

Dynamic, Flexible Security Architecture

By *Andreas M. Antonopoulos*
SVP and Founding Partner, Nemertes Research

Executive Summary

The Rube Goldberg approach to security is not just unmanageable, it is insecure: even a cherished concept like defense-in-depth needs to be rethought in light of mobility and agility. There is a significant gap between the speed with which companies are moving, and the speed of IT security. The only way to close the gap is to abandon the old ways and implement a new, more agile and more secure architecture. This new security architecture is identity-centric rather than location-centric, with centralized policy and distributed enforcement, across all types of networks, wired, wireless, mobile wireless, remote access, and VPN.

The Issue: Old Security Models vs. New Enterprise IT Models

While enterprise IT has gradually become more dynamic, mobile, open, outward-facing and collaborative, IT security has languished with a static, location-based and inflexible architecture. Location-centric security hampers companies, leaving them no option but to create inflexible rules that constrain innovation and mobility. In a location-centric security infrastructure where the level of security depends on layered perimeters, users' access depends on *where* they are, rather than *who* they are. As a result, when users go "mobile," they are penalized by the infrastructure because they can no longer access the same applications.

Furthermore, rigid security controls slow down the rate of innovation across the entire company. IT teams cannot easily deploy new applications, or introduce new types of devices, or extend access to new partners, suppliers and customers. Security becomes the gating factor for innovation: the slowest cog in the system that slows everything else down. Worse yet, the slowdown in innovation is not a necessary compromise that results in better security—location-centric and static security is less responsive to change, more complex and more susceptible to the introduction of human error. So in practice, companies trade off innovation and mobility for the illusion of security, not actual security.

Companies need to transform their security infrastructure, to remove location-specific controls and introduce an identity-centric architecture that offers dynamic, flexible, mobile, policy-based security.

Challenges in the dynamic mobile enterprise

Business has changed dramatically over the last few years. Companies create a significant amount of value in collaboration with partners and customers. Collaboration across ad-hoc teams spread around the world is now the norm. Employees and customers want to access data anytime, from anywhere, using whichever device form-factor is most appropriate and convenient. Consumer devices such as tablets and smartphones are streaming in the company doors, introduced by employees who juggle consumer gadgets, corporate IT applications and leisure activities seemingly at the same time.

All of this can easily be dismissed as a fad. Yet, progressive IT organizations have used these trends to expand working hours, increase productivity and make their businesses ultra-competitive. These companies use consumer devices, social media, consumer applications, and mashups to build new, agile business practices and to empower employees to communicate effectively and efficiently both internally and with customers and partners.

The traditional business software landscape is also transforming. For many companies, corporate IT used to be about homogeneity: a single OS, single desktop type, fixed “approved” applications and uniform business processes. IT may still *officially* support only Microsoft Windows, but look around the employee desks (especially executives’ desks) and another picture emerges: Dozens of devices, “semi-supported,” offering a rainbow of desktop operating systems (Windows 7, Mac OS X, RedHat, Ubuntu and other flavors of Linux), tablet operating systems (iOS, Android, WebOS), smartphone operating systems (Blackberry OS, iOS, Android, Symbian, Windows Mobile, WebOS). That’s only counting the operating systems connected to the corporate network. Beyond that, many applications are migrating into the cloud, with Software-as-a-Service (SaaS) solutions increasingly expanding beyond CRM and ERP to include other corporate applications, both simple (expense reporting) and complex (social media and collaboration).

An impossible choice

Today’s enterprise security teams are faced with an impossible choice: they must choose between a vibrant, mobile, dynamic business or a business that is “secure.”

“Traditional” security can be the antithesis of dynamic, flexible, and mobile. If security teams mostly use layer 3 ACLs, VLANs, perimeter firewalls, and other location-centric controls those controls will become a huge impediment to the business. Every time a business unit requests a new application, the security team must adjust the security architecture to incorporate it. New end-user devices cannot be introduced, because they cannot be secured. Mobile users must compromise, accessing only a subset of applications over complicated remote access solutions. For the security team, this “impossible choice” results in impossible constraints – if all they have is antiquated security architecture, the only answer they can give is “NO.” Every new development hits a barrier as soon as security is brought into the conversation.

Business units are demanding more flexible IT and are frequently opening their wallets to get their way: buying cloud infrastructure and SaaS applications without too much IT involvement. A worst nightmare for security, this trend is simply a result of IT security's failure to keep up with business demands.

At the end of the day, the role of information security in a business is to “enable” the business to take *reasonable* risks without losing control. But when security cannot adapt to change, it becomes a ball and chain, slowing down change and gradually making the business less and less competitive. Businesses that are able to resolve this “impossible choice” can turn security from an impediment to an enabler of change. A flexible and dynamic security infrastructure allows security officers to say “YES” to mobility, to tablets, to mashups, to collaboration and to rapid application deployment, while even reducing risk.

