

# Forrester Consulting

HELPING BUSINESS THRIVE ON TECHNOLOGY CHANGE

January 21, 2009

## DDoS: A Threat You Can't Afford To Ignore

A commissioned study conducted by Forrester Consulting on behalf of VeriSign

FORRESTER®

FORRESTER®

### Headquarters

Forrester Research, Inc., 400 Technology Square, Cambridge, MA 02139 USA  
Tel: +1 617/613-6000 • Fax: +1 617/613-5000 • [www.forrester.com](http://www.forrester.com)

## Table Of Contents

Executive Summary .....	4
Introduction .....	5
Survey Methodology.....	6
Businesses Stand To Lose Big With DDoS.....	7
Organizations Are Not Prepared To Deal With DDoS Threats.....	8
Cost Model: Companies Underestimate The Cost Of DDoS Threats .....	10
Estimating Overprovisioning Cost .....	10
Bandwidth Cost .....	11
Calculating Server Replication Cost And Management Cost.....	12
Estimate Service Disruption Cost.....	12
Calculating Loss Of Revenue .....	12
Repair And Forensics Cost.....	13
Estimate Probability Of Successful DDoS Attacks .....	14
Final Cost Model .....	15
Case Studies .....	16
Case No. 1, A Large Financial Institution: Company Expected To Spend Nearly A Million Dollars A Year On DDoS Threats .....	16
Case No. 2, A Medium-Size eCommerce Company: Expenditures Reach \$795,000 Annually On DDoS Protection .....	16
A Note About The Estimated Cost .....	17
Overprovisioning And On-Premise Alternatives Are Not The Answer For DDoS Threats .....	18
In-The-Cloud DDoS Protection: Best Positioned For Battle With Cybercriminals .....	19
Conclusions .....	22
Appendix A: DDoS Attack Probabilities .....	23

Endnotes.....24

© 2009, Forrester Research, Inc. All rights reserved. Unauthorized reproduction is strictly prohibited. Information is based on best available resources. Opinions reflect judgment at the time and are subject to change. Forrester®, Technographics®, Forrester Wave, RoleView, TechRadar, and Total Economic Impact are trademarks of Forrester Research, Inc. All other trademarks are the property of their respective companies. For additional information, go to [www.forrester.com](http://www.forrester.com).

### Executive Summary

Distributed denial of service, otherwise known as DDoS, refers to a form of attack whereby a number of source IPs simultaneously send an abnormally large number of packets to a particular destination, thereby overwhelming the bandwidth or the processing power of the destination. Over the years, the industry has seen many large-scale DDoS attacks, including some that took down the services of well-known, major global enterprises.<sup>1</sup> Despite the efforts of the security community, DDoS continues to wreak havoc on the Internet.

In October 2008, VeriSign commissioned Forrester Consulting to conduct a study to understand the state of and the need for DDoS protection in enterprises. Forrester conducted in-depth qualitative interviews with 19 IT and security professionals who are responsible for networking, operations, and security for organizations with a significant online presence. The organizations we interviewed are chosen from enterprises with an annual revenue of \$1 billion or above and a service availability requirement of at least 99.9%. They provide Internet services ranging from electronic banking, online eCommerce, and ISP services.

As a result of the study, we found that disruption of services for these organizations would cause devastating consequences, both financially and in terms of damage to business reputation. One large eCommerce company we surveyed would see a \$19 million loss if its services were disrupted for an hour. Generally speaking, companies with high availability requirements would stand to lose substantial business revenue if their online services were disrupted, due to attack or other reasons.

The study also found that companies today do not have adequate protection against DDoS attacks; many overprovision their bandwidth to account for unexpected traffic, including DDoS. As a result, it is not unusual for companies to spend 75% more on extra bandwidth. Bandwidth overprovisioning is an approach that optimizes for the worst case, and as such is far from the most economical or effective solution. Moreover, bandwidth overprovisioning is losing the battle against DDoS: Recent attack statistics show that some of the latest DDoS attacks now carry well more than 1 million packets per second (MPPS), with the largest nearing 5 MPPS. This type of attack can quickly overwhelm even the most well-provisioned enterprise networks.

The trends of DDoS attacks clearly point to a need for better DDoS protection. The current practices, which rely primarily on premise-based measures — including bandwidth overprovisioning — will soon become inadequate or prohibitively expensive. Rather than going it alone with on-premise only strategies, firms at risk for DDoS should look to cloud-based services that eliminate unwanted traffic in the core of the Internet and deliver legitimate traffic to the edge of the enterprise network. By virtue of being a service, these offerings share threat visibility, expertise, and assets across the client base while taking advantage of a broad view of Internet traffic, without requiring upfront capital expenditure.

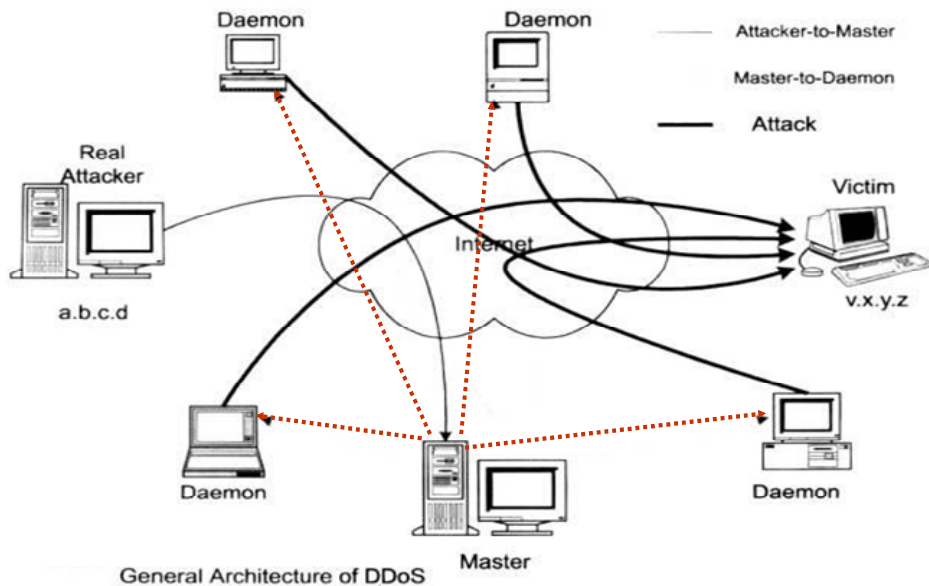
## Introduction

Companies operating critical online business, including those that provide electronic banking and eCommerce services, depend on their services to be “always up” for business success. For those companies, service disruption often means substantial financial loss as well as damage to business reputation.

