Junos Operations

Security, Associate (JNCIA-SEC)

KNOWLEDGE CHEATSHEET

CONFIGURATION WORKFLOW

Candidate Edit → **Validate** Commit check → **Activate** Commit → **Recover** Rollback

01 Junos fundamentals and operations

CLI and configuration control

- **Operational mode** uses the > prompt for monitoring. **Configuration mode** uses # and edits a candidate configuration. The root shell commonly uses %; enter **cli** to start the Junos CLI.
- **edit** enters a hierarchy; **up** moves one level; **top** returns to [edit]. **exit** leaves the current hierarchy, or leaves configuration mode when already at its top.
- **configure exclusive** locks configuration editing against other users. **show | compare** displays candidate changes; **commit check** validates them; **commit** activates them.
- **rollback 0** discards uncommitted candidate changes. Saved indices **0 through 49** represent the active configuration plus up to 49 earlier configurations. After loading an earlier configuration, commit to activate it.
- **commit confirmed** automatically reverts unless confirmed before its timer expires. The **super-user** login class has full permissions; factory initialization commonly shows the hostname **Amnesiac**.

Interfaces and diagnostic commands

- **ge-0/1/2.0** means media type ge, FPC slot 0, PIC slot 1, port 2 and logical unit 0. **Link speed** and duplex are physical properties; IP addressing and VLAN encapsulation belong to logical configuration.
- **family inet** enables IPv4 addressing; **family inet6** enables IPv6. Compare administrative and link state with **show interfaces terse**, including the logical unit assigned to the security zone.
- **show chassis hardware** displays inventory and serial numbers. **show chassis routing-engine** displays Routing Engine health and utilization. **show system license** checks installed licenses and feature usage.
- In **ping**, the target is the destination; **source** selects a local source address and **routing-instance** selects the routing table. In J-Web, management-address settings configure device management connectivity.

Junos Operations / SRX Architecture

Security, Associate (JNCIA-SEC)

KNOWLEDGE CHEATSHEET

NEW-FLOW DECISIONS

Ingress Filters and screens → **Translate** Destination → **Decide** Route and policy → **Process** Source NAT and services

Logging and management scope

- **System logs** describe device events; **security logs** describe traffic and security activity. A common factory syslog destination is the local **messages** file; inspect the actual syslog configuration.
- Enable **security logging** and policy session logging to supply traffic data for security reports. Session-init records creation; session-close records termination and final traffic statistics.
- **J-Web** manages an individual firewall. **Sky Enterprise** provides cloud management of distributed devices; **Junos Space Security Director** centralizes security policy management. **Secure Analytics** correlates security events.

02 SRX packet flow, zones and screens

Stateful flow and stateless forwarding

- **Flow mode** is the normal SRX security-processing mode. The first packet creates session state after route, policy and required service checks; later packets use the session's forward or reverse flow.
- **Packet mode** bypasses flow-based session processing. Routing, stateless firewall filters and CoS can operate without stateful security; platform support determines features such as MPLS.
- **Stateless firewall filters** evaluate packets individually and can coexist with flow mode. Ingress per-packet filters and policers act before the flow security processing stages.
- **Exception traffic** needs special handling: examples include packets addressed to the device and expired-TTL packets that generate traceroute responses. A transit SSH or HTTP session is not an exception merely because of its application.
- **show security flow session** exposes session endpoints, policy name, interfaces, timeout and packet counters. The generic **TCP idle timeout is 1800 seconds**; application-specific settings can change it.

Security zones and device access

- A **security zone** groups logical interfaces with shared security requirements. An interface belongs to one zone; zone objects can include an attached address book, screens and host-inbound permissions.
- The built-in **null zone** contains interfaces not assigned to another zone. Removing a zone assignment returns the interface to null; traffic to or from such an interface is discarded.
- The **management functional zone** supports out-of-band management and does not forward transit traffic or participate in transit security policies. **junos-host** represents traffic originating from or terminating on the firewall.

