

White Paper

Secure and Assured Networking with an Enterprise Infranet

Andrew Harding
Director

Roslyn Rissler
Sr. Manager, Product Marketing



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

Part Number: 200144-002

Contents

Contents.....	2
Introduction: A Secure and Assured Network.....	3
What is working – what is not.....	6
Layer 2 Only – Access Port Approaches	6
Client-centric Approaches.....	8
Limited, Complex, Single Vendor Solutions	8
A Secure and Assured Network.....	10
Introducing Infranets.....	10
An Enterprise Infranet Solution	11
The Emerging Need for Secure Transport	13
An Enterprise Infranet for the Extended Enterprise	13
An Enterprise Infranet for the Distributed Enterprise	13
An Enterprise Infranet for the WAN Gateway	13
An Enterprise Infranet for the Data Center	14
An Enterprise Infranet for the Campus LAN.....	14
Enterprise Infranets: Why Choose Juniper	14
Real Time Security	14
Cost-Effective.....	15
Compliance	15
Reliability	15
Built for Continuous Improvement	15

Introduction: A Secure and Assured Network

Today's enterprise - and, increasingly, today's economy - is dependent upon networks. This trend has served to speed innovation and remove barriers to efficiency, such as physical separation that can limit access to enterprise resources. Business partnerships are closer than ever before, because parties at distributed sites across the globe can access, use and share resources and applications as easily as those across the campus or those working from home using an SSL VPN. Employees are no longer tied to desks; they can get full access to network resources and applications from virtually anywhere. The growing trend to work closely with a wide range of business associates, along with moving production and development to remote locations, requires ubiquitous connectivity that goes past enterprise employees. Contractors, consultants and other business associates can be made productive immediately with complete connectivity, whether they are onsite using the campus network, or they access resources remotely.

Today's enterprise networks have evolved beyond the constraints of the traditional network infrastructure. This evolving network is increasingly heterogeneous and encompasses applications, endpoints and the users themselves. Different constituencies connect across different network segments, and much of their traffic traverses WAN gateways as they connect to follow-on networks over VPNs, or use business-critical applications that require Internet access. A diverse population connects to the campus network and uses applications deployed at Internet facing data centers. Network endpoints at distributed sites need access to resources at the campus headquarters and at those same data centers. Network infrastructure has grown to support this traffic, and overall productivity and efficiencies have grown in kind, but so have some security risks.

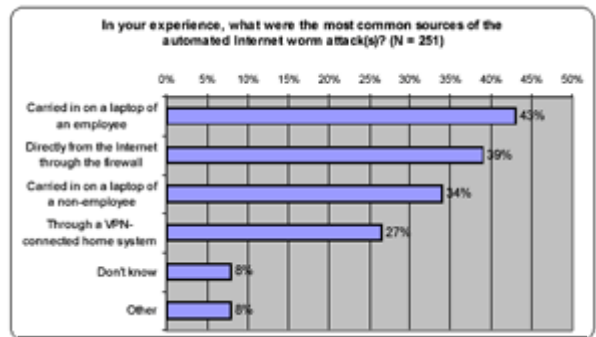
As the enterprise becomes critically dependent on network services, a downside of this trend emerges. Due to the increasing need *and* ease of providing network access for all audiences, and the resultant dependency upon it, today's network is both more mission-critical and more vulnerable than ever before. Providing access to any resource, be it data in classified documents, patient health information, or intellectual property, has always required the balance of the value of access against the risk of security breach. Traditional networks provided that protection against breaches through layers of security, which gave increasing levels of protection focusing on perimeters, critical segments, and application-oriented authorizations. That model provides comprehensive access to the network and protects resources from external threats and unauthorized application access, but it falls short. The assumption used to be that if the endpoint is in the building, it is acceptable to provision broad access and count on application controls to provide security. This is no longer the case. As the trust level that we can instill in devices decreases, providing comprehensive security and assured delivery of services grows more difficult. Some access is controlled and many assets are defended, but the perimeter, and each successive layer, provides less protection as users, devices and applications change over time. The network services and connected applications that enterprises need to deliver can leave the enterprise retreating toward the center of the network, with challenges arising from both inside and outside.