One distinct threat against the continued operation of critical online services is DDoS. DDoS attacks are designed to overwhelm the victim with an abnormally large volume of traffic, so that it overflows the victim’s bandwidth or its CPU or memory capacity. The victim of a successful DDoS attack would lose its ability to operate its online services for the duration of the attack. Today’s DDoS attacks are often criminal activities — they are orchestrated by organized crime rather than random individuals, for purposes of cyberterrorism, industry espionage, and even extortion.

Figure 1 shows a general architecture of DDoS. For each attack, there is typically a control center (the master) that controls a large number of daemons (also known as bots). The master issues attack commands to the bots, which then send attack packets to the victim. The master is in turn controlled by the real attacker in the background, who may be several times removed from the attack bots so he’d be difficult to locate.

**Figure 1: Architecture Of DDoS**



Source: Forrester Research, 2008

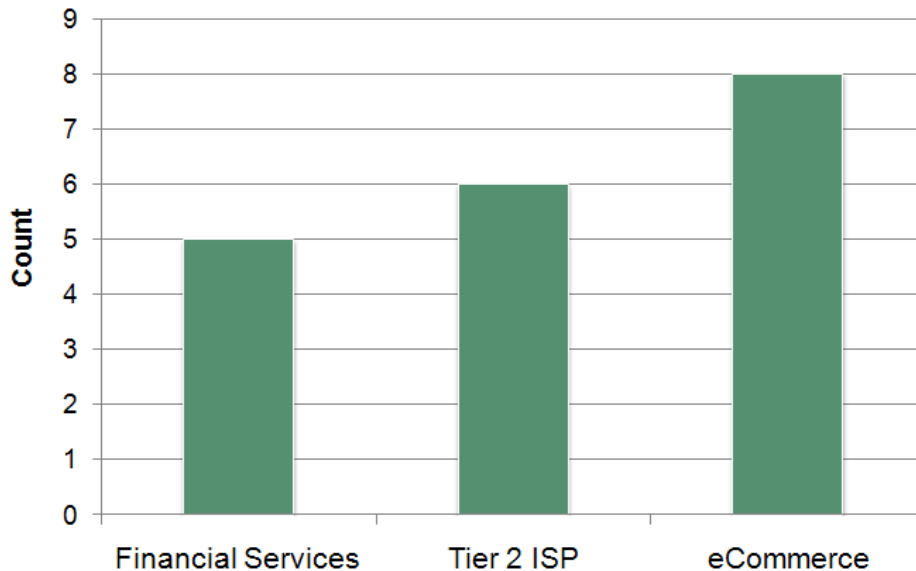
## Survey Methodology

VeriSign commissioned Forrester Consulting in September 2008 to conduct a qualitative study of DDoS threats and current mitigation mechanisms used by large organizations. The purpose of the study was to estimate the cost of current attack mitigation mechanisms and ascertain their efficiency and efficacy.

Forrester chose 19 companies in three industry verticals as our survey targets: ISP, financial services, and online retail (see Figure 2). Companies in these verticals operate services that must be available 24x7. As such, they are extremely sensitive to service disruptions. To these companies, a successful DDoS attack would cause significant business damage.

We also selected companies with annual revenue of \$1 billion and above. Companies of this size have significant online revenue, and therefore are more likely to be the target of DDoS attacks.

**Figure 2: Survey Respondents By Industry**



Base: 19 Enterprise IT Professionals operating significant external online services

Source: "DDoS Protection," a commissioned study conducted by Forrester Consulting on behalf of VeriSign, September 2008

The interviews uncovered a number of interesting details, which led to the construction of a model to estimate the expected cost of DDoS mitigation to organizations. In addition to the interview data, we used a number of auxiliary data sources, including data from Arbor Networks, Shadowserver, and US-CERT. While the model is based on the interview data, the elements in the model and the logic behind it are general enough to extend to other industry segments.

While a few places throughout the paper present numerical data, we caution against a generalized use of these numbers, as our interview sample size for this study is not statistically significant. Forrester makes no assumptions as to the potential cost of DDoS threats [insert other areas of calculation] that other organizations will receive. Forrester strongly advises that readers should use their own estimates within the framework provided in the study.

## Businesses Stand To Lose Big With DDoS

As discussed earlier, Forrester interviewed 19 companies that operate significant online services. These companies reported an availability requirement of at least 99.9% (also known as a three-nine requirement), including five companies with a five-nine requirement and six with a four-nine requirement. In availability parlance, a five-nine availability requirement means that every year, there should be no more than 5.25 minutes of downtime. A four-nine requirement translates to 53 minutes of downtime per year, and three-nine means 8.76 hours of downtime.

Many of the companies we interviewed stand to lose a substantial amount of revenue if their services are down for an extended period of time. Table 1 shows the estimated revenue loss for some of our survey respondents who reported five-nine service availability requirements.

**Table 1: Loss Of Revenue Due To Service Disruption (Services With Five-Nine Requirement)**

	Company A	Company B	Company C	Company D
<b>Loss of revenue per hour</b>	\$19 million	\$240,000	\$650,000	\$190,000
<b>Line of business</b>	E-banking	E-banking	eCommerce	eCommerce

E-banking companies A and B would lose \$19 million and \$240,000, respectively, for each hour their services are down. eCommerce companies C and D stand to lose \$650,000 and \$190,000 each per hour when their services are not available. These figures only account for loss of revenue; they do not include any service repair or incident response cost.

According to Arbor Networks, a typical DDoS attack can last anywhere between 2 and 6 hours.<sup>2</sup> For Company A (shown in Table 1), an average 4-hour DDoS attack will trigger a revenue loss of \$76 million! For companies that operate significant online business that require real-time or 24x7 access, it is not difficult to see that a successful DDoS attack will bring devastating consequences.

