Celestica Wired Campus Solution Guide

EXECUTIVE SUMMARY

Modern organizations, from educational institutions to corporate campuses, rely heavily on robust and adaptable network infrastructure. Celestica recognizes that a campus network extends beyond simple connectivity, serving as the backbone for seamless communication, stringent security, comprehensive network visibility, and future-proof growth. The increasing reliance on digital technologies necessitates a secure, scalable, and reliable foundation, which Celestica addresses through its Wired Campus Fabric Reference Architecture.

This architecture introduces a proven methodology for designing and deploying medium-sized campus networks. It leverages Celestica's advanced enterprise switching platforms, powered by technologies such as BGP EVPN VXLAN, to deliver unparalleled flexibility, integrated security, exceptional scalability, and cost-effectiveness. A fundamental aspect of this approach is the decoupling of hardware from software, which enables organizations to select open hardware designs, open Network Operating Systems (NOS) such as SONiC Distribution by Celestica, and custom orchestration tools. This disaggregated model provides the freedom to construct a network precisely aligned with specific organizational needs, thereby reducing vendor lock-in and fostering greater agility and cost efficiency.

THIS DOCUMENT WILL HELP YOU:

- Deploy a modern, open, disaggregated campus network architecture with Celestica switches and SONiC Distribution by Celestica NOS as the foundation.
- Enhance campus responsiveness, cost efficiency, security, reliability, and network visibility.
- Create a future-proof network that drives your innovation and growth, enabling new possibilities.

EVOLVING NEEDS OF WIRED CAMPUS NETWORK

Today's campus networks, from higher educational institutions to corporate campuses, are the critical backbone of these organizations. The increasing reliance on digital technologies, the proliferation of IoT devices, and the increasing adoption of cloud-based applications place immense pressure on network infrastructure to be inherently secure, highly scalable, and exceptionally reliable. Traditional networks, often built on monolithic, proprietary equipment, often struggle to meet these accelerating bandwidth, security, and management requirements, leading to operational complexities and significant capital expenditures.

The industry is witnessing a significant paradigm shift from these traditional, vendor-locked networking systems towards open, disaggregated solutions. This movement is driven by a desire to break vendor lock-in, reduce the total cost of ownership (TCO), and foster community-driven innovation and agility within network operations. Celestica leads in the movement toward open networking by delivering flexible and adaptable infrastructure that enables customers to build scalable networks tailored to their specific requirements. This fundamental commitment to openness is not merely a technical specification; it represents a core market differentiator. For customers often burdened by vendor lock-in, limited choice, higher costs due to proprietary licensing, and slower innovation cycles inherent in traditional models, Celestica's open stance directly addresses these pain points.

REFERENCE ARCHITECTURE PRINCIPLES

Celestica's wired campus fabric reference architecture is meticulously designed to deliver on core customer outcomes through specific architectural principles and their practical implementation.

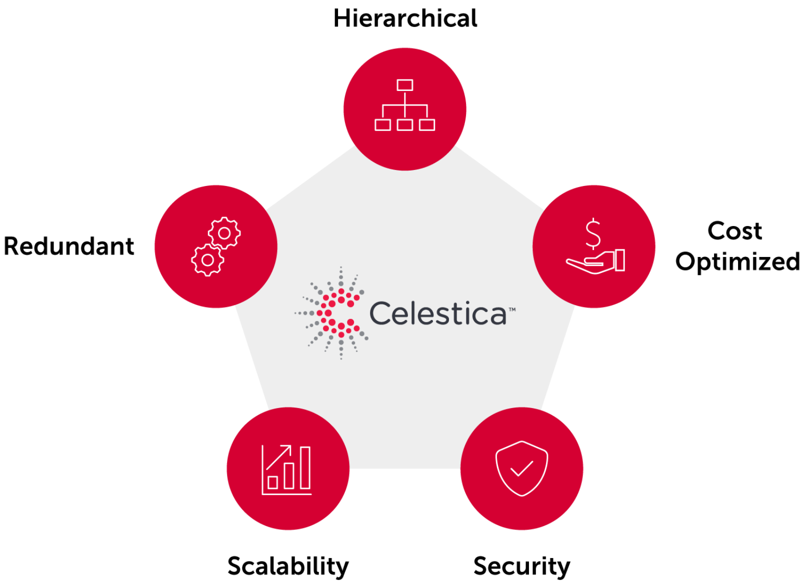

Hierarchical Design

Celestica's architecture supports and evolves traditional hierarchical network designs, adapting them for modern campus demands. While traditional models often feature distinct Core, Distribution, and Access layers, Celestica's approach can leverage a collapsed core model, merging distribution and core functionalities into a single layer for small to medium-sized campuses. For larger deployments, a spine-leaf (Clos-based) IP fabric can serve as the underlay, providing predictable performance and a consistent, scalable foundation, typically with every leaf connected to every spine. This design ensures all paths are active and optimizes East-West traffic flow, a critical requirement for modern applications and IoT devices.

Redundancy and Fault Tolerance

Reliability is built into the Celestica campus architecture through multi-level resiliency.

- **Device-Level Resiliency:** Achieved through features such as Multi-Chassis Link Aggregation (MCLAG), the architecture enables two physical switches to operate as a single logical entity, providing active-active redundancy for connected devices.

- **Link-Level Resiliency:** The architecture incorporates multiple redundant pathways, ensuring uninterrupted connectivity even during individual link failures. In VXLAN EVPN deployments, Equal-Cost Multi-Path (ECMP) routing in the underlay utilizes all available links for traffic forwarding, thereby enhancing resilience and improving bandwidth utilization.
- **Network-Level Resiliency:** Leveraging robust routing protocols (e.g., BGP EVPN) and network virtualization, the solution provides comprehensive network-level resiliency, minimizing downtime and maximizing uptime for critical services. VXLAN EVPN Multi-homing further enhances this by providing active-active redundancy for endpoints connected to multiple leaf switches.

Scalability and Flexibility

Celestica's architecture is engineered for seamless adaptation to increasing bandwidth demands and expanding user bases.

- **Scalable Network Infrastructure:** The design philosophy ensures the network can effortlessly scale as the user base grows and new applications and technologies emerge, protecting investment and preventing costly rip-and-replace scenarios.
- **Network Virtualization:** Through VXLAN, the architecture supports over 16 million unique virtual network segments (VNIs), vastly exceeding traditional VLAN limits. This enables extensive logical segmentation and multi-tenancy, allowing efficient sharing of infrastructure while maintaining isolated data and control planes for different user groups or applications.
- **Openness and Choice:** By decoupling hardware from software, Celestica provides the freedom to choose open hardware designs and an open NOS (SONiC Distribution by Celestica), thereby eliminating vendor lock-in and empowering organizations to build networks tailored precisely to their needs.

Security by Design

Integrated security is paramount, achieved through identity-driven services that enable granular visibility of users and devices, facilitating robust access control and sophisticated network segmentation.

- **Identity-Driven Services:** Enables precise identification of users and devices for robust access control and network segmentation based on identity.
- **Link-Level Security (MACSec):** Built-in MACSec capabilities provide an additional layer of encryption and authentication at the data link layer, safeguarding network traffic from eavesdropping and tampering.
- **Microsegmentation:** Isolates applications and workloads at a granular level within the network, enforcing strict communication policies and limiting lateral movement of threats.
- **Port-based Network Access Control (PNAC):** Supports 802.1X and MAC Authentication Bypass (MAB) for robust device authentication at the network edge.
- **Virtual Routing and Forwarding (VRF):** Provides strong logical separation for different departments or user groups, enhancing security and compliance by isolating Layer 3 domains.

Cost Optimization

Celestica's open networking approach directly contributes to significant cost optimization.

- **Reduced Capital Expenditure (CapEx):** Decoupling hardware from software enables procurement from a diverse supply chain, fostering competition and potentially leading to lower initial hardware costs compared to proprietary vendor hardware.

- **Reduced Operational Expenditure (OpEx):** Automation features, such as Zero Touch Provisioning (ZTP), streamline deployments and reduce manual errors. Leveraging open-source tools for NetDevOps (Ansible, Python) and monitoring (sFlow, Prometheus, Grafana) reduces reliance on expensive proprietary management systems and specialized vendor-specific skillsets, leading to lower ongoing operational costs.

CAMPUS NETWORK DESIGN MODEL: COLLAPSED CORE

This Celestica campus architecture document emphasizes a collapsed core topology ideal for medium-sized campus environments. This model merges the traditional core and distribution layers into a single, high-performance switch or a pair of switches. This simplifies the network topology, reduces the number of devices to manage, and minimizes latency by shortening the path between access and core services.

