Stratecast A Division of FROST & SULLIVAN

DDoS Mitigation to the Rescue

Business Communication Services (BCS)
Volume 4, Number 4
March 2010

DDoS MITIGATION TO THE RESCUE

INTRODUCTION[1]

Cyber threats that have garnered the most public attention have been those that attempt to infiltrate and infect the networks and systems of businesses and governments. Yet, this is only a piece of the ever evolving threat landscape. Denial of Service (DoS) and Distributed Denial of Service (DDoS) attacks are threat vectors that have existed for years, but with a different objective. Rather than infiltrate and infect, the objective of the perpetrators of these attacks is to cripple or demonstrate the ability to cripple a public-facing Web operation (e.g., eCommerce, online recreation and entertainment, and public safety) by overloading the website delivery infrastructure. Ultimately, the perpetrators' objective is more insidious, such as, financial gain via extortion or industrial warfare, or to make a political statement by hampering the victim organization's ability to promote its agenda online.

With Web properties being so engrained in how business is conducted, positions articulated, and citizens served, complacency in defending against DoS and DDoS attacks is an unacceptable strategy for many organizations. Their on-going Web presence and the

[1] In preparing this report, Stratecast conducted interviews with representatives of the following companies:

- Arbor Networks – Mat Mathews, VP of Product Management
- Cisco – Bill McGee and Fred Kost, Security Marketing
- F5 – Ido Breger – Product Manager, Application Security
- IntruGuard Devices – Hemant Jain, CTO
- Prolexic Technologies – Paul Sop, CTO; and Greg Burns, VP of Marketing
- RioRey – Kwok Li, CEO; and Bill Wilson, COO
- TippingPoint – James Collinge, Director, Product Line Management
- Top Layer – Mike Paquette, VP and Chief Strategy Officer
- VeriSign - Jason Malo, VIDN (VeriSign Internal Defense Network) Product Manager; and Ben Desjardins, Marketing Manager

Please note that the insights and opinions expressed in this assessment are those of Stratecast and have been developed through the Stratecast research and analysis process. These expressed insights and opinions do not necessarily reflect the views of the company executives interviewed.

critical functions their Web properties support cannot afford to be disrupted; the cost is simply too high. For these organizations, being proactive in blunting the effects of DoS and DDoS attacks is not an option, it is an absolute.

Fortunately there is a growing array of mitigation solutions. The purpose of this bulletin is to describe the types of mitigation solutions that are in the market today and their differences. Our focus will be on DDoS mitigation solutions as DDoS attacks are more difficult to defeat due to their nature of being launched from multiple sources (i.e., IP addresses) rather than a single source with DoS. In light of Cloud Computing being an online, always available environment, we will also examine the unique implications that DDoS attacks place on the business model of Cloud Service Providers.

PROTECT WHAT'S IMPORTANT

An important consideration regarding DDoS attacks and migration solutions are the website owner's motivations in investing in a DDoS mitigation solution and how the solution's attributes align with those motivations. Based on our conversations with several existing and upcoming providers of DDoS mitigation solutions,[2] we concluded that there are three primary motivations for investing in a DDoS solution.

1. Service availability
2. Resource cost optimization
3. Stability in security operations

Service Availability

First and foremost, the motivation to invest in a DDoS mitigation solution is driven by the need to maintain uptime in a Web property. As a critical extension of a non-online operation or the primary operation of a business or governmental entity, blocking legitimate use of the website is detrimental, possibly severely detrimental. Operators of online gaming and gambling sites, for example, lose the ability to serve their constituents and earn profits if access by legitimate users is crowded out by a DDoS attack. Therefore, and logically, investments in a DDoS mitigation solution are made to combat the negative consequences of a DDoS attack. A correlated motivation that ties into the capabilities of the DDoS mitigation solution is the avoidance of false positives, that is, blocking website traffic that is incorrectly identified as part of the DDoS attack. The less tolerant a website owner is of false positives, the more the website owner will appreciate the effectiveness of the mitigation solution that minimizes the occurrences of false positives. Not to be overlooked, an alternative to a DDoS mitigation solution is the investment in sufficient website delivery infrastructure to not only serve legitimate users but also accommodate the usage demands of illegitimate users (attackers). This too has consequences as explained in the next motivation.