A new dynamic, mobile, flexible architecture

When most business was “local”, most security was location-centric. Location-centric security is security that depends primarily or exclusively on the location of a network device to make security enforcement decisions. So, if you have a firewall that defines a “perimeter,” security decisions are made by the firewall based on whether a device is “inside” or “outside” the perimeter. Essentially, the decisions are based on the assigned IP address and network subnet of the device. That works well when users have a single device that doesn't move much (a desktop), in a fixed infrastructure. Over time, security teams have expanded this model to layer successively more “trusted” perimeters deeper and deeper into the network. Thus, you have 3-tier application architectures, where the “core” of the application is within the innermost perimeter and layers of security form successive concentric circles around the core.

Layered security and defense-in-depth are not broken per-se. Rather, it is *static* layers and *static* perimeters that are broken. The disadvantages of this model become evident when you try to move a device, either a server (eg. VMotion or live-migration), or a client (a mobile device). When mobility is introduced, the architecture of concentric layers of security instead becomes an impediment to change. To solve that, security teams often resort to network gimmicks, like punching holes through firewalls, using VLANs and VPNs to link security “zones” and constantly changing firewall rules to adapt to new applications. The reason security teams resort to “kludgy” solutions is because mapping a dynamic set of users and applications to a static and location-based security architecture is impossible. Inevitably the system becomes too complex to manage.

Today's dynamic, mobile and flexible businesses require equally dynamic, flexible and mobile security. The security model that achieves this is an identity-centric model that is location-independent, in favor of user and application identification. In such a model, security controls are applied based primarily on *whom* you are not *where* you are in the network. Location can be taken into consideration, of course, but it is only one of the attributes considered by the security policy. The most important attribute is the user identity, which becomes

the key policy component for controlling access to applications and the infrastructure. User identity is also critical to logging, monitoring and audit functions that can be used to track who access what and when, for regulatory compliance or policy reasons.

An identity centric infrastructure consists of 3 major architectural components: endpoints, policy enforcement points, and policy decision points. The core of the architecture is a policy engine (the policy decision point) that allows central management of the user-centric policies for the business. The policy engine is where policy decisions are made by comparing attributes (such as user identity, device type, location, time of day, application etc.) to dynamic access policies.

As an example, a policy may state that a specific group of users (e.g. nurses) can access a specific application (e.g. patient records) from specific stations, at specific times, over specific networks. By centralizing policy management, the policy engine allows companies to apply a consistent policy across their entire network and even expanding to encompass external (e.g. SaaS) applications.

The other two components of a dynamic, mobile, and flexible security architecture are the endpoints and the enforcement points. Endpoints are the end user devices, but also applications, databases, servers, and other devices connected to the network. Enforcement points are various security devices or network devices (routers, switches, etc.) that can enforce access control decisions. While enforcement is distributed across the network, policy is centralized in a policy engine. Centralization of policy means security teams can modify policies in one place and rapidly introduce new applications, new logical zones and trust zones, new devices and new business processes. Since enforcement is distributed, devices can move from network segment to network segment, jumping from wired to wireless to mobile wireless networks while maintaining consistent access policies and application experience.

Example of identity-centric security and the implications for business

By implementing identity-centric security based on central policy management, companies can ensure that their security architecture is as flexible as their work practices. Policies remain attached to users so, as users move around, switching from wired to wireless networks, from desktops to tablets and from application to application, their security policies “follow them.” Let’s consider, for example, an identity-based policy for a healthcare organization. The policy defines “who” can access “what”, “when” and from “where”. The attributes are defined dynamically and in a flexible manner that is independent of the network architecture, endpoint device, or geography.