In a collapsed core model, Celestica's core switching platforms, such as the Celestica DS3001, DS4000/DS4001, or DS5000, serve as the central point for both Layer 2 aggregation and Layer 3 routing. This design can then be augmented with either MCLAG for Layer 2 high availability or VXLAN EVPN for advanced Layer 2/3 virtualization and segmentation, depending on the specific campus requirements.

CAMPUS CORE NETWORK LAYER: HIGH AVAILABILITY AND SCALABILITY

The core network layer forms the backbone of the campus fabric, providing high-speed connectivity, centralized routing, and critical services. Celestica's core switching platforms are designed to support both MCLAG and VXLAN EVPN for robust and scalable deployments.

Celestica Core Network Platforms

Celestica's campus core switching platforms, such as the Celestica DS3001 (64-port 100GbE switch in a 2U form factor with 6.4Tbps bandwidth), DS4000/DS4001 (32-port 400GbE switch in a 1U form factor with 12.8Tbps bandwidth), or DS4101 (32-port 800GbE switch in a 1U form factor with 25.6Tbps bandwidth), serve as the core of the underlay network. These platforms are purpose-built to support high-density, high-bandwidth requirements and are fully compatible with SONiC Distribution by Celestica, offering flexible deployment and management capabilities.

Multi-Chassis Link Aggregation (MCLAG)

Multi-Chassis Link Aggregation (MCLAG) is a crucial technology for building highly available and resilient Layer 2 campus networks. It enables two separate physical switches to appear as a single logical switch to a connected device (e.g., an access switch, server, or firewall). This provides device-level redundancy and high availability, ensuring continuous operation even if one of the MCLAG peer switches fails. In a campus environment, MCLAG is typically deployed at the core (small campus) or distribution layer to provide redundant uplinks from access switches or directly connected servers, minimizing downtime and maximizing network uptime.

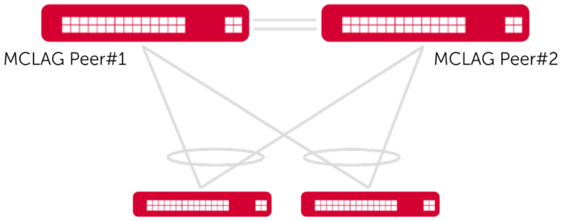

MCLAG and SAG – High Availability and Resiliency: MCLAG in SONiC Distribution by Celestica utilizes LACP to establish a peer relationship between two SONiC switches, enabling them to act as a single logical gateway. They share a virtual IP address (VIP), which serves as the default gateway for clients. Static Anycast Gateway (SAG) works in conjunction with MCLAG to provide Layer 2 gateway redundancy without requiring specialized protocols, such as VRRP. Both MCLAG peers learn and maintain MAC address entries of clients connected to either switch, and these entries are checkpointed between peers. Both MCLAG peers use the same MAC address for the virtual IP address (VIP) or SAG IP, which is unique and serves as the default gateway for clients. If one MCLAG peer fails, the other seamlessly takes over gateway duties, eliminating the need for MAC address relearning or upstream network reconfiguration due to synchronized MAC and ARP tables.

Configuration commands for MCLAG (using SONiC Klish)

Ensure the iccpd Docker state is enabled before configuring MCLAG:

```
Celestica-DS3001(config)feature iccpd state enabled
```

To configure MCLAG domain & source IP address:

```
Celestica-DS3001(config)# mclag domain 1
Celestica-DS3001(config-mclag-domain-1)# interface PortChannel 100
Celestica-DS3001(conf-if-PortChannel100)# ip address 10.41.128.4/28
Celestica-DS3001(conf-if-PortChannel100)# exit
Celestica-DS3001(config)# interface Ethernet 48
Celestica-DS3001(conf-if-Ethernet48)# channel-group 100
Celestica-DS3001(conf-if-Ethernet48)# mclag domain 1
Celestica-DS3001(config-mclag-domain-1)# source-ip 10.41.128.4
```

To configure MCLAG domain peer IP address:

```
Celestica-DS3001(config-mclag-domain-1)# peer-ip 10.41.128.3
```

To configure MCLAG domain keepalive interval:

```
Celestica-DS3001(config-mclag-domain-1)# keepalive 1
```

To configure MCLAG domain session timeout:

```
Celestica-DS3001(config-mclag-domain-1)# session-timeout 10
```

To configure MCLAG domain peer link interface:

```
Celestica-DS3001(config-mclag-domain-1)# peer-link PortChannel 100
```

To configure the interfaces as part of PortChannel:

```
Celestica-DS3001(config-mclag-domain-1)# interface ethernet 16
Celestica-DS3001(conf-if-Ethernet16)# exit
Celestica-DS3001(config)# interface PortChannel 2
Celestica-DS3001(conf-if-PortChannel2)# interface ethernet 16
Celestica-DS3001(conf-if-Ethernet16)# channel-group 2
```

To configure MCLAG member interface (on Port-Channel):

```
Celestica-DS3000(conf-if-Ethernet16)# interface portchannel 2
Celestica-DS3001(conf-if-PortChannel2)# mclag 1
```

Equivalent SONiC Click CLI commands

```
admin@DS3001:~$sudo config feature state iccpd enabled
admin@DS3001:~$sudo config interface ip add PortChannel100 10.41.128.4/28
admin@DS3001:~$sudo config portchannel member add PortChannel100 Ethernet48
admin@DS3001:~$sudo config mclag add 1 10.41.128.4 10.41.128.3 PortChannel100
admin@DS3001:~$sudo config mclag keepalive-interval 1 2
admin@DS3001:~$sudo config mclag session-timeout 1 10
admin@DS3001:~$sudo config portchannel member add PortChannel2 Ethernet16
```

To view MCLAG configuration:

```
Celestica-DS3001# show mclag brief
-----------------------------------------------------------

Domain ID            : 1
Role                 : active
Session Status       : up
Source Address       : 10.41.128.3
Peer Address         : 10.41.128.4
Peer Link            : PortChannel100
Keepalive Interval : 1 secs
Session Timeout      : 10 secs
System Mac           : 00:00:00:02:03:04
Number of MCLAG Interfaces:1
-----------------------------------------------------------
MCLAG Interface Local/Remote Status
-----------------------------------------------------------
PortChannel2       up/up
```

SAG configuration (using SONiC Klish)

The IP address on the VLAN should be configured prior to adding the SAG-IP. The SAG MAC address must be the same across all MCLAG peer devices.

To configure SAG MAC address:

```
Celestica-DS3001(config)# interface vlan 609
Celestica-DS3001(conf-if-Vlan609)# ip address 10.41.6.1/29
Celestica-DS3001(conf-if-Vlan609)# exit
Celestica-DS3001(config)# sag mac 00:BB:BC:02:01:00
```

To configure SAG IP address on VLAN interface:

```
Celestica-DS3001(config)# interface vlan 609
Celestica-DS3001(conf-if-Vlan609)# sag-ip 10.41.6.1
```

Equivalent SONiC Click CLI commands

```
admin@DS3001:~$sudo config vlan add 609
admin@DS3001:~$sudo config interface ip add Vlan609 10.41.6.1 10.41.6.4
admin@DS3001:~$sudo config sag mac add 00:bb:bc:02:01:00
admin@DS3001:~$sudo config sag interface add Vlan609 10.41.6.1
```

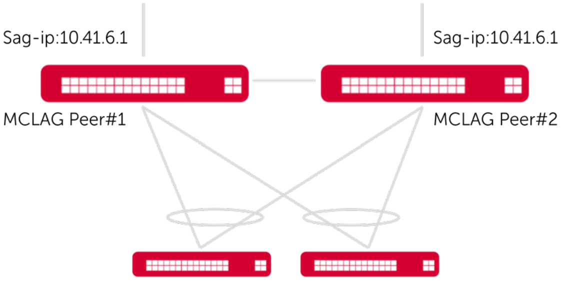

To view SAG Configuration:

```
Celestica-DS3000#  show running-config | grep sag
sag mac 00:BB:BC:02:01:00
sag-ip 10.41.6.1
```

To display SAG interface information:

```
Celestica-DS3001# show sag interface
Vlan                IP address
Vlan609             10.41.6.1

Number of SAG Interfaces:1
```

To display the configured SAG MAC address:

```
Celestica-DS3001# show sag mac
---------------------
SAG MAC
---------------------
00:BB:BC:02:01:00
---------------------
```

Supported Platforms for MCLAG and SAG: DS2000, DS3000, DS4000, DS4100, DS4101, DS5000.

