

Results of Testing: Juniper Branch SRX Firewalls

by Joel Snyder / Opus One

prepared for Juniper Networks
June 2012



Table of Contents

Introduction	1
Firewall Feature Set and Role-based Firewall	2
UTM Feature Set: URL Filtering	6
UTM Feature Set: Anti-Malware	7
UTM Feature Set: Intrusion Prevention	9
Virtual Private Networks and Remote Access	13
Wireless Support	15
IP Version 6 Support	17
Management	18
Junos Routing and Switching	20

Introduction

In May 2012, Opus One tested Juniper's branch SRX firewall¹ product line running recently-released Junos 12.1 software. The goals of the testing were to evaluate the SRX firewall from a security point of view, and to determine whether the SRX was ready to deploy into enterprise branch offices as a multi-service security device.

Opus One tested the branch SRX firewall², running literally hundreds of tests in 30 broad evaluation areas such as UTM capabilities, next-generation firewall, role-based firewall and IPSec VPNs.

Our tests show that the Juniper branch SRX firewall is fully ready for deployment in most enterprise branch office environments.

Opus One also tested Juniper's Security Threat Response Manager (STRM), a log collection and correlation tool, as an integral part of the SRX firewall. Our tests show that STRM offers valuable additional information, especially in deployments using the built-in IPS. We feel that STRM is a critical component of any mid-sized (or larger) SRX firewall deployment.

Enterprise customers and existing ScreenOS customers will find that the SRX exceeds the capabilities of the ScreenOS platform in areas such as UTM capabilities (especially IPS), clean integration of VPN and routing, IPv6 support, next-generation firewall, and advanced networking. While there are a few limited areas where the SRX and its supporting tools have not reached the sophistication level of older Juniper products, the SRX should be on the short list and test bench of every enterprise customer and existing ScreenOS customers.

Network managers with competitors' branch office firewall products will find that the branch SRX represents a new approach that complements the fusion of networking and security in organizations. Because the SRX UTM firewalls are built on top of the Junos routing platform, network managers don't have to surround the firewall with additional routing and switching devices to build a reliable security boundary. And in the branch environment, the SRX UTM firewalls have enterprise-class switching and routing capabilities, making a one-box-in-the-branch solution possible.

About Juniper SRX

The Juniper SRX firewalls are high-performance security, routing and network solutions for the enterprise. These devices pack high port-density, advanced security with application visibility and control, and flexible connectivity into a single, easily managed platform that supports fast, secure and highly-available operations.

While some of the low-end SRX platforms such as the SRX100/ 110/200 are better suited for the branch environment, many of the SRX devices, particularly the SRX240/550/650, are deployed in mid-enterprises where the SRX is not a branch device but is the enterprise security device.

The SRX firewalls are based on Junos, Juniper's proven operating system which delivers security and advanced protection services. Junos also supports rich routing and switching capabilities; Junos' unique architecture provides reliable service operations and manageability, even under the highest loads.

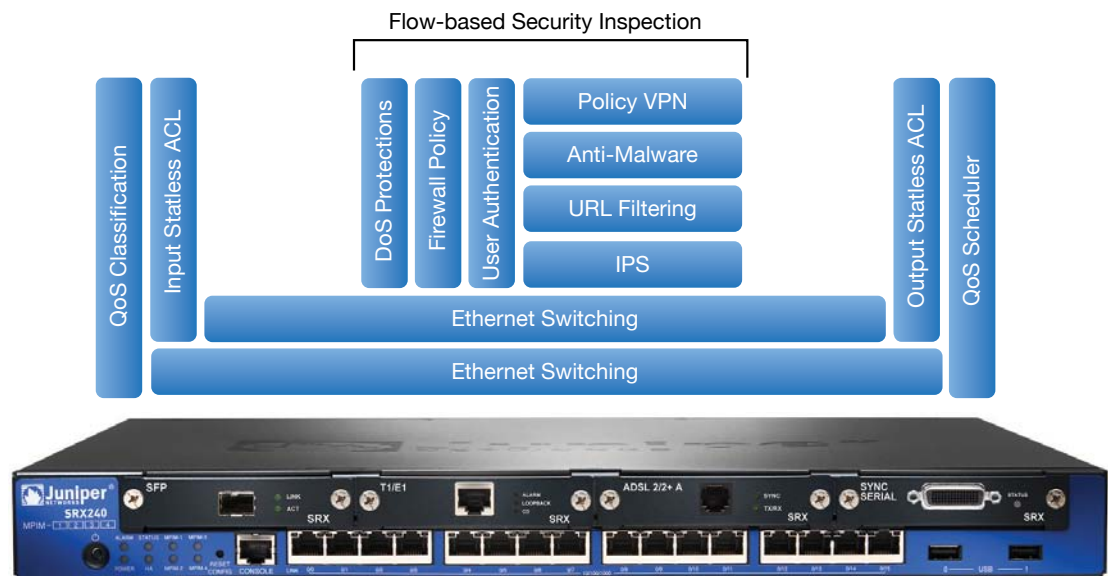


Figure 1. The Juniper SRX firewall is a flexible platform that combines routing, switching and security, with advanced threat mitigation technologies including anti-malware, URL filtering, application security and IPS.

¹ Juniper refers to the SRX product line as "SRX Series Services Gateways." For the purposes of this test, we'll simply call them "SRX firewalls."

² Juniper's SRX firewalls extend from the small SRX100, designed for limited bandwidth scenarios, up to the massive SRX5800, designed as a core firewall for data centers or service providers. The "branch" firewalls are based on a common architecture, running the same software image, and are marketed using a three-digit family, ranging from 100 to 650. The larger "four digit" SRX firewalls have a different hardware architecture and use a software image which has greater functional capabilities as well as higher performance.

Firewall Feature Set and Role-based Firewall

Firewall vendors are busily adding both breadth and depth to their product lines, but any enterprise class firewall must have some basic security and networking features to form a solid foundation. The bar for a “baseline enterprise firewall” has been raised over time so that devices are now expected to go beyond stateful inspection to include:

- Site-to-site VPNs
- Network address translation
- High availability
- QoS/CoS features such as bandwidth management
- Enterprise networking including VLANs and link aggregation
- IPv6 support
- Global management and log aggregation
- Basic dynamic routing
- Role-based (user authentication) rule base

Key Issues and Test Criteria

With its heritage as an enterprise routing platform, Junos—the technology underlying the SRX branch platforms—is not an obvious choice as a firewall. Opus One evaluated the SRX firewall to ensure that network managers choosing the SRX would not be disappointed by its firewall feature set.

A firewall has to start with the traditional rule-based security policy. There’s a reason that every firewall looks similar in this regard: the rule-based access controls are a proven way to represent security policy in a firewall. However, the simplicity of the rules needs to hide a set of application layer gateways (ALGs) that are responsible for making sure that protocols such as FTP, SIP and SCCP (voice over IP), H.323 (video conferencing), and IPSEC (VPN) operate properly and securely.

The rise of NAC (network access control) in enterprise networks has made clear the advantages of tying access controls to users, rather than simply to IP addresses. Enterprise firewalls are following the same trend, which leads to a new universal requirement: role-based firewall policies. Any firewall policy must provide the capability to be qualified by user and group information to provide identity-aware access control. The difficult part of this is not building the policy, but making the firewall aware of user identity information without adding an undue burden to the end user.

Firewalls must also have integrated site-to-site VPNs based on IPSec and network address translation (NAT) features as part of the basic security policy enforcement. Better products allow for a variety of VPN topologies to meet the needs of complex WANs in large enterprises.

While firewalls are primarily security appliances, they also have a role in enterprise networks, even in the branch office. They serve as routers between segments and may need to participate in some basic dynamic routing protocols such as OSPF and BGP to properly integrate in a reliable way to existing infrastructure. Features such as high availability, VLANs and link aggregation are needed to ensure that firewalls can be cleanly integrated into branch office and headquarters networks. Because there is increasing demand for IPv6, any device installed into an enterprise network today must be IPv6-compatible as well.

In branch environments, where enterprises may have hundreds or thousands of firewalls to manage, a centralized management system and log aggregation tool is a requirement. While many firewalls will be initially configured using a command line to simplify high volume deployments, continuing management requires an easy way to make consistent changes to many units in a short period of time. Enterprises also often require some basic Quality of Service/Class of Service configuration to help prioritize delay-sensitive traffic (such as voice and video) over bandwidth-limited and congested data circuits.

Results of Testing

Over a two week period, Opus One tested the Juniper SRX branch firewall by using both the J-Web GUI and the CLI to create and modify firewall policies. To validate proper operation of ALGs, FTP, SIP and H.323 protocols were run over the firewall both in NAT and non-NAT modes. The firewall, configured using both link aggregation and VLANs, demonstrated excellent support for these enterprise-class networking features.

