VOLUME V

WORLDWIDE INFRASTRUCTURE SECURITY REPORT

2009 REPORT

ARBOR NETWORKS

Table of Contents

OVERVIEW 3
KEY FINDINGS 3
DEMOGRAPHICS OF SURVEY RESPONDENTS 4
SURVEY METHODOLOGY 5
MOST SIGNIFICANT OPERATIONAL THREATS 6
SCALE AND EFFECTIVENESS OF ATTACKS 7
ATTACK VECTORS 9
FREQUENCY OF ATTACKS 10
ATTACK DETECTION AND TRACEBACK 11
ATTACK MITIGATION 12
- Attack Mitigation Techniques 12
- Time to Mitigation 14

MANAGED SECURITY SERVICES 14
LAW ENFORCEMENT, CERTS AND CSIRTS 16
INFRASTRUCTURE PROTECTION TECHNIQUES 17
- Access to Network Elements 18
- Manage and Monitor Routers 18
- Application of Anti-Spoofing Techniques 18
- Protection of Routing Protocol Transport Connections 20
- Use of Internet Routing Registries 20
- Tools for Event Correlation 23
- ACL Revision Control 23
- Open Recursive DNS Resolvers 23
- On Disclosure 23
- Most Limited Vendor Features 24

SECURITY TEAM CHARACTERISTICS 24
- Composition of Security Teams 24
- Management- and Executive-Level Support 25

ISPS: BOTS, BOTNETS, AV AND MALWARE 26
Botnet Activities 26
Tracking Botnet Activities 26
Quarantine, Walled Gardens and Cleanup 27
IPV4 ADDRESS EXHAUSTION AND MIGRATION TO IPV6 27
ADDITIONAL QUESTIONS AND MISCELLANEOUS INFORMATION 28
ADDITIONAL INFORMATION, FREE-FORM COMMENTS 30
CONCLUSIONS 30
ABOUT THE AUTHORS 31

List of Figures

Figure 1: Largest DDoS Attack – 49 Gigabits Per Second 3
Figure 2: 2009 Respondent Organization Type 4
Figure 3: 2009 Respondent Geographic Distribution 5
Figure 4: Largest Anticipated Threat – Next 12 Months 6
Figure 5: Largest Attack Observed – Past 12 Months 7
Figure 6: Largest Observed Attack Vectors 9
Figure 7: Customer- and Infrastructure-Impacting Attacks 10
Figure 8: Primary Attack Detection Techniques 11
Figure 9: Primary Attack Mitigation Techniques 13
Figure 10: Revenue-Generating Network-Based Service Offerings for Customers 15
Figure 11: Attacks Referred to Law Enforcement 16
Figure 12: Mechanics Used to Access and Configure Network Devices 18
Figure 13: BCP 38/Strict uRPF Application 19
Figure 14: Routing Protocol Transport Protection 20
Figure 15: Route Registration by ISPs and Customers 21
Figure 16: Route Hijack Monitoring Tools or Techniques 22
Figure 17: Size of Dedicated Security Staff 25
Figure 18: Observed Bots – Past 12 Months 26

Overview

Arbor Networks, Inc., in cooperation with the Internet security operations community, has completed the fifth edition of an ongoing series of annual operational security surveys. This survey, covering roughly a 12-month period from 3Q 2008 through 3Q 2009, is designed to provide industry-wide data to network operators. This data can enable more informed decisions about the use of network security technology to protect mission-critical Internet and other IP-based infrastructures. The survey is also intended to serve as a general resource for the Internet operations and engineering community, recording information on trends and employment of various infrastructure security techniques.

Operational network security issues—the day-to-day aspects of security in commercial networks—are the primary focus of survey respondents. As such, the results provided in this survey more accurately represent real-world concerns than theoretical and emerging attack vectors addressed and speculated about elsewhere.

Key Findings

DDoS Bandwidth Growth Slows: Over the last six years, service providers reported a near doubling in peak distributed denial of service (DDoS) attack rates year-to-year. Figure 1 illustrates that peak attack rates grew from 400 Mbps in 2001 to more than 40 Gbps in 2007. This year, providers reported a peak rate of only 49 Gbps (a more modest 22 percent growth over the previous year). As we discuss later in the survey, the slowing in DDoS flood growth likely reflects attacks reaching underlying Internet physical constraints and a migration to other more effective denial of service attack vectors.

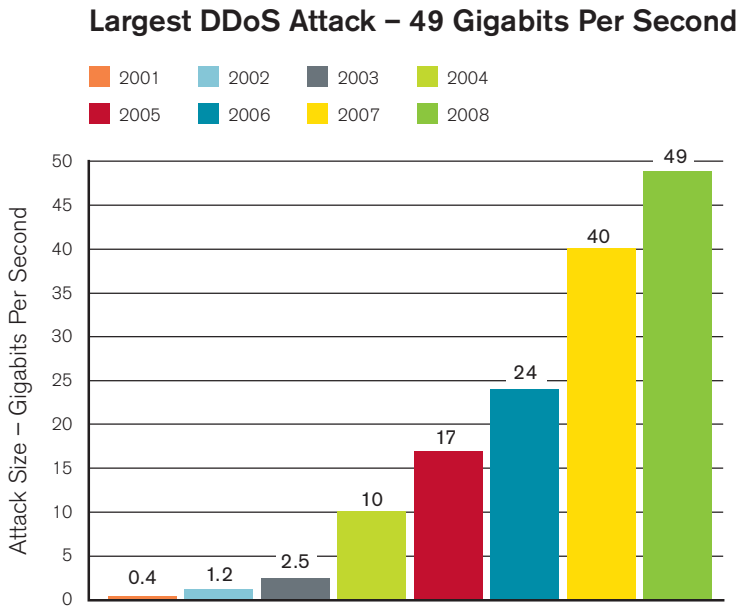

Figure 1: *Largest DDoS Attack – 49 Gigabits Per Second*

Source: Arbor Networks, Inc.

Attacks Shift to the Cloud: Again this year, more than half of the surveyed providers reported growth in service-level attacks at gigabit or less bandwidth levels. Such attacks are specifically designed to exploit service weaknesses, like vulnerable and expensive back-end queries and computational resource limitations. Several ISPs reported prolonged (multi-hour) outages of prominent Internet services during the last year due to application-level attacks. These service-level attack targets included distributed domain name system (DNS) infrastructure, load balancers and large-scale SQL server back-end infrastructure.

The Internet Is Not IPv6-Ready: A majority of this year's surveyed providers reported concerns over the security implications of IPv6 adoption and the slow rate of IPv4 to IPv6 migration. As in previous years, providers complained of missing IPv6 security features in routers, firewalls and other critical network infrastructure. Other providers worried the lack of IPv6 testing and deployment experience may lead to significant Internet-wide security vulnerabilities.

IPv4 Address Exhaustion, IPv6 Migration, DNSSEC Migration, 4-Byte ASN Migration: The 'perfect storm' of looming IPv4 address exhaustion, concerns surrounding migration to IPv6, concerns surrounding migration to Domain Name System Security Extensions (DNSSEC), and concerns surrounding migration to 4-byte ASNs is a source of uncertainty for respondents with regards to their ability to operate, maintain, secure and defend their networks.

Lack of Skilled Resources: Non-technical factors such as lack of skilled resources, internal/external communications siloing, lack of clearly defined operational responsibilities, lack of clearly defined policies, and lack of management understanding and commitment are the most significant obstacles to reducing mitigation times and proactively strengthening operational security postures.

Demographics of Survey Respondents

Survey participants included 132 self-classified Tier 1, Tier 2 and other IP network operators from North America, South America, Europe, Africa and Asia. All survey participants are directly involved in network security operations at their respective organizations. This year's participation doubles the 66 respondents to last year's survey and represents a notable increase in geographic diversity, albeit with some observable increase in participation among more local and regional network operators globally. While this increase represents a wider spectrum of respondents and in general is very positive, it also has an impact on various results and observable trends over the five-year survey lifetime—something we attempt to highlight accordingly where considered pertinent.

As illustrated in Figure 2, this year's respondent pool saw a major demographic shift away from self-described Tier 1 and Tier 2 providers and towards Tier 2/3 regional providers, educational networks, content providers/content delivery networks (CDNs), hosting providers and enterprise/hybrid operators. The "Other" category was composed of wireless operators, voice application service providers (ASPs), DNS top-level domain (TLD) operators and Internet exchange point (IXP) operators.

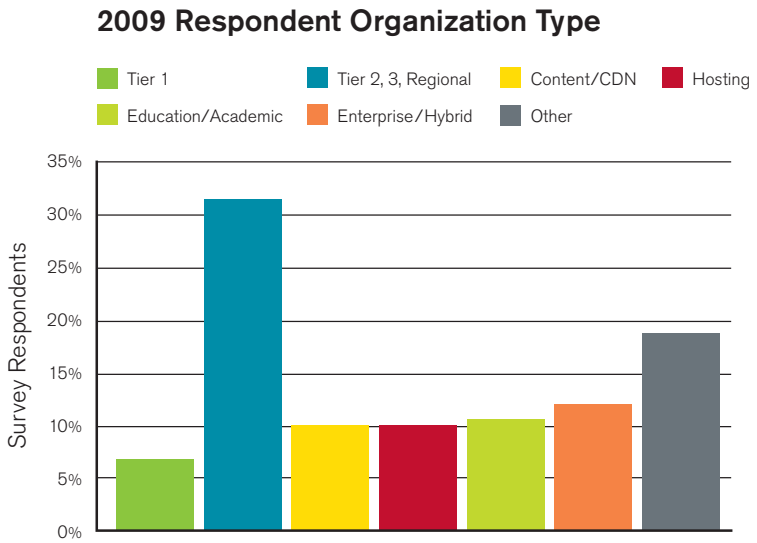

Figure 2: *2009 Respondent Organization Type*

Source: Arbor Networks, Inc.

As previously noted, this most obvious shift was primarily the result of this year's far wider pool of respondents rather than a decrease in self-described Tier 1 participation, which remains in line with that of previous survey reports.

Figure 3 depicts the geographic distribution among 2009 survey respondents.

We continue to believe this data is successful at highlighting the global representation afforded by a geographically diverse survey respondent pool, and note that this year's survey reflects the broadest geographical and organizational diversity of respondents to date.