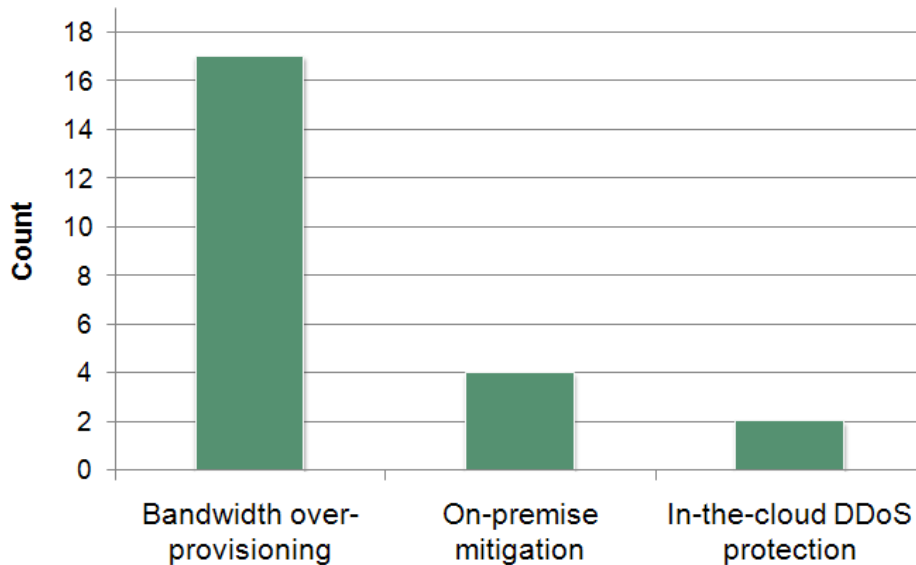
## Organizations Are Not Prepared To Deal With DDoS Threats

In the past 12 months, the Internet has seen a substantial increase of DDoS attack size.<sup>3</sup> The largest DDoS attack has now reached 5 MPPS, which consumes approximately 40 Gbps of bandwidth. The rate of the growth made many believe that the Internet will soon see 100 Gbps attacks.

In our study, we found that many organizations do not have specific DDoS protection mechanisms. The most popular anti-DDoS measure is bandwidth overprovisioning. Figure 3 shows the reported DDoS mitigation techniques used by our survey respondents; 17 out of 19 companies we interviewed cited bandwidth overprovisioning as their measure to guard against DDoS. Four companies utilize edge-based, on-premise DDoS mitigation, including DDoS-aware IDS and firewalls. Only two companies we interviewed reported the use of in-the-cloud or hosted DDoS mitigation services offered by an ISP or a third party.

Bandwidth overprovisioning is a commonly practiced method to accommodate peak flow. Many companies go a step further and purchase additional bandwidth beyond peak flow as a means of DDoS protection. For example, if a company's normal bandwidth usage is 100 Mbps, the company would purchase 400 Mbps from its ISP. Typically, that includes 100 Mbps for regular traffic, 100 Mbps to accommodate peak traffic, and the remaining 200 Mbps for potential attacks and other unforeseen circumstances.

Figure 3: Types Of DDoS Mitigation Techniques Employed (Multiple Answers Allowed)



Base: 19 Enterprise IT Professionals operating significant external online services

Source: "DDoS Protection," a commissioned study conducted by Forrester Consulting on behalf of Verisign, September 2008

## DDoS: A Threat You Can't Afford To Ignore

---

Bandwidth overprovisioning can provide protection against small DDoS attacks, but is powerless against attacks above its bandwidth tolerance. Within the companies we interviewed, the most well provisioned have a 7.5 Gbps overpeak bandwidth allocation. Even with such a generous provision, a 40 Gbps attack would leave these companies underprovisioned nearly by a factor of 6; they would in fact need to procure an additional 32.5 Gbps of bandwidth to handle this specific attack. Even if the traffic got through their pipe unperturbed, it would most likely overwhelm their edge firewall or application servers.

It is important to note that the criminal organizations behind these DDoS attacks have, at their disposal, a massive amount of botnet resources to launch DDoS attacks. It is not unheard of to have a botnet with millions of hosts. Groups in possession of such computing power can easily wield an attack that devastates the victim's bandwidth provision, no matter how generous it is.

## Cost Model: Companies Underestimate The Cost Of DDoS Threats

Of the companies we interviewed, many with five-nine availability requirements overprovision their bandwidth by an average of 75%. Those that have four-nine availability requirements typically overprovision by 62%, and three-nine organizations by 55%.

A 75% bandwidth overprovision may not sound like much — a company with a 100 Mbps bandwidth need would purchase instead 175 Mbps, which on the surface costs only an additional \$750 per month. However, this calculation ignores many factors, such as additional management cost, idle bandwidth, and more importantly, potential revenue loss in the event of a DDoS attack.

We propose a cost model to estimate the overall economic impact of DDoS mitigation, including capital cost, operational expenditures, and potential service disruption cost. The model is shown below:

*Total DDoS mitigation cost = overprovisioning cost + (probability of successful DDoS attack \* service disruption cost)*

The model includes these elements:

- **Overprovisioning cost.** This cost includes both capital expenditures (cost you pay to ISP for bandwidth) and operational expenditures. The latter includes human time to manage the bandwidth, etc.
- **Service disruption cost.** This is the combined cost of loss of revenue, repair, and forensics cost due to service disruption.
- **Probability of successful DDoS.** This is the probability that the organization experiences a DDoS attack that its bandwidth cannot absorb.

In other words, your total cost of using bandwidth overprovisioning to defend against DDoS is the operational cost you pay due to overprovisioning, plus the service disruption cost when an attack is successful.

### Estimating Overprovisioning Cost

Intuitively, the overprovisioning cost consists of bandwidth cost, server replication cost, and management cost. More specifically:

- Bandwidth cost is the cost of extra bandwidth used for DDoS protection. We are only interested in the amount of bandwidth that is considered over the traffic peak.
- Server replication cost is the cost of extra servers used to handle DDoS traffic. This is typically a capital expenditure cost.
- Management cost includes human time and other noncapital cost of operations in handling DDoS attacks.

***Bandwidth Cost***

Bandwidth overprovisioning can take many forms. A company may purchase redundant bandwidth from multiple ISPs, in which case the company may use the two ISP lines in a load sharing or active-passive mode. Overprovisioning can also happen on a single line, which means simply purchasing a larger pipe.

