



WHITE PAPER

CORE ELEMENTS OF THE CISCO SELF-DEFENDING NETWORK STRATEGY

Thanks in part to a Cisco® advertising campaign to introduce the Cisco Self-Defending Network (such as the TV commercial in which a young girl unknowingly downloads a virus onto her father's computer), many consumers are becoming aware of the need for integrated network security.

But can a network really defend itself?

The short answer is, "Yes, it can." Network security has evolved from independently deployed products such as firewalls into the realm of system-wide solutions. And Cisco Systems® is at the forefront of the technology development that is making self-defending networks a reality.

The reason is simple: For today's companies, especially in this era of regulatory activity, preserving the integrity, confidentiality and longevity of corporate information is critical to success. As we move further into an information-driven global economy, the value of information, and controlled access to that information, has never been greater. The goal of IT infrastructure therefore is to create systems that can detect and protect against unauthorized access while providing timely access to legitimate users. Simply denying access in the face of an attack is no longer acceptable. Today's networks must be able to respond to attacks in ways that maintain network availability and reliability and allow a business to continue to function. In many respects, the goal of security is to make networks more resilient by making them more *flexible*. Rather than succumb, networks must be able to absorb attacks and remain operational, much in the same way the human immune system allows us to keep functioning in the presence of viruses and related bacterial infections.

This paper outlines the rationale for the Cisco Self-Defending Network, its foundation, and the incremental approaches Cisco Systems has adopted to deliver these capabilities.

THE CHANGING LANDSCAPE OF SECURITY

Whether we like it or not, the future of security technology has changed more in the last three years than it did in the prior 10. The extent of these changes, as well as the rate of change, has made it difficult for security IT departments to keep up. Before we can regain control, we must better understand this changing landscape.

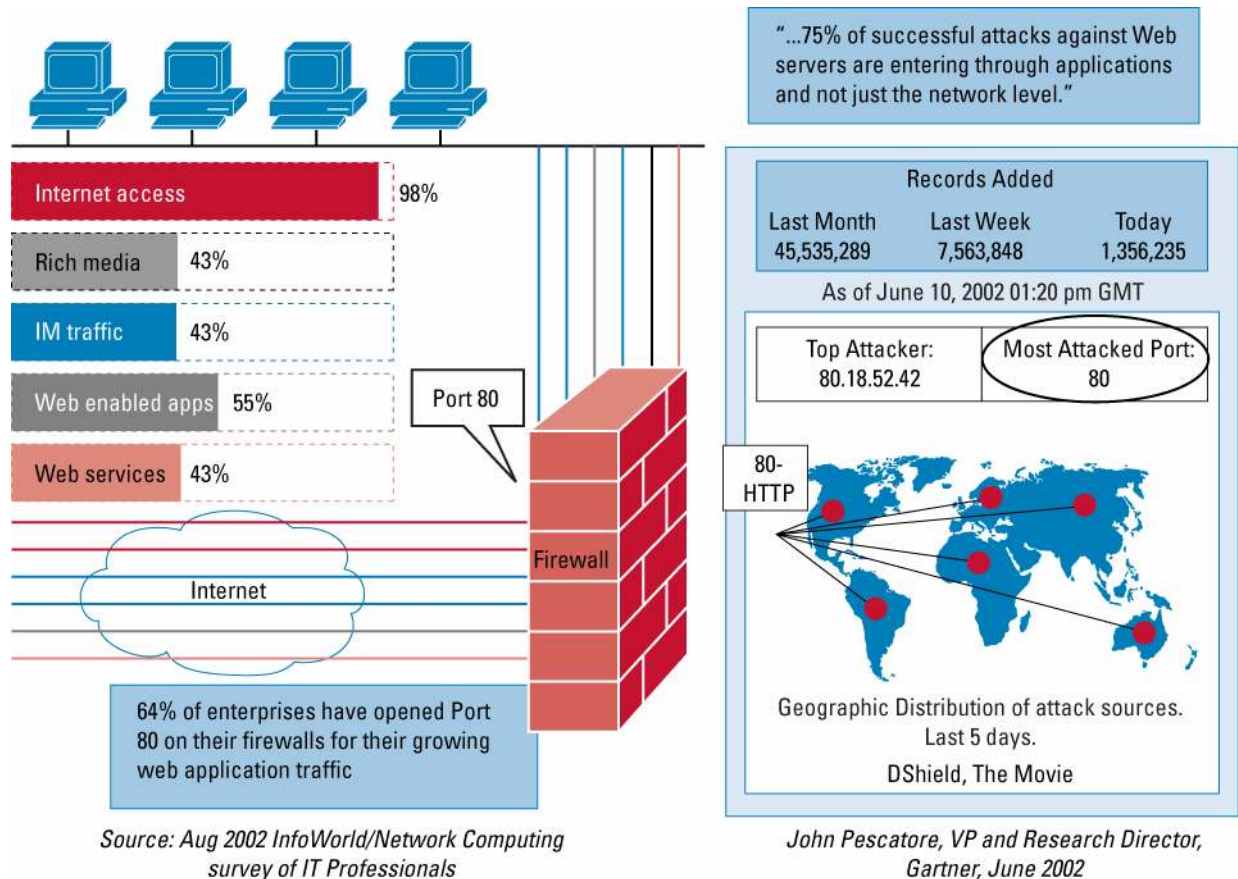
The Secure Network Perimeter—Perhaps one of the greatest changes to the industry's approach to network security has resulted from the changing nature of the network itself. No longer can a network be secured by simply securing the network perimeter; as corporations have consolidated their data centers, converged internal networks, and embraced the Internet, what was once a self-contained, controlled environment is now typically open to partners through business-to-business extranets, retail outlet connections, and home-based employees, to name but a few examples. Extending the corporate network in this way extends the trust boundary across untrusted intermediate networks and into uncontrolled environments. Devices that connect into the corporate network through these pathways are frequently not in compliance with corporate policies. And devices that are compliant frequently are used to access other uncontrolled networks prior to connecting into the corporate network. As a result, devices on these external networks can become conduits for attacks and related misuse.

