



Scaling Network Security

Version 1.3

Released: July 26, 2018

Author's Note

The content in this report was developed independently of any sponsors. It is based on material originally posted on [the Securosis blog](#), but has been enhanced, reviewed, and professionally edited.

Special thanks to Chris Pepper for editing and content support.

This report is licensed by Gigamon.



www.gigamon.com

Gigamon is the company leading the convergence of network and security operations to help organizations reduce complexity and increase efficiency of their security stack. The Company's GigaSECURE® Security Delivery Platform is a next generation network packet broker that helps customers make threats more visible across cloud, hybrid and on-prem environments, deploy resources faster and maximize the performance of their security tools. Global 2000 companies and government agencies rely on Gigamon solutions to help stop tool sprawl and save costs. Learn how you can make your infrastructure more resilient, more agile and more secure at www.gigamon.com, on our [blog](#) and [Twitter](#), [LinkedIn](#) and [Facebook](#).

Copyright

This report is licensed under Creative Commons Attribution-Noncommercial-No Derivative Works 3.0.

<http://creativecommons.org/licenses/by-nc-nd/3.0/us/>



Scaling Network Security

Table of Contents

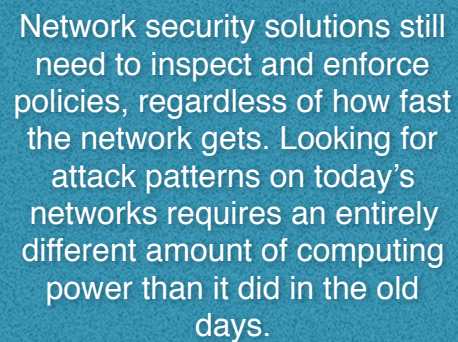
RIP, the Moat	4
The New Network Security Requirements	8
The Scaled Network Security Architecture	11
Summary	15
About the Analyst	16
About Securosis	17

RIP, the Moat

Young people today laugh at folks with a couple decades of experience when they rue about the *good old days*, when networks snaked along the floors of offices (shout out for [Thicknet!](#)), and trusted users were on the corporate network, while untrusted users were not.

Suffice it to say the past 25 years have seen some rapid changes to technology infrastructure. First of all, in a lot of cases, there aren't even any wires. That's kind of a shocking concept to a former network admin who fixed a majority of problems by swapping out patch cords. On the plus side, with the advent of wireless and widespread network access, you can troubleshoot your network from the other side of the world.

We've also seen continuing insatiable demand for network bandwidth. Networks grow to address that demand every year, which stresses our ability to protect them. But network security solutions still need to inspect and enforce policies, regardless of how fast the network gets. Looking for attack patterns on today's networks requires an entirely different amount of computing power than it did in the old days. So an essential requirement is to ensure that your network security controls can keep pace with network bandwidth, which may be Mission: Impossible. Something has to give at some point to keep the network secure.



Network security solutions still need to inspect and enforce policies, regardless of how fast the network gets. Looking for attack patterns on today's networks requires an entirely different amount of computing power than it did in the old days.

In this "Scaling Network Security" paper, we will look at where secure networking started and why it needs to change. We'll present requirements for today's networks which will take you into the future. Finally we will wrap up with some architectural constructs we believe can help scale up your network security controls.

The Moat

Let's take a quick tour through the past 20 years of network security. We appreciate the digression — we old network security folks get a bit nostalgic thinking about how far we've come. Back in the day, a modern network security industry started with a firewall to provide access control. Then a seemingly never-ending set of additional capabilities were introduced in the decades since.

Next was network Intrusion Detection Systems (IDS), which looked for attacks on the network. Rather than die IDS morphed into IPS (Intrusion Prevention Systems) by adding the ability to block attacks based on policy. We also saw a wave of application-oriented capabilities in the form of

Application Delivery Controllers (ADC) and Web Application Firewalls (WAF), which applied policies to scale applications and block application attacks.

