White Paper:

Building High Traffic Firewall Gateways

January 2001

NOKIA

CONNECTING PEOPLE

Abstract 2
Scope 2
Why Clustering? 2
Clustered firewalls 3
Limits of Security State Synchronization 5
Load Distribution with Equal-Cost Multipath Routing 7
Other Methods of Connection Distribution 10
The Symmetric Security Core Architecture 10
Distributing connections with hashing in the presence of NAT 13
The Nokia Network Application Platform and FireWall-1 14

Abstract

Use of the Internet as a communications platform has created demand for reliable, high-performance network systems. For institutions that depend on Internet communications for transaction purposes, network designs must incorporate the twin goals of continuous availability and security. Customers and business partners must always be able to reach information servers, but they must also have assurances that confidential data is protected. Meeting these goals while coping with the explosive traffic increases on the Internet has proven to be a major challenge for network architects.

Internet traffic is increasing at a rate that quickly makes single-box solutions inadequate to meet the demand for secure connections. To meet the demand for both security and throughput, architects have turned to clustered solutions, which combine the capabilities of several independent devices into a single virtual device. Clustered solutions meet the immediate needs for high-throughput security gateways, while protecting the investment in security solutions by enabling currently deployed solutions to scale to the needs of tomorrow. To achieve the maximum benefit from clustered solutions, the individual devices in a cluster must be easy to manage and have a low total cost of ownership. Building clusters often goes hand in hand with the use of appliance-type devices in the cluster to create administrative scalability.

Scope

This document describes a cluster architecture referred to as the "Symmetric Security Core Cluster" within Nokia. It is intended for network engineering staff at sites with requirements for high-reliability, fault-tolerant security systems. It describes the rationale and architecture of a resilient and scalable firewall gateway, but it does not include configuration details for any vendor's equipment. Supplemental white papers, application notes, and configuration guides available from the vendors referenced in this paper can provide hands-on configuration advice.

Why Clustering?

Explosive Internet traffic has been an engine for economic growth and new business opportunities, but the sheer growth in traffic has a dark side for network planners. New users and unprecedented financial opportunities have also created growth in the number of network attackers. Protecting against these ever-increasing attackers, against a backdrop of similarly increasing legitimate traffic, is a treadmill for security officers.

Legitimate traffic is becoming more important from a business perspective. Best-effort packet delivery is a core service, but a well-established one. Networks have grown in value and business importance, both for internal-facing applications as well as external connections. Staff productivity can be boosted by electronic communication and remote access. Exchanging data with business partners has become commonplace, both for supply chain management, marketing and sales efforts, and financial automation. Networks are increasingly carrying both the financial transactions and the information that is the lifeblood of today's business, which makes network outages (or even degraded performance) less and less acceptable. Even planned outages are quickly becoming unacceptable. Global business ties mean that the traditional early morning maintenance window must always interrupt systems during peak hours halfway across the globe.

Copyright @ 2001 Nokia Inc.
All rights reserved.

Clustering has been developed by a variety of vendors for a variety of applications to address these business needs. Distilled to its essence, clustering allows several independent hardware platforms to join together for a common goal as one virtual machine. In addition to processing network traffic in parallel, cluster members share information about the context of that traffic to enable the cluster to survive the failure or degradation of any of its members. By dividing and conquering, clustering can allow several members to work in concert to take on a task that would be beyond any single member. Additionally, strength in numbers allows for easy scalability. As traffic processing needs grow, network administrators can add cluster members to divide the increased load among more devices, ensuring that every device can handle the load assigned to it.

Resiliency and fault tolerance in clusters is based on the statistical improbability of multiple simultaneous failures. On the rare occasions when problems develop with a node, its workload can be transparently redistributed to the surviving cluster members without disrupting communication through the cluster. Transparent workload redistribution also makes maintenance possible. Administrators can perform transparent "rolling upgrades," in which nodes are gracefully removed from the cluster, upgraded, and re-inserted, all without any disruption to end-user operations.

Fault tolerance and transparent upgrades are only half the story. Networks must evolve quickly to handle the increasing traffic placed on modern networks. Yesterday's high-performance single-box solutions quickly become obsolete in such a climate. Using clustered solutions to share the load is the only solution when the demands are beyond the capabilities of any single box. By its nature, clustering also adds scalability. When the cluster is reaching its capacity limitations, additional cluster members can be added to increase throughput.

Clustered firewalls

Headlines throughout the year 2000 reminded the world that Internet security is now everybody's business. Downstream liability for security breaches is becoming an established legal precedent, though businesses can take steps to insure against that risk. Consumers are becoming increasingly aware of the risks posed to personal information and credit card numbers as they traverse the unsecured Internet. Some analysts have suggested that security will become a differentiating factor in the future of e-commerce as consumers vote with their dollars against lax security practices.