Wireless and Mobility—Tied to the notion of a secure perimeter, the wireless and mobile network within enterprises now supports laptop PCs, personal digital assistants (PDAs), and mobile phones that have more than one network connection. These multihomed hosts are capable of establishing ad-hoc wireless networks to enable peer-to-peer communication. In addition, packets can effectively be forwarded across devices

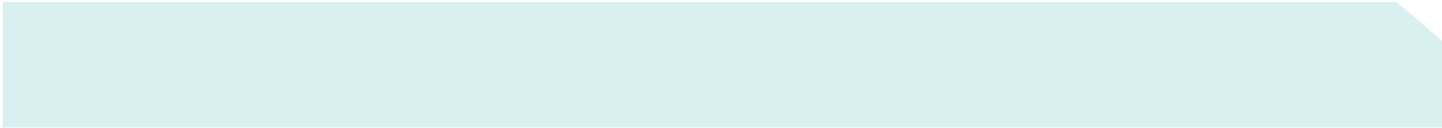
at the application level. As a result, where a network boundary begins and ends becomes much more ambiguous. Corporations need to be able to extend a control point onto these mobile devices in order to manage secure system and maintain network availability.

E-Commerce, Extranets, and Conducting Web-Based Business—The emergence of common application interfaces based on messaging protocols—such as Extensible Markup Language (XML) and Simple Object Access Protocol (SOAP)—has been a boon to e-commerce and corporate productivity. But as with most new technologies, these new message protocols have introduced an entirely new set of vulnerabilities and attack vectors with which corporations must contend. Data that was once spread across multiple network protocols and could be fairly easily filtered through firewall policies is now combined within a few, if not a single transport protocol (such as HTTP on TCP Port 80). As a result, much of the data that used to reside in packet headers now resides in the packet payload. This creates significant processing challenges that make it easier for an attacker to evade classic network defenses (Figure 1). Further, in order to meet corporate data confidentiality and integrity requirements, more and more of this application-level traffic is now being encrypted through the Secure Socket Layer/Transport Layer Security (SSL/TLS) and HTTP Secure socket (HTTPS) protocols. A side effect of this trend is that it makes it much harder for IT departments to enforce corporate access policies at the network edge because they cannot inspect the packet payloads of those encrypted flows.

Figure 1. Networks Face New Vulnerability Through Port 80



Virus, Worms, and the Rate of Propagation—The number and variety of viruses and worms that have appeared over the past three years is daunting in its own right. But two factors have had enormous impact on businesses and their operational efficiency: the shrinking window between the time a vulnerability is detected and the time an exploit appears, and the *rate* at which many of these attacks spread across an enterprise. This has led to unacceptable levels of business outages as well as expensive remediation projects that consume staff, time, and funds not originally budgeted for such tasks.