VXLAN EVPN Fabric

A VXLAN EVPN-based campus fabric offers a highly scalable, flexible, and secure network architecture by decoupling the virtual network topology from the physical infrastructure. This approach is ideal for modern campus environments that require extensive network segmentation, seamless mobility for users and devices, and efficient handling of East-West traffic. By building a Layer 2 overlay network over a Layer 3 IP underlay, VXLAN EVPN eliminates the limitations of traditional VLANs (e.g., 4094 VLAN ID limit) and Spanning Tree Protocol (STP), providing active-active forwarding paths and simplified network management.

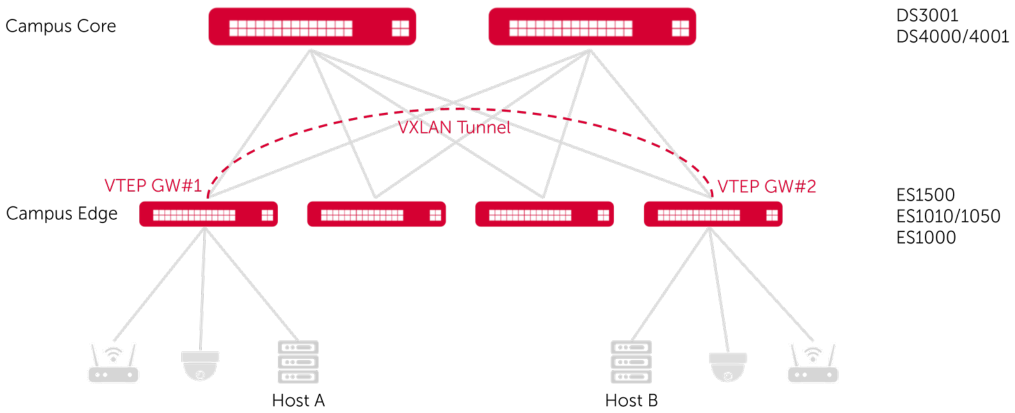

In this spine-leaf architecture, every leaf switch (acting as a VTEP) has a Layer 3 link to every spine switch. The spine switches forward traffic between leaf switches. Logical VXLAN tunnels are established between VTEPs over the Layer 3 underlay. EVPN (MP-BGP) acts as the control plane, exchanging MAC and IP reachability information, enabling Layer 2 segments (VLAN A, VLAN B) to span across the physical fabric, identified by unique VNIs.

Deep Dive into VXLAN EVPN

- **VXLAN Tunnel Endpoints (VTEPs) and VXLAN Network Identifiers (VNIs):** Celestica's switching platforms (e.g., DS2000) serve as VXLAN Tunnel Endpoints (VTEPs). These devices are the crucial "gateways" that connect the Layer 2 overlay network to the Layer 3 underlay network. VTEPs encapsulate original Ethernet frames with a VXLAN header (followed by UDP and outer IP headers) when Layer 2 traffic needs to traverse the VXLAN fabric. Conversely, upon reaching the destination VTEP, the outer headers are removed, and the original Layer 2 Ethernet frame is restored and forwarded to the target device. Celestica VTEPs intelligently route and bridge packets into and out of VXLAN tunnels, supporting both Layer 2 and Layer 3 forwarding within the fabric.

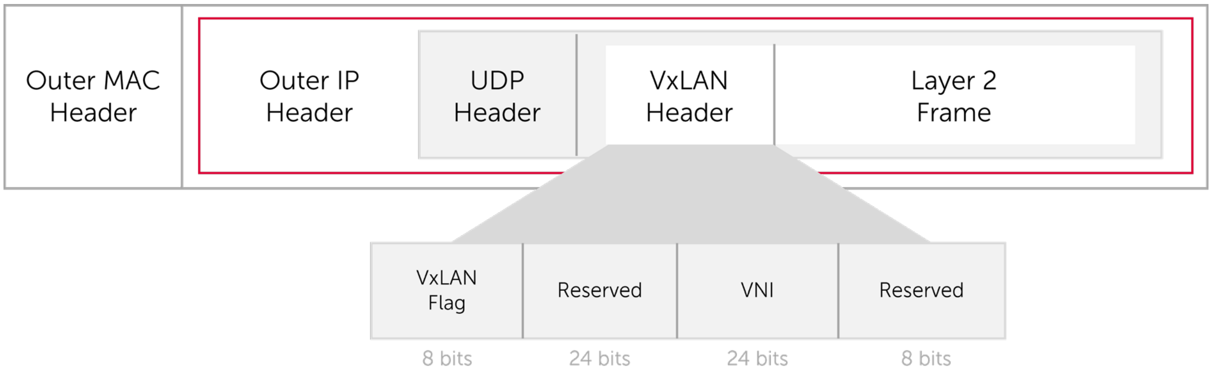

The Virtual Network Identifier (VNI) is a 24-bit identifier that uniquely distinguishes each virtual network segment (VXLAN) within the overlay. Functionally similar to a VLAN ID but with a significantly expanded scale, the VNI allows for the creation of over 16 million unique VXLANs. This vast address space easily accommodates the segmentation needs of large campus environments, supporting numerous departments, research groups, or IoT device categories with isolated logical networks.

- **Ethernet VPN (EVPN) Control Plane – MP-BGP for Scalable Learning:** Ethernet VPN (EVPN) functions as the intelligent control plane for VXLAN tunnels, leveraging Multi-Protocol Border Gateway Protocol (MP-BGP). EVPN's role is critical for efficient MAC/IP address advertisement, enabling VTEPs to advertise endpoint information to other VTEPs across the fabric proactively. This eliminates the need for inefficient flood-and-learn mechanisms for unicast traffic, significantly improving network scalability and performance. EVPN also facilitates dynamic VTEP discovery, simplifying configuration and management by allowing VTEPs to discover each other and establish VXLAN tunnels automatically. By separating the control plane from the data plane, EVPN provides a robust, scalable, and dynamic mechanism for managing the forwarding state of the overlay network.

- **Layer 2 VXLAN – Extending Department/Business Unit VLANs:** Within the Celestica campus fabric, traditional Layer 2 VLANs, often referred to as Department/Business Unit (BU) VLANs, provide dedicated Layer 2 broadcast domains for specific user groups or applications at the access layer. Each Tenant VLAN is mapped to a unique L2 VNI. When traffic from a Department/BU VLAN enters a VTEP, it is encapsulated within a VXLAN header that includes its associated L2 VNI. This process ensures that the traffic remains logically isolated and is delivered only to the correct tenant segment on the remote VTEP.

  Layer 2 VXLAN Mapping: Each Tenant VLAN is mapped to a unique L2 VNI, extending its Layer 2 broadcast domain across the fabric while maintaining isolation.

```
Celestica-DS2000(config)# interface vxlan vtep22
Celestica-DS2000(config-if-vxlan-vtep22)# source-ip 10.1.1.1
Celestica-DS2000(config-if-vxlan-vtep22)# map vni 1010 vlan 10
Celestica-DS2000(config-if-vxlan-vtep22)# map vni 2020 vlan 20
```

i Equivalent SONiC Click CLI commands

```
admin@DS2000:~$sudo config vxlan add vtep22 10.1.1.1 11:22:33:44:55:66
admin@DS2000:~$sudo config vxlan map add vtep22 10 1010
admin@DS2000:~$sudo config vxlan map add vtep22 20 2020
```

