High Availability Solutions & Technology for NetScreen's Security Systems

Features and Benefits

A White Paper By

NetScreen Technologies Inc.

http://www.netscreen.com

Introduction

Service providers and enterprises that deliver revenue-generating services over the Internet face a myriad of performance and security challenges. However critical those challenges may be, high availability (HA) remains the paramount concern. If the network does not remain available to customers and end-users, then the financial foundation of the company is placed at risk. A well-designed infrastructure security system needs to offer HA tools to create a resilient, scalable, and easy to manage solution.

NetScreen Technologies delivers a line of purpose-built security systems that integrate firewall and VPN functions together with a set of HA tools, all within a single, comprehensive, high-performance platform. This white paper describes the main features and benefits of the NetScreen Redundancy Protocol (NSRP) version 2. NSRP is the name given to the set of protocols, features and tools that NetScreen devices use to achieve High Availability. The NSRP v2 protocol will be supported on NetScreen's security systems: the NetScreen-1000 and the NetScreen-500. The existing NSRP protocol will continue to be supported on the NetScreen-100.

NetScreen's Approach

NetScreen has delivered some of the industry's highest performance firewall/VPN solutions. NetScreen now takes another leadership position by providing one of the industry's most comprehensive solutions for HA and scalable bandwidth within the network's security layer. Within NetScreen ScreenOS, the firmware that powers the NetScreen Security Systems and Appliances, NetScreen has developed a suite of protocols and tools that empower next generation network architects with the best HA and scalability solution available.

Key Points

NetScreen has focused on these three areas of design for NSRP v2:

- **Resilience** – Designed for environments where the goal is maximum uptime on production networks – Even during two, opposite path points of simultaneous failure
- **Scalable Performance** - The clustering of NetScreen systems to secure multiple gigabit throughput
- **Sophisticated, yet easy to use Management** - Changes propagate across the entire cluster configuration and reporting provided by WebUI, CLI, NetScreen-Global PRO, SNMP, Syslog

Resilience

NSRP, combined with redundancy in the surrounding networking devices, enables the security layer in the network to be very resilient, operating with five 9s (99.999%) uptime. This is achieved in the following manner:

Redundancy for Stateful Connections

NSRP ensures the network security function never disappears, regardless of the types of failures that may occur. NSRP was designed to maintain the secure flow of traffic through a redundant network even in the unlikely event of two simultaneous device failures on opposite sides of the network. An example is depicted below in Figure 1.

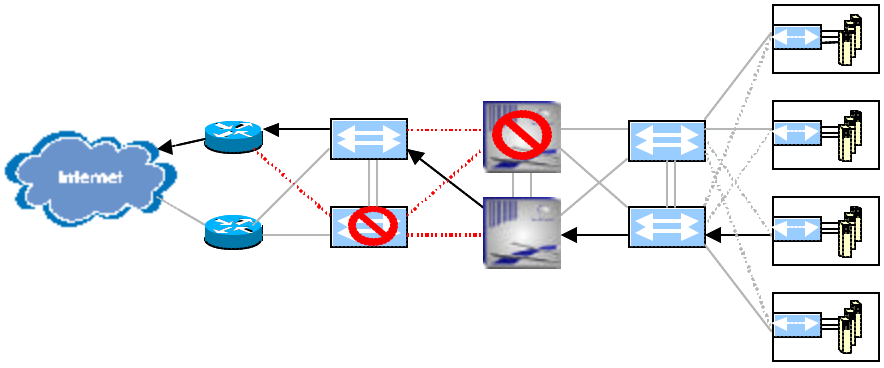

Figure 1: Active Connections Maintained During a Simultaneous, Two Device Failure

In such an event, active connections through the NetScreen device will not be lost because NSRP synchronizes all necessary session state information from one NetScreen device to its Mirror. This includes all established IPSec Security Associations (SAs) and keying material so even site-to-site and remote access VPN connections will be maintained during a failover event.

Policies and configurations are also synchronized across the NSRP Cluster so each device in the cluster knows the job it is to perform should it become a Master.