[2] Suppliers of existing DDoS mitigation solutions that we spoke with are listed on the first page. To protect the confidentiality and strategic direction of solutions being contemplated for future availability by the other providers, the names of those providers are not listed. Similarly, upcoming shifts and enhancements in the solutions of the listed suppliers will not be identified by supplier in this bulletin.

Resource Cost Optimization

Over-provisioning of the website's delivery infrastructure – the Web servers, routers and switches, load balancers, and Internet connectivity – is part of being on the Web. Precise predictability in the volume and characteristics of website traffic does not exist for most if not all public-facing websites. Furthermore, incurring the cost of surplus infrastructure to address anticipated traffic growth or surges, an insurance policy of sorts, is considered a better approach than adding infrastructure capacity once a threshold has been reached.[3] Similarly, maintaining surplus infrastructure capacity can be an approach to combat the crowding out of legitimate traffic by DDoS attack traffic. From a tactical perspective, the website owner concludes it is better to serve some bad versus the alternative of not serving all the good. The reasonableness of this tactic, however, diminishes as the volume and computational demand of DDoS attacks increase. In other words, a point is reach in which the cost of surplus capacity in website delivery infrastructure to bend but not break when faced with DDoS attacks exceeds the cost of a DDoS mitigation solution.

Stability in Security State

Security personnel already have a significant work effort just to maintain a good fight against the routine flow of threats; enforcing current security policies; and managing existing security platforms such as firewalls, Intrusion Prevention Systems (IPS), email and Web content filters, and desktop security software. With a lack of new and qualified security staff, the addition of spontaneous and continuously evolving DDoS attacks can be detrimental to maintaining a consistent level of security state in other areas. Further, a crafty DDoS attacker may use a DDoS attack as a decoy hoping that in the redirection of personnel to address the DDoS attack their guard will temporarily be let down in other areas, providing the attacker with an opening to infiltrate and infect the victim's network and systems. The quest to maintain stability in security state is the third motivation for investing in a DDoS mitigation solution.

DDOS MITIGATION SOLUTIONS

At the risk of oversimplifying a complex security solution, we compartmentalized DDoS mitigation solutions into just two functional areas, detection and mitigation, with further compartmentalization based on the location where each function is performed. Our definition of detection is the collection of data or clues which signal that a potential DDoS attack is underway or imminent. Mitigation is the process of examining the Web traffic to confirm the existence of the DDoS attack, assess its characteristics, calculate the implications, and follow through with a course of action to mitigate the effects of the DDoS attack. Mitigation efforts are unlikely to be uniform and will vary based on several factors including: the criticality of website availability; the severity of the attack; the level

[3] Incidentally, serving surges and peaks in website traffic through a shared infrastructure is one of the original value propositions of a content delivery network; cache content at the edge of the CDN's network nearer to the requestor (i.e., upstream) rather than tax the website delivery infrastructure (i.e., downstream).

of confidence in attack assessment accuracy and precision in mitigation procedures; and the threshold of DDoS attack traffic which the website infrastructure can accept before that attack becomes impactful on legitimate website usage.

Detection

There are essentially three approaches to collecting data that would tip off that a DDoS attack is underway.

Inline Appliances

Generally deployed near the network firewall and in the direct flow of network traffic, inline appliances have the beneficial properties of viewing all inbound network traffic and, if equipped, conducting deep packet inspection to gain a granular traffic perspective. Similar in concept to Intrusion Prevention Systems and Intrusion Detection Systems (IPS/IDS), a DDoS inline appliance needs to mimic the attributes of IPS/IDS of low latency processing, high throughput capacity, and rock solid reliability. Considering these parallel attributes, several IPS/IDS vendors position their solutions as having DDoS mitigation capabilities (e.g., TippingPoint); and conversely, vendors that originally designed a singularly focused DDoS mitigation appliance have advanced their products to support IPS/IDS (e.g., Top Layer). Separately, there are specialists in DDoS inline appliances, such as IntruGuard and RioRey. Both of these vendors highlight their capabilities in behavioral analysis to detect a wide range of DDoS attacks in real time.

