

Product Brief

Key Highlights

- Open networking with Standard Linux and ONIE support
- Zero Touch Provisioning (ZTP) support
- IPv4/IPv6 switching and routing with eBGP and iBGP
- Container-based architecture
- Programmability and visibility with gNMI/REST/Openconfig
- Advanced silicon driven telemetry such as Buffer Statistics Tracking and Inband Flow Analyzer
- SONiC Unified Manageability Framework
- Platform Development Kit (PDDF and PDE)
- Automation support: ZTP, Chef, and Ansible Playbooks
- L3 protocols such as iBGP, eBGP, VRRP, BFD, and VRF-Lite
- Network Address Translation (NAT)
- Full support for VXLAN and Logical VTEP (LVTEP) using BGP-EVPN control plane
- Multi-chassis LAG (MLAG)
- Multicast features (IGMP snooping, PIM-SSM)
- Extending SONiC to Campus use case
- Data Center Interconnect

Enterprise SONiC by Broadcom[®]

Overview

Enterprise SONiC by Broadcom[®] is an open source network operating system based on Linux that runs on merchant silicon-based platforms. The open source SONiC project is available at GitHub (<https://github.com/Azure/SONiC/wiki>).

Enterprise SONiC is in production today at multiple web-scale companies for Data Center fabric deployments and has a thriving developer community and vendor ecosystem. The underlying architecture of SONiC is described at GitHub (<https://github.com/Azure/SONiC/wiki/Architecture>).

Enterprise SONiC is a commercial offering based on open source SONiC with feature enrichment and hardening that is targeted at Data Center leaf, spine, and super-spine deployments. Enterprise SONiC supports ODM and OEM platforms based on the StrataXGS[®] family of silicon from Broadcom.

Enterprise SONiC offers benefits such as cloud performance, simplicity based on industry leading merchant-silicon and standards based IP-Clos architecture. It also provides agility driven by a Unified Manageability Framework with programmatic APIs and an extensible, container-based architecture. Its open source foundation and standardized ecosystem provide strong economic benefits for a Data Center fabric solution.

Enterprise SONiC 4.0.0 General Availability (GA) uses Broadcom SAI Adapter version 4.3.2.1. The Broadcom SDK version being used for this release is 6.5.21.

The Enterprise SONiC 4.0.0 packages include the following:

- **Cloud_Base Package** – Includes features and functionality (such as, eBGP, ZTP, programmatic API, QoS, ACL, and security features (such as TACACS+) needed for DC Fabric underlay leaf, spine, super-spine use cases, and base telemetry features (such as Thresholds and Snapshots [BST])
- **Cloud_Advanced Package** – Includes all features in Cloud_Base Inband Flow Analyzer (IFA, version 2.0), Tail Stamping and Drop Monitor Linux PTP (KNETSync)
- **Enterprise_Base Package** – Includes underlay features (such as, eBGP, ZTP, programmatic API, QoS, ACL, and so on) needed for DC Fabric underlay Overlay features (BGP EVPN, VxLAN, and so on) for DC fabric overlay use cases (for PINs leaf, spine, super-spine), and Enterprise Features (such as RPVST+, IP Multicast, and so on), and base telemetry features (such as Thresholds and Snapshots [BST])

- Enterprise_Advanced Package - Includes all features in Enterprise_Base plus Inband Flow Analyzer (IFA, version 2.0), Tail Stamping and Drop Monitor Linux PTP (KNETSynC) Broadcom Debug Tool
- Campus Package - Includes all L2/L3 features plus POE, POE+ and POE-bt, features related to Port Access Control (802.1x, MAB, Dynamic ACLs, RADIUS assigned VLANs, RADIUS support), LLDP-MED, Port Security, Digital Optical Monitoring and Time Domain Reflectometry, and EVPN VXLAN Campus package is supported on campus-specific platforms only

The Campus package does not include Telemetry Features, gNMI, PTP, DCI, NAT, and Multi Site Data Center Interconnect (DCI) ACL-based replication (Packet-Too-Big [PTB] ICMPv6 Hashing enhancement).

All the packages are supported on the platforms listed below. Enterprise SONiC 4.0.0 Packages are available as a 1-year, 3-year or 5-year subscription license.

