WHITE PAPER

The Network Security Architecture (NSA) Meets Web 2.0

By Jon Oltsik
Principal Analyst

June, 2009

Table of Contents

Executive Summary

The first phase of the Internet was relatively prosaic. Browsers connected to Web sites to view relatively static Web pages and HTML forms—much to their delight. More recently, those calm days of static connections and predictable network utilization have been replaced by a new breed of dynamic, bandwidth-hungry, and increasingly vulnerable Internet applications. What will this transition mean for corporate networks? This paper concludes:

- **New consumer-focused Internet applications impact business networks.** Web 2.0 sites like Facebook, Twitter, and YouTube continue to gain popularity inside and outside the corporate network. These sites create new risks for corporate IT assets as they open a new delivery channel for malicious code and push critical business traffic aside by consuming vast amounts of precious bandwidth.

- **Security managers can't stop the music.** Users now depend upon these new applications for business benefit, personal communication, and entertainment—it is too late for security managers to advocate blocking this traffic altogether.

- **Today's new Internet applications predict the future.** Consumer-focused Web 2.0 applications presage future business uses for resource sharing, file distribution, and collaboration. As such, large organizations that master new Internet application management today will have a head start on future network challenges.

- **Managing new Internet applications isn't easy.** Point tools can provide tactical help in small segments of the network, but can't integrate with existing firewalls, IDS/IPS devices, routers, and switches. Given these restrictions, their value is extremely limited.

- **Large organizations need an end-to-end Network Security Architecture (NSA).** To protect against threats and manage resource utilization, security must be integrated into the network in an enterprise architecture: the NSA. ESG believes that an NSA will provide the intelligence to help organizations control access to Internet applications, protect against malicious code attacks, and work with L2/L3 devices to prioritize business-centric traffic.

An Avalanche of Web 2.0 Applications

CIOs who examine the applications crossing over their networks may be overwhelmed by the results. They are certain to find multiple generations of applications spanning the mainframe, client/server, and Internet Computing eras. What may surprise them is the tremendous number of "other" applications traversing the wires in addition to Instant Messaging (IM), Peer-to-peer file (P2P) file sharing, and desktop telephony.

Actually, IT executives shouldn't be surprised by the presence of these kinds of applications on employee systems. Over the past few years, the use of these kinds of applications has skyrocketed. For example:

- According to company sources, Facebook now has a population of over 200 million active users, with 100 million users spending 3.5 million minutes on Facebook each day. The volume of online Facebook content is also staggering: more than 850 million photographs and 8 million videos are uploaded each month.

- Over 1 billion "tweets" have been posted on Twitter since its inception, with over 3 million new "tweets" posted per day.

- As of 2008, there were over 70 million videos posted on YouTube, with over 100 million videos viewed per day and over 13 hours of new videos uploaded every minute of every day. It would take of 400 years to view the entire library of YouTube videos.

- There over 130 million blogs in 81 different language indexed by the web site Technorati. Over 900 thousand new blogs are posted each day. Blogs are also widely consumed: 77% of active Internet users claim to be regular blog readers.

Of course, this is just the tip of iceberg. Application growth goes hand in hand with the increasing use of home broadband networks and mobile devices. Users are rapidly moving to an "always on" Internet experience that includes personalization, multi-channel communications, and ubiquitous Web access. In the near future, users will demand constant communications across all types of devices, regardless of whether they are working, traveling, or sitting in their living rooms.

Web 2.0 Traffic Impacts Business Processes and Network Utilization

Since many of these Web 2.0 sites are most often associated with consumers, many IT managers dismiss them as mere Internet toys. IT professionals claim to have bigger concerns than YouTube, Digg, and MySpace. Unfortunately, ESG believes that disregarding the ever-growing use of these applications would be a big mistake. Why? Employees often use corporate laptops as their primary computer, regularly downloading and using Web 2.0 applications—whether they are working remotely from Starbucks or sitting at their desks. This situation will soon be exacerbated by internal Web 2.0 development and production applications. In a recent survey of IT professionals around the world, nearly one-fourth expect that Web 2.0 applications will be in widespread production use within the next 24 months while another 25% anticipate limited production use (see Figure 1). Given the proliferation of external AND internal Internet applications, network utilization and security issues will only increase moving forward.

