VERISIGN™

DDoS MITIGATION

BEST PRACTICES FOR A RAPIDLY CHANGING THREAT LANDSCAPE

CONTENTS

ALTHOUGH DISTRIBUTED DENIAL OF SERVICE (DDoS) ATTACKS HAVE BECOME A MAINSTAY OF HACKERS' ARSENALS, THEIR PROFILE HAS CHANGED CONSIDERABLY IN THE PAST YEAR OR SO, MAKING THEM AN EVEN GREATER THREAT TO COMPANIES THAT CONDUCT BUSINESS ONLINE OR HAVE SIGNIFICANT INVESTMENTS IN THEIR ONLINE BRAND AND REPUTATION. DDoS ATTACKS ARE LARGER, STEALTHIER, MORE TARGETED AND MORE SOPHISTICATED THAN EVER. INCREASINGLY, EVEN AMATEURS CAN EXECUTE ATTACKS THEMSELVES OR CHEAPLY RENT BOTNETS TO DO THE JOB FOR THEM.

Given the extraordinary and rapid changes in the DDoS terrain, traditional DDoS mitigation tactics such as bandwidth overprovisioning, firewalls and intrusion prevention system (IPS) devices are no longer sufficient to protect an organization's networks, applications, and services. Verisign has successfully defended its global DNS infrastructure against DDoS and other attacks for more than 12 years and has maintained 99.99 percent availability of its critical infrastructure during that time. In addition, Verisign has maintained 100 percent availability of its .net and .com infrastructure and resolves more than 60 billion DNS transactions per day. Drawing on this success and hands-on engagements with customers in a range of industries, Verisign has identified a set of best practices that enables organizations to keep pace with DDoS attacks while minimizing impact on business operations. This paper describes these practices.

SMARTER, STEALTHIER, AND MORE ADAPTIVE

In a recent Forrester survey[1] of 400 IT decision-makers in the United States and Europe, 74 percent of respondents reported experiencing one or more DDoS attacks in the past year—even though they had security measures in place to prevent such an attack. Thirty-one percent of these attacks resulted in service disruption. In some cases, millions of dollars per hour were lost while the organization fought to restore its online services. The following changes have made DDoS attacks more prevalent and more virulent:

Increasingly sophisticated tactics
The notorious July 4 attacks, which were orchestrated by a custom bot, included SYN, PING, and GET floods.

1 Forrester Consulting, The Trends and Changing Landscape of DDoS Threats and Protection, July 2009

These attacks on U.S. and South Korean networks targeted more than 47 government and private institutions. Although the magnitude of the attacks was fairly low (averaging 39 megabytes per second),[2] more than 200,000 bots were employed, greatly amplifying the impact and reach of the attacks. Besides this type of direct flood attack, in which a high volume of spoofed packets is sent directly to the victim(s), attackers are increasingly using reflection flood attacks. In reflection flood attacks, attackers use recursive DNS servers to bounce attacks to their victims, and in the process amplify the attack and make it more difficult to track down the attack source.

Targeted vs. random victims

Whereas most attacks of the past were random, today's attacks often focus on a single organization or a small subset of businesses. In the Forrester survey mentioned above, nearly half of all attacks were targeted attacks. Worse, attackers are now taking down Web sites and applications with which they have no direct quarrel in order to inflict damage on a third-party target. For example, analysts believe that the 2009 DDoS attacks on Facebook and Twitter™ were designed to silence a single user of these social networking sites. Facebook and Twitter—and all their users—were simply collateral damage.

Surreptitious, application-level exploits

Instead of using brute-force volume to bring down an entire network, cyber criminals can execute subtle application-level attacks that mimic legitimate traffic. These attacks operate within an application's (or an application server's) normal thresholds of activity, making them difficult to detect with threshold-based detection tools. Until recently, the main application targets were known flaws in commonly used software and networking technologies. However, attacks on custom-built applications are on the rise.[3]

Novice-level tools

Even individuals with minimal technology skills can orchestrate DDoS attacks. Low-cost botnet rentals are advertised on the Internet, with one site offering botnets capable of launching DDoS attacks of 10–100 Gbps for as little as $200 per 24 hours. Would-be attackers can also unite with others to use "crowd-sourcing" tactics, such as those shared on Twitter during the recent Green Revolution in Iran. In this case, Twitter users posted links to attack tools (e.g., highvolume page reloads) with the goal of enlisting protestors' help in taking down pro-government Web sites. Although crowd-sourcing tactics require the continuous engagement of many people and are difficult to sustain, they illustrate the ease with which anyone can attempt an attack.