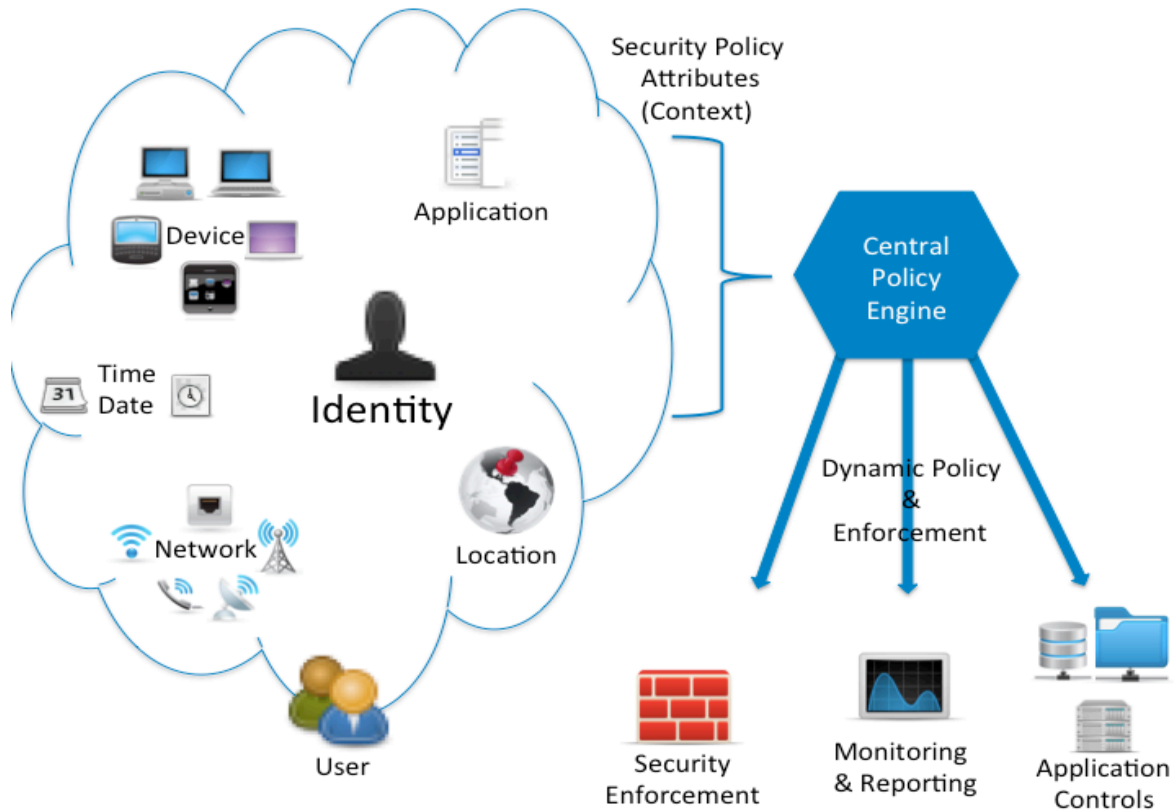


Figure 1: Dynamic, Flexible Security Architecture

Who – a user identity based on role or group. Unlike a traditional security model, in an identity-centric model the user identity follows the user across every device, every network and every application. All network flows to and from that user are clearly linked to the user identity and can be controlled by central policy. A nurse for example, can be identified and have the appropriate policy applied, regardless of whether he or she is using a kiosk/terminal, a wired desktop or a wireless tablet. The user identity is also critical for logging, monitoring and auditing user access, regardless of device or location.

What – an application profile, defined not in terms of TCP ports or IP addresses, but to the identity of the application, and dynamically mapped to specific servers, ports and IP addresses on the fly. So instead of allowing access to 10.1.2.3, the policy allows access to “Customer Relationship Management”, whether that application is running on one server or 10, whether it is dynamically load-balanced or geographically routed, or whether it is failed-over to a secondary data center or is delivered as a SaaS application using SAML (Security Assertion Markup Language) authentication. Applications can be added to this policy without having to make changes to every firewall and access control device throughout the infrastructure

When – policies in an identity-centric security model can also control for attributes such as time-of-day. Therefore, it is possible to specify that a nurse can only access the patient medication dispensing system between 1 and 3 PM.

Where – a flexible identity-centric policy engine can still use location or network medium (wired vs. wireless) as one of the policy attributes to control access. But unlike a location-centric security policy, the location is only one of the attributes considered and is not the defining attribute. This allows the flexibility to control access to data by location (for example, to restrict patient records from crossing a national border for HIPAA related reasons) without making location a needless constraint (you can only access “pharmacy” apps when in the eastern campus LAN behind the firewall).

Moving from location-centric to identity-centric security

If you want to move from location-centric to identity-centric security, you must change the way you evaluate and buy security solutions. To ensure you are buying a solution that incorporates identity-centric principles. Here are some key questions:

- What is the primary security attribute in access control policies (is it IP address, or user?)
- Does the security solution allow central control of policies for all networks (wired, wireless LAN, mobile wireless, remote/VPN)
- Can a user move from network to network, under the same policy (or do you need to define policy in multiple places)
- Can a server move without changing security policy
- Can a new application be added without making changes in enforcement points (e.g. firewalls)
- Can a user’s access be monitored end-to-end on multiple devices regardless of IP address or network location?

Conclusions and Recommendations

Static security architectures are getting in the way of your highly dynamic, mobile and interconnected business. Security may act as a barrier to your business’ innovation, technology adoption and competitiveness. More security spending won’t solve the fundamental mismatch between security that is location-centric and business that strives to operate anywhere, anytime, on any device. You must start planning a transition to identity-centric, context aware, and dynamic security based on an architecture that centralizes policy and distributes enforcement. Only then can your users be liberated to be productive and collaborate without barriers, increasing your company’s rate of change, innovation and competition.

About Nemertes Research: Nemertes Research is a research-advisory and strategic-consulting firm that specializes in analyzing and quantifying the business value of emerging technologies. You can learn more about Nemertes Research at our Website, www.nemertes.com, or contact us directly at research@nemertes.com.