Defending Against Network-based Distributed Denial of Service Attacks

Version 1.5

Released: April 17, 2014

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.

This report is licensed by A10 Networks,
which allows us to release it for free.
All content was developed independently.

www.a10networks.com

A10 Networks (NYSE: ATEN) is a leader in application networking, providing a range of high-performance application networking solutions that accelerate and secure data center applications and networks of thousands of the largest enterprise, service provider and hyper-scale web providers around the world. Founded in 2004, A10 Networks is based in San Jose, CA, and serves customers globally with offices worldwide.

Copyright

Defending Against Network DDoS Attacks

Table of Contents

Introduction

Back in 2013 volumetric denial of service (DoS) attacks on networks were all the rage. 'Hacktivists' first used them effectively against Fortune 500 class banks, largely knocking down major banking brands for days at a time. But these companies quickly adapted and gained proficiency at defending themselves, so attackers shifted targets and bifurcated their tactics. They went after softer targets like public entities and smaller financial institutions. They also took on content delivery networks like CloudFlare with multi-hundred-gigabyte volumetric attacks leveraging new protocols (including DNS and NTP) and amplification techniques, just because they could.

In our Defending Against Denial of Service Attacks[1] research we described network-based DoS attacks:

> Network-based attacks overwhelm the network equipment and/or totally consume network capacity by throwing everything including the kitchen sink at a site — this interferes with legitimate traffic reaching the site. This volumetric type of attack is what most folks consider Denial of Service, and it realistically requires blasting away from many devices, so current attacks are called Distributed Denial of Service (DDoS). If your adversary has enough firepower it is very hard to defend against these attacks, and you will quickly be reminded that though bandwidth may be plentiful, it certainly isn't free.
>
> Application-based attacks are different — they target weaknesses in web application components to consume all the resources of a web, application, or database server to effectively disable it. These attacks can target either vulnerabilities or 'features' of an application stack to overwhelm servers and prevent legitimate traffic from accessing web pages or completing transactions.

The motivation behind these attacks hasn't changed much. Attackers tend to be either organized crime factions threatening to take commerce sites down, or hacktivists trying to make a point. We do see a bit of competitor malfeasance and Distributed DoS (DDoS) to hide exfiltration activities, but those don't seem to be primary use cases any more.

---

[1] https://securosis.com/research/publication/defending-against-denial-of-service-dos-attacks

Regardless of motivation attackers now have faster networks, bigger botnets, and increasingly effective tactics to magnify the impact of their DDoS attacks — organizations can no longer afford to ignore them. After digging deeper into the application side of denial of service in Defending Against Application Denial of Service Attacks[2], we now turn our attention to the network side of the house.

It's Getting Easier

It is easier than ever to launch large-scale network-based DDoS attacks. There are a few main reasons:

- **Bot availability:** More devices are compromised every day. Sophisticated malware kits are available to make attacking the devices even easier. As a result there are millions of compromised devices (predominately belonging to consumers), available to be brought to bear in DoS attacks.
- **Faster consumer Internet:** Network speeds at homes and small offices continue to climb. This enables consumer bots to blast targets with growing bandwidth, and this trend will continue as networks get faster.
- **Cloud servers:** It is uncommon to see 50Mbps sustained from a consumer device, but quite feasible at the server level. Combine this with the fact that cloud servers are Internet-facing, and it's no surprise that attackers now use compromised cloud servers to blast DDoS targets. This kind of activity is harder to detect because these servers *should* be pumping out high volumes of traffic.
- **Magnification:** Finally, attackers are getting better at magnifying the impact of their attacks, manipulating protocols like DNS and NTP to provide order-of-magnitude magnification of traffic hitting the target site. This makes far better use of attacker resources, allowing them to use each bot sporadically and more lightly to evade detection.

Limitations of Current Defenses

Before we dive into specifics of how these attacks work, we need to mention why existing network and security devices aren't well-suited to defend against current DDoS attacks. It is not a question of core throughput — we see service provider network firewalls processing upwards of 500gbps of traffic, and getting faster rapidly. But the devices aren't architected to deal with floods of legitimate traffic from thousands of devices, while maintaining state and providing the packets per second throughput required to handle legitimate connections.

