

White Paper

The Components of a Strong Security Stance

Understanding Risk and the Requirements for a Secure Network

Sarah Sorensen
Product Marketing Manager



Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
408 745 2000 or 888 JUNIPER
www.juniper.net

Part Number: 200115-001 Apr 2005

Contents

Introduction.....	3
Basic Security Stance - Achieving an “Acceptable Level” of Risk	4
The Value of the Network	4
Network Threats.....	5
Understanding the Network’s Vulnerabilities.....	5
Security is a Process	6
Analysis.....	6
The Importance of Network Awareness.....	7
Mitigation	7
Where to Start – Understanding How Traffic Flows Through the Network	8
Management.....	10
Requirements for a Strong Security Stance.....	10
Network Awareness for Accurate Analysis.....	10
Security Enforcement throughout the Network.....	12
Security for Endpoint Clients.....	12
Security for the Network	14
Security for Endpoint Servers	15
Security Management.....	16
The Role of Policies	16
Logging, Auditing and Reporting for Security Event Resolution	17
Discovery, Feedback and Adaptive Controls	18
The Juniper Networks Approach – Conclusion	19

Introduction

Over the past decade organizations have seen a disappearance of the "trusted network." Initially, organizations worked off the general premise that internal (coming from the corporate LAN) meant trusted and external (everything outside the corporate LAN) meant untrusted. If a user/device was deemed part of the trusted network, there were very few restrictions to their access to information and application resources. Now, however, an organization needs to re-evaluate how trust is granted. No longer can security be a binary decision of allow/deny; rather there is a new security paradigm that requires a multi-level approach that incorporates different levels of trust.

This security evolution is driven by the ever-expanding requirements of the business. To remain competitive and increase productivity, organizations are extending communications and business tools to a much broader range of people. In addition to full time employees, the network must efficiently and securely support customers, contractors, business partners, etc. and the increasing types of applications and services these different people require. In addition, the network must support all the different devices used to communicate on the network (e.g. laptop, wireless device, kiosk, etc.) and the communications must be enabled on a global scale.

Because of the increasing ways in which the businesses are interconnected to facilitate business operations for such diverse constituencies, organizations are facing increasing threats to the information flowing through the network and business systems. Organizations want to provide a good experience and provide each user with just the information and applications they need to conduct business.

Organizations need to create a strong security stance, which is the overall security level of the network, to ensure the network is not exposed to undue risk. Ultimately, a strong security stance does not equate to shutting the network down, rather it is about appropriately controlling the use of the network, based on the user's relationship to the business, and minimizing any threats. The security stance should support the business priorities of an organization and effectively map those priorities to the appropriate security measures throughout the network. These measures include a combination of policies, processes and technologies that appropriately prevent access and control communications.

To create a security stance that meets the needs of the organization, the risks to the network must first be understood. Once identified, an organization must figure out how best to control the appropriate level of use for resources (information, applications) on the network, as well as mitigate the risks associated with using those resources. Ultimately the goal is to create a security stance that effectively protects the integrity of the network and the information flowing through it.

This paper describes, at a high-level, how organizations can start to establish their security stance. It includes discussions on risk assessment, the security process and the considerations organizations should make to mitigate risks throughout the network. It concludes with a brief synopsis of some of the ways in which Juniper Networks can help organizations enforce their security stance.

Organizations are faced with increased regulatory requirements that require greater control over information and processes — often restricting the flow of information across legal jurisdictions — while being driven by market pressures to globalize and reduce costs. The two are contradictory and require careful oversight.

--Forrester Research, Inc., "Enterprise Risk Management." by Michael

Basic Security Stance - Achieving an “Acceptable Level” of Risk

Risk is what an organization would lose if an asset was “compromised,” taking into consideration how likely it is that the asset could be exploited (which is based on how vulnerable it is). Understanding there are risks to everything that is in or connected to the network, it is important to establish what is deemed an “acceptable level” of risk for the organization.

Traditional belief is that making an asset available on the network automatically places it at odds with the security that will restrict or inhibit its access and availability. The key to security is finding a balance and being able to both facilitate and secure network communications. It is important to understand the trade-offs that need to be made to ensure an asset is available and/or secure. For example:

- If a patch needs to be applied to several devices in the network, does it happen immediately, during an acceptable “downtime,” or as part of a regular maintenance schedule? The decision will depend on the criticality of the device and the severity of the vulnerability exposed by not patching the device. This is a clear case of security of the device versus availability of the network.
- If a business unit needs a particular application delivered to geographically distributed locations, for example CRM applications to outsourced customer care centers, what level of risk is an organization willing to assume in order to ensure that particular application is available? Do all user constituents get complete access to the network, or is access limited based on what business function the user performs?