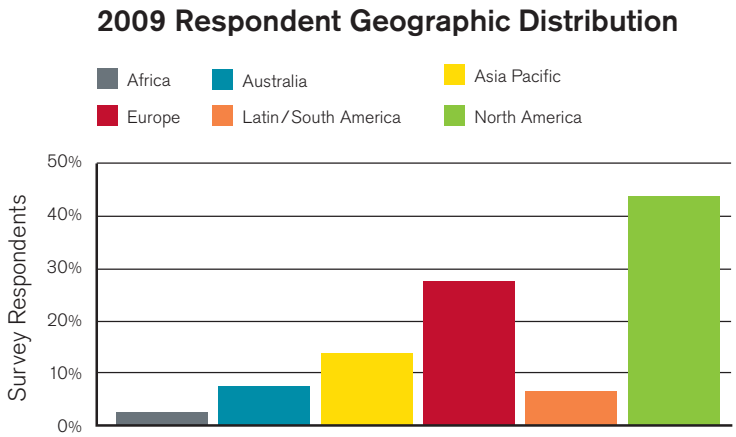

Figure 3: *2009 Respondent Geographic Distribution*

Source: Arbor Networks, Inc.

Survey Methodology

This edition of the survey consisted of 98 free-form and multiple-choice questions, representing an array of issues facing network security operators today. Questions addressed such topics as threats against backbone infrastructure and individual customers; techniques employed to protect network infrastructure itself; and mechanisms used to manage, detect and respond to security incidents.

All data represented here is presented in an aggregate and anonymous manner, and is provided with the permission of the respondents. Individual respondents were typically senior network security architects or operations engineers at their respective organizations. Standard mathematical methods to weight responses have been applied where incomplete answers were provided for a given question. Several refinements occurred in this edition of the survey, primarily based on respondent feedback. Some questions were deleted, some added and many simply honed in an attempt to capture more of the desired data sets. Again this year, several of the additional questions were added verbatim as provided by respondents to a previous survey, or as a result of direct feedback from one of the many polled network security or operations forums from which survey review was expressly solicited.

Arbor Networks intends to continue conducting this survey annually and sharing all results with the global Internet security and operations communities. Our goals are: 1) to continually refine the questionnaire in order to provide more timely, detailed and relevant information in future editions; and 2) to increase the scope of the survey respondent pool to provide greater representation of the global Internet network operations community, as illustrated above.

Most Significant Operational Threats

Respondents were asked to rank which threats they believe would pose the largest operational problems over the next 12 months (Figure 4). Displacing bots and botnet-enabled activities from last year, services, host or link DDoS threats took the top spot at nearly 35 percent, followed by botnets and bot-enabled activities at 21 percent. Additional concerns, in descending order, included credentials theft, DNS cache poisoning, route hijacking, system or infrastructure compromise, and worms.

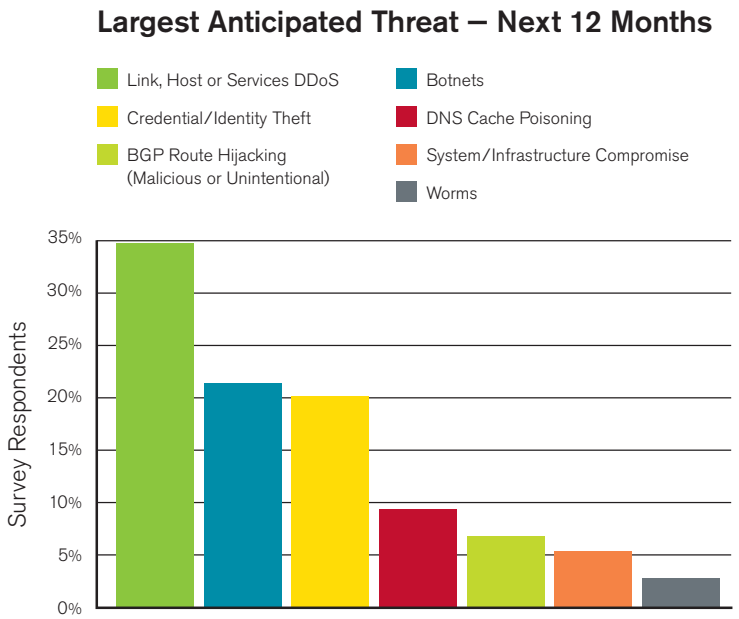

Figure 4: *Largest Anticipated Threat – Next 12 Months*

Source: Arbor Networks, Inc.

Network-based worms have declined significantly as a perceived threat over the last several years. This is to be expected given a concerted and effective effort by operating system vendors to decrease "wormable vulnerabilities," and in part reflects the continued shift to client-side infections and Web 2.0 worms affecting popular sites and services, such as Twitter and Facebook. Coincidentally, the 20th anniversary of the Morris Worm in November 2008 coincided with an out-of-cycle patch from Microsoft to address a "wormable vulnerability" described in Microsoft Security Bulletin MS08-067, a vulnerability for which exploits were seen nearly immediately in the wild. It cannot be understated that MS08-067 in late 2008 was considered an anomaly.

DNS cache poisoning dropped observably in the rankings as a primary concern, perhaps in large part because of little observed exploit activity in the wild, and certainly because the previous year's survey feedback period squarely overlapped with the disclosure of new cache poisoning techniques. The increase in relative prioritization of system/infrastructure compromise and credentials theft reflects the growing awareness of and emphasis on security vulnerabilities in infrastructure components in general and within the security research community in particular. The increasing numbers of vulnerabilities and fixes announced by major infrastructure vendors is the single largest factor in raising public consciousness of this threat category.

As previously indicated, respondents were also provided with a free-form text entry field under this line of questions, as well as many others, in order to provide some additional color and input regarding their chief concerns. The introduction of Domain Name System Security Extensions (DNSSEC), IPv6 deployment, and routing system threats and extensions (e.g., 32-byte AS numbers) were resounding themes.

Respondents were also asked to indicate which of the following areas they believe consume the largest amount of operational resources today:

- Distributed Denial of Service (DDoS) Attacks
- Spam
- Peer-to-Peer (P2P)
- Ongoing (Constant) Security Events
- Bots/Botnets
- Other

Spam retains the number one spot from last year, with background "clutter" from automated activities such as port scanning/host scanning/worm propagation and bots/botnets tying for second place. DDoS attacks, followed by peer-to-peer (P2P)-related activities such as law enforcement engagement, round out the operational resource consumption rankings. It is apparent from the feedback that background radiation and constant/ongoing security probing, while not consuming an enormous amount of bandwidth, do consume a sizeable amount of security resources.

Another 15 percent of respondents indicated that something else ("Other") consumes the largest amount of operational resources. Several respondents listed the operational burden of patching and vulnerability remediation as their primary time sink, along with end-customer interaction and responding to law-enforcement warrants, subpoenas, etc., for information related to ongoing investigations and prosecutions.

Scale and Effectiveness of Attacks

Respondents were asked again this year about the scale of the largest attacks either they or their customers endured at any given instant over the previous 12-month period. Figure 5 illustrates their responses. Only 19 percent reported the largest attacks they observed as being within the 1-4 Gbps range this year, as opposed to some 30 percent in 2008. However, last year 57 percent of survey respondents had reported observing attacks larger than 1 Gbps, while that number this year increased to just over 60 percent, even with a much wider pool of respondents.

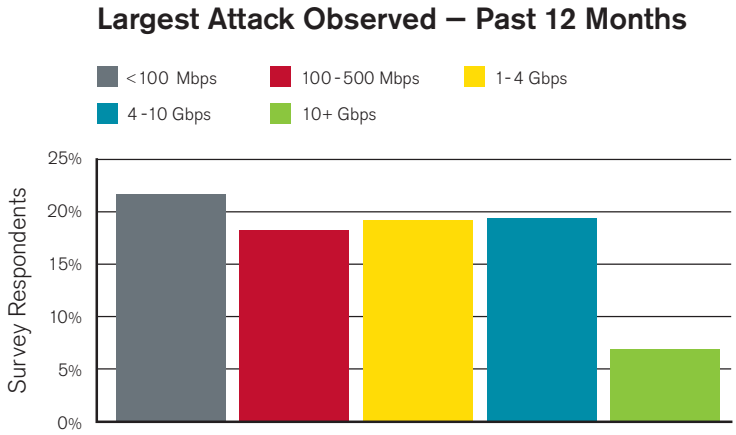

Figure 5: *Largest Attack Observed – Past 12 Months*

Source: Arbor Networks, Inc.

Additionally, there was a small increase in the number of respondents who reported attacks over 10 Gbps this year, with the largest attack reported by a European network operator as being 49 Gbps. The largest sustained attacks reported in the previous two editions of the survey were 40 Gbps and 24 Gbps, respectively. This represents a 23 percent increase in attack scale over last year, and an increase of nearly 150 percent of the largest attack reported in 2007. Figure 1, *Largest DDoS Attack* (page 3), illustrates the continued growth in attack scale since 2001.

It should be observed that DDoS attack scale growth has slowed over the past 12 months in comparison to previous years. The authors believe this to be the result of several factors. These include gating the upper bounds of IP backbone network capacity (e.g., Nx10 Gbps backbone link rates, awaiting upgrades to 100 Gbps rather than 40 Gbps deployment); the overwhelming effectiveness of attacks at smaller rates, leading miscreants to conserve bot resources and reserve firepower; better management of DDoS bot resources, etc.; employment of pulsing and similar attacks; and a pronounced move to more effective lower-rate application-layer attacks. We expect DDoS attack rates to continue to grow, but note that most enterprises are still connected to the Internet at speeds of 1 Gbps or less, and any attack near that rate or larger is typically sufficiently effective.

We again asked the respondent this year who reported the largest sustained attack (49 Gbps) if he could provide any details on the attack for inclusion in the survey; we are still awaiting his response.

Of particular note, reflective amplification attacks—a common large-scale DDoS attack vector—were responsible for the largest reported attacks in terms of bandwidth consumption during 2007 and 2008. These attacks exploit IP address spoofing (e.g., the reflective part) and protocol query/response behaviors by spoofing queries with a given target as the source address. This elicits much larger responses (e.g., the amplification part) from third-party systems (e.g., DNS resolvers), which respond to the target with the large payloads. Attacks of this nature employed against Internet root and TLD name servers in early 2006 achieved a 1:76 amplification factor. The attacks reported in 2008 reportedly had an even larger amplification factor. With such levels of amplification, a small number of well-connected hosts are capable of generating large amounts of attack traffic, easily overwhelming most organizations connected to the Internet today. We discuss some of the anti-spoofing techniques that operators can implement later in the *Infrastructure Protection Techniques* section (page 17), and survey where and why anti-spoofing mechanisms, which would largely mitigate this type of attack, are not universally deployed today. Furthermore, expressed concerns indicated that DNSSEC deployment will increase both network and server-side resource consumption, and may enable even more effective DDoS attacks and amplification vectors in the near future.