FIGURE 1. ENTERPRISES ARE EMBRACING WEB 2.0 FOR INTERNAL APPLICATIONS

To what extent does your organization use or plan to implement Web 2.0 applications within the next 24 months? (Percentage of users, N=602)

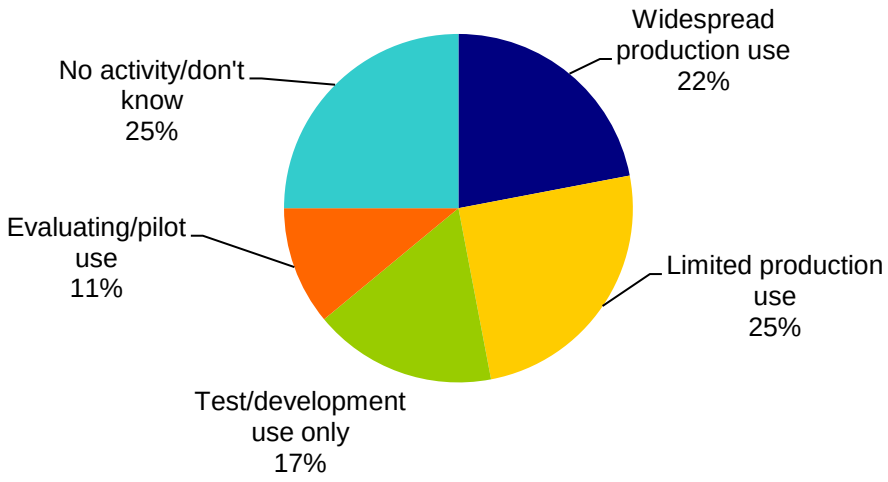

Source: Enterprise Strategy Group, 2009

Consumer applications like these are not an anomaly, but rather a sign of things to come. ESG believes that large organizations will likely implement a significant number of Web 2.0-based business applications over the next 3 to 5 years. Why? Web 2.0 technology will act as the underpinning for a multitude the next-generation collaborative applications (see Table 1). In this regard, today's "fad" is actually an extremely accurate model for future networking resource requirements.

TABLE 1. BUSINESS USES OF WEB 2.0 TECHNOLOGY

Application	Examples	Value
Internal communication	Blogs, wikis, social networking, video, IP telephony integration	Real-time personal contact opens lines of communication
Employee training	Video, wikis, Twitter, social networking	New tools can be constantly updated. Employees use experience to help others.
External communications	Video, IP telephony integration, Twitter, social networking	Various tools can be easily customized for business partners to enhance communication and productivity

As the saying goes, the next-generation Internet application train has already left the station—and there is no turning back. This situation leaves IT managers with two choices: 1) Get dragged into Web 2.0 applications kicking and screaming, or 2) Get to know the good and bad about Web 2.0 applications in order to use them for business benefit while minimizing operational problems. ESG strongly suggests that the latter course is far more prudent than the former.

New Applications, New Network Challenges

Clearly, new Internet applications are already in widespread use—and are here forever—so it is important to understand how these systems impact IT assets like the network. It turns out that the effect can be quite profound: Web 2.0 applications such as Facebook, Twitter, and YouTube can greatly influence network security and performance because of:

- **Bursty, unpredictable traffic patterns.** New Internet applications can be incredibly chatty, resulting in numerous network flows and bursty bandwidth utilization that can quickly grab 50% to 75% of total network capacity. Earlier this decade, schools like the University of Delaware, Salem State College (MA), and Stanford Universities were forced to more than double Internet bandwidth capacity in order to address students' appetite for P2P file sharing and music downloads. ESG finds that adding more and more bandwidth is the most common response to this growing network traffic.

- **Added vulnerabilities and new threat vectors.** Network managers face a security Hydra here. First, new Internet applications bridge the untrusted outside world and the internal network. They could be used by hackers to gain network access, attack IT assets, or steal confidential data. These applications can also introduce specific software vulnerabilities. This has happened in the recent past with popular applications such YouTube (SQL injection vulnerabilities), Facebook (cross-site scripting vulnerabilities), and MySpace (private data exposure vulnerability) Finally, Web 2.0 applications can open a back door for data leakage as innocent or malicious users send private data to outsiders in cleartext.