These questions and more will need to be answered by the organization as they gauge where they are on the risk spectrum. They will then need to map their risk threshold to the resources on the network. They will need to identify what they place value on in the network, understand the potential threats associated with those assets and ascertain how vulnerable those assets are to those threats. Then they can start to determine what to do in the network to protect their assets and create a security stance that reflects their tolerance for risk. These next sections discuss how to evaluate the value, threats and vulnerabilities that could feed into an organization’s approach to securing their network.

The Value of the Network

In order to understand risk, the organization must understand what they would lose if they were compromised. Once the loss value is determined, the next thing to do is understand their priorities, developing a corresponding criticality of their assets to their business. What can the organization never live without and what needs to be protected at all costs, versus what is a nice to have or convenience?

The organization will need to ask themselves, what their proprietary code is worth, their financial data, customer profile information, business/product plans, etc.? How much would it cost them in lost revenues, brand erosion, etc. if their public Web site suddenly was not available, or their transaction portal could not be reached, or their streaming data feeds were stopped? What productivity and operational costs would they incur if e-mail was brought down or their order tracking system corrupted?

The organization will need to rank where their Web site, customer database, financial statements, e-mail, etc. sit in order of importance. This is also where potential compliance (Sarbanes-Oxley, HIPAA, Gramm-Leach Biley, Data Protection Act) issues come into play, which may dictate which assets require more stringent security measures; and where customer data versus partner data versus internal data gets evaluated. Understanding where an organization places the highest value is the starting point for developing a security plan that can protect those assets. Some of the key value considerations organizations must weigh are itemized in the list below:

- Loss of proprietary information, intellectual property
- Loss of personal data, customer data, financial data
- Degraded or damaged brand, loss of goodwill
- Downtime, service degradation, compromises in performance
- Failure to comply to standards and/or regulations

Network Threats

Threats to the network are greater than ever. This is due to a variety of factors including:

- An increasing number of network users
- An infinite number of trust levels driven by different business purposes and relationships
- The growing ways in which the network can be accessed (e.g. PDA, wireless, etc.) – opening up new entrances into the network
- An increase in network, computer and software vulnerabilities – more applications, more network use – more ways to exploit the network
- The proliferation of “attack tools” and the ease by which they can be attained (e.g. Web sites, books, etc.) – making it possible for almost anyone to try to attack the network

The threats can come from anywhere; they can be launched both externally and internally. They can be malicious in nature, with the trend towards attacks designed for financial gain - phishing, spyware – or overall network disruption - worms and viruses. They can also be unintentional, such as an employee accessing something or deploying a device/application they didn't know they shouldn't. The key to assessing risk is in understanding where and how those threats are most likely going to succeed, which stems from being aware of what is going on in the network.

Understanding the Network's Vulnerabilities

Organizations need to understand the vulnerabilities associated with their network to accurately assess where they are at risk. To do that, requires specifics. It requires a clear understanding of exactly where the assets are in their network, how they are being used and who is using them. The second step is to determine exploits and vulnerabilities associated with each of these assets.

Organizations need an inventory of what the organization has –clients and servers –, the state of those devices – OSes and applications loaded/used, patch levels, etc. –, and associated activity – user access, downloads, requests, etc. This can be done in an absolute manner, gathering information all at once, at predetermined intervals or by

subscribing to scanning services. It can also be done in a continuous manner, gathering information from the network traffic on a regular basis.

However gathered, the information must be correlated in a way that gives it meaning. It is the relationship of what is on the network with how it is being used -- up to the application and content level -- and by whom that enables organizations to understand and minimize the risks to their business.

Knowing there are five Apache servers on the network is not significant by itself. Knowing exactly where they are, what software versions they are running, how, when and who is using them represents the type of information an organization needs to assess the availability versus risk trade off.


Once assets are gathered and identified, specific vulnerabilities, exploits or patches can then be mapped to these assets to determine the assurance and security level needed.

Take an example of an OS vulnerability and ensuing patch announcement. An organization with good awareness of their network, will know exactly which devices are affected, what applications are running on those devices and who accesses them. They can then focus on creating a patch plan that minimizes the chance of attack, accommodates potential downtime of key applications and notifies users, so they can prepare for any disruption. Without good visibility, an organization is left potentially blind to the risks that some systems in their network pose and scrambling to create a patch plan, which provides ample window of opportunity for a successful exploit.

Security is a Process

A security stance is reliant on a combination of people, policies and technologies; all of which are continuously evolving. The threats of today, may not be the threats of tomorrow and what protected the organization a year ago, may need to be supplemented by different policies or solutions today.

Security is a on-going process that can be broken down into three types of activity: analysis, mitigation and ongoing management. The organization should analyze their risks, put appropriate technologies and policies in place to mitigate those risks, and then manage them to optimize their success. The organization will then repeat these efforts to continue to refine their security to meet the ever-changing needs of their organization.



Analysis

Analysis of an organization's risks is key to a strong security stance. Unfortunately, the analysis part of the security process has traditionally been reactive, meaning organizations learn of a security "event" and then spend their time analyzing the network. As a result, analysis has been tied to investigating an event to understand the breakdown and subsequent extent of the damage/success, in an effort to mitigate the effects and try to prevent it from happening again.