As described in previous years, most individual core Internet backbone links today are no larger than 10 Gbps, with standardization efforts and some networking products coming to market to enable 40 or 100 Gbps IP links. As such, most of the larger attacks today still easily inflict collateral damage on infrastructure upstream from targets themselves, while completely overwhelming the actual targets. Furthermore, given that most enterprises and other network properties quite likely do not have more than 1 Gbps of aggregate Internet access capacity, organizations concerned with Internet availability must plan accordingly with their ISPs to be prepared to respond to attacks of such scale.

Even with a much larger and geographically distributed respondent pool this year, the scale of attacks observed over the previous 12-month period increased notably. In addition, attack vectors, which are covered in the following section, employed both large-scale, brute-force flooding attacks, as well as more virulent and sophisticated application-layer attacks. With these reported attack rates, cooperation between enterprises and ISPs, and all network operators and other stakeholders, is the only way to mitigate these attacks effectively and minimize the impact of target and collateral damage.

Attack Vectors

Respondents were asked what attack vector was employed for the largest attack they observed over the past 12 months, with responses as provided in Figure 6.

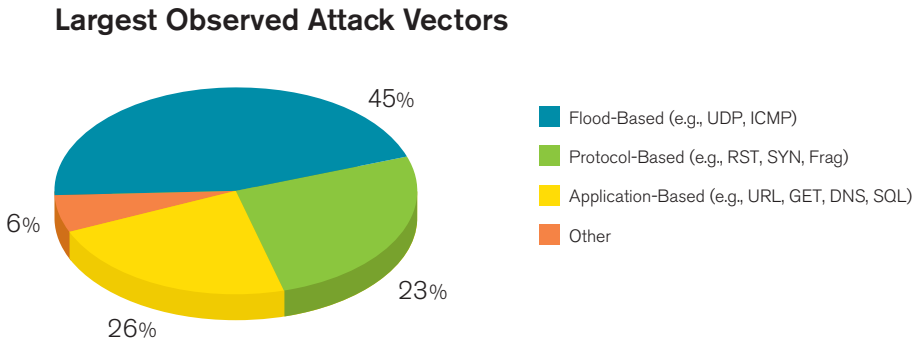

Figure 6: *Largest Observed Attack Vectors*
Source: Arbor Networks, Inc.

Flood-based attacks remain the most predominant attack vectors reported in the survey, accounting for nearly half of the vectors employed in the largest observed and reported attacks. Many of the respondents mentioned again this year that they are seeing an increase in application-based attacks aimed expressly at triggering back-end transaction activity and resource state:

- The respondents note that the application-level attacks, while not the largest in traffic volume, are some of the most sophisticated and operationally significant attacks they have observed year-over.
- When asked if they have observed any trends in attacks moving from brute-force to more complex attacks over the past year, 58 percent of the respondents indicated they have observed such trends. Brute-force attacks against DNS, Secure Shell (SSH) and Hypertext Transfer Protocol (HTTP) are the most common attack types cited in this category.

Frequency of Attacks

Actionable attack frequency remained somewhat constant on aggregate again this year. However, the distribution of attack frequency was somewhat wider yet again, an effect we anticipated as we continue expanding the respondent pool. Figure 7 indicates the number of attacks per month that impact customers and network infrastructure, respectively.

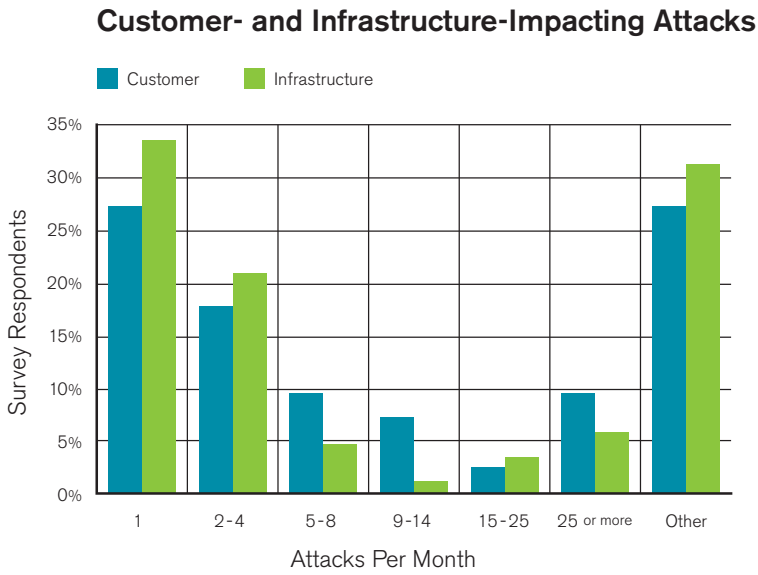

Figure 7: *Customer- and Infrastructure-Impacting Attacks*

Source: Arbor Networks, Inc.

When respondents were asked what attack vectors have been employed when attacks target their infrastructure, an array of responses were received. These included:

> "Brute-force attacks, packet-floods targeting router interfaces and network elements."
>
> "Known vulnerability probes."
>
> "Application-based attacks targeting network services."
>
> "Lots of SSH brute-force login attempts to all network elements and management systems."
>
> "Flood-based attacks targeting external BGP 'peering' addresses."
>
> "Multi-mode attacks, from SYN and DNS, evolving to application-layer attacks targeting customer-facing systems and networks elements."
>
> "DNS and related name hijacks, mainly through the provisioning side of DNS (i.e., Registrant or Registrar credentials compromise)."
>
> "Continued increases in scans expressly targeting router OS vulnerabilities."
>
> "Attacks seemingly crafted to require CPU-intensive filtering on infrastructure devices."

Several respondents indicated that infrastructure-impacting attacks they observe are not generally expressly targeting their infrastructure, but instead, are simply the result of collateral damage, as previously discussed.

When asked where infrastructure and internal security incidents occurred in the past, respondents indicated that these were the primary threat vectors:

60% – External Brute-Force Attacks
12% – Known Vulnerability
3% – Misconfiguration
3% – Social Engineering
2% – Insider Threat
0% – Zero Day
20% – Other

External brute-force attacks seem to remain high, particularly SSH brute-force login attacks, while insider threats are still quite low on the chief concern index. When asked about attacks towards customers in the past, respondents indicated that little to no variation exists across attack vectors.

Attack Detection and Traceback

Respondents were asked to identify the primary tools and techniques they employ for attack detection and traceback. Figure 8 illustrates the distribution of responses and compares it to responses in three previous editions of the survey.

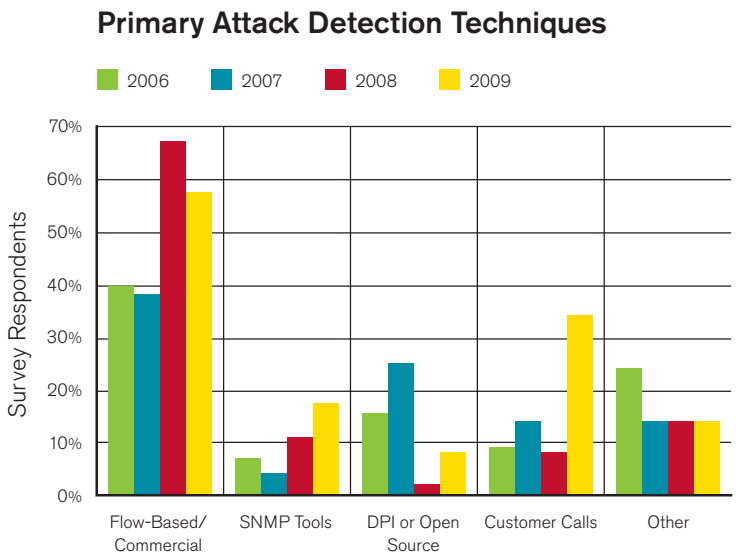

Figure 8: *Primary Attack Detection Techniques*

Source: Arbor Networks, Inc.

We expect that the wider distribution of smaller respondents accounts for the decrease in use of flow-based tools that can help to enable attack detection and traceback. An increase in absolute numbers was observed this year in terms of number of respondents employing these tools, although the overall percentage declined in relation to other tools due to the larger respondent pool. An increase in the number of reported "Open Source" tools was also observed, perhaps as many smaller network types were well-represented in this edition of the survey compared to previous editions. There was also a considerable increase in the number of respondents who rely upon a customer call to serve as the attack detection mechanism—from 8 percent to 34 percent—which also seems to reflect the shift in demographics of survey respondents this year. Finally, it should be noted that in the 2009 questionnaire, respondents were able to select multiple attack detection techniques. This accounts for the aggregate percentage (140 percent) factor exceeding 100 percent.

With regard to attack detection, respondents were asked to identify their mechanism for tracing attacks back to network ingress interfaces and upstream or downstream networks from an inter-domain interconnect perspective. This should obviously be considered an integral part of attack incident response, especially given the scale of attacks that have emerged over the past several years.

As illustrated in Figure 8 (page 11), 58 percent of the respondents indicated that they use flow-based tools to trace attacks back to network ingress interfaces. Another 18 percent of respondents indicated that they use Simple Network Management Protocol (SNMP)-based tools, 8 percent reported using deep packet inspection (DPI) or other tools, and 11 percent indicated that they have no current solutions or tools to trace attacks back to network ingress. The delta between this year's responses and last year's responses in this category also appears to be largely the result of the aforementioned shift in survey respondent demographics.

Attack Mitigation

Attack Mitigation Techniques

The number of respondents who employ either source or destination-based access control lists (ACLs) as their primary attack mitigation technique increased from 30 percent last year to 75 percent this year (Figure 9, page 13). The use of source and destination-based Border Gateway Protocol (BGP) remotely triggered blackholing (RTBH) also increased considerably. Intelligent filtering—or "scrubbing" using tools such as the Arbor Peakflow® SP Threat Management System (Peakflow SP TMS)—increased from 14 percent in 2008 to 18 percent over the past year even with a wider respondent distribution pool and size. Such intelligent filtering tools provide more effective mitigation techniques without effectively completing the attack or further disrupting services to the target. We believe the reported increase in intelligent filtering is related to continued growth in the number of network operators offering DDoS detection and mitigation services, as well as increased enterprise preparedness to DDoS incident response, as discussed in later sections. Again note that respondents this year were able to select multiple techniques as part of their primary mitigation methodology, so the aggregate percentages below will exceed 100 percent.