Calculating the bandwidth cost is straightforward. Using the standard bandwidth cost of \$10 per Mbps per month, an extra 100 Mbps of bandwidth would cost \$1,000 per month.

To see how the industry typically overprovisions, we introduce a concept called overprovisioning index or OPI. An OPI is the overpeak bandwidth, as a percentage of the peak bandwidth. For instance, if a company's peak traffic is 100 Mbps, but it has 200 Mbps worth of bandwidth, then the company is 100% overprovisioned, and its OPI is 100%.

Every company has a slightly different OPI. In our study, we endeavored to find an OPI that is representative of the industries we interviewed. What we found is that the OPI for companies with five-nine requirements is typically higher than those with four-nine or three-nine requirements. As such, for the remainder of this paper, when we speak of average OPIs, they will always be for a specific availability requirement level.

Tables 2, 3, and 4 show the respective OPIs for five-nine, four-nine, and three-nine companies. For the five-nine companies we interviewed, the average OPI was 75%, while average for four-nine companies was 62%, and 55% for three-nine companies. However, we do caution against the general usage of these average numbers, as they are derived from a small sample size.

**Table 2: Five-Nine Companies' Overpeak Provisioning Index**

Five-nine companies	Company A	Company B	Company C	Company D	Company E	Average OPI
Overpeak provisioning percentage	43%	100%	100%	33%	100%	75%

**Table 3: Four-Nine Companies' Overpeak Provisioning Index**

Four-nine companies	Company I	Company II	Company III	Company IV	Company V	Company VI	Average OPI
Overpeak provisioning percentage	43%	50%	100%	18%	100%	Not disclosed	62%

**Table 4: Three-Nine Companies' Overpeak Provisioning Index**

Three-nine companies	Company 1	Company 2	Company 3	Company 4	Company 5	Average OPI
Overpeak provisioning percentage	40%	35%	50%	100%	50%	55%

The OPI represents a metric that allows us to reason about the level of overprovisioning.

### *Calculating Server Replication Cost And Management Cost*

Other cost elements for overprovisioning are the cost of additional server capacity to process the traffic load and the cost to manage these servers. For redundancy reasons, high availability applications are typically replicated onto multiple servers. The servers can be used in a load sharing or standby mode. However, if overpeak traffic occurs, and the company has extra bandwidth to deal with the traffic, the servers must also have extra capacity to process the traffic. This extra capacity is the cost that interests us.

Many of our interviewees replicate applications onto two to five servers. Unfortunately, it is difficult to ascertain the extra capacity cost due to a number of reasons. The chief obstacle is that our interviewees were not able to tell us at what capacity level their servers run while peak traffic occurs. In addition, server replication is needed for reasons other than DDoS protection (e.g., redundancy). By the same token, it is difficult to ascertain the extra management cost due to overprovisioning.

We made a further observation that a company that consumes substantial bandwidth and has an OPI greater than 50% would spend substantially more on bandwidth than on server replication or management. As such, we decided to simply use loss of revenue to represent service disruption cost.

### **Estimate Service Disruption Cost**

Service disruption cost is the cost to the company when its services are disabled. This includes loss of business revenue, increased support cost, and repair/forensics cost. As many of the companies we interviewed handle their own customer support, a short duration of support call surge will not result in direct cost increase. We therefore deem the total service disruption cost as the sum of loss of revenue and repair/forensics cost.

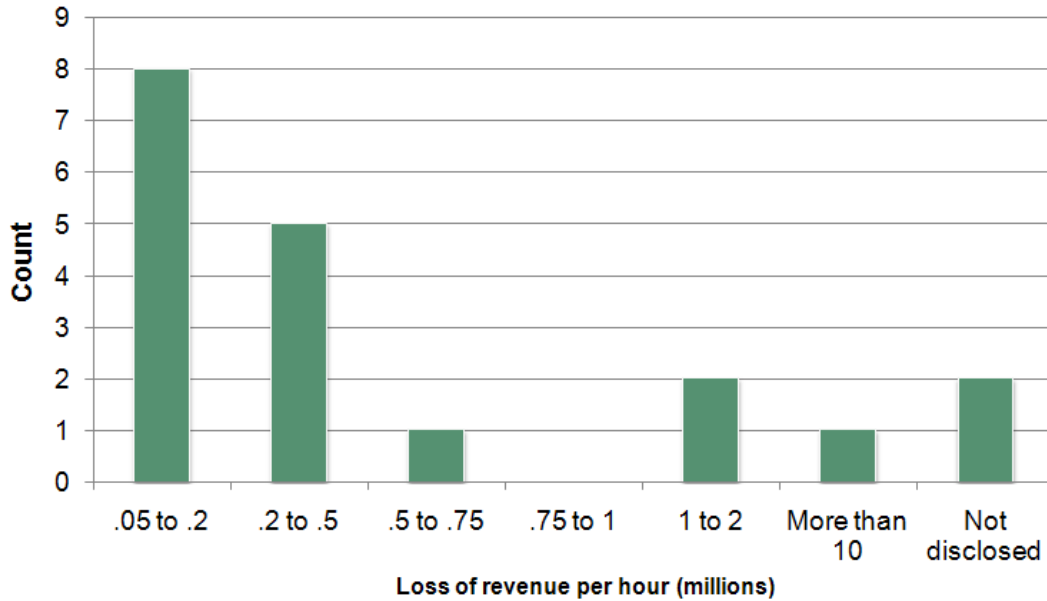
### *Calculating Loss Of Revenue*

Revenue loss due to service disruption is somewhat easy to pin down. However, it varies widely from company to company, depending on the line of business, company size, and service usage. Companies that provide real-time services (e.g., stock trades) will see immediate and drastic revenue loss if the service is disrupted, while those that operate non-real-time services would see moderate loss of revenue, as their customers would often resubmit orders later when the services come back online.

In addition to the figures already presented in Table 1, we summarize the hourly disruption cost for our interview respondents in Figure 4. Seventeen out of the 19 companies we interviewed disclosed their service disruption figure. One company stands to lose more than \$10 million in revenue each hour its services remain offline. Two of our survey respondents would lose between \$1 million and \$2 million per hour, five indicated that they would lose between \$200,000 and \$500,000 per hour, and eight would lose between \$50,000 and \$200,000 per hour.