SRX Architecture / Security Policies

Security, Associate (JNCIA-SEC)

KNOWLEDGE CHEATSHEET

POLICY DECISION LAYERS

Zone Context rules → **Global** Shared rules → **Default** Fallback → **Session** Reverse flow

- **host-inbound-traffic** allows selected system services, such as SSH or ping, and routing protocols to reach the device. Configure it at zone or logical-interface scope; explicit interface settings take precedence for that interface.
- Keep **host-inbound permissions** separate from transit policy. Allowing SSH to the SRX does not permit SSH through it. A security zone can contain interfaces associated with different routing instances; routing still controls reachability.

Screens and first-packet processing

- **Screens** detect common network attacks such as SYN floods, ICMP floods, malformed packets and IP spoofing. Attach a screen profile to the ingress security zone.
- Screens inspect **incoming traffic** in the forwarding security path. They complement policy and application inspection; antivirus and malware attachment analysis belong to other services.
- For a new flow, screening precedes the policy decision. **Destination translation** occurs before route and policy lookup; the route determines egress context. **Source translation** and enabled services follow policy processing.

03 Security objects and policies

Address books and applications

- **Address books** name hosts, networks and address sets. A zone-attached book supplies addresses for that zone; the **global address book** supports shared objects across zones.
- Use a host /32 or a correctly aligned network prefix. For example, **198.51.100.0/24** is a network; a /24 network address has all eight host bits zero.
- A Junos **wildcard-address mask** uses 1 bits for positions that must match and 0 bits for positions to ignore. Noncontiguous masks belong to wildcard syntax, rather than ordinary subnet-prefix syntax.
- **Application objects** describe services using IP protocol, ports and related parameters. Use predefined **junos-ssh** or **junos-ping**, custom applications, and application sets as needed. A source port is usually ephemeral.

Policy matching and actions

- A zone policy has **from-zone** and **to-zone** context. Its primary matches are **source-address**, **destination-address** and **application**, followed by an explicit action.
- **permit** allows matching transit traffic and can invoke additional services. **deny** silently drops it. **reject** drops it and sends an appropriate response, such as a TCP RST or an ICMP unreachable.

Security Policies / NAT

Security, Associate (JNCIA-SEC)

KNOWLEDGE CHEATSHEET

POLICY DECISION LAYERS

Zone Context rules → **Global** Shared rules → **Default** Fallback → **Session** Reverse flow

- Policy lists use **first match**, in configured order. Place specific rules before broader ones with **insert ... before ...**. Renaming a policy does not change its position.
- The normal lookup hierarchy is **zone policies**, then **global policies**, then **default policy** when no earlier rule matches. Global policies can match zone context but do not require a separate list for every zone pair.
- The unconfigured fallback is **deny**. Factory configurations may install explicit trust-side permits; inspect active policies instead of assuming that a zone name grants access. Return traffic uses the permitted session's reverse flow.

Application and identity awareness

- **AppSecure** adds application awareness. **AppID** identifies applications beyond simple port numbers; **AppFW** controls them. Early packets may need to pass under initial policy handling before identification enables a final application decision.
- **Unified security policies** combine traditional address and service matches with dynamic application criteria and security services. Policy order and context still matter; the firewall does not simply choose the most restrictive rule.
- **Integrated user firewall** maps user identities to IP addresses for user-based policy. Directory integration and identity collection supply mappings without requiring an 802.1X supplicant on every endpoint.
- **IDP/IPS** inspects application-layer attack patterns; AppFW enforces application usage. Identity-aware policy identifies users, while zone membership identifies network security context.

04 Network address translation

Source NAT and pool behavior

- **Source NAT** changes the source address of a new outbound flow. **Interface NAT with PAT** lets many internal hosts share the egress interface address using distinct translated ports.
- A **source pool** supplies translated addresses. **PAT** multiplexes sessions through ports; **no-PAT** preserves ports and needs sufficient addresses. No-PAT alone does not define a fixed host-to-address mapping.
- **Address shifting** provides a deterministic offset between an original address range and a translated range. Distinguish this mapping from dynamic allocation out of a pool.
- **Address persistence** keeps a host on the same translated pool address across sessions. Without it, the same pool address is not guaranteed. Persistent NAT bindings and their permitted remote endpoints are separate controls.