Regulatory Compliance—Well-publicized breaches and internally-generated corporate malfeasance have forced regulatory bodies in many industries to create rules for corporate information risk management. In the United States, the resulting regulations—the most well known of which are the *Sarbanes-Oxley*, *Gramm-Leach-Bliley* (GLB) and *Health Insurance Portability and Accountability Act* (HIPAA)—have forced fundamental changes in the manner in which corporate networks, servers, databases and hosts are organized.

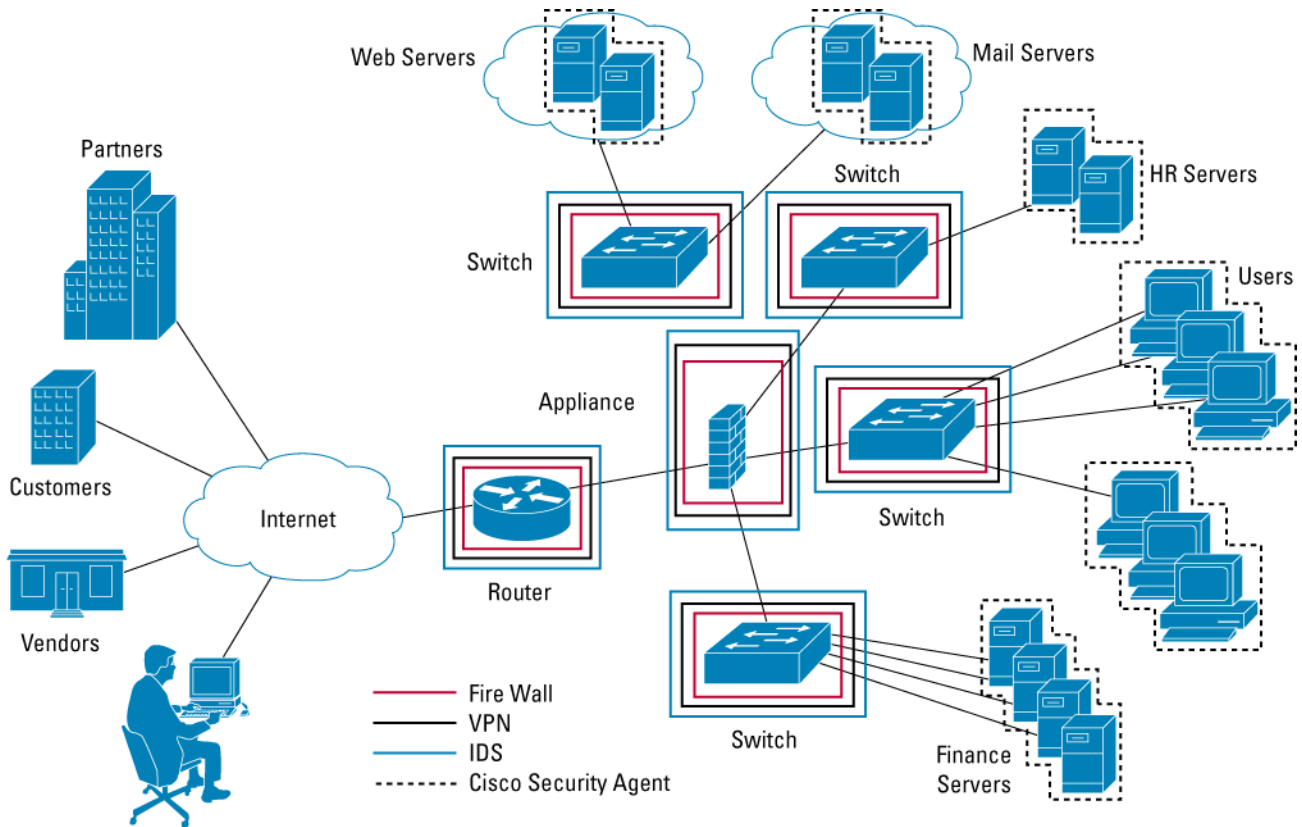
Although many organizations mistakenly assume that if they comply with regulations, their infrastructure is more secure, this frequently is not the case. Following the law of unintended consequences, the very act of creating compliance may introduce new vulnerabilities. For example, worms and viruses may spread more effectively in a network supporting end-to-end VPNs, given that the intermediate nodes have no visibility into the traversing traffic. Such traffic may carry worms to sensitive corporate servers in a secure, encrypted packet. In addition to taking longer to diagnose such an attack, these end-to-end VPNs can make it more difficult to remediate the problem.