Although VPN, IPv6, dynamic routing and global management capabilities are part of any enterprise firewall, we tested those separately. See “VPN” below for the results of our remote-access and site-to-site VPN testing; “IP Version 6 Support” below for the results of our IPv6 testing; “Junos Networking and Switching” below for the results of our dynamic routing testing; and “Management” below for the results of our management testing.

Security Policies, NAT, and Role-Based Firewall

NAT configuration on the SRX firewall is handled separately from firewall security policies. Common NAT cases, such as “interface NAT” (the hiding of an entire subnet behind a firewall’s external IP address), are very easy to configure, as is inbound NAT. Network managers who are used to the complex configuration of incoming NAT in older ScreenOS firewalls will find the simplicity of the SRX firewall a welcome change—even though the NAT capabilities of the two platforms are very similar.

The SRX firewall also supports user-based (role-based) security policies, extending the 5-tuple match criteria to include user roles; however, these types of security policies are not currently configurable in the J-Web GUI and must be done through the CLI or Juniper’s Security Design. Users logged in through these policies can be viewed in Juniper’s UAC (User Access Control) Web UI or the SRX firewall CLI. Role-based firewall functionality in the SRX is handled in conjunction with Juniper’s UAC product, part of Juniper’s NAC product suite. Although UAC has a strong feature set for LAN-based authentication, we simply used the UAC to provide the captive portal to users who need authentication, and then sent the credentials over to the SRX firewall.³ We were able to use the captive portal to authenticate, and then see that the correct role-based rules were applied to our user.

The SRX firewall also supports two additional authentication methods when a role-based security policy is needed. Enterprises that have selected Juniper SSL VPN or UAC products will have Juniper’s Pulse client (available on Windows, Mac OS X, iOS, Windows Mobile, Blackberry and Android platforms) already installed on their devices, and this can be used to provide a transparent login to the firewall. Additionally, a Microsoft authentication mechanism called “SPNEGO” can be used to transparently pass domain credentials from supported browsers using Kerberos, minimizing the intrusion on the end user when Windows Active Directory is in use.

High Availability

We tested high availability in several different ways. We started with an active/passive pair of firewalls in a traditional high-availability cluster using default timing settings for failover and several sessions running through the firewalls. Then, we plugged and unplugged interfaces and blocked different types of connectivity without dropping the Ethernet link. In each case, we were able to verify that traffic continued to flow with short (less than ten seconds in every case, and much less in some cases) interruptions. These timings could be reduced even further with more aggressive timers at the cost of potentially increased “false” failover events due to packet loss or network conditions.

It’s worth noting that the management of the cluster is done via the primary cluster member, and that any changes committed are automatically synchronized with the backup device.

An equally important high availability test involved detecting the failure of a communications path. In our testing, we had two upstream Internet connections, which is a common branch topology. Our goal was to detect the failure of one path and change routing in the firewall to prefer the other backup path. To do this, we used two features of the SRX firewall: IP Monitoring and Real-time Performance Monitoring (RPM). RPM is a comprehensive monitoring toolkit that

³ Network managers familiar with ScreenOS-based firewalls may be aware that ScreenOS has a small captive portal and authentication system itself. This feature has been carried forward into the SRX for network managers who want compatibility with the ScreenOS-based feature set. However, using this built-in portal and authentication do not actually result in a role-based firewall, because the users cannot be given different access depending on authentication information.

lets a Junos device (such as the SRX firewall) actively probe the network using multiple protocols (ICMP, UDP, TCP and HTTP get) to determine statistics such as latency, jitter and packet loss. We defined RPM probes, and then used IP Monitoring to take an action, changing the static routing table to modify the default route⁴ when the RPM probes failed. In our testing, we simply looked for connectivity failures to an upstream router. However, this same technique could be used to change routing between a branch office and headquarters, for example, if the path to headquarters became very lossy or had a high latency.

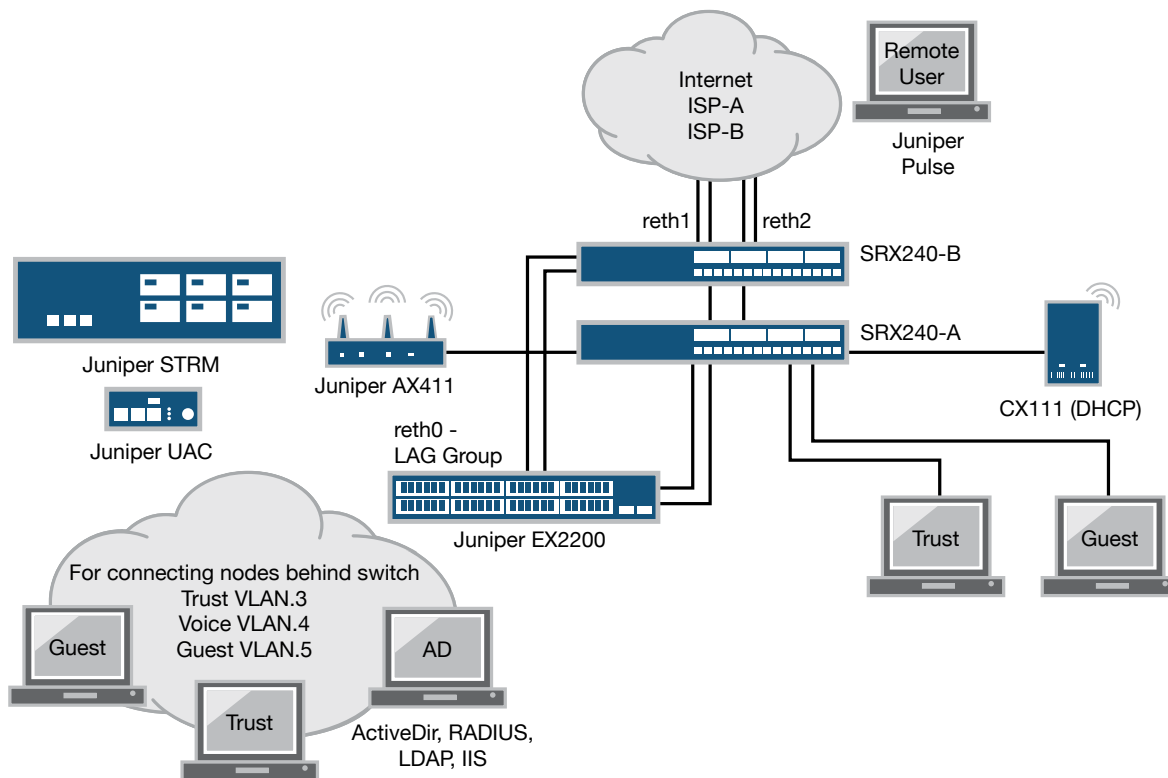


Figure 2. When deployed in a high availability cluster, the Juniper SRX firewall quickly recovers from both device failures and network failures, providing continuity of operations to end-users.

Juniper staff also showed us how high availability using IP monitoring could interact with the Quality of Service/Class of Service mechanisms in the SRX firewall. For example, they showed us how to configure a change in QoS/CoS behavior (using Junos scheduler queues) when forwarding traffic over a low-speed 3G wireless connection during a failover event. Additionally, separate security, web filtering or UTM policies could be implemented when running on the backup link. The goal here was to maintain the bandwidth allocated to services such as VPN tunnels, voice and video while reducing allowed Internet bandwidth when the backup connection was in use, thereby maintaining essential corporate applications even during a failover.

⁴Other actions are possible, such as changing the state of an interface.

Conclusions

Criteria	Notes
Create firewall security policies using normal criteria (source and destination IP, port, and protocol)	PASS SRX is a zone-based firewall, giving greater options for policy definition as well.
Create role-based firewall security policies to give different user groups different levels of access	PASS SRX role-based firewalls work in conjunction with Juniper's UAC appliance, extending the 5-tuple to match criteria to include user or group matching.
Validate correct operation of ALGs for VoIP and video, including when NAT is in place	PASS SRX firewall has a large set of media ALGs, including SCCP and MGCP in addition to the SIP and H.323 we tested.
Configure NAT using typical network scenarios	PASS SRX firewall NAT is configured separately from firewall policy, a more flexible approach than used in ScreenOS.
Build site-to-site VPNs using IPsec	PASS SRX firewall supports both IKE v1 and IKE v2. A VPN wizard is able to simplify the process of creating VPNs. More VPN testing appears below.
Configure and test active/passive high availability	PASS SRX firewall kept existing sessions alive and passed traffic when failure of the active node was detected during every test.
Use high availability features to switch outbound routes even when cluster connectivity is normal	PASS SRX firewall Realtime Performance Monitoring and IP Monitoring work together to change routing based on many possible criteria, providing high availability during network and IPS failures.
Configure QoS/CoS to manage bandwidth for different types of traffic	PASS SRX firewall QoS/CoS is configured separately from firewall policy, and may require some duplication of firewall rules and objects to match policies. SRX firewall QoS/CoS is not TCP-aware, so it throttles bandwidth by discarding packets.
Verify support and correct operation for VLANs and link aggregation (multiple Ethernet links)	PASS SRX firewall builds on Junos routing and switching platforms, and is discussed below in "Junos routing and switching."
IPv6 support	PASS SRX firewall IPv6 support is discussed below.
Dynamic routing support	PASS SRX firewall dynamic routing is discussed below in "Junos Routing and Switching."
Evaluate support for global management systems and log aggregation	PASS Juniper STRM log management is discussed in detail below in "Management."