Leaderless Clustering with No Common Resource

No one element within an NSRP Cluster is a single point of failure. Devices share no common resources, no special switching apparatus, no fixed "brain" of the operation. On the contrary, each individual device holds all that will be needed to fulfill the duty of the cluster on its own. It is true that at any point in time there will be a Master for a given backup group. But that Master has no special characteristics other than being configured with the best preferences and all its Path Monitors are reporting solid health in the surrounding network. That Master can be lost at any point, and another in the cluster will immediately pick up its role and continue to process the active connections appropriately.

Sub-second Failover

NSRP enabled devices can register a failure, react, and the Primary Backup device can commence processing active connections in one second or less. This is achieved for failures where the Master can alert the rest of the NSRP cluster and step down. An example of such a failure is the loss of Ethernet link connectivity on a monitored interface, either due to a cable failure, NetScreen port failure, or adjacent device failure.

Highly configurable Path Monitoring features also allow NetScreen devices to detect failures in the surrounding network and take corrective action immediately. Path Monitors can be set to detect datalink (Layer 2) failures in Ethernet connectivity or network (Layer 3) failures in IP connectivity. Administrators can tune the definitions of "failure" as appropriate for their networks. Once failure thresholds have been triggered, a failover event will occur. Path Monitoring's ability to actively assess the health of the surrounding environment provides a critical first step to minimizing network downtime.

In addition, the master election process in a backup group will pre-elect a Primary Backup. By pre-electing the new master before a failover event actually occurs, the time required to failover is only slightly longer than the time required to detect the master's ineligibility. Detection and recovery can be anywhere from less than one second, if a Master voluntarily relinquishes its role, to just over a second, if a Master fails without notifying the group and the thresholds are set to their minimums. Administrators can further engineer the expected failover time by altering heartbeat intervals and failure thresholds.

Full Mesh Topology, Redundant Physical Paths

Even with sub-second failover a network architect still wants to avoid device-to-device failover whenever possible. NSRP v2, combined with the redundant physical port options on NetScreen-1000 and NetScreen-500 systems, allows network operators to provide redundant physical paths through ingress and egress switching fabrics. When NetScreen devices are connected to switching fabrics using only one link each, a failure in the switching fabric, or the connectivity to it, necessitates a failover between the NetScreen devices, leaving one device operational. Such an example is shown below in Figure 2 and Figure 3.