In an era of lower traffic loads, corporations could deploy hot-standby solutions in which one active firewall would handle the entire incoming traffic load with a backup firewall ready to assume its functions in case of problems. However, single-box solutions are no longer able to handle the load of many large corporations, let alone server farms and colocation centers.

Two basic approaches to firewall clustering exist. Software running on firewalls may assign workload and keep track of which cluster member handles which connection; or, external switches may be used to assist in the connection distribution process. The two processes are illustrated in Figure 1.

Copyright @ 2001 Nokia Inc.
All rights reserved.

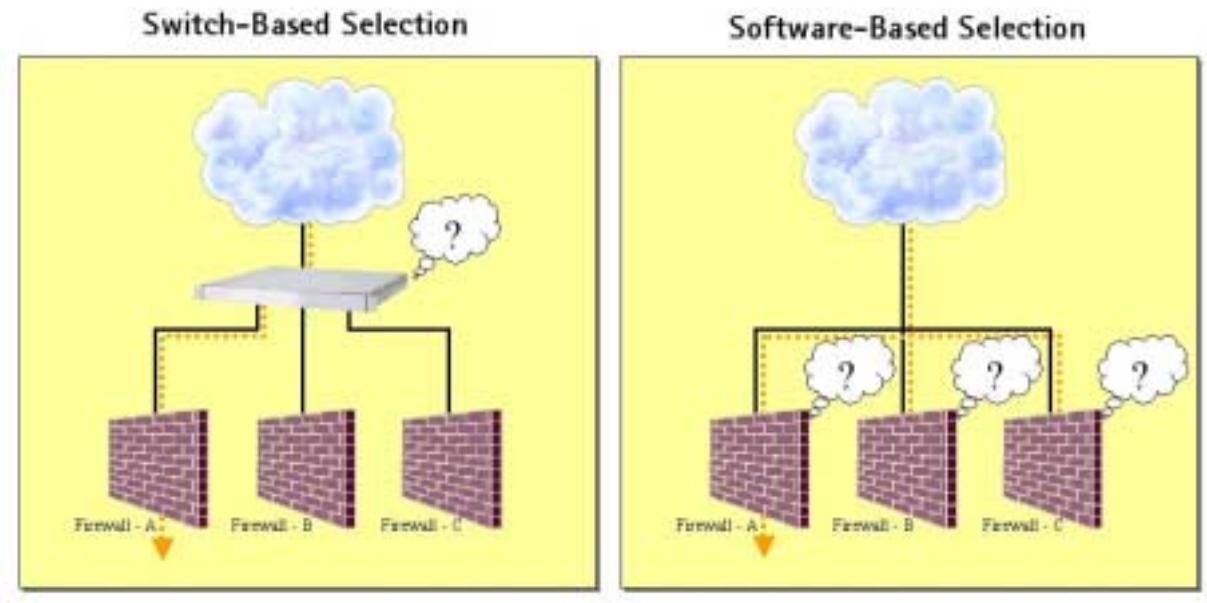

Figure 1: Connection Distribution Methods

For firewall solutions, using external switching elements has several advantages. Switches can trunk several external links which are faster than the wire speed to a single firewall. By distributing connections among several cluster members, the cluster as a whole may process traffic faster than any single member's interface can.[1] The only restriction on switching solutions is that traffic flow through the cluster elements must be symmetric to avoid overloading synchronization processes. Several switch vendors have symmetric forwarding code:

- Alteon (now Nortel Networks)
- Arrowpoint (now Cisco)
- Cabletron
- Cisco[2]
- Foundry[3]
- Nokia
- Riverstone

Additionally, the firewall software and hardware used in a clustered setup must support synchronization.

- Check Point FireWall-1
- Cisco PIX (does not support VPN synchronization)
- NetScreen-100

[1] Distributing all packets to every cluster member makes much greater sense for VPN applications. VPN packets are encrypted, so the cost of processing a packet is quite expensive compared to the cost of pulling the packet out of the interface card buffer to examine it. Firewalls present a less attractive trade off. Examining a packet to see if it is worth full security inspection can be almost as expensive as the pre-inspection check, which makes the process far less attractive.
[2] Cisco has announced a Route Switch Module for the Catalyst 6500 which incorporates symmetric flows.
[3] Foundry offers a novel "least connection" approach well-suited to NAT and VPN situations.

Copyright @ 2001 Nokia Inc.
All rights reserved.

Limits of Security State Synchronization

Network security devices provide security by filtering unacceptable traffic before it reaches the internal network. Advanced packet filters store a security context for each connection, with high availability provided by sharing the security state between two devices. Incoming connections are placed into the state table of one firewall. Synchronization processes monitor the state table for new entries and send any new connection information to peer firewalls.

Synchronization processes fail when connections take asymmetric paths through a cluster. In an asymmetric routing condition, inbound security and outbound security are handled by two separate firewalls. In many asymmetric routing setups, the state synchronization process is not fast enough to maintain a consistent security state between multiple firewalls.