The fact is that edge perimeters alone have been insufficient for some time. Edge security will not protect the internal segments of the networks or be a replacement for host security. Viruses and email worms, Trojans, denial of service attacks and other malicious activities frequently utilize end-user machines to penetrate enterprise environments, even when perimeter security mechanisms like firewalls are in place. New technologies, such as instant messaging or peer-to-peer applications, represent an emerging threat vector, which may

bypass the perimeter devices completely. Enterprises have attempted to control the use of these technologies with policy, but such policies are virtually impossible to enforce in a traditional network. If you have e-mail, Web content, public facing servers, or mobile users, then your network is inherently open to a constantly changing variety of attacks. And the threats and vulnerabilities are more subtle and dangerous than they first seem—there is mobile code that can be brought into a network via a secure (but mobile) PC. As a leading analyst has observed:

[N]etworks have many back doors that provide opportunities for circumventing perimeter defenses: employees who access public e-mail accounts and instant messaging systems; WLANs that have been misconfigured, compromised, or installed against policy; contractors with access to internal systems; links to partner networks that bypass firewalls; holes that have been opened in firewalls to support business applications such as conferencing; remote employees who have these same network vulnerabilities at home; and the encrypted tunnels formed during normal Web browsing that bypass enterprise controls and allow downloaded executable content to run on user desktops”¹

Traditional security approaches, designed to protect a traditional network of well-defined users with fixed workstations, were not created to protect the virtual network. Analysis of recent worm outbreaks demonstrate the point. Most recently the ZOTOB worm, which exploits the plug-n-play vulnerability in Windows 2000, has affected such large enterprises as CNN, the ABC (American Broadcasting Corporation) television network, Visa, American Express and the New York Times. The worms have been spreading widely due to several factors, including the rapid appearance of exploits and infected laptops. While organizations may block port 445 on their firewalls, if employees' laptops become infected elsewhere, when they bring them inside the perimeter and connect them to the network, the worm begins to spread internally. The editor of the SANS Newsletter comments that *“The time from vulnerability announcement to release of a worm was one of the shortest seen in recent times. Patch announced August 9th (Tuesday); exploit code posted publicly August 11th (Thursday); worm started to hit on August 13th (Saturday). Because worms spread over 139/tcp or 445/tcp, ports that cannot be firewalled without breaking some functionality in Windows environment. That means that even a single infected laptop brought inside an enterprise will infect all the other machines.”*



Enterprise Strategy Group, January, 2005

Over 80% of infection and propagation of recent worms (SQL Slammer, Blaster, etc) has occurred though mobile and remote, managed and unmanaged PCs. Gartner has found that *“...best-of-breed organizations have control of only approximately 80 percent of the endpoints on their networks. This means that on the order of 20 percent are unmanaged. Also mismanagement of the other 80 percent often results in additional vulnerable endpoints.”⁴* While threats from internet content that exploits managed PCs owned by the enterprise and used only by employees certainly need to be addressed, mobile PCs and unmanaged devices are also a critical problem because they likely are not running the standard security software deployed within the enterprise. New threats are emerging daily, blending the characteristics of worms, Trojans and viruses. In the past 4 years, high proliferation worms have cost over \$31 billion dollars, either in opportunity cost, recovery and cleanup costs, or permanent damage to digital assets that were either impossible or impractical to replace or repair. London-based consultants Mi2g estimate damages caused by viruses like Sobig or Klez to be in the billions of dollars. The Blaster virus

alone is reported as having infected more than 300,000 computers in 24 hours, and causing \$525 million in damages. Complicating matters is the fact that these mobile and unmanaged PCs are as critical to the productivity benefits and efficiencies as the network to which they connect. These client hosts are the main tool of the users, employees and business associates that rely on the enterprise network. Enterprises need to address the threat of ubiquitous access and the problem of mobile and unmanaged PCs to extract the very benefits of enabling every user – employee, consultant, or other business associate--who has a legitimate business need to get and remain connected.

But the need to address the problem is about more than reducing downtime, increasing productivity and protecting intellectual property and other confidential information. A wide range of new regulations, including the Health Insurance Portability and Accountability Act of 1996 (HIPAA), the Gramm-Leach-Bliley Act of 1999, and the Sarbanes-Oxley Act of 2002 each pertain at least in part to maintenance of security, privacy and integrity of networked information. For example, Sarbanes-Oxley (SOX) states that weaknesses in internal controls add up to risk. If such a risk impacts (or can allow impact to) financial data, it is considered a material risk. Under SOX, material weaknesses must be publicly reported, and audits get both more expensive and more intrusive with each reported weakness.

