

Security for the Borderless Network: Making Web 2.0 and 3.0 Safe for Business

This chapter includes the following topics:

- Security Policies for the New Open Networked World
- The Borderless Network Security Architecture
- Super-Charged Scanners
- Security Everywhere in the Network
- Collaboration with Confidence

Now is not the time to *plan* for a mobile Web 2.0 world. It's already here, like it or not. Internet-based collaborative work tools. iPhones, BlackBerrys, and growing numbers of other “smart” mobile wireless devices with multimedia and PC-like capabilities. Social networking, instant messaging, and Skype. Virtualization, cloud computing, and SaaS. Cisco WebEx and TelePresence. Remote and mobile working. These trends, technologies, and tools enable people to work far more efficiently, and be more connected to information and one another, than ever before.

The case studies and the comments from individuals featured in this book reveal that many organizations are proactively trying to find their place in the Web 2.0 world. They are truly eager to leverage new tools and technologies that can increase their competitive edge. They realize the potential these tools have to enable them to streamline their work processes, keep their workforce motivated and productive, connect with their customers, generate significant cost savings, and remain relevant in today's fast-changing global business landscape.

But we have also come to understand that two things are standing in the way of today's workers achieving even greater heights of productivity and creating tremendous value for the enterprise: management's fear and IT's inflexibility.

Many organizations—including those considered leaders in their industries—have been slow to embrace mobile and remote working and Web 2.0, or accept the inevitability of

consumerized IT. Some have rejected these powerful forces of change outright, clinging to outdated, irrelevant, and often irrational policies that do little to enhance security. Their policies also make their employees unhappy. And that discontent among the rank and file, more so than any threat that may arise from unauthorized use of Facebook or an iPhone, could turn out to be the most damaging to an organization's success and survival.

Security Policies for the New Open Networked World

As previously discussed, the evolution of technology demonstrates that when people find a tool or process they like or that makes their life easier, keeps them connected to what and whom are most important to them, and gives them a feeling of greater autonomy and empowerment, there is simply no going back to “the way it was before.” Today's workers also reject the idea that their “personal” and “professional” use of technology should not be allowed to intersect, especially because they already do—and have for some time.

The irony here is that companies are aware their current security policies are weak. They know their employees are either willingly or accidentally bending or breaking the rules to access whatever they need to do their jobs better. These violations range from the use of unauthorized web-based applications to misuse of corporate computers. Meanwhile, employers also seem resigned that there is little if anything they can do to prevent or control this behavior.

Today's organizations need a new set of tools to express and enforce a more intelligent and relevant security policy that deals with the realities of Web 2.0, social networking, and the anywhere-anytime-any device nature of the borderless enterprise. Without these tools, policies are as effective as the ugly mattress tags that warn: “Do not remove under penalty of law.” Most are removed.

At the same time, security that gets in the way of the end user will rarely be successful. Consider this classic picture: A gate is placed across a road with flat, grassy shoulders. Either side of the gate has been worn down by the tire tracks of countless vehicles that have driven *around* the barrier. That's exactly how many users view security measures. They are incredibly resourceful at finding ways to bypass such restrictions in the name of getting their jobs done or living their lives more efficiently. (Consider the admission of a chief information officer of a 60,000-person tech company in Silicon Valley, who reports the organization has 50 “official” iPhone users but 5000 to 6000 “unofficial” users who have found “backdoor ways” to connect into the corporate infrastructure.)

The most effective security solutions can actually *enhance* the end-user experience; if they can do that, users will be inclined to adopt them. To achieve this, modern security systems need sufficient intelligence and granularity, so they can enforce policies that are much more user-oriented—for example, not bluntly blocking access to social networking sites such as Facebook or MySpace, but instead, putting controls on the type of content that can be uploaded from these sites.

As we have clearly seen through the case studies presented in this book, companies create material gains of *productivity* and *efficiency* due to greater workforce mobility and the use of cloud computing and collaborative Web 2.0 tools and technologies.