The broad range of revenue loss figures in Table 1 and Figure 4 has much to do with the type of online services these companies offer. The company showing losses of more than \$10 million per hour if service is disrupted is a financial services company that conducts high-valued, real-time financial transactions online. This loss of revenue numbers may also include cost due to failure to sign on new customers during the downtime.

Figure 4: Distribution Of Loss Of Revenue Per Hour



Base: 19 Enterprise IT Professionals operating significant external online services

Source: "DDoS Protection," a commissioned study conducted by Forrester Consulting on behalf of Verisign, September 2008

### Repair And Forensics Cost

When a service disruption happens, an organization would typically carry out an incident response routine, which includes steps to restore the services and conduct forensics analysis. Such an incident response routine would consume human time and hence has an inherent cost associated with it.

We asked our survey respondents to disclose the typical number of man-hours that go into an incident response procedure to restore online services. Table 5 summarizes our findings with the five-nine companies.

Table 5: Repair And Forensics Cost

	Company A	Company B	Company C	Company D	Company E
Incident handling man-hours	38	200	Don't know	70	Don't know
Region	US	US	UK	UK	UK
Cost*	\$1,486	\$7,823		\$3,307	

\*This assumes that US fully loaded network engineer salary is \$78,230 and UK fully loaded network engineer salary is \$94,500.<sup>4</sup>

Although the cost varies from company to company, the repair and forensics cost is typically thousands of dollars. Compared with the loss of revenue, this cost is insignificant. As such, we simplified the model for service disruption cost to include just the loss of revenue, that is:

*Service disruption cost = Loss of business revenue due to disruption*

### Estimate Probability Of Successful DDoS Attacks

Over the years, DDoS attacks have increased both in volume and frequency. In 2002, a typical DDoS attack might constitute 500 packets per second (PPS).<sup>5</sup> In 2007, we saw attacks topping 5 MPPS.<sup>6</sup> Assuming a standard packet size of 1.5 KB, a 1 MPPS attack carries a bandwidth of 10 Gbps of attack traffic.

In addition, the frequency of attacks has increased tremendously — the Internet in 2005 might see a few hundred attacks per day. In the middle of 2007, reports indicate that as many as 8,000 attacks were seen on a daily basis. Forty-nine percent of all companies experienced a DDoS attack in a 12-month period between 2006 and 2007, according to the 2007 E-Crime Watch Survey, conducted by CSO Magazine, the US Secret Service, CERT, and Microsoft (September 11, 2007).<sup>7</sup>

It is difficult to ascertain the actual probability that a particular organization might be the target of a DDoS attack. However, if you look across a long-enough 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. The more interesting question is therefore how big the attack is and if it would impact an organization's infrastructure and operations.

In order to answer this question, we set out to gather attack data and samples from a variety of sources, including Arbor Networks, Shadowserver, and US-CERT. A 255-day traffic study from Arbor Networks collected statistics of more than 140,000 DDoS attacks. Some of the attack sizes are shown in Table 6.

**Table 6: DDoS Attack Statistics**

Bandwidth	Number
More than 5 MPPS	12
Between 4 MPPS to 5 MPPS	21
Between 3 MPPS to 3.99 MPPS	20
Between 2 MPPS to 2.99 MPPS	38
Between 1 MPPS to 1.99 MPPS	60
Less than 1 MPPS	140,000

Using these statistics and other correlating statistics from other sources, we derived a probability density function of general DDoS attacks on the Internet. The function and the detailed derivation can be found in Appendix A. Using this probability density function, we can estimate the probability that an organization may experience a DDoS attack of a certain size.

For example, an overprovision of 1.2 Gbps, which is a very generous bandwidth allocation, can tolerate attacks up to 800,000 PPS. However, according to the attack size probability density function in Figure 8, the probability that there is an attack larger than 800,000 PPS is approximately 11%. In other words, the probability that the company will succumb to a DDoS attack, even with a 1.2 Gbps overpeak provision, is approximately 11%.

It should be noted that the probability distribution in Figure 8 is a general distribution function, derived from statistics of overall Internet traffic. The probability of a particular organization experiencing a DDoS attack could be different. For high-profile organizations, it's likely that they might expect a higher rate of attacks than depicted here. As such, the probability distribution function in Figure 8 is a conservative estimate for organizations with a substantial online presence. Using this model, we can estimate the baseline probability that an organization would experience an attack that its bandwidth allocation cannot absorb.

### Final Cost Model

We are now ready to present the final model of total DDoS mitigation cost:

*Total DDoS mitigation cost = overprovisioning bandwidth cost + (probability of successful DDoS attack \* loss of revenue due to service disruption)*

We can easily extend this model to estimate the cost of other forms of DDoS protection. For example, if an organization uses an on-premise firewall or IDS that includes DDoS protection or source ACL capabilities, the cost model will become:

*Total DDoS mitigation cost = firewall/IDS operational cost + (probability of successful DDoS attack \* loss of revenue due to service disruption)*

Here the firewall/IDS operational cost is the capex and opex of operating the DDoS protection and source ACL part of the firewall or IDS. The probability of attacks, however, will change, because source ACL technology isn't set up to deal with large capacity of attacks — the probability that the organization will experience a debilitating attack will likely increase.

## Case Studies

In this section, we present two real case studies to illustrate the use of the cost model. For confidentiality reasons, the company's names are removed from this study. Forrester selected these two organizations because they represent two vastly different profiles of companies we interviewed: One is a large financial organization with a significant online presence, while the other is a much smaller eCommerce company with moderate online revenues and moderate resources. The calculations here estimate the expected cost to these organizations when they use bandwidth overprovisioning to guard against DDoS attacks.

### **Case No. 1, A Large Financial Institution: Company Expected To Spend Nearly A Million Dollars A Year On DDoS Threats**

This financial institution operates electronic banking and other financial transactions online. Its services have a stringent five-nine availability requirement. The bank's peak bandwidth usage is 10 Gbps, and the company uses an OPI that is consistent with the average OPI for five-nine companies. For each hour of service disruption, the company would lose \$650,000 of business revenues.