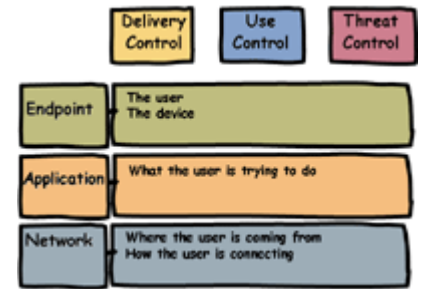
Part of the difficulty in securing the network is that it is now made up of so many different components. One of the first areas to be considered as vulnerable was the extended enterprise, whereby remote and mobile users needed to access the enterprise network, typically using a remote access VPN. The use of SSL VPNs, with sophisticated endpoint assessments and policy enforcement capabilities, has done much to control that threat. However, many threats remain, and each poses unique challenges. The WAN gateway is a key source of vulnerability. Users on unsecured or inadequately secured endpoints can inadvertently bring in network threats such as viruses, Trojans, worms and more without even being aware of it. Any of these threats can then quickly propagate to the point of bringing down the network. Another area of potential vulnerability is the data center. Key resources are often protected by a firewall that has no information about the users or the client endpoints that seek to access the data center. When contacted by a compromised endpoint, sensitive internal servers are at risk. Still another area for concern is the distributed enterprise, made up of remote or branch offices. An exploit that compromises a client at one of these offices can propagate via a site-to-site VPN.

The challenge for the enterprise is to find a way to leverage their existing infrastructure as much as possible, while providing a simple, secure way for a diverse audience to access resources and applications. According to leading industry analysts, “the goal is to achieve a balance between restricting users, enhancing and extending current perimeter security measures, and assuring long-term compatibility through the use of open standards.”¹ This goal must be achieved while ensuring the delivery of services and applications in such a way that the enterprise can depend upon them. In order to realize this goal, the enterprise network must provide overall control in the forms of use control, threat control, and delivery control. Such a structure allows the enterprise to go from knee jerk reactions to each event, to the creation of a network with sustainable compliance and enforcement literally built into the infrastructure.

Use Control – Determine what use is allowed, deny the rest

As network connectivity becomes increasingly ubiquitous, control of who can get access to resources and applications and the circumstances in which they can get it is vital. User identity and authentication must be combined with endpoint assessment and must be applied on a per-session and intra-session basis. It must also be considered that in most security deployments, the user and their habits are

typically the weakest links. As such, use controls must also meet the seemingly opposite needs of being consistently applied, yet transparent to the user. One example is user passwords, in which users often choose weak passwords, or write them down. This problem can be solved with the use of dual factor authentication, consistently applied. Use control also goes farther than simply identifying the user – it must bind the user to their endpoint device, and ensure that the device is and remains compliant with the required security posture.



Delivery, Use, and Threat control must permeate all facets of the network to be effective

Threat Control – Detect and remove malicious content

The network perimeter has effectively disappeared, making network threats a fact of life. These must be proactively factored in to the enterprise security stance, not simply reacted to. Endless retreat to the middle ultimately results in no protection at all.

Delivery Control – Assure application delivery: performance, confidence, predictability

Delivery can no longer be considered best effort, nor can it introduce additional risk. End-to-end secure transport is required for many transactions to assure the controls are applied, based on the composite of users, endpoints and infrastructure. Also, because enterprises can no longer assume their internal networks can be trusted, they must protect confidential information in transit by using encrypted channels between clients, servers and other network-attached devices.ⁱⁱ

What is working – what is not

While the description of a Secure and Assured network is relatively simple, the actual creation of such a solution has thus far been elusive. Some approaches are nothing more than proposals for a time-consuming and expensive replacement of large portions of network infrastructure. Other proposals, while valid, meet only part of the overall challenge of a truly Secure and Assured network. Some solutions are functional, but aren't standards-based and cannot operate in heterogeneous environments. Others solve only current security issues without looking to the future, ensuring that by the time such solutions are deployed, they are already practically obsolete. Technologies currently under discussion are discussed in the following sections.

Layer 2 Only – Access Port Approaches

Extensible Authentication Protocol, or EAP, is a general protocol for authentication. 802.1x implements EAP for port-based authentication within the Layer 2 network. The purpose of EAP is to ensure interoperability and compatibility of authentication methods, enabling an intermediate connection device to repackage and forward these packets.

The 802.1x standard was developed in part to secure wireless LANs. The client that wants to be authenticated, referred to as the supplicant, sends a request to the intermediate connection device, which, in the case of WLANs, is usually a wireless access point. Using EAP, the authentication request is simply repackaged and forwarded to the authentication server. The authentication server and the access point simply passes the packets on. If the supplicant provides valid user credentials, the authentication server responds with a success message and allows access to the network. 802.1x-based solutions are being offered primarily by switch vendors, and addresses user authentication. The switch acts as an authentication backstop, relaying the request to a RADIUS server for authentication. This combination is being used as part of various overall security schemes, by enabling the authentication server to route the supplicant to a “quarantine” portion of the LAN based on authentication.