Increasingly, business is conducted without borders. Thus, the security community needs to adapt and evolve policies so that businesses can embrace these new technologies while not exposing themselves to harm through increased malware infection, acceptable use problems, or data loss.

The Borderless Network Security Architecture

We have determined that security policy needs to evolve to the point where businesses can collaborate with confidence. However, before these more intelligent policies can be developed and fully embraced, the tools to enforce them must also undergo a significant evolution.

As discussed in the previous chapter, next-generation, network-based security systems need to *understand* content, applications, and end-user identity, not just the lower-level constructs of IP address, network port, or network protocol. Cisco has a vision for the next-generation security architecture that will secure the borderless network. These tools are identity-, application-, and content-aware. Simply put, they can tell the difference between CNN and Skype and Oracle. They recognize users and directory structures: John Smith is in sales and can access sales data in a cloud-based application. Joe Smith is in engineering and can access source code on an internal server.

This architecture is comprised of five major components:

- **Scanning engines:** These are the foundation of security enforcement and can be viewed as the workhorses of policy enforcement. They are the proxies or network-level devices that examine content, identify applications, and authenticate users. A scanning engine can be a firewall/IPS, a proxy, or an interesting fusion of the two. Scanning engines can run multiple layers of antimalware signatures, behavioral analyses, and content inspection engines.
- **Delivery mechanisms:** These are the mechanisms by which scanning elements are introduced into the network. This includes the traditional network appliance, a module in a switch or a router, or an image in a Cisco security cloud.
- **Security intelligence operations (SIO):** We're talking now about the "brains" that can identify good guys from bad. The Cisco SIO (described in detail later in this chapter) encompasses multi-terabyte traffic monitoring databases, thousands of servers in multiple data centers, and hundreds of engineers and technicians with a single purpose—identifying and stopping malicious traffic.
- **Policy management consoles:** These consoles are separate from the scanners that enforce policy. By separating policy creation and management from enforcement, we make it possible to have a single point of policy definition that spans multiple different enforcement points such as email, instant messaging, and the Web.
- **The next-generation endpoint:** This is the critical piece that ties everything together. The role of the next-generation endpoint is to reside on a wide variety of devices and make sure all connections coming on or off a device are routed through one of the network-based scanning elements previously described.

Big and Little: The Evolution of Endpoint Security Traditional network security consists of two major components: a heavy endpoint protection suite (antivirus, personal firewall, and so on) and perimeter-based, network-scanning devices (firewalls, web proxies, and email gateways). This architecture worked well in a world of high-powered PCs that were mainly on the LAN and behind the firewall. But in the Web 2.0 world, enterprise computing devices are changing rapidly.

As discussed in previous chapters, iPhones, BlackBerrys, netbooks, and thousands of other devices are becoming powerful substitutes or complements to the traditional PC. In these highly distributed or lightweight, portable, heterogeneous computing platforms, the traditional “antivirus” endpoint suite is no longer relevant. (In short, we don’t want a traditional AV client running on our phones.)

The Cisco security architecture for the borderless network relies on a lightweight, pervasive endpoint. Its role is not to scan content or run signatures. Instead, its sole focus is making sure every connection coming on or off the endpoint is pointed at a network scanning element somewhere in a Cisco security cloud.

These scanning elements are now capable of running many more layers of scanning than a single endpoint possibly could: five layers of malware signatures, data loss prevention and acceptable use policies, content scanning, and more.

The endpoint of tomorrow won’t be an antivirus suite, but an intelligent connection manager that sits on the edge of every device imaginable. It is the new perimeter of the de-perimeterized network.

Super-Charged Scanners

The impact that multicore processors are having on modern security cannot be underestimated.

With well-written software, network security applications scale linearly on multicore systems. For example, the throughput obtained from the current generation of Cisco IronPort web security appliances with eight cores is 800 percent faster than the throughput from a single core system just a few years ago.