While investigative analysis will always be a necessary part of the security process, the goal should be to shift the role of analysis from an event resolution activity to an event prevention activity. Strong analysis of the network could provide the organization the

information they need to be proactive in their security stance. The analysis phase could then be used to identify problems, so exploits can be avoided.

The effectiveness of an organization's analysis is based on the organization's level of understanding of the network's composition and activity and their ability to turn that knowledge into action.

The Importance of Network Awareness

Keeping track of what is going on in the network may be the biggest problem that organizations have in understanding and establishing their network security. Organizations should try to achieve a real-time, complete picture of the network to effectively analyze their security stance. They should identify network devices and activity and understand the role they play in the network.

Everything needs to be considered in the context of what it could mean to the organizations' overall network security posture. The organization must have a way to tie the client/server and application use to their users to understand if the activity is truly allowed. For example, someone connecting to a server they have never connected to before may indicate malicious/unauthorized activity or may indicate a new/expanded use of that application well within normal/accepted business practices.

Network awareness enables organizations to accurately analyze their risks, so they can avoid being blindsided by attacks and be proactive in their efforts to deploy the appropriate measures to enforce appropriate use and minimize threats. Once deployed, organizations should continue to analyze the network activity to ensure proper enforcement and enable ongoing refinement of the organization's security stance.

Mitigation

Security is all about controlled access. If employed properly, it can ensure access is granted to only those who should have it and to only the things they need and eliminate dangerous or malicious traffic; all without impeding the overall availability of the network. The main problem organizations face is how to deploy the appropriate enforcement and controls.

An organization can begin by balancing security techniques, so an appropriate level of security is enforced throughout the network. The key is to put more intelligent enforcement points into the network and enforce security controls at the endpoint (both client and server) and as the traffic traverses the network. How this is done will vary based on the network topology, asset allocation, user requirements, etc. Some of the general considerations are discussed in the following sections.

Where to Start – Understanding How Traffic Flows Through the Network

The first thing to do is determine what security/availability trade-offs to make within the network. This starts with an understanding of how all the components work together.

From an elementary view, the network is made up of users, endpoint devices and network devices. The users leverage endpoint devices, which “talk” over a network of interconnected network devices, to other endpoint devices/servers to request information/applications (services). The receiving endpoint will then respond. Network devices, which have knowledge of where all of the requests and responses should go, transport these communications from one endpoint to another.

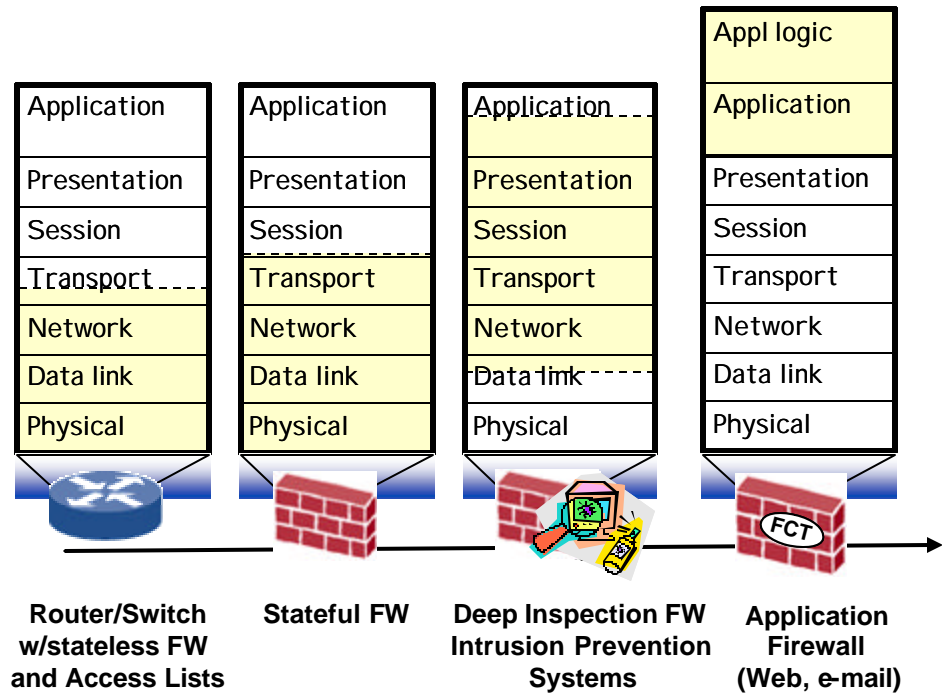
As traffic flows from one endpoint to another, there are a lot of “participants” that see and touch it. Each one of those participants must make a traffic forwarding decision and determine how they want to handle that traffic. The number of checks and the level of “detail” that goes into each decision determines the level of intelligence associated with the forwarding decision, which dictates the level of security and assurance for that traffic.