" …if you look across a longenough period, such as a year, it is highly likely that an organization, particularly one that has a substantial presence on the Internet, will experience at least one DDoS attack."[1]

1 Forrester Consulting, DDoS: A Threat You Can't Afford to Ignore, January 2009

2 Arbor Networks, Quick Stats Around the US – KR DDoS Attacks, July 10, 2009, http://asert.arbornetworks.com/2009/07/quick-stats-around-the-us-kr-ddos-attacks
3 SANS Institute, SANS Top 20 Internet Security Risks of 2007 Point to Two Major Transformations in Attacker Targets

Millions of packets per second (Mpps) Cyber criminals can harness the processing power and bandwidth of thousands of compromised computers to form "botnets" capable of sending millions of packets per second (Mpps) to disable even the largest networks. Attack magnitudes are more than 100 times greater than they were in 2001; one of the largest reported attacks of 2009 peaked at 69 gigabits per second (Gbps).[4]

DDoS MITIGATION CHALLENGES:
WHY TRADITIONAL TACTICS ARE NOT SUFFICIENT

While many organizations are increasingly concerned about the DDoS threat, few organizations have specific DDoS protection mechanisms in place. Those that do address DDoS often rely on approaches that lack the capacity and agility to mitigate attacks rapidly—and preferably before they reach the network.

Despite popular belief, the following measures, when implemented within most organizations, are insufficient to mitigate today's diverse, largescale attacks:

Over-provisioning of bandwidth
Although over-provisioning of bandwidth is one of the most common anti-DDoS measures, it is neither cost efficient nor highly effective for most organizations. It is not uncommon for organizations to spend an extra 75 percent for bandwidth beyond what they need to handle peak loads, and overprovisioning becomes useless as soon as an attack exceeds the amount of bandwidth that has been provisioned. In addition, over-provisioning only addresses network-level attacks, not application- or OS-level attacks. With attacks now capable of carrying more than one million packets per second (Mpps), even the most well-provisioned network can be overwhelmed.

Firewalls
Whereas firewall management used to be a sufficient strategy to manage denial of service (DoS) attacks, botnets and reflectors have since reduced the effectiveness of blocking attacks at the network edge. Using a firewall for mitigation may cause the CPU to spike and deplete memory resources. In addition, firewalls do not have anomaly detection capabilities.

Intrusion detection system (IDS)
An IDS device typically sits behind the firewall and links to a router in front of the firewall. Like an IPS (discussed in the next bullet), an IDS is designed and fine-tuned to inspect for single malicious packets. Neither IDS nor IPS devices are designed to handle high-volume attacks. Using them for DDoS mitigation can impact performance in their intended role of intrusion mitigation. In addition, by the time an IDS detects an anomaly and issues an alert, attack traffic is already consuming Internet bandwidth, potentially saturating the network, causing the CPU to spike, and depleting memory resources.

DDoS DEFINED
A denial-of-service (DoS) attack occurs when traffic is sent from one host to another computer with the intent of disrupting an online application or service. A distributed denial-of-service (DDoS) attack occurs when multiple hosts (such as compromised PCs) are leveraged to carry out and amplify an attack. Attackers usually create the denial-of-service condition by either consuming server bandwidth or by impairing the server itself. Typical targets include Web servers, DNS servers, application servers, routers, firewalls and Internet bandwidth.

4 Arbor® Networks, The Internet Goes to Way, Dec. 14, 2010

Intrusion prevention system (IPS)

An IPS has the capability to work as an anomaly detector; however, it can require a few weeks to understand "normal" traffic patterns and then organizations (or their IPS vendors) must spend several more days on manual tuning to specify which traffic is allowed and which should be alerted or blocked. For this reason, threat signature updates often occur too late to block a DDoS attack. In addition, many IPS devices rely on vendor-specific threat information, so they are not tuned and updated to address the full range of threats, which may include DDoS attack signatures. Finally, IPS devices are limited in the number of TCP sessions and amount of bandwidth that they can handle at a given moment. When overloaded, they shut down.