Additionally, and by being deployed at the website owner's network edge or further downstream in front of the server farm, inline appliances have the strategic placement to observe traffic bi-directionally while also maintaining state. These are valuable attributes for gaining the clarity needed to detect Layer 7/application layer and transaction-based DDoS attacks. Noting this, providers of Web Application Firewalls (WAFs), such as F5, have developed mechanisms to automatically flag certain types of DDoS attacks within the F5 BIG-IP Local Traffic Manager and Application Security Module.

As with IPS/IDS and other inline devices, DDoS inline appliances can suffer from a similar challenge of maintaining low latency processing as the number of DDoS attack signatures and profiles increase. Also, throughput capacity or scalability can be an issue as the nature of DDoS flood attacks is to overwhelm the delivery infrastructure with traffic volume. With DDoS flood attacks measured in the multi-gigabits per second level, inline appliances can themselves be a traffic bottleneck and, in a perverse way, be an involuntary contributor to the attacker's objective of slowing or disrupting the normal flow of user-to-website interactions. To combat this risk with inline appliances, website owners require and would incur the premium cost associated with high-end appliances that can accommodate high traffic volume DDoS attacks. With the assumption that high traffic volume DDoS attacks will continue to ramp upward, the website owner is potentially caught in a continuous cycle of upgrading its inline appliances. Alternatively, the website owner can hope that the highest volume DDoS attacks do not get aimed at

its website and, if they do, their duration is short. Both are scenarios that do not have any guarantees.

Passive Out-of-Band Listening Appliances

Another approach to examining traffic for the purpose of detecting DDoS attacks is through passive listening appliances. In comparison to inline appliances, passive listening devices can have similar granular detection capabilities but without the risk of being a potential traffic bottleneck. Conversely, these devices do not support the dual functions of inline appliances to detect and mitigate DDoS attacks. Consequently, redirection of website traffic to a scrubbing platform for mitigation is required.

Connect into the NetFlow Stream

A popular method of DDoS attack detection with traffic redirection and scrub solutions is to leverage the IP traffic information communicated through the NetFlow protocol supported in routers and switches. Like passive listening devices, NetFlow does not interfere with network routing and switching and, by itself, Netflow does not constitute a potential traffic bottleneck. Also similar to passive listening devices, the routers and switches from which NetFlow is generated do not support DDoS attack mitigation; that function is left to a separate scrubbing platform. The principal difference between NetFlow and passive listening devices is in traffic granularity. The granularity of NetFlow is limited to detecting anomalies in the flow of network traffic. While beneficial, detecting the subtleties of application layer DDoS attacks can be missed.

Mitigation

Like detection, there is no single location where DDoS attack mitigation can occur. And like detection, location has a bearing on the attributes of the DDoS solution. Prominent in the market today, we count three locations where DDoS mitigation occurs. We are hopeful that a fourth location, nearer the attack origins, will become more prominent in the future.

Inline Appliances

As previous indicated, one of the positive properties of DDoS inline appliances is that they do double duty, detection and mitigation. For some organizations, the retention of full control, logical and physical, of their DDoS mitigation solution is also a positive attribute.

Bearing the full cost of appliance management and maintenance, and the need to purchase appliances that excel in low latency processing and high volume scalability are drawbacks. Furthermore, the risk and cost of technology obsolescence are also present. As the nature of DDoS attacks evolve, there are no guarantees that the appliance investment that was once considered best of breed and could address a meaningful range of attack types will retain these distinctions in the future. Subsequently, investing in dedicated DDoS appliances from multiple or replacement vendors may be a required course but one that adds cost and complexity for the organization.

In addition, inline appliances do not alleviate the potential bottleneck in Internet connectivity. Even with a highly scalable inline DDoS appliance and equally capable website backend, if access bandwidth is overwhelmed with legitimate and DDoS traffic, the attacker has a win under his/her belt.