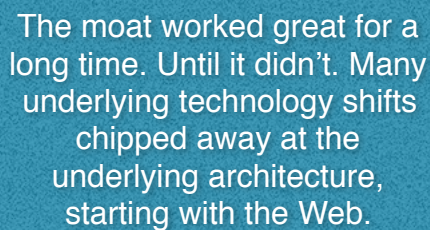
What did all these capabilities have in common? They were all based on the expectation that attackers were **out there**. Facing an external adversary, you could dig a moat between them and your critical data to protect it. Default Deny best illustrated this concept, a central concept in secure networking for many years. It held that if something wasn't expressly authorized it should be denied. So if you didn't set up access to an application or system it was blocked. That enabled us to dramatically reduce attack surface, by restricting access to only those devices and services which should be accessed.

Is Dead...

The moat worked great for a long time. Until it didn't. Many underlying technology shifts chipped away at the underlying architecture, starting with the Web. Yeah, that was a big one.

The first was encapsulation of application traffic into web protocols (predominately ports 80 and 443) as the browser became the interface of choice for pretty much everything. Firewalls were built to enforce access controls by port and protocol, so this was problematic. Everything looked like web

traffic, which you couldn't block, so the usefulness of traditional firewalls was dramatically impaired, putting much more weight on deeper inspection using IPS devices.



The moat worked great for a long time. Until it didn't. Many underlying technology shifts chipped away at the underlying architecture, starting with the Web.

But the secure network would not go quietly into the long night, so a new technology emerged a decade ago, which was unfortunately called the *Next Generation Firewall* (NGFW). It provides far more capabilities than an old access control firewall, including the ability to peek into application sessions, profile them, and both detect threats and enforce policies at the application level.

These devices were really more *Network Security Gateways* than firewalls, but we don't usually get to come up with category names, so NGFW stuck.

The advent of NGFW was a boon to customers who were very comfortable with moat-based architectures. So they spent the last decade upgrading to the NGM: Next Generation Moat.

Scaling Is a Challenge

As described above, networks have continued to scale, which has increased the compute power required to implement an NGM. Yes, network processors have gotten faster, but not as fast as packet processors. Then you have the issue of the weakest link. If you have network security controls which cannot keep pace you run the risk of dropping packets, missing attacks, or more likely both. To address this you would need to upgrade all your network-based security controls at

the same time as your network to ensure protection at peak usage. That seriously complicates upgrades. So your choice is between:


1. **\$\$\$ and Complexity:** Spend more money (multi-GB network security gateways aren't cheap) and complicate the upgrade project to keep networks and network security controls in lockstep.
2. **Oversubscribe Security Controls:** You can always take the bet that even with an upgraded network, bandwidth consumption takes some time to scale up beyond what network security controls can handle.

Of course you don't want all your eggs in one basket, or more accurately all your controls focused on one area of the environment. That's why you implemented compensating controls within application stacks and on endpoint devices. But you still need to figure out an approach to ensure network security can scale.

The Cloud Is the Final Nail

Then the cloud happened. It's the final nail in the coffin of the Moat architecture. Protecting critical assets was difficult, but at least you used to know where the data was. We'll again point to our 1870s Man meme, which reminds us of the olden days, when data was in the data center. It's no longer even clear what a *data center* is.

As organizations started adopting Software as a Service (SaaS) for things like customer relationship management, service desks, and even HR and accounting functions, network traffic dynamics began changing, compounded by the widespread adoption of collaboration SaaS including Office 365 and G Suite. Employees started hitting web services for critical business functions, not necessarily needing to be on the corporate network at all. Some organizations did force employees to route their traffic through the VPN (and therefore the corporate network) to ensure inspection and policy enforcement, but that requires you to have sufficient ingress and egress bandwidth and capacity on the inspection points for all this traffic. It's just not the best way to protect those employees and data.