TENETS OF MODERN, SECURE NETWORKS

Corporations attempting to navigate this changing security landscape can only accommodate so much change before it becomes operationally unacceptable. Ideally, security enhancements should have a minimal impact on existing routing and switching infrastructure, segmentation and access control techniques, and the related organizational structures that support these systems. This section explains the foundational elements of a Self-Defending Network in support of this goal: *Presence*, *Context*, *Linkages*, and *Trust*.

Presence—Fundamental to a secure system is the concept of control points, which we will define as *presence* (Figure 2). Much like our bodies' immune systems rely on detectors and responders distributed throughout the body to achieve presence, a network relies on the availability of certain capabilities within discrete nodes on the network. These capabilities include classic identity, access control, data inspection, and communication security technologies, as well as newer application-aware capabilities that deal with the growth in the exchange of peer-to-peer content, Web services, voice services, and dynamic mobile content.

Figure 2. Presence in the Self Defending Network



Context—When a user signs on to the network, the network requests and gains access to a set of credentials for both the user and the host that constitute an endpoint *entity*. It is important to note that these credentials may change over time in response to the host’s actions while connected to the network. Taken together, this information represents *context*. Whereas existing network security systems tend to focus only on permissions at the time a user enters a network, a Self-Defending Network grants or revokes permissions based on changes in behavior and associated context for the duration of the entity’s association with the network. For example, if the network detects that a host has been infected by a virus, it reacts by quarantining the host into a remedial network segment. Because information can be spoofed, securing a system may require obtaining context from other systems in order to accurately assess the host’s rights and privileges at a given point in time.

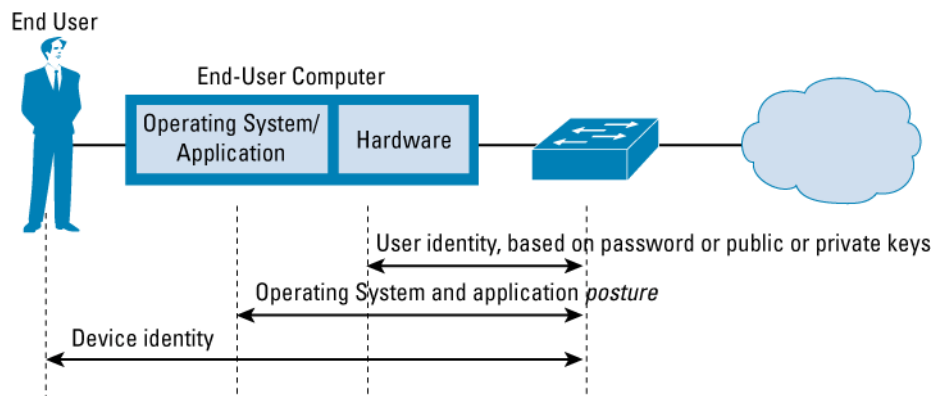
Linkages—Linkages between discrete entities is what allows for sharing of context and is what creates a “system”. Traditionally, networks have established linkages between devices through routing protocols such as the Border Gateway Protocol (BGP). In order to deal with the latest forms of threats and misuse we must now extend these network linkages all the way to the source and the destination of network traffic. In addition, given the growing presence of multi-homed mobile devices, these linkages have begun to cross demarcations until recently viewed as outside the purview of traditional networks. The privileges an entity receives upon entry to a network, and how they may change over the duration of that connection, are tied to the context of that entity and its linkages into a network, or networks.

Trust—A secure system is only as good as the information that is put into it, and it functions much more effectively when accompanied by comprehensive *trust* relationships. In the past trust has been tied primarily to the identity of a device or user. Recent advances have shown that secure systems must be augmented to include understanding the state or *posture* and location of a device.

In many respects, the activities of users and devices within a network can be compared to how we drive an automobile on the road. Just as we obtain a drivers’ license that entitles us to operate a certain class of vehicle, users must possess some type of identity in order to log in to a network. And

just as a car contains a Vehicle Identification Number (VIN) and must be registered to a particular municipality, network and endpoint devices will all soon contain a digital certificate installed at the time of manufacture and then require some type of registration when deployed within a company. But because the appropriate credentials may not always be present at a given location or point in time, a Self-Defending Network uses innovative forms of *inferred trust* and *best effort* to authenticate and authorize an entity. At minimum, a Self-Defending Network must be able to acquire credentials for each device and user identity, must be able to assess device posture, and ascertain the location of the device in the environment (Figure 3). The technology required to do this ultimately will be ubiquitous and will be enabled through well-defined, standards-based message formats and protocols such as 802.1x and Extensible Authentication Protocol (EAP) methods.