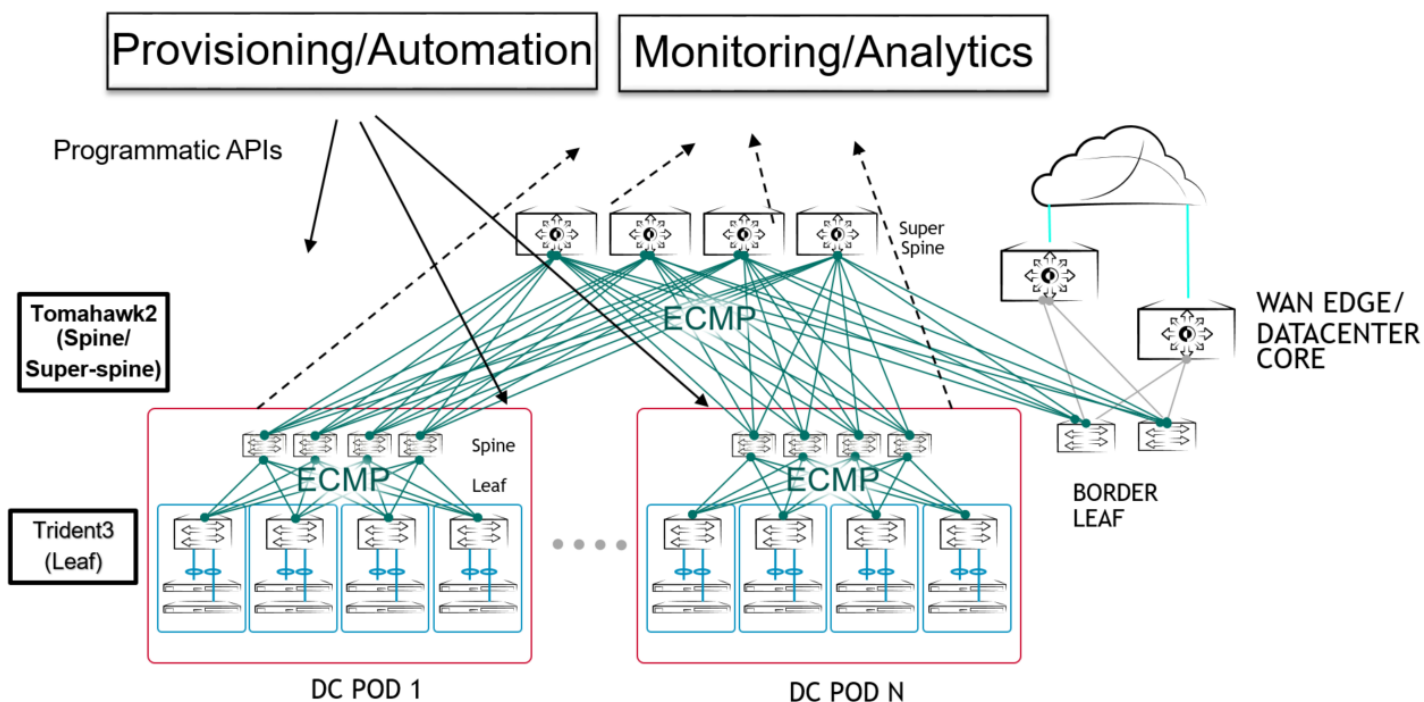
Customer Use Cases

Data Center L3 CLOS (Underlay) Use Case

Enterprise SONiC Cloud editions are targeted for Data Center Fabric deployments (Public, Private, and Edge compute). Enterprise SONiC can be deployed at various Places-In-Network (PIN) - ToR, Leaf, Spine, Super spine, and Border Leaf PINs.

Enterprise SONiC-based data center fabrics can be deployed in an underlay use case for web-scale data center architectures, or for data center PODs in enterprises or service providers for select workloads, such as Hadoop, that require an underlay network. Enterprise SONiC can also be used in enterprises and service providers as an underlay data center fabric for VMware-based PODs deploying VMware ESX, NSX, vSAN, and other VMware solutions.