A fuller description of this problem is contained in the following example. Given the network topology in Figure 2, inbound connections to *Web* may be dropped frequently. When *Client* initiates a connection to *Web*, it will send out TCP SYN messages that enter through Firewall A. If permitted by the security policy, the connection from *Client* to *Web* will be placed in A's state table and allowed to pass through the security perimeter to *Web*. If *Web* has a default route to Firewall B, then problems occur. On a server farm, the network between the firewalls and *Web* is fast, so moving the TCP SYN/ACK reply packet between *Web* and one of the firewalls may take only one millisecond, giving the synchronization process perhaps 5 ms to send the state information from A to B.

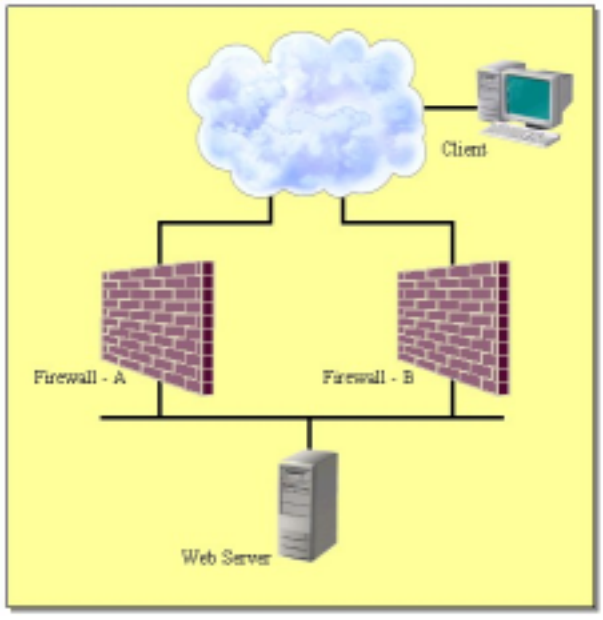

Figure 2: Asymmetric Routing

The SYN/ACK reply from *Web* to *Client* will be dropped because B does not yet have the state information for this connection. After an implementation specific time, typically a few seconds, *Client* will resend the SYN, which will be accepted again. Because the second SYN is identical to the first SYN, the state information about the connection sent by A after the first SYN allows B to accept the connection. Multi-second delays are introduced for each connection, which can make a modern, graphically-rich Web page painfully slow.

Copyright @ 2001 Nokia Inc.
All rights reserved.

Many solutions use a dedicated link for synchronization messages. Each message must uniquely identify a connection for use in subsequent filtering. To estimate the maximum possible number of connections on a synchronization link, consider the following types of connections:

- Plain Connection: IP protocol (1) + source IP address (4) + destination IP address (4) + source TCP/UDP port (2) + destination TCP/UDP port (2) = 13 bytes
- Sequence number-tracked connection: Plain Connection + forward sequence number (4) + reverse sequence number (4) = 21 bytes
- Address-translated connection: Sequence number-tracked connection + translated source IP address (4) + translated destination IP address (4) + translated source port (2) + translated destination port (2) + new forward and reverse sequence numbers (16) = 49 bytes

To determine the maximum number of messages possible, simply divide the connection speed by the size of the record used by the synchronization process.[4] Real-world implementations will always have a lower message throughput because of link-layer headers and framing, network-layer packet headers, and the overhead of the synchronization process itself. Table 1 shows the calculation for several common link types which are used for synchronization in common products.

Table 1: Maximum Synchronization Message Throughput on Common Link Types

		Maximum Possible Synchronization Messages			
Type	Size (bytes)	Typical Serial (57,600 bps)	Fast Serial (115,200 bps)	Ethernet (10 Mbps)	Fast Ethernet (100 Mbps)
Plain Connection	13	554	1,108	96,154	961,538
Sequence-tracked Connection	21	343	686	59,524	595,238
Address Translated Connection	49	147	294	25,510	255,102
"Real-world" implementation (Check Point FireWall-1)	70	103	206	17,857	178,571

Table 1 shows that synchronization must be done over fast network links to meet the throughput needs of even reasonably busy sites.

> **Conclusion:** State synchronization is an excellent tool to survive failure of a firewall, but it is not a real-time process.
>
> **Design Goal 1:** Avoiding asymmetric routes through firewalls is of paramount importance to assure maximum throughput.

Further problems with state synchronization may arise when network address translation (NAT) is deployed. A guiding principle of IP was end-to-end communication: intermediate devices need only move packets based on addresses, not based on the context of a particular packet. NAT places an additional burden on

[4] Real-world stateful packet filtering implementations will always have larger state entries than the theoretical sizes calculated above. Several reasons exist for this. Recordkeeping and identification of connection type add storage overhead. Underlying memory allocation routines may work much more quickly with memory chunks which are an even power of two bits in size. Most importantly, speed is of the essence when processing new connections and comparing inbound packets to the state information. Maximum acceleration is obtained when a hash table enables very fast table look-ups. As an example of real-world sizes, Check Point's Stateful Inspection engine uses a plain connection entry of approximately 70 bytes; NAT entries are approximately 200 bytes.