Our study shows that the average bandwidth OPI for five-nine companies is 75%. This means that for a 10 Gbps peak bandwidth, the company will likely provision a total of 17.5 Gbps of bandwidth (you can of course replace this with the actual overprovisioning bandwidth number, if that's available). Consequently, the company can tolerate attacks up to 7.5 Gbps (which is approximately 625,000 PPS).

The probability distribution function in Appendix A indicates that the probability for there being an attack of more than 625,000 PPS is approximately 2.5%. Assuming an average length of DDoS attack is 4 hours, and every Mbps bandwidth costs \$10 per month, we can estimate the total expected cost to the company as follows:<sup>8</sup>

*Total DDoS mitigation cost = overprovisioning bandwidth cost + (probability of successful DDoS attack \* service disruption cost)*

$$= (7,500 * \$10 * 12) + ((0.025 * 4) * \$650,000)$$

$$= \$965,000$$

The company is expected to spend nearly a million dollars annually on DDoS threats.

### **Case No. 2, A Medium-Size eCommerce Company: Expenditures Reach \$795,000 Annually On DDoS Protection**

This company operates eCommerce applications, with a three-nine availability requirement. The company's normal bandwidth consumption is 10 Mbps, and the peak bandwidth sometimes reaches 25 Mbps. The company purchases a 50 Mbps link from its ISP. For each hour of service disruption, it would lose \$220,000 of revenues. Compared with the company in case study No. 1, this company is moderate in scale and revenue size.

As the company reserves an additional 25 Mbps over its peak bandwidth, its OPI is 100%. Its annual bandwidth overprovisioning cost is  $25 * \$10 * 12 = \$3,000$ , assuming \$10 per month is the per-Mbps market price.

This additional bandwidth allows the company to tolerate attacks up to 25 Mbps, which is approximately 1,666 PPS. Using the attack probability values estimated in Appendix A, we see that the probability that there is an attack greater than 1,666 PPS is 90%. Therefore:

*Total DDoS mitigation cost = overprovisioning bandwidth cost + (probability of successful DDoS attack \* service disruption cost)*

$$= \$3,000 + ((0.9 * 4) * \$220,000) \text{ (assuming a DDoS attack lasts 4 hours)}$$

$$= \$795,000$$

The company is expected to spend at least \$795,000 annually in dealing with DDoS threats.

### **A Note About The Estimated Cost**

It is worth noting that the cost model we developed is conservative — it estimates the lower bound of expected cost. More specifically, the model predicts the expected cost for one successful DDoS attack per year. If the organization encounters multiple attacks, the cost will go up.


## Overprovisioning And On-Premise Alternatives Are Not The Answer For DDoS Threats

The analysis thus far, including the cost model and the case studies above, shows that overprovisioning bandwidth (and related server resources) is a costly and an ineffective method for DDoS protection, especially for those organizations that are likely targets for DDoS attacks. Between the overprovisioning cost and the probability of service disruption, organizations that bring in significant revenues from online services can be expected to spend hundreds of thousands, if not millions, of dollars every year on DDoS-related expenditures.

As a user organization, there are a limited number of alternatives available for DDoS mitigation. Other on-premise techniques include the use of IDS and firewalls with rate-limiting or source ACL capabilities — once DDoS attack sources are identified, IDS or firewalls can start dropping or rate-limiting packets from the sources. Application firewalls with rate-limiting capabilities can also be used to limit the number of concurrent sessions for a particular IP or an application interface.

On-premise DDoS detection and protection technologies can defend against small attacks, but will be increasingly ineffective as DDoS continues to grow in size. Historical DDoS data suggests that DDoS attack size has been on a steady increase for the past six years. Figure 5 shows the largest attack size observed since 2001.<sup>9</sup> At this rate, it is entirely possible that we could see 100 Gbps attacks by 2009.

Figure 5: Attack Size Increase Trend Over The Years



Source: Arbor Networks

Today, most organizations' network pipe to the outside world is no more than 10 Gbps, which is similar to the bandwidth of some, so are the high-end firewalls and IDS devices. Modern attacks can quickly overwhelm such connectivity and the processing device before any DDoS detection can be done.

## In-The-Cloud DDoS Protection: Better Positioned For Battle With Cybercriminals

An alternative to an on-premise treatment to DDoS protection is in-the-cloud services. An ISP, network operator, or a third-party provider with large-enough capacity can provide such a service. Essentially, an in-the-cloud DDoS protection service means that packets destined for an organization (in this case, the end customer of the service) is first sent through an Internet scrubbing center, where bad traffic like DDoS packets is dropped and the cleansed traffic is then delivered.

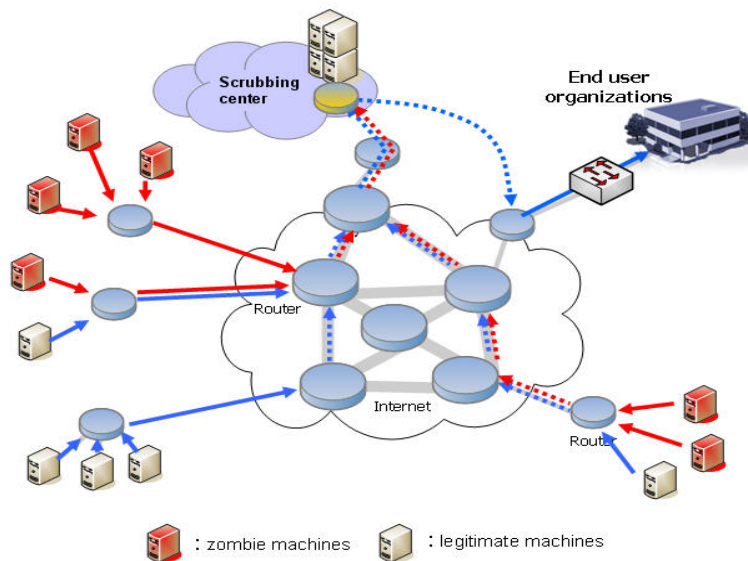
Large attacks are a rare event, but dealing with them requires specialized skills, technology, and bandwidth — yet there is no competitive advantage in maintaining those capabilities in-house if they are available from a service provider. The in-the-cloud DDoS mitigation service admittedly needs a substantial infrastructure, with adequate bandwidth and capacity to deal with traffic from multiple customers. But once the infrastructure is built, the service provider can share the skills and capacity across many clients, without clients having to build out their on-premise capacity. There are several advantages performing DDoS mitigation in the cloud:

- The service provider has a broad view of the Internet traffic across multiple clients and networks that it can learn from and apply mitigation to. For example, by looking across multiple clients' traffic, the service can quickly recognize malicious sources that participate in DDoS activities. As a result, this type of DDoS detection is much more effective and timely than any end user organization can do standalone.
- By virtue of sharing the service, the costs should be lower and the service better than a go-it-alone effort.
- The end user organization need not invest any on-premise resources, either capital or operational, to deal with traffic that is not wanted in the first place. The service requires only an ongoing service expense.
- The scrubbing center would typically have core Internet connectivity and therefore has a large capacity to deal with traffic, much larger than a typical enterprise network. This means that it can deal with attacks larger than any single user organization can handle.
- By virtue of being a service, the service provider can be easily swapped out for another if the client's needs change.