- **New malicious code attacks.** While there hasn't been a widespread global network attack caused by a Web 2.0 application, many experts believe it is only a matter of time. In April of 2009, a Twitter JavaScript vulnerability was exploited for worm propagation. In December 2008, social networking site Friendster members were inundated with e-mail messages containing links that downloaded Trojans to unfortunate victims' systems. Unfortunately, incidents like these are becoming increasingly common, placing valuable corporate assets and data at risk.

Managing New Internet Applications can be Difficult

Since new Internet applications add business value, blocking them would pit IT against business managers. Sagacious CIOs know that the business folks usually win these battles. From a technology perspective, cutting

off all Web 2.0 traffic would also prove to be a lost cause. Addressing this traffic can be difficult for several reasons:

- **Legitimate web sites can be used for malicious code distribution.** The unprecedented growth of Web 2.0 applications has come with some unfortunate baggage: web application vulnerabilities. According to several industry sources, over 60% of Web sites have security critical, high, or urgent security issues. Little wonder, then, that approximately 80% of Web sites hosting malicious code are actually legitimate. This situation creates a difficult network security challenge. Large organizations can no longer simply block individual websites; rather, they must inspect traffic on a URL by URL basis.

- **All or nothing solutions don't make sense.** Many solutions propose the old draconian measures of simply blocking any unwanted Internet traffic from proceeding beyond the gateway, but ESG believes these tools have business and technical limitations. The outright blocking of Facebook and MySpace may cast IT in the role of "big brother"—how many CIOs want this label? On the technical side, blocking all Internet applications means identifying every Internet ingress/egress point and then implementing security countermeasures at each one. Given the "disappearing perimeter" in today's networks, this step backward would prove extremely cumbersome to implement.

- **Tactical solutions don't provide integration for enterprise needs.** Networking technologies like WAN acceleration can be effective at identifying Internet application traffic and applying policies. This works well across WAN links, but fails to integrate with internal LAN switches, routers, and appliances. In this way, WAN acceleration systems act in isolation rather than across the enterprise.

Include Internet Application Management in the Network Security Architecture (NSA)

Today's Internet application security products have two fundamental problems. First, they adhere to a strict security "grant or deny" access metaphor. This may be appropriate for dealing with malicious code, but there may be a middle ground here where some employees are granted access and others are not. The other issue is that most security products act as isolated islands on a particular network segment. This, too, makes no sense. Managing Internet application access, policy, content, and traffic demands an end-to-end strategy across the enterprise, not a secluded network add-on.

If security point tools don't work, what can large enterprises do to address new Internet application security? ESG believes that this is the wrong question and part of the fundamental problem. Controlling the use of new Internet applications is NOT an issue of security alone; rather, it requires a multi-layered approach that encompasses enterprise network security, management, operations, and policies (see Table 2). When it comes to new Internet applications, large enterprises must be able to identify applications, block malicious code attacks, control access on a user-by-user basis, and manage bandwidth utilization across the enterprise.

TABLE 2. MANAGING NEW INTERNET APPLICATIONS

Networking Discipline	Requirement
Access Control	Restrict use of Internet applications to specific users who gain business benefit. Deny access to all others.
Security	Filter traffic for malicious code and data leakage.
Network Management	Monitor utilization of network traffic and devices.
Network Operations	Consolidate operational tasks and controls.
Reporting	Centralize alerts, reports, and audit information.
Policy Management	Enforce business policies across the network.