Further exacerbating this change in traffic dynamics is the adoption of Infrastructure as a Service (IaaS) offerings, to initially supplement and then likely replace corporate data centers. You get more control over how traffic is routed through IaaS and which security controls are in place, but the native cloud providers offer strong network security capabilities which don't require bottlenecks and inspection points, and many cloud-native applications leverage these embedded security capabilities.

Remember that the Moat requires inline network security controls, able to inspect traffic and enforce access control policies and detect threats. And once again the secure network isn't going quietly, so many network security vendors now offer *virtual* network security devices for deployment in the cloud. We consider virtual firewalls a stopgap, not a long-term solution. Broad network access and elasticity are fundamental tenets of cloud architecture. Forcing network traffic through a series of inspection points is suboptimal.

We consider virtual firewalls a stopgap, not a long-term solution. Broad network access and elasticity are fundamental tenets of cloud architecture. Forcing network traffic through a series of inspection points is suboptimal.

But we don't see the moat going away tomorrow, because on-premise networks are not going away any time soon. As long as organizations have data in their data centers and devices connecting to it from outside the organization, traditional network security controls (firewall, IPS, web filter, etc.) are required, but *are not sufficient* for tomorrow's network security requirements (or even today's).

The New Network Security Requirements

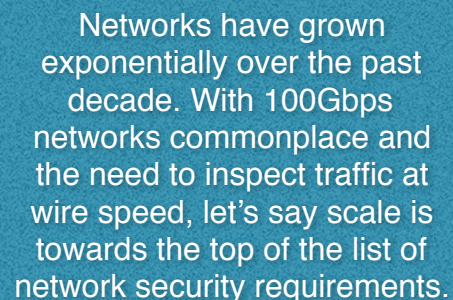
Depending exclusively on perimeter-centric security (The Moat) is no longer a viable network architecture, given the encapsulation of almost everything into standard web protocols and the movement of critical data to an expanding set of cloud services. Additionally, the insatiable demand for bandwidth further complicates how network security scales. So it is time to reassess the requirements of the new network security. As we rethink network security, what do we need it to do?

Scale

Networks have grown exponentially over the past decade. With 100Gbps networks commonplace and the need to inspect traffic at wire speed, let's say scale is towards the top of the list of network security requirements. Of course, as more and more corporate systems move from data centers to cloud services, traffic dynamics change fundamentally. But pretty much every enterprise we run into still has high-speed networks, which need protection. So you can't punt on scaling up your network security capabilities.

How has network security scaled so far? Two techniques:

1. **Bigger Boxes:** The old standby is to throw more iron at the problem. At some point, the security controls just aren't going to get there — whether in performance or cost feasibility, or both. There is undoubtedly a time and a place for bigger and faster equipment; we don't dispute that. But your network security strategy cannot depend on the never-ending availability of bigger boxes to scale.
2. **Drop (and don't inspect) Packets:** The other option is to drop packets all over the (proverbial) floor since your security controls cannot keep up, thus you end up not inspecting the traffic (missing attacks) or slowing down the application, both bad options. At least if you ask your senior management.



Networks have grown exponentially over the past decade. With 100Gbps networks commonplace and the need to inspect traffic at wire speed, let's say scale is towards the top of the list of network security requirements.

The need for speed isn't just pegged to increasing network speeds — it's also dependent on the types of attacks you'll see and the amount of traffic preprocessing required. For example with today's complicated attacks, you may need to perform multiple kinds of analyses to detect an attack, which requires more compute power. Additionally, with the increasing amount of encrypted traffic on networks, you need to decrypt packets before inspection, which consumes tremendous resources.

Even if you are looking at a network security appliance rated for 80Gbps throughput for threat detection, you need to understand the kind of inspection the box does, and whether it would detect the attacks that concern you.

