VOLUME IV

WORLDWIDE INFRASTRUCTURE SECURITY REPORT

OCTOBER 2008

ARBOR®
NETWORKS

Table of Contents

List of Figures

Overview

Arbor Networks, Inc., in cooperation with the Internet security operations community, has completed this fourth edition of an ongoing series of annual operational security surveys. This survey, covering a 12-month period from August 2007 through July 2008, is designed to provide data useful to network operators so that they can make informed decisions about their use of network security technology to protect their mission-critical infrastructures. It is also meant 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

The ISP Security Battlefront Expands

In the last three surveys, ISPs reportedly spent most of their available security resources combating distributed denial of service (DDoS) attacks. For the first time, this year ISPs also describe a far more diversified security landscape, including significant concerns over domain name system (DNS) spoofing, border gateway protocol (BGP) hijacking and spam. Almost half of the surveyed ISPs now consider their DNS services vulnerable. Others expressed concern over related service delivery infrastructure, including voice over IP (VoIP), session border controllers (SBCs) and load balancers.

Attacks Now Exceed 40 Gigabits

From relatively humble megabit beginnings in 2000, the largest DDoS attacks have now grown a hundredfold to break the 40 gigabit barrier this year. The growth in attack size continues to significantly outpace the corresponding increase in underlying transmission speed and ISP infrastructure investment. Figure 1 shows the yearly reported maximum attack size.

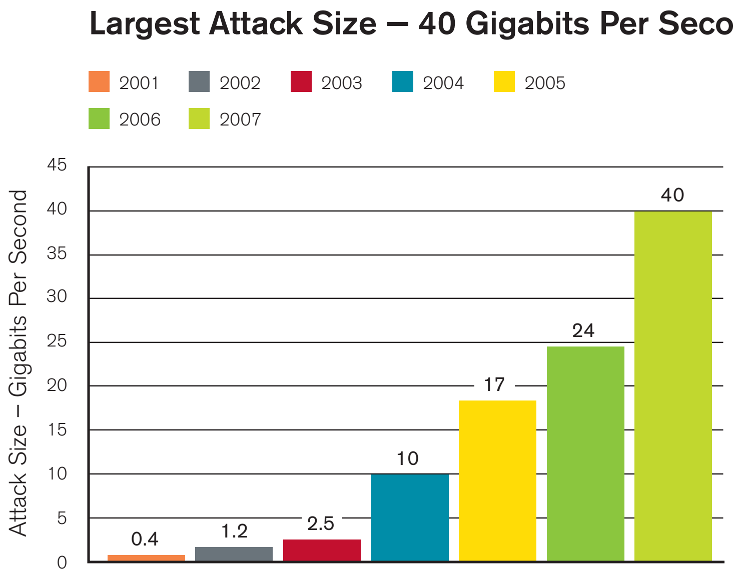

Figure 1: *Largest Attack Size – 40 Gigabits Per Second*

Source: Arbor Networks, Inc.

Services Under Threat

Over half of the surveyed providers reported growth in sophisticated service-level attacks at moderate and low bandwidth levels—attacks specifically designed to exploit knowledge of service weakness 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.

Fighting Back

The majority of ISPs now report that they can detect DDoS attacks using commercial or open source tools. This year also shows significant adoption of inline mitigation infrastructure and a migration away from less discriminate techniques like blocking all customer traffic (including legitimate traffic) via routing announcements. Many ISPs also report deploying walled garden and quarantine infrastructure to combat botnets.

Survey Methodology

This edition of the survey consisted of 90 free-form and multiple-choice questions, representing an array of issues facing network security operators today. Questions addressed topics such as threats against backbone infrastructure and individual customers, techniques employed to protect network infrastructure itself, and mechanisms used to detect and respond to security incidents.

All data represented here is provided 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 weigh 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 to previous surveys. Some questions were deleted, some added and many simply honed in an attempt to capture more of the desired data sets. Several of the new questions were added verbatim as provided by respondents during a previous survey, or as a result of direct feedback from one of the many polled network security or operations forums.