Existing network and security devices aren't well-suited to defend against current DDoS attacks, the devices just aren't architected to deal with floods of legitimate traffic from thousands of devices

[2] https://securosis.com/research/publication/defending-against-application-denial-of-service-attacks

Even with NGFW capabilities providing visibility into web and other application traffic; dealing with millions of active connection requests can exhaust link, session, and application handling capacity on security devices — regardless of maximum possible throughput.

IPS devices are in the same boat but their job even is harder because they are actively looking for attacks and profiling activity to find malicious patterns. They are far more compute-intensive, and as such have an even harder time keeping pace with DDoS traffic volumes. Many attackers actually target firewalls and IPS devices with DDoS attacks, knowing they typically fail closed and disable the target network — neatly achieving the goal.

You should start by over-provisioning your networks. This has always a been a popular tactic among networking folks: throw more bandwidth at the problem. Unfortunately you probably can't compete with a botmaster leveraging the aggregate bandwidth of all their compromised hosts. And it gets expensive to provision enough unused bandwidth to deal with a DDoS traffic spike.

You can also look at CDNs (Content Delivery Networks) and/or DoS scrubbing services, as we will describe later in this paper. Unfortunately CDN offerings may not offer full coverage of your entire network and are increasingly DDoS targets themselves. Scrubbing centers can be expensive, and shifting traffic to the scrubbing center still entails downtime. Finally, part-time scrubbing is inherently reactive — you are likely to be down before you even discover you have a problem.

Further complicating things is the fundamental challenge of simply detecting the onset of a DDoS attack. How can you tell the difference between a temporary spike in traffic and a full-on blitzkrieg? It is hard to be sure and some approaches (particularly those based on looking for anomalies) make it difficult to strike a good balance between pulling the alarm (and shifting traffic to the scrubbing center) in time, against waiting too long and incurring significant downtime. Even if you have devices in place to deal with DDoS traffic, you can still be taken down if you don't time things correctly.

So what to do? You need a multi-pronged approach to adequately protect networks against DDoS attacks.

So what to do? You need a multi-pronged approach to adequately protect networks against DDoS attacks. This paper will delve into the types of attacks and how they are changing. We will also help you understand how attackers are increasing the impact of their tactics to blast targets with unprecedented amounts of traffic. Finally we will discuss specific tactics for defending yourself.

The Attacks

DDoS is a blunt force instrument for many adversaries. Easy to launch and hard to defend against, so it will remain a staple in attacker toolkits for the foreseeable future. There is not much elegance in a volumetric attack — adversaries impair network availability by consuming all the bandwidth into or out of a site and/or knocking down network and security devices, overwhelming their ability to handle traffic. But attackers don't get points for elegance — they only need to achieve their mission, which is to knock down networks.

There is not much elegance in a volumetric attack. But attackers don't get points for elegance — they only need to achieve their mission, which is to knock down networks.

Today's traditional network and security devices (routers, firewalls, IPS, etc.) were not designed to handle modern attacks. Nor were network architectures built to decipher this volume of attack traffic and keep legitimate traffic flowing. So an additional layer of products and services has emerged to protect networks from DDoS attacks. Before we dig into ways to deal with these attacks, let's review the types of attacks and how attackers assemble resources to blast networks to virtual oblivion.

The Attacks

The first category of DDoS attacks is the straightforward flood. Attackers use tools that send requests using specific protocols or packets (SYN, ICMP, DNS, and NTP are the most popular) but don't acknowledge responses. If enough computers send requests to a site its bandwidth will quickly be exhausted. Even if bandwidth is sufficient, on-site network and security devices need to maintain session state for each bogus connection, while continuing to handle additional (legitimate) inbound session requests — which can consume all resources on the device. Despite their simplicity floods continue to be very effective.

Increasingly we see the DNS (Domain Name System) infrastructure, which serves as the map of the Internet, targeted by DDoS attacks. This prevents the network from routing traffic from point A to point B. As with floods, attackers can overwhelm the DNS by blasting it with traffic, especially because its decades-old design didn't anticipate today's traffic volumes.

DNS has other frailties which make it an easy target for DDoS. Like the shopping cart and search attacks highlighted in Application DoS, legitimate DNS queries can also overwhelm DNS servers, impacting legitimate visitors' ability to find sites. Attacks target weaknesses in the DNS system, where a single request for resolution can trigger 4-5 additional downstream DNS requests. This

leverage can overwhelm domain name servers. Similarly, attackers may request addresses for hosts that do not exist, causing targeted DNS servers to waste resources passing on requests and polluting caches with garbage to further impair performance.