Routers

Routers cannot block spoofed IP sources (which are a leading source of DDoS packets) or manually trace back to thousands of IP addresses, rendering access control lists (ACLs) useless against DDoS attacks.

Black hole routing

Black hole routing an IP address or a range of IP addresses (i.e., intentionally causing packets coming from a specific IP address to be discarded rather than forwarded) can protect your resources from the ill effects of DDoS, but can also result in legitimate packets being discarded along with malicious attack traffic, effectively ensuring that the attack is successful in disrupting your operations.

Reliance on Internet service provider (ISP) mitigation

Many organizations assume that their ISP provides DDoS protection without inquiring specifically about service level agreements, attack reporting, bandwidth capabilities, black hole routing, and other important details of third-party DDoS mitigation.

BEST PRACTICES FOR DDoS MITIGATION

At the most basic level, successful DDoS mitigation involves knowing what to watch for, watching for those symptoms 24/7, having the technology capability and capacity to identify and deflect attacks while allowing legitimate traffic to reach its destination, and possessing the skills and experience to address issues appropriately in real time. The following best practices reflect these principles and draw on Verisign expertise gained through hands-on customer engagements, industry best practices, and successful defense of the global Verisign DNS infrastructure against numerous DDoS attacks:

- Centralize data gathering and understand trends

- Define a clear escalation path

- Use layered filtering

- Build in scalability and flexibility

- Address application and configuration issues

BEST PRACTICE 1:
CENTRALIZE DATA GATHERING AND UNDERSTAND TRENDS

At the most basic level, successful DDoS mitigation involves knowing what to watch for, watching for also allow Verisign to create synthetic events based on system logs that indicate malicious or inappropriate internal activity.

Centralize monitoring
Develop a centralized monitoring capability that allows you to see your entire network and traffic patterns all in one place; limit traffic oversight to a small team for consistency and continuity of oversight.

Understand normal network traffic patterns
To establish a baseline for normal traffic entering your organization, regularly collect sample packets and other pertinent information from switches, routers, and other devices. Know what types of traffic come in (e.g., SMTP, HTTP, and HTTPS), when (e.g., every Friday, early morning, the first of each month), from where, and how much. Establish a rolling 13-month (at least) view of what normal traffic looks like and incorporate this information into a correlation engine for threat detection, alerts and reporting.

Track worldwide, historical trends and threat intelligence
Conduct ongoing tracking and analysis of attack patterns around the world to identify and validate potential/emerging attacks more rapidly and to extract lessons learned into the appropriate incident response. Use existing intelligence to look for predefined deviations (i.e., analyze signatures) that signal a DDoS attack. Complement in-house intelligence-gathering with subscriptions to third-party intelligence service providers and participation in industry security groups and forums, where information-sharing can help reveal similarities in types of probes or unusual activity.

Implement DDoS-specific alerting, logging, and reporting systems
Make sure that alert notifications specifically alert you to signs of DDoS attacks, including attacks that are not necessarily volumebased. Implement a logging and correlation system to collect detailed attack data, which can be used to avert future attacks. Implement a clear process for generating and evaluating transaction, traffic summary, application, protocol, and event reports. Keep in mind that transaction reports are as important as traffic reports. A significant decrease in the expected number of transactions, for example, can be an even more significant indication of suspicious activity than increased traffic.

Work with experienced security researchers
The best monitoring, detection, alerting, logging, and reporting devices are useless unless organizations know what to do with the data. Security researchers should have hands-on expertise in distinguishing suspicious traffic from legitimate traffic, dealing with botnets, managing and defending against DDoS attacks (e.g., by infiltrating or taking down DDoS command-and-control servers), and changing mitigation tactics rapidly as circumstances dictate.

More than half of the respondents in a recent Forrester survey reported that their Internet service providers' (ISPs') services had been disrupted by DDoS attacks in the past 12 months.[1]

1 Forrester Consulting, DDoS: A Threat You Can't Afford to Ignore, January 2009

BEST PRACTICE 2: DEFINE A CLEAR ESCALATION PATH

Systematic processes and methodology are essential to effective DDoS attack mitigation.

Define standard operating procedures (SOPs) for incident response
Take into account internal infrastructure, services, and applications as well as the resources of customers and partners that may be impacted. If necessary, craft individual SOPs to address specific types of attacks or specific types of resources being attacked. Review SOPs on a regular basis and conduct periodic "fire drills" to make sure SOPs are up-to-date and functioning properly.

