Switching & Availability

Aruba Certified Switching Professional

KNOWLEDGE CHEATSHEET

RESILIENT SWITCHING

- **VSF** — Logical stack
- **VSX** — Independent peers
- **ISL** — Sync and data
- **Keepalive** — Peer liveness

01 Switching architecture and high availability

Platform, stacking and power

- **VSF** combines supported switches into one logical stack with a shared management and control model. **VSX** pairs two switches while retaining independent control planes and configurations.
- For this campus platform family, **CX 6300** uses VSF; **CX 6400** is a modular access/aggregation platform with VSX. **8325** targets data center switching; **8400** targets modular core/aggregation deployments. Match PoE and resilience requirements to the platform.
- **Always-on PoE** can preserve power to attached phones and APs during supported switch reboots. Power continuity does not guarantee uninterrupted packet forwarding; redundant links and upgrade sequencing address data connectivity.
- **LLDP/CDP** exchanges device and power information. On authenticated ports, check discovery-protocol handling and PoE negotiation as well as the chassis power budget; allocated power is not the same as actual consumption.

VSX links and failure handling

- The **inter-switch link (ISL)** carries peer synchronization and traffic that must cross between VSX members. Use a resilient, high-bandwidth LAG with consistent VLAN and MTU configuration.
- The **keepalive** independently verifies peer liveness. Keep its path separate from the ISL, preferably through a dedicated routed connection; it is not a backup user-data link.
- Configure complementary **primary/secondary roles** and a stable shared **system MAC**. Preserving that MAC during hardware replacement avoids changing the pair's LACP identity.
- If the **ISL fails but keepalive remains up**, the secondary suppresses its VSX LAGs to prevent dual-active forwarding. Reachability then depends on remaining primary paths and whether endpoints are dual-homed or single-attached.
- Use **show vsx status** and LACP state to check peer roles, ISL health, synchronization and member consistency before adjusting timers. Retain default timers unless a validated design requires changes.

Switching & Availability / Unicast Routing

Aruba Certified Switching Professional

KNOWLEDGE CHEATSHEET

ROUTING DECISIONS

- **OSPF** — Areas and cost
- **BGP** — Attributes and policy
- **VRF** — Separate tables
- **PBR** — Selected flows

Multichassis forwarding and gateways

- A **VSX LAG** is a Layer 2 multichassis LAG. Configure matching LAG identity, VLANs and LACP behavior on the peers; LACP active mode is valid. Local traffic normally uses local member links when available.
- **Active gateway** gives both peers the same virtual gateway IP and MAC on a client SVI. Each peer retains its own SVI IP; either can route traffic addressed to the shared gateway.
- **Active forwarding** handles routed traffic arriving at one peer for the other peer's router MAC. It serves a different purpose from the shared client default gateway; use explicit routed adjacencies for core transit designs.
- **VSX configuration synchronization** is selective and flows from primary to secondary for enabled features. It does not merge every local change; unsynchronized objects can differ between peers.
- For upgrades, verify **dual-homing**, LACP, routing neighbors and gateway reachability before taking one peer down. Independent control planes allow staged maintenance, but every endpoint still needs a surviving forwarding path.

02 Unicast routing and path control

OSPF areas and advertisements

- **Area 0** is the OSPF backbone. An **ABR** connects areas and advertises inter-area reachability; an **ASBR** redistributes routes from outside OSPF. Ordinary nonbackbone areas connect through the backbone.
- **Type 1** LSAs describe a router within its area; **Type 2** describe a broadcast transit network and are originated by its DR. **Type 3** carry inter-area prefixes, **Type 4** identify an external ASBR, and **Type 5** carry external routes.
- A **stub area** excludes Type 5 externals and uses an ABR default route. A **totally stubby area** also suppresses ordinary Type 3 summaries, retaining its internal topology and default route.
- An **NSSA** permits local redistribution as Type 7 LSAs; an eligible ABR translates them into Type 5 outside the NSSA. A **totally NSSA** additionally suppresses ordinary inter-area summaries.
- Configure **area range summarization** at an ABR to replace component inter-area advertisements with an aggregate. LSDB entries and installed IP routes are different counts; summarization does not erase intra-area topology.
- An **OSPF virtual link** can restore logical backbone connectivity through a suitable transit area. Its endpoints are ABRs; it cannot cross a stub area. Prefer a direct backbone connection when the topology permits.