Due to a variety of factors, which will be explored in greater detail in subsequent sections, the way access is granted to the network is evolving; additional levels of intelligence need to be applied throughout the network and be of a dynamic nature to ensure that any changes to the user, endpoint, resource or application can be accommodated into the access controls. It is important to note that different devices throughout the network leverage different levels of information when making their transport decisions. Not all devices anchor their decisions based on all of these intelligence levels.

Traditionally, access was granted based on the IP address of the source and destination and protocol being used. This is sometimes called the 5-tuple (source IP address, source port, destination IP address, destination port, and the protocol) policy and is most commonly associated with network-level security. It may be the function of stateless firewalls, which are often associated with routers/switches using access control lists, as well as stateful firewalls that track the sessions, or state of the connection, associated with the 5-tuple policy. Making access decisions based solely on network-level information leaves the network open to misuse and vulnerable to attacks at the application-level that are brought in by “trusted” or “allowed” traffic. This phenomenon is evidenced by the proliferation of attacks and worms over the last several years.

Recent innovations have brought application-level aware security devices into the network, allowing them to enforce application level traffic decisions. There are several types of devices that can provide this level of intelligence, including deep inspection firewalls, intrusion prevention devices and application firewalls. Each of these provide the same type of intelligence (application level), but perform different levels of interpretation and contextual extraction.

The graphic below depicts the “level” of information generally used in the traffic decisions by these different network and application-level devices:



If a user is involved in the communication, the user credential and the state of the endpoint could be another important element to the traffic decision. Since only network and application data is transported across the network, user and endpoint information must be requested or gathered as an adjunct mechanism. There are generally three mechanisms in which user and endpoint intelligence is gathered:

1. Implying a user community and trust level based on an IP – this is what most network security devices do
2. Requesting user credentials as needed for access to specific resources – note it is a one time challenge by an enforcement point in the network, so the user/device information is not incorporated into ensuing communications
3. Requiring an integrity check of the user and device (enforcing certain standards-patch levels, AV signature files, OS versions, etc.) as a pre-authorization mechanism used by an enforcement point when granting access to the network

The new “security paradigm” requires greater intelligence and multiple trust levels that correspond to different levels of access. What and where enforcement/control points are deployed will depend on what the organization is trying to do. Because there is no single “right way” to mitigate risks, what is put where and the way in which the enforcement points are managed will determine the security stance of the organization.

Management

Management is where the analysis and mitigation phases of security come together, providing a way to both understand and change what is going on in the network. Management is about ensuring that each person who plays a role in the organization's security can quickly get the information they need to do their job. It is also about being able to create a security stance that continuously maps to the dynamic nature of the network and the needs of the business.

For simplicity sake, security activities can generally be separated into two types, operation and threat management. Operation management relates to the deployment, configuration and ongoing maintenance of the network and its security. Threat management relates to the event monitoring, data correlation, analysis and incident handling aspects of security. It is important the management systems have a user-friendly, intuitive way to accomplish all of these activities to ensure an optimal security stance.

There are two levels of management systems, those that provide an all-encompassing view of the network and those that are designed to manage a specific device's functionality. Each provides different "lenses," and both are often necessary in a large, distributed organization.

"Umbrella" management systems gather information from almost everything in the network. These overarching systems gather logs from routers, switches, security devices, clients and servers and present that information in several "lenses" for the organization to analyze, often focusing on overall network performance, availability, uptime, etc.

The information from a specific device's management system may feed into the larger management systems, however, the robust functionality of the device is best, and often only, managed by these specific systems. If there are multiple, different devices, it is easy to see how important it is for the management systems to be able to simplify the security activities and not add to the complexity.

Requirements for a Strong Security Stance

The real challenge in establishing a strong security stance is often in the implementation. These following sections iterate through some of the requirements for each phase of the security process – analysis, mitigation and management – that will help an organization identify key implementation requirements.

Network Awareness for Accurate Analysis

Protecting information assets starts with tight integration of corporate policies (what is acceptable use to meet overall business goals) with rigorous, regular checks of all network activity (to identify any aberrations and/or violations). To identify what is on or what is happening on the network, the organization may use a variety of tools, which are listed below and followed by a brief overview describing each:

- Vulnerability assessment and scanning tools - active
- Traffic analyzer, network IDS, network recorders - passive
- Profiling and passive fingerprinting - passive
- Security Event Mgmt, Security Incident Management Systems

Vulnerability assessment and scanning tools are designed to actively initiate communications to network IP addresses and analyze the responses to determine what IP ranges, ports and protocols are available. Further assessments could consist of sending a series of known exploits to the network to see which devices respond and are susceptible to the series of canned attacks. The detailed information they provide about the devices, services/applications (including versions and patch numbers) on the network can help an organization determine how to protect their resources. For example, the information can be used to create a patch management plan, identifying which vulnerabilities are best protected by patches and which require other mechanisms/actions (e.g. services that need to be shut down, deeper analysis by enforcement points, etc.).