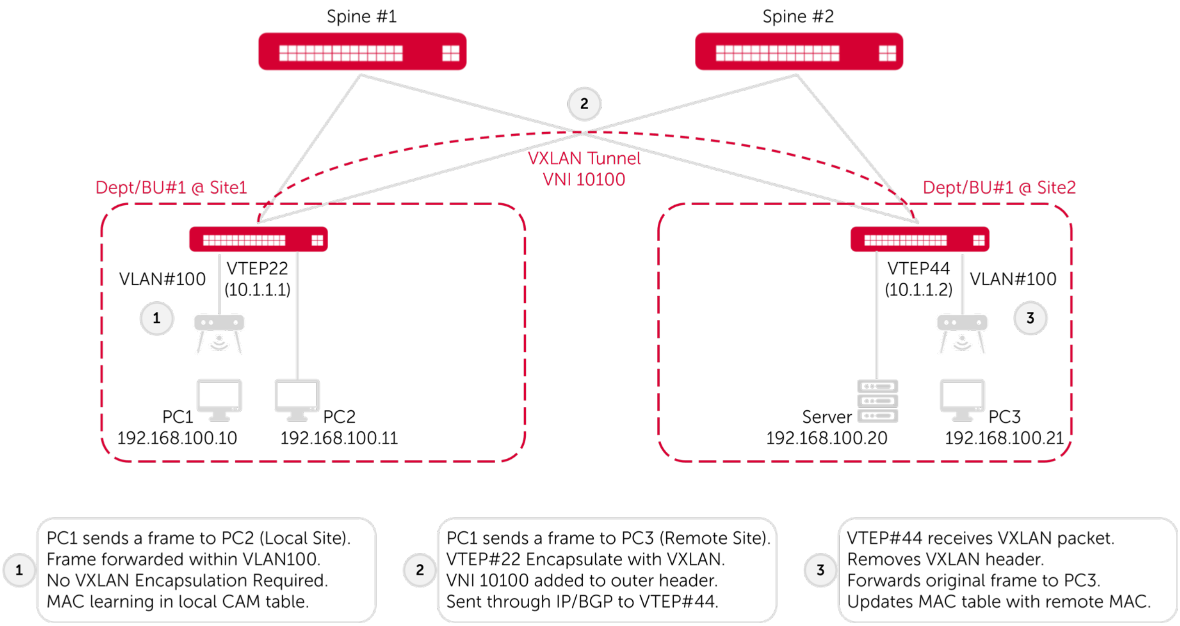

1 PC1 sends a frame to PC2 (Local Site).
Frame forwarded within VLAN100.
No VXLAN Encapsulation Required.
MAC learning in local CAM table.

2 PC1 sends a frame to PC3 (Remote Site).
VTEP#22 Encapsulate with VXLAN.
VNI 10100 added to outer header.
Sent through IP/BGP to VTEP#44.

3 VTEP#44 receives VXLAN packet.
Removes VXLAN header.
Forwards original frame to PC3.
Updates MAC table with remote MAC.

To view VXLAN configuration:

```
Celestica-DS2000# show vxlan interface
VTEP Name              : vtep22
VTEP Source IP         : 10.1.1.1
VTEP MAC               : 11:22:33:44:55:66
EVPN NVO Name          : nvo1
Source Interface       : Loopback0

Celestica-DS2000# show vxlan tunnel
Name                SIP                DIP           source     operstatus
EVPN_10.1.1.1       10.1.1.1          10.1.1.2         EVPN       oper_up

Celestica-DS2000# show vxlan remote mac
Vlan           Mac             RemoteVTEP      VNI       Type

Vlan10   b4:db:91:fd:7c:f6      10.1.1.2       1010     DYNAMIC
Total count :    1

Celestica-DS2000# show vxlan remote vni
Vlan      RemoteVTEP      VNI
Vlan10     10.1.1.2       1010
Vlan20     10.1.1.2       2020
Total count :    2
```

Equivalent SONiC Click CLI commands

```
admin@DS2000:~$sudo show vxlan interface
admin@DS2000:~$sudo show vxlan tunnel
admin@DS2000:~$sudo show vxlan vlanvnimap
```

- **Layer 3 VXLAN – Tenant VRFs and Inter-Subnet Routing:** For deployment scenarios requiring Layer 3 routing between different subnets or isolated routing domains, the Celestica campus fabric utilizes Tenant VRFs (Virtual

Routing and Forwarding instances). Each Department/BU is assigned a separate VRF to isolate its Layer 3 domains. These VRFs are mapped to unique L3 VNIs, allowing tenant-specific routing information to be encapsulated and transported across the VXLAN fabric. This enables inter-subnet communication while maintaining strict logical separation between tenants.

For Layer 3 routing between different subnets or isolated routing domains, Tenant VRFs are mapped to unique L3 VNIs.

```
Syntax:
vrf <vrf-name> vni <1..16777215>
vxlan vni-map <vrf-name> <1..16777215>
```

Example L3 VXLAN mapping (Leaf1)

```
interface Vlan 201
 vrf Vrf1
 ip address 172.16.0.1/24
!
interface Vlan 300
 vrf Vrf1
!
vrf Vrf1
 vni 30000
 ip router-id 1.1.1.1
 exit-vrf
!
interface vxlan vtep1
 source-ip 10.10.10.3
 map vni 20001 vlan 201
 map vni 30000 vlan 300
!
vxlan vni-map Vrf1 30000
```

Equivalent SONiC Click CLI commands

```
$sudo config vrf add Vrf1
$sudo config vlan add 201
$sudo config vlan add 300
$sudo config interface ip add Vlan201 172.16.0.1/24
$sudo config interface vrf  bind Vlan201 Vrf1
$sudo config interface vrf bind Vlan300 Vrf1
$sudo config vrf add_vrf_vni_map Vrf1 20001
$sudo config vxlan evpn_nvo add nvo1 vtep1
$sudo config vxlan add vtep1 10.10.10.3
$sudo config vxlan map add vtep1 201 20001
$sudo config vxlan map add vtep1 300 30000
```

Example L3 VXLAN mapping (Leaf2)

```
interface Vlan 204
 vrf Vrf1
 ip address 172.16.2.1/24
!
```

```
interface Vlan 300
 vrf Vrf1
!
vrf Vrf1
 vni 30000
 ip router-id 1.1.1.1
 exit-vrf
!
interface vxlan vtep1
 source-ip 10.10.10.5
 map vni 20004 vlan 204
 map vni 30000 vlan 300
!
vxlan vni-map Vrf1 30000
```

Equivalent SONiC Click CLI commands

```
$sudo config vrf add Vrf1
$sudo config interface ip add Vlan204 172.16.2.1/24
$sudo config interface vrf bind Vlan204 Vrf1
$sudo config interface vrf bind Vlan300 Vrf1
$sudo config vrf add_vrf_vni_map Vrf1 30000
$sudo config vxlan evpn_nvo add nvo1 vtep1
$sudo config vxlan add vtep1 10.10.10.5
$sudo config vxlan map add vtep1 204  20004
$sudo config vxlan map add vtep1 300  30000
```

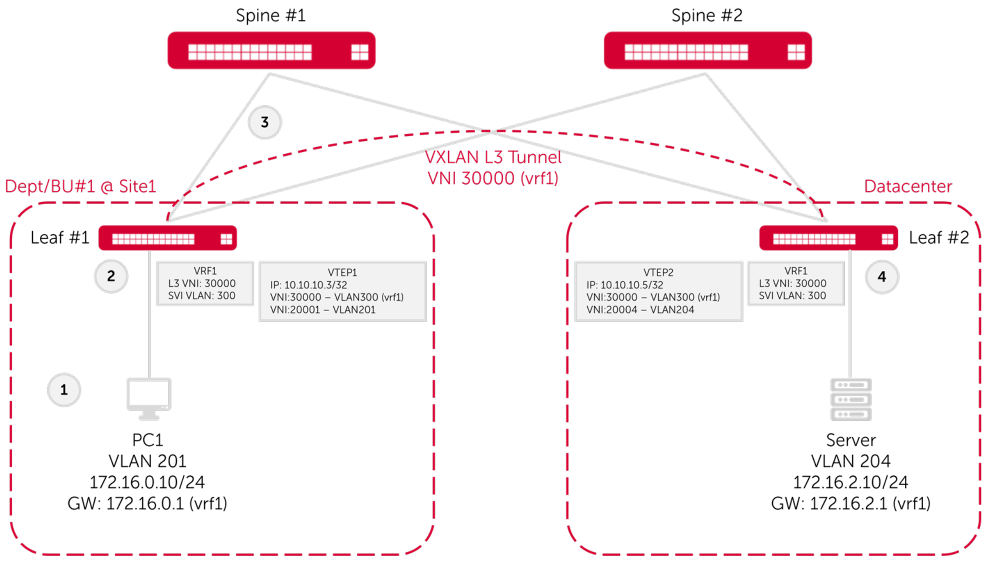

1. PC1 sends a frame to Server (172.16.2.10). Frame forwarded to GW (172.16.0.1). Leaf1 receives the frame on its VLAN int. Associated with vrf1.

2. Leaf1 performs a route lookup for destination within context of vtf1. Leaf1 learned a route 172.16.2.0/24 subnet via Leaf2.

3. Leaf1 takes original IP Packet from PC1. Leaf1 adds VXLAN Header w/ L3 VNI 30000. Leaf1 utilizes BGP (Underlay) to send the packet via Spine1 to Leaf2.