Unicast Routing

Aruba Certified Switching Professional

KNOWLEDGE CHEATSHEET

ROUTING DECISIONS

- **OSPF** — Areas and cost
- **BGP** — Attributes and policy
- **VRF** — Separate tables
- **PBR** — Selected flows

OSPF adjacency and convergence

- For neighbor formation, check **area ID/type, hello/dead timers, authentication, network type and addressing**. Use unique router IDs. An MTU mismatch commonly stalls database exchange in ExStart or Exchange.
- On a broadcast segment, **DROTHER-to-DROTHER 2-Way** can be normal. Full adjacencies form with the DR and BDR; a different OSPF priority alone does not prevent adjacency.
- **BFD** detects forwarding-path failures rapidly and notifies routing protocols. **Graceful restart** preserves forwarding during a supported control-plane restart with neighbor assistance; it does not detect broken paths.
- **E1** external cost includes the seed metric plus internal cost to the ASBR. **E2** primarily compares the external metric, using internal cost when needed to distinguish equal external paths.
- Use **route maps** to control redistribution, seed metrics and external type. At stateful firewall exits, coordinate forward and return routing; metric choice on one side alone cannot guarantee symmetry.

BGP peering and policy

- **eBGP** connects different AS numbers; **iBGP** connects peers in the same AS. Verify IP reachability, TCP 179, local/remote AS values and activation of the peer in the required address family.
- A directly connected eBGP peer normally needs neither **multihop** nor a loopback update source. If loopbacks are used, configure the source, reachability and hop allowance deliberately.
- Early best-path comparisons commonly favor **higher weight**, **higher local preference**, locally originated routes, **shorter AS_PATH**, lower origin type and lower comparable **MED**. The next hop must also be reachable.
- **Weight** affects only the local router. **Local preference** influences outbound choice across the local AS. **MED** suggests a preferred entry to a neighboring AS, usually with lower values preferred; remote policy controls acceptance.
- **AS-path prepending** makes a path less attractive to remote networks. **Route reflectors** reduce the iBGP full-mesh requirement; peer groups simplify repeated configuration without providing route reflection.
- Prevent unintended **transit-AS** service with outbound prefix filters and route maps that advertise only authorized local prefixes. **aggregate-address** creates an aggregate; suppress component advertisements when only the summary should be exported.

Unicast Routing / Multicast

Aruba Certified Switching Professional

KNOWLEDGE CHEATSHEET

MULTICAST DELIVERY

IGMP Receiver membership → **PIM** Distribution tree → **RPF** Incoming-path check → **Forward** Interested receivers

VRFs and policy-based routing

- A **VRF** provides a separate routing table. Attach a routed port or SVI with **vrf attach** in interface context; a Layer 2 VLAN definition alone is not the Layer 3 VRF boundary.
- **Policy-based routing (PBR)** classifies incoming traffic and can direct selected flows toward a next hop or tunnel, such as a guest path to a DMZ. Ordinary destination-based routing handles traffic outside the policy.
- A policy **next hop** overrides normal routing for a matching flow when usable. A **default next hop** is a fallback after normal route lookup; configured alternatives and their order matter.
- If a policy next hop fails, inspect the configured **fallback or drop action** and platform behavior. Do not assume every PBR policy drops traffic or always bypasses the routing table.

03 Multicast membership and forwarding

IGMP, snooping and RPF

- **IGMP** communicates IPv4 receiver membership on a local subnet. **IGMP snooping** observes that membership to restrict Layer 2 forwarding; a querier refreshes membership state.
- Enable **IGMP on receiver-facing routed interfaces** and **PIM on multicast routing paths**. IGMP does not build the router-to-router distribution tree, and a source need not join its own destination group.
- **Reverse path forwarding (RPF)** uses unicast reachability to identify the expected incoming interface toward the source or RP. A missing or incorrect unicast route can break multicast even when receivers have joined.
- Check **group membership, querier presence, router ports and unknown-multicast handling** when traffic floods unexpectedly. IPv4 groups map into a smaller Ethernet multicast address space, so different IP groups can share a destination MAC.

PIM modes and VSX behavior

- **PIM dense mode** initially floods and then prunes unused branches; expired prune state can allow flooding again. It assumes receivers are relatively widespread.
- **PIM sparse mode** explicitly joins distribution trees. A **rendezvous point (RP)** connects sources and receivers through a shared tree; eligible flows can switch to a source-specific shortest-path tree.
- A **bootstrap router (BSR)** distributes candidate-RP information in a dynamic RP design. Static RP configuration is another option. RP reachability and consistent group-to-RP mapping are essential.
- Traditional **VSX PIM-SM** runs independent PIM instances on the peers, with a DR forwarding toward a receiver LAN. A VSX pair is not automatically one shared multicast control plane.
- The optional **PIM active-active** mechanism coordinates DR and proxy-DR forwarding state for faster recovery. Verify that feature explicitly, retain keepalive, and inspect membership and multicast routes on both peers.

