Identity-based Policy Enforcement

A LAN Security Best Practices Whitepaper

By Nevis Networks
Persistent LAN Security Solutions

Executive Summary

This whitepaper focuses on the evolving nature of LAN security in today's enterprise in light of a dissolving network perimeter and the need for an identity-based solution to address new requirements. Network security policies arise from compliance and risk management initiatives across multiple lines of business throughout the organization. Security, compliance and business requirements are articulated in a readable policy built up from basic identity, role and group definitions, or can be read as network security access decisions that are mapped to user profiles. Identity is at the core of enterprise policies.

Network infrastructure, and network security solutions built on top of the infrastructure, is not identity-aware, since network packet headers provide information about machine addresses and location, not user information. Enforcing identity-based policies with identity-blind systems has proven to be a futile endeavor, in light of increasingly complex security policies, open networks, mobile systems, and unmanaged endpoints. The dilemma facing network security administrators has become an insurmountable obstacle and cash drain, resulting in poorly designed security models being implemented at the wrong places in the network. Exacerbating the problem is that without an identity-aware network infrastructure, it is almost impossible to demonstrate compliance with the identity-based policy initiatives. The events of interest are occluded in the network cloud of machine-address-based technology.

The solution is to build user identity knowledge into the network fabric, and enforce identity-based policies within the secure network. Network security policies can then be easily mapped from the definition stage into the network security architecture, with clear visibility to user activity through the enforcement, remediation and reporting phases. This offers a clear ROI by greatly improving network administration and user management costs, reducing the complexity of ill-fitting network security infrastructure, as well as reducing the costs of managing policy breaches and compliance reporting. Security policy enforcement is moving into the network to address the dissolving network perimeter problem, and when it does, the network infrastructure and the security policy enforcement layer must be identity-aware.

1. THE DISSOLVING NETWORK PERIMETER AND THE EVOLUTION OF SECURITY POLICIES

Traditionally, network security has focused on the external perimeter. The vast majority of threats to the enterprise network have arisen from outside the organization and from the open nature of the Internet. This naturally led to securing the perimeter of the corporate network very early on. Internal users and systems were generally “trusted” to be secure and there were little internal security mechanisms beyond securing the endpoints themselves. Under such an environment, an employee who had access to the network could get around on the corporate network unchecked. Access to resources such as file servers, application servers, etc. was controlled at the server directly via access privileges.

With the advent and pervasive deployment of mobile devices, portable storage media, wireless communications, and ubiquitous access points, security policies grew more complex and more security infrastructure proliferated as point solutions to diverse risks. The once-perceived secure internal perimeter now has to accommodate a myriad of “untrusted” users such as guests, contractors, business partners, customers, mobile systems, employee-owned systems, and other unmanaged endpoints. The ability to create a secure DMZ has effectively become outdated, and the entire LAN or internal corporate network has become the new DMZ. Administrators now have to defend the entire network from untrusted endpoints with conflicting and diverse security policies and access requirements, while dealing with conflicting goals of productivity and security.