ISP Scrubbing Platforms

To address the cost of single ownership and the risk of an access bandwidth bottleneck, several ISP's offer a managed DDoS mitigation service in which inbound website traffic, once detected as potentially containing DDoS traffic, is redirected to a scrubbing platform located within the ISP's network. As a shared platform serving multiple subscribers, the cost of the hardware platform and maintenance is spread over many subscribers – an often cited argument for network or cloud-based services. In addition, the access bandwidth robbing effect of DDoS traffic is muted if that traffic is blocked within the ISP's network.

Several years ago, ISPs started to offer managed DDoS mitigation services that followed this model. Cisco, with its Cisco Guard product line and Arbor Networks, with its Peakflow product line were frequently identified as vendors behind this ISP managed service. With Cisco announcing that Cisco Guard will be entering an end-of-life progression and backing ISP client transitions from Cisco Guard to Arbor Peakflow, Arbor Networks is in a favorable position to deepen its ISP relationships. The same is true for Cisco Guard's enterprise deployments.

We were not able to confirm whether ISPs are following a multi-vendor approach for scrubbing. However, we do believe that this would be a favorable service attribute, but not necessarily an easy one. The ISP would need to identify DDoS mitigation vendors that have a material degree of complementariness, one excelling in certain types of DDoS attacks and another a different set, and then build a service layer that seamlessly integrates multiple platforms. This latter feature is important not only for the provider's operational effectiveness and efficiency with its managed DDoS mitigation service but also in providing subscribers with a comprehensive sense of control and transparency. Given a primary objective to blunt DDoS attacks while minimizing collateral damage to legitimate traffic, we believe that most subscribers will want deep collaboration with the service provider on mitigation decisions. Having a common dashboard for provider and subscriber to share we believe is essential in meeting this objective.

Cloud-based Scrubbing Centers

Cloud-based Scrubbing Centers operate in a similar fashion to ISP scrubbing platforms. Once an attack is detected, inbound website traffic is redirected to a scrubbing platform for mitigation. Rather than redirect to a scrubbing center located in the ISP's network, website traffic is redirected to an Internet-based scrubbing center unassociated with the website-serving ISP. In this manner, the providers of cloud-based scrubbing centers are marketing to a larger market—any website owner—rather than only those website owners that also procure their Internet connectivity from the ISP. From our perspective, the providers of cloud-based scrubbing centers do, however, face a steeper challenge in gaining customers, as the ISP has the beneficial position of an existing customer

relationship from which to offer a complementary service, possibly in a packaged engagement.

This challenge notwithstanding, providers of cloud-based scrubbing centers are thriving, particularly with online business owners that have the least tolerance for lapses in website availability. One such company is Prolexic. Several attributes of the Prolexic service contribute to the company's success and differentiated position versus ISP scrubbing platforms and premise-based appliances. First, the company equipped its regional cloud-based scrubbing centers with DDoS mitigation appliances from multiple vendors, thus allowing the company to surgically leverage the unique DDoS mitigation capabilities of each vendor's solution. Second, the company supplements these appliances with its own proprietary DDoS mitigation capabilities to create a superset of DDoS fighting tools. Third, the company structured its service to serve a variety of website owners' consumption preferences, from on-demand usage-based pricing to a fixed price for continuous service. Fourth, the company fosters provider-client collaboration and service transparency through 24 x 7 client access to Prolexic service technicians. Fifth, the company backs its service with Service Level Agreements (SLAs). While no provider of DDoS mitigation service or vendor of DDoS mitigation appliances can provide retribution to its clients for the impact of a DDoS attack to its business (e.g., brand impairment or lost revenue-generation transactions), Prolexic's SLAs, covering time-to-mitigate and remedy, highlight the company's commitment to minimizing the business impacts of DDoS attacks.

Prolexic is not the only provider of this brand of cloud-based DDoS mitigation service. Late last year, VeriSign announced the availability of its Internet Defense Network. Originally built to protect the company's Domain Name Services, the company took the next step and created a commercial DDoS monitoring and mitigation service. We view this entrance by a highly recognizable brand in trusted Internet communication and commerce as testament to the validity of a cloud-based model for fighting DDoS attacks. We are also anxious to learn where else VeriSign may extend its cloud-based service delivery model. DDoS monitoring and mitigation, we believe, is only the start of additional cloud-based security offerings by the company.

