecture retains the independence of each individual device and does not rely on a merged centralized control plane. As a result the solutions shown eliminate the risks and are not prone to these types of split-brain situations and the network wide side-effects.

Arista Cognitive Campus Portfolio

Arista's Enterprise campus portfolio consists of a broad variety of platforms designed to meet different workspace use cases. As all Arista platforms run a common, single image operating system (EOS) and deliver consistent features, there are no restrictions as to which platforms can be deployed in a campus environment.

For typical campus deployments, the CCS-7xx product family provides a broad set of common capabilities, including support for Multi-Gig (mGig) and high capacity Power Over Ethernet (POE); other Arista products may be deployed as required and are fully interoperable, running the same EOS operating system and sharing the same protocol stack.

Together with CloudVision, Arista Wi-Fi, Awake Security and the DANZ Monitoring Fabric, the CCS portfolio provides solutions for enterprises of all sizes and end-user requirements.

Table 1 below lists the Cognitive Campus Switches and their high level interface configurations:

Table 1: Overview of Arista CCS Systems

Models	xBASE-T Ports	Port PoE Capability	QSFP and SFP Ports
CCS-758	8 slot chassis with line card options: 48 X 10M/100M/1G 48 X 100M/1G/2.5G 48 X 1G/2.5G/5G/10G	30W and 60W Options	Up to 2 Supervisor modules each supporting either: 2 x 100G-QSFP or 4 x 25G-SFP
CCS-755	5 slot chassis with line card options: 48 X 10M/100M/1G 48 X 100M/1G/2.5G 48 X 1G/2.5G/5G/10G	30W and 60W Options	
CCS-720XP-96ZC2	80 x 100M/1G/2.5G and 16 x 1/2.5/5G	60W	4 x 1/10/25G SFP and 2 x 40/100G QSFP
CCS-720XP-48ZC2	40 x 100M/1G/2.5G and 8 x 1/2.5/5G	30W/60W mix	4 x 1/10/25G SFP and 2 x 40/100G QSFP
CCS-720XP-24ZY4	16 x 100M/1/2.5G and 8 x 100M/1/2.5/5G	30W/60W mix	4 x 1/10/25G SFP
CCS-720XP-48Y6	40 x 10M/100M/1G and 8 x 100M/1/2.5G	30W	6 x 1/10/25G SFP
CCS-720XP-24Y6	16 x 10M/100M/1G and 8 x 100M/1/2.5G	30W	6 x 1/10/25G SFP

Note, for flexibility all 100G ports support both 40G and 100G modes and break-out into 4 x 25G or 4 x 10G links, providing the ability to connect high speed hosts directly to the switch or to use speeds below 100G for connectivity to the distribution or spine layers with scope to increase in future without changing hardware. Any port on an Arista switch can be used for either host connectivity or for uplinks to another switch, with the full suite of Layer-2 and Layer-3 features.

The Arista 7050X3 and 7300X3 family of switches are ideal for BDF/MDF aggregation in traditional 3-tier or collapsed 2-tier topologies requiring high densities of connectivity from 1G to 100G. All these platforms support rich functionality for Layer-2, Layer-3 and EVPN/VXLAN networking using common, open standard features that are consistent with the entire Arista networks portfolio.

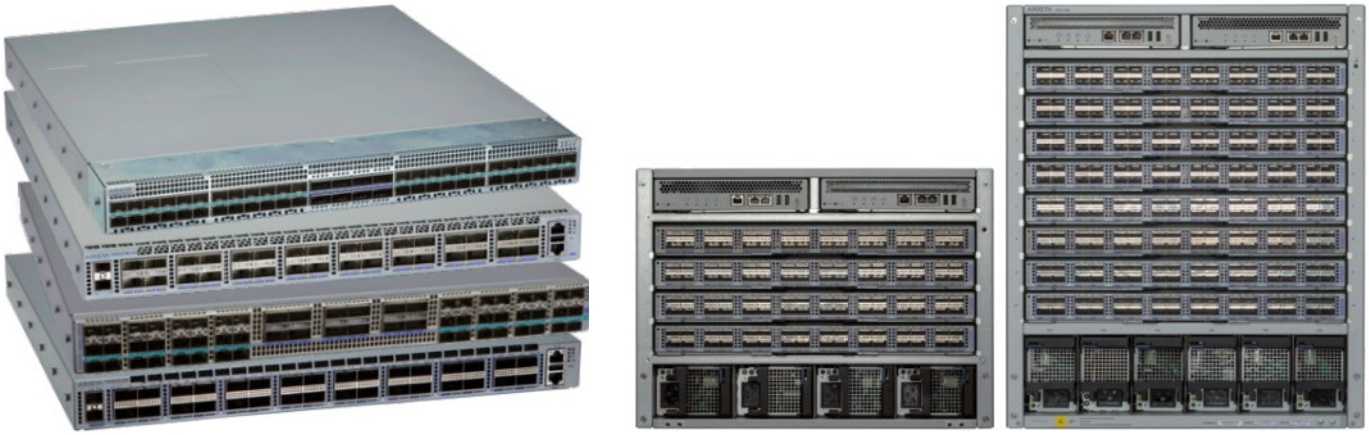

Figure 11: 7050X3 and 7300X3 for high speed aggregation

Summary

Traditional approaches to campus architecture no longer meet the needs of today's modern, IoT enabled, always on campus networks. Optimisations made during an era when lowest cost-per-port was the primary design goal have proven fragile, opaque and operationally expensive in the face of contemporary expectations for business critical infrastructure.

Arista's state of the art campus portfolio addresses these limitations, implementing a modern and open approach that builds on the experience of over a decade working with the world's most demanding Enterprise, Service Provider and Cloud customers.

Replacing traditional, proprietary approaches, the Arista campus portfolio with CloudVision, Wi-Fi, DANZ and Awake eliminates the shortcomings in manageability, security, performance and visibility with industry standard protocols, modern, reliable software and real-time streaming telemetry.

Modern campus management is increasingly critical for consistent day-1 provisioning and day-2 operational management across the enterprise. Arista's Cognitive Campus philosophy provides single point administrative control and oversight, the ability to consistently apply configuration at any level of defined hierarchy, proactive compliance management of software defects and security threats, while offering unprecedented client visibility with telemetry tracking to guard against malware.

Arista's cognitive campus platforms, Extensible Operating System, CloudVision and Awake platforms improve reliability, scalability and security while reducing operational complexity and costs, enabling campus administrators to shift focus from break-fix to pro-actively increasing the productivity of the distributed campus workforce.

Santa Clara—Corporate Headquarters
5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500
Fax: +1-408-538-8920
Email: info@arista.com

Ireland—International Headquarters
3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office
9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office 1390
Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office
Global Tech Park, Tower A & B, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office
9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office
10 Tara Boulevard
Nashua, NH 03062