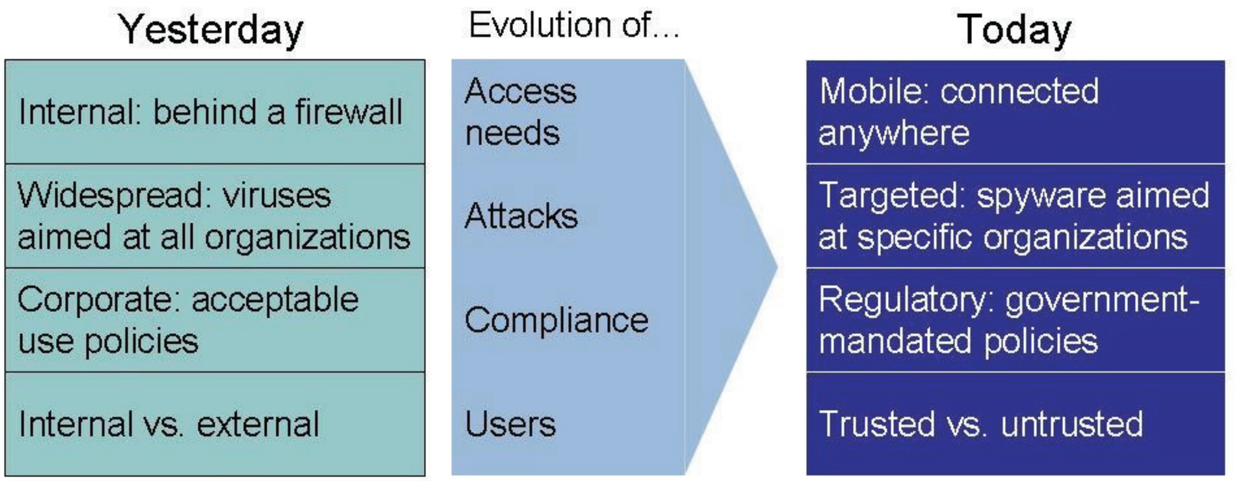

Yesterday	Evolution of…	Today
Internal: behind a firewall	Access needs	Mobile: connected anywhere
Widespread: viruses aimed at all organizations	Attacks	Targeted: spyware aimed at specific organizations
Corporate: acceptable use policies	Compliance	Regulatory: government-mandated policies
Internal vs. external	Users	Trusted vs. untrusted

Source: Forrester Research

Figure 1. – The changing requirements of network security policies in light of a dissolving perimeter

Internally, access to sensitive network resources has been controlled at the resource itself by various access mechanisms, such as passwords and Acess Control Lists (ACL). However, that model is proving woefully inadequate today. The nature of threats today exploits the easy access to resources to mount denial of service (DoS) attacks, propagate worms, probe for vulnerabilities, and other malicious behavior. A better model to defend against an onslaught of potential attacks from internal systems is to secure the fabric of the network, not merely the endpoints (servers and user PCs). This would defend against such attacks long before they reach the critical resource, ideally as close to the source of the malicious traffic as possible.

Perhaps the biggest change as a result of the dissolving network perimeter is the resulting explosion in the sophistication of network security policies. Whereas once the core of the policy was to differentiate internal versus external users, now administrators have to deal with a much higher degree of granularity in their trust model of internal users, depending on the nature of their work, their role in the organization, what processes they use, how mobile they are, and what compliance guidelines affects their role and responsibilities.

The Role of Policy and the Policy Lifecycle

Network policies arise from compliance and risk management initiatives. "Policy" is a poorly defined and perhaps overused term, and certainly implies very different concepts to a compliance officer and a network security administrator. In order to provide some context and definition, we usually associate an organization's policy with their defined processes and procedures in order to align with business objectives, as well as compliance and risk management goals. Policies are frequently abstract and because they arise from many initiatives from different parts of the organization, it may not exist in a single format or location, and in many cases, may not even be written down at all.

Policy, however, manifests itself throughout the design of the network, and is enforced throughout in applications, servers, firewalls, password management systems, desktop security software, and a wide range of other network security products. In fact, policy is best applied when it is carried consistently through the definition and management phase, through the enforcement phase, to the policy breach remediation and compliance reporting phase. Together, we think of these as the complete policy lifecycle. Today, very few organizations can carry their intended policy consistently through all phases of the lifecycle in a cost-effective fashion.

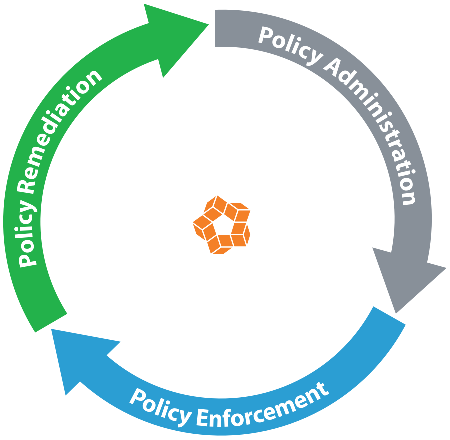

Figure 2. – The Complete Policy Lifecycle includes a policy definition and management phase, a policy enforcement phase, and a procedure and processes phase to deal with policy breaches and audit and compliance reporting activities

We now see a trend that the increasingly sophisticated policies are built on a more granular view of roles and identities throughout the organization. Policies are defined and are effectively mapped onto roles, which are in turn associated with groups of specific user identities.