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 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.

Demographics of Survey Respondents

Survey participants included 66 self-classified Tier 1, Tier 2 and other IP network operators from North America, South America, Europe and Asia. All survey participants are directly involved in network security operations at their respective organizations. While the number of survey respondents dropped slightly from 70 to 66 this year, the geographic diversity of respondents increased.

As illustrated in Figure 2 (page 5), this year's survey respondent pool was once again dominated by self-classified Tier 1 and Tier 2 organizations, which accounted for nearly half of all responses. Content, hosting and academic networks represented a large percentage of the response base as well. The "Other" category was composed of government, wireless and voice ISPs, some regional network providers and Internet exchange point (IXP) operators.

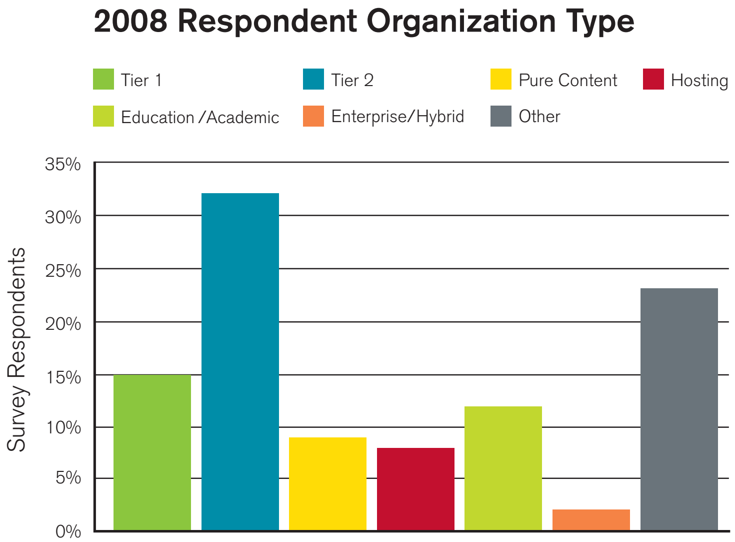

Figure 2: *2008 Respondent Organization Type*

Source: Arbor Networks, Inc.

A new question was added to this year's survey to determine the primary geographic region within which the respondent's organization operates. While this question was intended to provide insight into and illustrate the geographic diversity of the respondent pool, we quickly discovered that adding a "Global" option within the available choices was undesirable. That is, most of the Tier 1 and Tier 2 respondents, as well as several others, believe that their primary region of operation is "Global," thereby skewing the results we intended to illustrate. Figure 3 depicts the geographic distribution among 2008 survey respondents.

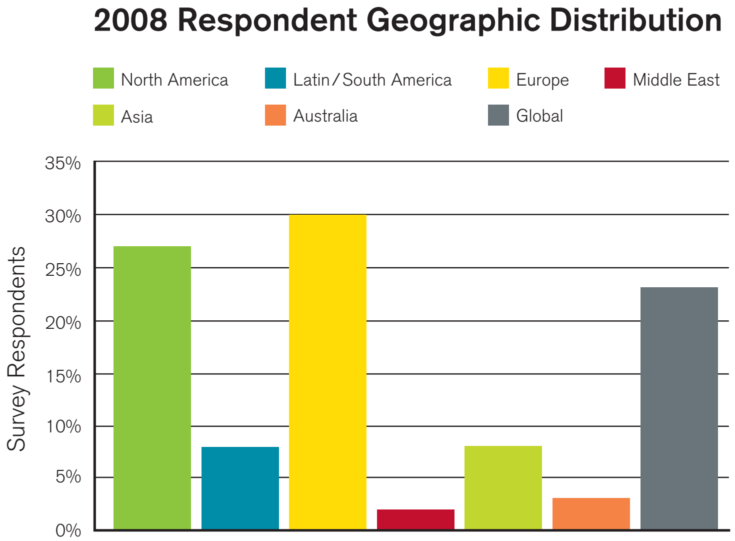

Figure 3: *2008 Respondent Geographic Distribution*

Source: Arbor Networks, Inc.