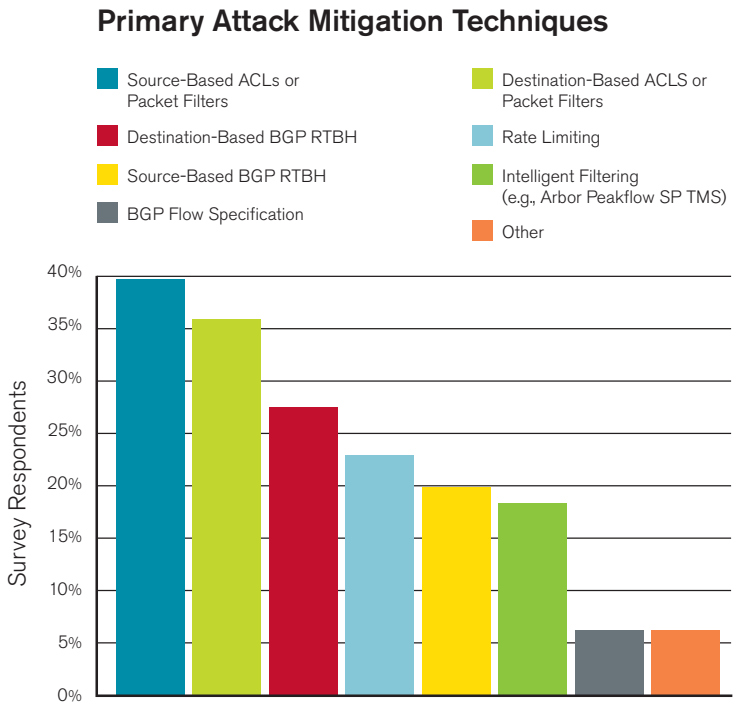

Figure 9: *Primary Attack Mitigation Techniques*

Source: Arbor Networks, Inc.

One disturbing statistical trend was a substantial increase in the reporting of rate-limiting as a primary mitigation mechanism, from 11 percent last year to 23 percent this year. Rate-limiting is almost always an iatrogenic strategy for mitigating DDoS attacks.

This is because constraining resources during a DDoS attack—which is aimed at exhausting capacity and/or state—ends up serving the goals of the attacker, as programmatically generated attack traffic tends to "crowd out" normal traffic from legitimate users. As with many other significant changes in this year's survey results, we believe this change to be largely a result of demographic shifts in survey respondents.

For more information on each of the above mitigation techniques, see the *Attack Mitigation* section on page 12 of the 2008 edition of this survey.[1]

1 www.arbornetworks.com/report

Time to Mitigation

Respondents were asked to indicate how long it typically takes to mitigate an attack once they have detected it, and what the largest obstacle is in reducing time to mitigation. Over 18 percent of the respondents indicated they typically mitigate attacks in less than 10 minutes once they have been detected, while another 17 percent said "less than 20 minutes," and some 15 percent said "less than 30 minutes." About 26 percent of respondents indicated that once detected, it takes about an hour to mitigate an attack, while 24 percent of respondents said they require more than one hour to mitigate an attack, even after it has been detected.

When explaining what the primary obstacle is in reducing time to mitigation, an array of input was provided. Some input varied widely, although common themes did indeed emerge. Some of the challenges identified with mitigating attacks include:

"Shortage of skilled operational personnel."

"Delays in internal escalation to capable staff, obtaining senior management authorization."

"Poor quality of data/tools in-house, identification and classification after detection."

"Lack of clearly defined policies surrounding attack mitigation."

"Lack of clearly defined and assigned responsibilities for attack mitigation."

"Inefficient, compartmentalized internal communications."

"Inefficient internal processes, bureaucracy."

"Difficulty in verifying an attack—ensuring it's not a flash crowd or other legitimate customer traffic."

"Inadequate budget for infrastructure to surgically mitigate attacks."

"Lack of cooperation from and coordination with internal groups."

"Difficulty contacting peers/upstreams/downstreams/customers."

"Managing capacity of dedicated mitigation devices."

"Language barriers with regards to international peers/upstreams/downstreams/customers."

It was also noted by more than one respondent that many of the abovementioned difficulties do not apply to attacks directed towards customers who subscribe to managed DDoS mitigation services and other network-based managed security services. One respondent also stated that he hopes more customers can be brought into the managed services fold solely because he would then have both the mandate and resources to assist them in defending their properties in a timely, efficient manner.

Managed Security Services

As discussed in the previous sections, opportunities involving managed security services (MSS) are causing many network service providers to invest a great deal in intelligent filtering or "scrubbing" infrastructure. With more mission-critical services being converged onto IP-based networks, and more revenue being tied to customer network availability, a DDoS MSS market emerged several years ago—purely out of necessity. Many organizations generate a majority and oftentimes all of their revenue from Web or other network service transactions, and their Internet "presence" and availability are critical to their fiscal well-being. Any disruption of network service has a direct impact on the financial performance and stability of these organizations. As a result, many enterprises have demanded that their service providers offer "Clean Pipe" services. These enterprises now consider a subscription to such services as an everyday cost of doing business on the Internet, and budget for these services just as they would disaster recovery, data backups and traditional network redundancy.

Figure 10 illustrates the number of respondents who currently offer network-based traffic visibility (i.e., Internet, MPLS VPN, etc.), DDoS detection and/or DDoS mitigation services to customers. The number of respondents who offer attack detection and reporting services was up to 36 percent from 24 percent last year, while attack mitigation services in aggregate dropped from 44 percent last year to 36 percent this year. Both of these numbers were skewed by the larger respondent pool down market, we believe, but this data still highlights continued deployment and offerings in this area by network services providers.

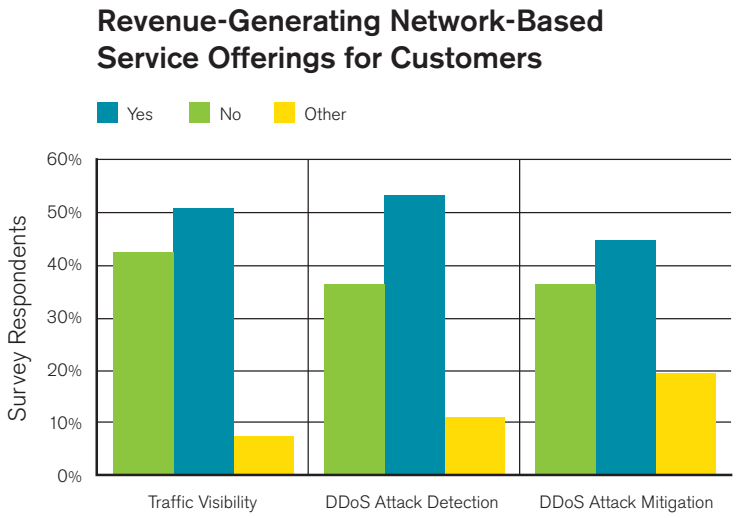

Figure 10: *Revenue-Generating Network-Based Service Offerings*

Source: Arbor Networks, Inc.

Note: Over half of the respondents provided no input to this line of questions, primarily because they do not offer network services (e.g., end-site, education, government or academic organization type). Incomplete or null responses to these questions were factored out when calculating percentages. As a result, these numbers primarily represent ISP survey respondents.

The majority of Tier 1 and Tier 2 respondents indicated that they currently offer DDoS detection and attack mitigation services. The "Other" category largely included hybrid enterprises, academic and other network types that deploy their own scrubbing infrastructure and/or subscribe to MSS offerings from their network services provider(s), as well as dedicated managed security service providers (MSSPs). Interestingly, many ISPs indicated that they offer traffic visibility services already, seemingly in parallel with those that offer DDoS attack mitigation services. This may in large part be attributable to the fact that market-leading solutions for DDoS detection and mitigation provide traffic reporting and visibility capabilities as well, and traffic visibility provides an easily obtained incremental revenue source with little or no additional capital expenditures or network infrastructure changes. Only 13 percent of the respondents indicated that they currently provide any type of service level agreement (SLA) for DDoS attack protection services, while 35 percent indicated they are actively considering the formulation of such service level agreements.

Law Enforcement, CERTs and CSIRTs

When respondents were asked how many attacks they have referred to law enforcement over the past year (Figure 11), responses varied widely based on organization types. Sixty-four percent of respondents indicated that they had referred no incidents to law enforcement over the past year, while another 27 percent said they had only referred five or fewer incidents to law enforcement. With 91 percent of the respondents indicating that they referred five or fewer (or none) responses to law enforcement over the past 12 months, this number nearly parallels the 88 percent reported in last year's survey, even with a much broader respondent pool.

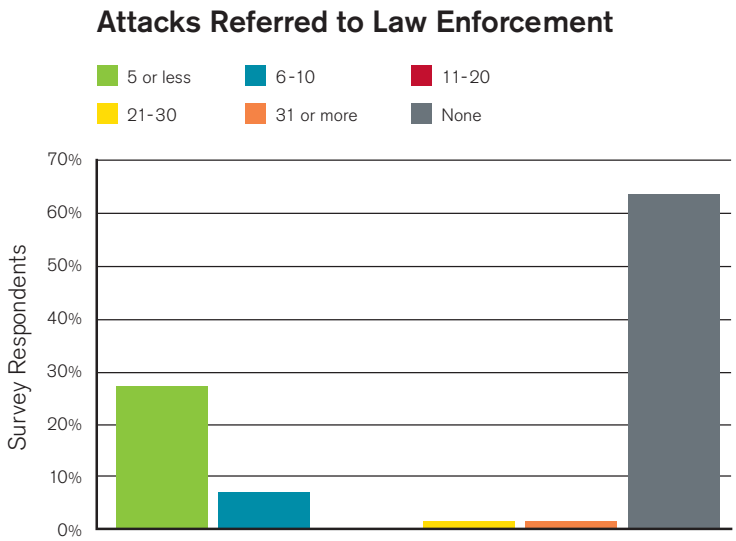

Figure 11: *Attacks Referred to Law Enforcement*

Source: Arbor Networks, Inc.

When asked what limited the number of referred attacks, responses were fairly well distributed and nearly identical to those received during last year's survey:

> "Law enforcement has limited/nil capabilities to respond."
>
> "Expect customers to report, will not report on their behalf."
>
> "Disbelief in utility of reporting."
>
> "Lack of clear jurisdictional delineation."
>
> "Sheer volume of attacks."