The data generated by the scans represents a single moment in time and requires on-going assessments to create time based analysis of what happened when and where. Even with on-going scans, system vulnerabilities could occur in between scans. During the assessment or scan, a load is put on the network which could burden the network and disrupt normal operations. There is also little context on what is actually being used in the network, and without this context, it is hard to know how to prioritize or assess the most immediate threats to the organization. Also, vulnerability assessment/scanners generally focus on servers, leaving the organization potentially blind to the client vulnerabilities.

Traffic Analyzers, Network IDSes, Network Recorders are devices that passively monitor and extract information from the network. Depending on the device, they then analyze and present the data in a variety of ways to create varying levels of awareness. Traffic analyzers are primarily used for network diagnostics. They capture the traffic and decode the protocols, (e.g. translating the bit pattern to ASCII characters) for analysis by the organization. A network IDS goes further, reassembling the traffic, up to the application-level, and then applying a set of rules to look for suspicious activities. Once identified, the network IDS can trigger an alert for additional analysis by the organization. The network recorder captures the traffic with the intent of storing the information so it can be replayed by the organization to determine what happened. Each of these devices provide information on what is happening in the network, however, due to the amount of additional analysis required by the organization to gain a lot of value, they often play niche roles in creating network awareness

Profiling and passive fingerprinting tools have recently emerged to passively collect data from the network and extract contextual information on the communications of clients and servers on the network. They can provide organizations with different levels of information: some provide flow-based information, some provide details on all network communications, others focus on network and application data, while others focus on the characteristics of the client (operating system, etc.). These tools provide ongoing visibility, enabling organizations to determine what hosts (clients and servers) are using the network and the types of activity (services/applications) taking place.

There are typically two approaches to delivering profile information; the heuristic approach and the data points approach. The heuristical approach maintains a traffic baseline that can be used to identify trends or major changes based on activity that is "statistically significant." The data points approach can present specific details on users, devices, the network and applications that can be used to understand network activity and any changes.

Security Event Management (SEM), Security Incident Management (SIM) Systems provide a way to gather logs and events from a variety of sources, including routers, switches, firewalls, VPNs, desktop management tools, IDSes, IPSes, etc. They provide a picture of who is coming in and out of the network, as well as who is trying to get into the

network but being blocked. Reviewing logs within a single device provides one perspective, while gathering and correlating logs from different devices provides another. (Data correlation often offers greater context to what is happening on the network.) These solutions (SEM/SIM) provide additional functionality including heuristic and/or rule-based triggers to notify the user of certain conditions or threats. While these functions are valuable, the base function is to increase awareness of what is happening in the network.

Organizations may use any number of these technologies, both individually and in concert, to gain visibility into their network activity. Achieving an understanding about what is going on in the network enables organizations to effectively identify risks, align security and availability objectives, prioritize activities and ultimately deploy enforcement points throughout the network that are able to establish a security stance that maps to the network goals of the organization.

Security Enforcement throughout the Network

Organizations need technologies in place to control the use of the network and remove threats from the network. What is implemented is based on an organization's tolerance for risk. Organizations need to deploy the appropriate measures to ensure they can achieve a security stance commiserate with the organization's security and availability goals.

Regardless of the organization's approach, to achieve a uniform security stance requires the network be considered as a whole, with measures incorporated in the following locations:

- Security for endpoint clients
- Security for the network
- Security for the endpoint servers

Security for Endpoint Clients

Endpoints can introduce threats and vulnerabilities both within the endpoint device, itself, and in the network, adding to an organization's risk level. As businesses grow and expand their reach, they are extending the connectivity. Opening up the network creates tension between provisioning the network to make it available and controlling risk.

To ease the tension, an organization should consider the types of users they have, what those users need, and how they access the resources to configure the network to facilitate business operations, while mitigating risks. Usually, the users of an endpoint client are requesting information, with different users requiring different things. An organization should try to deliver the appropriate experience level and grant access to content and applications based on the identity of that user (for both user productivity and security reasons).

Organizations should take into account what the user is using to initiate access. The range of what constitutes an endpoint client has grown in the last couple of years. A user can employ a laptop, kiosk, phone, PDA, among other devices to access the network, and they can be coming from almost anywhere, such as a home office, a café, an airplane, etc. No longer does an endpoint client have to be physically connected or tied to a single location within the network to gain access. The degree to which the organization can check and trust the endpoint device, should influence what an organization grants access to for that user.

In order to mitigate risks posed by the endpoints to the network, organizations can employ a variety of technologies to do threat removal, use enforcement and integrity checks. It may be enforced at the individual device level or incorporated into an overall network security enforcement policy.