These attributes of an in-the-cloud DDoS service are a great example of the industry buzz around the concept of cloud computing or cloud services. DDoS mitigation in the cloud is a virtual extension of one's enterprise infrastructure, which handles a particular networking and security function.

Figure 6 shows a general architecture for such an in-the-cloud traffic filtering service. As depicted, attack traffic and legitimate traffic alike are routed to an Internet scrubbing center, where attack traffic is discarded and only legitimate traffic is forwarded to the user organization's edge router.

Figure 6: A General Architecture For Cloud-Based Scrubbing



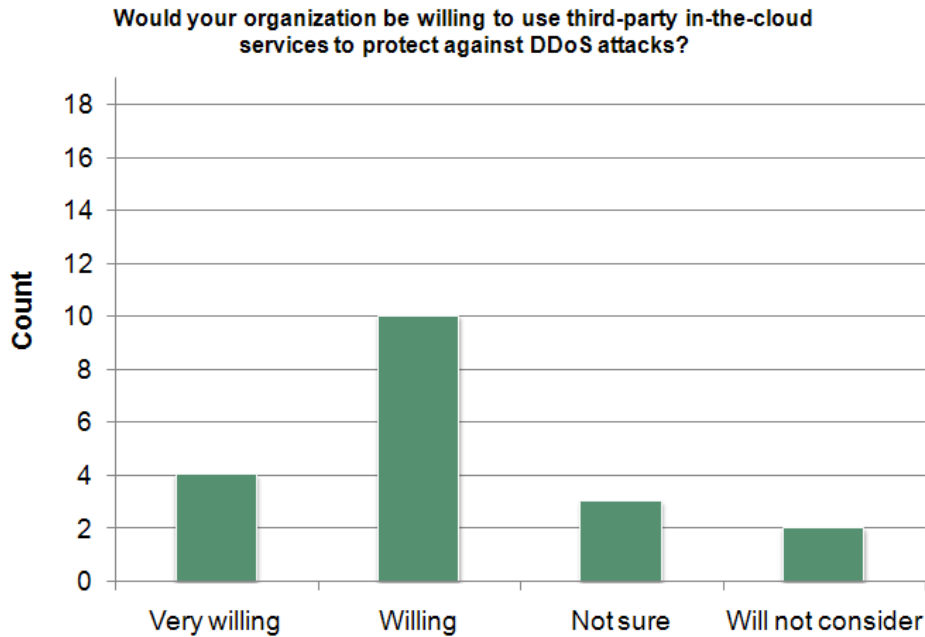
Source: Forrester Research, 2008

To be successful, the DDoS mitigation service must meet a number of requirements. They are:

- **The scrubbing service itself must be highly available.** If clients have a five-nine requirement for their services, the scrubbing center must also afford at least five-nine availability.
- **The service must provide timely DDoS detection.** It is not sufficient to wait until customers notify you that they are under attack. The service must proactively identify attack packets and route them through the scrubbing center. It is also a best practice that the scrubbing center deals with as little legitimate traffic as possible to minimize resource consumption.
- **Detection should be done as close to the core network as possible.** This means that the service provider can utilize core routing techniques, such as BGP RTBH, to curb bad traffic. Core routing is much more efficient than dealing with traffic at the edge of the network.
- **ISP neutrality is a plus.** Many organizations today use multiple ISPs for redundancy reasons. An ISP-neutral service provider therefore has many benefits, including the ability to work with multiple ISPs' traffic and immunity to network-specific attacks.

In our study, we found that many interviewees are enthusiastic about the prospect of such an in-the-cloud DDoS mitigation service. When asked: "Would you be willing to use third party in-the-cloud services to protect against DDoS attacks?" four out of 19 responded "very much so," while many others — 10 out of 19 — responded "willing" (see Figure 7).

Figure 7: Many Are Willing To Consider An In-The-Cloud Service Provider For DDoS Mitigation



Base: 19 Enterprise IT Professionals operating significant external online services

Source: "DDoS Protection," a commissioned study conducted by Forrester Consulting on behalf of Verisign, September 2008

How would an in-the-cloud DDoS mitigation service impact your cost expenditure? Using an in-the-cloud traffic scrubbing service, the DDoS mitigation cost becomes the sum of the scrubbing service cost plus the service disruption cost if the scrubbing service fails to stop an attack.

$$\text{Total DDoS mitigation cost} = \text{scrubbing service cost} + (\text{probability of scrubbing failure} * \text{service disruption cost})$$

In this model, we assume that the organization is not provisioning extra bandwidth overpeak traffic — using the in-the-cloud DDoS mitigation allows one to minimize its bandwidth cost.

Furthermore, if the scrubbing center does its job right, the probability of scrubbing failure should be close to zero. That is, almost all attack traffic will be discarded during the scrubbing process; only the legitimate packets are sent to the final destination. Therefore, the total DDoS mitigation cost becomes the scrubbing service cost. As the service is provided to multiple clients, costs should be lower than on-premise alternatives.

### Conclusions

DDoS is a persistent and growing risk that threatens organizations that operate online businesses. As attacks continue to grow in size and frequency, many are simply not prepared to deal with the threat.