We also asked respondents if they believe law enforcement has the power and/or means to act upon information provided by network operators. Of the 74 who responded, only 20 percent said "Yes," while nearly 58 percent said "No," and 15 percent offered varying input on this question. Some 28 percent of respondents did indicate that they believe law enforcement is becoming more useful to Internet security operations, while only 4 percent said they believe law enforcement is becoming less useful. Of course, 33 percent said they "have seen no noticeable change in law enforcement activities," while another 27 percent said "What law enforcement?" (e.g., very little presence).

When respondents were asked if they have a computer emergency response team (CERT) or computer security incident response team (CSIRT), only 45 percent responded "Yes" (the same as last year), while 53 percent responded "No." Correspondingly and unsurprisingly, 50 percent of respondents indicated that they work frequently with a government or national CERT or CSIRT, and some 70 percent of respondents indicated that they do believe government CERTs/CSIRTs have a role and responsibility in operational security. Interestingly, this number is down from 77 percent last year.

Some 21 percent (down from 27 percent last year) of respondents indicated that they believe government has a role in enabling and assisting in infrastructure protection, but that it fails because of lack of knowledge. Another 11 percent (down from 18 percent last year) said government organizations "fail because of lack of cooperation with network operators," while 17 percent (up from 15 percent last year) said they "fail because of lack of regulation, policy or legislation." Nearly 15 percent (down from 23 percent last year) said they "fail because they're slow and far too political," while 13 percent said they "seem to be doing a decent job."

All in all, considering the wider respondent pool, it seems government and national CERT participation in operational Internet security incidents over the last year has been viewed in a more positive light by survey respondents. We would certainly like to see an increase in the number of organizations that have expressly defined internal CERT roles and procedures, and look forward to trending these numbers as more exposure is given to the role of CERTs in Internet security and incident response.

Infrastructure Protection Techniques

Again in this edition of the survey, we asked questions regarding several well-known infrastructure security techniques for the mitigation of spoofing and protection of routing protocols, as well as questions regarding management and monitoring of the infrastructure itself. The responses to these questions are indicated in the following sections.

When asked which infrastructure components are the most vulnerable, DNS services took the top spot, followed closely by session border controllers (SBCs), load-balancers and routers, and then IPTV infrastructure and services.

When respondents were asked if they have any tools or techniques to monitor for threats against DNS resolvers and/or authoritative name servers, 58 percent indicated they do have tools in place for detecting such threats. Fifty-seven percent of respondents indicated that they have tools in place to detect threats against voice over IP (VoIP) infrastructure or services.

Of the 75 respondents who answered questions related to the preparation and/or process of deploying DNSSEC for their infrastructure or customers, some 45 percent indicated they are currently working on DNSSEC deployment, testing or implementation, while 43 percent said they are not currently deploying or preparing for DNSSEC deployment. Another 12 percent of respondents indicated that they are uncertain if their organization is preparing for DNSSEC deployment. When asked about DNSSEC-specific concerns, many respondents expressed doubts about performance, the ability to deploy, maintainability, lack of tools, lack of operational experience, whether it will materially impact their security posture in a meaningful way, and various issues and concerns surrounding signing and key management.

Access to Network Elements

Respondents were asked what mechanisms they use to access and configure network devices (Figure 12). Their responses are as follows:

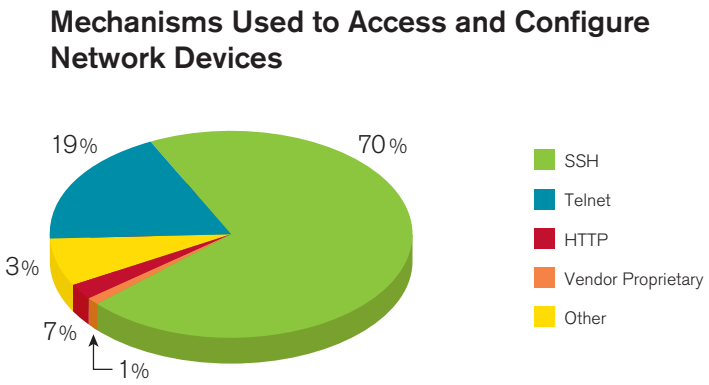

Figure 12: *Mechanisms Used to Access and Configure Network Devices*
Source: Arbor Networks, Inc.

Respondents indicated that SSH is preferred to Telnet by almost 4:1 (which is a good thing, considering Telnet is clear-text on the wire and easily snoopable). This is an improvement over last year's results. We believe that this is largely the result of increased awareness on the part of respondents and are delighted that even with the much wider distribution of respondents, SSH represents a larger percentage of the respondent's pool.

Manage and Monitor Routers

In addition to what tools respondents use for command line access to their network devices, we asked what protocols they use to manage and monitor their routers. Nearly 14 percent of the respondents indicated that they still use SNMPv1, while 65 percent reported making use of SNMPv2/v2c and only 12 percent indicated that they have migrated to SNMPv3 (more secure). Nine percent of respondents employ unspecified "Other" mechanisms.

Nearly 20 percent of respondents said they "Do" enable SNMP write access on network devices (nearly the same as the previous year, but with a much larger respondent pool), while 75 percent said they "Do Not" and 5 percent were unspecified. We do hope that operators continue to migrate to SNMPv3 and understand the security considerations of using SNMPv1 or enabling SNMP write access without putting proper techniques in place to mitigate threats from spoofed packets or other similar attacks.

Application of Anti-Spoofing Techniques

IETF BCP 38/RFC 2827[2] provides an overview on anti-spoofing measures that should be employed by network operators. Essentially, it recommends that a network operator should not accept packets into the network on a given interface for a given source address unless the packet's source destination address is considered reachable through that interface. This policy is traditionally implemented by provisioning ingress ACLs on each interface, statically defining which destinations are considered reachable through the interface, and therefore, which source address should be permitted to ingress the network through that interface.

2 www.ietf.org/rfc/rfc2827.txt

Unicast RPF (uRPF)[3] provides a mechanism to automate guidelines provided in BCP 38. There are multiple modes in which uRPF may be implemented to provide anti-spoofing measures, depending on the topology of the network and network equipment in use. Note that some of these modes permit spoofing within address spaces known in the local routing systems, and are therefore less effective than more explicit BCP 38 and "strict-mode" uRPF style policies. Furthermore, some types of uRPF can create a false sense of protection because, even when implemented, they can still allow attacks such as the reflective amplification attacks described earlier in this report. IETF BCP 84/RFC 3704[4] provides some additional information on these techniques.

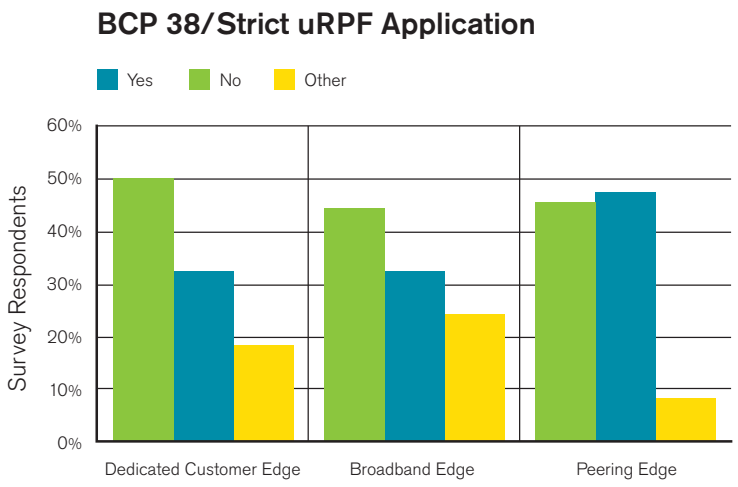

Figure 13: *BCP 38/Strict uRPF Application*

Source: Arbor Networks, Inc.

According to this data, in general, application or anti-spoofing techniques are implemented 28 percent of the time or less by this year's pool of respondents (Figure 13). Given that the most effective DDoS attacks today (e.g., reflective amplification attacks) and some of the common targeted attacks (e.g., DNS cache poisoning) are most effective when spoofing is possible, this data remains discouraging. That said, we believe that a large portion of the downward trend of percentages from last year is due to the aforementioned demographic shift in the pool of respondents. We also suspect that some level of anti-spoofing might be applied further upstream in many configurations.

Application of anti-spoofing techniques is without question one of the more effective ways to squelch many of today's most successful attacks. ISPs not implementing these techniques today should be focusing on doing so in the near future.

3, 4 www.ietf.org/rfc/rfc3704.txt

Protection of Routing Protocol Transport Connections

Network operators were also asked whether they use Transmission Control Protocol (TCP MD5) signature option, Generalized TTL Security Hack (GTSH) a.k.a. Generalized TTL Security Mechanism (GTSM), or Internet Protocol Security (IPSec) to protect BGP transport connections in the network, and whether they employ MD5 protection mechanisms available with their Interior Gateway Routing protocols (IGP) as well. Responses are illustrated in Figure 14.

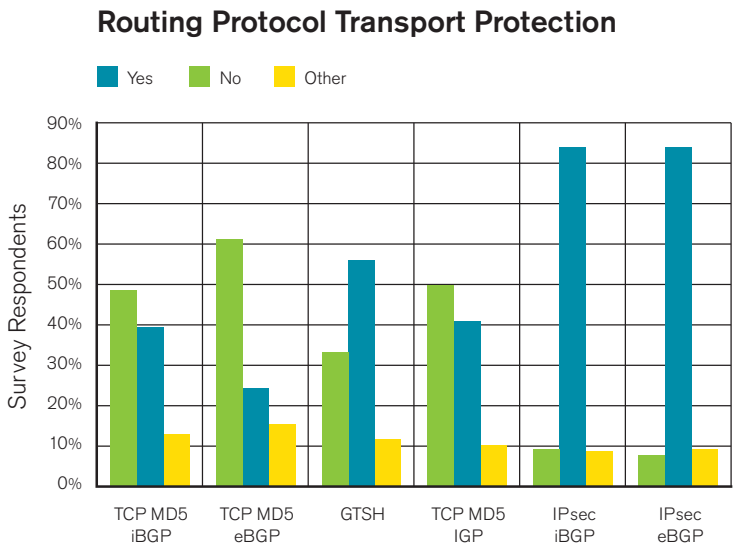

Figure 14: *Routing Protocol Transport Protection*

Source: Arbor Networks, Inc.