Upstream Mitigation

From our perspective, mitigation at the website owner's edge—within the ISP network near egress into website owner's datacenter, or at an Internet-connected scrubbing center—all face a common challenge of scalability. Can each scale to combat the largest of DDoS attacks now and in the future? Only time will tell. Nevertheless, what is predictable is that built-in, ready-now scalability, regardless of approach or approaches chosen (including adding more capacity in the website delivery infrastructure), adds to the cost of DDoS mitigation or, stated from the perspective of a business objective, the cost of maintaining a highly reliable and available Web presence. Certainly technological and operational improvements made by vendors and providers will lessen an escalation in DDoS mitigation costs, but not totally.

Considering this structural cost element of current DDoS mitigation solutions, Stratecast speculates on whether an alternative or perhaps supplemental approach exists that can tame

or even reverse this cost spiral. Rather than battle attackers at one or more downstream traffic concentration points, might taking the battle upstream, closer to the attack origins, be an improvement? This concept is already being pursued by Arbor Networks through its ATLAS traffic and routing sharing consortium (100+ ISP strong) and its Fingerprint Sharing Alliance. With these initiatives, Arbor Networks operates as an information clearinghouse for participating ISPs. By leveraging near real-time DDoS-related information among participating ISPs, ISPs are better equipped to block certain types of DDoS attack traffic before reaching ISP traffic exchange points and, as a positive outcome, reduce the overall cost incurred among ISPs in transporting traffic that is, in theory, unwanted by their website customers.

However, we do not have a measurement of the extent that DDoS attack blocking occurs before ISP traffic exchanges and the cost savings that this blocking produces. In fact, pre-exchange blocking may not be the preferred use of information shared among ISPs, as the incentives may be weak relative to the alternatives and the presence of collateral risk. For example, if DDoS attack traffic is a relatively small portion of an ISP's total traffic, building out network capacity may be a more economically rational approach than being hopeful that partnering ISPs will only send traffic that is partially filtered of DDoS traffic. Also, the ISP must consider the potential risk of false positives—blocking legitimate traffic that is identified as DDoS attack traffic. An embarrassing conversation could arise between ISP and a client if the ISP's inter-ISP DDoS traffic management arrangements result in the blocking of legitimate traffic without the client's approval. However, bringing in the client to approve blocking of suspected DDoS attack on another ISP's network may prove to be more complex and less timely approach than other DDoS mitigation approaches. It could be that the more common but still beneficial use of the information shared among ISPs through Arbor Networks' ATLAS and Fingerprint Sharing Alliance is in serving managed DDoS mitigation subscribers better. At least then the ISP offering the managed service is getting direct payment by its clients for service delivered.

Our conclusion is that in order for meaningful upstream blocking of DDoS attack traffic to occur, there must be reliable economic incentives for both the website owner and the ISP. Additionally, the website owner must be a participant in the decision process to block traffic. This scenario has not yet materialized in the market.

MULTIPLE OPTIONS GOOD; MORE IS BETTER

The beneficial aspect of DDoS mitigation is that several options exist, all with proven track records. Each solution type we outlined has its pros and cons. For website owners, they can mix and match options and, through that, adhere to the security best practices concept of multiple layers of defense. A mixed approach may also be a more economical approach, for example, blocking flood type attacks before they reach the website owner's Internet access connection and thwarting application layer attacks in the owner's environment. Another perspective on this mixture is balancing containment and control—contain the attacks outside website owner's environment where individual

website owner control is less of a requirement and contain the attacks in the website owner's environment where more individual website owner control and intimate knowledge is required.

Clearly website owners' need for DDoS solutions will vary. The attention and spending that website owners will devote to DDoS detection and mitigation will be correlated to their aversion to disruptions in website availability by their legitimate website users.