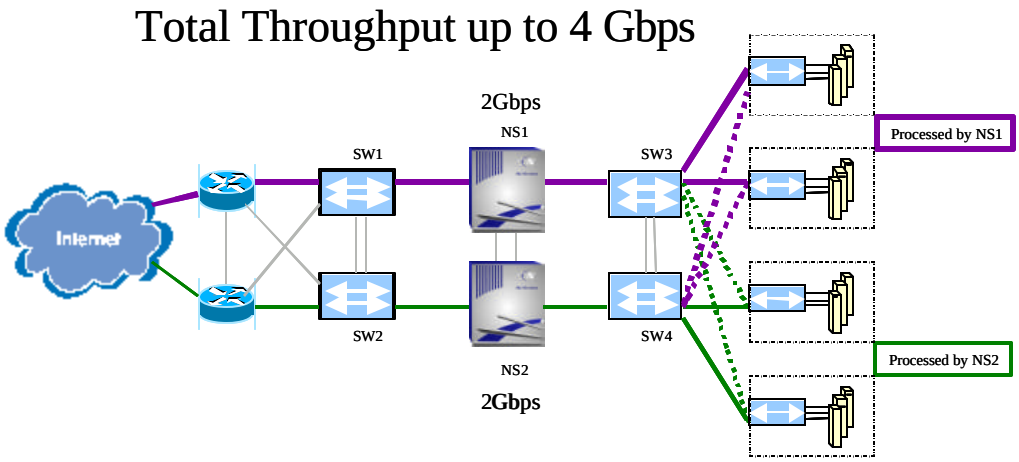

Figure 2: Up to 4 Gbps Total Throughput using Active/Active and Single Links

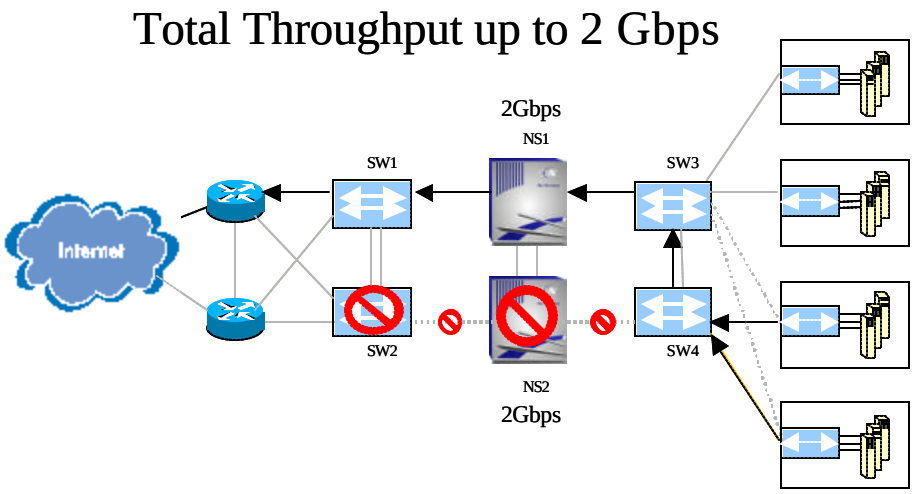

Figure 3: Throughput Halved During Switch Failure in Single Port Topology

Using redundant links in NSRP v2, both NetScreen devices will remain active and functional when a device, cable, or port in the adjacent switching fabrics fails. In Figure 4, SW2 has failed. Using Path Monitoring, the bottom NetScreen device, NS2, detects the loss of Ethernet link on its lower port, the one hosting the connection to SW2. Knowing that its link on the upper port, connecting it to SW1, is operational, it immediately begins using this second port to send and receive traffic associated with its virtual interfaces. By simply choosing the backup port, NS2 remains active in processing its share of the network load, and network performance continues without degradation, as long as SW1 can handle the total bandwidth.

With the ability to use redundant physical interfaces, NSRP ensures all NetScreen devices continue their work of enforcing network security and maintaining optimal throughput even when failures in the surrounding topology occur.

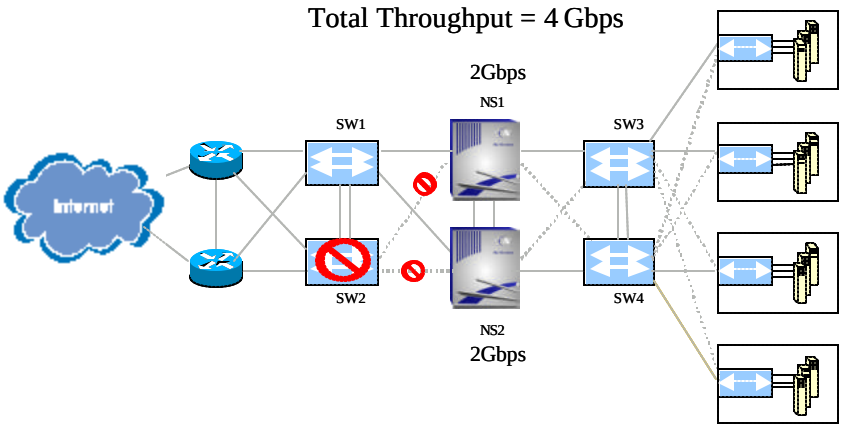

Figure 4. Failure with Dual Physical Ports in a Mesh Topology

With dual physical ports in a mesh topology, if an adjacent link or device fails, the backup link is used and full throughput is maintained.

Scalable Performance

NSRP provides total resilience to the security solution, and does so without sacrificing speed and throughput in high performance, high bandwidth networks. In fact, NSRP's design allows for scalable throughput in the security layer of the network by using clustering features.