Sixty-one percent of the respondents use MD5 mechanisms to protect external BGP sessions in particular, with 48 percent using it to protect internal BGP sessions and 49 percent using MD5 protection mechanisms available within their IGPs. Nearly 9 percent of respondents reported using IPsec for iBGP and 8 percent for eBGP in this edition of the survey, in contrast with none reporting using it in 2008. Thirty-three percent of the respondents reported using GTSH/GTSM.[5]

Use of Internet Routing Registries

While one of the five Regional Internet Registries (RIRs) are responsible for allocation and assignment of IP addresses and Autonomous System (AS) numbers to ISPs or end sites, one or more (of many) Internet Routing Registries (IRRs) provide a database for ISPs to register routes and associated routing policies (Figure 15, page 21) in order to enable routing policy generation, obtain operational contact information, etc. There currently exists no strict linkage between RIRs and IRRs today with the exception of some Réseaux IP Européens (RIPE) mechanisms and some new work underway by the American Registry for Internet Numbers (ARIN). Nor does there exist any linkage between RIRs and the actual routing system itself. While some of this work is underway[6] and many holes exist in the security of the current Internet routing system, employing IRRs for routing policy generation and filtering of customers and peers alike is currently the best way to protect yourself and your peers from routing configuration errors and some malicious attacks.

5 www.ietf.org/rfc/rfc5082.txt
6 http://asert.arbornetworks.com/2008/05/using-rpki-to-construct-validated-irr-data

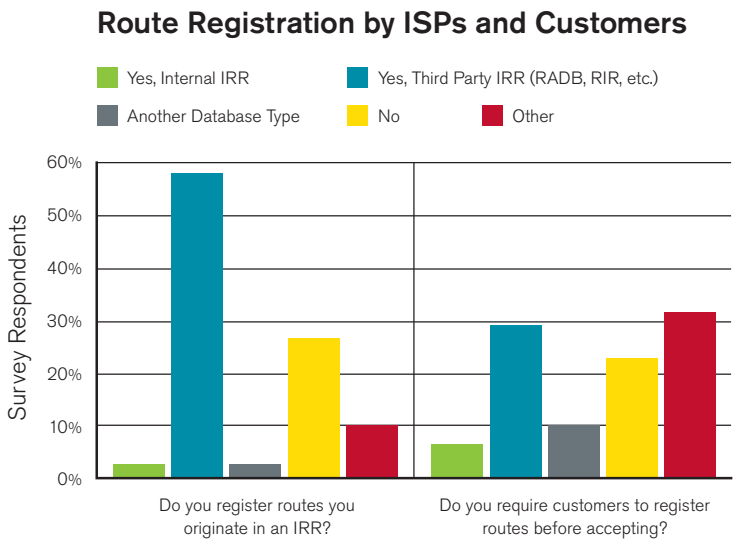

Figure 15: *Route Registration by ISPs and Customers*

Source: Arbor Networks, Inc.

Nearly 61 percent of respondents register routes they originate in any IRR, and only 35 percent require customers to register their routes before they will accept them. Both of these numbers are down from last year's edition of the survey (80 percent and 42 percent, respectively), but again this is likely an artifact of a much wider respondent distribution pool.

The responses above seem to only highlight the fact that very little or no inter-domain route filtering is applied on the Internet today, the state of which has only diminished the security of the Internet routing system over the past decade. One of the primary reasons for this almost certainly has to do with the insecurities and awkwardness of the IRRs themselves. However, it should also be attributed to the fact that no authoritative database exists today for verifying who has been allocated what IP address space, and what AS is authorized to originate or transit that address space reachability announcement. Fortunately, work is underway by several of the RIRs to develop a Resource Public Key Infrastructure (RPKI) that will provide just this allocation and route announcement authorization database, which can then be used either to populate IRRs or other tools, or in the future possibly directly by the routing protocols themselves, in order to enable more secure routing on the Internet.

Along the same lines, we asked respondents if they currently have tools that monitor for hijacking of their routes or those belonging to their customers. Fifty percent of respondents said they do have tools to monitor for route hijacks, while 35 percent indicated that they currently have no tools or services to monitor for such route hijacks (Figure 16, page 22). Given the frequency and exposure of route hijacks and insecurities of the routing system over the past year, we expect to see a continued increase in both commercial and open source tools and services for monitoring route hijacking over the coming years.

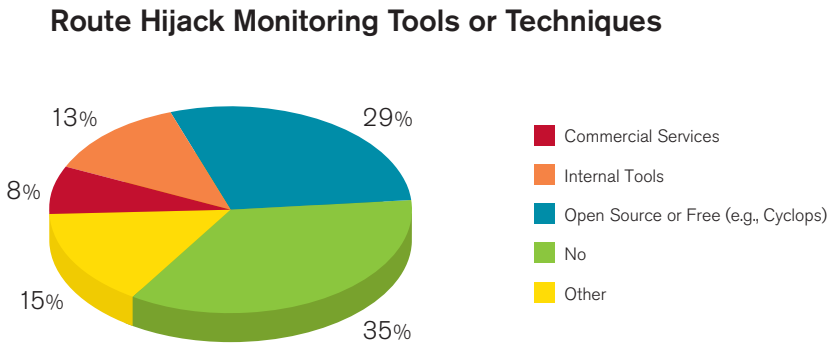

Figure 16: *Route Hijack Monitoring Tools or Techniques*
Source: Arbor Networks, Inc.

Survey respondents were asked if they are familiar with the Secure Inter-domain Routing (SIDR) work of the Internet Engineering Task Force (IETF) and the RPKI system being developed to provide an authoritative cryptographically verifiable database for confirming who owns what address space, and which networks are authorized to originate what prefixes. Most respondents (89 percent) indicated that they are not familiar with the work taking placing around SIDR. Those that are (or are not) familiar offered some of the following comments:

> "No, how can I get involved?"
>
> "Interesting idea, but the devil is in the details with multiple levels of sold address space."
>
> "Yes, seems like a waste to me. Too utopian to ever get implemented correctly without far larger support."
>
> "I do not follow them closely. I think PKI is a must for any security protocol, though key management is always a difficult thing to manage."
>
> "Too complex and takes control away from the operators. So it will only take a little while longer than IPv6 took to begin to be implemented."
>
> "Yes, but I believe it to be an overly complex solution to the problem. If RIRs could force registrants to (keep their data current and) explicitly state what upstream networks they will announce a given block to, a map could be created and hijacks identified."
>
> "Yes, somewhat familiar. I hope the machinery being invented is not too heavy-weight. Prefix-filters generated from IRRs are already cumbersome and the value of the data they provide (accuracy and completeness) is at doubt. RPKI work should improve this situation, not make it worse."
>
> "PKI is crazy and unwieldy. There has to be a better mechanism than that. It was unsuccessful in the enterprise and took over a decade to get traction. Why would I try this in the core? There has to be a better way—more likely an offline service that can be input back into an existing/conventional BGP service."
>
> "I have not been following them closely, but a way to verify route origination would be a great benefit."

More details on the SIDR work group are available.[7]

Finally, we asked respondents if they have experienced any unintentional configuration changes that produced the same effect as a DDoS attack; if they had any outages, traffic loss or performance impacts due to traffic or routing changes with one of their peers; and if they are concerned about monitoring for such changes and threats. Sixty percent of the respondents indicated that they have indeed observed outages or adverse effects from such incidents over the past 12 months alone.

7 www.ietf.org/dyn/wg/charter/sidr-charter.html

Tools for Event Correlation

When respondents were asked if they have tools in place to provide event correlation, 43 percent of respondents indicated they do have such tools or systems in place. However, the tools or systems in use vary widely, including:

- Developed in-house
- Commercial (e.g., ArcSight, NetCool Micromuse, netForensics, Splunk, HP Openview)
- AIRT or other Open Source
- In-house build on FLOSS (freely available community maintained software)
- Mix of commercial and custom
- Basic tools, but nothing automated
- Highly modified version of Open Source Security Information Management (OSSIM)

Several respondents mentioned they are developing or in the evaluation stage at this time, or said that they "wished they had tools or resources along these lines."

ACL Revision Control

We also asked respondents if they have tools to generate, maintain and deploy access control lists (ACLs) to network devices and firewalls. Obviously, these tools could be used to install persistent security policies, respond to DDoS attacks or other security incidents, or equally, be used for routing policy application. Forty-five percent of respondents indicated they have some tool to perform the functions, including:

- In-house, concurrent versions system (CVS), subversion, etc.
- Open source (e.g., Rancid, tool)
- Commercial (e.g., Arbor for DDoS response, FOSS)

Most indicated they use some mixture of in-house developed tools, CVS and open source tools available for most all router and network device configuration management, not just ACLs.

Open Recursive DNS Resolvers

Of the 63 respondents who answered questions regarding open recursive DNS resolvers, some 78 percent indicated that they "Do" restrict access to their recursive resolvers, while 16 percent indicated that they "Do Not." Previous editions of the Worldwide Infrastructure Security Report discuss in more details the benefits, offshoots and considerations of limiting resolver access to off-network resolvers or enabling off-net clients to local resolvers recursively.

On Disclosure

Given a great deal of the activities surrounding the DNS cache poisoning vulnerability and the three-phased disclosure method (pre, partial, full) employed in 2008, we asked what impact network operators believe such techniques have on scanning and attempted exploits. Of the 72 respondents who replied to this line of questioning, some 58 percent (down from 72 percent last year) indicated that they "Do" believe partial disclosures such as those with DNS cache poisoning simply result in more scanning, reverse engineering and exploit activities, while 32 percent said they "Do Not."

Most Limited Vendor Features

When respondents were asked what are the most critical (due to scale, functionality, stability, etc.) or missing vendor security features, a large array of responses were received (all vendor names have been removed):

> "Lack of feature/functionality parity and consistency across platforms."
>
> "Anything and everything to do with IPv6."
>
> "ACL length/granularity limitations."
>
> "Number and complexity of ACLs supported."
>
> "Poor documentation and understanding of security-related features/functionality and BCPs."
>
> "No sensible defaults for control- and management-plane self-protection."
>
> "Inadequate ACL performance at 10gb/sec and above."
>
> "Lack of ASLR, buffer/heap overflow protection, and other 'anti-hacking' measures."
>
> "IPv6 filtering and security capabilities."
>
> "Lack of hardware support for key features."
>
> "Lack of pps-based rate-limiting vs. bps-based rate-limiting."
>
> "High complexity, low usability."

The main themes that seem to resonate are forwarding policy mechanisms (e.g., ACLs), router self-protection capabilities and all things IPv6.