Post-analysis of the respondents' geographic location indicates that most respondents selecting "Global" as the primary geographic region were from North America, with Europe being a close second. Nonetheless, we believe this data is successful at highlighting the global representation afforded by a geographically diverse survey respondent pool.

Most Significant Operational Threats

When asked to rank threats that they believe would pose the largest problems over the next 12 months, bots and botnets again took the top spot, followed closely by DNS cache poisoning and BGP route hijacking (Figure 4).

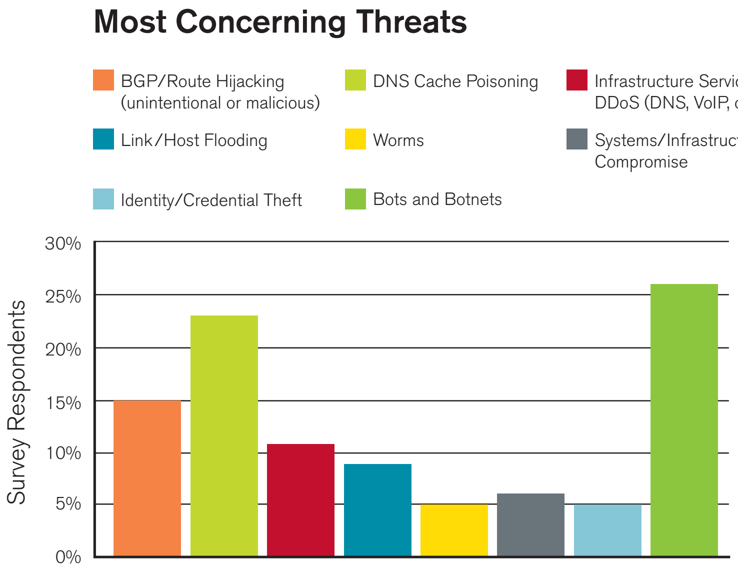

***Figure 4:** Most Concerning Threats*

Source: Arbor Networks, Inc.

As discussed later in this report, growth of the largest botnets continues to outpace containment efforts and infrastructure investment. Despite this trend, bots/botnets fell slightly in the concern rankings from 29 percent last year as providers now face a broader range of threats.

One growing attack vector is DNS cache poisoning. Techniques considerably enhancing vulnerability to, and effectiveness of cache poisoning were disclosed in mid-2008 by a security researcher named Dan Kaminsky. As a result, DNS cache poisoning gained significant attention in July of this year, shortly before the start of the survey.[1] This year DNS/infrastructure attacks hold spot two in the concern rankings—up from number three a year ago at 19 percent. Though providers reported a number of serious DNS attacks in the last few months, aggressive patching and widespread publicity appear to have thwarted some of the worst-case scenarios predicted in the trade press.[2]

Concern over DDoS flooding of links and hosts fell in the rankings from 24 percent last year to 11 percent this year, likely reflecting the increased ability of ISPs to mitigate these types of attacks. ISPs report that improved communication, collaboration and previous infrastructure investment have significantly advanced their capability to protect network services.

Representing a significant change in concern rankings, BGP route hijacking jumped from the lowest spot last year at 1 percent to the third most serious ISP concern. BGP hijacking also received significant media attention following several high-profile outages, such as the "apparent" misconfiguration of YouTube routes by Pakistan Telecom[3] and newly announced routing system monitoring and proof-of-concept tools for anonymous route hijacking.[4,5,6]

[1] www.us-cert.gov/cas/techalerts/TA08-190B.html
[2] www.securityfocus.com/news/11526
[3] http://asert.arbornetworks.com/2008/02/internet-routing-insecuritypakistan-nukes-youtube
[4] BGPMon: http://bgpmon.net
[5] Watch My NET: www.watchmy.net
[6] http://blog.wired.com/27bstroke6/2008/08/revealed-the-in.html

Finally, worms and identity/credential theft came in last at 5 percent each (compared to 9 and 11 percent respectively last year). We were somewhat surprised that worms garnered any attention at all in the current survey; perhaps their ranking is due to the residual pain from Slammer and other network-impacting worms of yesteryear.