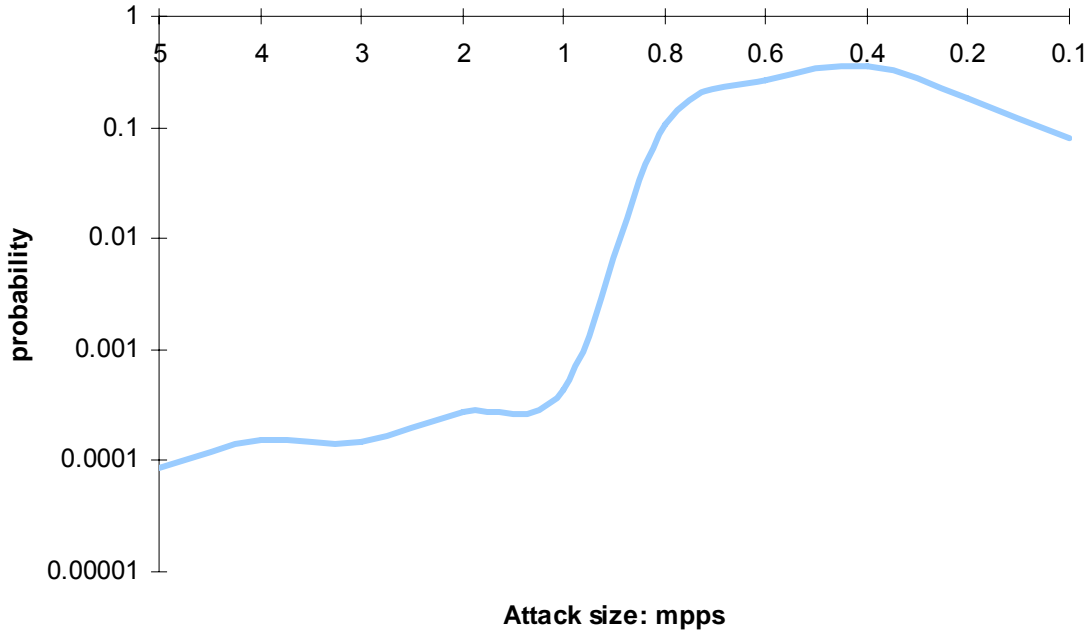
Forrester conducted a survey of 19 organizations that operate significant online business. In our study, we found that companies predominately rely on bandwidth overprovisioning to guard against DDoS attacks. As 40 Gbps attacks have been observed in the wild, bandwidth overprovisioning will quickly become ineffective. In today's climate, while many are asked to "pair down" resources and cut spending, relying on overprovisioning bandwidth to provide DDoS protection is an inefficient approach that will undoubtedly become obsolete soon.

Today, many on-premise technologies, such as firewalls and IDSeS, fall short in the face of large DDoS attacks. Organizations that are likely DDoS attack targets should consider cloud-based protection — in-the-cloud traffic cleansing to be more specific. In-the-cloud DDoS protection, when done close to the Internet core, presents a powerful countermeasure to DDoS threats. Not only is it more efficient to filter out unwanted traffic closer to the source of the attack, but it can also result in significant bandwidth and resource savings downstream, both for ISPs and end organizations.

There are, however, a number of challenges that the provider of the traffic filtering services must address. Among those are adequate infrastructure and the capacity to deal with flooding bandwidth, the ability to detect session-level attacks, and most importantly, the ability to keep up line performance while performing DDoS detection and mitigation.

## Appendix A: DDoS Attack Probabilities

Figure 8: Estimated Probability Density Function Of DDoS Attacks



Source: Based on data from Arbor Networks, Shadowserver, and US-CERT

Figure 8 shows a general probability distribution of DDoS attacks of varying sizes, based on data gathered from Arbor Networks, Shadowserver, and US-CERT. Note that the Y-axis in Figure 8 is on a logarithmic scale. Using this density function, we can calculate the probability of an attack greater than a certain size. For instance, the probability of experiencing an attack greater than 800,000 PPS is:

$$P(x \geq 800,000 \text{ PPS}) = 1 - P(x < 800,000) = 1 - \int_{800,000}^{\infty} f(x) dx$$

Where  $f(x)$  is the density function shown in Figure 8 and  $x$  is the variable that denotes attack size. The values in Figure 8 indicate that  $P(x \geq 800,000 \text{ PPS})$  is approximately 11%.

## Endnotes

---

<sup>1</sup> In February 2000, a DDoS attack crippled Yahoo!, Amazon.com, and CNN. Source: Corey Rice, "How a basic attack crippled Yahoo," *CNET News*, February 7, 2000 (<http://news.cnet.com/2100-1023-236621.html>). In March 2007, Internet domain registrar GoDaddy.com was hit by a DDoS attack that took it offline. Source: Dan Kaplan, "GoDaddy hit by DDoS attacks, not daylight-saving time issue," *SC Magazine*, March 13, 2007 (<http://www.scmagazineuk.com/GoDaddy-hit-by-DDoS-attacks-not-daylight-saving-time-issue/article/106078/>). A more recent attack targeted BBC's Web site in November 2008. Source: "DDoS attack crippled the BBC website," *Pingdom*, November 11, 2008 (<http://royal.pingdom.com/2008/11/11/ddos-attack-crippled-the-bbc-website/>).

<sup>2</sup> Source: Arbor Networks, "Worldwide Infrastructure Security Report, Volume IV," November 11, 2008 (<http://www.arbornetworks.com/report>).

<sup>3</sup> Source: Arbor Networks, "Worldwide Infrastructure Security Report, Volume IV," November 11, 2008 (<http://www.arbornetworks.com/report>).

<sup>4</sup> Figures cited are from Salary.com.

<sup>5</sup> Statistics can be found within the CERT advisory archive. Source: CERT (<http://www.cert.org/advisories/>).

<sup>6</sup> According to Arbor Networks, available through Arbor Networks Security Blog: Security To The Core. Source: Arbor Networks (<http://asert.arbornetworks.com/>).

<sup>7</sup> According to the E-Crime Watch Survey published in September 2007, by CSO Magazine, US Secret Service, CERT, and Microsoft. Source: CERT (<http://www.cert.org/archive/pdf/ecrimesummary07.pdf>).

<sup>8</sup> According to Arbor Networks and Shadowserver, a typical DDoS attack lasts anywhere between 2 and 6 hours.

<sup>9</sup> Source: Arbor Networks, "Worldwide Infrastructure Security Report, Volume IV," November 11, 2008 (<http://www.arbornetworks.com/report>).