Copyright @ 2001 Nokia Inc.
All rights reserved.

synchronization because the security context and the address translation context must be synchronized between multiple devices.

Consider the case in Figure 2 where Firewall A and Firewall B are performing address translation for a protected network such as a Web farm. Most firewalls can be configured to pass the "unsolicited" SYN/ACK from *Web*. When Firewall B must perform address translation, however, there is no mechanism other than state synchronization for transmitting the appropriate address mapping information from Firewall A to Firewall B. The address translation applied by Firewall B must reverse the address translation applied by Firewall A or *Client* will not accept the return packets.

Exterior routing protocols such as BGP provide only crude controls over inbound traffic. It is not uncommon for customer networks connected to multiple ISPs to observe packets from a connection exit through one ISP link and return via a different ISP link. BGP provides only rough tools to influence the flow of inbound traffic, and a service provider's policy may override any configuration that attempts to influence your inbound traffic pattern.[5] Symmetry can be insured if one link is used strictly as a backup when a primary link fails, but then transmission capacity is paid for but not used. Symmetry and load sharing with multiple ISPs are opposing goals and must be balanced according to each specific site's goals. Interior routing does not pose the same threat to symmetry as exterior routing. Interior routing in many cases will be a default route to the Internet gateway (or the nearest Internet gateway on global networks).

> **Conclusion:** In the absence of failure, internal routing and external routing must select the same cluster member to handle a connection. Left to their own devices, there is no guarantee that connections will be symmetric.
>
> **Design Goal 2:** Routing must be carefully designed to ensure symmetry.

Load Distribution with Equal-Cost Multipath Routing

Modern routers support using multiple next hops with equal IGP cost to improve throughput. Equal-cost multipath routing can be a powerful tool for load distribution, especially if the next hop selection algorithm was designed for traffic symmetry. In Figure 3, a pair of redundant, synchronized firewalls protects a single network, which uses a router to connect to the outside world.

[5] BGP's multi-exit discriminator (MED) attribute is the best tool for influencing inbound traffic. However, in the absence of contractual provisions (and payment), ISPs will ignore MEDs. An enterprise has very little control over the Internet routing. If a network is advertised to an ISP link, traffic to that network will be received on that ISP link.

Copyright @ 2001 Nokia Inc.
All rights reserved.

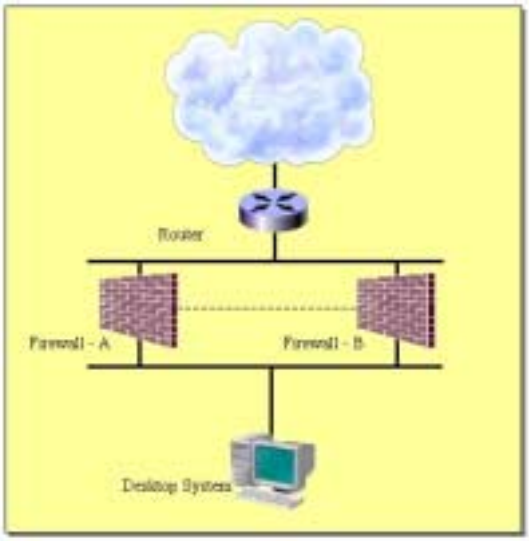

Figure 3: A Topology to Illustrate Problematic Equal-Cost Multipath Routing

The router may use both Firewall A and Firewall B as next hops to the protected network. If Firewall A and Firewall B both are configured with equal costs to the protected network, then the router will attempt to share the load between the two next hops.

Current Internet best practices strongly recommend using all equal cost next hops to a destination, but no standard exists for sharing the load among a set of equal cost next hops. Different vendors have taken different approaches to path splitting. The most common are:

1. Round robin: the router sends a packet to the first next hop, then the second, then the first again, and so on. Round robin load sharing is guaranteed to create asymmetric paths for every other packet when used with a synchronized firewall setup, and should not be used.[6]
2. Source hash: the router takes the source IP address of the packet, runs a hashing algorithm on it, and uses that next hop exclusively for that source IP address.
3. Destination hash: the same as source hash, but using the destination IP address.
4. Source/destination hash: the same as source hash, but using both the source and destination IP addresses.

Using the wrong type of equal-cost path splitting can be devastating to throughput, as in the case of round robin forwarding. Say that the router in Figure 3 is configured for round robin next hop selection and four ICMP echo requests to the PC arrive from the Internet, and that the round-robin algorithm will start by sending the first ICMP packet to Firewall A. Also assume the PC's default route is to Firewall B.