Figure 1: Data Center L3 CLOS (Underlay) Use Case



Data Center L3 CLOS Overlay Use Case (with VXLAN and BGP-EVPN)

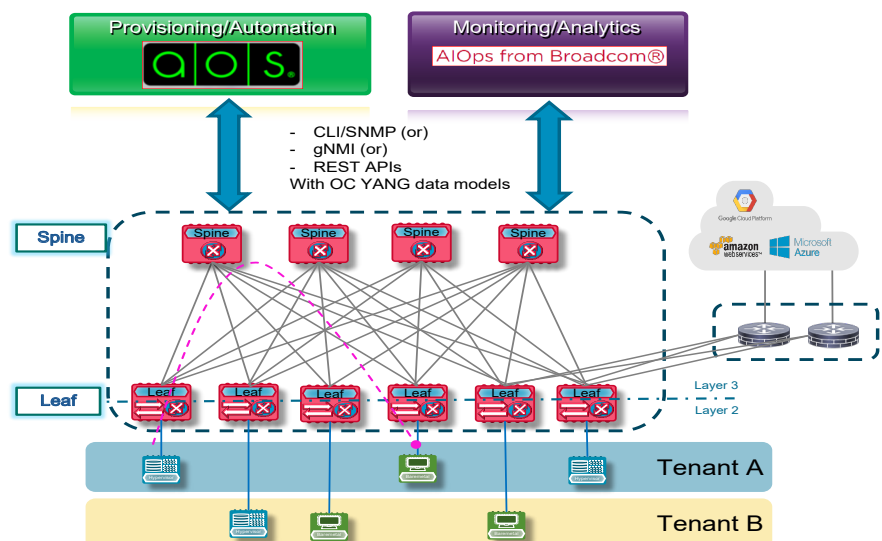
Starting with the 3.4.0 release, Enterprise SONiC can also be deployed in enterprises or service providers for select workloads, such as Hadoop, that require an overlay to support multi-tenancy.

Using an overlay architecture in the data center allows end users (network administrators) to place endpoints (servers or virtual machines) anywhere in the network and remain connected to the same logical Layer 2 network, enabling the virtual topology to be decoupled from the physical topology. This decoupling allows the data center network to be programmatically provisioned at a per-tenant level.

Overlay networking generally supports both Layer 2 and Layer 3 transport between servers or VMs. It also supports a much larger scale. SONiC overlay networks use a control-plane protocol (BGP-EVPN) to facilitate learning and sharing of endpoint information, and use VXLAN tunneling protocol to create the data plane for the overlay layer.

Figure 2: Data Center L3 CLOS Overlay Use Case (with VXLAN and BGP-EVN)

Multi-tenancy with EVPN VXLAN + Provisioning/Monitoring



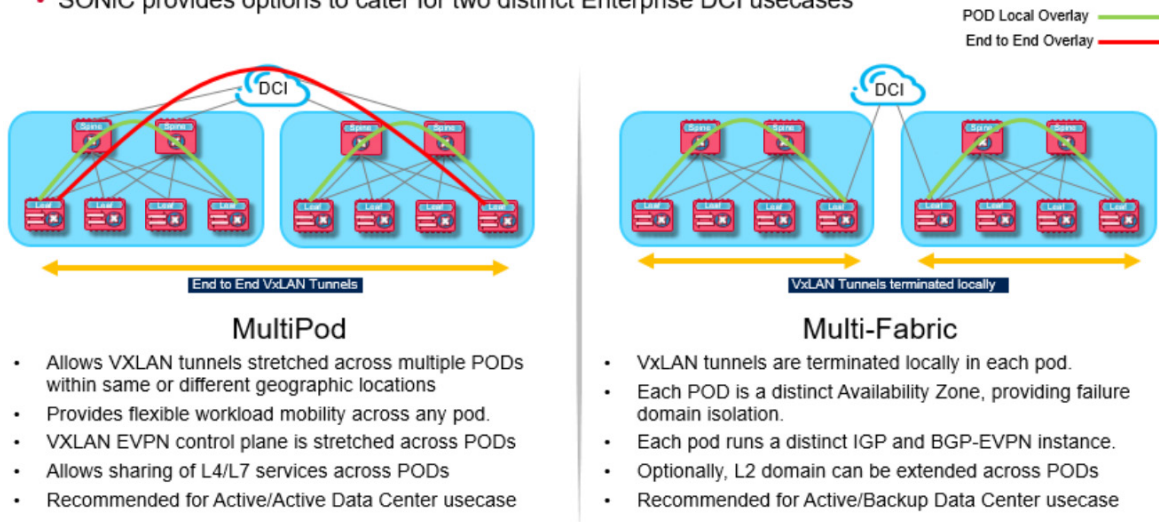
Key Integrations

- Broadcom provides northbound APIs (gNMI or REST APIs with Open Config YANG data models) that integrate with
 - Provisioning, Automation tools such as **Apstra**
 - Monitoring, Analytics tools such as **AIOps from Broadcom (formerly CA technologies)**