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
- Ongoing (Constant) Security Events
- Peer-to-Peer (P2P)
- Spam
- Other

Somewhat surprisingly, spam (both inbound and outbound) took the top spot at nearly 32 percent, followed closely by security events from constant background activity (i.e., scans, worms, etc.) at 27 percent. DDoS attacks came in at 21 percent, closely in-line with its position last year. Note that spam was not an available option in previous versions of the survey, and therefore, no historical information exists regarding how respondents viewed its level of resource consumption in the past.

Nearly 5 percent of respondents indicated that P2P (i.e., its impact on operations, law enforcement engagement, etc.) consumed the largest amount of operational resources. Unfortunately, we were not more explicit in this line of questions and did not provide for free-form entry; therefore, it is difficult to determine whether these respondents were referring to basic traffic management issues, copyright and related issues, or something else.

Another 15 percent of respondents indicated that something else ("Other") consumed the largest amount of operational resources. Again, given the lack of a free-form entry option, we were unable to capture just what these other operational issues are.

Scale and Effectiveness of Attacks

Respondents were asked again this year about the scale of the largest attacks either they or their customers have endured. Figure 5 illustrates their responses. Nearly twice as many respondents reported observing attacks in the 1 to 4 gigabit range this year as opposed to 2007. Last year 36 percent of survey respondents had reported observing attacks larger than 1 gigabit. This year, that number has nearly doubled, with 57 percent of respondents reporting attacks larger than 1 gigabit.

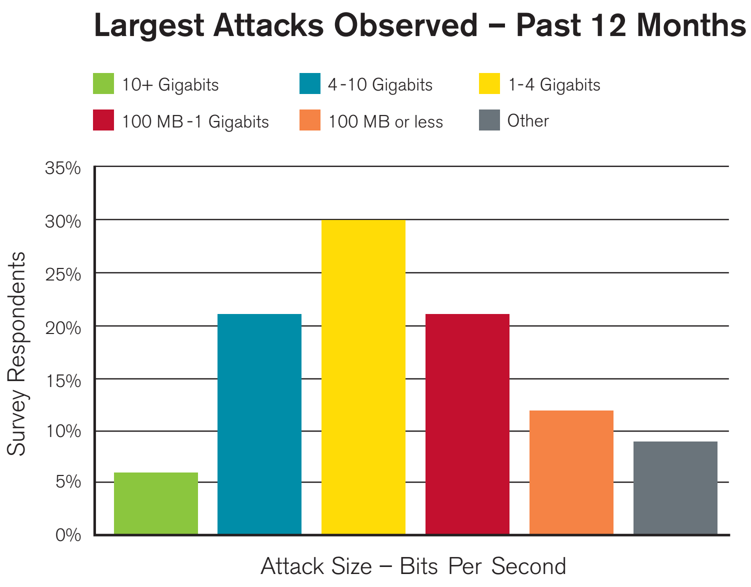

***Figure 5:** Largest Attacks Observed – Past 12 Months*

Source: Arbor Networks, Inc.

Comparing the last two surveys, about the same percentage (6 percent) of respondents reported attacks over 10 gigabits, with the largest attack reported this year being "just north of 40 gigabits." The largest sustained attacks reported in the previous two surveys were 24 gigabits and 17 gigabits respectively. This represents a 67 percent increase in attack scale over last year, an increase of nearly 2.5 times the size of the largest attack reported in 2006, and a 100-fold increase since 2001. Figure 1 (page 3) illustrates the continued growth in attack scale since 2001.

