Threat Intelligence for Ecosystem Risk Management

Version 1.6
Released: September 9, 2013

Securosis, L.L.C. 515 E. Carefree Highway Suite #766 Phoenix, AZ 85085 **T** 602-412-3051
info@securosis.com www.securosis.com

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.

Licensed by BitSight Technologies

BitSight Technologies is transforming how companies manage information security risk with objective, evidence-based security ratings. The company's SecurityRating Platform continuously analyzes vast amounts of external data on security behaviors in order to help organizations make timely risk management decisions. Based in Cambridge, MA, BitSight is backed by Globespan Capital Partners, Menlo Ventures, Flybridge Capital Partners, and Commonwealth Capital Ventures. For more information, please visit www.bitsighttech.com or follow @BitSight on Twitter.

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/

Threat Intelligence for Ecosystem Risk Management Table of Contents

The Risk of the Extended Enterprise

A key aspect of business today is the extended enterprise. That's a fancy way of saying no organization goes it alone any more. Organizations have upstream suppliers who help produce whatever it is they make. Downstream distribution channels help sell whatever needs to be sold. Business processes are outsourced to third parties who handle them better and more cheaply. With the advent of advanced communication and collaboration tools, distributed multi-enterprise teams work on projects, even though they don't work for the same company or on the same continent. Jack Welch coined the term "boundaryless organizations" back in 1990 to describe an organization that is not defined by or limited to horizontal, vertical, or external boundaries imposed by a predefined structure. 23 years later, Welch was right. This kind of organization is common today.

In order to make the extended enterprise work, your business partners need access to your critical information. And that's where security folks tend to break out in hives.

In order to make the extended enterprise work, your business partners need access to your critical information. And that's where security folks tend to break out in hives. It's hard enough to protect your networks, servers, and applications while making sure your own employees don't do anything stupid to leave you exposed. Imagine your risk — based not just on how *you* protect your information, but also on how well all *your business partners* protect their information and devices as well. Actually, you don't need to imagine that — it's reality.

Let's try a simple thought exercise to get a feel for the risk involved in one of these interconnected business processes. For cost reasons a decision was made to outsource software maintenance on legacy applications to an offshore provider. These applications run in your datacenter, and maintenance is limited to simple bug fixes. You can't shut down the application but it's clearly not strategic. What's the risk here? Start getting a feel for your exposure by asking some questions:

- Which of your networks do these developers have access to? How do they connect in? Where are the connecting from?

- Who are the developers? Has the outsourcer done background checks on them? Are those checks trustworthy?
- What is the security posture of the outsourcer's network? What kinds of devices do they use? Assuming if the developers are trustworthy, how do you know their machines are not compromised by external attackers?
- Ultimately, can you make fact-based risk decisions about your ecosystem of business partners?

Even if you could answer those questions, an answer is only good at that point in time. Business partner connectivity requires a leap of faith, although fortunately you aren't totally powerless to protect yourself. You can segment your network to ensure developers only have access to the servers and code they are responsible for. You can scan devices to check its security posture when they connect to your network to ensure they aren't pwned. You can check developer backgrounds yourself. You can even audit the outsourcer's network.

And you can still get compromised via business partners, because things move too fast to really stay on top of everything. It takes seconds for a machine to be compromised. With one compromised machine your adversary can gain presence on your network and then move laterally to other devices within your environment that have privileged access. This happens every day.

You have very little visibility into business partner networks, which means additional attack surface you don't control. No one said this job was easy, did they? These interconnected business processes happen whether you like it or not. Even if you think they pose unacceptable risk. You can stamp your feet and throw all the tantrums you want, but unless you can show a business decision-maker that the risk of maintaining the connection is greater than the benefit of providing that access, you are just Chicken Little. Again.

You need due diligence to understand how each organization accessing your network increases your attack surface. You need a clear understanding of how much risk each business partner presents.

So you need due diligence to understand how each organization accessing your network increases your attack surface. You need a clear understanding of how much risk each business partner presents. You need to assess each partner and receive notification of any issues which appear to put your networks at risk. We call this an Early Warning System, and external threat intelligence can give you a head start on knowing which attacks are heading your way.