1. Firewall A accepts the packet and sends it to the PC, but because the PC's default route is to Firewall B, the ICMP exchange is asymmetric. If the protected network is an Ethernet, the PC's reply will reach Firewall B too quickly for state synchronization to have updated Firewall B's security information, and the ICMP echo reply will be dropped.
2. The router sends the second ICMP echo request to Firewall B. When the PC sends the ICMP echo reply, it will go to Firewall B. Firewall B has a record of the ICMP echo request and will allow the reply through.

[6] Based on the TCP acknowledgements sent by the PC back to the source, the source will infer packet loss due to congestion and throttle back the data rate, which will severely reduce throughput. Sequential packet delivery is a major goal for any modern network design.

Copyright @ 2001 Nokia Inc.
All rights reserved.

3. The third ICMP echo request will be treated as the first and will be dropped due to asymmetry.
4. The fourth ICMP echo will be treated like the second and will successfully return.

The sender of the four ICMP echo requests will see an alternating pattern of success and failure.[7] In this scenario, equal-cost multipath causes traffic to be dropped. Rather than distribute the load between the two firewalls in an intelligent manner that avoids dropped traffic, round-robin equal-cost multipath can cause congestion and network performance problems by preventing symmetry.

To achieve traffic symmetry, a second router must be added, as in Figure 4. Both routers in the figure are capable of equal-cost path splitting; that alone, however, does not alleviate the asymmetry. When choosing a next hop, the algorithm must be *deterministic* so the flow path will be *symmetric*.[8]

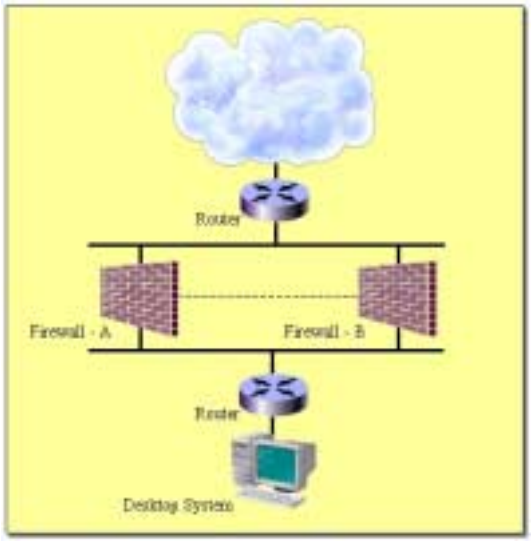

Figure 4: Using A Second Router to Assure Symmetry

Determinism means that when given a hash input, the router will always select the same next hop to forward packets to.[9] If, for example, a source IP address of 10.3.23.47 and a destination IP address of 192.168.93.26 result in using the second next hop in the list, any packets going from 10.3.23.47 to 192.168.93.26 taking the second next hop. In the context of Figure 4, this means that packets from 10.3.23.47 to 192.168.93.26 are directed towards Firewall B.

Symmetry is maintained by both the determinism of the hash function and the network design. When 192.168.93.26 replies to 10.3.23.47, the bottom router will receive the packet and calculate a hash value to select its next hop for the return traffic. A symmetric hash algorithm will also select the second next hop.

---

[7] Note that synchronization of security information makes the problem worse. Round robin forwarding in the absence of state synchronization could result in out-of-order delivery, but the use of synchronized firewalls guarantees that every other packet will be dropped.

[8] This paper uses the term "flow" rather than the term "connection" to describe a sequence of packets that are logically associated with each other. TCP connections are flows, but other types of exchanges are also flows. For example, the ICMP echo request/echo response transaction or DNS queries are also flows.

[9] Not all hashing algorithms are deterministic. Cisco IOS calculates a hash value, but does not select a deterministic next hop from that hashing value. Given the same pair of addresses, Cisco routers will calculate the same hash value, but not necessarily select the same next hop.

Copyright @ 2001 Nokia Inc.
All rights reserved.

Because hash results are used to select next hops, both Firewall A and Firewall B must be configured consistently by giving Firewall A lower IP addresses than Firewall B on each network.

A deterministic source/destination hashing algorithm guarantees that all traffic between a particular source and a particular destination will use the same firewall. When flows are symmetrically distributed through firewalls, state synchronization is only needed in the comparatively rare case of a firewall failure. NAT may complicate the next hop selection process. For the common case of translating either the source or destination address, two complementary mechanisms are required. Not all vendors support a hashing algorithm with the required properties: Nokia, Cabletron, and Foundry are known to.[10]

Other Methods of Connection Distribution

Equal-cost multipath routing is by no means the only method of distributing connections to multiple firewalls. To cope with server farms of all purposes, including firewall farms, switch vendors have developed several features to distribute connections to multiple servers.

Path monitoring: Some switches allow administrators to define several paths through multiple firewalls, and connections are distributed between the available paths. Health check features can be used to verify functionality beyond simply detecting the presence of an Ethernet link. Path monitoring has one important advantage over simple hash-based distribution. When a firewall fails, all connections will be redistributed based on hash values, and it is likely that every connection will be reassigned. Path-monitoring switches need only distribute the connections assigned to the failed path.