Multicore-powered systems can now run three layers of antivirus scanning, advanced acceptable use filtering, DLP algorithms, and reputation filtering for 10,000 users on a single two-rack unit appliance. All this processing is done with just eight cores. The security delivered in the network far exceeds what a traditional endpoint can offer: Imagine running three separate AV engines on your PC. It would cause significant disruption to the end-user experience, yet it can be done invisibly in the network.

And there appears to be no limit in immediate sight. The next revision of hardware will support 32 cores, which will be 3200 percent faster than a single core system. This massive surge in throughput far surpasses the needs of traffic increases and makes possible an entirely new class of security scanning that can provide far more sophisticated content analysis. These scanners can distinguish between different types of traffic. They will be

able to peer into an application such as webmail and understand the text in the message bodies—looking for information such as credit card or Social Security numbers—with no added latency.

Super-charged scanning elements are at the heart of the new, more intelligent policies required to embrace Web 2.0 and the borderless network safely. They can identify users and their roles in the organization and make decisions on a high level by understanding applications and content. They can enforce a policy that says, “Allow YouTube video streaming for marketing, but make sure it doesn’t interfere with WebEx sessions.”

Intelligent progressive policies give end users the opportunity to harness the power of Web 2.0, but with logical safeguards. Multicore processors running in the next generation of firewalls and web gateways are the scanning engines making these policies feasible. As these scanners evolve, the distinction between networking devices such as firewalls and web proxies will melt away. They are all simply scanners that reside in multiple points in the network.

Security Everywhere in the Network

Apple’s iPhone and Cisco TelePresence are examples of the sexy new technologies that are significantly changing the way we work and communicate. They are also transforming how the network is architected, leading to decentralized Internet access. The traditional architecture of the hub-and-spoke network with a small number of access points is giving way to the borderless network with a large number of access points. As such, modern security needs to be widely distributed across the network.

The Cisco security architecture for the borderless network relies on a spectrum of delivery mechanisms that make it possible to put security in more and more places to accommodate the decentralized, borderless network. The intelligent scanning elements previously described will run as software in a virtualized data center infrastructure. A form factor will be a traditional hardware appliance, such as those commonly used today, and a module in a Cisco networking device, such as a branch office router or a powerful data center switch. Or it will be available as an image in the steadily growing Cisco security cloud, powered either directly by Cisco or one of the many Cisco service provider partners.

Regardless of the delivery mechanism, the scanning capabilities, policy enforcement and management, and the reporting system will be consistent. A customer might choose to put appliances in its headquarters, integrated security modules for its branch office routers, and hosted cloud images for mobile users. These flexible delivery options coupled with higher-level application- and identity-aware scanners enable policy enforcement to be abstracted from the physical network. Users will get the same policy enforcement regardless of whether they are on an iPhone in India or a desktop in Denver.

Security Intelligence Designed In

As discussed in Chapter 10, “Signs of Hope,” the next generation of security intelligence requires a broader look at traffic patterns. Using techniques like Cisco’s Global Threat Correlation, the ever-sophisticated waves of attacks can be stopped based on the nature of the attacker, not just the nature of the attack. The foundation of this approach is having security telemetry—statistical data about the behavior of the network—built in to all scanning elements in a bidirectional exchange. Traffic data is sent into the Cisco SIO and new rules are pushed out, almost in real time.

The Cisco SIO is a collection of data, machines, and people that work together to identify and stop malicious traffic on the Internet. Now in its tenth year of operation, Cisco SIO is the largest traffic monitoring network in the world. It processes more than 5 billion web requests and 100 million email messages daily, sampling more than 35 percent of the world’s email traffic. It has more than 500 engineers, technicians, and researchers working in five facilities around the world. It uses more than 1000 servers and stores more than 2 terabytes of data. This sophisticated infrastructure generates more than 875,000 rules per day. That’s about 10 new rules every second of every day.

The scale of this system is remarkable—and it is growing. As Cisco adds telemetry into more networking devices such as firewalls and IPS systems, followed by switches and routers, the sample of Internet traffic will continue to expand rapidly. For every packet that traverses the Internet, there’s a good chance it will hit a piece of Cisco equipment somewhere along the way.