UTM Feature Set: URL Filtering

URL filtering is used to control web browsing by blocking traffic to web sites based on policy criteria. URL filtering can be used to control the types of sites that users can browse (for instance, blocking time-wasting traffic such as games and personal shopping, or work-inappropriate traffic such as hate speech or pornography) and can help protect users against Internet security threats such as malware and phishing attacks.

Key Issues and Test Criteria

No URL filtering product is 100% effective, so products should be evaluated on additional criteria, including comprehensive logging information for both blocked and allowed traffic. Full logging helps when exploring compliance issues, and in debugging and resolving complaints.

Better products allow the network manager to configure different URL filtering policies for different networks (such as differentiating “staff users” from “guest users”) and different user groups (such as differentiating “marketing group” from “administrative group”).

Testing has shown that combining reputation-based information with more traditional lists of URLs provides a higher level of threat protection to end users. Better products incorporate both types of technologies as an option in their URL filtering engines.

Enterprise UTM products also give the network manager a choice of URL filtering engine, allowing the network manager to pick the best product based on their requirements for performance, coverage and scalability. These products not only have well-respected security vendors represented, but also afford choice among the vendors included in the stack. This is particularly important for anti-malware services, where an enterprise needs to pick products that properly complement its server-side and desktop security strategies.

Results of Testing

The Juniper SRX firewall includes a choice of four URL filtering “engines.”⁵ Only one of the four engines may be active at one time, giving network managers the choice of:

- Juniper Enhanced – Cloud-based URL filtering which combines both traditional categories and site reputation information. This approach is preferred when URL filtering is also used to block malware and phishing attacks.
- Websense – “Off-box” URL filtering done via integration with an existing enterprise Websense URL filtering engine. Most URL filtering policies are configured on the Websense engine, leaving the SRX firewall with minimal configuration required.
- Local – URL filtering based entirely on white-lists and block-lists defined by the network manager. Typically used where just a few sites should be blocked, or in an environment where only white listing is needed (such as a kiosk). Does not require a separate subscription license.
- Surf Control – a traditional URL filtering engine with category lists and the option of a different policy on each firewall rule.

To test URL filtering, we created multiple URL filtering policies to be applied on different firewall rules. We wrote these rules based on a user’s identity, but policies could be based on any other policy criteria, such as IP subnet, security zone or failover status. We verified that the policies were properly applied by testing each. We believe that the ability to apply different policies in different situations is an important feature that distinguishes enterprise branch office devices such as the Juniper SRX from lower-end products without this flexibility.

We also used Juniper’s Security Threat Response Manager (STRM) reporting platform to validate that both blocked and allowed web browsing traffic was being logged with information that a network manager would need to help discuss results with end users or to debug problems. The SRX firewall sends log entries including source and destination IP and user identity information, the URL being checked, the category, and the policy and action matched to the STRM platform. Using the STRM GUI, both individual log entries and aggregated statistics can be explored.

⁵All but the “local” engine require a separate subscription license. The SRX firewall can create its own 30-day demonstration licenses so that network managers can try each engine and choose the solution that best meets their needs.

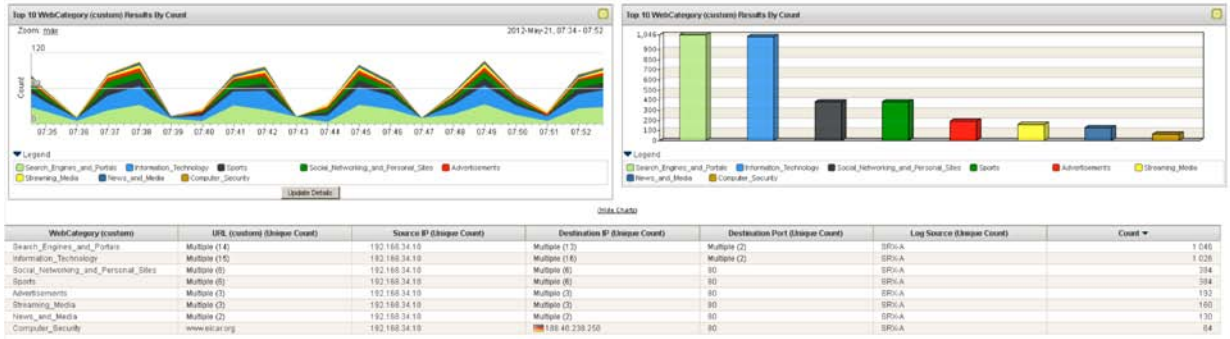


Figure 3. Juniper’s STRM reporting console shows both allowed and blocked URLs, and can display this information by URL Category or threat level.

Conclusions

Criteria	Notes
Ability to define separate policies for different users, groups, and networks	PASS URL filtering is applied on a per-firewall rule basis, and each rule can invoke a different policy.
Choice among URL filtering engines	PASS SRX firewalls have four different options, including integration with existing enterprise Websense engine (if present). Only one engine can be active at a time.
Inclusion of reputation-based filtering as an option in URL filtering	PASS The Juniper Enhanced engine includes reputation services to increase the effectiveness of anti-malware protection through URL filtering.
Accurately blocks and allows URLs per policy	PASS The Juniper SRX firewall has excellent URL filtering capabilities.
Fully logs both blocked and allowed traffic with category and URL info	PASS SRX, combined with STRM, gives access to logs as well as usage graphs.

UTM Feature Set: Anti-Malware

Enterprise best practices call for anti-malware software to be running on every desktop. However, network managers frequently deploy anti-malware on edge devices such as UTM firewalls to provide an additional layer of protection. Edge anti-malware helps protect branch users if their desktop anti-malware falls behind in updates or is shut off entirely.

Key Issues and Test Criteria

Anti-malware in UTM firewalls is a challenging problem. Traffic flying through the firewall doesn’t look the same as a virus sitting in a file on a disk, and network managers should consider edge anti-malware to be secondary protection, backstopping desktop protection.

In Opus One’s testing of UTM firewalls, we find that the same anti-malware engine that misses some viruses passing through the firewall will then successfully identify the virus once it is downloaded to an end-user’s PC. This is due to the differences in scanning a file as it is passing over the wire, where it will be encoded and often compressed, and scanning the same file sitting on a disk with all the pieces together in one place in a native format.

Configuration options for anti-malware tools are normally very constrained, as features such as quarantines and virus removal are not important in this context. However, better products should allow the network manager to configure different anti-virus policies for different types of traffic, both to optimize performance and to focus protections where they are most needed.

The common case for edge anti-malware calls for scanning of web browsing and file transfer (FTP) or instant message (IM) download traffic. Better UTM firewall products allow the network manager to configure web scanning across multiple ports and, if SSL decryption is supported, in encrypted traffic. To help protect guest users, network managers should have the ability to enable scanning of other traffic, especially mail protocols.

Enterprise-quality UTM firewalls also give the network manager a choice between multiple anti-malware engines, provided

by well-respected security vendors. Choice is important because it allows the network manager to pick the engine based on their requirements and existing anti-malware products. Best practices usually call for selection of a complementary engine when edge anti-malware is in place, rather than duplicating existing server-side and desktop products.

Because of the potential for false positives or other anomalous behavior, anti-malware products must also provide full logging whenever malware is detected. The logging must include sufficient information to help the help desk or network manager troubleshoot complaints, such as the virus detected, the URL being scanned, and all user identification and IP address information.

Results of Testing

The Juniper branch SRX firewall includes a choice of three anti-malware engines. Any one engine may be active at a time, and use of an engine requires a continuing subscription fee and connectivity to the Internet. Network managers have a choice of:

- Sophos anti-malware streaming engine, using in-the-cloud technology rather than on-device pattern matching. The Sophos engine also utilizes URL checking to give additional protection against malware.
- Kaspersky Labs anti-malware engine, an on-device anti-malware engine using traditional pattern signatures downloaded via scheduled updates.
- Express AV anti-malware engine, also an on-device signature-based engine. The Express engine trades off some anti-malware capabilities in exchange for lower memory and CPU usage in firewalls that may be performance-constrained.

To test anti-malware, we created multiple anti-malware policies. Then, we applied these policies to different traffic flows through the SRX firewall. We also switched between engines to see if different engines had different coverage. Because only one engine can be active at a time and the anti-malware license is specific to a particular engine, we licensed the firewall separately for each anti-malware engine we tested.⁶

Policy configuration options in the SRX firewall for the anti-malware are appropriate for the branch environment. Fallback rules (what to do when an anti-malware scan does not complete for some reason), timeouts, size and decompression limits, special handling and white-listing for some file types, and notification messages and rules are all configurable, hitting the most commonly configured and important options.