Figure 3. Credentials are Fundamental to Network Security



None of these concepts are especially noteworthy in their own right. But they become very powerful when combined in the Cisco Self-Defending Network. In the remainder of this document, we will show some of the ways these concepts work within the Self-Defending Network framework.

THE NECESSITY OF THE SELF-DEFENDING NETWORK

Corporate networks, and the attacks used to exploit them, are no so complex that no single mechanism can be relied upon to keep them secure. This has led to the concept of “Defense in Depth.” Until recently, this concept had been built on the notion of *proactive* defenses. But given the type of vulnerabilities and attacks that have accompanied our ever-changing networks, Cisco Systems believes in building better *adaptive* solutions. As a result, Cisco has begun looking to other, seemingly unrelated real-world examples such as our immune system as a model for the Self-Defending Network. Other real-world systems that have also proven to be instructive can be found in the field of epidemiology and in the way communities police themselves. A common theme with all of these systems is that they employ *adaptive* as well as *proactive* defenses.

Taking this observation a step further, the defenses to guard a system of this nature tend to be built in to every functional block. The key abilities of these adaptive defenses are that they:

- Remain active at all times
- Perform unobtrusively
- Minimize propagation of attacks
- Quickly respond to as-yet unknown attacks

These systems are built on the premise that resources are finite and must be marshaled carefully to avoid resource exhaustion. They also are designed to take full advantage of existing infrastructure with a minimum of disruption to a customer’s IT operations.

The Self-Defending Network from Cisco Systems provides systems-based solutions that allow customers to use their infrastructure in new ways to reduce windows of vulnerability, minimize the impact of attacks, and improve overall infrastructure availability and reliability. It is also helps create autonomous systems that can quickly react to an outbreak with little to no human intervention. This type of rapid response is required to thwart the latest forms of misuse that are much more virulent than their predecessors.

The Cisco Self-Defending Network continues to improve its ability to respond to new threats. The first phase, **Integrated Security**, incorporates security elements in network elements such as switches and routers. The second phase, **Collaborative Security**, involves building linkages between network security elements and extending the network presence out onto endpoints that connect into a network. The **latest phase** of the Cisco Self-Defending Network introduces **Adaptive Threat Defense (ATD)** capabilities which enhance the ability of a network to respond to threats based on a new set of *Anti-X* technologies.

THE BUILDING BLOCKS OF THE CISCO SELF-DEFENDING NETWORK

Most customers will not adopt all of the components of the Cisco Self-Defending Network at one time, as it may be difficult to overhaul all of the required subsystems at once without disrupting the integrity of the IT services. Some customers may hesitate to turn over security controls to an automated system until they are confident the system will operate dependably. The Cisco Self-Defending Network initiative deals with these concerns by first providing products that can be usefully deployed independently of one another and then by offering solutions that can link these products together as confidence builds in each product and subsystem—a successful approach based on a combination of product development, product acquisitions, systems development, and partnering. With this in mind, it is worth reviewing the key milestones of the Cisco Self-Defending Network initiative to date.

Endpoint Protection—One of the realities of viruses and worms is that they frequently create network congestion as a byproduct of rapid propagation as well as infection of endpoints. Cisco realized it could begin addressing both problems by offering its customers an endpoint intrusion prevention called the Cisco Security Agent. Cisco Security Agent uses novel forms of behavioral security to detect and prevent viruses and worms from gaining a foothold on an endpoint system, and prevents these viruses and worms from propagating across a network. In effect, Cisco Security Agent becomes a *first order dampener* to the virus and worm propagation effect.