Here is an excerpt from our EWS paper to illuminate the concept.

> *You can shrink the window of exploitation by leveraging cutting-edge research to help focus your efforts more effectively, by looking in the places attackers are most likely to strike. You need an Early Warning System (EWS) for perspective on what is coming at you.*
>
> *None of this is new. Law enforcement has been doing this, well, forever. The goal is to penetrate the adversary, learn their methods, and take action before an attack. Even in security there is a lot of precedent for this kind of approach. Back at TruSecure (now part of Verizon Business) over a decade ago, the security program was based on performing external threat research and using it to prioritize the controls to be implemented to address imminent attacks. Amazingly enough it worked.*

Following up our initial EWS research we expanded into a few different aspects of threat intelligence, which provides the external content for the EWS. There is Network-based Threat Intelligence and Email-based Threat Intelligence, but both those sources focus on what's happening on your networks and with your brands. The kind of threat intelligence required to help you understand what's happening on your partner networks is a different animal, which we call *Ecosystem Threat Intelligence* (EcoTI). And the application of EcoTI to impact decisions and reduce/manage risk is what we call "Threat Intelligence for Ecosystem Risk Management."

In this paper we delve into both the general concept of Ecosystem Threat Intelligence and how to gather your own intelligence regarding the security of your business partners. By systematically assessing your extended network of business partners you are much better equipped to understand the risks they present to your organization. Armed with that information you will finally have sufficient data to understand and manage the real risk of connecting to business partners. It's definitely a good thing when you can make a case to senior management with facts, as opposed to conjecture or fear.

Assessing Partner Risk

As we discussed above, today's business environment features increasing and permanent use of an extended enterprise. Integrating systems and processes with business partners can benefit the business at the cost of dramatically expanding the attack surface. A compromised business partner, with trusted access to your network and systems, gives their attackers that same trusted access to *you*. To net out the situation you need to assess the security of your partner ecosystem and be in a position to make risk-based decisions about whether each connection (collaboration) with a business partner makes sense, and what types of controls are necessary for protection in light of the potential exposure.

Now let's consider how to quantify the risks of partners for better (more accurate) decisions about the levels of access and protection appropriate for them.

When assessing risks to your ecosystem, penetration tests or even vulnerability scans across all partners are rarely practical. You can certainly try (and for some very high-profile partners with extensive access to your stuff probably should), but you need a lower-touch way to perform ongoing assessments of the vast majority of your business partners. As with many other aspects of security, a leveraged means of collecting and analyzing threat intelligence on partners can identify areas of concern and help you determine whether and when to go deeper to perform active testing with specific partners.

A leveraged means of collecting and analyzing threat intelligence on partners can identify areas of concern and help you determine whether and when to go deeper to perform active testing with specific partners.

Breach History

Investors like to say past performance isn't a good indicator of future results. *Au contraire* — in the security business, if an organization has been compromised a number of times, they are considerably more likely to be compromised in the future. Some organizations use publicly disclosed data loss as a catalyst to dramatically improve security posture… but most don't. There are various sources for breach information, and consulting a few to confirm the accuracy of a breach report is a good idea. Aside from breach disclosure databases, depending on your industry you might have an ISAC (Information Sharing and Analysis Center) with information about breaches as well.

But we need to point out some limitations of this approach. First of all, many of the public breach reporting databases are volunteer-driven and can be a bit delayed in listing the latest breaches, primarily because the volume of publicly disclosed breaches continues to skyrocket. Some organizations (think military and other government) don't disclose their breaches, so there is no publicly accessible information about those organizations. And others play disclosure games with what is material and what isn't. Checking out public disclosures is not comprehensive, but it's certainly a place to start.