Using a small sample of recent viruses, we verified that the SRX firewall was able to block traffic in the most critical protocols, HTTP and FTP. We verified that different anti-malware policies were properly selected by the firewall based on our firewall rules.

We also tested the ability of the SRX firewall to block viruses in email traffic, finding that it successfully blocked viruses in SMTP and POP3 traffic, but did not catch the same viruses sent over the IMAP protocol in the release tested (even though IMAP is listed as a supported protocol). Juniper acknowledges this limitation and is addressing this in an upcoming release. The branch version of the SRX firewall does not currently support SSL/TLS decryption of traffic passing from protected client to the Internet, so we did not test viruses in encrypted traffic. This is a feature that is available in the larger data center SRX firewalls (SRX1400 – SRX5800) and will be added to the branch SRX software in 2013. One area which could be improved in the SRX firewall is coverage of HTTP scanning over non-standard ports. In this version (12.1) of the SRX firewall, only configured ports (such as 80 and 8080) will be scanned for malware. However, as a next-generation firewall, the SRX is able to identify applications (such as HTTP) running on non-standard ports. This application identification knowledge is planned to be integrated with the UTM anti-malware scanners to give greater coverage.

We used the Security Threat Response Manager (STRM) tool from Juniper to verify that blocked malware was properly logged, and had all the information that a network manager or help desk needed to debug problems and identify blocked traffic.

⁶The SRX firewall can create its own 30-day free demonstration licenses so that network managers can try each engine and choose the solution that best meets their needs.

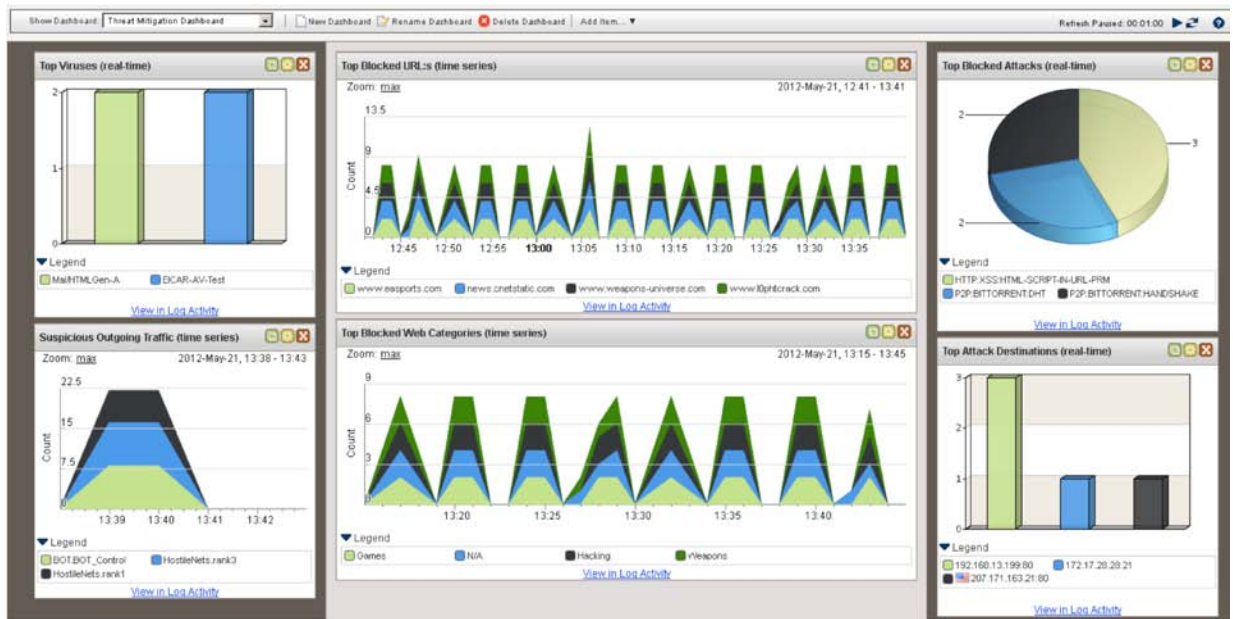


Figure 4. Juniper’s STRM reporting console logs blocked viruses and malware, correlating user identity with other log information.

Conclusions

Criteria	Notes
Ability to define separate policies for different networks and traffic types	PASS Anti-malware scanning is applied on a per-firewall rule basis, letting you turn scanning on and off for different subnets and user groups. Configuration options appropriate for this type of device and network location are available.
Choice among anti-malware engines	PASS SRX firewall has three different options, including an in-the-cloud option, an on-device engine, and a performance-optimized engine.
Accurately blocks traffic in critical protocols (HTTP and FTP)	PASS SRX firewall blocked malware found in HTTP and FTP traffic.
Accurately blocks traffic in secondary protocols (SMTP, POP3 and IMAP4)	PASS SRX firewall blocked malware in SMTP and POP3 traffic, but not in IMAP4 traffic.
Blocks application traffic on non-standard ports (such as HTTP on port 1234)	FAIL SRX firewalls do not identify selected protocols running on non-standard ports or encrypted traffic.
Logs all blocked malware	PASS SRX firewalls contain a number of alerting mechanisms to log malware, including syslog and Juniper’s STRM.

UTM Feature Set: Intrusion Prevention

Intrusion prevention technology in edge firewalls helps to move the protections provided by the firewall towards the application layer. With an attack surface closed off by the firewall at the transport layer and below, the application layer remains a huge target with what seems like an unlimited number of vulnerabilities to be exploited. Adding intrusion prevention technology between the enterprise and the Internet helps to reduce overall risk by blocking thousands of known attacks. As security researchers at Microsoft have shown, more than 80% of infected Windows systems are from malware more than six months old. While zero-day threats remain an issue, 180-day threats are a bigger operational problem for network managers.

Key Issues and Test Criteria

Intrusion prevention systems⁷ (IPSeS) are fundamentally different when deployed in the branch as opposed to large data centers. Network managers generally place protections at the edge of the network, the Internet-to-branch boundary, rather than internally between network segments. This means that internal attacks from user to server (and vice versa) within a single office are not a significant concern. Branch offices rarely have significant numbers of Internet-facing servers, which makes the primary function of IPS to protect end-users from attack and malware.

Branch offices also represent a significant IPS management challenge, as few enterprises have the resources (or interest) to track down every IPS alert from hundreds to thousands of different offices. What is needed in branch offices is a fairly conservatively designed IPS that offers moderate levels of protection against known vulnerabilities, a low false positive rate for traditional enterprise end-user traffic, an easy-to-manage rule set, and reasonable performance impact.

Policy management of IPSeS in the branch must offer flexibility to easily exempt traffic and flows from the IPS, and to bring different policies into play for different user communities. For example, streaming video and audio (such as voice over IP or videoconferencing) in the branch should sidestep the IPS to minimize performance impact and service interruptions. At the same time, a branch with a guest network may want to apply significantly stricter protections to the guest network, risking a high false positive rate, than the protections assigned to enterprise end-users in the branch.

Ideally, even though branch office IPSeS are not managed with the same intensity as core network devices, IPS logging tools should be linked with all other logs and policy management. This will enable alerts to be quickly put into context and be easily correlated across offices. Better IPS management systems provide a full spectrum of background information about alerts as well, including both vendor and open-source documentation, to help speed the network manager's understanding of what an alert means. A tight integration between IPS alerts and policy management helps reduce the number of clicks and keystrokes needed to turn a false positive into a needed policy change.

Intrusion prevention is more than just running a set of signatures across network traffic and blocking unwanted attacks. Intrusion prevention also includes other protections, such as Distributed Denial of Service (DDoS) attack avoidance and mitigation. Although these features are primarily of interest in data centers, branch offices with even a single Internet-exposed server (such as a webmail server) can come under DDoS attacks that can interrupt operations.

Results of Testing

Opus One tested the SRX firewall using a Spirent (formerly Mu Dynamics) Security Studio test tool to validate efficacy (coverage) of the IPS signature set. The SRX firewall does not have a built-in policy that is optimized for branch offices, so we created a policy selecting all IPS signatures marked as "critical," "major" and "minor" severity across all applications. Although creation of the policy was very simple, applying the policy was more challenging. Because the management system (control plane) in the branch SRX firewall is intentionally separated from the firewall's forwarding engine (data plane), compiling the policy on the firewall and installing it into the IPS engine did stress the management system for several minutes. Network managers using SRX firewalls with the IPS feature set should be prepared for longer-than-usual times to commit changes to IPS policies.