This huge footprint gives Cisco unique insight into global Internet traffic patterns. The Cisco SIO is collecting traffic data about every publicly routable IP address on the Internet. Even if all the SIO has to say about a server is “you are new; we have never seen you before,” that is still useful information. What it means is don’t block the server, but apply maximum content and signature inspection, and feed the results back into the SIO database. If the server is passing legitimate traffic, its reputation steadily climbs. If it is found to be passing malware, its reputation drops.

Cisco’s Global Threat Correlation, meanwhile, is the technique of analyzing reputation across multiple realms—email, web, IM, FTP, and so on. The simplicity and robustness of reputation analysis and Global Threat Correlation have enabled these techniques to stop malware, on average, 12 hours ahead of signature availability, and to boost the catch rate of a Cisco IPS system by more than 300 percent. These sustained advances are a result of the broad context that reputation provides.

If a security scanner is attempting to analyze a traffic flow, it is difficult to analyze the bits going by and make an accurate determination based solely on the data traveling past. However, if the system can “look” at where the bits are coming from and going to, make assessments of the client and the server, and identify the application and the content being transmitted, a more accurate decision can be made. This is similar to the manner humans use to analyze threats. When looking narrowly at a brown paper-wrapped package, it can be difficult to know if it is good or bad. But looking at where it is being sent,

where it is from, who is delivering it, and what's inside of it provides a much more complete picture. The Cisco SIO provides a comprehensive view of the Internet-facing side of any transaction.

In the borderless network security architecture, all Cisco security devices will have this contextual capability built in. This means every device will know more about who is sending and who is receiving the traffic, what application it is, and what the content is. With this broad contextual information, the scanners can make far more accurate decisions, stopping even the stealthiest threats.

It also means that all Cisco security devices will be tethered to the SIO. Participation in this network will always be opt-in and default off. But after nearly a decade of operations, experience has shown that the participation rate will be extremely high, keeping the global defense systems of the Cisco SIO pulsing with new data feeds and more accurate rule generation.

The Line Between Policy and Enforcement

Today, most policies are developed around a corporate directory. Marketing can have one set of privileges, engineering another. Individuals can have policy exceptions tied to their presence in a directory server. Directories are a foundational capability in any security policy, yet they tend to be static and lack context. Directories answer the fundamental question of *Who are you?*

As we move to a borderless business, we need security policies that will have a broader context. We want to know not just who you are, but also where are you coming from, when you are attempting to access something, what device you are doing it from, and how you are attempting to access a particular piece of information. Therefore, we need to transition from the “who” of a static directory to the “who, what, where, when, and how” of a borderless network security system.

If each scanning element in the network needs to assess the broad context of who, what, when, and how, collecting and coordinating this information becomes unwieldy. It makes more sense to create a new policy server that stores these dynamic attributes and makes them available to a variety of different devices—firewalls, web proxies, endpoints, and NAC servers. Then, each of these devices can make a more robust contextual enforcement decision from a centralized, dynamically updated set of data.

This architecture has another advantage: It allows an IT team to create policies that span multiple different enforcement points. Thus, a policy can be created around access control that would be available to both a firewall and a web proxy. Or a policy can be created around credit card numbers that would span email, web, and IM. Today, this exchange of policy happens manually. In the borderless network security architecture, the policies flow from the centralized policy store through the operations console of the individual devices.

The Cisco vision is to make the interface between policy and enforcement systems open and built on industry standards. Therefore, if a customer chooses to use a Cisco application entitlement system and an enterprise DLP policy manager from RSA, for example, both should work smoothly with existing network infrastructure.

This is not a small undertaking, however, and it will take years to fully develop. However, centralized policy management and coordinated, dynamic policy enforcement provide the flexibility needed when deploying intelligent policies in the borderless enterprise.

Redefining the Endpoint

A critical component of the secure, borderless network is the security system on the endpoint. As we move to a world in which mobile devices take up a larger percentage of our enterprise computing time, the traditional antivirus client suite on the endpoint doesn't go away, but it does become less relevant.