Mapping Your Ecosystem

The next step is to figure out whether the partner has current active security issues. But until you understand the extent of their systems and networks you have no idea of the attack surface you're dealing with. So you need to associate devices and IP addresses with specific business partners. If you have the proverbial "big bat" with a partner — meaning you do a lot of business with them and they have significant incentive to keep you happy — you can ask for this information. They may share it, or perhaps they won't — not necessarily because they don't want to, but more likely because they don't have it. It is very difficult to keep topology and device information accurate and current.

If you can't get it from your partner you need to build it yourself. That involves mining DNS and `whois` among other network mapping tactics, and is resource intensive. This isn't brain surgery, but if you have dozens (or more) of business partners it can be a substantial investment. Alternatively you might look to a threat intelligence service specializing in third-party assessment, which has developed such a mapping capability as a core part of their offering.

Another question on network mapping: how deep and comprehensive does the map need to be? Do you need to map every single network in use within a Global 2000 enterprise? That would be a very large number of networks to track, and you need to apply realistic constraints on time and investment when looking at a partner's attack surface. We recommend you start with specific locations that have direct access to your networks and stay focused on operational networks. Many organizations have large numbers of networks but use very few of them.

> If your partner can't protect stuff that is obviously under attack (Internet-facing devices), they probably don't do a good job with other security.

Public Malaise

Now that you have a map associating networks with business partners, you can start analyzing security issues on networks you know belong to business partners. Start with Internet-accessible networks and devices — mostly because you can get there. You don't need permission to check out a partner's Internet-facing devices. In-depth scanning and recon on those devices is bad form, but hopefully attackers aren't doing that every day, right?

If you find an issue it usually indicates a lack of security discipline. Especially if the vulnerability is simple. If your partner can't protect stuff that is obviously under attack (Internet-facing devices), they probably don't do a good job with other security. That is a coarse generalization but issues on Internet-facing devices fail the sniff test for good security practices.

Where can you get this information? Several data sources are available:

1. **Public website malware checking:** These services check for malware on websites — mostly by rendering pages automatically on vulnerable devices and seeing whether bad stuff happens. Often business partners buy these services themselves to assess their own security posture, but nothing says you can't quickly assess your partners as well. Especially a high-profile business partner. As an aside, it's not a bad idea to scan your own sites for malware as well, just to make sure you aren't throwing stones from a glass house.

2. **Phishing hosts:** If a site on a business partner's network is a known phishing host, the partner has a compromised Internet-facing server being used by a phisher, and they likely don't know about it.

3. **C&C nodes:** Like a phishing host, if your business partner is unwittingly hosting a C&C node at least one of their devices has been compromised, and they probably don't know.

4. **DNS issues:** If there are rogue DNS resolvers within your partner network or some other means of tampering with DNS results, that correlates with a higher likelihood of other security issues.

A lot of this analysis falls under the more comprehensive IP reputation analysis that many security vendors use in security and malware protection offerings. Access to the IP reputation databases of your strategic security providers can streamline the process of looking for issues with your business partners. But remember that IP reputation should not be the definitive arbiter of your security posture. You also need to look a bit deeper into business partner networks.

Sickness from Within

The other major category of threat intelligence helpful for assessing business partners is internal compromise. If you can get a feel for the number of compromised devices or bots participating in command and control networks, you can get a pretty good sense of the effectiveness of an organization's security program. Another reason to include internal compromise in your ecosystem risk assessment is that outward-facing devices are sitting ducks, so many organizations prioritize protecting them over internal devices. Not that that's a bad choice, but good external security doesn't guarantee the security of internal devices.

There are a few sources for this kind of data:

1. **IP reputation:** As discussed above, IP reputation factors in a lot of information that provides perspective on the intent of traffic originating from an IP address. With advanced NAT isolating the true IP address can be challenging, but you probably don't need it for useful threat intelligence. If you have significant bot traffic coming from an IP address within a partner network, whether that represents 2 or 20 bots is interesting but ultimately a detail.

2. **Botnet intelligence services:** A new type of threat intelligence is emerging which provides visibility into specific devices on botnets in one of two ways. One approach involves penetrating botnets and monitoring their command and control sources and destinations. The other tactic is to monitor partner networks and DNS for indicators of botnet traffic.