Finally, HTTP continues to be a popular target for floods and other application-oriented attacks, targeting inherent weaknesses in the protocol. We discussed slow HTTP attacks in the Application Denial of Service paper so we won't rehash the details here, but remediations for volumetric attacks alleviate slow HTTP attacks as well.

Assembling the Army

To launch a volumetric attack an adversary needs devices across the Internet to pound target networks with traffic. Where do these devices come from? If you were playing Jeopardy the correct response would be "What is a bot network, Alex?" Consumer devices continue to be compromised and monetized at an increasing rate, thanks to increasingly sophisticated malware and the lack of innovation in consumer endpoint protection. These compromised devices generate the bulk of DDoS traffic.

Where do these devices come from? If you were playing Jeopardy the correct response would be "What is a bot network, Alex?"

Of course attackers need to be careful — Internet Service Providers are increasingly sensitive to consumer devices streaming huge amounts of traffic at arbitrary sites, and take them off the network when they find terms of service violations. Bot masters use increasingly sophisticated algorithms to control compromised devices, protecting them from both detection and remediation. Another limitation of consumer devices is constrained bandwidth, particularly upstream. Even as bandwidth continues to grow around the world, DDoS attackers hit capacity constraints as they continue scaling up attack volumes.

DDoS attackers work around the limitations of consumer devices by adding compromised servers to the arsenal. Given the millions of businesses with vulnerable Internet-facing devices, it remains unfortunately trivial for attackers to compromise servers with much higher upstream bandwidth. This makes servers much better at serving up malware, commanding and controlling bot nodes, and launching DDoS attacks.

Attackers are currently moving a step beyond conventional servers, capitalizing on cloud services to improve their economics. Cloud servers — particularly Infrastructure as a Service (IaaS) servers — are usually Internet-facing and often poorly configured, and of course their bandwidth is substantial. For network attacks a cloud server is like a conventional server on steroids — DDoS attackers realize major gains in both efficiency and leverage. To be fair, better cloud providers take great pains to identify compromised devices and notify customers when they notice something amiss. Unfortunately by the time misuse is detected by a cloud provider, server owner, or other server host, it may be too late. It doesn't take long to knock a site offline.

Attackers without the resources or desire to assemble and manage botnets can just rent them. DDoS attack service costs from $2/hour for short attacks up to $1,000 to take a site down for a month.

Attackers without the resources or desire to assemble and manage botnets can just rent them. Yes, a number of folks offer DDoS as a service (DDoSaaS for the acronym hounds), so it couldn't be easier to harness resources to knock down a victim. And it's not expensive[3] according to McAfee, which reported that a DDoS attack service costs from $2/hour for short attacks up to $1,000 to take a site down for a month.

It is a bit scary to think you could knock down someone's site for 4 hours for less than the cost of a cup of coffee. But when you take a step back to consider the easy availability of compromised devices, servers, and cloud servers, DDoS is a very easy service to add to an attacker's arsenal.

[3] http://www.mcafee.com/uk/resources/white-papers/wp-cybercrime-exposed.pdf

Magnification

The predominant mechanism for network-based DDoS attacks is to flood the pipes with common protocols like SYN, ICMP, DNS, and NTP. But as ISPs and other scrubbing services learned to deal with these floods, attackers ratcheted attacks to the next level of volume. They now leverage protocol weaknesses to magnify the impact of floods by an order of magnitude. This makes each compromised device far more effective and allows attackers to scale attacks over 400gbps (as recently reported by CloudFlare). Only a handful of organizations in the world can handle attacks of that magnitude, so DDoS + magnification is a potent combination.

Fat Packets

The most obvious way to beef up a volumetric attack is to simply send larger packets, clogging the pipes that much faster. For example, simple SYN packets can crush the computational capabilities of network/security devices. Combining small SYNs and large SYNs can also quickly saturate the network pipe, so we see often them combined in volumetric attacks.