4. Leaf2 receives packet from Spine1 and looks at VXLAN header & sees VNI=30000. Leaf2 knows VNI is associated to vrf1. Leaf2 performs route look for 172.16.2.10. Leaf2 forwards original frame w/ VLAN204.

Verification commands for VXLAN mappings

To show VLAN to VNI mapping information:

```
Celestica-DS3000# show vlan
Flags: d - Dynamic config done by system
       c - User configuration
------------------------------------------------------------
Name            Id     Members             Mode
------------------------------------------------------------
Vlan201         201    Ethernet120 (c)     tagged
------------------------------------------------------------
Vlan300         300
------------------------------------------------------------
Celestica-DS3000#

Celestica-DS3000# show vrf
VRF       Interfaces
-----     ----------
Vrf1      Vlan201
          Vlan300
Celestica-DS3000# show vxlan vlanvnimap
 VLAN           VNI
 ====           ====
 Vlan201        20001
 Vlan300        30000
 Total count : 2
```

- Configuring VXLAN and EVPN Peering: Celestica's access and collapsed core switches, powered by SONiC Distribution by Celestica, function as VTEPs within the campus fabric, enabling the establishment of VXLAN tunnels and EVPN peering.
- Underlay Configuration: The underlay network forms the foundational Layer 3 IP network. Before configuring the overlay, ensure the Layer 3 network is established and loopback IP interfaces (used as VTEP source IPs) are reachable across all devices. This can be achieved using routing protocols such as OSPF or BGP.

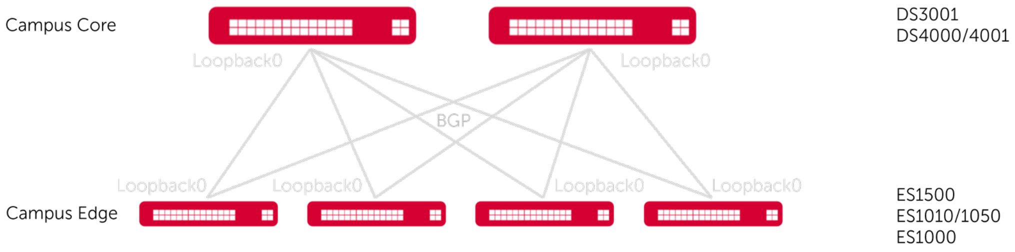

Example underlay loopback configuration

```
interface Loopback 0
 ip address 10.10.10.3/32
```

Equivalent SONiC Click CLI commands

```
$sudo config interface ip add Loopback0 10.10.10.3/32
```

- Overlay Configuration (VXLAN and EVPN): Overlay configurations are performed on devices acting as VTEPs.

Configuring VXLAN interface and source IP

This command configures the VXLAN interface and specifies the VTEP source IP address.

```
Celestica-DS3000(config)# interface vxlan vtep1
Celestica-DS3000(config)# source-ip 100.1.1.1
```

Equivalent SONiC Click CLI commands

```
$sudo config vxlan evpn_nvo add nvo1 vtep1
$sudo config vxlan add vtep1 100.1.1.1
```

Configuring EVPN BGP peering

EVPN uses MP-BGP to exchange MAC and IP address information between VTEPs.
Syntax:

```
router bgp <asn>
 bgp router-id <router-id>
 neighbor <remote-ip> remote-as <remote-asn>
 address-family l2vpn evpn
  neighbor <remote-ip> activate
  advertise-all-vni
```

Example EVPN BGP configuration (Leaf1)

```
router bgp 65001
 bgp router-id 10.10.10.3
 no bgp ebgp-requires-policy
 bgp bestpath as-path multipath-relax
 neighbor Ethernet32 interface remote-as external
 neighbor Ethernet56 interface remote-as external
```

CAMPUS ACCESS NETWORK LAYER: EDGE DEVICE CONNECTIVITY AND SECURITY

The access network layer is where end-user devices, IoT sensors, and wireless access points connect to the campus network. Celestica's access switches provide robust connectivity, power delivery, and granular security at the network edge.

Celestica Access Network Platforms

Celestica's enterprise switching platforms, such as Celestica ES1500 (1U 8/24/48-port 2.5GbE with optional PoE++), Celestica ES1010/ES1050 (1U 48-port 1GbE or 32-port 1GbE plus 16-port 2.5GbE with optional PoE and 25GbE uplinks), and Celestica ES1000 (1U 24/48-port 1GbE with optional PoE and 25GbE uplinks), function as the access layer switches. These switches provide the necessary port density and uplink capacity to connect a diverse range of end devices.

Zero Touch Provisioning (ZTP)

Zero Touch Provisioning (ZTP) is a cornerstone of operational efficiency in the Celestica campus fabric, automating the initial setup of network devices without manual intervention. When a device with SONiC Distribution by Celestica boots for the first time, its ZTP service initiates a DHCP request to obtain the location of a boot file, which then guides the configuration process. This process significantly reduces manual configuration errors and accelerates deployment.

Devices with SONiC Distribution by Celestica can acquire their ZTP JSON configuration file in two primary ways:

1. **Local File:** The ZTP JSON file (ztp_data_local.json) can be placed directly on the switch at the designated path (/host/ztp/) before boot-up.

2. **Remote Server:** The ZTP JSON file can be retrieved from a specified TFTP, HTTP, or HTTPS server. The DHCP server provides the URL of this file using DHCP Option 67.

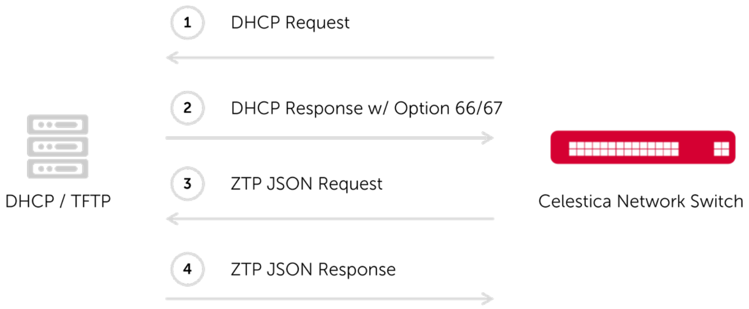

Configuration commands for ZTP

To enable the ZTP feature:

```
Celestica-DS1000 (config)# ztp enable
```

To disable/stop the ZTP feature:

```
Celestica-DS1000 (config)# ztp disable
```

To manually start the ZTP process:

```
Celestica-DS1000(config)# ztp run
```

Equivalent SONiC Click CLI commands

```
$sudo config ztp enable
$sudo config ztp run
```

Viewing ZTP configuration and status

To display the status of the ZTP process:

```
Celestica-DS2000# show ztp status
ZTP Admin Mode      : True
ZTP Service         : Inactive
ZTP Status          : SUCCESS
ZTP Source          : dhcp-opt67 (eth0)
ZTP Runtime         : 08m 12s
ZTP Timestamp       : 2024-10-09 05:47:00 UTC
ZTP Service is not running
01-firmware         : SUCCESS
```

To display a detailed status of the ZTP process:

```
Celestica-DS2000# show ztp status detail
ZTP Admin Mode     : True
ZTP Service        : Inactive
ZTP Status         : SUCCESS
ZTP Source         : dhcp-opt67 (eth0)
ZTP Runtime        : 08m 12s
ZTP Timestamp      : 2024-10-09 05:47:00 UTC
ZTP JSON Version   : 1.0
ZTP Service is not running
01-firmware Status : SUCCESS
Runtime            : 02s
Timestamp          : 2024-10-09 05:47:00 UTC
Ignore Result      : False
```

ZTP is supported on all platforms with SONiC Distribution by Celestica except the ES1000 (ARM models).

Power Over Ethernet (PoE)

Power over Ethernet (PoE) technology streamlines network installations by delivering both electrical power and data over a single Ethernet cable. This eliminates the need for separate power outlets, significantly reducing cabling complexity and overall project costs within an enterprise network. PoE is particularly beneficial for powering distributed end devices such as access points, VoIP phones, ID card scanners, cameras, and various IoT devices. Centralized power management via PoE also enhances system reliability through easier monitoring, control, and proactive maintenance of powered devices. The core components of a PoE system are the Power Sourcing Equipment (PSE), typically a PoE-capable Celestica campus access switch that provides power, and the Powered Device (PD), which receives the power.