This complexity in turn manifests itself in increasingly sophisticated network security policy enforcement requirements. Whereas traditional policies were simple and networks could be designed without inherent knowledge of user identity, today's policies can no longer be easily mapped to security solutions that are not identity aware. For example, if access policy is enforced at the server by requiring a user id and password, that policy is identity-based and straightforward to implement. However, as policy enforcement migrates off the endpoints and servers and into the network fabric to compensate for the dissolving perimeter, in-network policy enforcement solutions also need to be identity-aware. To date this has not been the case, because network packets do not carry the notion of user identity. They are based on machine and network address identifiers that facilitate transmission without the notion of user identity.

Compensating for increasingly sophisticated identity-based policies with in-network security solutions, such as internal firewalls and VLANs that are not identity-aware, have caused a nightmare scenario for network administrators, and led to increasingly complex network designs and overly burdensome management costs. As we shall see, this is driving a new generation of LAN security solutions that deliver identity-based, in-network policy enforcement.

2. THE CASE FOR IDENTITY IN SECURING THE NETWORK

Given the dissolving network perimeter described above, the need to enforce more sophisticated access policies is becoming increasingly important. Greater emphasis is being placed on enforcing a model that ensures accessibility to information and other resources on the network based on the role of the user within an enterprise. In effect, the network is required to enforce a policy where employees, contractors, guests and external business partners are subject to tighter, more granular policies dealing with access control and threat mitigation.

As a result, the network administrator must often resolve inconsistent goals: The need to foster greater collaboration and productivity leads to more open networks. The need to secure the network due to the dissolving LAN perimeter and the corresponding increase in vulnerability to threats, intuitively leads to a policy that naturally tightens up access to the network.

Until now, network security has focused around L2/L3 type controls. Port-based, MAC-based and VLAN-based security policies have been the foundation of network architectures. However, actual policies are expressed in terms of user roles and applications, not as L2/L3 type controls. There has been no easy or intuitive way to bridge this gap between user/role based policies and the L2/L3 based controls in the network infrastructure without having network security solutions that were identity-aware. The reason why L2/L3 controls are inadequate can be easily seen by taking a few examples:

1. Consider the case where a contractor connects to the enterprise network to do collaborative development. The contractor needs access to a set of resources such as file servers, development servers and printers within the network. Employees who work with the contractor, need access to the same set of resource, but at the same time should be able to access other resources that the contractor should not have access to. Typically both the contractor and employee would work closely doing collaborative work. Putting the contractor along with those resources in a separate VLAN requires re-wiring and re-configuration of the network. Additionally, if the employee were to connect to that same VLAN when doing collaborative work with the contractor, the employee would be severely restricted in what he could do as he would now have access to only the resources the contractor has. Other approaches, such as writing complex ACLs, again results in the same situation. What gets applied to the contractor will also get applied to the employee unless restrictions are put on which physical port someone can connect on or which subnet someone can connect on. This is an onerously complex task that doesn't scale well to large and frequently changing user groups.
2. Consider the case where guests in conference rooms need access to the Internet. Again, this has typically been accomplished by wiring conference room ports directly out to the Internet. If during a meeting an employee wishes to access the resources on the network, he would be forced to go out to the Internet as well, then have to VPN back in to the corporate network in order to access local resources. Not only is this wasteful, but is also completely avoidable if the underlying network infrastructure could detect and enforce usage based on the user's identity, wherever they connected from. Increasingly mobile users will exacerbate this problem going forward.
3. Another example involving only internal employees is the case where departmental segregation is required between finance and engineering servers. In order to do that, an administrator may put application servers into different subnets based on the department the servers belong to, and restrict activity between the two subnets by ACLs or other standard legacy mechanisms. What that means is that if an employee from the finance department "roams" to the engineering area, he would no longer have access to the finance department servers. Again, in an age where mobility and portability is fast becoming the norm, this type of static segregation clearly proves very counter-productive and even archaic.