Define incident response teams
Do not wait for a 3:00 a.m. incident to decide who should be contacted. Prepare, publish, and regularly update an escalation contact list that includes information for the in-house team as well as relevant customers, vendors, partners, and upstream providers such as application service providers (ASPs). If you rely on an Internet service provider (ISP) for DDoS mitigation, recognize that unless you are a very large company, your service requests will likely go into a queue along with thousands of other customers' requests.

Address functional silos and areas of cross-function
Make sure that DDoS mitigation as it relates to business continuity is a universal goal. Identify functional silos and areas of overlapping ownership or responsibility. Break down silosbetween different groups (e.g., the network team and the information security team), clarify incident response roles and responsibilities, and enforce accountability.

Prepare for downtime
Understand which systems are vital to your business, and then develop and test contingency plans for short-term (e.g., 1 hour), mediumterm (e.g., 24 hours), and long-term (e.g., multiple-day) network or service outages.

BEST PRACTICE 3: USE LAYERED FILTERING

The goal of DDoS mitigation is to exclude only unwanted traffic while allowing legitimate traffic to enter the network with minimal delay. The most effective means to accomplish this is to use a multi-layered verification process that employs all the practices mentioned here.

Filter traffic in layers
Inspect incoming packets using signature analysis, dynamic profiling (based on monitoring and analysis of normal behavior), anti-spoofing algorithms and other technology to progressively filter harmful traffic upstream of the network.

Return legitimate traffic to the network with minimal latency
Ideally, legitimate traffic should continue to route through the network with little to no impact on end users, even during a large attack.

Apply filters at multiple levels of the OSI stack
Although some attacks can be mitigated by implementing filters at the network layer, complex attacks now require analysis and filtering up through the application layer.

Rate limit traffic, as needed
To prevent "low-tolerance" resources from being overwhelmed, have the capability to limit traffic rates according to the number of concurrent connections or bandwidth.

Be able to change and customize filters quickly
Have the capability to apply and remove standard filters (i.e., signatures) as needed, as well as generate custom filters in response to an attack or changes in your network.

Enhance rule sets over time
Analyze global intelligence as well as monitoring, alerting and reporting logs to identify attack vectors, and use this information to continually update rule sets.

BEST PRACTICE 4: BUILD IN SCALABILITY AND FLEXIBILITY
To make sure systems will function properly under attack conditions, organizations must have a highly scalable, flexible infrastructure.

On-demand capacity
Capacity includes bandwidth as well as the hardware processing power and scalability required to process the traffic load traveling over the bandwidth. Sufficient capacity is vital—yet very difficult to maintain and often impractical—within a single organization. Over-provisioning to absorb highmagnitude attacks, for example, requires significant expenditures for extra bandwidth (and additional servers) that may be needed rarely, if ever. In addition, there is no guarantee that the overprovisioned amount will suffice in today's environment, where DDoS attack magnitudes are increasing at an alarming rate and the typical organization's network connection to the Internet is one Gbps or less.[5]

Limit test every component and know your break points
Know how your infrastructure behaves under attack. Profile traffic scenarios and identify which components will not work under a heavy load. For example, know at which point a firewall or Web server fails, and know which packets or queries are harder on the system than others. Test various scenarios in a mirror production environment instead of merely forecasting, and retest every time you change any part of the infrastructure.

Load-balance the infrastructure
Once break points have been identified, load-balance the infrastructure to optimize traffic flow for normal and peak-load scenarios.

Consider the scalability of monitoring tools
Make sure that monitoring tools can continue to work during times of high load. In some bandwidthconsuming DDoS attacks, monitors slow down, cease to function or, worse yet, report bad data. For example, a monitor may keep reporting the same value because it cannot report anything higher.

Enforce hardware and software diversity
To protect against known vulnerabilities in any single vendor's DDoS mitigation applications, source tools from a variety of vendors.

Use a distributed model
If possible, use a distributed model to create and maintain redundancy for high-value applications and services.