Figure 3: Campus Use Case with Access/Aggregation Layers

Use Case : Data Center Interconnect (DCI)

- SONiC provides options to cater for two distinct Enterprise DCI usecases

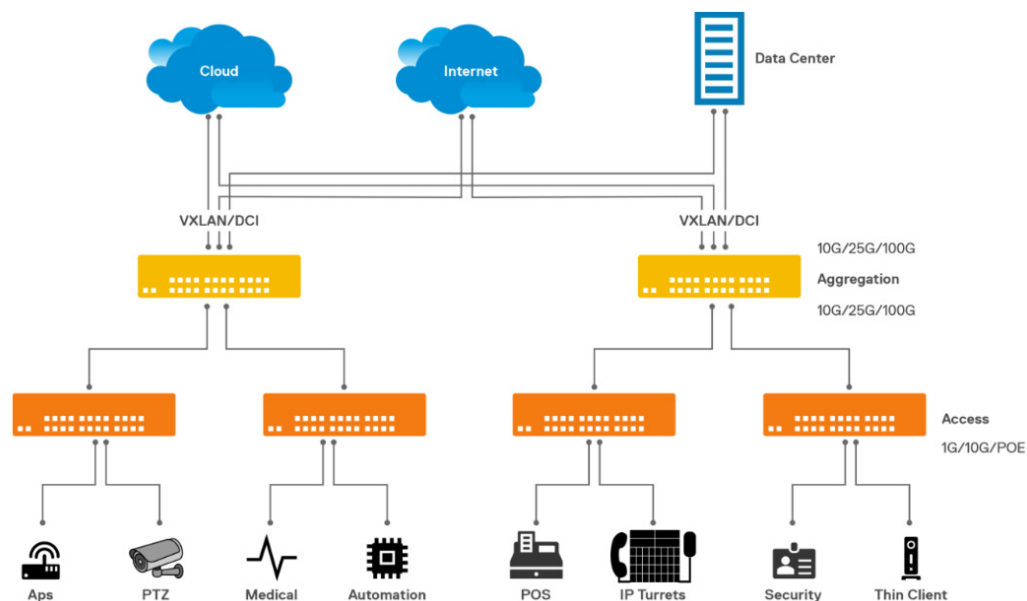


Campus Use case with Access and Aggregation Layers that can be Connected to Existing DCs

Enterprise SONiC can serve as the unified Network Operating System (NOS) that you use to connect edge devices (Campus devices such as POS, thin clients, and security cameras as shown in the following figure) and a data center. The Campus bundle extends a DC fabric to remote locations using the same DC NOS. Additionally, the Campus package works with the traditional three-tier architecture of Access-Aggregation-Core.

- A two-layer fabric is implemented where there is an aggregation and an access layer. Customers can use VxLANs to stretch the fabric.
- The CLOS network spine-leaf architecture allows for future scaling. Each leaf-layer access switch is connected to each spine-layer aggregation switch in a full-mesh topology.
- In the aggregation layer, VxLAN EVPN supports multi-tenancy and multi-site data center interconnection (DCI).
- You can leverage the VxLAN EVPN automation and management tools in the data center to configure and maintain edge switches.

Figure 4: Spine-Leaf Architecture



Full List of Features

Ethernet

- IEEE 802.3ae 10G Ethernet
- IEEE 802.1Q VLAN Tagging
- IEEE 802.1p Class-of-Service Prioritization and Tagging
- IEEE 802.1v VLAN Classification by Port
- IEEE 802.1AB Link Layer Discovery Protocol (LLDP)
- IEEE 802.3x Flow Control (Pause Frames)
- IEEE 802.3z 1000BASE-X

Layer 2 and Layer 3 Protocols

- Border Gateway Protocol (v4, v6)
- iBGP
- eBGP
- FRR BGP for EVPN (with MLAG)
- IGMP Snooping (v1, v2, v3) (with MLAG)
- Route Reflector
- Layer 3 ACLs
- Static routes
- IPv4 ACL
- Bidirectional Forwarding Detection (BFD)
 - BFD IS-CLIs
 - BFD with VRF
- 64-Way ECMP
- VRF Lite
- VRF Support for SSH.in
- Management VRF Hardening
- VRRP (IPv4)
- VRRPv3, VRRP/VRRPv3 over VRF IPv4/IPv6 Dual Stack
- ICMPv6 Route-Advertisement
- Route Policies
- BGP-Allow AS