4. Yet another example would be the case where an employee inadvertently has picked up malicious code that scans the network. In such cases, quarantining the user entirely can be severe, particularly if the user is engaged in time critical activity. In such situations, it may be desirable to simply drop just the malicious traffic and allow the legitimate traffic to continue. In other cases, if that same user were a contractor or guest, it may be okay to quarantine the user entirely. Once again, the ability to express those types of enforcement policies in terms of the user and the user role, not the machine location or network packet parameters, becomes critical.

The above issues arise because the underlying network infrastructure has no notion of the identity of the user. Nor does it have the ability to enforce policies based on the identity of the user. A guest is a guest regardless of which port he plugs into. Similarly an engineer is an engineer regardless of whether he is plugging into a port in a guest conference room, or in the engineering section, or the finance section, or whether he is roaming from one access point to another and in the process acquiring different IP addresses. If the network has the intelligence to detect the identity of the user and enforce the policies associated with that user, then fostering collaborative activity and securing the network now become complementary goals rather than conflicting.

The fundamental advantage of creating an identity-aware network is that the implementation of identity-oriented policies can be enforced in the network rather than on the now untrusted endpoint. It is also a simple matter to move the policy statements and rules to a centrally-managed directory, which can be easily managed and updated, without making expensive, unmanageable changes to the network itself. Adds, changes, and deletes to the user base become trivial changes in the central directory, as do changes in access policies for whole classes of users, or the creation of entirely new user classes, all of which can be completely mobile throughout the internal network.

The next section addresses some of these tangible ROI advantages at each phase of the previously discussed policy lifecycle.

3. ROI IMPROVEMENTS IN THE POLICY LIFECYCLE WITH IDENTITY-BASED ENFORCEMENT

Carrying a consistent view of identity through the entire policy lifecycle can lead to a very tangible ROI. We'll now look at some of the efficiency and cost advantages inherent with this approach at each phase of the cycle.

The Policy Definition Phase

The first phase of the lifecycle is the ability to define policies based on the identity of the user. This is important because the vast majority of compliance and risk management are designed around users, roles and functional groups. This user and group information is already present within an enterprise via their current AAA infrastructure. The logical next step is to associate this user/group information with the appropriate policies for those groups of users. Managing policies in this way takes advantage of how users are already being managed in the network, and makes policy changes trivial going forward.

Policy rules would frequently specify which users and groups could access which applications, resources, data repositories, servers, or services. Migrating this information off of the distributed resources themselves and into the network has the clear advantage that policies can now be centrally-managed in one or a small number of locations more efficiently. Policies can then be pushed to enforcement points throughout the network and be enforced consistently without extensive audit procedures.

The Policy Enforcement Phase, i.e., the Implementation of Network Security

The main ROI improvements come from carrying identity-based policies through to the enforcement phase. As previously shown, cost reductions come from not having to modify network infrastructure anymore to reflect changes in policies. Changing VLANs and ACLs are no longer required. Policy rules can be pushed out to the network policy enforcement engines from the central repositories.

In addition, the overall network architecture and the hardware requirements are greatly simplified. VLANs that were once enforced with internal firewalls simply go away. Subnets that hosted defined groups of users no longer need supporting network hardware. Networks become more flexible, easy to manage, with less hardware required.

In addition, a number of benefits arise from the ability to drop spurious and malicious network traffic near the source. Those packets are dropped in the access layer and are no longer running around the network trying to access unauthorized resources, only to be dropped at the destination. In the case of worm propagation or denial of service attacks, which may both be low level and hard to detect, the amount of spurious traffic over time can be significant. As a result of detecting and dropping this traffic, internal bandwidth becomes much more efficient, and quality of services improves.

The Monitoring, Reporting and Incident Response Phase

The third and final aspect of the identity-policy lifecycle is the reporting and incident response phase. In order to truly gain visibility into the network, the reporting and monitoring of the network also needs to reflect the user's activity in terms of the user as well as the network resources and applications that the user is accessing. Traditionally, monitoring, tracking and reporting have been done based on the IP address or MAC address of the user, etc. To an administrator, this can be quite frustrating and non-intuitive. Tracking down IP addresses and Mac addresses can be time consuming. Moreover a user may have one IP address in one area and another IP address in another area which would further compound the frustration.