We don't like to compromise on either spending a crap ton of money to buy the biggest security box you can find (which still might not be big enough) or deciding not to inspect some portion of the traffic. The scaling requirements for the new network security are:

1. **No Security Compromises:** You need the ability to inspect traffic which may be part of an attack. Period. To be clear that doesn't mean all traffic on the network, but you need to be able to enforce security controls where necessary.
2. **Capacity Awareness:** I think I saw a bumper sticker once which said: "TRAFFIC HAPPENS." And it does. So you need to support a peak usage scenario without having to pre-provision for 100% usage. That's what's so attractive about the cloud. You can scale up and contract your environment as needed. It's not easy on your networks, but that's the capability we need. Understand that security controls are capacity constrained, and make sure those devices are not overwhelmed with traffic and don't start dropping packets.

So what happens when you upgrade network speeds, which does occur from time to time? You want to improve your security controls on *your* timetable. You can't compromise on security just because network speeds increased. And a network upgrade represents a legitimate burst, potentially at all times. If you can satisfy those two requirements, you'll be able to gracefully handle network upgrades without impairing your security posture.

Intelligent and Flexible

The key to not compromising on security is to apply the required controls intelligently. For example not all traffic needs to run through the network-based sandbox or the DLP system. Some network sessions between two trusted tiers in your application architecture require access control. In fact, you might not need security inspection at all on some connections. In all cases **you** should be deciding where security makes sense, not being forced by the capabilities of your equipment.

You require the ability to enforce a security policy and implement security controls where needed.

1. **Classification:** Figuring out which controls to apply to the network session depends first on understanding the nature of the connection. Is it associated with a specific application? Is the destination a segment or server you know holds sensitive data?

2. **Policy-based:** Once you know the nature of the traffic, you need the ability to apply an appropriate security policy. That means some controls are in play and others aren't. For example, if it's an encrypted traffic stream you'll need to decrypt it first, so off to the SSL decryption gear. Or as we described above, if it's traffic between trusted segments, you can likely skip running it through a network sandbox.
3. **Multiple Use Cases:** Security controls are used both in the DMZ/perimeter and within the data center, so your new network security environment should reflect those differences. There is likely more inspection required for inbound traffic from the Internet than for traffic from a direct connection to your public cloud. Both are external networks, but they generally need different security policies.
4. **Cloud Awareness:** You can't forget about the cloud, even though network security can differ significantly from your corporate networks. So whatever kinds of policies you implement on-premise, you'll want an analogy in the cloud. Again, the controls may be different, as will deployment, but the level of protection must be consistent regardless of where your data resides.



The new network security architecture is about intelligently applying security controls at scale, with a clear understanding that applications, attackers, and technology infrastructure continually evolve.

The new network security architecture is about intelligently applying security controls at scale, with a clear understanding that applications, attackers, and technology infrastructure continually evolve. Your networks will look significantly different in 3 years, but you don't want your level of protection to differ, nor do you want your security environment to need forklift upgrades every 18 months.

The Scaled Network Security Architecture

These new scaled network security requirements are all well and good, *but how do you get there?* Let's dig into a couple key architectural constructs which will influence how you build your future network security architecture.

But before we go into specifics let's state a few caveats around the architecture. Not everything works for every organization. There may be cultural impediments to some of the ideas we

Not everything works for every organization. There may be cultural impediments to some of the ideas we recommend. You will need to decide which ideas are suitable for your current problems, and which battles are not worth fighting.

recommend. We point this out because any new way of doing things can face resistance from folks impacted by the change. You will need to decide which ideas are suitable for your current problems, and which battles are not worth fighting.

There may also be technical challenges, especially with extensive networks. Not so much conceptually — faster networks and increased flexibility are already common, regardless of the size of your network. The challenge comes in phasing migration. But nothing we will recommend requires a flash cutover, nor are any of these ideas incompatible with existing network security

constructs. We have always advocated *customer-controlled* migration, which entails *deciding* when you will embrace new capabilities — not accepting an arbitrary requirement from a vendor or any other influencer.

Access Control Everywhere

Our first construct is *access control everywhere*, which is pretty fundamental because network security is about controlling access to critical resources. Duh. We have been pointing out that [segmentation is your friend](#) for years. But in traditional networks it became tough to do true access control scalably because data flows weren't predictable, workloads and data move around, and users need to connect from wherever they are.