Access Security

Aruba Certified Switching Professional

KNOWLEDGE CHEATSHEET

IDENTITY TO POLICY

Authenticate 802.1X or MAC → **Assign** User role → **Forward** Local or tunnel → **Enforce** Switch or gateway

04 Access security and dynamic segmentation

Port authentication and client modes

- **802.1X** uses a supplicant, switch authenticator and AAA server. **MAC authentication** supports devices without supplicants, but a MAC address is spoofable; profiling adds context rather than cryptographic identity.
- Configure the **RADIUS server group**, shared secrets, reachability and port-access authentication. Repeated EAP start without progress calls for checks of supplicant behavior, server assignment and RADIUS responses.
- In **client mode**, individual clients authenticate separately. Increase the **client limit** when a phone and PC share a port; account for the expected endpoints rather than only the physical port count.
- In **device mode**, an authenticated attached device can authorize traffic behind it, useful for supported AP designs. This trusts the attachment device and should be confined to the intended ports.
- **RADIUS server tracking** probes server availability to avoid repeatedly waiting for an unavailable server. Tracking credentials are health-check inputs, separate from downloadable-role credentials and CoA configuration.

Roles, captive portal and CoA

- A **local user role** stores VLAN and policy settings on the switch and can be selected by an external RADIUS server. **Downloadable user roles (DUR)** obtain role definitions from ClearPass over HTTPS after authentication.
- Successful authentication does not prove a successful **role download**. Check TCP 443, DNS, server credentials, certificate trust, clock validity and role syntax; inspect the reported failure stage before choosing a fix.
- **Captive portal** gives guests browser-based authentication after initial network access. Coordinate the initial role, redirect policy, portal/AAA service and access to DNS, DHCP and the portal itself; MAC authentication commonly establishes the initial guest session.
- **RADIUS accounting** reports session start, stop and interim usage. **Change of authorization (CoA)** changes or disconnects an existing session; authorize the server and allow its messages through intervening filters.
- When a server sends a disconnect but access remains, verify **CoA enablement, server identity, shared secret and session matching**. Accounting measurements and enforcement messages are separate parts of a guest quota design.

Access Security / Traffic Policy & QoS

Aruba Certified Switching Professional

KNOWLEDGE CHEATSHEET

IDENTITY TO POLICY

Authenticate 802.1X or MAC → **Assign** User role → **Forward** Local or tunnel → **Enforce** Switch or gateway

User-based tunneling

- **Dynamic segmentation** assigns policy by identity or device role, enabling a **colorless port** design. A role can select local forwarding or **user-based tunneling (UBT)** to an Aruba gateway for centralized firewall enforcement.
- A **UBT zone** identifies the gateway destination and backup configuration. The switch role references a **gateway role** that selects policy at the gateway; that role can align wired and wireless access.
- **GRE (IP protocol 47)** carries tunneled traffic; **PAPI (UDP 8211)** provides control communication. Permit required ICMP reachability as well, and allow underlay MTU room for encapsulation. GRE itself supplies no encryption.
- Distinguish the switch-side **transit/client VLAN** used for tunnel handling from the **user VLAN at the gateway**. The relationship depends on VLAN extend or local VLAN mode; these are not VXLAN identifiers.
- Use switch **UBT statistics** and gateway tunnel counters to locate a failed forwarding stage. Authentication, role assignment, tunnel establishment and gateway policy must all succeed.
- For **voice**, prefer a design that authenticates at the switch and locally forwards voice when appropriate, avoiding unnecessary tunnel latency. Device profiles can apply VLAN and QoS settings from LLDP/CDP or MAC matches without a full AAA deployment.

05 Traffic filtering and quality of service

ACL matching and direction

- An **ACL** is an ordered set of ACEs evaluated by ascending sequence number. The first match decides permit or deny; unmatched traffic meets the implicit deny. Define and apply the ACL at the intended attachment point.
- Match the correct **source, destination, protocol and service port**. A client request normally uses the server's destination port; return traffic reverses addresses and ports. A stateless ACL does not automatically permit a session's replies.
- AOS-CX IP ACEs use **prefix lengths or subnet masks**, such as 10.1.6.0/24. Do not copy wildcard-mask assumptions from another vendor; use explicit hosts or groups when a compact mask could include unwanted addresses.
- To block one subnet pair while allowing other traffic, place the specific **deny before a broad permit**. For an allowed range with exceptions, arrange exception rules before any permit that would otherwise match them.
- Check **ingress/egress direction** from the switch interface's perspective and the supported port, VLAN or control-plane scope. An ACL that never sees a flow cannot enforce the intended policy.