Once the IPS policy was installed, we used the Security Studio test tool to generate attacks designed to stress the client-protective IPS separately from the server-protective IPS. The SRX firewall blocked an impressive 98.5% of attacks on clients. To put this number in perspective, identical tests run by Network World over the past year gave block rates of 57% to 97.3% with similar competitive products, putting the SRX firewall ahead of the pack in protecting clients. Server-protective IPS had a lower block rate with the out-of-the-box policy we implemented, but well in line with competitive products.

⁷ Juniper uses the acronym "IDP" for Intrusion Detection and Prevention interchangeably with the term "IPS."

The table below summarizes the SRX firewall's IPS performance with our test policy and shows the outstanding coverage across known vulnerabilities.

Type of Test	SRX block rate using Critical, Major, Minor signatures	Comparable Network World results
Client-protective IPS	98.5%	57% to 97.3%
Server-protective IPS	87.1%	64% to 94.5%

Enterprise management of the SRX firewall is in transition at Juniper. The older Network and Security Manager tool (NSM) has full IPS management capabilities and meets all of the test criteria listed in this section. Juniper is preparing Juniper Security Design (SD), an application that runs on top of the Junos Space management platform as its next generation management offering. Security Design has extensive functionality for configuring firewall, IPS and VPN policies. Full UTM (including IPS) configuration in Security Design is scheduled for release to customers in late 2012. Since evaluating Security Design would have been a significant project in itself, we decided to use the CLI to manage the SRX during these tests.⁸

We found that management of the IPS using the SRX's built-in CLI was fairly easy. Although one rarely thinks of managing an IPS using a CLI because of the massive number of signatures, the SRX has many built-in IPS signature groups, cutting across different criteria (such as direction, severity and application). Since most IPS policies will make use of the built-in groups, selecting them with the CLI is not difficult. In addition, the SRX supports dynamic signature groups which let the network manager define a set of signatures based on criteria that are evaluated every time the signature set is updated. These criteria include common matches such as application, direction (client-to-server or server-to-client) and category, as well as more sophisticated criteria such as expected false positive rate and performance impact. These tools together made a daunting task not particularly difficult.

IPS is triggered as an action in a firewall security policy. This allows the network manager to control which traffic will be sent to the IPS engine for scanning, and which will bypass the IPS engine. Tying IPS policies to firewall policies is a very flexible approach, and it simplifies common IPS tasks such as exempting corporate VoIP and video conferencing traffic from IPS inspection or triggering IPS based on user, group or security zone information. Although you can achieve similar specificity in the IPS policy itself, Juniper's approach here helps to make the relationship between security policy and firewall configuration more transparent and self-documenting.

Sending IPS events to Juniper's STRM logging system, as we did in our test, gives additional benefits. STRM puts events into context by providing an IPS "super-console" that makes it easy to pivot across different types of events and explore patterns and correlations. STRM also integrates open-source material when displaying many events, such as vendor vulnerability bulletins. This gives additional layers of information to help the network manager quickly differentiate between relevant and unimportant events.


We also tested the SRX firewall's DDoS mitigation features. We set up a DDoS policy (called a "Screen" in Juniper's vocabulary) that applied to multiple zones on our test firewalls. In addition to rate-based DDoS protection, the SRX firewall's Screen can be configured to identify and block different types of network scanning (such as port scanning and IP range sweeping), IP spoofing, and many types of protocol anomalies (such as illegal TCP flag combinations, unusual IP options or IP source routing).

The Screen technology in the SRX firewall both includes server protection using SYN cookies (a DDoS avoidance tool that provides a liveness test of every connection before passing the connection onto the end server) and simpler "drop offending packet" protections when the attack rate gets too high. We activated these protections and verified that our attacking test tools were properly blocked by the SRX firewall.

⁸ Most UTM management, including anti-malware and URL filtering, is also possible in the on-board J-Web GUI, but IPS management is not fully available.

Conclusions

Criteria	Notes
Ability to define separate policies for different networks and traffic types	PASS IPS is enabled on a per-rule basis within the firewall, and the IPS policy also can apply specific signatures and signature groups on a per system or subnet basis.
Accurately blocks attacks when directed against clients	PASS SRX firewall blocked over 98% of attacks against clients.
Accurately blocks attacks when directed against servers	PASS SRX firewall blocked over 87% of attacks against servers even though server protections not commonly used in the branch.
Policy management is straightforward and fast	PASS IPS policy management will improve when Juniper Security Design is available and integrated with STRM, but existing CLI-based tools are sufficient for branch policy management.
Logs attacks (when configured), and provides context and central correlation features	PASS SRX firewall combined with Juniper STRM gives sufficient information to identify and research attacks.
Protects against DoS/DDoS attacks	PASS SRX firewall "Screens" tested to block flood attacks.

Event Information					
Event Name:	HTTP:XSS:HTML-SCRIPT-IN-URL-PRM				
Low Level Category:	Web Exploit				
Event Description:	HTTP: HTML Script Tag Embedded in URL Parameters				
Magnitude:	 (7)	Relevance:	6	Severity:	9
Credibility:	5				
Username:	N/A				
Start Time:	2012-03-30 00:54:56	Storage Time:	2012-03-30 00:55:56	Log Source Time:	1987-08-09 12:34:53
Action (custom):	CLOSE				
Application (custom):	N/A				
Bytes From Client (custom):	N/A				
Bytes From Server (custom):	N/A				
Destination Zone (custom):	untrust				
Nested Application (custom):	N/A				
Packets From Client (custom):	N/A				
Packets From Server (custom):	N/A				
Service (custom):	SERVICE_IDP				
Source Zone (custom):	trust				
Total Sessions (custom):	N/A				

Additional Information			
Protocol:	tcp_ip	QID:	6251898
Log Source:	SRX-A	Event Count:	1
OSVDB IDs:	40269 - CA eTrust SiteMinder Agent forms/smpwservices.fc SMAUTHREASON Parameter XSS 37630 - Microsoft SharePoint PATH_INFO (query string) XSS 76962 - HP Network Node Manager Unspecified XSS (2011-4155)		
Custom Rules:	BB:PortDefinition: Web Ports BB:CategoryDefinition: Exploits Backdoors and Trojans Magnitude Adjustment: Context is Local to Local Magnitude Adjustment: Destination Network Weight is Low Magnitude Adjustment: Source Network Weight is Low BB:NetworkDefinition: Client Networks BB:PortDefinition: Authorized L2R Ports BB:BehaviorDefinition: Compromise Activities System: Load Building Blocks		
Custom Rules Partial Matched :	Exploit: Chained Exploit Followed by Suspicious Events Exploit: Recon Followed by Exploit Exploit: Multiple Exploit Types Against Single Destination Exploit: Exploit/Malware Events Across Multiple Destinations		
Annotations:	Relevance has been increased by 8 because the context is local to local. Relevance has been decreased by 2 because the destination network weight is low. Relevance has been decreased by 2 because the source network weight is low.		

Figure 5. When showing expanded details on IDS attacks, Juniper's STRM will correlate multiple events and display additional context information including severity, relevancy, vendor-provided remediation information and OSVDB (Open Source Vulnerability Database) IDs.

Virtual Private Networks and Remote Access

Virtual private network (VPN) technology has been considered a “must-have” in firewalls for over a decade. Linking secure communications and access controls makes VPNs and firewalls a logical combination. In the enterprise environment, where each network element represents a support liability and potential failure point, combining VPN, remote access and a UTM firewall in a single system makes even more sense.

Key Issues and Test Criteria

The years that have passed since the IETF adopted the IKE and IPSEC protocols for IP layer VPNs have not made those protocols any less complicated. However, firewall vendors have worked hard to build tools to simplify and ease deployment of large VPNs. Firewall GUIs, wizards and the accompanying global management systems are there to help the network manager build large VPNs with complex topologies. Better products combine VPN design, deployment and monitoring. At the very least, any firewall should have built-in deployment wizards to assist the network manager in designing and configuring their VPN tunnels.

In most large branch office environments, some type of dynamic routing protocol is needed to provide reachability information throughout the network and allow for alternate path routing of traffic. As more enterprises elect to use less-reliable services such as the Internet (rather than private MPLS-type services) for their VPNs, failover from one ISP or circuit to another based on some type of dynamic routing protocol is a near-universal requirement. Better products include optimized routing protocols that reduce chattiness and suppress routing updates over high-cost circuits such as 3G wireless backup links.

The combination of the IPSec Security Policy Database (called for in the IETF’s VPN standards) and any firewall’s own security policy can be very limiting, especially in dynamic routing or multicast network environments. Firewall vendors have developed strategies such as route-based VPNs (VPNs that encrypt traffic based on IP routing tables rather than firewall security policies) to help simplify overall deployment. Network managers have a clear need for proprietary and semi-proprietary workarounds to handle the deficiencies of the IKE and IPSEC protocols. Firewall vendors cannot slavishly hide behind the RFCs, but must adapt and extend this work to make deployment of VPNs realistic. While every product should be interoperable with other standards-based products, necessary extensions and interpretations of IPSEC and IKE are usually needed for overall usability of VPNs.