The advent of software-defined everything (including networks) has given us an opportunity to manage who gets access to what, and when, more effectively. The key is setting policy. Yes, you start with critical data, and who can & should access it from where, to set your baseline. But the larger the network and the more dispersed employees and resources (including mobility and the

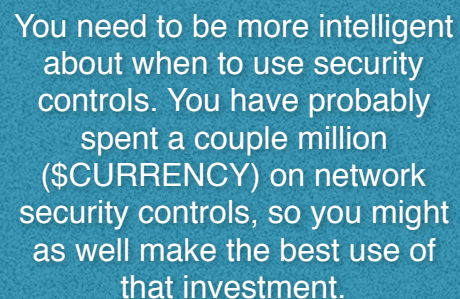
cloud) are, the tougher it is. So you do the best you can with the initial set of policies and then hit it again from the other side. Your new network security should be able to monitor traffic flows and suggest a workable access control policy. You'll need to scrutinize and tune the policy while comparing it with the initial cut, but this will accelerate your efforts.

Returning to the need for flexibility, you should be able to adapt policies as needed. Sometimes even on the fly, within parameters defined by policy. That doesn't mean you need to embrace machines making policy changes without human oversight or intervention, at least at first. In a customer-controlled migration *you* determine the pace of automation, enabling you to get comfortable with policies and ensure maximum uptime and security.

Applying Security Controls

With segmentation reducing the attack surface by preventing unauthorized access to critical resources, you still need to ensure authorized connections and sessions are not doing anything malicious. But devices still get compromised, so we can't forget the prevention and detection tactics we've been using on our networks for decades. As described under Requirements, you need to be more intelligent about when to use security controls. You have probably spent a couple million (\$CURRENCY) on network security controls, so you might as well make the best use of that investment.

Once again we return to the importance of policy-based network security. Depending on the source, destination, application, time of day, geography, and about a zillion other attributes (okay, we may be exaggerating a bit), we want to leverage a set of controls to protect data. Not every security control applies to every session, so the network security platform needs to implement controls selectively.



You need to be more intelligent about when to use security controls. You have probably spent a couple million (\$CURRENCY) on network security controls, so you might as well make the best use of that investment.

Decryption

Before you start worrying about which controls to apply, you need to make sure you can inspect sessions. With more and more network traffic encrypted nowadays, before you can implement security controls you will likely need to decrypt. We wrote about this at length in [Security and Privacy on the Encrypted Network](#), but things have changed a bit over the past few years.

The standard approach to network decryption involves intercepting the connection to the destination (called person-in-the-middle) and then decrypting the session using a master key. The decryption device then routes the decrypted stream to the appropriate security control per policy and then sets up a separate encrypted connection to the destination server. And yes, our political correctness may be getting the best of us, but we're pretty sure that network security equipment is not gender-binary, so we like 'person' in the middle.

Any network security platform will need to provide decryption capabilities as required. But that's getting more complicated, as described in the [TLS 1.3 Controversy](#). A person in the middle weakens the overall security of a connection because any organization (some good — like your internal security team; and some bad — like adversaries) could theoretically get in the middle to sniff the session. The TLS 1.3 specification addresses that weakness by implementing Perfect Forward Security, which uses a different key for each session to prevent a single master key which could monitor everything.

Not being able to get in the middle of network sessions eliminates your ability to inspect traffic and enforce security policies *on the network*. To be clear, it will take a long time for TLS 1.3 to become pervasive; in the meantime your connections can negotiate down to TLS 1.2, which still allows person-in-the-middle. But we need to start thinking about different, likely endpoint-centric, approaches to inspecting traffic before it hits the encrypted network.

Contextual Protection

Assuming we can inspect traffic on the network, we want to implement a policy-centric security approach. That means identifying the traffic and determining which security control(s) are appropriate based on the specifics of the connection. Context helps ensure you are using the proper security controls, which improves security posture and helps to optimize control capacity (as we'll discuss below).