What should not be lost in this review of DDoS mitigation solutions is that DDoS mitigation, like most other categories of security, is an information intensive operation and one where timeliness and accuracy in mitigation decisions is critical. But relying solely on the experience and knowledge of a single organization will limit any organization's ability to meet that quick and accurate objective. Therefore, essential solution elements that a website owner should consider in choosing a DDoS solution is the level of technical and professional support they are entitled to receive from their vendor or service provider, and the richness of supporting information related to combating DDoS attacks they have at their disposal. Tools to detect and mitigate DDoS attacks are essential but it is the effective application of those tools that will matter the most.

THE CLOUD DILEMMA

The risk of a DDoS attack creates a dilemma for cloud service providers as the "disrupt" objective of DDoS perpetrators is a direct assault on the "always available" attribute of public cloud-based environments. This potential of placing the value proposition of cloud providers at risk makes cloud providers attractive victims.

Furthermore, a DDoS attack does not need to be aimed at the cloud provider to be effective, it can be aimed at just a single cloud tenant. Given the shared infrastructure architecture of a cloud environment, attacking a single tenant can cause collateral damage for other tenants. The cloud provider could react by shutting down the tenant being attacked in order to protect its other tenants, a practice that is not uncommon in shared hosting environments, but this practice, again, represents a contradiction to the value proposition and objective of cloud environments to replicate the core attributes of private data centers, such as high availability.

We believe that the risk of DDoS attacks and the implications for cloud providers could have a bearing on the structure of cloud-based services. Following are two potential scenarios.

Managed DDoS Mitigation Service

Considering that once a business operation is placed on the Web it becomes vulnerable to DDoS attacks, regardless of its hosting location—within the business premise (private datacenter), with a hosting provider, or in a cloud environment—a cloud provider can turn this risk of DDoS attacks into a business opportunity by offering a managed DDoS mitigation service to its tenants. In keeping with the on-demand pricing structure of

other services offered from the cloud, this managed DDoS mitigation service could also be structured and priced as an on-demand service. Furthermore, if the cloud provider can build into its tenant contract agreement that access to its Web properties hosted in the provider's cloud will be turned off if a DDoS attack reaches a specified threshold, this contract language could become a strong incentive for tenants to subscribe to this optional security service. Additionally, this incentive has the potential to gain strength after a DDoS attack occurs and service has been suspended, as the tenant will then have a more tangible understanding of the true cost of a DDoS attack to its business.

A less binary approach for a managed DDoS mitigation service in a public cloud environment is a tiered approach. In a possible tiered approach, the service provider automatically blocks DDoS attack traffic where there is little doubt of the DDoS traffic authenticity. This automatic blocking would be a standard service element for cloud tenants. Further granularity in blocking would be accomplished with the cloud provider and tenant-DDoS victim working collaboratively to reduce the volume of DDoS traffic that is mixed in with legitimate traffic. The objective with the provider-tenant collaboration is to reduce DDoS traffic while minimizing false positives. As an incentive for the tenant to collaborate actively with the provider is to keep the cloud usage meter on such that the tenant pays for the processing of legitimate and DDoS traffic.

Privileged Public Cloud

Not all business operations considered for placement in cloud environments exist for the purpose of interacting with the public. Many business operations are internal-facing such that only privileged or credentialed users are granted access. Because user communities differ, a potential service direction for cloud service providers is to establish two types of public cloud environments, one for public-facing applications and another for privileged user applications. The value in creating dual clouds is analogous to network service providers' public and private IP networks. Private IP networks were developed for the community of communication service subscribers that want to take advantage of the benefits of IP convergence but without the security risks inherent in the Internet.

While this dual cloud approach may not completely eliminate the risk of a DDoS attack, as other privileged cloud tenants could be part of a DDoS attack, we believe that the risk or exposure is greatly reduced. In addition, the ability for the tenant and provider to accurately detect a DDoS attack, determine the attack origins, and block attack traffic will be more straightforward than accomplishing the same for DDoS attacks aimed at public-facing applications.