While 802.1x, EAP, and AAA servers are important components, they do not form a comprehensive solution, nor do they represent an immediately realizable one. According to Forrester Research, “...only about 15% of all enterprises are underway with 802.1x-enabled switches.”ⁱⁱⁱ This is partly because actually using 802.1x is fraught with costs. First, the supplicant software must be installed and configured on all endpoints. According to Forrester, “Although Microsoft includes 802.1x in all versions of Windows XP, only 17% of enterprises have actually deployed Windows XP to all desktop PCs. Also Microsoft’s supplicant isn’t robust enough for all enterprises...as a result many enterprises will need to purchase a standalone 802.1x supplicant.” And there are additional costs, including the cost for IT to configure the supplicant, the installation of RADIUS on the “back end” for security and scalability, and a backup authentication mechanism for the interim. But the cost and complexity are not limited to client hosts or to the AAA enhancements required to support such systems. Additional challenges emerge at the switches themselves, which will need to be located, assessed, and potentially updated. This job alone can be very taxing to enterprises without a homogeneous, centrally managed switching infrastructure.

There are other drawbacks to this method. Often, the endpoint device is checked only at the beginning of the session. Because users and their devices are mobile, as are attacks, any infection that occurs after the session has begun would be completely undetected. In addition, this implementation focuses on Layer 2, with minimal consideration of the higher layers. In fact, the network and application layers are where today’s blended threats wreak havoc, and have been able to do so, because the majority of networks and endpoints cannot protect themselves against them. Some traditional networks do not even see that such attacks are taking place until the damage is already done. These systems do not address mobile and remote PCs, unmanaged PCs, or exploits that result during a connected session due to Web-based threats. They amount to authentication backstops and are limited to the Layer 2 network and an initial authentication transaction, not the entire user session. This is an unacceptable posture for a security control in today’s environment. While port-based quarantine can create physically or logically segment network zones for non-compliant endpoints, it requires an extensive inventory of all networking equipment along with comprehensive upgrades to succeed. It is also worth considering how robust such a quarantine network would be, and how much infrastructure will be required to support it. With port controls alone, the resulting network is inherently fragile; a missed port, a NAT box, or a simple router, and even virtual machine technology can circumvent such systems without much effort.

Central to these Layer 2 approaches is the notion of network quarantine. Network quarantine is defined by one leading analyst as the restriction of client system’s access to the network based on its compliance with policy. Layer 2 approaches, however, are not the only method that “quarantine” systems. Another method, known as server-based quarantine, uses server

software and mechanisms like DHCP address assignment to enforce quarantine measures on non-compliant endpoints. While this implementation attempts to move users to another part of the network, it does not protect against trivial circumventions. Both approaches are incomplete and do not give the complete protection that is needed.

Client-centric Approaches

Client-centric implementations presume that security and assurance will be provided by a security client or combination of clients on the endpoint device. In a homogeneous network of identified, managed endpoints with a true, defensible perimeter, this would be a sound assumption. One could rest assured that only managed devices would be connecting to the network, and that such managed devices would all have approved, up-to-date security applications running on them.

One quick look at today's network, however, dispels this notion. Consider the trusted consultant (or, ironically, a compliance auditor!) that arrives on site. Their PC will almost certainly not be carrying the enterprise-approved security suite, yet they must have access to the network in order to do their work. While security applications are indeed important – if not vital – to network security, they cannot provide a complete solution for two reasons:

1. SSL VPNs provide access to remote users, and provide the required user, endpoint, and application controls
2. Campus networks, however, lack those controls and are exposed to mobile PCs and unmanaged devices that business associates and employees use every day. Network protection begins at the endpoint, but endpoint protection is only the beginning of network access management.

Such applications are not distributed to every endpoint that will connect to the network. The endpoint cannot be the last line of defense, but should instead be part of a solution with a scheme for implementation. Alone, client-centric solutions provide limited benefits, because they offer no protection for--or from--unmanaged and mobile PCs.

Limited, Complex, Single Vendor Solutions

An alternative approach to the problem of balancing access with security requires a homogenous single vendor solution, which uses some of the partial solutions above as well as extensive proprietary technology that requires equally extensive replacement of--and modifications to--network infrastructure. In these implementations, some of which are led by Cisco Systems, the vendor uses a combination of installed heavyweight software clients and new network infrastructure. The installed clients must run on all endpoints. They utilize proprietary methods to interact with the vendor's switches and routers throughout the network.

While the concept behind such an implementation is understandable – at least from the single vendor's viewpoint – it is also unrealistic. There is virtually no such thing as a truly single-vendor network today, nor should there be, as such an approach can lead to the creation of proprietary security and service delivery mechanisms that limit the enterprise in the future. Should a company try to deploy a single vendor solution anyway, they are literally dooming the solution from the start by creating a gaping hole for unauthorized, unsecured traffic unless every single switch in the network is replaced, and every single user (including those that are not employees) have the 802.1x client software installed. Single vendor, heavyweight client- oriented approaches leave unmanaged, mobile or remote PC unaddressed, since there

is no way to provision them on all endpoint devices that may enter the network. If one endpoint is missed, the network is at risk. This is sure to be discovered by compliance auditors, or worse yet, by an attack, leaving the enterprise worse off than it was before the exercise began.