Reflection + Amplification

Another technique for magnifying DDoS attacks is reflection. This entails sending requests to a large number of devices (think millions), using the spoofed IP address of a target site as their origin. The replies to those millions of requests hammer the target. This works because the UDP-based protocols used (ICMP, DNS, NTP, etc.) don't require handshaking to establish new sessions, so origin IPs can be spoofed.

DNS can provide about 50x amplification: responses are about 50x larger than requests. To scale attacks further, NTP has recently become a DDoS protocol of choice because it offers almost 200x magnification.

The latest wave of DDoS attacks uses reflected DNS and NTP traffic to dramatically scale the volume of traffic hitting targets. Both protocols provide good leverage: DNS and NTP responses are typically much bigger than requests. DNS can provide about 50x amplification: responses are about 50x larger than requests. And the number of open DNS resolvers which respond to any DNS request from any device make this attack easy and scalable. Until the major ISPs get rid of these open resolvers DNS-based DDoS attacks will continue.

NTP has recently become a DDoS protocol of choice because it offers almost 200x magnification. That makes attacks much more effective. Again, this is enabled by a protocol 'feature': clients can

request a list of the last 600 IP addresses to access a server. To illustrate its magnitude, the CloudFlare folks reported being targeted with an NTP reflection attack using 4,529 NTP servers, running on 1,298 different networks, each sending about 87mbps to the victim. The resulting traffic totaled about 400gbps. Even more troubling, all those requests (to 4,500+ NTP servers) could be sent from one device on one network, quickly and with little chance of detection.

Other UDP-based protocols offer even greater amplification. An SNMP response can be 650x the size of a request, for better than 3 times the leverage of NTP. This could theoretically be weaponized to create 1gbps+ DDoS attacks — awesome, so long as you aren't the target.

Stacking Attacks

Of course none of these techniques exists in a vacuum — sometimes we see attackers pounding a target directly, while other times they combine reflection and amplification to hammer the target even harder. These attacks are enabled by sloppy network hygiene on the part of Internet service providers who allow spoofed IP addresses for these protocols and don't block floods. Unfortunately these issues are largely beyond the control of enterprise targets, leaving victims with little option but to respond with a bigger pipe to absorb attacks and put mitigations in place to defend against them.

Mitigations

Let's move on to defenses. To illustrate what you are up against we'll take a small excerpt from our Defending Against Denial of Service Attacks paper.

> First the obvious: you cannot just throw bandwidth at the problem. Your adversaries likely have an unbounded number of bots at their disposal and are getting smarter at using shared virtual servers and cloud instances to magnify the amount at their disposal. So you can't just hunker down and ride it out. They likely have a bigger cannon than you can handle. You need to figure out how to deal with a massive amount of traffic, and separate good traffic from bad while maintaining availability.

Your first option is to leverage existing network/security products to address the issue. As we discussed earlier, this is a problematic strategy because those devices aren't built to withstand the volumes or tactics in a modern volumetric DDoS. Next, you could deploy a purpose-built device on your network to block DDoS traffic before it melts your networks. This is an important step but if your inbound network pipes are saturated an on-premise device cannot help much — applications will be unavailable whether or not any attack traffic reaches the on-premise device on the internal network. Finally, you can front-end your networks with a service to scrub traffic before it reaches your network. But this is no panacea either — it takes time to move traffic to a scrubbing provider, during which window you are effectively down.

So the answer is typically a combination of tactics deployed in a complimentary fashion to give you the best chance at maintaining availability in the face of a network volumetric attack.

Do Nothing

Before we dig into alternatives we need to acknowledge another option: *doing nothing*. Many organizations must go through an exercise after being hit by a DDoS attack, to determine what protections are needed. Given the investment required for any of these mitigations you need to weigh the cost of downtime against the cost of (hopefully) stopping attacks. You can spend a bunch of money on defenses, and then get hammered with a 100gbps attack which crushes your network.

So you face another security tradeoff. If you are a frequent or high-profile target then doing nothing isn't an option. If you got hit with a random attack — which may happen when attackers are testing new tactics and code, or perhaps if you are just unlucky — and have no reason to expect to be targeted again, you may be able to get away with doing nothing. If that is your chosen risk, you need

to make sure all the relevant parties are aware of what it actually means, and to manage expectations so they understand your network will crumble if you get attacked again.

We don't advocate the do-nothing approach, but we do understand that tough decisions need to be made in the face of scarce resources. Assuming you want to put some defenses in place to mitigate the impact of a DDoS, we will now work through the alternatives.