Active/Active

In NSRP v2 NetScreen devices may be run in Active/Active load sharing mode. In the previous NSRP implementation each NetScreen device could only participate in one redundancy group at a time. Thus, with a two device group, the device operating as master would be actively processing all the production traffic, while the backup device remained a hot stand-by, poised to take over if needed, but not sharing any of the processing burden. This mode of operation is called Active/Passive (See Figure 5).

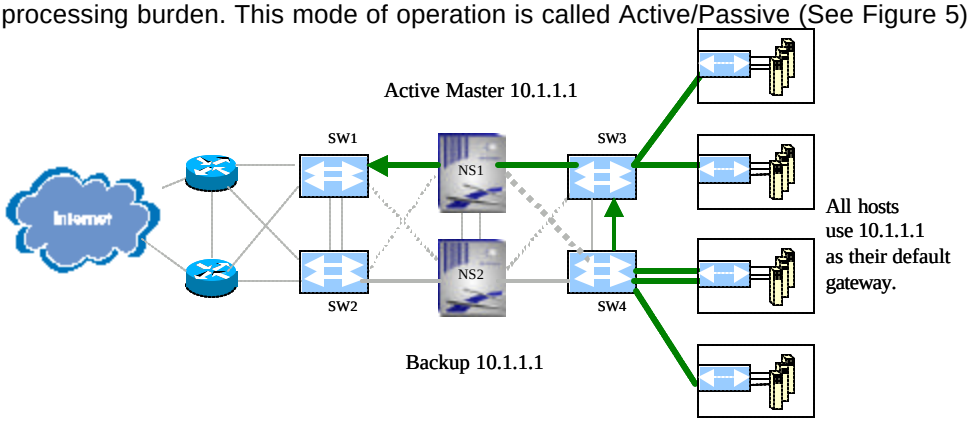

Figure 5: Active/Passive Configuration

Active/Passive configuration protects against all single points of failure, however all network traffic is borne by one device.

NSRP now allows for a single device to operate in multiple backup groups (called a VSD group) simultaneously. Each device may act as a Master in one backup group, while simultaneously serving as a backup in others. The network would then be engineered so that half of the protected hosts use NS1 as their default gateway while the other half use NS2 as their default gateway. In this way, the two devices will both be processing active network traffic load at the same time.

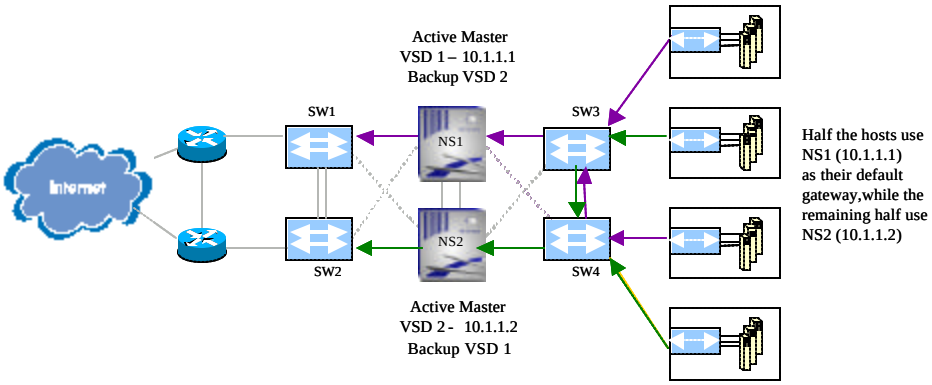

Figure 6: Active/Active Configuration

Active/Active configurations allow for the network traffic load to be shared across multiple devices. Half the devices use NS1 (10.1.1.1) as their default gateway, while the remaining half uses NS2 (10.1.1.2).

Performance Gains

Since NetScreen devices may be run in Active/Active mode the network security layer may now be scaled to achieve linear throughput gains. Adding another NetScreen device and configuring the cluster to operate in Active/Active fashion will lead to a significant bandwidth increase.