The process of implementing such an approach is also far from simple. According to industry analysts, it requires the following steps at a minimum:

1. Audit the current network infrastructure for non-Cisco, non-supported hardware
2. Map all supported hardware to the supporting version of IOS
3. Upgrade or replace your infrastructure according to the audit-map matrix
4. Install Cisco's Access Control Server (ACS)
5. Integrate Cisco's policy extensions throughout existing AAA and access systems
6. Re-map or create a restricted network segment for quarantined users
7. Install a new version of the Cisco Trust Agent on all of your clients
8. Install the Cisco Security Agent
9. Maintain the software and infrastructure inventory over time

The problems with any single vendor implementation are not unique to Cisco's implementation. Port-based approaches also require that a great deal of money be expended to provide ubiquitous access control deployments. In reality, single-vendor solutions have all the weaknesses of traditional approaches, including the fact that they are incomplete, either

limited to Layer 2 or with impractical Layer 3 support, and easily circumvented, leaving the network vulnerable to intra-session exploits. They do not function at all in heterogeneous environments, leaving the network open to the greatest threats: mobile/remote and unmanaged endpoints. In addition to being complex and difficult to maintain, this approach can cost tens of millions of dollars to deploy in a large enterprise network. Even more worrisome, because it requires numerous installed clients, such a deployment leaves much of the problem untouched. In a Cisco environment, mobile and remote PCs, whether managed or unmanaged, can remain an unaddressed vulnerability.

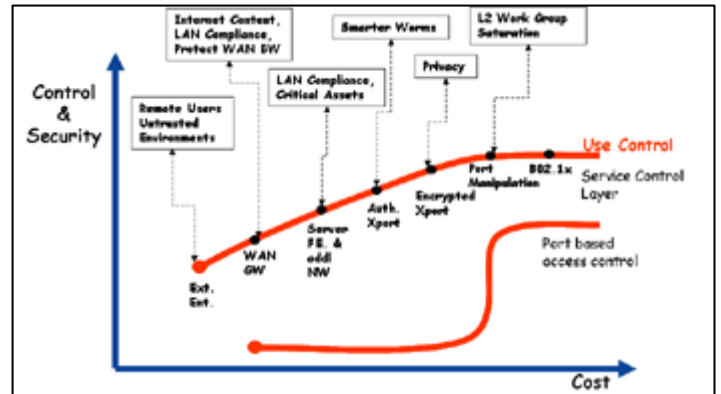
Still another roadblock to the adoption of a proposed single-vendor solution is the scope of the deployments. Close to 100% deployment must be achieved for the solution to be effective. This means a major investment to achieve any meaningful level of use control, and does not even consider the enormous "hidden" IT costs to upgrade servers and install/configure software on each client.



Port-based access control requires an "all or nothing" approach

A Secure and Assured Network

A Secure and Assured network must go beyond costly patchwork of partial solutions that come together only in concept. A true solution will reduce risk while remaining cost-effective. One way to facilitate cost effectiveness is to leverage the existing infrastructure, which also minimizes the introduction of new risk factors that are inherent in any new deployment. The solution must have the throughput required in today's LAN, and network policy must be incorporated to determine which traffic needs to be encrypted.




Costs vs security for unified access control and port-based access control

A secure and assured network will demonstrate consistent uptime and be inherently resilient to even the newest attack. Such a network injects intelligence and visibility into policy decisions and enforcement, and grants not only guaranteed, but session-appropriate access to users, based on dynamically determined criteria. It will work for all users and for all client hosts, as well as at all network attachment points, regardless of manufacturer, providing visibility into network access and control events.


Introducing Infranets

Juniper Networks headed the creation of the IPSphere Forum (www.ipsphere.org), whereby carriers and providers of carrier infrastructure work with content providers to establish trusted transport and federated policy. An IPSphere is a model for using and producing structured IP services based on the Service Oriented Architecture (SOA). Juniper's recipe for the creation of IPSpheres, in both the carrier and enterprise environments, is the Infranet^{iv}. The secure and assured networking solution discussed below is the first implementation of a reference Enterprise Infranet.



An Enterprise Infranet Solution

This Enterprise Infranet implementation has taken the Infranet concept from carrier to enterprise scale, combining the best of existing technologies, many of which are deployed today, with new innovations that will connect both Juniper and third-party technologies effectively. The result is a scalable, standards-based, cost-effective network that can be implemented and immediately and progressively built upon as open standards mature and are adopted.