DDoS Defense Devices

These appliances are purpose-built to deal with DoS attacks; they include both optimized IPS-like rules to detect and prevent floods and other network anomalies, and simple web application firewall capabilities to protect against application-layer attacks. Additionally, they feature anti-DoS features such as session scalability and embedded IP reputation capabilities to discard traffic from known bots without full inspection.

To understand the role of IP reputation let's recall how email connection management devices enabled anti-spam gateways to scale up to handle spam floods. It is computationally expensive to fully inspect every inbound email, so immediately dumping messages from known bad senders focuses inspection resources on email that might be legitimate to keep mail flowing. The concept applies here as well. Keep the latency involved in checking with a cloud-based reputation database in mind — you want the device to aggressively cache bad IPs to avoid a lengthy cloud lookup for every incoming session.

DDoS Defense Devices are purpose-built to deal with DoS attacks; they include both optimized IPS-like rules and simple web application firewall capabilities. Additionally, they feature anti-DoS features.

For kosher connections which pass the reputation test, these devices additionally enforce limits on inbound connections, govern the rate of application requests, control client request rates, and manage the total number connections allowed to reach servers or load balancers. Of course these limits must be defined incrementally and tuned over time to avoid shutting down legitimate traffic during peak usage.

Speed is the name of the game for DDoS defense devices, so make sure any device you consider has sufficient headroom to handle your network pipe. Overprovision to ensure the device can handle bursts and keep up with future bandwidth increases. Also be sure to test the device in your environment with traffic representative of your network. Let's just say data sheet performance numbers don't always equal real world throughput.

CDN/Web Protection Services

Another popular option is to front-end web applications with a content delivery network or web protection service[4]. This tactic only protects web applications you route through the CDN, but if such an applications is targeted a CDN can scale to handle very large DDoS attacks cost-effectively. Though if the attacker targets other unprotected addresses or ports, you are out of luck. DNS servers, for instance, aren't protected.

Our research indicates that CDNs can be effective for handling network-based DDOS in smaller environments with a small external web presence. There are plenty of other benefits to a CDN, including caching content to improve website performance and shielding your external IP addresses. But for stopping DDoS attacks a CDN is a limited answer.

External Scrubbing

The next level beyond a CDN in sophistication (and cost) is routing your traffic through an external scrubbing center. These services allow you to redirect *all* your traffic or just the protocol under attack through their network when you are attacked, with the scrubber absorbing the excessive DDoS traffic. You trigger the switchover using either a proprietary switching protocol (if your perimeter devices or DDoS defense appliances support the carrier's signaling protocol), redirecting DNS, or via a BGP request. Once the determination has been made to move traffic to the scrubbing center it takes some time for the network to converge, impacting availability while the Internet routes traffic properly through the scrubbing center. Once the service kicks in you receive clean traffic through a tunnel from the scrubbing center.

The question with a scrubbing center is when to move the traffic. Finding that balance is a company-specific decision based on the perceived cost of downtime compared to the cost and value of the protected service.

The question with a scrubbing center is *when* to move the traffic. Pull the trigger too soon and your resources stay up, but at excessive cost. Do it too late and you can suffer additional downtime. Finding that balance is a company-specific decision based on the perceived cost of downtime compared to the cost and value of the protected service.

Another blind spot in the scrubbing approach is *hit and run* attacks, when an attacker blasts a site briefly to take it down. Once the victim moves the traffic over to a scrubbing center, the attacker stops, not even trying to take down the scrubber. But the attack has already achieved its goals: disrupted availability and increased latency.

These factors have pushed scrubbing centers to advocate for an *always on* approach where the customer runs *all* traffic through the scrubbing center all the time. Obviously there is a cost, but if you are a frequent DDoS target or cannot afford downtime for any reason, it may be worthwhile.

---

[4] https://securosis.com/research/publication/quick-wins-with-website-protection-services

All of the Above

As we stated in Defending Against DoS Attacks, the best answer is often all the above. The choice of network-based DoS mitigations inevitably involves trade-offs. It is no good to over-generalize but most organizations are best served by a hybrid approach involving both an on-premise appliance and a contract with a CDN or anti-DoS service provider to handle more severe volumetric attacks. It is rarely cost-effective run all traffic through a scrubbing center, and many DoS attacks target the application layer — in which case you need an on-premise device anyway.