PoE relies on several IEEE standards that define power levels and delivery mechanisms, ensuring compatibility and safety across devices. These include IEEE 802.3af (PoE), providing up to 15.4W for low-power devices, IEEE 802.3at (PoE+), offering up to 30W for more demanding applications such as high-end wireless access points, and IEEE 802.3bt (PoE++), which further expands power capabilities with Type 3 (60W) and Type 4 (90W) for devices such as laptops and video conferencing systems. PoE systems also incorporate various power-management modes, such as Static-with-priority (guaranteed power for critical devices) and Dynamic-with-priority (efficient allocation with prioritization), allowing for flexible and efficient power distribution based on network needs.

Supported Platforms: PoE is supported on select Celestica ES1000, ES1010, and ES1500 models.

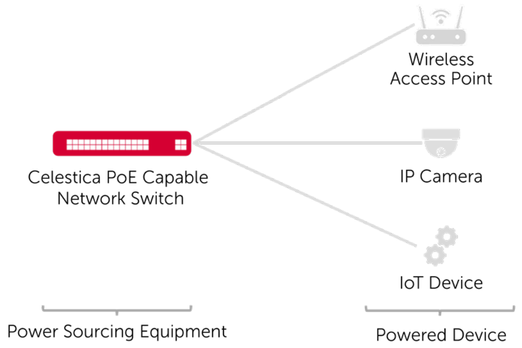

By default, PoE is globally enabled in Celestica campus access switches. If it is disabled, use the following command to enable it globally:

```
Celestica-ES1000-48C(config)# no poe disable
```

To configure power management:

```
Celestica-ES1000-48C(config)# poe power-management dynamic-with-priority
```

To configure guard band:

```
Celestica-ES1000-48C(config)# poe guard-band 55
```

To enable PoE at the interface level:

```
Celestica-ES1010-48CP(conf-if-Ethernet0)# poe enable
```

To configure PoE interface detection type:

```
Celestica-ES1000-48C(conf-if-Ethernet0)# poe detection-type two-point-dot3af
```

To configure PoE interface power threshold:

```
Celestica-ES1010-48CP(conf-if-Ethernet0)# poe power-threshold user-defined 30
```

To configure PoE interface powerup mode:

```
Celestica-ES1010-48CP(conf-if-Ethernet0)# poe power-up-mode dot3bt-type4
```

To configure PoE interface priority:

```
Celestica-ES1010-48CP(conf-if-Ethernet0)# poe priority critical
```

Equivalent SONiC Click CLI commands

```
$sudo config poe global enable
$sudo config poe global power-management dynamic-with-priority
$sudo config poe global guard-band 55
$sudo config poe port Ethernet0 enable
$sudo config poe port Ethernet0 detection-type two-point-dot3af
$sudo config poe port Ethernet0 power-threshold user-defined 30
$sudo config poe port Ethernet0 power-up-mode dot3bt-type
$sudo config poe port Ethernet0 priority critical
```

Viewing PoE configurations

To display PoE global configuration:

```
Celestica-ES1000-48C# show poe global summary
Global PoE Configuration:
  PoE Admin State:            Enabled
  PoE Power Management:       dynamic-with-priority
  PoE Guard Band:             55
```

To display PoE interface configurations:

```
Celestica-ES1010-48CP# show poe port detail Ethernet 0
PoE Interface Configuration:
  Port:                       Ethernet0
  PoE Admin State:            Enabled
  PoE Detection Type:         two-point-dot3af
  PoE Power Threshold:        user-defined 30.0
  PoE Power-Up Mode:          dot3bt-type4
  PoE Power Via MDI:          Enabled
  PoE Priority:             critical
  PoE Reset:                auto 30
```

To display PoE port status:

```
Celestica-ES1010-48CP# show poe port status Ethernet 0
PoE Port Status:
Port: Ethernet0
Power Status: Delivering Power
Voltage: 54.0 V
Current: 0.5 A
Power Consumption: 27.0 W
Max Power: 30.0 W
Class: Class 4
Detection Status: Success
Error Status: None
```

Enhanced LLDP (LLDP-MED)

Enhanced LLDP (LLDP-MED), also known as Link Layer Discovery Protocol – Media Endpoint Discovery, is a crucial feature in campus networks for simplifying the deployment and management of IP-based endpoints, particularly Voice over IP (VoIP) phones. LLDP-MED enables network devices, such as Celestica switches, to automatically discover and configure connected media endpoint devices. This eliminates the need for manual configuration of each phone port for specific voice VLANs, Quality of Service (QoS) parameters, or power requirements. By enabling IP phones to advertise their network policy information (including desired voice VLAN and QoS settings) to the switch, LLDP-MED ensures that voice traffic is automatically placed on the correct VLAN and receives appropriate prioritization, improving voice quality and reducing configuration errors.

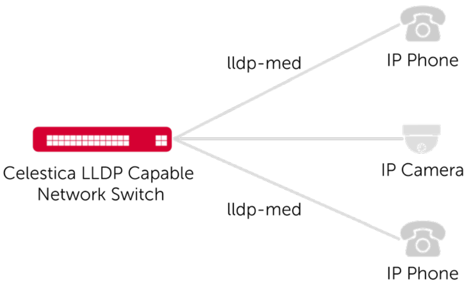

To configure LLDP MED capabilities on an interface:

```
Celestica-DS1000(conf-if-Ethernet0)# lldp med network-policy application voice dscp 3
priority 2 vlan-id 100 tagged
```

Equivalent SONiC Click CLI commands

```
$sudo config lldp med network-policy application modify Ethernet0 voice 3 2 100 tagged
```

To remove the configuration, use no lldp med network-policy.

To view LLDP MED configuration:

```
Celestica-DS1000# show lldp med

Port MED-Application VlanId DSCP Dot1p Tagged

Ethernet0 voice 3094 15 7 false
```

To view detailed LLDP neighbor information, including LLDP-MED policies:

```
admin@DUT2:~$ show lldp neighbor Ethernet11

LLDP neighbors:

Interface: Ethernet11, via: LLDP, RID: 12, Time: 0 day, 00:00:20
Chassis:
  ChassisID: ip 0.0.0.0
  SysName: AVX057AB8
  MgmtIP: 0.0.0.0
  Capability: Bridge, on
  Capability: Tel, off
Port:
  PortID: mac 50:cd:22:05:7a:b8
  PortDescr: Eth11/1
  TTL: 120
  PMD autoneg:
    supported: yes, enabled: yes
  MAU oper type: 1000BaseTFD - Four-pair Category 5 UTP, full duplex mode
LLDP-MED:
  Device Type: Communication Device Endpoint (Class III)
  Capability: Capabilities, yes
  Capability: Policy, yes
  Capability: Inventory, yes
  LLDP-MED Network Policy for: Voice, Defined: yes
    VLAN: 555
    Priority: Voice
    PCP: 5
    DSCP Value: 45
  LLDP-MED Network Policy for: Voice Signaling, Defined: yes
    VLAN: 555
    Priority: Internetwork control
    PCP: 6
    DSCP Value: 34
Inventory:
  Hardware Revision: J139D01B
  Software Revision: 4.0.12.0.6
  Serial Number: 23WZ08500192
  Manufacturer: Avaya
  Model: J139
Unknown TLVs:
  TLV: OUI: 00,40,0D, SubType: 1, Len: 7
    00,00,00,00,00,00,00
  TLV: OUI: 00,40,0D, SubType: 3, Len: 4
    00,00,00,00
  TLV: OUI: 00,40,0D, SubType: 4, Len: 12
    00,00,00,00,00,00,00,00,00,00,00,00
  TLV: OUI: 00,40,0D, SubType: 5, Len: 4
    00,00,00,00
  TLV: OUI: 00,40,0D, SubType: 6, Len: 4
    00,00,00,00
  TLV: OUI: 00,40,0D, SubType: 7, Len: 1
    01
```

Secure Network Access

SONiC Distribution by Celestica integrates robust security features to control network access and protect data.

Edge Security (Port Security)

Edge Security, also known as Port Security, restricts network access on an interface by limiting and identifying the MAC addresses of allowed workstations. If the maximum number of secure MAC addresses is reached and a new, unauthorized MAC address attempts to access the port, a security violation occurs.

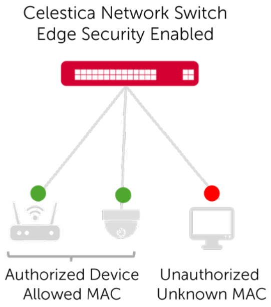

SONiC Distribution by Celestica supports three types of sticky MAC addresses for edge security:

- **ConfiguredSticky:** MAC addresses explicitly configured by the user.
- **DynamicSticky:** Dynamic entries that are converted to sticky entries at a specific point in time.
- **DynamicLearnt:** Dynamic entries learned after sticky MAC is configured, which are tracked under the maximum count limit.