Infranet Controller, Agent, and Enforcers in the network

The explicit consideration of delivery, use and threat control is a cornerstone of Juniper's Infranet architecture. The Enterprise Infranet solution considered here for unified access control begins with the most critical element for immediate success – Use Control. Use Control is already a key element of Juniper Network's market-leading SSL VPN, which is designed to provide endpoint assessments before and during a session, ensure strong authentication/authorization, leverage the existing infrastructure, dynamically create session-specific roles, and make use of extremely granular resource policies. This functionality is now extended throughout the enterprise network, tying policy visibility to existing enforcement points^v, using the Infranet Controller^{vi}. The Infranet Controller (IC) communicates with Infranet enforcement points via the Infranet Agent. These components combine to create a service control layer on top of the existing enterprise infrastructure, integrating end point/user intelligence with network and application intelligence to ensure adaptive and granular levels of Use Control. Additional components delivering Threat Control and Delivery Control can also be layered on this solution, without requiring any forklift upgrades to the existing infrastructure.

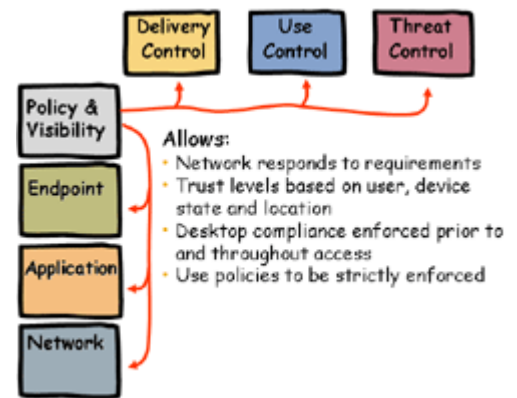
When the user first logs in to the IC, the IC dynamically downloads the Infranet Agent (IA). This is very similar to the way that a remote or mobile user would log in to the enterprise using an SSL VPN. The IA is a lightweight software agent that determines the endpoint's compliance to enterprise security policy, similar to Juniper's SSL VPN Host Checker. Also like Host Checker, the IA can be configured to check predefined and customizable criteria, including running processes, port activity, application authenticity, and the presence/version of best-of-breed third-party security software applications. If the user is found to be noncompliant, due to the absence of endpoint security applications, such as antivirus or malware protection, associated data files or settings, the user is directed to a site where the posture can be remediated – many times, on the fly. The dynamically provisioned IA also provides optional authenticated and encrypted transport to enforcement gateways where it is necessary.

The information gathered by the Infranet Agent is then passed back to the Infranet Controller. The Controller leverages Juniper's Secure Access SSL VPN policy and control engine to provide seamless communication with authentication servers, identity and authorization directory stores. The Controller then enables session-specific, conditional access to networks based upon user authentication and a real-time check of the device's security characteristics. This implementation will provide benefits similar to those of Juniper's SSL VPNs for remote users and partner extranets, at the scale of enterprise LANs. The combination of the Infranet Agent and Controller effectively prevents non-compliant hosts from attaching to the enterprise network and provides use control for enterprise network traffic.

Access information that reconciles authentication and authorization information, as well as identity and endpoint assessment, is then passed to an Infranet enforcement point. Enforcement capabilities are added to platforms via a software upgrade, enabling already-deployed infrastructure to take on this role without requiring any kind of forklift upgrade. The first enforcement points available for the extended enterprise are Juniper's SSL VPNs. The second iteration will leverage the scale and deployment of Juniper's best-in-class NetScreen firewall platforms. These firewalls are broadly deployed in key locations within the network and make granular traffic decisions. By consuming signals from the Infranet Controller and terminating connections from agents, enforcers can provide end-to-end transport security and greater policy control and visibility without requiring significant infrastructure changes.

This Enterprise Infranet solution also complements Juniper's work with Microsoft on Network Access Protection (NAP) as well as standards-based efforts, such as the Trusted Computing Group's Trusted Network Connect Subgroup. The effort represents the next phase in integrating user, application and network policy and enforcement in a seamless manner through open, multivendor solutions that facilitate a secure and assured network.

Use Control – Use control should determine what use is allowed and deny the rest, with consideration for the user, the endpoint and the resources. This goes far beyond what can be done at Layer 2, and must encompass all aspects of the virtual network and the infrastructure, including firewalls, Web filtering, authorization and directory stores and the network log on. This Enterprise Infranet solution incorporates all elements of use control, including endpoint defense, where the host is assessed for compliance, remediation is provided, and dynamically determined, session-specific access management for users, endpoints, session roles and network resources/applications is the result.