3. **Honeypots:** You can set up honeypot devices and networks to detect attacks from partner networks.

4. **Black/grey market accounts:** Finally, you can mine black and grey market sites peddling stolen email addresses and authentication credentials. If you see a bunch of your business partner's accounts for sale, that's a good indication of a bad security problem.

> If you can get a feel for the number of compromised devices or bots participating in command and control networks, you can get a pretty good sense of the effectiveness of an organization's security program.

You probably won't invest in all these different information sources, aggregate all this data, analyze it all yourself, and draw conclusions about the health of each business partner. But we work through the details to help you become an educated buyer of Ecosystem Threat Intelligence. This enables you to ask the right questions about how any intel provider gets their data and make sure it's relevant to your situation.

> Bots happen (we should sell t-shirts). But bots that stay operational over long periods of time indicate security underperformance.

Time to Remediate

Another way to leverage botnet intelligence is to track when specific devices are remediated. If you are mining a C&C network and see the same IP address for weeks or months at a time, that indicates that specific organization either doesn't

know the device is pwned or is inept at remediating the issue. Perhaps they are on top of things and you found their honeypot, but don't hold your breath.

The point isn't to harshly penalize every business partner for every bot that shows up on a specific network. Bots happen (we should sell t-shirts). But bots that stay operational over long periods of time indicate security underperformance. Similarly, Internet-facing devices that show indications of compromise for an extended period of time fall in the same category of security failure.

Quantifying Partner Risk

Once you have data on the devices and networks that show signs of being compromised, via the intelligence gathering process above, you can get a feel for the numbers of security issues at your business partners. Cross-reference your list of compromised devices against the map of business partner networks developed at the beginning of this process. This gives you data to draw a somewhat objective conclusion about how secure each partner is, and track each partner's security score over time.

When we talk about quantifying we are referring to some kind of objective, consistent, and repeatable means of generating a security posture metric. We could go into detail about what a good metric is and isn't, but we covered that ground in Security Benchmarking, so check that paper out for a deep discussion of metrics.

As long as the metrics are derived in a fair, objective, and consistent method — and you buy into the scoring algorithm and specific components of deriving the metrics — this relative approach to quantification offers value for decision support.

For Ecosystem Threat Intelligence we aren't overly concerned with specifics of how the metric is generated, so long as the approach is logical and defensible. We know the risk quantification crowd is likely to violently disagree, but we focus on how to make better security decisions — not on defending your metrics for a dissertation. As long as the metrics are derived in a fair, objective, and consistent method — and you buy into the scoring algorithm and specific components of deriving the metrics — this relative approach to quantification offers value for decision support. Keep your eye on the prize, which is understanding the relative risk of each business partner connecting to your network.

Use Cases and Selection Criteria

Now let's apply these concepts to a few use cases to make them a bit more tangible. We will follow a similar format for each use case, talking about the business needs for access, then the threat presented by that access, and finally how Ecosystem Threat Intelligence helps you make better decisions about specific partners.

Single Partner Business Process Outsourcing Use Case

Let's start simply. It is often cheaper and better to have external parties fulfill non-strategic functions. They could be anything from legacy application maintenance to human resources form processing. But almost all outsourcing arrangements require you to provide outsiders with access to your systems so they can use some of your critical data.

For any kind of software development an external party needs access to your source code. And unless you have a very advanced and segmented development network, developers have access to much more than just the applications they are working on. So if any of *their* devices are compromised attackers can gain access to *your* developers' devices, code repositories, build systems, and a variety of other things that would be bad.

If we are talking about human resources outsourcing, those folks have access to personnel records, which may include sensitive information such as salaries, employment agreements, health issues, and other stuff you probably don't want published on Glassdoor. Even better, organizations increasingly use SaaS providers for HR functions, which moves that data outside your data center and removes even more of your waning control.

The commonality between these two outsourcing situations is that access is restricted to just one business partner. Of course you might use multiple development shops, but for simplicity's sake we will just talk about one for now. In this case your due diligence occurs while selecting the provider and negotiating the contract. That may entail demanding background checks on external staffers and a site visit to substantiate sufficient security controls.