ESG believes that providing this type of granular control demands a more holistic strategy based on a model called the Network Security Architecture (NSA) that overlays security on top of network switching and routing. The NSA is made up of four integrated network layers:

- **Network Identity Context (NIC).** The NIC acts as an intelligent network edge that enforces access policies and tags Ethernet frames so the network can enforce rules or provide QoS by prioritizing specific types of traffic.
- **Ubiquitous packet filtering and firewalling.** Networking equipment must be able to differentiate between productive and malicious packets, regardless of location or protocol. To accomplish this, the network must become super intelligent and perform L4-7 tasks that improve network efficiency and security.
- **NSA visibility and analysis.** NSA security and operations management is based upon knowledge of network traffic behavior across the enterprise and on a device-by-device basis. This means identifying traffic on an application-by-application basis, monitoring flows, detecting anomalies, and reacting with alerts and automated corrective actions.
- **Policy-based rules enforcement.** Today's networks are a morass of standard and proprietary ACLs and queuing algorithms with no coordination across the network. The NSA provides consistent network traffic management based upon business and security requirements not device-centric technologies.

FIGURE 2. THE NSA ARCHITECTURE

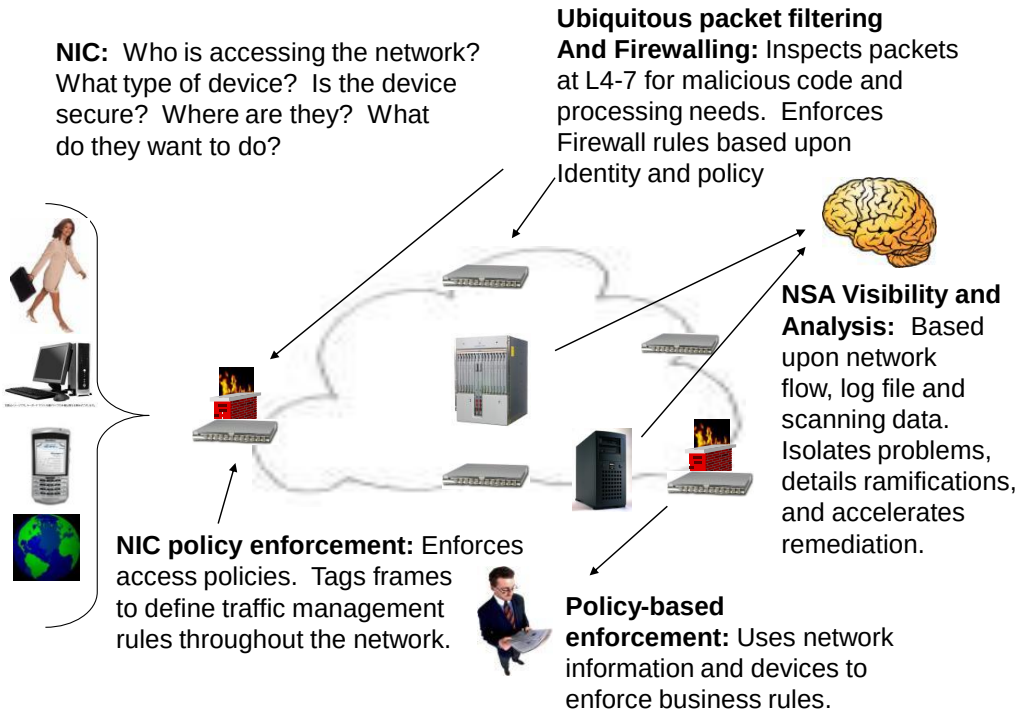

Source: Enterprise Strategy Group, 2009

The NSA and Internet Applications

The key to managing new Internet applications is providing secure access to the right employees while controlling bandwidth utilization. These are dynamic activities that can change on a moment's notice based upon characteristics like where the employee is physically located and what other traffic has bandwidth priority. The NSA can accommodate these kinds of requirements. When users log on to the network, the NSA immediately

determines who the user is and where they are located. With the user's profile and location as background, the NSA can either grant or deny access to applications, like Twitter. If access is granted, the NSA remains vigilant as it scans traffic for malicious code or security policy violations (i.e., data leakage) and manages bandwidth utilization based upon simultaneous traffic priorities and requirements.

Unlike today's amalgamation of one-off security point tools, the NSA depends upon a foundation of integration and the establishment of a two-way communications channel. Networking and security devices act as the "feet on the street" by providing a constant flow of data on their current state. Management platforms consolidate this information to get a picture of the overall network status. Finally, network policy management engines analyze the network status in the context of business requirements. To enforce these business requirements, network policy engines hand out instructions to individual devices and coordinate aggregate activities (see Figure 3).