Branch offices are not usually major entry points for enterprise remote access, but it is often desirable for local staff or technical support personnel to be able to connect to the office directly. Remote access VPN capabilities to simplify this type of connection are a characteristic of better enterprise-class products.

Results of Testing

We tested the Juniper SRX firewall by building IPSec VPN tunnels between our test cluster and a remote Juniper firewall.⁹ To speed the configuration of the VPN tunnel, we used the wizard included with the J-Web GUI, reducing total time for configuring and establishing the VPN to only a few minutes. VPN tunnels established with the wizard are easily re-configured in J-Web or using the CLI.

To establish dynamic routing over the VPN tunnel, we used J-Web to start up OSPF routing over the tunnels, and verified that reachability information from both sides was being accepted by the routers at each end of the tunnel. Next, we built a second tunnel and brought OSPF up through that tunnel.

In addition to OSPF’s own timers and link failure detection, the SRX firewall has several ways of monitoring tunnel status. These include Juniper’s own VPN Monitor tool (which checks IPSec reachability), IETF-standard Dead Peer Detection, IETF-standard Bidirectional Forwarding Detection, and Juniper’s Realtime Performance Monitoring (RPM) tool, which can reach beyond the tunnel end point into the network at the other end to detect true end-to-end tunnel failures. We selected VPN Monitor to detect tunnel failure because it was the simplest approach—a single check box in the wizard.

⁹ We used IKE v1 for our testing, but IKEv2 is also supported.

We caused a failure of one tunnel and observed that the traffic flowed down the other tunnel within 10 seconds without impacting existing sessions.

We stressed the routing by configuring the two firewalls at either end of the tunnel differently so that the failover across the tunnel was not symmetric. With this in place, we arranged for traffic to leave the SRX firewall cluster we were testing through one tunnel, and be sent back through the other tunnel. We observed that this asymmetric routing did not interrupt traffic or impact existing sessions.¹⁰

Although we used route-based VPNs for our testing, the SRX firewall also supports policy-based VPNs, which are usually needed for interoperability with non-Juniper devices or in certain higher-security environments. Policy-based VPNs are configured using the J-Web or CLI on the SRX firewall. The Junos base that the SRX firewall is built on also supports other more types of VPNs, including Layer 2 and Layer 3 MPLS VPNs and multicast traffic over VPNs using the IETF GDOI (RFC 3547). Typical enterprise branch offices would not use these types of VPNs, so we did not test them, but service providers selecting the SRX firewall as their Customer Premises Equipment (CPE) for delivering MPLS services should take note of this high-end feature set in a low-cost security device.

We also tested remote access VPN using the SRX firewall as a VPN concentrator. Juniper refers to this type of a VPN as a "Dynamic VPN" and both traditional end-device IPsec clients and Juniper's own Pulse client are supported. We tested this using the Pulse client¹¹ and were able to easily design and deploy remote access VPN tunnels. In our deployment, end users connected to a web page hosted on the SRX firewall that provided authentication and automatically downloaded and launched a pre-configured copy of the Pulse client.

Remote access VPN users on the SRX firewall may either be authenticated locally to a small on-device database, or can be separately authenticated through a RADIUS server. We tested local authentication using the VPN wizard to define the remote access VPN.

Conclusions

Criteria	Notes
Test for full support of IPsec-based site-to-site VPN tunnels	PASS SRX firewalls support both IKE v1 and IKE v2.
Validate easy design and deployment of VPNs using simplified on-device configuration	PASS SRX wizards support both site-to-site and remote access VPN definition.
Test for dynamic routing support over VPN tunnels	PASS SRX firewalls include support for OSPFv2/v3 (tested), as well as RIP/RIPng, reduced-traffic RIP, BGP/MBGP, IS-IS, multicast routing, and Bidirectional Forwarding Detection (BFD, RFC 5880).
Validate asymmetric routing over VPN tunnels does not interrupt traffic	PASS Asymmetric routing over VPN tunnels is a common problem during routing table convergence in some products, but not with the SRX.
Configure policy-based and route-based IPsec site-to-site VPNs	PASS Most SRX firewall users will choose route-based IPsec VPNs for their management simplicity and flexibility.
Build remote-access VPNs for end-user network connection to branch offices	PASS Both traditional IPsec remote access clients and Juniper's own Windows Pulse client are supported. Users without the Pulse client will automatically download it simply by browsing to the SRX firewall and logging in.
Validate support for central design and deployment of VPNs	Not Tested: Juniper Security Design is Juniper's central management solution for SRX firewalls with full centralized VPN configuration capability, but was not tested.

¹⁰ A key requirement for this capability is placement of both tunnels in the same security zone, which is the most common configuration.

¹¹ Although Pulse client can run on Windows, Mac OS X, Unix, iOS, Windows Mobile, Blackberry and Android platforms, only the Windows Pulse client is supported for Dynamic VPN with SRX firewalls today. Juniper's SSL VPN product line supports all Pulse client platforms.

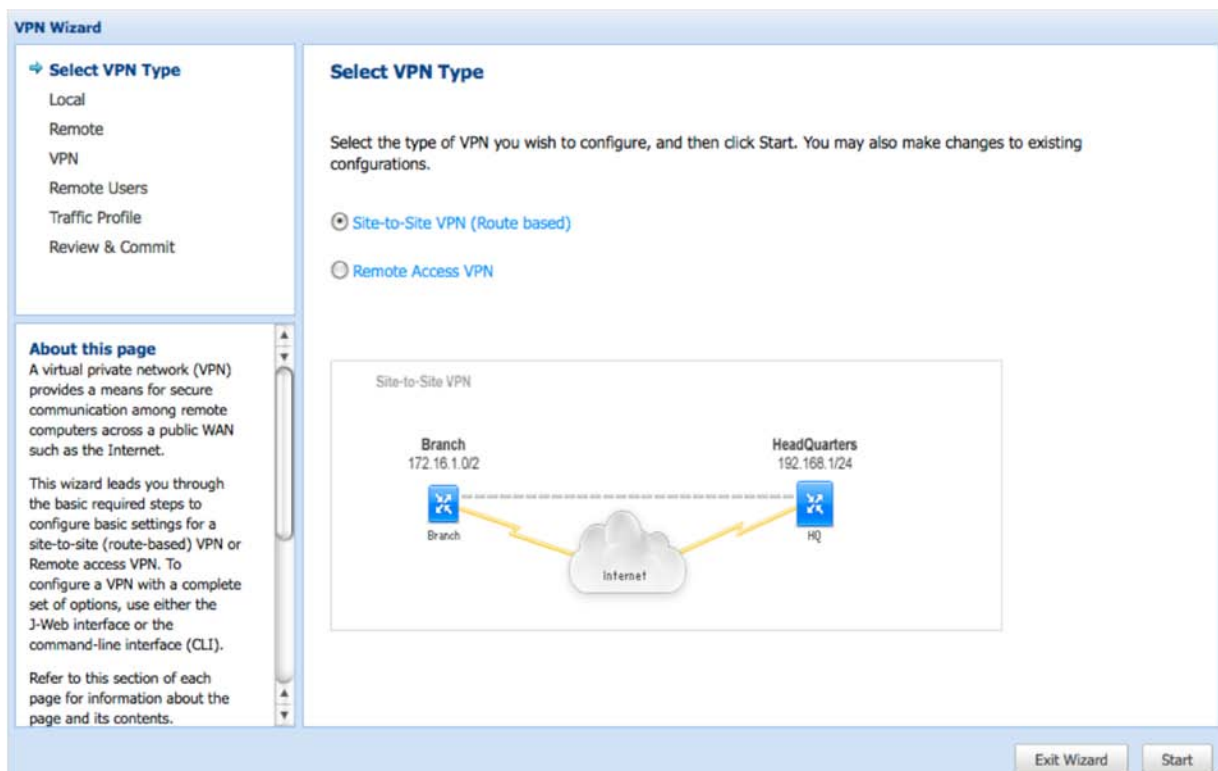


Figure 6. The SRX onboard GUI includes a VPN Configuration Wizard that simplifies the creation of common site-to-site or remote access-based VPNs.

Wireless Support

Wireless LANs are a key enabling technology for mobility, yet many enterprises have handicapped themselves by avoiding Wi-Fi because of security concerns. Securing Wi-Fi and integrating it with firewall policy enforcement can help break down administrative objections and barriers to adoption.

Key Issues and Test Criteria

Securing wireless networks means more than encrypting the channel between the end user and the access point—although that's a basic requirement as well. Users need to be steered to the correct wireless network and securely authenticated. Traffic must be segregated between trusted and guest types of Wi-Fi networks, and security policy and bandwidth controls must be applied to all traffic as it leaves the wireless network. Finally, UTM protections such as anti-malware, intrusion prevention and URL filtering may need specific wireless-focused policies. In many enterprises, wireless intrusion detection systems are also appropriate, searching for unauthorized access points and RF-specific performance issues.