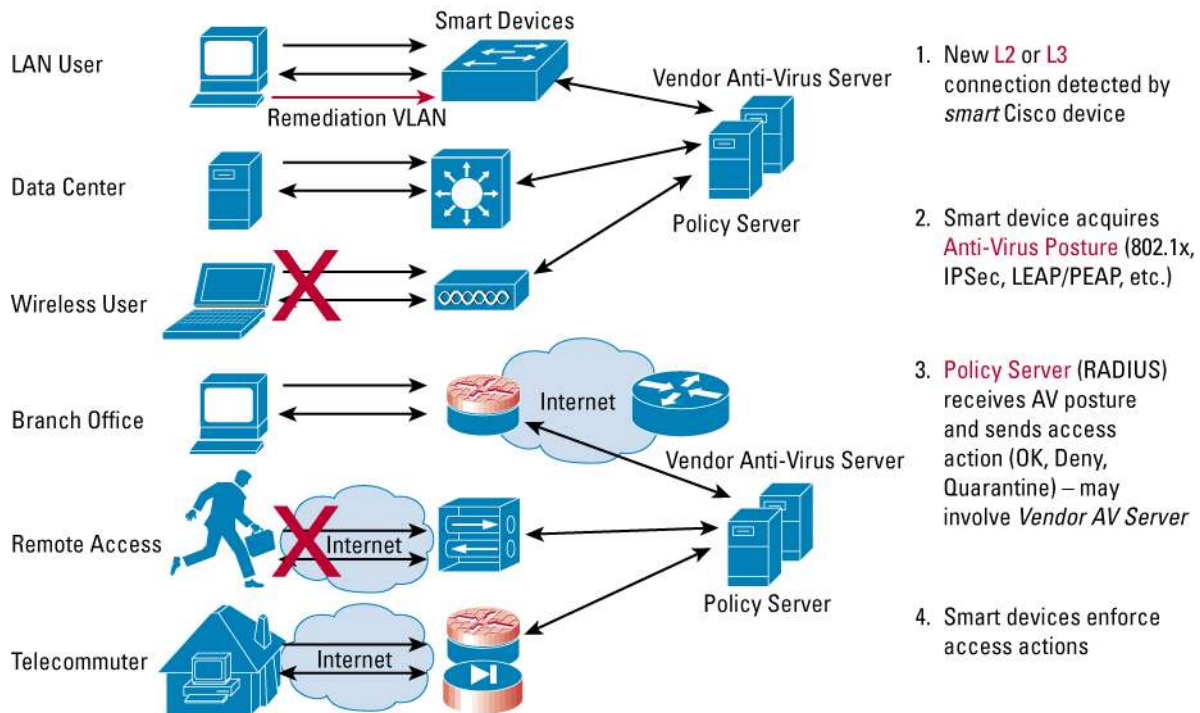
A second and equally compelling reason for deploying Cisco Security Agent is that it establishes a presence on endpoints that can be used to establish a feedback loop between the endpoint and the network resulting in a network that rapidly adapts to emerging threats.

Admission Control—One of the most high-profile Cisco Self-Defending Network initiatives to date is the Cisco Network Admission Control (NAC) program. NAC allows customers to determine what level of network access to grant to an endpoint based on its security *posture*, which is based on the security state of the operating system and associated applications. In addition to controlling access, NAC gives IT administrators a way to automatically quarantine and remediate noncompliant endpoints. Making sure endpoints are compliance with OS patches and antivirus software updates is an effective *second order dampener* to the virus and worm propagation effect. Another way of looking at NAC is as an *on-demand* vulnerability assessment and patch management tool.

A distinguishing feature of NAC is that it provides both client and backend AAA interfaces that allow customers to plug in products from their preferred endpoint security and policy vendors to drive the NAC plane.

Today, more than 30 industry-leading vendors are actively integrating NAC into their technologies.

Figure 4. Controlling Admission into the Network



It is important to extend NAC capability to small and medium-sized businesses (SMBs). To that end, Cisco recently acquired Perfigo, Inc., a developer of packaged network access control solutions that provide endpoint policy analysis, compliance, and access enforcement capabilities and now offers the bundled compliance solution under the name Cisco Clean Access to this market segment.

Infection Containment—Strong network admission policies are not a cure all, and do not eliminate the need to continue monitoring devices once they enter a network. Determined attackers can evade just about any admission check, and the network cannot always rely on, or *trust*, an infected element to turn itself in. Compliant devices also can become infected through a variety of vectors once they are members of a network—e.g., a USB key with infected content. To further help protect the network, Cisco Self-Defending Network is designed to extend the security checks performed at the time of admission for the duration of the network connection. In addition, the Self-Defending Network can rely on *other* network elements, including other endpoints, to detect when another endpoint(s) is no longer trustworthy, much as a police force monitors crime in a community via a 911 call center. Cisco regards Infection Containment as a *third-order dampener* to the virus and worm propagation effect.

Existing authentication protocols are not designed to work beyond the initial exchange, however. So the Self-Defending Network must provide new ways of communicating the state of a device (context) and new ways of gauging the veracity of that information based on inferred as well as direct forms of trust. For example, an administrator can create a rule stating that notification received from an endpoint running Cisco Security Agent is more trustworthy than a notification from an unprotected endpoint. As a result Cisco has begun development of new types of correlation and feedback based on inferred attributes.