Security Team Characteristics

This section always proves to be a favorite data set for respondents, as it typically helps them demonstrate to management where their staffing and related expectations fall in respect to other organizations.

Composition of Security Teams

Respondents were asked where their network security team resides within their organization. Most indicated that it is either part of Network Engineering (29 percent) or falls within the Network Operations (24 percent) group. While some 8 percent of the respondents said that their network security team is part of IT Security, most of those organizations are not what you might refer to as "traditional" ISPs. Another 5 percent of respondents indicated that their network security team is part of a larger Managed Security Services (MSS) team, and 8 percent of respondents said their network security team is an independent organization.

When respondents were asked about the size of their network security team, responses again varied widely (Figure 17, page 25). Nearly 16 percent of the respondents indicated that their organization has no staff dedicated explicitly to network security, with another 10 percent indicating that just one individual ("me") is responsible. However, 30 percent of the respondents indicated that their network security team consists of 2-4 team members, 11 percent have 5-8 team members, and nearly 15 percent have 9 team members or more. As you might suspect, the "Tier 1s" and teams that are independent or part of a larger MSS team have more human resources dedicated to network security.

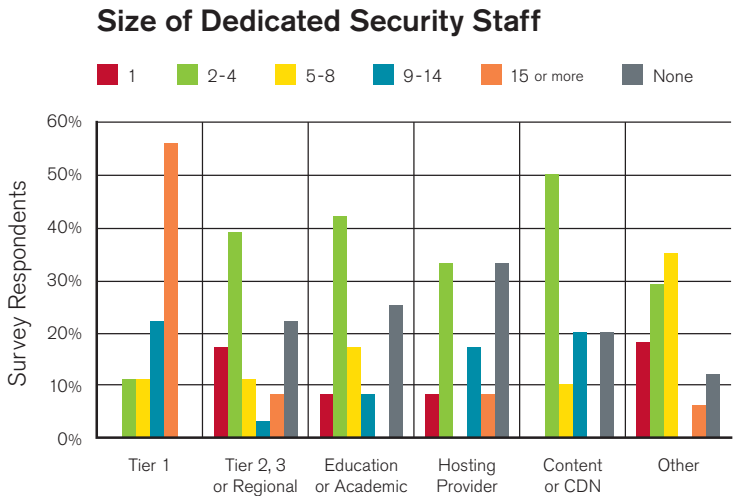

Figure 17: *Size of Dedicated Staff*

Source: Arbor Networks, Inc.

Respondents were also asked if their organization has a dedicated Security Operations Center (SOC), either as part of an MSS offering or for network and infrastructure security services specifically. Only 10 percent of the respondents indicated they have a dedicated SOC as part of an MSS offering, while 17 percent indicated that there is a dedicated SOC in their organization focused on network infrastructure and security operations of network services. Nearly 45 percent of respondents indicated that there is no dedicated SOC within their organization.

Management- and Executive-Level Support

When respondents were asked if they believe their team receives adequate management-level support for security initiatives and projects, only 50 percent of the respondents indicated they do believe they receive adequate support. Respondents were also asked if they believe they receive adequate executive-level support for security projects and initiatives, and only 47 percent of respondents indicated they believe they do.

Given the loose nature of these questions, we will not attempt to form any conclusions based on the responses received, although we suspect that similar response rates would occur for such a question in any industry. We also suspect that the global economic crisis is in part responsible for some underlying financial constraints that tend to exacerbate these concerns.

ISPs: Bots, Botnets, AV and Malware

We asked respondents an array of questions ranging from botnet sizing, to distribution of anti-virus (AV) and malware, to walled garden and quarantine techniques. Some of the data sets returned are clearly more useful than others, but we will share the lot of it here nonetheless. Most of the information in this section is shared as is, with very few author conclusions provided. As with the rest of the survey, it is simply meant to be somewhat representative of the network operator perspective on the issue.

Botnet Activities

Respondents were asked what activities they have personally observed bots performing over the past year (Figure 18). Not surprisingly, spam and DDoS share the top spot, followed by click fraud, ID theft and an array of other nefarious activities.

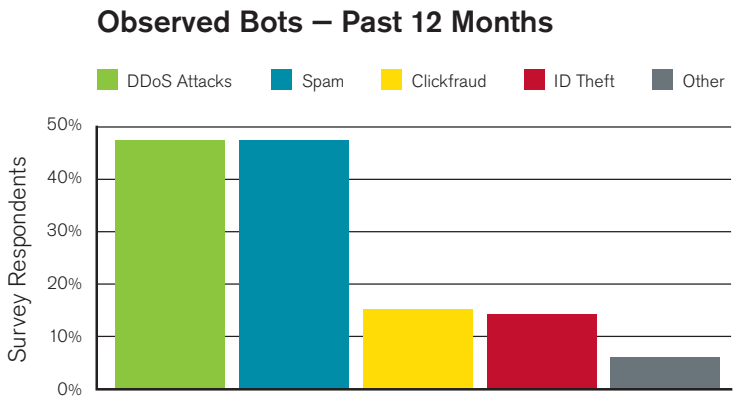

Figure 18: *Observed Bots – Past 12 Months*

Source: Arbor Networks, Inc.

The "Other" category included phishing, drop sites and an array of other malicious activities.

Tracking Botnet Activities

When asked what the most effective tools to detect, measure and monitor botnets, responses were as follows:

- Flow telemetry analysis (e.g., IRC and HTTP C&C detection)
- Rely on CSIRTs, abuse complaints, security groups and data sharing
- Honeypots and darknet monitoring
- Snort with Bleeding-Edge rules
- Collaboration
- Internally developed tools
- DPI at the customer edge

Arbor Networks' Active Threat Level Analysis System (ATLAS®) tool was mentioned several times as well. Information sharing and collaboration were common responses and their value cannot be overemphasized; many of the respondents reiterated this point.

We also asked if respondents believe that detection and monitoring of botnets is a role for which ISPs should be responsible. Forty percent of respondents said "Yes," while 10 percent disagreed and 8 percent responded "Yes, with some criteria." Forty percent of the respondents did not provide an answer.

When respondents were asked how successful they believe anti-botnet tools and techniques have been over the past year, only 10 percent responded they were sufficient, while 43 percent indicated they were insufficient (and 40 percent of the respondents again did not answer this question).

Quarantine, Walled Gardens and Cleanup

Some 28 percent of the respondents surveyed said they "Do" employ automated techniques for quarantine or walled-garden infected or malicious subscribers, while 24 percent said they "Do Not." Twenty-seven percent of all respondents—nearly all who provide automated walled gardens—indicated that their walled garden solutions include some form of notifications (email, text, Web redirect or other) to the customer once they have been quarantined.

When respondents were asked if they offer infected subscribers assistance in cleaning their compromised or infected systems, responses were as follows:

21% – No
18% – Yes, via internal resources, for free
5% – Yes, via internal resources, for a fee
3% – Yes, via third party, for a fee
1% – Yes, via third party, for free
10% – Other

Note: The percentages above are representative of all respondents. Many do not have traditional subscribers and therefore selected "No" when asked this and other questions that are subscriber- or access-related. Furthermore, nearly 42 percent of all respondents did not answer this question.

Interestingly, the respondents who indicated they have automated quarantined and customer notification systems in place seem to be the ones that assist customers with system cleanup via internal resources for free. About one-third of the ISPs that offer walled gardens today seem to have monetized customer cleanup via internal resources.

IPv4 Address Exhaustion and Migration to IPv6

This year we added several survey questions specifically on the impact of imminent IPv4 address exhaustion and the deployment and interoperation with IPv6, as this topic is increasingly in the forefront of discussions among network operators.

When asked whether organizations are preparing for or deploying IPv6 internally and/or for customer-facing services, 42 percent answered "Yes," 11 percent answered negatively and 42 percent did not respond or indicated they are not aware of any activities related to IPv6 within their organization.

When asked about IPv6-related concerns specific to the respondents' operating environment, we received many negative comments:

"Vendors aren't ready. There is no feature parity and there is no CPE to handle it."

"[Lack of] In-depth knowledge of IPv6 by the Ops team."

"Lack of support from management."

"Complexity of dual-stack, training IT personnel."

"The reintroduction of legacy bugs from IPv4 into IPv6. Also, hardware handling of this traffic/ACLs."

"Complexity involved with new address scheme, routing concurrently with v4, Internet peering via v6, cost of allocations from ARIN. Take rate by our customer base. Basically, nothing about IPv6 that I am not concerned about."

"It's mostly done in hardware and many features used to protect our routers today are not implemented in v6 by the vendors yet."

"Vendor support and interoperability. Security components not as sophisticated."

"Not many (any?) devices work on IPv6 the way they do on IPv4, especially with respect to packet headers, flags and speed…Major impact to databases, DNS, provisioning and other systems. Large organizations will feel a major impact in tool-rewriting as well."

"Not all equipment supports IPv6, customers demand is low, IPv6 migration priority is low or stuck because not all equipment supports it, vendor beat-downs are few because customer demand is low."

"Complexity (it's not just the network, OSS, BSS, carrier-grade NAT, etc)."

"Performance issues with our network equipment."

"Lack of flow-based monitoring tools available for IPv6; lack of operational experience dealing with attacks because they are not currently widespread."

"I suspect, and have some anecdotal evidence, that all of the old IPv4 bugs are present in IPv6. IPv6 simply hasn't been 'field-tested' to the same degree as IPv4."

In summary, most of these comments relate to the fact that IPv6 is still viewed as unproven; that there is a lack of IPv6 tools and knowledge in operations; that IPv6 network infrastructure functionality lacks parity with IPv4; and that management does not understand the need to invest in preparation for IPv6 interoperation and support.

When asked about concerns regarding IPv4 exhaustion, IPv6 readiness, RIR policies or other issues surrounding IPv4 address depletion, we received overwhelmingly negative comments, such as those listed below:

"Organizations should have been IPv6-ready and enabled a long time ago; there will be chaos when space actually runs out because people didn't believe it and/or procrastinated."

"IPv6 is not widely accepted/developed by Internet communities and application developers."

"Mobile enterprises require a lot of IPs, which will cause an imminent exhaustion by the end of 2010."

"I'm very concerned that the current model (one that we participate in) is unsustainable for the long term."

"IPv6 still seems to be in flux, IPv4 is about gone, yet no one is trying to reclaim huge chunks of over-allocated space to certain institutions."

"Yes, but I have concerns about IPv6 functionality, too."

"RIR policies need to be tightened, and more operators should start to consistently use an authoritative routing database."