NAT / Threat Protection

Security, Associate (JNCIA-SEC)

KNOWLEDGE CHEATSHEET

TRANSLATION CHOICES

- **Source** — Outbound identity
- **Destination** — Published service
- **Static** — Fixed mapping
- **Proxy ARP** — Local reachability

Publishing services and fixed mappings

- **Destination NAT** rewrites an incoming destination to an internal server. Match destination ports to publish separate services, such as web and FTP, through one public address and different NAT rules or pools.
- **Static NAT** creates a fixed, bidirectional mapping, commonly one public address per internal host. Reverse static translation lets that host originate traffic using its mapped public identity.
- Static NAT can also support explicit **port mapping**. Do not equate static NAT with a blanket prohibition on port translation; its central distinction is the fixed mapping and reverse translation.
- Use **destination NAT plus separate source NAT** when an inbound server address and its outbound internet identity must differ. Use a suitable static mapping when both directions should use the same public identity.

NAT order and reachability checks

- For a new inbound flow, **static NAT precedes destination NAT**. Route and security-policy lookup follow destination translation. **Reverse static NAT takes precedence over source NAT** for the corresponding outbound mapping.
- An inbound policy matches the **translated destination address** and original source address. Source NAT is applied after policy lookup, so outbound policy matches the original internal source.
- NAT does not grant access. Configure the required **security policy**, correct zone assignments and routes, plus a valid return path. NAT rule-set contexts must agree with the actual ingress and egress topology.
- Configure **proxy ARP** when the SRX must answer ARP for translated addresses on a directly connected public subnet. If the upstream router routes that prefix to the SRX, the reachability design differs.
- **Hairpin NAT** lets an internal client reach an internal server through its public identity. Source translation may be needed to keep both directions passing through the firewall.

05 UTM and advanced threat protection

UTM feature selection and attachment

- **Unified Threat Management (UTM)** includes antivirus, antispam, Web filtering and content filtering. These inspect supported application traffic, including HTTP, FTP and SMTP; feature and protocol support varies.
- **Antivirus** detects malicious files. **Antispam** evaluates unwanted email, particularly SMTP. **Content filtering** controls file types, extensions, MIME types and other supported content attributes.
- Use **content filtering** to block executable downloads and **Web filtering** to control browsing categories. A simple requirement to block all POP3 connections can be enforced by a security policy matching that service.

Threat Protection

Security, Associate (JNCIA-SEC)

KNOWLEDGE CHEATSHEET

PROTECTION LAYERS

- **Screens** — Network attacks
- **UTM** — Content and URLs
- **IDP** — Attack patterns
- **ATP Cloud** — Malware intelligence

- Build a **feature profile**, reference it in a **UTM policy**, then attach that UTM policy to a permitting **security policy**. UTM runs in the services stage after the policy and NAT processing stages.
- For antispam, define local allow or block entries as **custom objects** and reference them in the antispam feature profile. Object creation alone does not activate enforcement.

Web filtering methods

- **Local Web filtering** uses administrator-maintained URL allowlists and blocklists. **Enhanced Web Filtering** queries an internet categorization service and applies category-based actions.
- **Websense Redirect** delegates supported URL decisions to a Websense server. Distinguish an external categorization server from a locally maintained URL list.
- Web filtering primarily decides from the **URL or its category**. Content filtering examines supported content attributes. Together they support acceptable-use controls and restrictions on downloads.

ATP Cloud analysis and onboarding

- **Juniper ATP Cloud** supplies cloud malware analysis, including static inspection and dynamic sandbox execution to uncover previously unknown threats. It complements local signature-based protection.
- The analysis pipeline starts with **file-hash cache lookup**, followed as needed by **antivirus scanning**, **static analysis** and **dynamic analysis**. A known verdict can avoid repeated analysis.
- Onboarding uses an **ATP Cloud account or realm** and enrollment of the SRX, commonly through the generated enrollment script. Available services depend on platform support, connectivity and subscription entitlement.
- **Connected Security** shares threat intelligence and coordinates enforcement across security components. Malware detection and blocking integrate into a broader defense; IPsec connectivity serves a different purpose.