- BGP Peer Auto-shutdown
- Dynamic BGP Neighbor
- IPv6 routing
- Wirespeed routing for IPv4 and IPv6

Quality of Service

- Class of Service (CoS) IEEE 802.1p
- DSCP to Traffic Class Mapping
- Random Early Discard
- Scheduling: Strict Priority (SP), Deficit Weighted Round-Robin (DWRR)
- Priority Flow Control (PFC)
- Explicit Congestion Notification (ECN)
- DSCP Marking Preservation for VxLAN
- QoS Map Support for Remarking and SVI
- BUM Storm Control
- Port and Priority Shaping

Manageability, Automation, and Monitoring

- Zero-Touch Provisioning (ZTP)
- IPv4/IPv6 management
- Industry-standard Command Line Interface (CLI)
- SSH/SSHv2
- Link Layer Discovery Protocol (LLDP) IEEE 802.1AB
- MIB II RFC 1213
- Syslog
- SNMP v1, SNMP v2C, and SNMP v3
- SNMP CLI
- SNMP Trap Infra and MIBs Support
- Out-of-band management
- Network Time Protocol (NTP)

- Management Access Control Lists (ACLs)
- ACL DSCP map/remarking
- ACL rate limiting
- DHCP Relay Enhancements
- gNMI Subscription Support for Limited YANG Paths (OnChange, Interval, Once, Poll, Target defined)
- Bulking support in both REST(YANG patch) and gNMI
- Query parameter for REST and filtering support for gNMI
- Scalar encoding support for gNMI
- Industry Standard CLI (IS-CLI)
- REST and gNMI Interfaces through OpenConfig YANG (OC-YANG)
- Role-Based Access Control (RBAC)
- SNMP Configuration Traps and OIDs
- Configuration Services – Chef for EVPN

Infrastructure

- Infrastructure level changes have been made extensively to make sure SONiC can run in Campus platforms (with lower memory 8G)
- POE, POE+, and POE-bt
- Port Security
- Aruba Clearpass, Cisco ISE, FreeRadius Interoperability
- LLDP-MED
- Digital Optical Monitoring (DOM)
- Time Domain Reflectometry (TDR)
- EVPN VXLAN (scale will be different for VXLAN on Campus Platforms, pls refer scalability matrix)

Port Access Control

- 802.1X
- MAC Authentication Bypass (MAB)
- RADIUS Support
 - RFC 2865 – RADIUS Client
 - RFC 2868 – RADIUS Attributes for Tunnel Protocol Support
 - RFC 3579 – RADIUS Support for EAP
- Authentication Tiering
- Dynamic ACL
 - Named ACLs
 - Per-session ACLs
 - Filter-Id
- Guest VLAN
- Unauth VLAN
- Open VLAN
- Monitor Mode

ACL and Flow-Based Services

- ACL-based Packet Replication
- ACL Consistency Checker
- PBR Enhancements for Service Chaining
- ACL-based CoPP
- ACL DSCP Map/Remarking
- ACL Rate Limiting
- Control Plane ACL
- Policy-based Routing (IPv4 and IPv6)

System and Platform Infrastructure

- Link Statistics Enhancements
- Link-Down Reason Codes
- Link Flap Error-Disable
- Forwarding Plane Drop Counters
- Time zone Command Support
- Broadcom Debug Tool
- Memory Histogram
- System Ready for Services and Applications

- Secure Boot Process and Reference Implementation
- Syslog High Threshold notifications and clear for CPU/Temperature
- Per Platform CoPP
- Dynamic Port Breakout
- DOM Information Display
- Locator LED Support (Beacon)
- CMIS 4.0 Optics Support
- Hardware Watchdog
- 1G/10G BASE-T copper BASE-T Support (On select platforms)
- Control Plane Policing (CoPP)
- Transceiver Parameter Tuning
- Third-Party Container Management
- PDDF and PDK Framework
- Interface Aliasing (IS-standard Interface Naming)
- Kdump Support
- Maintenance Mode
 - LACP Graceful Shut
 - BGP Graceful Shut
 - OSPFv2 Maximum Metrics
- Multi-Instance Redis DB
- Hardware Resource Allocation and Reservation
- Zero Touch Provisioning (ZTP)