Again, the role of the next-generation endpoint is not to scan content or run signatures. It simply manages connections and makes sure all content coming on or off a device is connected to one of the scanning elements deployed in one of the form factors and running the policies of the system previously described. The client ties the whole security system together.

The next-generation connection manager takes care of all network connections, enhancing the end-user experience. From an end user's perspective, he or she is always on the LAN—it all “just works.” When a user is behind the firewall, the connection manager senses this and does not attempt to create a VPN tunnel or connection. Instead, it takes care of basic network functions such as the 802.1x authentication required to make wireless and wired Ethernet switches work smoothly. It also handles posture checking to make sure the client has up-to-date patches and AV signatures.

When a user closes her laptop at work, and then goes home and opens it back up, the connection manager “wakes up” and realizes it is no longer behind the firewall. It then automatically finds the nearest network attach point. Web traffic gets pointed to the nearest web proxies, application traffic is tunneled via IPSec or SSL VPN into the nearest remote access concentrator, and voice and video traffic are pointed at the Session Initiation Protocol (SIP) gateways.

These attach points can be any of the form factors previously described: an appliance at corporate headquarters, a module in a branch office router, or an image in the worldwide Cisco security cloud. All this is invisible to the end users; they know only that it works and it's easier than before. No more fumbling with passwords and authenticating repeatedly. No more struggling to establish a VPN connection.

From the perspective of end users, they feel that they are always on the LAN. From IT's perspective, they now have control and policy enforcement no matter where an end user might go—on a PC behind the firewall or on a smartphone in Timbuktu. The borderless network security architecture offers consistent policy enforcement, reporting, and

inspection. This is a perfect example of improving security while also improving the end-user experience.

The connection manager, working with the other components of the borderless network architecture, yields a next-generation security architecture that enables advanced high-level policy enforcement across the network, independent of the physical infrastructure underneath. It understands users, applications, and content. It has the robustness to make security decisions in the broadest context, looking at the who, what, where, when, and how of a transaction. It has policy that is managed separately from devices. And it has an endpoint solution that works on every major enterprise computing platform—Windows, Macs, handhelds, and tablets. It is a complete solution for the borderless enterprise.

Collaboration with Confidence

Although we have yet to experience Web 2.0's full potential, or to know exactly what working and living in a Web 3.0 world will be like, economists, historians, and industry leaders all agree that handheld computing and advanced collaboration technologies will drive another decade or more of productivity enhancements.

In a recent report, global financial services firm Morgan Stanley concluded that the rise of the mobile Internet is likely the fastest adoption of a new technology in history.¹ Two years after its launch, the mobile Internet has eight times the adoption rate that AOL's dial-up Internet service had two years after its launch. This massive adoption is having a profound impact on the way we work and share information, and how our networks will be designed moving forward.

At the same time, new forms of collaboration such as Cisco TelePresence are forcing companies to rethink how they design networks and where they define the “perimeter.” And the highly compelling economics of cloud computing mean more of our data and applications are now—or will soon be—on the move.

All these important trends are driving what will be a substantial change in the way that security policies are expressed and enforced in the future. And the underlying technology that powers advanced security—including multicore processors and Global Threat Correlation—is having a massive impact on our ability to create and enforce more intelligent, robust policies.

The driving force for all this change? The relentless corporate need for competitiveness. Companies are finding significant advantages from embracing mobility and Web 2.0 technologies. But the security policies and the technologies underneath it all need to evolve to keep pace with change, so companies can take advantage of these new capabilities while still maintaining basic controls.

The security architecture for the borderless network is what will finally enable businesses large and small to open their networks, transcend traditional borders, and collaborate with confidence.

Endnote

¹ “The Mobile Internet Report,” Morgan Stanley, December 15, 2009.
<http://www.scribd.com/doc/24129386/The-Mobile-Internet-Report>.

Reference

Cisco 2009 Annual Security Report.
http://cisco.com/en/US/prod/collateral/vpndevc/cisco_2009_asr.pdf.