Intelligent Correlation and Incident Response—In order for steady-state feedback mechanisms such as Infection Containment to work effectively, the Self-Defending Network needs to provide services such as real-time correlation of events, quick assessment of the security impact of an event, the ability to decide what action to take, the ability to identify the closest control point to implement a response, and more. To achieve this, Cisco recently announced the acquisition of Protego Networks, whose MARS family of products provides methods for overlaying feedback from a variety of points of presence (POPs) in the network (firewalls, network intrusion detection systems [NIDS], routers, switches, and hosts) with context it obtains from learning the L2 and L3 network topology. This ability helps the Security Incident Response team to rapidly identify where attacks are occurring in the network.

Cisco is also working with netForensics and other partners to enhance their correlation capabilities to better audit the Self-Defending Network.

Inline IDS and Anomaly Detection—An important area of ongoing security development at Cisco has been in the area of network intrusion detection systems (NIDS). One of the first Cisco innovations in this area was to integrate NIDS into its router and switching platforms. But in order for NIDS to fully deliver on its capabilities it needs to transform into an intrusion *prevention* system (IPS) with inline filtering capabilities. This provides a mechanism to remove unwanted traffic with fine-grained programmable classification engines.

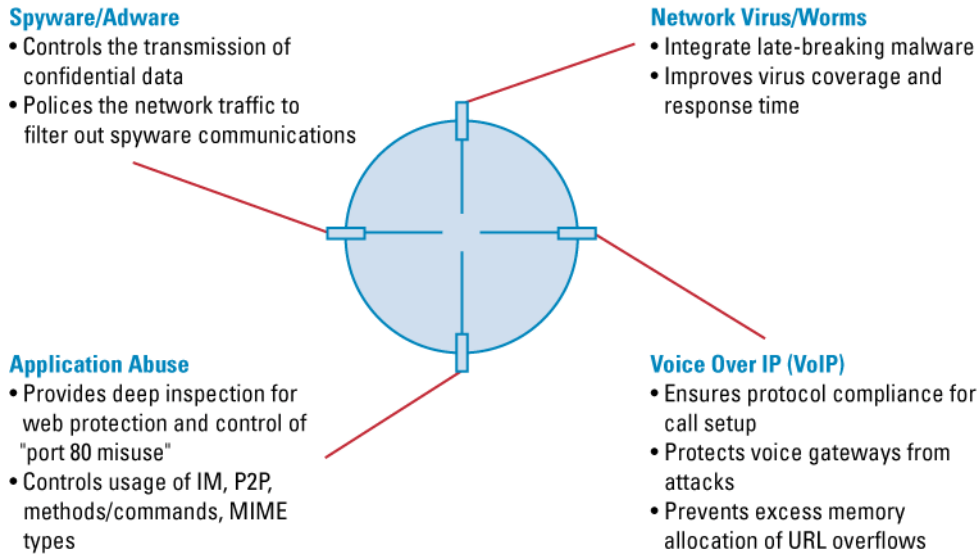
Unfortunately most NIDS generate too many false positives to reliably operate as inline security services. Part of the problem is that a large amount of information (context) must be assembled and processed within a fairly short time window—particularly with the amount of context that must be amassed for application-based protocols. This is especially true for applications such as IP Telephony, which is very sensitive to packet transmission delays. To address this, Cisco is developing several methods that provide high-fidelity signaling to these inline classification engines.

Many legitimate activities can be mistakenly regarded by the network as anomalous, particularly in networks where the number of variables is significant. As a result Cisco intentionally adopted a conservative, incremental approach to anomaly detection, beginning with Cisco Security Agent, because it was acknowledged that operating systems are easier to model than network environments. Cisco then acquired Riverhead, which is effectively an inline prevention system with low false positive rates because denial of service (DoS) activity clearly stands out from other network activity. DoS attacks can be identified with greater accuracy by deploying the technology Cisco gained through its acquisition of Riverhead network. This again results in low false positives when identifying a DoS attack.

Based on lessons learned from Cisco Security Agent and the DoS capabilities acquired with the Riverhead family of DoS Guard products, Cisco is introducing an IPS that reduces the false positive rate through the application of innovative anomaly detection techniques and the sharing of state (context) between endpoints and network elements (linkages). It also provides multivector threat identification and meta-event correlation to assess both the vulnerabilities and exploits quickly. By incrementally introducing these technologies into the market, Cisco establishes greater customer confidence in its capabilities and by extension the Self-Defending Network.

Application Security and Anti-X Defense—Over the past several years, a number of new application-layer network products have emerged to help address new classes of threats that were not adequately addressed by classic firewall and NIDS products, including *viruses and worms, e-mail based SPAM and phishing, spy-ware, Web services abuse, IP Telephony abuse, and unauthorized peer-to-peer activity* (Figure 5). Cisco has developed the next generation of packet- and content-inspection security services to deal with these types of threats and misuse. This convergence brings granular traffic inspection services to critical network security enforcement points, thereby containing malicious traffic before it can be propagated across the network.