When asked for details of the largest reported attack, the respondents provided a considerable amount of information, as shown in the direct quotes below:

> "This was, initially, criminal-on-criminal crime though obviously the greatest damage was inflicted on the infrastructure used by the criminals."
>
> "The attack exceeded 40 gigabits in aggregate at one time."
>
> "The attack used DNS amplification."
>
> "The attack stopped only because the attacker was paid."
>
> "The attacker remains at large and active."
>
> "No bots were used in this attack. The attacker had a small number of compromised Linux boxes from which he'd launch the spoofed source DNS query. The DNS servers were all DNS servers open to recursion."

Reflective amplification attacks responsible for the largest attacks during 2006 and 2007 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., open DNS resolvers), which respond to the target with the large payloads. Attacks of this nature employed against Internet root and top-level domain (TLD) name servers in early 2006 achieved a 76-fold amplification factor. The largest reported attacks this year 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 16)—surveying where they are implemented and why anti-spoofing mechanisms, which would largely mitigate this type of attack, are not universally deployed today.

As described last year, most individual core Internet backbone links today are no larger than 10 gigabits per second, with standardization efforts and some networking products coming to market to enable 40 or 100 gigabit IP links. As such, most of the larger attacks today still easily inflict collateral damage on infrastructure that is upstream from the 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 gigabit 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.

Attack Vectors

Respondents were asked what attack vector was employed for the largest attack they observed over the past 12 months. Their responses are summarized in Figure 6.

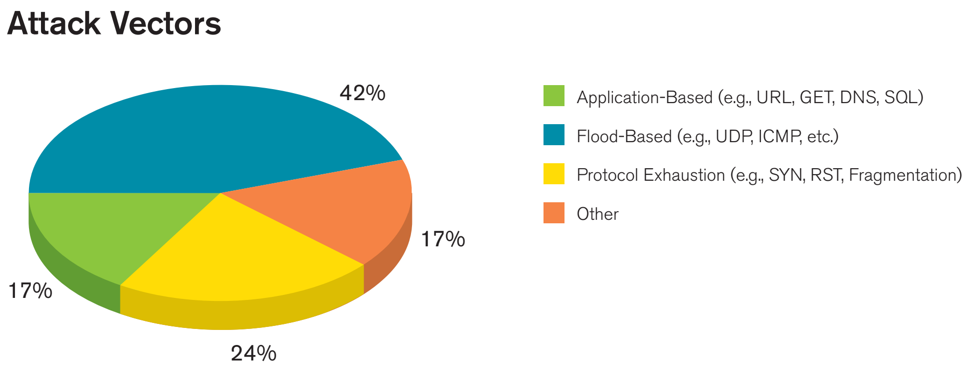

Figure 6: *Attack Vectors*

Source: Arbor Networks, Inc.

Protocol exhaustion and flood-based attacks are still the most predominant. Many of this year's respondents mentioned that they are seeing an increase in application-based attacks aimed expressly at triggering back-end transaction activity and resource state, and that these attacks, while not the largest, were certainly the most sophisticated and devastating attacks they had observed over the past year.

When respondents were asked if they have observed any trends in attacks moving from brute force to more sophisticated attacks over the past year, about half of the respondents said they have observed such trends, though few details were provided. One-fourth of the respondents indicated that they have observed attacks that were more sophisticated in the manner in which they targeted network services, either as second-stage attacks or attacks aimed at impacting other systems by affecting adjacent network services (i.e., DNS, load balancers, VoIP systems, routers, etc.).

Frequency of Attacks

Actionable attack frequency remained somewhat constant on aggregate 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.