Least connections/weighted least connections: Some switches maintain information on connections that are assigned to each firewall. New connections are assigned to the firewall handling the fewest connections. Variants of this approach allow administrator-assigned weights to allow administrators to build clusters of different capacity.

Response time: Some switches poll cluster members for a response and distribute new connections to the firewall with the fastest response time.

The Symmetric Security Core Architecture

In Figure 4, routers are used to assure traffic symmetry through the firewalls. While modern routers are quite reliable, they may occasionally fail. Enhancing Figure 4 with redundant pairs of routers plus redundant ISP links leads to a significantly more resilient solution, named the Symmetric Security Core, as shown in Figure 5.

[10] Cisco has announced a symmetric flow feature for the high end Catalyst switches with RSMs as this paper was written, but its availability date is unknown.

Copyright @ 2001 Nokia Inc.
All rights reserved.

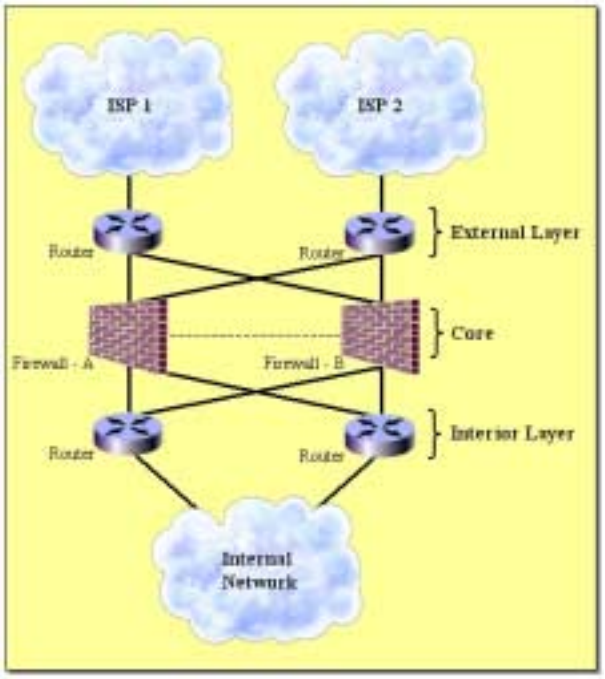

Figure 5: Symmetric Security Core Architecture

Security is provided by the firewall core at the center of the diagram. Load sharing and symmetry arise from the use of equal-cost multipath routing in the interior and exterior layers. Routing information to other networks is maintained by the interior and exterior layers, each of which is free to make forwarding decisions independently.

The *core layer* is responsible for providing resilient network security. By selecting the appropriate hardware platform, the core can be sized according to throughput, concurrent connection and fault tolerance requirements.

IP addressing in the core layer: The main objective in addressing the core is to provide an IP addressing scheme that guarantees symmetry. As discussed earlier, a deterministic hashing algorithm will always pick the same next hop for packets belonging to the same flow between a given source and destination. The goal in assigning IP addresses is to make sure that the exterior layer's next hop for a flow will be the interior layer's next hop for that flow.

Next hops are sorted by IP addresses. The key, therefore, is to make sure the core firewalls are numbered in the same order on the interface to both the interior and exterior layers. If crossover Ethernet cables are used between the routers in the external layers and the core firewall devices, a number of 30-bit networks will be used. One sample numbering is shown in Figure 6. Very simplified routing tables for the four exterior routers are shown in the following table.

Copyright @ 2001 Nokia Inc.
All rights reserved.

Table 2: IP Addressing in the Core to Assure Symmetry

Upper Left Router next hop to internal network	Upper Right Router next hop to internal network
.2 (left firewall), .10 (right firewall)	.6 (left firewall), .14 (right firewall)
Lower Left Router default route	Lower Right Router default route
.18 (left firewall), .26 (right firewall)	.22 (left firewall), .30 (right firewall)

If a deterministic hashing algorithm is used, this addressing scheme will ensure symmetry. The left firewall is always numbered lower than the right firewall, so the first next hop always corresponds to the left firewall. Likewise, the right firewall is always numbered higher than the left firewall, so the second next hop always corresponds to the right firewall.