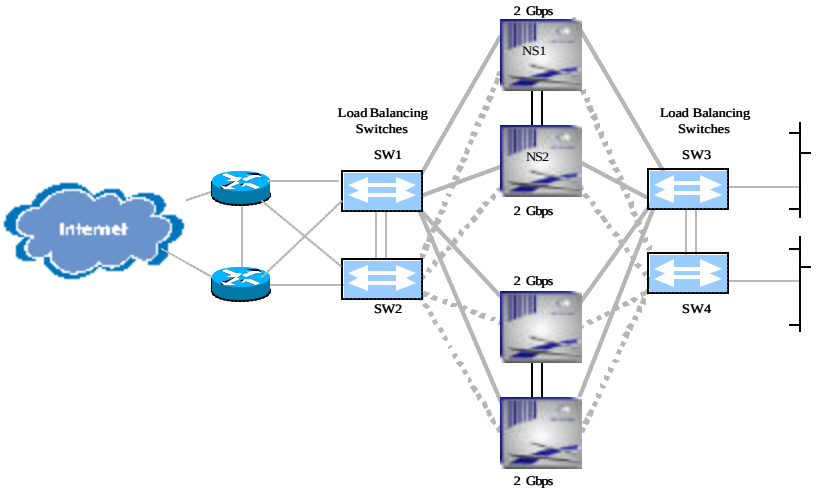

Figure 7: Four NetScreen-1000 Systems Running Active/Active Provides 8 Gbps

Using Active/Active, adding a 3rd and 4th NetScreen device raises the total security capacity from 4 to 8 Gigabits per second. Load balancing switches may be used to ease the effort of distributing traffic across all four devices.

In Active/Passive configuration the maximum bandwidth possible for the security layer is equal to the maximum bandwidth potential of a single device. While NetScreen-1000 devices perform stateful inspection access control at a full 2 Gbps (700 Mbps for the NetScreen-500), many sites still need faster throughput. With Active/Active configuration, each device added offers its full bandwidth potential. As long as the total allowable capacity parameters are not exceeded (for example, the total session count across the Mirror pair does not exceed 500,000, which is the single device maximum for a NetScreen-1000), an additional NetScreen device can be added to a cluster Group and provide significant throughput gains.

Load Sharing & Load Balancing

Load sharing refers to the ability for two devices to somehow split the traffic load so that both are simultaneously contributing to the cause (Active/Active configuration). At any given time the distribution

will most likely not be equal, unless specific engineering pains have been taken to make it equal. The distribution in load sharing may vary over time samples. It may be something like 70/30 or 60/40 or 50/50, depending on the network implementation and use patterns at the time of measurement.

Load balancing, on the other hand, can be more technically advantageous than load sharing. Load balancing aims to distribute the processing load as evenly as possible across all the available devices at all times. If, for example, four NetScreen devices exist in the security layer and process an average of N connections per second, load balancing seeks to constantly maintain an allocation of N/4 concurrent connections per second to each NetScreen device.

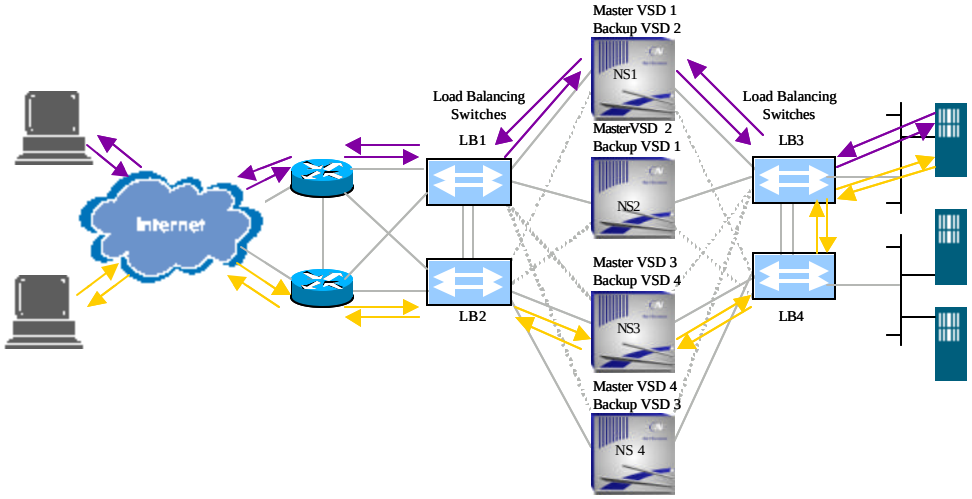

Figure 8: Load Balancing NetScreen Systems

Load balancers "sandwich" a set of four NetScreen devices, evenly distributing the connections between them.