In addition, information gathering within the network traditionally has also been in terms of IP addresses, packet statistics, byte counts, and other network-specific parameters. Not only is this difficult to correlate back with the application level policies that have been defined for the user, but the ability to identify scans, threats and breaches becomes severely limited due to the lack of stateful reporting. Packet counts and byte statistics do not really give a clear picture on whether a user is genuinely accessing valid data or whether the user is intentionally or un-intentionally scanning or mounting an attack on the network. What may seem normal traffic load for a user may actually be masking malicious activity by that user. The visibility that is truly needed into the user's real activity can only be obtained if the reporting and monitoring is based on the user's application level traffic monitoring rather than raw byte or packet statistics.

By monitoring and reporting activity on a user id by user id basis, it becomes easy to demonstrate policy compliance, rather than translating from machine, address or network-specific patterns. Incident response is more efficient because anomalies and compliance issues are immediately traced to specific users and can be correlated with the user's activity from multiple access points or with other users across the network. Incident response also is more efficient with an in-line response where much of the remediation takes place directly in the network rather than the network security device acting in an advisory mode.

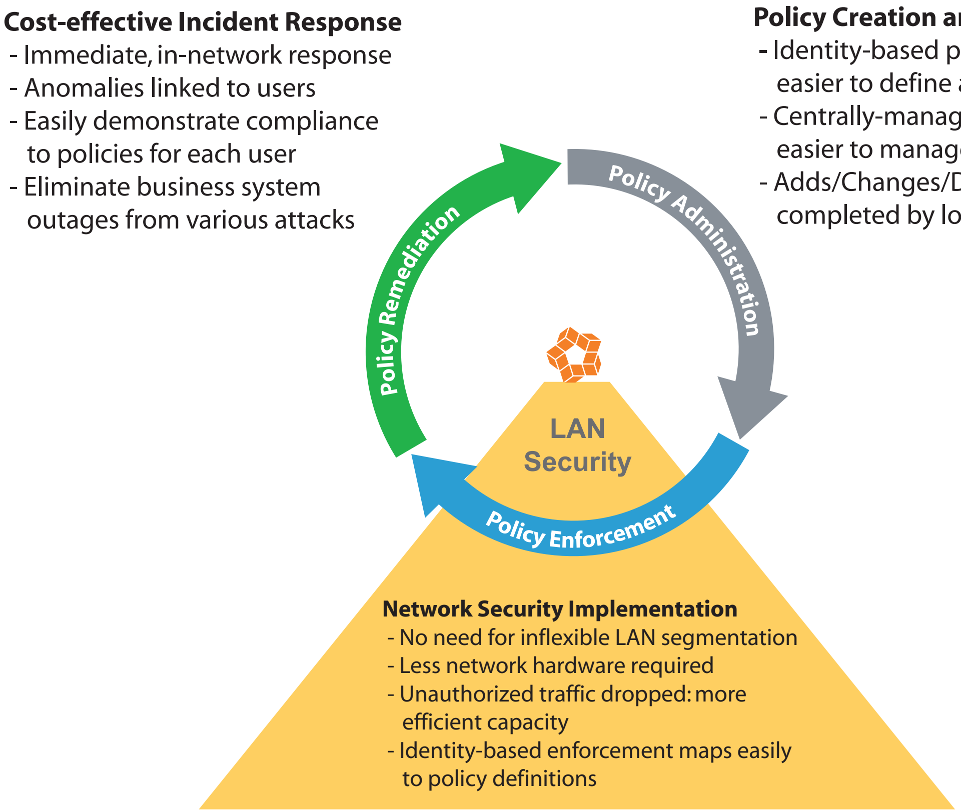

Cost-effective Incident Response
- Immediate, in-network response
- Anomalies linked to users
- Easily demonstrate compliance to policies for each user
- Eliminate business system outages from various attacks

Policy Creation and Management
- Identity-based policies are easier to define and update
- Centrally-managed policies are easier to manage
- Adds/Changes/Deletes can be completed by lower-cost staff

Network Security Implementation
- No need for inflexible LAN segmentation
- Less network hardware required
- Unauthorized traffic dropped: more efficient capacity
- Identity-based enforcement maps easily to policy definitions

Figure 3. – The benefits of identity-based policy throughout the policy lifecycle

In summary, having the ability to monitor, track and report usage based on actual user identity and applications, provides for quicker reaction time, easier reporting for compliance and audit as well as more intuitive visibility into the network. In addition, the ability to correlate user activity across the network regardless of where the user connects into the network, becomes a much more intuitive and simple exercise. Rather than dealing with packet counts, packets statistics, IP addresses, etc. the administrator can work off more intuitive user data which provides far more effective control of the network.