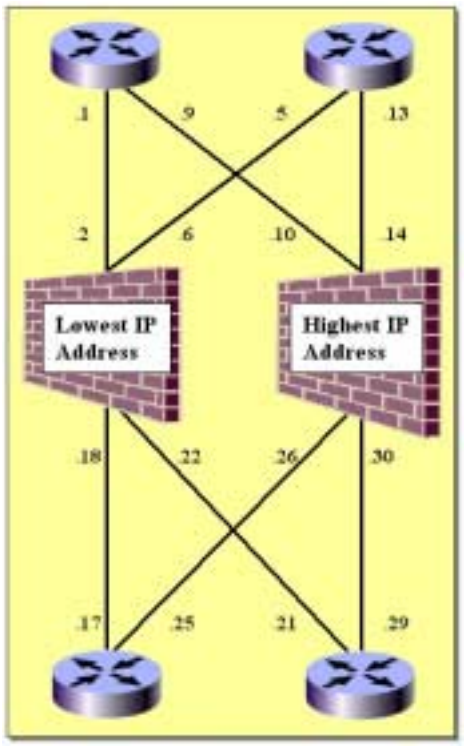

Figure 6: IP Addressing in the Core Layer

Equal-cost path splitting indirectly provides scalability because path splitting can be used to divide the load between several firewalls. Although Figure 6 shows only two firewalls, most equal-cost path splitting algorithms support more than two next hops.

Routing in the core layer: For the simplest core layer routing setup, the core layer uses static routing. Some operating systems provide the ability to have multiple prioritized static routes so that if the interface to primary static route to a destination goes down, then the system will start using a secondary static route. If the secondary fails, a tertiary route will be used.

In the topology shown in Figure 6, backup static routing could be configured on the firewalls to provide resilient routing to both the external world and the internal networks, as described in the following table.

Copyright @ 2001 Nokia Inc.
All rights reserved.

Table 3: Prioritized Static Routing in the Core

	Left Firewall	Right Firewall
Default route (the Internet)	.1 as primary, .5 as backup	.13 as primary, .9 as backup
Internal networks	.17 as primary, .21 as backup	.29 as primary, .25 as backup

Many operating systems, including the IPSO operating system on the Nokia network appliances, support route dampening. If the upper left router were to fail, the left firewall would immediately begin to use the upper right router as its default route. When the upper left router was restored, the Ethernet link integrity would return, but the left firewall would wait for an adminstratively-specified period before using the upper left router as its default router again.

There is no reason to run BGP in the core layer. Exit point selection is done by the external layer, so the core need only deliver packets to the external layer. Resources that would used by BGP to carry full Internet routing information would be better used by applications in the core layer, especially because the exterior layer will already be carrying full Internet routing information. If the number of internal networks is large, however, the core layer may instead use OSPF to learn about the internal networks from the interior layer.

The *exterior layer* is composed of at least two routers that provide the interface to the WAN and assure that inbound packets are distributed among the group of core firewalls symmetrically. If multiple ISPs are used, the exterior layer is responsible for selecting the appropriate exit ISP link through a protocol such as BGP.

The *interior layer* accepts traffic from the interior network and distributes it into the core layer. Like the exterior layer, it does not apply any security to the traffic. The interior layer follows a default route to the core. Each router in the interior layer should have multiple next hop default routes to each of the core firewalls, using the symmetric source/destination hash algorithm to select the particular core firewall. Like the core layer, multiple prioritized static routes can be used, as in the following table.

Table 4: Prioritized Default Routing in the Core

	Lower Left Router	Lower Right Router
Default route	.18 primary, .26 backup	.30 primary, .22 backup

If the internal network follows a default route to the interior layer, then the interior layer should use a protocol such as VRRP to present a resilient gateway address to the internal network routers. The interior layers would then need to be programmed with static routes to the internal networks. Alternatively, the interior layer can run OSPF and interact fully with the internal routing architecture.

Distributing connections with hashing in the presence of NAT

Source/destination hash cannot be used with NAT because the source address changes. In Figure 7, the NAT policy is to take the internal address *z* and translate it to *x*.

Copyright @ 2001 Nokia Inc.
All rights reserved.