Unified policy permeates every facet of the LAN

Threat Control – Threat control must move beyond basic protection of the endpoint to include the entire network, as well as analysis of the traffic flowing over it. This includes deep packet inspection at the firewall, and the ability to not only detect intrusions but prevent them. Threat control must also affect use control, by being able to detect and remove malicious content from allowed access. Juniper's Enterprise Infrastructure provides this protection, with a combination of best-in-class, Deep Inspection firewalls at the perimeter, and the award-winning Intrusion Detection and Prevention (IDP) appliances. The combination, along with anti-virus protection, provides complete security from attacks, whether they are launched from outside the network, brought in accidentally by a mobile employee or unmanaged PC, or launched from inside the network itself.

Delivery Control – Assured delivery is often taken for granted, but in reality it is a vital component that must be built into the infrastructure to succeed. Delivery control assures application delivery; this can include performance, Quality of Service (QoS), MPLS, Virtual Private Networks (VPNs), high availability, optimal path selection and more. Juniper Networks enables assured delivery control, with industry leading firewall/VPN platforms, as well as enterprise routing products built with many of the same capabilities found in Juniper backbone routers used by the world's top service providers.

The Emerging Need for Secure Transport

If one considers today's enterprise network to be both broadly dispersed and de-perimeterized, it follows that some transactions require secure transport, just as they do in the remote and extranet use cases. This transport would effectively bind the user and the endpoint to the infrastructures, providing a path to use control, threat control and delivery control. Juniper's Enterprise Infranet solution dynamically invokes secure transit where and when it is most needed. When used as an Enforcer, a firewall/VPN can selectively provision components of IPSec as required by policy, specifically, ESP with no encryption, or when required, full ESP. This selective deployment option gives an easy method to realize full encryption when it is required, without the cost, complexity or overhead seen in trying to deploy IPSec VPNs throughout the enterprise.

It is best to illustrate how this Enterprise Infranet solution can be realized today by considering the network components mentioned earlier: the extended enterprise, the distributed enterprise, the WAN gateway, the data center, and the campus LAN.

An Enterprise Infranet for the Extended Enterprise

An Enterprise Infranet solution already exists for the extended enterprise, via Juniper's award-winning Secure Access SSL VPNs. New innovations with the Network Connect network-layer access method combine the high performance of IPSec transport with the ubiquity of SSL to form an adaptive transport mode that ensures the best connection possible in every network environment. Core Clientless Web Access from Juniper remains the industry's gold standard, now supporting XML-based and Flash content, as well as providing a unique delivery platform for Java applets. In addition, Juniper's partners with best-in-class, third-party security applications deliver state-of-the art protection and enforcement for every session.

An Enterprise Infranet for the Distributed Enterprise

Use control and threat control can also easily be delivered to the distributed enterprise. Customers typically deploy smaller firewall/VPNs in these locations, such as the Juniper Networks NetScreen-500 or NetScreen-200 families. Users would typically come into the branch office, often from an insecure network or using an unmanaged endpoint, and be passed directly into the enterprise LAN via the site-to-site IPSec tunnel. Juniper's NetScreen firewalls already feature Deep Inspection and antivirus checking, so some of the threat is already mitigated. Empowering the branch office firewall to act as an Infranet Enforcer introduces policy into the equation. User's compliance will be checked before the connection is allowed, eliminating a major source of security problems.

An Enterprise Infranet for the WAN Gateway

Users with unmanaged or noncompliant endpoints going out to the Internet are also a traditionally weak point in the network. Juniper's industry-leading ISG firewall VPNs were built to handle the multi-gigabit traffic typically seen at the WAN gateway and to apply Deep Inspection to it. In addition, the ISG is now available with an integrated Intrusion Detection and Prevention (IDP) module and will become an Infranet Enforcer in early Q1 2006. In addition to multi-gigabit throughput, the IDP modules will add other important security features, such as Layer 7 inspection, spyware and shell code protections, as well as

consideration of both client-to-server and server-to-client traffic. When used as an Infranet Enforcer, these platforms will combine their intrinsic security measures with additional policy control and visibility into network traffic. This combination results in trusted transport between endpoints and the WAN gateway, ensuring that only those users that are protected sufficiently can go out to the internet.

An Enterprise Infranet for the Data Center

Another vital area to secure is the datacenter. As we've mentioned, these data centers are often protected only by the firewall. However, when that firewall is functioning as an Enterprise Infranet Enforcer, it has all the power needed to protect sensitive resources and applications. This protects the datacenter from non-compliant or possibly compromised endpoints – those endpoints that have been assessed and authenticated can get access, while those that are not cannot communicate with the Infranet clients or with internal resources. By shunning endpoints that are unmanaged or not compliant, Juniper's Enterprise Infranet creates security without replacing any infrastructure at all.