Restrictions for Edge Security

- A secure port cannot be a trunk port.
- Only Layer 2 interfaces with a single untagged VLAN member are allowed.
- VLAN changes on a port are blocked after edge security is enabled.
- A secure port cannot be a destination for SPAN (Switch Port Analyzer) or belong to an EtherChannel port-channel interface.
- Static MAC address configuration is mutually exclusive with a secure port.
- Dot1x/MAB (MAC Authentication Bypass) cannot be enabled on a sticky MAC interface.

Configuration commands for Edge Security

To enable edge security on a given port (the port must belong to an untagged VLAN):

```
Celestica-DS1000(conf-if-Ethernet0)# edge-security enable
```

To set the maximum number of dynamic clients allowed (default is 0):

```
Celestica-DS1000(conf-if-Ethernet0)# edge-security dynamic-learn-max 2
```

To determine the action taken after a violation occurs (restrict (default) or shutdown):

```
Celestica-DS1000(conf-if-Ethernet0)# edge-security violation-action restrict
```

To convert all current dynamic MAC entries on an interface to SecureSticky:

```
Celestica-DS2000# edge-security sticky-mac current-macs Ethernet 24
```

ℹ **Equivalent SONiC Click CLI commands**

```
$sudo config edge-security enable Ethernet0
$sudo config edge-security dynamic-learn-max Ethernet0 2
$sudo config edge-security violation-action restrict Ethernet0
$sudo config edge-security sticky-mac mac Ethernet0 A0:67:B2:78:9A:BC
```

Viewing Edge Security configuration

To show all port details for interfaces with edge security enabled:

```
Celestica-DS1000# show edge-security port

Port          Edge-security Status    Max Dynamic-learn-macs    Violation Action

Ethernet0     Enabled                 0                         restrict
```

To show MAC addresses learned or configured via edge security:

```
Celestica-DS2000# show edge-security mac

PORT        MAC                   VLAN ID    Type              Violation Action
TIMESTAMP

Ethernet24 00:11:01:00:00:01    216        LEARNT-DYNAMIC    None
26-02-2025 03:13:15
Ethernet24 00:11:01:00:00:02    216        LEARNT-DYNAMIC    Restrict-Max-Limit-Hit
26-02-2025 03:13:15
```

AAA (Authentication, Authorization, and Accounting)

SONiC Distribution by Celestica supports AAA services, including RADIUS and TACACS+, for centralized user authentication and authorization, enhancing overall network security.

Configuring RADIUS AAA

To configure a primary RADIUS server:

```
Celestica-DS4001(config)# dot1x radius server primary 192.168.1.1 1812 passkey mgmt
```

Equivalent SONiC Click CLI commands

```
$sudo config dot1x radius server add 192.168.1.1 1812 mgmt
```

Configuring TACACS+ AAA

To configure a TACACS+ server:

```
Celestica-DS4001(config)# tacacs-server host 192.168.1.1 passkey mgmt
```

Equivalent SONiC Click CLI commands

```
$sudo config tacacs add 192.168.1.1
```

MACSec (Media Access Control Security)

MACSec provides an additional layer of encryption and authentication at the data link layer, safeguarding network traffic from eavesdropping and tampering.

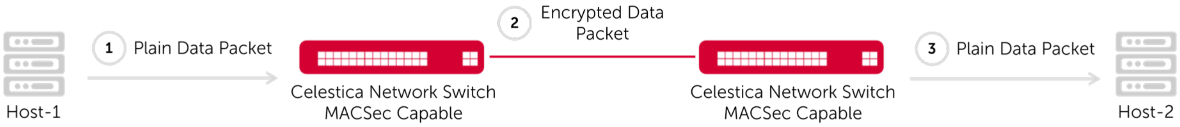

Configuring MACSec

To apply a MACSec profile to an Ethernet interface:

```
Celestica-ES1000-48CP# configure terminal
Celestica-ES1000-48CP(config)# interface Ethernet 40
Celestica-ES1000-48CP(conf-if-Ethernet40)# macsec-profile security_profile
```

Equivalent SONiC Click CLI commands

```
$sudo config macsec add Ethernet40 security_profile
```

Viewing MACSec Configuration

To display MACSec profile configurations:

```
Celestica-ES1000-48C# show macsec profiles
MACsec profile : security_profile
------------  ----------------------------------------------------------------
cipher_suite  GCM-AES-XPN-256
policy        security
primary_ckn   6162636465666768696A6B6C6D6E6F707172737475767778797A303132333435
priority      200
send_sci      False
------------  ----------------------------------------------------------------
```

To display MACSec configuration on a specific interface:

```
Celestica-ES1000-48C# show macsec interface Ethernet 36
MACsec port(Ethernet36)
_____________________ _______________
cipher_suite          GCM-AES-XPN-256
enable                True
enable_encrypt        True
enable_protect        True
enable_replay_protect False
profile               security_profile
replay_window         0
send_sci              False
_____________________ _______________
```

SONIC DISTRIBUTION BY CELESTICA: THE ENHANCED OPEN NOS ADVANTAGE

Modular Architecture: Containers, Redis DB, and SAI

SONiC Distribution by Celestica, a version of Community SONiC with value-added features and functions, is based on Debian Linux, and redefines network infrastructure administration through its modular architecture. Unlike traditional monolithic NOS, the design of SONiC Distribution by Celestica encapsulates major subsystems within Docker containers, each serving specific networking functions. For instance, the FRR routing stack operates within the 'bgp' container for routing protocols, while LLDP handles topology mapping and device discovery. This containerized approach enables independent development, upgrades, and fault isolation for individual network services, aligning with modern cloud-native software development principles resulting in increased stability, faster feature development, and easier customization to meet specific enterprise requirements.

A central component of the SONiC Distribution by Celestica architecture is the Redis database engine, which acts as a critical datastore for all SONiC components. It maintains several databases, including CONFIG_DB (for configuration), APPL_DB (for application-generated state), ASIC_DB (for driving ASIC configuration), STATE_DB (for cross-application dependencies), and COUNTERS_DB (for statistics). This centralized, key-value database enables efficient storage and retrieval of network states and configurations. The Switch Abstraction Interface (SAI) serves as a crucial translator layer, abstracting hardware-specific SDKs into a common API. This hardware abstraction is fundamental to the ability of SONiC Distribution by Celestica to run on diverse underlying ASIC platforms, providing unprecedented flexibility and choice in hardware procurement, which is a cornerstone of Celestica's open hardware design philosophy.

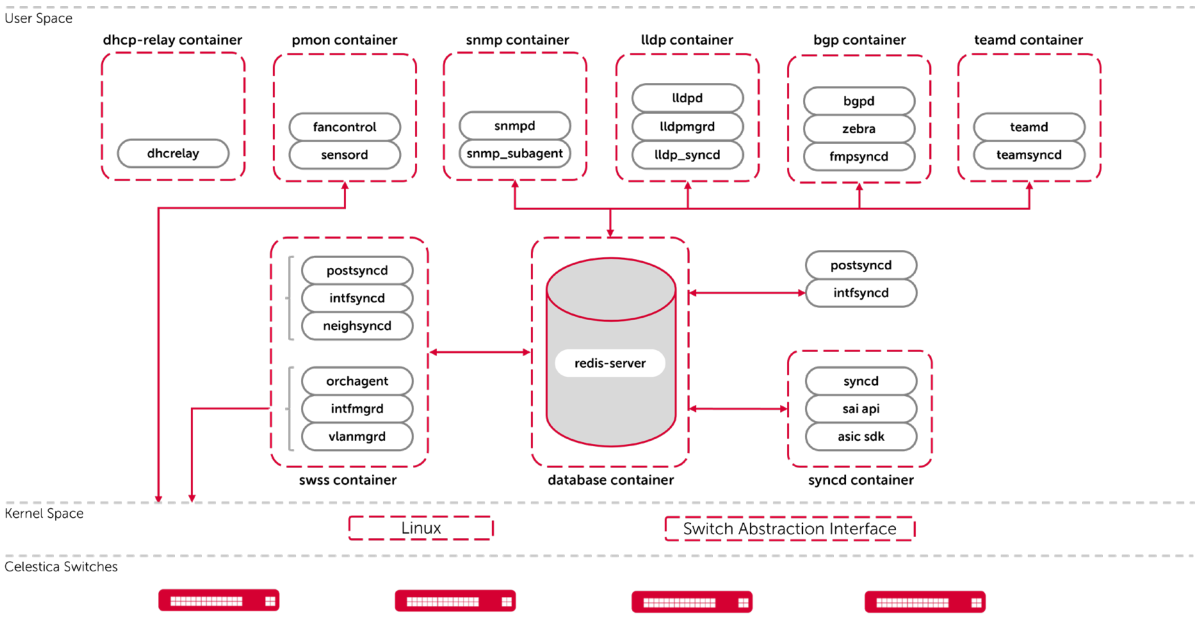

Control Plane Implementation with FRRouting (FRR)