A generation of SMB firewalls has been shipping with integrated wireless access points, providing a very nice package that combines wireless and firewall. Unfortunately, in all but the tiniest of offices, putting the Wi-Fi access point in the firewall doesn't provide very good service since end users are not usually clustered around a rack in a small telephone closet. The idea of tightly linking the access point to the firewall is a good one, but an important part of the linkage is the management and integration, not the physical co-location. Juniper's remotely deployed, centrally managed solution resolves this issue nicely.

While wireless vendors have struggled to one-up each other with feature after feature, the branch wireless deployment has a small number of requirements:

- Multiple SSIDs for each access point to segregate users by class
- Routing of traffic from the access point to the firewall via tunnel or private VLAN so that access controls and UTM protections can be applied
- Support for WPA Enterprise authentication for staff users and captive portal for guests
- Integrated management and monitoring

Other parts of Wi-Fi deployments, such as automatic channel selection, multiple radios per access point, and wireless intrusion detection systems are all desirable, but not strictly necessary in small branch deployments.

Note: This section discusses only Wi-Fi style wireless. Juniper also supports both on-board and external 3G/4G wireless links for wide-area connectivity. Where wireless signal strength at the firewall is strong enough, onboard support is available on the SRX100, SRX110 and SRX210 firewalls through a wide variety of supported USB 3G/4G modems. Juniper also supports an external 3G/4G solution using their CX111 device. The CX111 attaches to the SRX firewall via a standard Power-over-Ethernet port (or can be externally powered with a transformer) and can be placed wherever the wide-area wireless network signal is strongest. In some deployments, the CX111 could be a primary uplink for a small office, but the more common use case is as a backup connection.

Results of Testing

We used Juniper's AX411 access point for our testing which is suitable for branch office deployments with a small number of access points. Juniper also offers controller based wireless equipment, the WL product line, for more sophisticated or larger deployments.

The Wi-Fi access point was connected to our SRX firewall and configured to send management traffic untagged (without a VLAN tag) across the connection to the firewall. Traffic from Wi-Fi users, though, was VLAN tagged by the access point and sent across the same physical interface. We set up two SSIDs, one for trusted users (authenticated using Windows Active Directory credentials using WPA2 Enterprise) and one for guest users (without authentication). Traffic to and from each SSID on a single access point was segregated by VLAN across the same interface.

Because we had previously set up and tested the ability of the SRX firewall to force unauthenticated guest users to a UAC-hosted captive portal, we did not re-test this scenario.

An elegant part of the Juniper SRX firewall solution is the hands-off management of the access points. Pairing the access point with the SRX firewall simply required us to enter the MAC address of the access point into the SRX J-Web management interface. From there, the SRX firewall took over all management of the access point, including collecting logs, so that we never had to log into the access point or learn how to manage yet another device.

Conclusions

Criteria	Notes
Full management of access points	PASS SRX completely manages the APs it is responsible for, including collection of log data.
Creation of secure and unencrypted wireless SSIDs on the same radios	PASS AX411 supports up to 16 SSIDs per radio per access point, although we only tested two. Juniper calls these "Virtual APs."
Segregate traffic by SSID to different VLANs apply per-SSID access policies	PASS Traffic from each SSID is tagged with a VLAN, letting the SRX apply security policies based on SSID.
UTM protections can be applied to wireless traffic before it is sent to the rest of the network	PASS Traffic from the access point can be VLAN-trunked back to the SRX firewall to keep it off the network until after risk mitigation tools such as anti-malware, intrusion prevention and URL filtering.
Enterprise-class wireless security can be applied to the wireless channel	PASS We tested WPA Enterprise, but the access point also supports older encryption protocols and MAC authentication.
Allows physical separation of the AP from the firewall	PASS The AX411 is connected via Ethernet to the SRX firewalls and can be PoE powered by the SRX.

IP Version 6 Support

IP version 6 (IPv6) represents a complete change in the way that TCP/IP networks are designed and protected. Pseudo-security measures, such as NAT, disappear, bringing greater emphasis on correct design of security policies. At the same time, each component in the network and security infrastructure, including firewall UTM services (such as IPS), DNS and DHCP servers, switches and routers, may need to be modified or replaced. While some network managers are attempting to delay the inevitable, smart ones are making sure that current investments will support them in their migration to IPv6.

Key Issues and Test Criteria

IPv6 support in firewalls can be divided into two broad categories: security support and networking support.

Security support includes all of the security services used in the firewall. In UTM devices, this would include IPv6 support in:

- Security policy definitions
- VPN tunnels
- Anti-malware engine
- URL filtering
- Intrusion prevention
- NAT 6-to-4
- QoS/CoS

Networking support for IPv6 includes the network technology inside of the firewall. In a branch office environment, for example, the firewall is the default router for the office subnets. In some cases, the firewall is also the DHCP server or the DNS proxy. This means that the firewall's networking services should ideally include IPv6 support in:

- Interface addressing
- Static and any Dynamic routing protocols
- Local DHCPv6 service and DHCPv6 relay
- SLAAC (stateless automatic address configuration)
- DNS proxy

Not every firewall includes all of these services (such as DNS proxy or advanced dynamic routing protocols), but when these services are present, there is a need for them to fully support both IPv4 and IPv6.

Management and reporting tools may also need IPv6 upgrades to handle the longer addresses used in IPv6. If the reporting tool provides geographic reporting, perhaps based on one of the available IP-to-location databases, then this should also be updated to include IPv6 information.

Results of Testing

Opus One tested the Juniper SRX firewall by installing it as an IPv6 router connected directly to the production IPv6 Internet. We defined IPv6 security policies and applied UTM services to those policies. Then, we re-tested most of the features using pure IPv6.

Our results on the security side of the device were mixed. Most security policies worked well, with the exceptions of the IPS and URL filters, neither of which seemed to catch attacks or URLs that they caught under IPv4.

On the networking side, our results were strong. We found a native DHCP v6 server, as well as full support for the set of SLAAC (stateless IP address assignment) flags needed. The only missing piece was DHCPv6 relay service, used to catch and redirect DHCPv6 queries to a DHCP server on another network. The SRX firewall did not support DNS proxy, a common feature in very low-end SOHO firewalls, under IPv4 so we did not test it under IPv6. Juniper is adding support for DNS proxy in 2013.

Conclusions

Criteria	Notes
IPv6 support in security policy definitions	PASS IPv6 is supported in address objects both in the CLI and in the J-Web GUI.
Application Layer gateway correctly handles IPv6	PASS We tested FTP, the most commonly used ALG.
Anti-malware engine	PASS Catch rate for anti-malware on IPv6 was the same as for IPv4.
URL filtering	FAIL URL filtering is currently not supported for IPv6 URLs.
Intrusion prevention	FAIL IPS is not currently supported for IPv6 in branch SRX but is available in the higher end-SRX platforms.
Interface addressing	PASS Web GUI support for IPv6 is spotty, but the CLI worked well.
Static and dynamic routing protocols	PASS We tested BGP and static routes. Junos also supports OSPFv3 and RIPng.
Local DHCPv6 server	PASS DHCPv6 server was available and tested with Windows and MacOS.
DHCPv6 relay	FAIL DHCPv6 relay is not currently available on branch SRX Devices.
Stateless auto configuration	PASS IPv6 Neighbor Discovery is configured as a routing protocol through the CLI only.

Management

Well-designed and easy-to-use security device management is a critical part of any large deployment. In a branch office environment, where there may be hundreds or thousands of devices, effective management tools are needed to ease the tasks of updating security policies, VPN configurations, software images, and device settings. Regulatory regimes for many industries—and best practices overall—require tight control over firewall configurations and the ability to audit changes and compliance with policy. Reporting is also an important part of firewall deployments, with help desk troubleshooters, IPS managers, and security auditors all dependent on good log management and log analysis tools to support them in their daily tasks.

Key Issues and Testing Criteria

Successful firewall management requires tools at every point in the device lifecycle. Evaluations of these tools will include both objective and subjective criteria. The existence of a tool is easy to verify and is objective, but terms such as “easy to use” are highly subjective.

Many enterprises will choose a multi-vendor approach to their networking and security requirements, which makes support for industry standards such as SNMP very important. For example, monitoring network and security devices for performance and reachability will often be handled by third-party tools across an entire enterprise, rather than management platforms provided by each product vendor. In these common cases, support for external monitoring is a critical requirement.

Centralized management tools that automate configuration changes are often mixed with local device management, even in very large deployments. Some types of configuration, such as VPN definitions in large networks, are impossible to securely manage without automation. Other configuration information, such as basic platform security configuration, is easy to update and keep compliant with organizational policies with central management systems. On the other hand, some tasks such as debugging and policy development are much simpler to complete using local device management, either a CLI or a web-based GUI.