Once the administrator has this level of visibility into the network's operation, fine tuning policies, optimizing the network, as well as keeping the network operational, takes on a far more intuitive aspect since all aspects of running the network from policy definition, to enforcement to monitoring and back to policy definition and tuning, all happen at the context of the user, the user's role in the network and the user's application and traffic patterns.

In the next section, we will examine some high-level features of a product that brings this level of identity-awareness to network security policy enforcement.

4. THE NEVIS NETWORKS SOLUTION

In order to truly do identity-based network security policy enforcement correctly, the enforcement engine needs to accomplish the following:

1. Determine the identity of the user. A variety of methods can be used to achieve this including snooping Kerberos exchanges, presenting a captive portal to the users to authenticate against, using 802.1x, etc.
2. Based on the identity of the user, determine the group membership/role of the user within the organization. Again, this should tie into the existing AAA infrastructure of the network since that infrastructure is already present.
3. Based on the group membership/role, determine the set of policies and rules that apply to this user. Typically the policy definition and management platform would push this information to the policy enforcement devices. As mentioned earlier, it is important that the policy enforcement devices see the same consistent view throughout the network in order for the user experience to be transparent, consistent and accurate. The ability to treat the policies as network wide objects and push the same policies to all enforcement devices is important here as individually having to manage each device can be error prone and lead to inconsistent behavior within the network.
4. Enforce the rules specified in the policies. As mentioned earlier, legacy network infrastructure equipment lacks the capability to enforce policies at the granularity at which they are specified. Most network infrastructure equipment still operates their security model based on L2/L3 primitives such as Mac addresses, VLANs, etc. whereas policies would be specified at the granularity of applications that a specific user is allowed access to. However, in order to enforce policies as they are specified, network infrastructure devices need to start looking beyond the traditional L2/L3 headers and become more application aware and stateful. This requires the ability to do deep packet inspection on every packet. The information within the packet needs to be co-related back to the user's policies to make forwarding decisions in real time on a per packet basis. In effect the forwarding decision is no longer based solely on standard L2/L3 header information. Rather the forwarding decision needs to become far more intelligent in that it needs to take into account the nature of the traffic such as the application, etc., the policies associated with the traffic flow including the source and destination of the traffic, as well as the standard L2/L3 based forwarding information. Additionally every packet needs to go through a thorough inspection in order to do threat analysis based on the particular type of traffic and based on the user's traffic and usage patterns. All of this needs to happen at traffic forwarding rates that are typical of LAN environments, i.e. at 10Gig type of speeds and at price points that permit pervasive deployment of the technology.

In addition to the forwarding function becoming more intelligent and stateful, the statistics, event and alarm reporting, and troubleshooting functions of the network infrastructure devices also needs to start building user identity and application intelligence. This is critical as the ability of an administrator to quickly resolve issues, identify threats and take corrective action is only as good as the data available to the administrator. By enabling an administrator to gain insight into who is doing what on the network, when and where, the ability of the administrator to keep the network operation is greatly enhanced.

Nevis Networks offers a complete line of solutions that address all of the requirements outlined above. The Nevis product line is broken down into the LANenforcer family which is the in-network enforcement appliance (in either a switch or network appliance form factor), and the LANsight management and reporting console.