Stratecast

The Last Word

Whether a business should be concerned about Distributed Denial of Service (DDoS) attacks is a matter of two factors: (1) the implications to its business if an attack occurs and is successful, and (2) the likelihood that an attack would be launched. These factors are positively correlated factors as the more valuable a Web property is, the more attractive it is to be attacked as an attacker's reward has the potential to be higher than with less critical/valuable Web properties. Considering the directional change in cyber perpetrators' motivations from fame to fortune and the ongoing migration of external and internal business operations to be Web-based, it would be difficult to argue that the likelihood of being attacked is not on the rise.

Additionally, the means to launch a DDoS attack is likely to rise. For example, a prominent means to launch a DDoS attack is through botnets. The number and size of botnets are related to the number of Internet-connected, malware vulnerable devices. There is no doubt that the number of Internet-connected devices will be increasing; the only question is by how much. While projections vary, a conservative projection solely in the increase in mobile devices is one billion worldwide over the next three years. The popularity of smartphones, netbooks and, more recently, e-readers by consumers is supporting evidence that a substantial increase in the number of Internet-connected mobile devices will materialize. With these devices' always or frequently-connected mode, they are vulnerable to infection and being herded into a botnet army. Furthermore, the general tendency in cyber security for device users not to invest in protection unless they have a personal need, and the clandestine manner in which botnet -affected devices operate, without the awareness of the device users, also contributes to the overall likelihood that DDoS attacks will increase in number.

Considering these points, a Web-facing business is proverbially faced with a "pay me now or pay me later" decision. Taking steps to fight a DDoS attack before one occurs—pay me now—we view is the best and prudent choice.

Michael Suby

Director
Stratecast (a Division of Frost & Sullivan)
msuby@stratecast.com

CONTACT US

Silicon Valley
331 E. Evelyn Ave.
Suite 100 Mountain View, CA 94041
Tel 650.475.4500
Fax 650.475.1570

San Antonio
7550 West Interstate 10, Suite 400,
San Antonio, Texas 78229-5616
Tel 210.348.1000
Fax 210.348.1003

London
4, Grosvenor Gardens,
London SWIW ODH,UK
Tel 44(0)20 7730 3438
Fax 44(0)20 7730 3343

877.GoFrost
myfrost@frost.com
http://www.frost.com

Beijing
Bengaluru
Bogotá
Buenos Aires
Cape Town
Chennai
Delhi
Dubai
Frankfurt
Kolkata
Kuala Lumpur
London
Manhattan
Melbourne
Mexico City
Milan
Mumbai
Oxford
Palo Alto
Paris
Rockville Centre
San Antonio
São Paulo
Seoul
Shanghai
Singapore
Sydney
Tel Aviv
Tokyo
Toronto
Warsaw

ABOUT STRATECAST

Stratecast assists clients in achieving their strategic and growth objectives by providing critical, objective and accurate strategic insight on the global communications industry. As a division of Frost & Sullivan, Stratecast's strategic consulting and analysis services complement Frost & Sullivan's Market Engineering and Growth Partnership services. Stratecast's product line includes subscription-based recurring analysis programs focused on Business Communication Services (BCS), Consumer Communication Services (CCS), Communications Infrastructure and Convergence (CIC), OSS and BSS Global Competitive Strategies (OSSCS), and our weekly opinion editorial, Stratecast Perspectives and Insights for Executives (SPIE). Stratecast also produces research modules focused on a single research theme or technology area such as IMS and Service Delivery Platforms (IMS&SDP), Managed and Professional Services (M&PS), Mobility and Wireless (M&W), Multi-Channel Video Programming Distribution (MVPD), Network Infrastructure and OSS (NI&O), Secure Networking (SN) and Unified Communications (UC). Custom consulting engagements are available. Contact your Stratecast Account Executive for advice on the best collection of services for your growth needs.

ABOUT FROST & SULLIVAN

Frost & Sullivan, the Growth Partnership Company, partners with clients to accelerate their growth. The company's TEAM Research, Growth Consulting, and Growth Team Membership™ empower clients to create a growth-focused culture that generates, evaluates, and implements effective growth strategies. Frost & Sullivan employs over 45 years of experience in partnering with Global 1000 companies, emerging businesses, and the investment community from more than 30 offices on six continents. For more information about Frost & Sullivan's Growth Partnership Services, visit http://www.frost.com.