These conflicting requirements suggest that devices should have multiple avenues for configuration, including a CLI,

a web-based GUI and a central management system. Evaluating these tools can be subjective, but there are some objective criteria. For example, central management tools must have templating capabilities to simplify enforcement of organizational policies, and must have VPN tools to handle the coordinated configuration (and reconfiguration) of thousands of tunnels in complex topologies.

Reporting can be as important as configuration in security devices. For example, when IPS is activated on a security device, the management system should support an entire IPS lifecycle, including collection of alerts, correlation and aggregation of data, and direct linkage back to policy definition as alerts are analyzed. Reporting requirements for UTM protections (such as URL filtering) is a significant component of any management system.

Firewall reporting can also be useful for traffic analysis and capacity planning reasons. Better management tools will aggregate and summarize firewall logs to help network managers understand what is happening on their networks and plan for growth.

Results of Testing

The Juniper SRX firewall's management ecosystem is evolving rapidly. Although Juniper has been in the firewall business for many years, the SRX product line is new to Juniper and represents a merging of technology from Juniper's routing product line (Junos) and their firewalls (ScreenOS). Juniper has chosen to start from scratch in many areas, adding security features to their M Series GUI, J-Web, and building an entirely new management system, Security Design, for the SRX firewall. In other areas, such as the CLI configuration and the log aggregation and reporting tool (Security Threat Response Manager, or STRM), Juniper's products are quite mature.

Security Design is an application built on top of Juniper's Junos Space management platform. Junos Space can have more than one device management application (such as Ethernet Design, for managing data center networks), so organizations with other Juniper switch and routing devices can manage all of them through a common base. Opus One did not test Security Design as part of this report due to time limitations. Security Design will be the subject of a future test report.

The on-board CLI of the SRX firewall is a preferred tool by many Juniper engineers for managing the configuration of the firewall. Because the SRX firewall includes most of the features of the Junos routing platform, as well as all of the security features of the firewall, configurations can be long and complex. However, anyone with Junos experience will be comfortable configuring an SRX firewall.

End users who do not have any experience with Juniper switching and routing products will find the SRX CLI to be somewhat overwhelming. Certainly, it will take longer to learn and become proficient at the SRX CLI than the ScreenOS CLI or comparable products from other vendors. However, the SRX CLI offers greater power and features in exchange for its additional complexity. For example, the "commit confirmed" feature lets the network manager push a change to a remote firewall and have the firewall automatically rollback the change if no one logs into the firewall and confirms it. Every engineer who has pushed a change to a firewall, lost connectivity, and had to make an unscheduled visit to a remote site to reestablish connectivity will appreciate "commit confirmed."

The J-Web GUI supports the most common commands used to configure the SRX firewall. We found in our testing that we were able to use the GUI for most security policy, NAT and UTM definitions and other basic device setup. Wizards are available for common complicated tasks, such as VPN (site-to-site and remote access) definitions. However, the J-Web GUI is lacking in some areas and the release we tested had a number of limitations, such as in management of IPS and role-based security policy, and some important gaps, such as high availability feature configurations and in many of the underlying networking features. Network managers who are accustomed to using GUIs for firewall management must be prepared to use the SRX CLI for some configuration and monitoring tasks as Juniper continues to improve upon the J-Web GUI.

STRM is Juniper's log analyzer and security correlation tool (SIEM). STRM accepts log data from all Juniper devices, bringing together an entire network's worth of security information into a single display. Building on technology from Q1 Labs, STRM can support both Juniper and third-party devices. The strength of STRM is in its UTM event management, which is both powerful and easy to learn. However, STRM also has extensive network layer reporting, enabling network managers to quickly understand traffic and application usage passing through SRX firewalls.

Because we focused on STRM, we did not look in detail at the reporting features available in the J-Web GUI. However, J-Web does include a strong set of local reporting capabilities, over 60 different screens to both monitor and report on the status and traffic flowing through the SRX firewall. Of course, the amount of storage within the branch-class SRX firewalls limits the amount of long-term reporting available, but our system came with ample log space, sufficient for two to four weeks of traffic and IPS logs at a branch office with 10 heavy users.

Conclusions

Criteria	Notes
Has on-box web-based GUI	PASS SRX firewall GUI runs in most browsers, but some reporting is done in Flash.
Has on-box CLI	PASS CLI is based on Junos syntax, so will be very familiar to network managers with other Juniper network devices.
Has central management system	PASS Junos Space combined with Security Design can be used for central management ¹² .
Has on-box reporting and monitoring tools	PASS J-Web includes both logging and simple reporting tools to speed debugging and problem resolution.
Has centralized reporting and monitoring tools	PASS Juniper's STRM is a SEIM and log manager with extensive correlation and reporting capabilities.

Junos Routing and Switching

As network managers continue to incorporate security features more tightly throughout their networks, firewalls are expected to more actively participate in network infrastructure. The days of a firewall showing up with three interfaces (trust, untrust and DMZ) are long over. Today, firewalls—even branch devices—must integrate enterprise networking features such as high availability, Gigabit Ethernet, VLANs, link aggregation, QoS tagging and re-tagging, dynamic routing protocols and IPFIX¹³.

Key Issues and Test Criteria

Integrating networking features into a firewall—or integrating a firewall into a networking device—represents a challenge for product developers and designers. Both network and security devices have set the baseline quite high in terms of functionality and power, making a combined product that much more difficult to create.

Conflicting demands in areas such as device performance, default behaviors and even configuration models make the combination a difficult balancing act. For example, many network engineers are accustomed to using a combination of automated tools and CLI configuration to manage, debug and maintain network elements, while their peer security engineers prefer to use GUIs and centralized management tools for firewalls.

Certainly, advanced security products must also have a list of networking features including:

- Interface density for “one-box-branch” environments
- High performance to handle intra-building routing and security (such as backups running across the firewall)
- Network integration features, such as VLAN tagging, IEEE-standard link aggregation, and Layer 2/Layer 3 QoS tagging and re-tagging
- Dynamic routing protocols, including OSPF and BGP (IPv4 and IPv6)

¹² N.B. Junos Space and Security Design were not tested as part of this test plan.

¹³ IPFIX is a reporting protocol that provides summarized and sampled traffic information to an external analysis tool. Other vendors, including Juniper, have built similar proprietary protocols to solve the same problem. IPFIX is the IETF standards effort attempting to unify the many variations and provide a single solution. Juniper's JFlow, included in the SRX Series, is compatible with IPFIX.

In this test, our main focus was on the security elements of the SRX branch firewall, and testing many of these networking features was considered out of scope for this project. In any case, simply listing the networking and security features of the SRX firewall takes an 80 page manual, making even an enumeration of the options out of the question.

Instead, we looked at networking features that were natural complements to the parts of the SRX firewall we were testing.

Results of Testing

Our initial configuration of the SRX240 firewall was in a high-availability cluster. As part of this configuration, we verified correct operation of link aggregation and VLANs. When the SRX branch firewalls¹⁴ are linked into a high-availability pair, the devices form a stack (Juniper uses the term chassis cluster, although the devices are not in one common physical enclosure) and features such as link aggregation can be split across ports on different devices in the stack. This provides an even higher level of redundancy by removing the necessity for additional network elements (such as switches) to handle different device failure scenarios and simplifies the overall configuration of the high-availability solution.

When we tested the SRX firewall's Class of Service enforcement, we were also able to verify the correct operation of some of the many capabilities for QoS detection, enforcement, and marking. The SRX firewall is able to classify packets using Layer 2 (IEEE 802.1p and 802.1q values), Layer 3 (IPv4 DiffServ), and Layer 4 (TCP/IP addresses, ports, and flags) filters. Layer 7 (application) classification, what Juniper calls "AppQoS," is currently only supported on the SRX1400 and above. Support for SRX branch firewalls is planned for 2012. We found that once packets had been classified, we were able to manage the bandwidth using a common set of features (such as guaranteed minimum bandwidth and maximum bandwidth policing). We were also able to rewrite DiffServ QoS values on packets passing through the firewall. More details on our CoS/QoS testing appear above.

We were able to explore a rich part of the SRX firewall operating system called Junoscript when we tested dynamic DNS updating support. Junoscript is a "suite of tools used to automate operational and configuration tasks on network devices running Junos OS."¹⁵ Junos automation lets the network manager write short programs in a simple scripting language which can be brought into play either manually (for example, by being called from the CLI) or automatically (for example, by being bound to "commit" events, log messages or SNMP traps). In this case, we used a Juniper-provided script to handle dynamic DNS updating when a DHCP client on the firewall was assigned an address. Two large libraries of scripts are available, Juniper's Junoscript library and the open source "Scriptorium."



Figure 7. Juniper SRX is an all-in-one device solution providing consolidated networking and security, including firewall, VPN, and UTM features

¹⁴ Virtual Chassis technology is not available in the SRX100 and SRX110, but is in all the other branch SRX devices.

¹⁵ From Junos OS Configuration and Operations Automation Guide, release 12.1, published 13/March/2012.