SONiC Distribution by Celestica leverages the open-source FRRouting (FRR) stack, residing within the 'bgp' Docker container, to provide robust routing functionalities, including BGP. When a BGP update arrives, the Linux Kernel delivers it to the 'bgpd' process (part of FRR), which then notifies zebra, the Routing Information Base (RIB) daemon. 'zebra' validates the reachability of the new prefix and generates a route-netlink message to inject this new state into the Linux kernel's forwarding information base (FIB). Subsequently, 'zebra' sends a netlink-route message to 'fpmsyncd', which processes it and pushes the state into the APPL_DB.

The use of FRR as the routing stack provides a standards-compliant, widely adopted, and actively developed control plane for SONiC Distribution by Celestica. This open-source foundation significantly contributes to the cost-effectiveness and flexibility of the Celestica architecture by enabling organizations to benefit from community-driven innovation and avoid proprietary licensing costs. Celestica further enhances this with value-added features and dedicated support services for SONiC Distribution by Celestica, ensuring faster bug fixes and transparent code while reducing reliance on a single vendor's roadmap.

Data Plane Integration and Hardware Abstraction

The seamless integration of the control plane with the underlying hardware for data plane operations is achieved through a sophisticated process within SONiC Distribution by Celestica. The 'orchagent,' an essential component of the Switch State Services (SWSS) container, plays a pivotal role by translating configurations from the APPL_DB into instructions compatible with the SAI API, and then publishing these instructions into the ASIC_DB. The 'syncd' container is then notified of changes in the ASIC_DB. 'syncd,' which links with the ASIC SDK library, programs the new route and state directly into the network silicon (ASIC) via the SAI API.

This hardware abstraction, facilitated by SAI, is a key differentiator for SONiC Distribution by Celestica. It allows Celestica to offer open hardware designs, ensuring that the SONiC software can operate across different underlying ASIC platforms. This capability provides customers with unprecedented flexibility and choice in hardware procurement, enabling them to select devices based on performance and cost rather than being locked into a specific NOS. This directly underpins the "decoupling hardware from software" benefit highlighted in the Celestica architecture.

Management Plane: Configuration and Operational Data Flow

The SONiC Distribution by Celestica management plane offers diverse methods for configuration and operational data flow, indicative of its design for NetDevOps. Configuration can be managed through the Python Click-based SONiC CLI, which directly modifies the CONFIG_DB in Redis. Users can also define configurations using SONiC's JSON schema and push them directly to the CONFIG_DB using utilities like sonic-cfggen or the config load command. For Day 0 or full configuration replacement, a complete config_db.json file can be copied and a reload triggered. SONiC Distribution by Celestica also supports "patching" the running configuration with ADD, REPLACE, and REMOVE operations using config apply-patch.

While the FRR routing stack has its native vtysh shell for configuration, SONiC Distribution by Celestica allows routing configurations to persist either within FRR's domain or within /etc/sonic/config_db.json. Beyond CLI and JSON, SONiC Distribution by Celestica supports programmatic APIs, including gNMI (gRPC-based) and HTTP-based REST APIs for configuring network state and streaming telemetry data, leveraging native YANG models for the configuration DB. This multi-faceted approach to configuration, coupled with the central CONFIG_DB in Redis, enables infrastructure-as-code principles, facilitating version control of configurations and seamless integration with orchestration tools, thereby transforming network operations from traditional CLI-centric management to a more software-driven, programmatic approach.

DAY-TO-DAY OPERATIONS: AUTOMATION, MONITORING, AND TROUBLESHOOTING

Celestica's campus architecture, built on SONiC capabilities, natively supports industry-standard NetDevOps practices and tools, streamlining day-to-day operations.

Automation and NetDevOps with SONiC Distribution by Celestica

The Celestica campus architecture, leveraging SONiC Distribution by Celestica, is designed to enhance network automation and NetDevOps practices. Zero Touch Provisioning (ZTP) plays a crucial role in streamlining deployment by automating device setup without manual configuration. While ZTP significantly reduces manual configuration errors and accelerates deployment times, it is important to note that ZTP is often not enabled by default in community SONiC images, requiring either a custom image build or explicit enablement post-boot. This necessitates a higher level of in-house expertise or a robust NetDevOps pipeline for image management.

The Linux-native infrastructure of SONiC Distribution by Celestica makes it highly amenable to automation using widely adopted IT tools, such as Ansible and Python.

Network Visibility and Telemetry

Comprehensive network visibility and troubleshooting capabilities are essential for maintaining optimal campus network performance. The Celestica architecture supports advanced telemetry through sFlow, an industry-standard sampling technology that provides real-time insights into network traffic. sFlow agents embedded in SONiC Distribution by Celestica devices sample and collect data on packets, byte counts, traffic patterns, and flow information, which is then encapsulated in sFlow datagrams and sent to a designated collector.

CONCLUSION: FUTURE-PROOF YOUR CAMPUS NETWORKS WITH SCALABLE, OPEN, AND AI-READY SOLUTIONS FROM CELESTICA

The evolving digital landscape demands a campus network that is not just reliable, but inherently scalable, open, secure, and future-proof. Traditional, proprietary network infrastructures are struggling to support the exponential growth of IoT, cloud services, and the data-intensive applications that will form the foundation of next-generation, AI-driven operations.

Celestica's wired campus architecture, anchored by Celestica's suite of Data Center or Enterprise Access switches, provides a proven, disaggregated solution to meet these demands:

Future-Proof Scalability and Agility

- The Celestica family of Data Center switches (such as DS3001, DS4000/DS4001, DS4101) forms a high-performance core built on VXLAN EVPN. This virtualized architecture surpasses the limitations of traditional VLANs, providing over 16 million network segments (VNIs) for massive and flexible scalability.
- The use of Multi-Chassis Link Aggregation (MCLAG) ensures device and link-level redundancy, building a highly resilient, active-active foundation capable of seamless growth and minimizing costly rip-and-replace cycles.

Security at Network Edge Infrastructure

- The Celestica family of Enterprise Access switches, such as the ES1500, ES1010, and ES1000, is the intelligent access layer for the modern campus, delivering robust Secure Edge capabilities to protect the network's most vulnerable point.
- Security controls include Edge Security (Port Security) to restrict and authenticate connected devices, as well as support for AAA (Authentication, Authorization, and Accounting) via protocols such as RADIUS and TACACS+ for centralized access control (Port-based Network Access Control/PNAC).
- Furthermore, the Celestica family of Enterprise Access switches provides Power over Ethernet (PoE/PoE+/PoE++) to intelligently power and manage a dense ecosystem of IoT and media endpoints, ensuring automated traffic prioritization via LLDP-MED for superior Quality of Service (QoS) and lower manual configuration errors.
- Layer 2 security features, such as MACSec, provide encryption and authentication, safeguarding the data integrity necessary for critical network operations.

Open Networking Advantage

- Celestica's fundamental commitment to open networking with SONiC Distribution by Celestica decouples hardware from software. This eliminates vendor lock-in, reduces Total Cost of Ownership (TCO), and empowers organizations to select best-of-breed open hardware designs.
- Operational expenses are optimized through automation features, such as Zero Touch Provisioning (ZTP) and the use of open-source NetDevOps tools, fostering innovation and cost efficiency.

By embracing Celestica wired campus networks, organizations are not just modernizing their network; they are investing in a secure, transparent, and high-performance open foundation ready to support the next wave of digital transformation and intelligent, AI-driven solutions.

Celestica™

North America: (Toll Free) +1 888 899 9998

contactus@celestica.com | celestica.com

in Celestica | @Celestica_Inc

The information in this document is subject to change without notice. Certain features may evolve by working with the SONiC community. Celestica, Inc. assumes no responsibility for any errors that may appear in this document.