The best way to understand this is with a couple simple examples:

- **Ingress:** In case of an inbound connection you want to protect against malware coming into the network, as well as application attacks. So you can set a policy that routes email traffic through an email security gateway and then a network-based malware scanner. Or maybe you take email from an email security service and then run it through your malware scanner or IPS to ensure any links in the message aren't malicious. To protect application traffic first the connection goes through a WAF, but you can also run it through an IPS to detect more traditional attacks. Similarly, you'd like to be able to leverage different controls if the session originates in a hostile country which demands more scrutiny.
- **Egress:** Looking at it from the other end, if you are dealing with outbound traffic you first want to decrypt an encrypted session and then send it through a web filter to determine potential misuse, connecting to a malicious site, or showing patterns which may indicate command and control traffic. But depending on what kind of data is in the payload, you might also want that traffic to run through a DLP device to look for data misuse. You'll want to provide context for DLP inspection, both to improve accuracy and because DLP is very resource intensive.

These examples are deliberately oversimplified, but contextual protection enables you to use the controls you need to protect a *specific connection*.

Optimizing Capacity

You don't always have the luxury of upgrading network security controls at the same time as network bandwidth. Additionally, heavy-duty Deep Packet Inspection, as described in the examples of contextual protection above, may not be needed for all traffic — especially given the significant resources it requires. So when determining which controls to use on which connections, it's important to factor capacity into the mix.

You don't want to compromise security due to capacity. But a network security platform which can give you a sense of when specific security controls are at capacity, as well as potentially even the flow of traffic to prevent dropped packets (like a load balancer for inbound web apps), provides a graceful way to manage network capacity.

A network security platform which can give you a sense of when specific security controls are at capacity, as well as potentially even the flow of traffic to prevent dropped packets, provides a graceful way to manage network capacity.

For another example, if you recently upgraded your data center network to 100GB but don't have the security budget to increase the speed of your internal segmentation firewalls, you can buffer traffic with a network security platform while the firewalls enforce the policy. This is not a great answer because it impacts application traffic, but the alternative might be to either violate segmentation rules (which probably won't sit well with the auditors) or drop packets.

Another example is intelligently routing connections to authorized SaaS applications, but through your secure web gateway service rather than your internal DLP engine, because you already have a CASB monitoring activity in those SaaS applications. That can help your DLP device scale more effectively. Again, simple examples illustrate how intelligently selecting security controls per connection is useful.

Getting There

The key here is not to get wrapped up trying to boil the ocean. You can start small, perhaps implementing an SDN in front of your egress security controls to apply the policies we discussed. Or possibly introducing a packet broker in front of a critical application to make sure appropriate security controls are not overwhelmed in case of a traffic flood. You could start thinking about micro-segmentation in your virtualized data center, and map those capabilities to all new applications running in IaaS. Or you might be interested in a Zero Trust access control environment or a Secure Network as a Service offering for employee access, and roll out intelligent networks internally to provide access to some resources (which remote employees need) while segmenting everything else.

The possibilities for how to migrate to this kind of network security platform are endless, and there is no single right or wrong answer. There is only the reality that your security controls cannot scale at the same rate as your networks, which means you need to apply intelligence to deployment of security controls within your environment.

Summary

Existing network security architectures, based mostly on preventing attacks from external adversaries, don't reflect the changing dynamics of enterprise networks. With business partners and other trusted parties needing more access to corporate data and the encapsulation of most application traffic in standard protocols (Port 80 and 443), digging a moat around your corporate network no longer provides the protection your organization needs. Additionally, network speeds continue to increase putting a strain on inline network security controls that much scale at the same rate as the networks.

Successfully protecting networks requires you to scale network security controls while being able to enforce security policies flexibly. By applying context to the security controls used for each connection ensures proper protection without adding undue stress to the controls. The last thing you can do is compromise security in the face of increasing bandwidth.