FIGURE 3. NSA AND NEW INTERNET APPLICATION MANAGEMENT

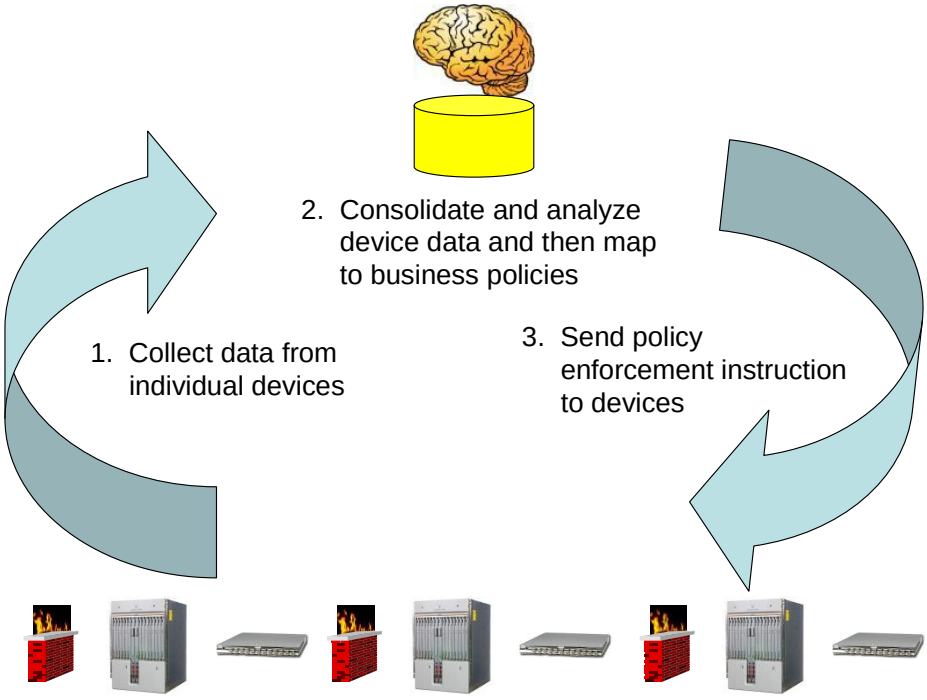

Source: Enterprise Strategy Group, 2009

Cynical networking executives may dismiss the NSA as science fiction, but the roots of the model are already taking shape. One example are the Adaptive Threat Management solutions by market leader Juniper Networks. Since its acquisition of NetScreen Technologies in 2004, Juniper has fastidiously introduced high performance devices while integrating networking and security functionality across its entire product line. In 2008, the company announced its Adaptive Threat Management solutions, which ties together its networking equipment along with its firewalls, intrusion prevention systems, dynamic services gateways (SRX), Universal Access Control, network and security management, and Security Threat Response Manager (STRM). Juniper's IDP Series works in harmony with its Unified Access Control infrastructure to enforce application and security policies based on user role information, allowing organizations to easily synchronize business needs with security needs. Each product maps to the components of NSA, making them compelling solutions on their own (see Table 3). To provide NSA-like end-to-end coverage, Juniper integrates the technologies together—making the value of the whole actually greater than the sum of its parts.

TABLE 3. MAPPING JUNIPER PRODUCTS TO THE NSA MODEL

NSA Component	Juniper Product
Network Identity Context (NIC)	Unified Access Control
Ubiquitous packet filtering and firewalling	SRX Series Services Gateways, ISG Series Integrated Security Gateways, IDP Series Intrusion Detection and Prevention Appliances
NSA visibility and analysis	Security Threat Response Manager
Policy-based rules enforcement	Network and Security Manager

The Networking "To-Do" List

ESG believes that all large organizations are bound to move to an NSA, but they are certain to have different priorities and timelines. Before buying any new networking technology, prudent networking executives must get a feel for current business requirements, Internet application usage, future plans, and technology options. ESG recommends that networking professionals take the following steps (see Figure 4):