Sample Historical Report for User Activity

What did this user access in August?
Destinations from host:
1 = domain controller
2 = HP printer
3 = application server
etc.
Sources talking to host:
1 = domain controller
2 = domain contoller
3 = exchange server

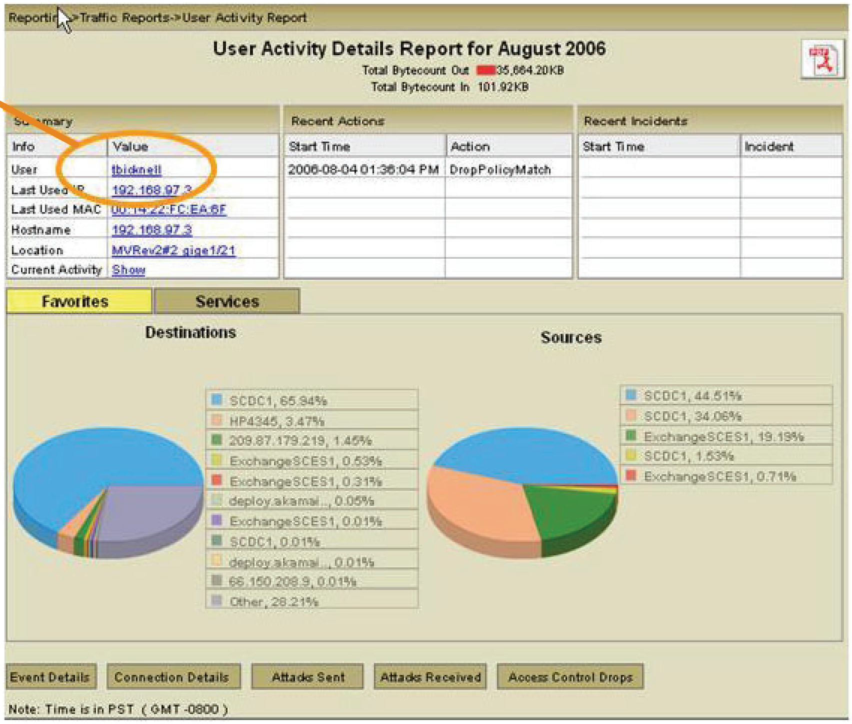

Reporting->Traffic Reports->User Activity Report

User Activity Details Report for August 2006

Total Bytecount Out 35,664.20KB
Total Bytecount In 101.92KB

Summary

Info	Value
User	tbicknell
Last Used IP	192.168.97.3
Last Used MAC	00:14:22:FC:EA:6F
Hostname	192.168.97.3
Location	MVRev2#2 gige1/21
Current Activity	Show

Recent Actions

Start Time	Action
2006-08-04 01:36:04 PM	DropPolicyMatch

Recent Incidents

Start Time	Incident

Note: Time is in PST (GMT -0800)

Figure 4. – User activity report showing historical activity for resource access (source and destination)

The LANenforcer product family

The LANenforcer (LE) family of products offers full stateful user identity aware enforcement within the network. The LE 1048 secure switch is a fully functional, 48 port 10/100/1000 secure switch. The LE 2024 is a 24 port, transparent, bump-in-the-wire security device that functions as 12 secure port pairs. Both devices offer full stateful policy-based enforcement of the user's policies as configured by the administrator. Every packet undergoes deep packet inspection to ensure full compliance with the user's policies, and is also scanned in real time for threats. The devices track user history and use that information to detect anomalies in the usage patterns of each user. All activity is correlated back to the user so that both policy enforcement as well as threat detection is localized to the user level within the LE. The LE family of products operates with full security at 10 Gbps speeds enabling, for the first time, a new breed of products capable of being deployed in-line within the LAN.

The LANsight product line

The LANsight (LS) appliances function as a central policy store for defining user and group based policies, as a centralized management platform for managing multiple LEs, as well as a centralized reporting platform that correlates events from all the LEs to bring a network wide view into both the user and the network's usage and traffic patterns. Every new flow established by the user is reported by an LE to the centralized LANsight appliance. All statistics and reporting functions are contextualized to the user and the user's activity. The LANsight platform can integrate with the existing AAA infrastructure as well as existing identity management platforms within the network.

The LS platforms offer intuitive and easy templates for defining policies and associating policies to users and groups. Policies are defined in terms of actual user traffic and behavioral patterns, from the network to the application layer, that make the policy definition piece intuitive and logical. The LS also offers extensive support for querying and reporting that can be use for audit and compliance.

Together the LE and LS form a complete identity-based policy management and enforcement system that completes the policy lifecycle and allows a consistent view of user identities to be carried through from the definition phase, to the enforcement phase, and the monitoring/reporting phase.