There is a class of agent and/or client technologies that reside on the client and perform threat removal, use enforcement and integrity checking functions that relate to the overall security posture of the endpoint. Some examples of the functions are described below (note that one or more of these functions may be delivered independently or integrated into a single “product” and not all endpoint risk mitigation functionality is listed):

Personal Firewall (pFW) – enforces a policy that dictates what services and applications are allowed on the host

Host IPS - protects the hosts against Trojans, spyware, worms and other known & unknown threats, as well as unauthorized or malicious applications

Malware scanner/cleaner – scans and protects the host against executables or “bad” programs, such as Adware, SpyWare, Key Loggers, Trojans, Dialers, Hijackers, Trackware, etc.

Anti-virus program – scans files using known attack patterns to protect the host against viruses and worms

Asset agent, patch agents – scans registry, filename and other computer information to determine asset or software installation, revision and patch level on the endpoint

Cache cleaner – deletes the contents in the browser cache, so session information is not retained on the system

Integrity checker – interrogates different components of the system (registry, running programs, installed programs, etc.), as well as other programs (pFW running status, AV scanner status, etc.) to identify the integrity of the endpoint at any given time

When looking at controlling the endpoint, these functions should be considered and implemented as appropriate to establish an endpoint security posture commiserate with the level of risk the organization is willing to assume. This endpoint posture can then be combined with who is using that device to access the network and where they are coming from to form the complete data set. This data set can be used to make a trust decision for that endpoint and help network devices make more intelligent traffic decisions.

In addition to the agent and client technologies mentioned previously, there are several other components that need to be considered in the context of endpoint security:

- An agent that gathers and transmits the endpoint data set to a policy controller
- A policy controller that authenticates the user, communicates with the appropriate backend identity management system, maps the user and endpoint data set into the policy store and then authorizes the level of assurance and security that should be performed in the network
- A network component that enforces the assurance and security, in accordance to the policy. (The network components in the next section.)

The last aspect of security for the endpoint is the understanding that conditions will change on the endpoint and that those changes need to be accommodated at the client, as well as in the network. There should be mechanisms to enable both interval and trigger driven scans of the system to monitor current endpoint state. When a change is identified, there should be a mechanism that can appropriately protect the endpoint from the misuse and threats that correspond with that change, in addition to adjusting the policy controller and the resulting network enforcement of the assurance and security policy.

Security for the Network

Devices deployed in the network need to provide the level of assurance and security the organization requires. To do this, network devices need to be able to make traffic decisions more intelligently, by increasing the depth of traffic inspection, and provide the appropriate controls. To that end, there are three general functions that must be considered in the network:

- **Quality controls** – ability to provide a predictable experience, optimized for any network environment - deliver what is expected
- **Use controls** – ability to control how the network is used – explicitly dictating what is allowed and how the allowed traffic can or cannot be used in the network (implicitly denying everything else)
- **Threat controls** – ability to remove “bad” stuff from the allowed traffic – traffic that is malicious in nature or behaving inappropriately

Network devices should have built-in **quality controls** that are in line with the connectivity requirements of the organization. Throughput, fault tolerance, availability, reliability, and specific network or application acceleration are some examples of attributes that should be considered. Organizations may also want to ensure particular applications are given appropriate priority, bandwidth or even accelerated to deliver an assurance-level that maps to the needs of the application or a particular user. Different attributes may be desired based on the location of the network device, the expected behavior and organizational constraints. For example, the fault tolerance requirements at a central site, versus a branch office, versus a remote office may be different based on the information at those sites, technical expertise, costs, tolerance for failover switching time, etc.

Organizations will need to enforce appropriate **use and threat control** which determines what is allowed to go in and out of the network as well as what is not allowed (use control) while removing unwanted or malicious traffic (threat control). Use and threat control enforcement encompasses intelligence at the network, application and user levels, as appropriate.

Use Control at:

- Network level - dictates what services are allowed in and out of the network
 - e.g. 5-tuple policy, allow anyone (IP) to access the Web server (IP), port 80, with HTTP/s traffic, but don't allow any other traffic to that server (SMTP, etc.)
- Application level - dictates what applications are allowed in the network and how those applications can or cannot be used

- e.g. specify which applications are allowed and which commands within an application are allowed/not allowed – allow Instant Messaging, but don't allow downloads
- It may even go to the content level, dictating what content can be accessed - e.g. preventing users from accessing blacklisted Web sites or content
- User level – grants access to applications based on the user/endpoint
 - e.g. full-time employee, coming from their corporate office desktop can access everything, same employee coming from their home office laptop can only access e-mail and the Intranet

To ensure that the “allowed” traffic in the network is behaving as it should, organizations should also have a way to understand the intent of the traffic and remove threats. At the network-level there are patterns and models that can be used to identify threats (e.g. requests coming from an external source, using an internal IP address are indicative of a spoof). At the application level, a device may reconstruct the traffic into the application message and then look for known patterns, characteristics and anomalies that indicate the traffic is a risk to the organization. Ultimately, threat control at different levels provides different types of protection:

- Network level – prevents network-level attacks
 - IP spoofing, Syn Floods, DoS, DDoS, etc.
- Application level – prevents application-level attacks
 - worms, viruses, Trojans, application vulnerability exploits, phishing, spyware, etc.
- User level – performs checks based on the trust level of the user/endpoint
 - perform “xxx” checks on the traffic of a contractor, perform “xxx” checks on the CEO's traffic

Having granular control and being able to dynamically enforce that control can further improve the assurance and security level. Since the level of traffic processing and traffic decisions occur at different logical levels (network/application/user), having intelligent enforcement points throughout the network that can complement any network infrastructure allows the appropriate assurance and security to be brought into the network in a cost effective manner. The following is a list of different types of network devices that provide quality controls, use controls and threat controls:

- Routers
- Stateless firewalls
- Stateful firewalls
- Deep inspection firewalls
- Application firewalls

Security for Endpoint Servers

The endpoint servers in the network often represent the assets of that organization. They are where the services - information and applications- are stored, whether it is financial records, customer databases or intellectual property. The servers are generally targets for attacks and where most “misuse” takes place, so organizations probably want to make sure they are protected by rigorous security measures, many of which have already been detailed in the endpoint client and network sections. In conjunction with the security provided throughout the network, however, there are a variety of other steps/technologies that can be used to help secure the server. The steps include:

- Hardening
- Host Intrusion Detection, Host Intrusion Prevention
- Keeping up with patches and vulnerabilities and protecting against them
- Auditing for vulnerabilities

Hardening: Most servers are created for general use, with general operating systems that are designed to maximize how users can use/access them. As a result, organizations should try to turning off or removing unnecessary or insecure services, utilities, administrative options that could compromise the integrity of the system (e.g. turn off most management protocols).

Host Intrusion Detection, Host Intrusion Prevention: There are programs that can be loaded onto the host server that can track activity that may indicate attacks or misuse. Some are based on logs, others on system calls, while others are based on network access/transmissions. Regardless of the method used to detect host intrusions, once detected, these clients can send an alert, quarantine, remove or even clean up the system to minimize exposure to the attack.

Keeping up with patches and vulnerabilities: Most attacks exploit vulnerabilities in an operating system or application, so one of the most important activities an organization can do to protect their servers is keep track of vulnerability alerts and any associated patches that are issued. Due to the ever-shortening time cycle between vulnerability announcement and exploit, it is particularly important to keep the patch levels on these servers up to date to minimize the chance of a successful attack.

Auditing: To effectively patch, an organization must understand exactly what these endpoint servers are running to ensure the risks can be appropriately assessed. The activity on these servers needs to be regularly audited to ensure nothing is running on them that should not be there. If an organization doesn't know exactly what is on a server, it can be a "way in" for an attacker.

Security Management

The overall management of an organization's security is tied together with governing principles that dictate the security stance. These governing principles are then translated into policies that can be enforced at the user, device and network level.

The Role of Policies

A security stance is founded on the security policies that define it. While policies can be expressed at different levels, they will ultimately get translated into device configurations. This is where the management system becomes critical. Organizations need to be able to translate "basic" policy logic (e.g. everyone should be able to access the corporate Web site) into configurations. They must also be able to interact with the management system in a way that is most appropriate for the business operations and response team.

There are two ways in which a system can be managed; organization can choose between a centralized policy approach where a single policy controls configuration of multiple or all devices in the system, and a point to point policy approach where individual policies map to individual configurations for different devices in the system.

Organizations also need to be able to extract information from the management system that can validate or invalidate the effectiveness of the policy and then adjust accordingly. Ideally, it would supply policy templates that simplify management tasks, making it easy to create and modify a policy for particular network segments/devices. Another consideration is how the policy is expressed. Management systems may offer a rich client, a web-based client, a command line interface (CLI) client or some combination.

Logging, Auditing and Reporting for Security Event Resolution

Once policies and configurations are established a variety of data needs to be extracted from the system. The data can represent different slices of the information and can be used for different purposes. Some of the activities that the system needs to enable include:

- Log and event collection and aggregation
- Investigation, analysis
- Reporting, exporting data to other systems (asset management system)

Log and Event collection and aggregation - Logs provide a myriad of information. The type of data generated will depend on the role of the device (e.g. router vs. intrusion prevention solution). In general, it can be split into two categories, performance (up-time, bandwidth utilization, resource allocation, etc.) and event (attack specifics, access information, etc.) related information.

The performance log information is used for operational purposes and is the type of data that can be rolled up into an overarching management system to identify network performance problems and trends. The event log information is instrumental in security incident management and is probably what administrators will spend most of their time working with to manage security.

A management system must not only be able to provide organizations with a way to drill into the details of the logs to get specific information on what happened, but it also must be able to correlate the logs and report out on larger trends.