Other Protection Tactics

Given that many DDoS attacks also target DNS, you will want to make sure your internal DNS infrastructure is protected by front-ending your DNS servers with a DDoS defense device. You will also want to perform some due diligence on your external DNS provider to ensure they have sufficient protections in place against DDoS, as they will be targeted along with you, and you could be impacted if they fall over — even if a different customer is targeted.

You don't want to be a part of the problem, so as a matter of course you should make sure you aren't responding to public NTP requests, as explained by US-CERT[5]. Additionally, you will want to remediate compromised devices as quickly as practical for many reasons, not least to ensure they don't blast others with your resources and bandwidth. It's living by the "Golden Rule," doing for others as you'd have them do for you.

A strong underlying process is your best defense against a DDoS attack. The good news is that the DoS defense process is quite similar to general incident response.

The Response Process

A strong underlying process is your best defense against a DDoS attack. Tactics change as attack volumes increase but if you don't know what to do when your site goes down it will be out for a while.

The good news is that the DoS defense process is quite similar to general incident response. We have already published a ton of research on this topic, so check out both our Incident Response Fundamentals series[6] and our React Faster and Better[7] paper. If your incident handling process isn't yet where it needs to be start there.

---

[5] https://www.us-cert.gov/ncas/alerts/TA14-013A

[6] https://securosis.com/blog/incident-response-fundamentals-index-of-posts

[7] https://securosis.com/Research/Publication/react-faster-and-better-new-approaches-for-advanced-incident-response

Building off the IR process already in place, think about what you need to do as a set of activities: before, during, and after an attack:

- **Before:** When you are not under attack, spend time figuring out attack indicators and ensuring you perform sufficient monitoring to provide adequate warning and collect enough information to identify the root cause. You might see increasing bandwidth volumes or a spike in DNS traffic. Perhaps your applications get flaky and fall down, you see server performance issues, or your CDN alerts you to a possible attack. Unfortunately many DDoS attacks come out of nowhere, so you may not know you are under attack until your networks are down.

- **During:** How can you restore service as quickly as possible? You can start by identifying the attack accurately and remediating effectively. You will also need to notify the powers that be, assemble your team, and establish responsibilities and accountability. Then focus on identifying root cause, attack vectors, and adversaries to figure out the best way to get the site back up. Restoring service depends on the mitigations at your disposal, as discussed above. Ideally your contracted CDN and/or anti-DoS service provider already has a team working on the problem by this point. In case you don't have one, hopefully the attack won't last long or your ISP can help you.

- **After:** Once the attack has been contained focus shifts to restoring normal operations, moving traffic back from the scrubbing center, and perhaps loosening anti-DoS/WAF rules. That doesn't mean you shouldn't keep monitoring for trouble. Then the focus turns to making sure this doesn't happen again. This involves asking questions: What worked? What didn't? Who needs to be added to the team? Who just got in the way? This analysis needs to objectively identify the good, the bad, and the ugly. Dig into the attack as well. What controls would have blunted its impact? Would running all your traffic through a scrubbing provider have helped? Did network redirection work quickly enough? Did you get the right level of support from your service provider? Then update your process as needed and implement new controls if necessary.

Summary

To wrap up this paper on network-based DDoS, let's recap a few key points:

- DoS today encompasses network attacks, application attacks, and magnification techniques to confuse defenders and exploit weaknesses in defenses.
- Organizations need a multi-faceted approach to defend against DDoS, which typically involves both deploying DDoS defense equipment on-site and contracting with a service provider (either a scrubbing center or a content delivery network) to handle excessive traffic.
- DoS mitigations in isolation are inadequate — on-premise devices and services are interdependent for adequate protection, and should communicate with each other to ensure efficient and transparent transition to the scrubbing service when necessary.

Of course there are trade-offs with DDoS defense, as with everything. Selecting an optimal mix of defensive tactics requires some adversary analysis, an honest and objective assessment of just how much downtime is survivable, and clear understanding of what you can pay to restore service quickly. If a few hours of downtime are survivable, defensive tactics can be much different compared to situations where no downtime is ever acceptable — which demands more expenditure and much more sophisticated defenses.

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/>.