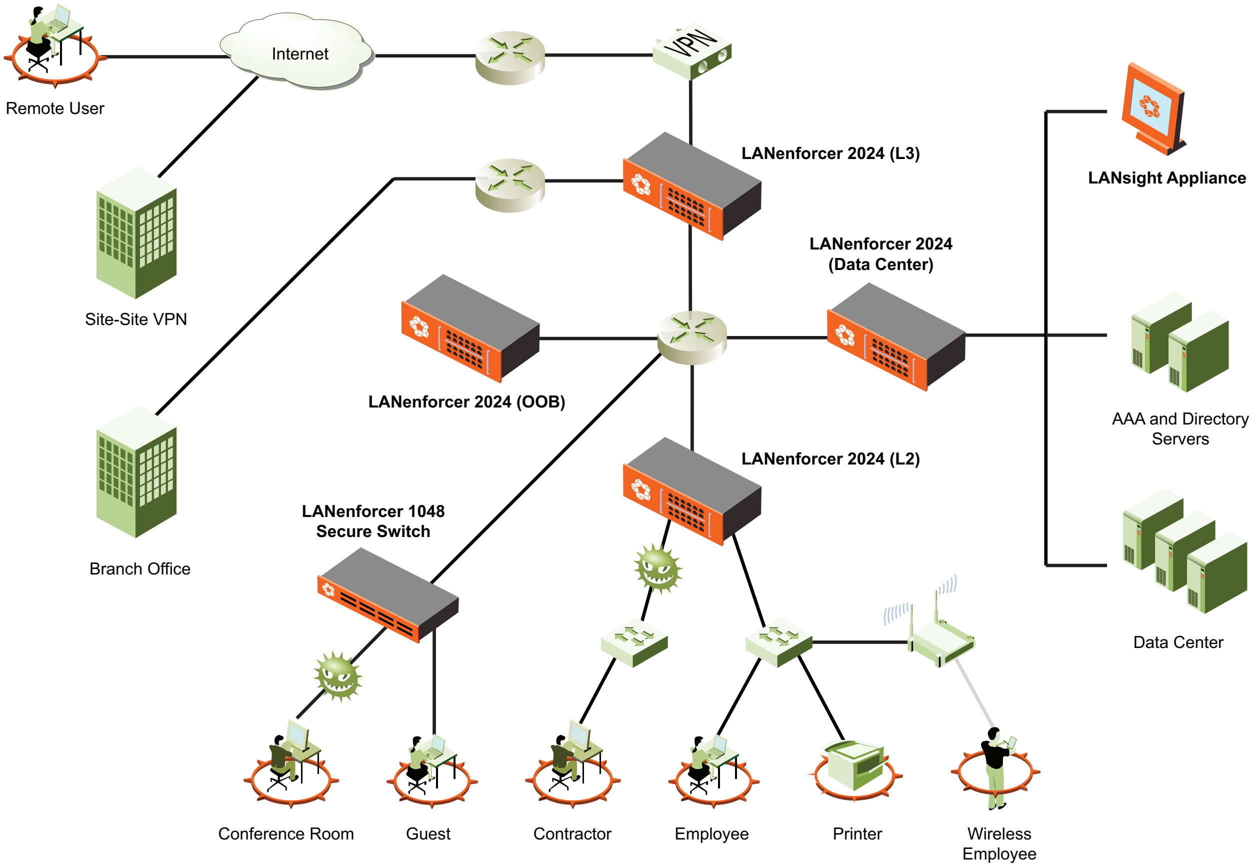

Figure 5. – The Nevis Product line can be deployed at various locations throughout the network, depending on where security policies are to be enforced. The LANenforcer 1048 secure switch is an access layer solution, whereas the LANenforcer 2024 can be deployed at layers 2 or 3, in front of the datacenter, or in an out of band mode. The LANsight management console appliance provides a network-wide view and correlation of all LANenforcer appliances.

Nevis Network, Inc.
295 Bernardo Ave., Suite 100
Mountain View, CA 94043
www.nevisnetworks.com

International:

Nevis Networks
Delegate House
30 Hart Street
Henley On Thames
RG9 2AL

Nevis Networks India
C301 Pune IT Park
Bhau Patil Marg
34 Aundh Road
Pune 411020, India