Investigation and Analysis - Organizations need to be able to quickly identify what happened during an “event” and then make changes to the policy, as necessary, to prevent further incidents. These closed loop investigative capabilities should enable an administrator to easily link between different data representations and levels of detail directly to the appropriate enforcement policy to make adjustments. In this way the security stance can be tuned to reflect the current risk assessment and ensure appropriate exploit mitigation.

The management system should relieve the burden that is placed on the administrator when an event occurs. The system should be able to facilitate understanding, enabling administrators to take the data points that are known and extrapolate what really happened. So whether they are investigating high activity on a given application, systems being accessed when they shouldn't, or a server becoming unavailable they are able to quickly drive to resolution.

Reporting - Once an organization has policies in place, it will need to monitor the effectiveness of those policies. The data contained within the logs may be used for auditing purposes, pulled into reports for trending and macro-level information and mined to get insight into a particular event on the network. Another key activity of security administrators is being able to provide information that supports audit and compliance reports. The management system should be able to provide granular records of who did what, when, etc. to help with regulatory requirements, as well as export data as necessary to support other business operations.

In summary, it is very important the management system provide correlation and links between information stores. It then must and present the data in such a way that it is easy for an organization to understand exactly what is going on in the network – both in terms of trends and specific security incidents/events – and enable administrators to easily make modifications and report out on the security stance of the organization if necessary.

Discovery, Feedback and Adaptive Controls

Organizations should be able to quickly adapt their enforcement when circumstances change. New information about the endpoint, network and things happening on the network should be accommodated and the appropriate responses activated. There are several ways to improve the adaptive nature of the security and assurance policies, such as:

Endpoint discovery and feedback – a collection of products and technologies that allow user and endpoint state information to be re-assessed on a regular basis and fed into the management system or policy control system. The collection mechanism can either be polling-based (regularly scheduled checks at time intervals on the endpoint state info) or trigger-based (checks on the endpoint state requested when particular criteria are met-something happens). The endpoint state information is the result of products described in the ‘security for endpoints’ section.

Network discovery and feedback – a collection of product and technologies that allow network data to be re-assessed on a regular basis and fed into the management system or policy control system. As detailed above, the collection mechanism can either be polling-based or trigger-based and the information is the result of products described in the ‘network awareness’ and the ‘security for network section.’

Network Device Event Feedback – alerts, logs, or other information that is sent from a device, as a result of a particular condition/s being met on the network, or some fault condition being detected. Most network devices have pre-built logging and event mechanisms, such as syslog, snmp, proprietary event logs, script triggers, etc. by which this information can be captured and sent.

All of this data should be incorporated into the security stance, so that it can evolve in the same way the network is evolving. This requires a level of communication among the devices throughout the network, which can be accomplished through a variety of communication mechanisms, some of which are open standards and others which are proprietary.

However accomplished, the key takeaway is the management solution should be able to adjust enforcement as soon as new information is discovered. By incorporating this feedback, organizations can appropriately respond to the dynamic nature of the network and prevent unnecessary exposure to risk.

The Juniper Networks Approach – Conclusion

Juniper Networks provides organizations with the right level of functionality throughout the network to ensure organizations have a secure and assured networking environment. Juniper helps organizations throughout the security process, delivering the tools organizations need to analyze their network activity, effectively mitigate their risks, and manage their network security in a manner that optimizes their resources.

Juniper Networks brings expertise, products, solutions and partnership to help organizations secure their network infrastructure. The Juniper product line includes:

- Juniper Networks M-Series Edge Routers
- Juniper Networks J-Series Customer Premise Routers
- Juniper Networks Integrated Firewall/IPSEC VPN Security Appliances
- Juniper Networks Intrusion Detection and Prevention Systems
- Juniper Networks Secure Access SSL VPN

With Juniper Networks, organizations don't have to compromise. They can achieve a level of security and assurance commiserate with the needs of their most demanding networks. For more information on Juniper Network solutions, please visit www.juniper.net.

Copyright © 2005 Juniper Networks, Inc. All rights reserved.

Juniper Networks, the Juniper Networks logo, NetScreen, NetScreen Technologies, GigaScreen, and the NetScreen logo are registered trademarks of Juniper Networks, Inc. NetScreen-5GT, NetScreen-5XP, NetScreen-5XT, NetScreen-25, NetScreen-50, NetScreen-100, NetScreen-204, NetScreen-208, NetScreen-500, NetScreen-5200, NetScreen-5400, NetScreen-Global PRO, NetScreen-Global PRO Express, NetScreen-Remote Security Client, NetScreen-Remote VPN Client, NetScreen-IDP 10, NetScreen-IDP 100, NetScreen-IDP 500, GigaScreen ASIC, GigaScreen-II ASIC, and NetScreen ScreenOS are trademarks of Juniper Networks, Inc. All other trademarks and registered trademarks are the property of their respective companies.

Information in this document is subject to change without notice.

No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without receiving written permission from:

Juniper Networks, Inc.

1194 N. Mathilda Ave. Sunnyvale, CA 95014 ATTN: General Counsel