At that point you should feel pretty good about the security of your business partner. But what happens after that? Do you assess these folks on an ongoing basis? What happens if they hire a bad apple? Or they are attacked and compromised due to some issue that has nothing to do with you? Thus the importance of an ongoing assessment capability. If you are a major client of your outsourcer you might have a big enough stick to get them to share their network topology.

In this scenario you are predominately looking for indicators of command and control activity (described under Sickness from Within above) because that's the smoking gun for compromised devices with access. Compromised Internet-facing devices can also cause issues so you need to consider them too. But in this use case it makes sense to prioritize internal issues over public-facing vulnerabilities when you calculate a relative risk score. In this limited scenario it is not actually a *relative* risk score because you cannot really compare the provider to anyone else with access to the same dataset.

Many Partners Use Case

To complicate things a bit you might need to provide access to multiple business partners to the same data. Perhaps external sales reps have access to your customer database and other proprietary information about your products and services. Or perhaps your chain of hospitals provides access to medical systems to hundreds of doctors with privileges to practice at your facilities. Or it could even be upstream suppliers who make and assemble parts for your heavy machinery products. These folks have your designs and sales forecasts because they need to deliver inventory just in time for you to get the product out the door (and hit your quarterly numbers).

> Specifics from your Ecosystem Threat Intelligence efforts enable you to make a fact-based case to senior executives that connecting to a partner is not worth the risk. Again, you can't make the business decision, but you can arm the decision-maker with enough information to make an informed decision.

Regardless of the details, you need to support dozens of business partners (or more), offering access to some of your most critical enterprise data. Sometimes it's easier for targeted attackers to go after business partners than target you directly. We have seen this in the real world, with subassembly manufacturers of defense contractors hacked for access to military schematics and other critical information on particular weapons programs.

In this situation, as in the use case above, the security team typically needs an airtight case to prevent a connection with a partner. Sales executives frown on the security team shutting down a huge sales channel. Similarly, security folks cannot tell the final assembly folks they can't get their seats because the seat manufacturer got breached. You cannot stop the business, but you can certainly warn the senior team about the risks of connecting with a specific business partner, and to substantiate those concerns you need *data*.

This is where relative risk scores for multiple business partners can help make your case. It is wise to assume that all business partners are compromised in some fashion. But which are total fiascos? Which partners cannot even block a SQLi attack on an online commerce site? Which has dozens of bots flooding the Internet with denial of service attacks? Specifics from your Ecosystem Threat Intelligence efforts enable you to make a fact-based case to senior executives that connecting to a partner is not worth the risk. Again, you can't make the business decision, but you can arm the decision-maker with enough information to make an informed decision.

Or you could suggest an alternative set of security controls for those specific partners. You might force them to connect into your systems through a VDI (virtual desktop) service on your premises so your data never leaves your network, and monitor everything they do in your systems. There are a number of ways to deal with vulnerable business partners, but you need to start with the knowledge that they are vulnerable.

Buy Versus Build

As we alluded earlier, gathering and analyzing this data on business partners is non-trivial. Especially if you are dealing with hundreds or thousands of partners. Just building the maps to understand your partner attack surface is resource-intensive. Then assessing networks, penetrating botnets, setting up honeypots, and the like, is a series of major endeavors. And given your list of things to do every day, building an EcoTI capability might not happen — even though it should.

As we alluded earlier, gathering and analyzing this data on business partners is non-trivial. Especially if you are dealing with hundreds or thousands of partners.

The good news is that emerging threat intelligence providers are focusing on third-party assessment. They gather the data to assess business partners and provide a quantified way to compare business partner security postures. There are not many providers today, but we expect growth in this capability over the next 2-3 years.

That said, you need some way to evaluate these services for applicability to your environment, so here are a few things to focus on in any discussion with an EcoTI provider.

1. **Data sources:** The first thing to understand is where the provider gets their data. Is it internal research? Do they buy external threat feeds? What kind of information can they detect? Don't be surprised if the provider requires you to sign a non-disclosure agreement for access to this information, because quality and breadth of data sources can be important competitive advantages.