Briefly explained, a set of load balancing switches makes (in this case) four NetScreen devices look like one device to both the outside and inside world. When the load balancers see a new connection they will use a distribution algorithm to decide which NetScreen device will receive the packet. Once the NetScreen device has processed the packet against the configured policy set, if permitted, the packet will be sent into the inside set of load balancing switches. Both the outside and inside load balancing devices will record the source/destination pair of the connection, and the port to which they forwarded the traffic, to ensure that the connection's response packets and subsequent packets follow the same return path. With NSRP v2, security load-balancing is no longer connectionless; the solution can be architected to preserve connections during a failover event.

Sophisticated, yet easy to use Management

High availability architectures add another layer of complexity to the already challenging task of managing network security solutions. The management tools of the products employed should ease the learning curve. Redundancy solutions must also provide instant and useful notifications when the cluster status changes, alerts that simplify an administrator's job of quickly identifying and resolving issues. NSRP contains such ease of use, monitoring and troubleshooting tools.

Easy to Use Interface Options

NSRP configuration and controls are all available from the graphical, point-and-click NetScreen ScreenOS WebUI, an HTML–based management console accessible from any standard browser. Seeing the configuration options laid out in graphical format eases the implementation of the administrator's network design.

However, some administrators work more fluidly on a Command Line Interface (CLI). Since the NetScreen ScreenOS CLI constitutes the foundation of commands from which all configuration scripts are produced, all NSRP options and actions may be accessed from within the CLI. The CLI also provides

several troubleshooting levels that allow administrators to watch all NSRP activity as they occur in real-time.

Secure Remote Management, Reporting and Alarming

NSRP v2 contains a full suite of private MIB objects for retrieving counters using SNMP and detailed messages for export to Syslog. Examples of logged events include counters that increment each time a device transitions to a specific state (like Master or Inoperable), encounters a conflict for Master or Primary Backup status, or fails to receive a Hello message from a partnering device. Together the counters, event messages, and alarms form a powerful set of data for alerting to and resolving issues. As with all NetScreen ScreenOS, any of the above methods of management may be accessed remotely and securely. Any method may be sent through an IPSec tunnel. In addition, the WebUI may be accessed using SSL and the CLI may be accessed using Secure Command Shell, which is ssh v1 compliant.

Summary

Increasingly, successful service providers and large enterprises are looking to leverage the public Internet to provide more value added services to their customers, and to gain more cost-effective business communications. Increased usage has lead to Internet connectivity at speeds of 1 Gbps and greater. Securing any Internet accessible network is essential, and even more so in those operating revenue generating services.

A well designed security solution for these demanding networks is one that provides security, complete redundancy, and scalable performance, all with exceptional ease of management. NetScreen's line of purpose-built, security hardened products deliver unparalleled, single-device performance. With NSRP v2, as part of the NetScreen ScreenOS, NetScreen also delivers the components necessary to build and secure and highly available infrastructure. Redundant links for full-mesh topologies, sub-second and stateful failover, Path Monitoring and a secured control protocol all join to provide complete resilience for the security layer. Scalable performance is delivered with Active/Active load sharing that yields significant efficiencies. SNMP MIBs and event messages, and large scale, secure remote management work together to create a solution that is easy to implement, manage, and monitor.

By providing an integrated solution for VPN and firewall, then tying it all together with the HA tools found in NSRP v2, NetScreen devices are uniquely suited to enable resilient, scalable, easy to manage, and totally secure network infrastructures.