An Enterprise Infranet for the Campus LAN

While it is typical for users in the extended enterprise to log in to an SSL VPN and thus be authenticated and verified as secure, users on the campus LAN are typically not required to do so. The Enterprise Infranet changes that paradigm. Now, users that do not yet have the Infranet Agent installed will be greeted by a portal page when they plug into the LAN. This portal page connects the user to the Infranet Controller and enables the download of the Infranet Agent. The Agent assesses the endpoint's security posture and can verify them as compliant, or direct them to a page where they can remediate. The Agent is configured to run these checks periodically throughout the session to ensure that security posture remains consistent. For users that already have the Infranet Agent installed, the day-to-day log on procedure is virtually unchanged.

Enterprise Infranets: Why Choose Juniper

The following sections detail the reasons for choosing Juniper's Enterprise Infranet solutions for secure and assured, and service-aware networking.

Real Time Security

Juniper Networks signaling capability with other network devices, in addition to the endpoint enforcement, allows the network to adapt to new threats in real time. Up-to-the-minute information from the endpoints is used across the network to make the right network services decision. This results in an effective, efficient solution that virtually secures itself through continuous feedback.

Cost-Effective

With the service control layer approach of Enterprise Infranets, deployment can be incremental and does not require a full-scale replacement of switching infrastructure. The enterprise may choose to focus on critical network segments, or particular network or security problems and build from there; for example, significant benefits can be provided by initial implementations, such as the extended enterprise, the WAN Gateway, data center and through using authenticated or encrypted transport where required. This allows the enterprise to add significant control and security to their infrastructure at an incremental cost.

Juniper's Enterprise Infranets will work in heterogeneous environments and are compatible with products from other infrastructure vendors. They are also compatible with other use control, threat control and delivery control approaches, providing a flexible solution that does not hamstring the enterprise. It is also interesting to note the overall expenditure is comparable, despite the wealth of features in the Enterprise Infranet solution.

Compliance


Another important point is in the area of compliance. Assessment and enforcement become integrated into everyday processes, so the enterprise does not need to deploy another infrastructure or insert another procedure to become compliant. These features are built into Enterprise Infranets.

Reliability

The introduction of unproven hardware and software inherently pose a risk to the enterprise, as they can constitute new vectors for intrusion. The Juniper Enterprise Infranet solution is built on existing and proven standards and technologies, much of which is already in use today. Juniper's SSL VPN solution, for example, has been deployed for some time, and is currently providing similar capabilities for the extended enterprise in securing remote/mobile users and partner extranets using Host Checker and Juniper Endpoint Defense Initiative (J.E.D.I.) modules. The extension of these capabilities provides a simple and consistent user experience in all scenarios.

Built for Continuous Improvement

In addition to Use Control, the Enterprise Infranet's service layer approach also provides a foundation for the addition of threat control and delivery control capabilities. Examples of such additions could include a feedback loop from Juniper's Intrusion Detection and Prevention platforms to the Infranet Controller that a particular endpoint is generating malicious traffic and needs to be forced to remediate. Another example could include user information, which signals the requirement for differentiated delivery requirements when particular applications are used.



This flexibility also positions the Enterprise Infranet to leverage the emergence of new standards and innovations from other vendors.

Juniper has created the Enterprise Infranet solution with a view not only to what you need the network to provide today, but what your needs will be tomorrow. This first reference implementation solves the most critical security needs of the enterprise, and the Enterprise Infranet architecture extends to meet additional requirements in a flexible and dynamic way. In this Enterprise Infranet deployment, all elements are designed to work in a heterogeneous environment and are based on open APIs and accepted or emerging standards. The solution does not require a forklift upgrade of any infrastructure, nor does it require the deployment, installation or configuration of any client software. Now and in the future, Enterprise Infranet solutions can be deployed gradually, thus eliminating the difficulty and expense of an “all-or-nothing” alternative.

ⁱ Securing the network perimeter is more important than ever, Gartner, March 23, 2005

ⁱⁱ Impact of the Disappearing Perimeter: Strategies for securing internal networks and endpoints, The Burton Group, September 23, 2004

ⁱⁱⁱ Beware The Hidden Costs of 802.1x, Forrester Research, March 14, 2005

^{iv} .An Infranet is based on the creation of three layers:

- *Service Structuring Stratum (SSS)*: allowing service capabilities to be joined together amongst Infranets
- *Network Policy and Control (NP&C)*: Includes policy servers and Operational Support Systems
- *Traffic Handling (TH)*: Routers, Switches, Firewalls, etc.

There are numerous foundational Infranet architectures in the carrier space (see *The Foundation Infranet: A Cookbook for Network Service Providers* at www.juniper.net/solutions/infranet/710013.pdf).

^v The existing enforcement points referred to are on the *Traffic Handling* layer of an Infranet)

^{vi} The Infranet Controller is on the *Network Policy and Control* layer of an Infranet)