Thus, the scaled network architecture involves applying access control everywhere to make sure only authorized connections have access to critical data and implementing security controls where needed, based on the requirements of the application. Moreover, security policies need to change as networks, applications and business requirements change, so the architecture needs to adapt without requiring forklift upgrades and radical overhauls.

In short, the time is now to start planning what your secure network strategy will be over the next ten years.

If you have any questions on this topic, or want to discuss your situation specifically, feel free to send us a note at info@securosis.com.

About the Analyst

Mike Rothman, Analyst and President

Mike's bold perspectives and irreverent style are invaluable as companies determine effective strategies to grapple with the dynamic security threatscape. Mike specializes in the sexy aspects of security — such as protecting networks and endpoints, security management, and compliance. Mike is one of the most sought-after speakers and commentators in the security business, and brings a deep background in information security. After 20 years in and around security, he's one of the guys who “knows where the bodies are buried” in the space.

Starting his career as a programmer and networking consultant, Mike joined META Group in 1993 and spearheaded META's initial foray into information security research. Mike left META in 1998 to found SHYM Technology, a pioneer in the PKI software market, and then held executive roles at CipherTrust and TruSecure. After getting fed up with vendor life, Mike started Security Incite in 2006 to provide a voice of reason in an over-hyped yet underwhelming security industry. After taking a short detour as Senior VP, Strategy at eIQnetworks to chase shiny objects in security and compliance management, Mike joined Securosis with a rejuvenated cynicism about the state of security and what it takes to survive as a security professional.

Mike published [The Pragmatic CSO](http://www.pragmaticcso.com/) <http://www.pragmaticcso.com/> in 2007 to introduce technically oriented security professionals to the nuances of what is required to be a senior security professional. He also possesses a very expensive engineering degree in Operations Research and Industrial Engineering from Cornell University. His folks are overjoyed that he uses literally zero percent of his education on a daily basis. He can be reached at mrothman (at) securosis (dot) com.

About Securosis

Securosis, LLC is an independent research and analysis firm dedicated to thought leadership, objectivity, and transparency. Our analysts have all held executive level positions and are dedicated to providing high-value, pragmatic advisory services. Our services include:

- **Primary research publishing:** We publish the vast majority of our research for free through our blog, and package the research as papers that can be licensed for distribution on an annual basis. All published materials and presentations meet our strict objectivity requirements, and follow our Totally Transparent Research policy.
- **Cloud Security Project Accelerators:** Securosis Project Accelerators (SPA) are packaged consulting offerings to bring our applied research and battle-tested field experiences to your cloud deployments. These in-depth programs combine assessment, tailored workshops, and ongoing support to ensure you can secure your cloud projects better and faster. They are designed to cut months or years off your projects while integrating leading-edge cloud security practices into your existing operations.
- **Cloud Security Training:** We are the team that built the Cloud Security Alliance CCSK training class and our own Advanced Cloud Security and Applied SecDevOps program. Attend one of our public classes or bring us in for a private, customized experience.
- **Advisory services for vendors:** We offer a number of advisory services to help our vendor clients bring the right product/service to market in the right way to hit on critical market requirements. Securosis is known for telling our clients what they NEED to hear, not what they want to hear. Clients typically start with a strategy day engagement, and then can engage with us on a retainer basis for ongoing support. Services available as part of our advisory services include market and product analysis and strategy, technology roadmap guidance, competitive strategies, etc. Though keep in mind, we maintain our strict objectivity and confidentiality requirements on all engagements.
- **Custom Research, Speaking and Advisory:** Need a custom research report on a new technology or security issue? A highly-rated speaker for an internal or public security event? An outside expert for a merger or acquisition due diligence? An expert to evaluate your security strategy, identify gaps, and build a roadmap forward? These defined projects bridge the gap when you need more than a strategy day but less than a long-term consulting engagement.

Our clients range from stealth startups to some of the best known technology vendors and end users. Clients include large financial institutions, institutional investors, mid-sized enterprises, and major security vendors. For more information about Securosis, visit our website: <http://securosis.com/>.