Figure 5. Enhancing Security through Anti-X Multivector Threat Identification



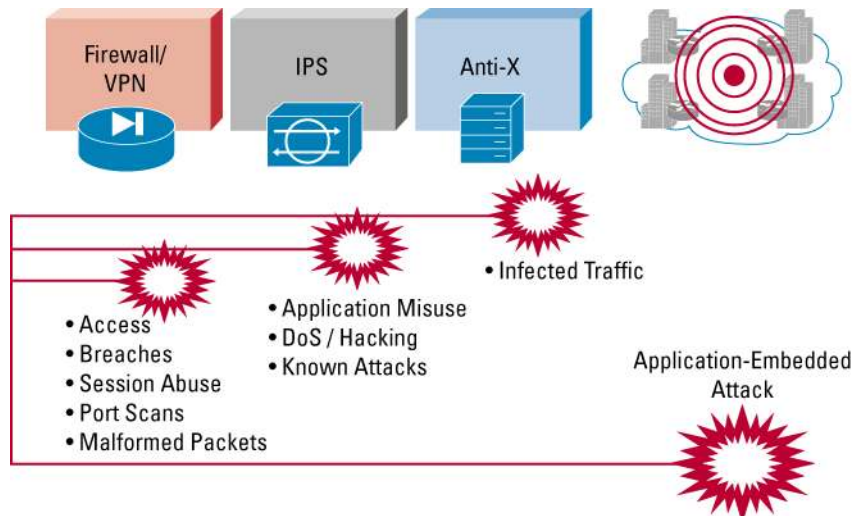
The consolidation of these services onto multiservice platforms represents an opportunity for vendor innovation as well as a reduction in customer cost of ownership. The Self-Defending Network will also become more application aware once these services are integrated into its framework.

Note that when applications employ end-to-end encryption, the Self-Defending Network can collect information from endpoints to compensate for the loss of visibility at the network edge.

NEXT STEPS

Cisco will continue to invest heavily in building Self-Defending Networks that create linkages between POPs across networks up to and including endpoint systems (Figure 5). By doing so, Cisco helps organizations to obtain greater visibility and control of devices, users, and applications that communicate across their infrastructures. This is a necessary and important evolution of what is classically expected of the network, much like the introduction of intelligent routing protocols were to packet forwarding through a communication fabric.

Figure 6. Cisco Continues to Develop the Abilities of the Self-Defending Network



The brief insights this paper provides on Cisco Self-Defending Network should be explored in more depth, both to better understand what is possible today and to lay the foundation for future security and network design projects. The choices each organization will make depend on the unique security, risk, and compliance issues they face:

- For perimeter security staff, the recently announced Cisco PIX[®] 7.0 Firewall and Integrated Service Router platforms offer detailed data inspection and control across a wide variety of protocols and their accompanying attack vectors, as well as many other security and networking features.
- Security operations groups that spend unreasonable and unbounded effort responding to incidents should investigate the newly acquired technology from Protego, as well as learn more about the new capabilities of inline intrusion protection in infrastructure and security appliances.
- Those responsible for high-data environments that deal with frequent DoS and DDoS attacks should review the Anomaly Guard technology acquired from Riverhead.
- For organizations that continue to be effected by worms and viruses, or those that require endpoint compliance solutions, investigate Cisco Security Agent, Network Admission Control, and Cisco Clean Access acquired from Perfigo.
- Auditors responsible for assessing regulatory compliance should also investigate NAC as well as CiscoWorks Security Information Management System (SIMS), a tool that provides detailed information regarding the use of the entire communication infrastructure.

Finally, IT professionals responsible for the design and deployment of the security systems and network infrastructure should contact their Cisco representative for further details on the Self-Defending Network and how it can have a meaningful positive impact on the IT environment.



Corporate Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters

Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters

Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the **Cisco Website** at www.cisco.com/go/offices.

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel
Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal
Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan
Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2005 Cisco Systems, Inc. All rights reserved. CCIP, CCSP, the Cisco *Powered* Network mark, Cisco Unity, Follow Me Browsing, FormShare, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, GigaStack, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, MICA, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, ScriptShare, SlideCast, SMARTnet, StrataView Plus, Stratm, SwitchProbe, TeleRouter, The Fastest Way to Increase Your Internet Quotient, TransPath, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0501R) 205227.f_ETMG_KL_2.05