"My fear is that ARIN (can't speak for the others) is trying to insert itself into the middle of IPv4 transactions (not just allocations from them to a user) and that there is a huge legal risk to this approach. My second fear is that no one is ready for IPv6, including the users and vendors. So, when it *has* to be deployed, people are going to be winging it with strict timelines—a recipe for network instability and unforeseen consequences."

"Major change creating new gaps and OSS systems not ready (nor certain network gear vendors). Possibly more infrastructure attacks once network IPv6 enabled."

"Not a fan of IPv6 at all."

"IPv6 is not ready and never will be."

Clearly, while there appears to be many negative perceptions regarding IPv6, operators have a great deal of work ahead of them to prepare for its imminent deployment over the next year or two. Whether folks like it or not, we believe "the ship has sailed" and it is time for folks to take stock of what they need to do to accommodate the associated new requirements IPv6 imposes in their operating environments.

Additional Questions and Miscellaneous Information

The following items were included at the request of industry survey questionnaire reviewers and/or previous survey respondents.

- When respondents were asked who should be responsible for covering losses associated with banking and related financial credential theft, 33 percent of respondents said "banks," 4 percent said "consumers or individuals" and 56 percent did not respond.

- Over 28 percent of respondents said they believe that AV and host security protection applications are not keeping up with Internet security threats, while 14 percent disagreed. Forty-five percent did not respond.

- Given the media surrounding the Russian Business Network (RBN) and alleged sponsorship of illegal activities, we asked respondents to indicate whether they believe hosting providers should be held accountable for selling services to known "shady" organizations. Eight percent said "No," while 31 percent said "Yes" and 54 percent provided no answer. Furthermore, 30 percent of respondents indicated that there are specific ASNs or networks on the Internet that they consider "evil" and would simply like to see them go away. Seventy percent provided no response.

- When respondents were asked whether they see the scale and frequency of security threats increasing or decreasing as IPv6 becomes more widely deployed, 27 percent believe threats will increase, while only 4 percent of respondents believe threats will decrease, and 45 percent provided no response.

- Regarding respondents who have deployed data retention methods, 18 percent said they are based on logs only (i.e., smtp, pop/imap, login/logout, etc.), 2 percent said they are based on "flow data," 31 percent said they are based on "both logs and flow data" and 3 percent said they are based on "other techniques."

- When respondents were asked what trends they have observed in network security over the past year, and how they spend their time on a daily (or nightly) basis as a result, responses were as follows:

  "Very little has changed."

  "Increasingly about money."

  "Attackers getting smarter, customer base clue decreasing."

  "SP management largely unconcerned with security unless/until a major event occurs."

  "Fewer clueful resources."

  "More high-quality phishing scams."

  "More effective attack methodologies."

  "Web 2.0 attacks on the increase."

  "Social networking and tech marketing practices give phishers a huge edge."

  "Increases in threat data/visibility translate into more work for security/abuse team, with the same or fewer resources available."

  "Increased interest in more automatic defenses."

  "More compromised customer hosts equates to more outbound DDoS."

  "Increased leverage of brute-force scanning and search engines to locate software vulnerabilities."

  "More ISP-centric phishing; more professional phishes with improved grammar and spelling."

  "Attackers have learned from mitigation efforts; so many attacks, one has to pick and choose which ones to defend against, numbers are overwhelming."

  "More scanning/probing of network infrastructure."

  "Increased demands from banks/financial industry."

- Respondents were also asked if they have deployed any DPI equipment, and if so, for what purpose. Twenty-eight percent said "No," they have not deployed any, while 6 percent said "Yes," for lawful intercept, 8 percent for service enforcement and protection, 6 percent for both lawful intercept and service enforcement and protection, and 7 percent responded with "Other." Forty-five percent provided no response.

Additional Information, Free-Form Comments

At the end of the survey, respondents were asked for any additional comments or anecdotal information they might want to see reflected in the survey report. Some of the input in that section includes:

> "4-byte ASN represents a major concern."
>
> "It would be nice to see questions further geared to enterprises as well as ISPs."
>
> "Fewer truly clueful security resources allowed and/or willing to participate in [inter-organizational] security groups, which means many security groups are filled with clueless management types."
>
> "Cooperation between ISPs, governments and equipment vendors is the most effective way to solve the problem."
>
> "The bad guys are beating us, badly."

Conclusions

This year's survey finds the Internet engineering and security community struggling with the rapid evolution of complex security challenges. While peak DDoS attack rates did not exceed the 2007 fears of 60-80 Gbps (see last year's survey),[8] providers report that gigabit attacks are now commonplace. Worse, the growing complexity of cloud and distributed infrastructure significantly increases the exposed vulnerability surface area of customer-visible services.

While attackers increasingly and successfully monetize DDoS, phishing and other illegal activities, many providers report struggles with budget and management support for security initiatives and investment. Any ISP optimism about security issues has been replaced by growing concern over a range of new threats, including DNS poisoning, route hijacking and service-level attacks. ISPs describe a double-edged struggle as they face increased cost and revenue pressure, along with attacks that are growing in size, frequency and sophistication.

Though a few providers believe they still have a technical advantage against attackers, this year's survey in part reflects a new general pessimism, articulated by one provider as, "The bad guys are beating us, badly." Many providers believe future Internet security challenges will require increased cooperation and coordination among providers and vendors.

Finally, the challenges represented by the "perfect storm" of IPv4 address exhaustion, IPv6 deployment, DNSSEC deployment, and 4-byte ASN support are a source of concern from an architectural, operational and security standpoint for many network operators. Taken both individually and collectively, the implementation of these technologies will undoubtedly alter the operational security posture of Internet-connected networks in many ways, and special care should be taken to ensure that organizations take all possible steps to maintain their operational security postures during this time of great change for the global operational community.

8 www.arbornetworks.com/report

About the Authors

Danny McPherson, Vice President and Chief Security Officer, Arbor Networks

danny@arbor.net

With nearly 20 years experience in the Internet network operations, security and telecommunications industries, Danny McPherson brings extensive technical leadership to Arbor Networks. Today he is a main contributor to the company's industry activities, overall strategy and product architecture. Prior to joining Arbor, he was with Amber Networks, and previously held network operations and architecture positions for nearly a decade at internetMCI, Genuity (acquired by GTE Internetworking), Qwest Communications and the U.S. Army Signal Corp.

Danny has been an active participant in Internet standardization since 1996. Currently he is a member of the Internet Architecture Board (IAB) and co-chairs the IETF's L3VPN WG. He also serves on the ICANN Security and Stability Advisory Council (SSAC), the FCC's Network Reliability and Interoperability Council (NRIC), and is quite active in the network and security operations and research communities.

Danny has authored a significant number of books, Internet protocol standards, network and security research papers, and other documents related to Internet routing protocols, network security, Internet addressing and network operations.

Roland Dobbins, Solutions Architect, Arbor Networks

rdobbins@arbor.net

Roland Dobbins has nearly 25 years of operational experience in the service provider (SP) and large enterprise arenas, designing, deploying, operating, securing, maintaining, troubleshooting and defending many of the highest-visibility networks in the world. He is a recognized industry leader in the fields of operational security (opsec) and network telemetry, and has an extensive background in security product/feature innovation, devising operational security requirements for network infrastructure devices and protocol design. His focus is on extending the availability, scalability and security of the network infrastructure and the applications/services it enables, with an emphasis on flexible and resilient global service delivery capabilities.

Michael Hollyman, Manager of Consulting Engineering, Arbor Networks

mhollyman@arbor.net

With more than 12 years in the network, security and telecommunications industries, Mike Hollyman brings extensive knowledge of service provider and large enterprise network design and security to Arbor. He provides leadership to the Arbor sales organization through his management of the company's Consulting Engineering team for North American service providers. Prior to joining Arbor, Mike was a network and security consultant, both independently and through his own consulting company. Prior to consulting, he worked as a network engineer for OneSecure, Qwest Communications and the University of Illinois.

Dr. Craig Labovitz, Chief Scientist, Arbor Networks

labovitz@arbor.net

Craig Labovitz brings extensive experience in network engineering and research to Arbor Networks. Before joining Arbor, he served as a network researcher and scientist for the Microsoft Corporation. Previously, he spent nine years with Merit Network, Inc. and the University of Michigan as a senior backbone engineer and director of the Research and Emerging Technologies group. His work at Merit included design and engineering on the NSFNet backbone and Routing Arbiter projects. He also served as the director of several multimillion dollar grants from the National Science Foundation for network architecture and routing protocol research. Dr. Labovitz received his Ph.D. and MSE from the University of Michigan.

Dr. Jose Nazario, Manager of Security Research, Arbor Networks

jose@arbor.net

Jose Nazario is senior security researcher within the office of the CTO at Arbor Networks. In this capacity, he is responsible for analyzing burgeoning Internet security threats, reverse engineering malicious code, managing software development and developing security mechanisms that are distributed to Arbor Peakflow platforms via Arbor's Active Threat Feed (ATF) threat detection service. Dr. Nazario's research interests include large-scale Internet trends such as reachability and topology measurement; Internet-scale events such as DDoS attacks, botnets and worms; source code analysis tools; and data mining. He is the author of the books "Defense and Detection Strategies against Internet Worms" and "Secure Architectures with OpenBSD." He earned a Ph.D. in biochemistry from Case Western Reserve University in 2002. Prior to joining Arbor Networks, he was an independent security consultant. Dr. Nazario regularly speaks at conferences worldwide, with past presentations at CanSecWest, PacSec, Blackhat and NANOG. He also maintains WormBlog.com, a site devoted to studying worm detection and defense research.

Corporate Headquarters
6 Omni Way
Chelmsford, Massachusetts 01824
Toll Free USA +1 866 212 7267
T +1 978 703 6600
F +1 978 250 1905

Europe
T-+44 208 622 3108

Asia Pacific
T +65 6299 0695

www.arbornetworks.com

WISR/EN/0110

About Arbor Networks

Arbor Networks® is a leading provider of secure service control solutions for global business networks. Its customers include a majority of the world's ISPs and many large enterprises. Arbor solutions deliver best-in-class network security and visibility, along with the power to improve profitability by deploying differentiated, revenue-generating services. By employing flow-based and deep packet inspection (DPI) technologies, Arbor solutions measure and protect the entire network—from the network core to the broadband edge. Arbor also maintains the world's first globally scoped threat analysis network—ATLAS—which uses technology embedded in the world's largest ISP networks to sense and report on comprehensive worldwide threat intelligence.