Traffic Policy & QoS / Management & Analytics

Aruba Certified Switching Professional

KNOWLEDGE CHEATSHEET

TRAFFIC TREATMENT

Match — ACL or classifier → **Mark** — Local priority → **Queue** — Map traffic → **Schedule** — Serve and shape

TCAM, counters and logging

- **Address and port groups** simplify repeated policy definitions. Their members can expand into hardware matches; shorter configuration does not guarantee lower **TCAM** consumption.
- Hardware resources depend on **ACE expansion, attachment scope and platform capacity**. Inspect resource utilization when adding large groups or applying policy widely; do not assume every ASIC compresses rules identically.
- **count** tracks matching packets; **log** generates diagnostic records subject to logging limits and processing. Log totals can differ from hardware counters, so use counters for volume and logs for context.

Classification, scheduling and congestion

- The QoS path is **classify → assign local priority → map to queue → schedule**. A frame's **802.1p PCP** or IP **DSCP** marking matters only when the ingress trust or classification policy uses it.
- Marking a VLAN as **voice** supports voice identification and advertisement. It does not alone create end-to-end priority; configure trusted markings, queue mapping and scheduling consistently. DSCP **EF 46** commonly identifies voice media.
- **Strict priority** serves a selected queue ahead of others, reducing voice delay but risking starvation. **DWRR/WFQ** distribute service by weights; weights describe relative service under contention, not a fixed transmission rate in every condition.
- **Policing/rate limiting** constrains traffic and may drop excess packets. **Egress shaping** buffers and paces output, using a rate and burst allowance. Neither creates more link bandwidth.
- **Virtual output queuing (VOQ)** separates ingress buffering by destination output, reducing head-of-line blocking when one egress is congested. Buffering absorbs bursts but cannot eliminate loss under sustained overload.
- Verify the platform's **queue count, scheduler support, trust state and drop counters**. Scheduler names and capabilities vary across switch families; do not infer them from another CX model.

06 Management, automation and troubleshooting

NetEdit and configuration plans

- **NetEdit / Switch Multi-Edit** centralizes multidevice configuration, validation and change tracking. NetEdit is deployed as a supported virtual appliance; size managed-node entitlement and support to the intended estate.
- For AOS-CX access, configure a real **administrator password**, HTTPS/REST API access with the required write permissions, credentials and management-VRF reachability. A blank factory password is not a usable automation credential.
- A plan progresses through **edit, validate, deploy and commit**. Deployment changes the running configuration; commit finalizes the plan and saves the configuration for persistence.

Management & Analytics

Aruba Certified Switching Professional

KNOWLEDGE CHEATSHEET

OPERATE AND VERIFY

- **NetEdit** — Configuration plans
- **NAE** — Observe and alert
- **ERSPAN** — Remote capture
- **Counters** — Verify behavior

- **Syntax validation** checks command form. **Semantic validation** checks configuration meaning and relationships, including through a temporary configuration-state copy. **Conformance checks** compare configuration against defined rules.
- Preserve **command context and indentation** when pasting hierarchical configuration. Validate the complete plan and review differences before deployment rather than treating each line as a global command.
- Discovery and topology use **IP reachability, SNMP and neighbor information**; third-party interaction can also use SSH. Devices may reappear while discovery still finds them. Use broader monitoring such as **AirWave** when mixed-device operational visibility is the main requirement.

Network Analytics Engine

- The **Network Analytics Engine (NAE)** runs on the switch and correlates state, events and time-series data for root-cause analysis. A **Python script** defines monitors, rules and actions; an **agent** instantiates that script with parameters.
- Use the switch **Web UI** to install scripts and create agents. A script file alone is not a running monitor. Select parameters and thresholds for the actual uplink, protocol or resource being observed.
- **Monitors** observe state or statistics; **rules** evaluate conditions; **actions** generate alerts or diagnostics. Choose observations that distinguish a cause from a symptom, such as link loss versus an application timeout.
- For agent errors or failed script installation, check **script/agent resource limits, compatibility, permissions and parameter validity**. Remove unused objects when capacity is exhausted; a red status icon alone is not a complete diagnosis.
- Synchronize switch **NTP** and verify the workstation clock and browser time zone when graphs appear shifted. Configuration-event history and AAA accounting complement NAE when investigating who changed a device.

Capture and operational verification

- **ERSPAN** mirrors selected traffic across an IP network using GRE encapsulation. Specify a reachable capture endpoint that can decode the tunnel; the destination is not automatically the switch nearest the analyst.
- Check **mirror direction, source ports, tunnel reachability and capture capacity**. Mirroring can oversubscribe its destination, so an incomplete capture does not prove packets were absent from the production path.
- The **AOS-CX mobile app** supports local setup and troubleshooting, including supported SSH access and file transfer. Use authenticated management access and confirm the switch model's connection options.
- Troubleshoot in order: **physical link and PoE → VLAN/LAG → authentication and role → routing/VRF → policy/QoS → application**. Compare counters and neighbor state on both sides of the failing boundary.