FIGURE 4. THE TO-DO LIST

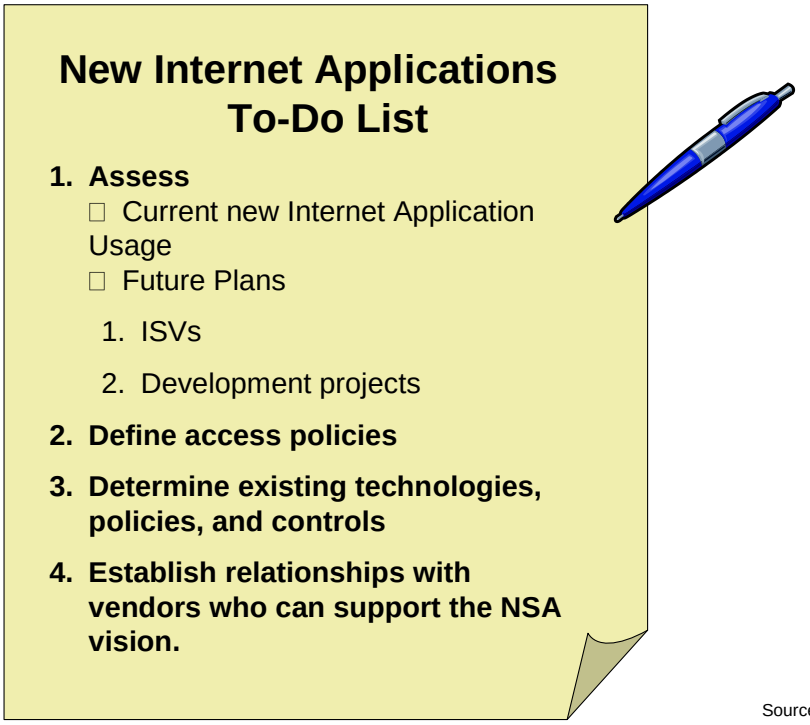

Source: ESG

1. **Assess new Internet application utilization.** It's important to start this process by understanding the scope of the problem. As previously mentioned, new Internet applications can be hard to identify, so don't depend upon network management tools for an answer. Rather, poll users to see if they are regularly accessing Web 2.0 applications. With this information in hand, see if network management tools can then provide some usage patterns for analysis. These two steps should give some indication of the scope of the issue. Finally, make sure to factor any upcoming Web 2.0-based business applications into the overall assessment and future plans.

2. **Caucus with managers on access policies.** Access policies around new Internet applications should be guided by a collective group composed of line-of-business, human resources, and IT management,

along with compliance managers and legal counsel. Networking and security managers' role in this process is to educate the rest of the team on network security and operational risks so they can base their policy decisions on the right factors.

3. **Determine existing tools, technologies, and processes.** It is likely that there are already defenses in the network based upon layers of tools and technologies like ACLs, firewalls, IDS/IPS devices, subnets, and VLANs. Take an inventory of these deployed defenses and note how and where they are redundant or in conflict with one another. Try to quantify the operational cost of these defenses in dollars and man hours as well. This map will point out where there are problems, holes, and operational inefficiencies and help the networking and security staff to build a good ROI model.

4. **Make NSA functionality a requirement in all new RFIs and RFPs.** The NSA can't be cobbled together based upon piece parts and custom integration. To avoid being trapped, base all future networking and security purchases on architectural as well as product considerations. Present your enterprise NSA vision to vendors and make sure they respond with details on how they can't integrate their devices into an integrated architecture.

The Bottom Line

Like death and taxes, new Internet applications are a fact of life—they can't simply be blocked or ignored. What's needed is a comprehensive plan for security and network management. Internet applications like Facebook, Twitter, and YouTube must be strictly controlled in terms of access, security, and bandwidth utilization. It is imperative that network managers develop a plan to accomplish this as soon as possible.

From a technology perspective, mapping business requirements to security and networking enforcement simply can't be done with today's tactical security solutions. Rather, it requires an end-to-end solution called the Network Security Architecture (NSA). Large organizations should assess their business needs and technology requirements and begin building an NSA with vendors like Juniper Networks as soon as possible.

20 Asylum Street
Milford, MA 01757
Tel: 508-482-0188
Fax: 508-482-0218

www.enterprisestrategygroup.com