2. **Network topology mapping:** How does the provider figure out which devices and networks are controlled by your business partner? How long does it take them to get coverage for your list of partners? How do they keep their lists current?

3. **Risk scoring:** If the provider makes a judgement about the security posture of a partner, how is that score derived? What is their quantification based on? Can you tune it for your specific situation and risk tolerances?

4. **Integration with enterprise systems:** Is there a clean way to get alerts or other data into your enterprise systems? It would be great if, when the EcoTI service finds a serious problem with a business partner, they could feed that directly into your SIEM or GRC platform to streamline your response.

5. **Viability:** In emerging markets you want to deal with organizations which will be there tomorrow (and hopefully the next day even after). So it is not out of bounds to get a feel for funding, senior team experience, level of resources, etc.

Look for shorter deals at this state of the market. We expect a lot of innovation and many new players to emerge so it makes sense to maintain your flexibility in EcoTI providers as the market matures — other services might better meet your needs soon.

Conclusion

The role of the security team is not to inhibit business operations — it's to ensure the security of enterprise data given the requirements of the business. Increasingly as the global drive towards "boundaryless organizations" continues, security teams are expected to not only understand the security posture of their own networks and devices, but also the security posture of any business partners accessing corporate data.

That is a tall order given the sheer number of business partners enterprises have to support. The challenge is compounded by the difficulty in collecting, analyzing, and alerting on the massive amount of data involved in assessing the risk of any organization's ecosystem. But just because it's hard doesn't mean it isn't worthwhile. Given the ongoing increases in data loss resulting from compromised business partners, Ecosystem Threat Intelligence is not something you can ignore forever.

This paper details what is involved in developing an Ecosystem Threat Intelligence capability to provide an early warning of security issues within business partner networks, which could eventually spread to your environment. Whether you decide to build your own EcoTI capability or buy a third party EcoTI service, factoring in the security posture of your extended enterprise can help you quantify the risks of partners for better decisions about the levels of access and protection appropriate for them.

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 or ask via the Securosis Nexus <http://nexus.securosis.com/>.

About the Analyst

Mike Rothman, Analyst/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/> 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:

- **The Securosis Nexus**: The Securosis Nexus is an online environment to help you get your job done better and faster. It provides pragmatic research on security topics that tells you exactly what you need to know, backed with industry-leading expert advice to answer your questions. The Nexus was designed to be fast and easy to use, and to get you the information you need as quickly as possible. Access it at <https://nexus.securosis.com/>.

- **Primary research publishing**: We currently release the vast majority of our research for free through our blog, and archive it in our Research Library. Most of these research documents can be sponsored for distribution on an annual basis. All published materials and presentations meet our strict objectivity requirements and conform to our Totally Transparent Research policy.

- **Research products and strategic advisory services for end users**: Securosis will be introducing a line of research products and inquiry-based subscription services designed to assist end user organizations in accelerating project and program success. Additional advisory projects are also available, including product selection assistance, technology and architecture strategy, education, security management evaluations, and risk assessment.

- **Retainer services for vendors**: Although we will accept briefings from anyone, some vendors opt for a tighter, ongoing relationship. We offer a number of flexible retainer packages. Services available as part of a retainer package include market and product analysis and strategy, technology guidance, product evaluation, and merger and acquisition assessment. Even with paid clients, we maintain our strict objectivity and confidentiality requirements. More information on our retainer services (PDF) is available.

- **External speaking and editorial**: Securosis analysts frequently speak at industry events, give online presentations, and write and/or speak for a variety of publications and media.

- **Other expert services**: Securosis analysts are available for other services as well, including Strategic Advisory Days, Strategy Consulting engagements, and Investor Services. These tend to be customized to meet a client's particular requirements.

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.

Additionally, Securosis partners with security testing labs to provide unique product evaluations that combine in-depth technical analysis with high-level product, architecture, and market analysis. For more information about Securosis, visit our website: <http://securosis.com/>.