5 Arbor Networks, Worldwide Infrastructure Security Report, Volume V, 2009

BEST PRACTICE 5: ADDRESS APPLICATION AND CONFIGURATION ISSUES

DDoS attacks have evolved from brute force attacks at the network layer to more sophisticated, difficultto-detect attacks at the application layer. Attackers can learn the acceptable threshold of activity for an individual application, and then sneak in as an unperceived increase in network traffic. In the overall context of the network, the increased traffic is not an issue, but if the targeted application has a low tolerance for high-volume traffic, the attack can take down the application.

Understand your applications
Know what each application does, how often it is used, what each application request looks like, and what the normal transaction levels are for each application-critical component. Determine the traffic threshold at which an application becomes flooded. If necessary, customize the traffic flow for individual applications.

Address simplistic configurations
This practice helps minimize resource-depletion attacks such as SYN, PUSH and ACK flood attacks. (For more information about these types of attacks, see the Verisign white paper, Distributed Denial of Service (DDoS) Attacks: Latest Motivations and Methods, available at www.verisigninc.com.

Address common application vulnerabilities
According to a recent report based on data from the SANS Institute,[6] the three most DDoS prone application vulnerabilities are improper input validation, buffer overflow, and incorrect calculation. All three of these vulnerabilities are prevalent in organizations, and attackers exploit them frequently. The good news is that the cost of remediating them is relatively low.

Be a good neighbor
Make sure that non-critical applications and systems cannot be exploited to attack other sites.

UNIQUE ADVANTAGES OF MANAGED DDoS MITIGATION SERVICES

At the most basic level, successful DDoS mitigation involves knowing what to watch for, watching for those symptoms 24/7, having the technology capability and capacity to identify and deflect attacks while allowing legitimate traffic to reach its destination, and possessing the skills and experience to address issues appropriately in real time. The following best practices reflect these principles and draw on Verisign expertise gained through hands-on customer engagements, industry best practices, and successful defense of the global Verisign DNS infrastructure against numerous DDoS attacks. Besides the significant advantages of reduced cost and complexity, managed DDoS mitigation services provide benefits that an in-house DDoS mitigation solution cannot:

Upstream location
Managed services operate "in-the-cloud," so packets destined for the organization travel through an Internet "scrubbing" center first. Malicious traffic is diverted, and then cleansed traffic is delivered to the network edge. Potential DDoS attacks are mitigated before they can ever reach the organization's network, protecting availability and performance and obviating the need to over-provision bandwidth as a DDoS mitigation tactic.

On-premise DDoS detection and protection technologies can defend against small attacks, but will be increasingly ineffective as DDoS continues to grow in size.[1]

1 Forrester Consulting, DDoS: A Threat You Can't Afford to Ignore, January 2009

6 MITRE Corporation and SANS Institute, 2009 CWE/SANS Top 25 Most Dangerous Programming Errors; www.cwe.mitre.org/top25

Core Internet connectivity

Because managed services have core Internet connectivity, they have an inherently higher capacity for traffic and can handle larger attacks than any single organization can handle. In addition, they can use core-routing techniques (e.g., Border Gateway Patrol [BGP]) to more efficiently divert malicious traffic.

Massive bandwidth

Managed services providers can afford to over-provision bandwidth and invest heavily in scalable infrastructure, allowing them to absorb larger attacks than most organizations. In addition, the best providers have multiple network operations centers, distributed globally, to provide redundancy and high availability.

24/7 expertise and staffing

In a recent Pricewaterhouse Coopers survey,[7] recovery time from a DDoS attack was proportional to the percentage of security staff in IT. Because DDoS mitigation is their primary focus, managed services providers can afford to invest in highly skilled, around-the-clock personnel who make it their business to know everything there is to know about DDoS detection and response.

Greater visibility

A standalone organization can rarely match the field of view that managed services providers gain by working with multiple carriers, clients, networks and peers worldwide. This wider view of Internet traffic helps providers accurately distinguish between normal and malicious traffic and more quickly recognize sources of malicious activity.

Carrier/ISP neutrality

Many organizations use multiple ISPs. A carrier-agnostic managed services provider can provide a single solution that encompasses all the ISPs being used by the organization. In addition, carrier neutrality increases the services provider's immunity to carrierspecific attacks.

SELECTION CHECKLIST FOR DDoS MITIGATION PROVIDERS

When evaluating a managed services provider for DDoS mitigation, consider the following differentiators:

Expertise in DDoS detection and mitigation

Does the managed services provider have a dedicated staff of DDoS mitigation experts who proactively assist customers in developing a long-term strategy? Does it employ a defined, systematic approach to organizing and managing traffic reports and handling DDoS attacks? Is DDoS mitigation a core expertise.