Warmboot

- For more information, contact the Enterprise SONiC team at Broadcom

DHCP Relay

- DHCP Relay over VxLAN Overlay Interfaces
- DHCP Relay Source Interface Selection (such as loopback)
- DHCP Relay over IPv6 Link-Local Interfaces with RFC5549 Routes
- DHCP Relay Hop Count Configuration

- DHCP Relay Over IPv4 Unnumbered Interfaces
- DHCP Relay Option 82, Sub Option 151 VRF Name/ID Option
- DHCP Relay Option 82, Sub Option 5 Link-Selection Option RFC3527
- Support for Circuit-Id Formats

Security

- RADIUS and TACACS
- RADIUS/TACACS Password Obfuscation
- NTP Server and NTP Authentication
- LDAP Integration

Multicast

- L3 Multicast with PIM operates on L3 interfaces only
- IGMP
- IGMP Snooping (v1, v2, v3) (with MLAG)
- IPv4 PIM-SSM Support

Debug and Service

- In-memory Debug Logging
- Audit Logging and Syslogs
- Command to Return Interfaces to the Default Configuration
- Port Mirroring on Port Channel and VLAN

Scalability

- L3 VLAN Scale to 4K (for TD3-X7 based Platforms)
- Host Table Resource Reservation for Local Hosts

Telemetry and Instrumentation

- sFlow on Management VRF
- Inband Flow Analyzer (IFA) 2.0
- Drop Monitor
- Tail Stamping
- BST – Watermarks, Thresholds, and Snapshots
- Linux PTP (KNETSync)

Other

- BUM/Storm Control
- Configuration Services – Chef for EVPN
- Drop Monitor
- Inband Flow Analyser (IFA, version 1.1)
- IP Fabric over IPv6 underlay RFC5549
- IP Helper
- IGMP
- IPv4 PIM-SSM Support
- IPv4 Unnumbered interfaces
- Instrumentation (Snapshots and Thresholds)
- L2 and L3 Multi-Chassis LAG (MCLAG)
- L2 and EP (LVTEP) support using BGP-EVPN control plane
- Link Aggregation Optimizations
- Link Aggregation – Static LAG support
- Linux PTP (KNETSync)
- Management Framework Enhancements
- OSPFv2
- PDDF BMC Support
- PDDF Support with SONiC 2.0 Platform APIs
- Port Mirroring on Port Channel and VLAN
- Port and Priority Shaping
- RADIUS
- REST and gNMI interfaces
- RPVST+
- Role-based Access Control (RBAC)
- Tail Time Stamping
- Transceiver parameter tuning
- UDLD
- Multi Site Data Center Interconnect (DCI)
- RIB/FIB Consistency Checker
- Next Hop Group (NHG) Support
- RIF Counters for L3 Interfaces
- BFD profile
- 4K L3 VLAN Interface Scale for SAG and Unique-IP Cases
- Nexthop resolution using default route
- MC-LAG Peer Gateway
- VLAN Auto-state
- Interface Hold-Down
- LACP Graceful Shutdown
- Uplink Tracking
- L2 and L3 MC-LAG
- L2 and L3 VxLAN
- L2 and Le LVTEP
- LACP Fast Rate and LACP Fallback
- Static LAG
- LLDP
- UDLD
- MC-LAG Fallback
- MC-LAG Graceful Shutdown
- xSTP over MC-LAG
- PVST and RPVST+
- PVST and RPVST+ over MC-LAG
- Dynamic BGP Neighbor
- Support for 4K L3 VLAN Interfaces
- IP Fabric over IPv6 underlay RFC5549
- IP Helper
- Routing Subinterface (on TD3-X5, TD3-X7, TD4-X11 platforms)
- Advertise PIP for both ACT-ACT and ACT-STBY on the Same Leaf Pair
- Route Leaking across VRFs including Management VRF
- BGP Docker Warm Restart
- Avoid Netlink for Handling IPv6 Link-Local Address
- BFD Optimizations to Support 5x100msec Aggressive Timers in SW
- IP SLA (ICMP and TCP tracker)
- IPv4 Unnumbered Interfaces
- RPVST+ Scaling to 3500 VLAN Ports
- Symmetric Hashing
- VxLAN over SVI Interface
- BGP for EVPN (with MLAG)
- BFD IS-CLIs
- BFD with VRF
- VRF support for syslog and SSH.in
- VRRPv3, VRRP/VRRPv3 over VRF
- Management VRF Hardening
- NAT
- OSPFv2

For more information, please contact the Enterprise SONiC team at Broadcom.