Threat Protection / VPNs

Security, Associate (JNCIA-SEC)

KNOWLEDGE CHEATSHEET

VPN BUILDING BLOCKS

- **IKE** — Authenticate peers
- → **IPsec** — Protect traffic
- → **Routing** — Select tunnel
- → **Policy** — Authorize access

Threat feeds and enforcement

- **Command-and-control (C&C) feeds** identify known botnet infrastructure and malicious download sources. Blocking these destinations helps stop compromised endpoints, including IoT devices, from contacting controllers.
- **GeoIP feeds** map IP prefixes to geographic locations and support country-based policy. Feed updates track changing address allocations; an IP location is not proof of a user's nationality.
- **Infected-host feeds** identify internal endpoints associated with detected threats. **Allow and block lists** provide explicit local intelligence; attach the applicable enforcement policy rather than merely downloading feeds.

06 IPsec VPNs and remote access

IKE negotiation and IPsec protection

- In **IKEv1 Phase 1**, peers authenticate and establish a **bidirectional IKE SA**. In **Phase 2**, they negotiate **unidirectional IPsec SAs**, normally a pair for two-way traffic.
- Peers need compatible **encryption**, **integrity**, **Diffie-Hellman group**, authentication and IKE mode. Local VPN object names and physical interface names do not need to match.
- **ESP** provides payload encryption and can supply integrity protection. **AH** provides integrity and authentication without encryption. **IKE** negotiates keys and SAs rather than encrypting application payloads itself.
- Recognize **MD5** and **SHA-1** as legacy integrity algorithms in older configurations. Prefer supported stronger algorithms for new deployments. **Perfect forward secrecy (PFS)** adds fresh DH key material for IPsec SA negotiation.
- IKEv1 **main mode** protects peer identity during negotiation; peer addressing and identity configuration must fit the static or dynamic gateway design. **Pre-shared keys** and **certificates** are peer authentication methods.

Tunnel routing and traffic selectors

- A **route-based VPN** binds an IPsec VPN to an **st0** logical tunnel interface. Routes steer remote prefixes into the tunnel; zones and policies authorize the protected traffic.
- A **policy-based VPN** selects a tunnel through policy. **Proxy IDs** or traffic selectors describe local and remote protected networks and services; the peers must agree on compatible selectors.
- When comparing selectors, local and remote roles reverse at the opposite peer. **Peer authentication** verifies identity; selector negotiation checks which traffic the IPsec SAs protect.
- In J-Web, configure the local and remote gateways, peer authentication, compatible proposals and protected networks. Verify the route, zone policy and return path after negotiation succeeds.

VPNs

Security, Associate (JNCIA-SEC)

KNOWLEDGE CHEATSHEET

VPN BUILDING BLOCKS

IKE Authenticate peers → **IPsec** Protect traffic → **Routing** Select tunnel → **Policy** Authorize access

Liveness, NAT traversal and checks

- **Dead peer detection (DPD)** tests peer liveness and removes stale state after repeated failures. The SRX default interval is **10 seconds**; the default failure threshold is **5**.
- **Optimized DPD** probes when outgoing traffic has no corresponding incoming peer traffic. **Always-send** probes periodically; **probe-idle-tunnel** also covers idle tunnels.
- **NAT traversal (NAT-T)** encapsulates IPsec ESP in UDP, typically port **4500**, when NAT is detected. Initial IKE negotiation normally uses UDP **500**.
- Use **show security ike security-associations** for the IKE control channel and **show security ipsec security-associations** for data SAs. An established IKE SA alone does not prove protected traffic can cross the tunnel.

Juniper Secure Connect

- **Juniper Secure Connect** provides client-based remote-access VPN connectivity to an SRX. Configure the gateway and install the client on supported platforms such as **Windows**, **macOS** and **Android**.
- The SRX includes **two built-in concurrent-user licenses**. Additional simultaneous users require sufficient additional licensing; count concurrent connections rather than only registered user accounts.
- Users can authenticate through a **local user database** or a supported **external authentication service**. Directory-backed authentication and MFA depend on the selected external service and configuration.
- Keep **gateway or peer authentication** separate from **user authentication**. Plan client address assignment, permitted internal networks, routing and security policy together so an authenticated user can reach only approved resources.