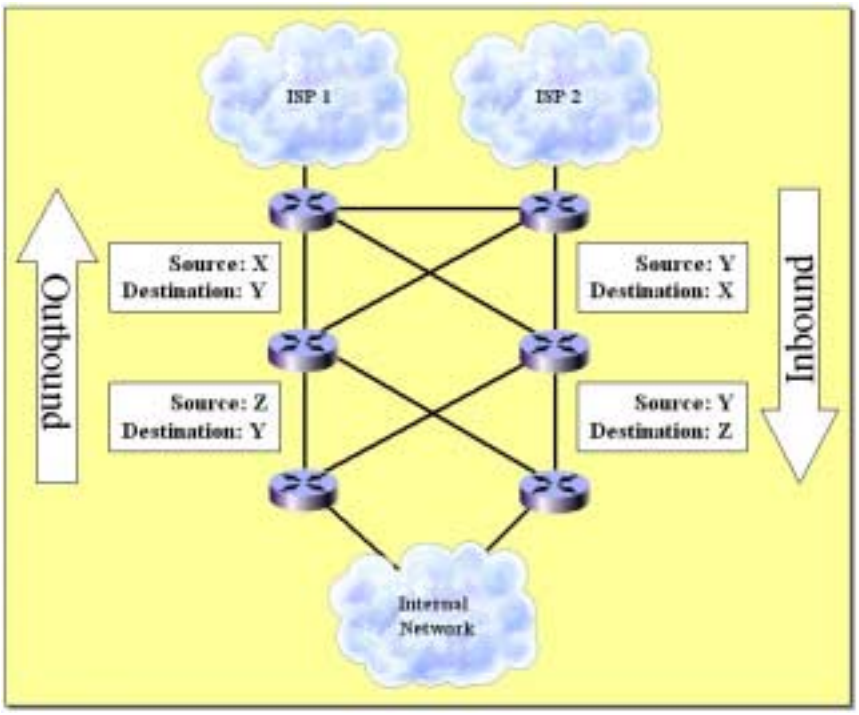

Figure 7: Deploying Address Translation for Symmetric Path Splitting

To use path splitting to achieve load balancing and symmetry, the input to the hashing algorithm must be identical on the interior layer and the exterior layer. In most address translation setups, this can be accomplished by using the address of the Internet host as the input to the hashing algorithm. On the interior layer, use the destination address as the hash input; the exterior layer should use the source address as the hash input. If the source hashing and the destination hashing algorithm are symmetric, then equal cost path splitting can be used for load distribution.

The Nokia Network Application Platform and FireWall-1

The Nokia family of Network Application Platforms offers an ideal combination of features for use in high-traffic firewall gateways. Switch-based firewall gateways are complex and require significant investment from the end-user, especially in terms of staff time. Using firewall appliances dramatically decreases the total cost of ownership by providing significant benefits in the areas of integration, scalability, and manageability.[11]

The value of the Nokia platform starts with the IPSO operating system. IPSO is designed for route processing and packet forwarding, which make it ideal for use within a firewall cluster. IPSO-based platforms can be configured with prioritized static routes or run OSPF to provide intra-cluster routing, while route dampening and link recognition delays can avoid race conditions which may be problematic in clustered environments.

IPSO is designed for use within demanding network security environments. IPSO is pre-hardened and pre-loaded with leading applications to meet the needs of administrators trying to handle rapidly growing

[11] A detailed study of the total cost of ownership of Nokia appliances is available from the Nokia Web site at http://www.nokia.com. Go to **Secure Network Solutions**, and click **Solutions**, **Main Benefits**.

Copyright @ 2001 Nokia Inc.
All rights reserved.

network traffic. Check Point's award-winning FireWall-1 package provides access control, authentication, and encryption capabilities well-suited to a wide range of network environments. Nokia's partnership with Check Point brings together expertise in high-speed networking and security, which has led to the development of new secure packet forwarding features like firewall flows. Firewall flows provides increased performance by integrating Check Point's security technology with the IPSO routing kernel to drastically reduce the trade-off between security and performance.

The Nokia appliance reduces total cost of ownership by reducing the support and administration costs of the joint solution. The Nokia Voyager interface and the upcoming Nokia Horizon Manager make it possible to administer large numbers of appliances in parallel. For relatively simple installations, the Check Point GUI interface enables administrators to quickly and easily develop and modify security policies. More demanding sites with very large numbers of firewalls and complex security policy requirements can build on the Check Point GUI by using the industrial-strength Provider-1 management solution. The value of the Nokia platform starts with the IPSO operating system. IPSO is designed for route processing and packet forwarding, which make it ideal for use within a firewall cluster. IPSO-based platforms can be configured with prioritized static routes or run OSPF to provide intra-cluster routing, while route dampening and link recognition delays can avoid race conditions which may be problematic in clustered environments.

Copyright @ 2001 Nokia Inc.
All rights reserved.

About Nokia

Nokia is the world leader in mobile communications. Backed by its experience, innovation, user-friendliness and secure solutions, the company has become the leading supplier of mobile phones and a leading supplier of mobile, fixed and IP networks. By adding mobility to the Internet Nokia creates new opportunities for companies and further enriches the daily lives of people. Nokia is one of the most broadly held companies in the world with listings on six major exchanges.

About Nokia Internet Communications

Nokia Internet Communications, headquartered in Mountain View, California, provides world-class Network Security, Virtual Private Network and Wireless Software solutions that ensure the security and reliability of corporate enterprise and managed service provider networks. Nokia is committed to enhancing the end user experience by bringing a new level of security and reliability to the network, enabling an Internet transaction that is personal and trusted—each and every time.

Nokia Internet Communications
313 Fairchild Drive
Mountain View, CA 94043
Tel: 1 877 997-9199
E-mail: internet.na@nokia.com

www.nokia.com

NOKIA
CONNECTING PEOPLE

Copyright © Nokia Internet Communications Inc. 2001. All rights reserved. Nokia and Nokia Connecting People are registered trademarks of Nokia Corporation. Nokia operates a policy of continuous development. Therefore we reserve the right to make changes and improvements to any of the products described in this document without prior notice. Under no circumstances shall Nokia be responsible for any loss of data or income or any direct, special, incidental, consequential or indirect damages howsoever caused.

Printed in Canada NI3018040