Capacity and scale

How large an attack can the provider absorb? What kinds of processing speed, memory speed, storage access, and latency can you expect under normal conditions and under attack conditions?

7 United Kingdom Department for Business Enterprise & Regulatory Reform (BERR), 2008Information Security
Breaches Survey, conducted by PricewaterhouseCoopers

Attack management
How does the provider manage attacks? Does it proactively detect attack packets or does it wait for your organization to report downtime or other signs of attack? How much legitimate traffic is blocked during an attack, and for how long? (Beware of black hole routing. This approach is used by many thirdparty vendors, and denies resources to legitimate users, thereby accomplishing the attacker's goals.)

Service transparency
Does the provider give you sufficient visibility into the traffic monitoring and mitigation process to understand circumstances during a DDoS attack as well as when no attack is occurring?

ISP neutrality
If you use multiple ISPs, can the managed services provider support DDoS mitigation across multiple carriers? Can it protect multiple points of intersection with the Internet?

Filtering capabilities
Can the provider progressively filter traffic during security incidents, and if so, what filtering capabilities can it deploy? Can it filter for application-, session- and OS-level attacks? How long will the provider leave filters in place? What is the SLA to deploy filters?

Service level agreements (SLAs)
How soon will the managed services provider notify you of an anomaly or DDoS attack? Do you have a dedicated service representative or rapid response times, or will your service request go into a general queue? Can the service levels scale to accommodate rapid growth, mergers and reorganizations?

Service availability
What is the SLA throughput when your network is under attack (e.g.,variable, burst rates or best effort)? What is the availability (e.g., 99 percent vs. 99.99 percent)? Does the provider use load sharing across multiple customers, or does it support dedicated systems?

Reporting
What kinds of reports does the provider generate after an event or attack occurs? How long does the provider retain logs and reports related to your organization?

Online security
Does the managed services provider have an established security infrastructure to protect customer data? Does it provide a secure, Web-based portal that includes an audit trail, data encryption (SSL) and password protection?

Physical security
Does the provider employ military-grade security for facilities that contain mission-critical systems and data? Does it have redundant systems for power, HVAC, and network services? Does it have redundant data stores?

Ease of use
Does the managed services provider offer extensions customizable interfaces, and reporting tools? Does the provider offer a Web portal to provide alerts, visibility into real-time network traffic and bandwidth utilization and the ability to detect anomalies?

ABOUT THE VERISIGN DDoS PROTECTION SERVICE

The Verisign® DDoS Protection Service is an in-the-cloud solution that provides scalable, cost-effective protection against Distributed Denial of Service (DDoS) attacks. The Verisign DDoS Protection Service combines people, processes, and technology to offer a massively fortified, comprehensive service that can effectively and efficiently mitigate the world's largest DDoS attacks. The service addresses the facets of DDoS mitigation—from assessment, monitoring, detection, and reporting to DDoS source analysis.

The Verisign DDoS Protection Service helps protect organizations from catastrophic DDoS attacks by detecting and filtering malicious traffic upstream of the organization's network. An international team of security experts staffs the globally distributed network operation centers and is available 24/7 to monitor, detect, analyze, and respond to malicious traffic. Participation in forums with Tier 1, 2, and 3 carriers; global intelligence and analysis from Verisign® iDefense® Security Intelligence Service; and interaction with hundreds of customers and partners give Verisign analysts broad visibility into Internet traffic.

This unique insight—along with highly refined layered-filtering technology that works at the network, application, and session layers of the OSI stack—enables the Verisign DDoS Protection Service to quickly and accurately distinguish between malicious and benign traffic, and to forward legitimate traffic to the organization. The military-grade network operations centers are provisioned to absorb the largest DDoS attacks, while proven procedures enable Verisign staff to react quickly and with agility to defend against the rapidly changing threat landscape.

LEARN MORE

For more information, please email Learnmore@verisign.com.

ABOUT VERISIGN

Verisign is the trusted provider of Internet infrastructure services for the digital world. Billions of times each day, companies and consumers rely on our Internet infrastructure to communicate and conduct commerce with confidence.

VerisignInc.com