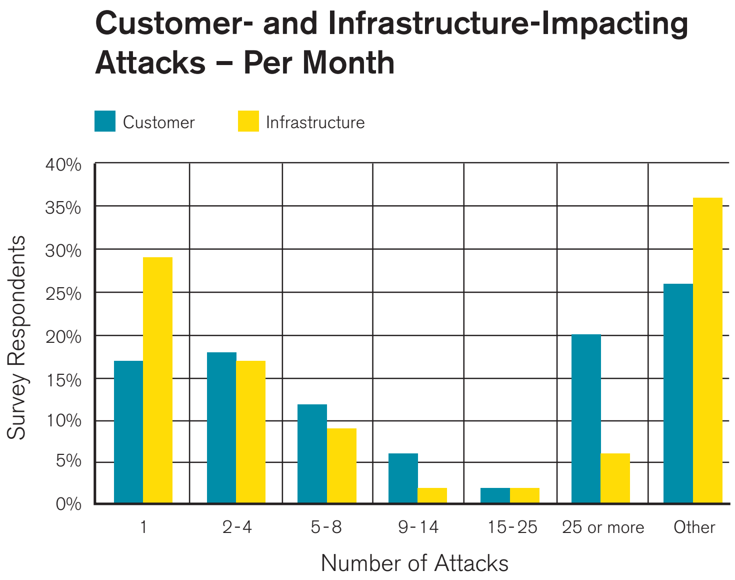

Figure 7: *Customer- and Infrastructure-Impacting Attacks – Per Month*

Source: Arbor Networks, Inc.

Respondents were asked what attack vectors have been employed when attacks target their infrastructure. An array of responses was received, including the following:

> “Brute force attacks, floods targeting router interfaces and network elements.”
>
> “Know vulnerability probes.”
>
> “Heavy VoIP scans—on the increase recently.”
>
> “Application-based attacks targeting network services.”
>
> “Lots of SSH [secure shell] brute force login attempts to all network elements and management systems.”
>
> “Multi-mode attacks, from SYN [synchronized] and DNS, evolving to application-layer attacks targeting customer-facing systems and network elements.”
>
> “Lots of attacks towards customers, resulting in collateral damage to infrastructure.”
>
> “More scans targeting router OS vulnerabilities expressly.”
>
> “Often getting DDoS directly as a result of mitigating attacks for customers.”

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

61%	External Brute-Force Attacks
12%	Known Vulnerability
3%	Social Engineering
3%	Misconfiguration
2%	Insider Threat
20%	Other

External brute-force attacks seem to remain high, while insider threats are still quite low. When asked about attacks towards customers in the past, little to no variation exists across attack vectors.

Attack Detection and Traceback

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

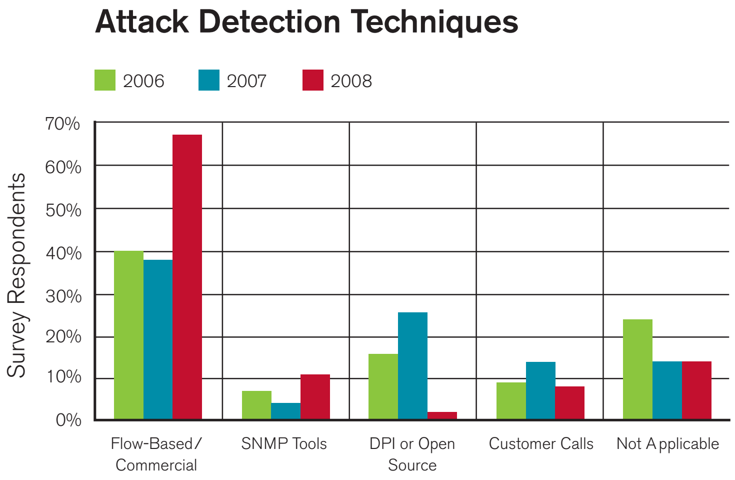

***Figure 8:** Attack Detection Techniques*

Source: Arbor Networks, Inc.

A considerable increase in flow-based tools that can help enable attack detection was observed this year over previous years. A decrease in the number of reported "open source" tools was also observed, perhaps because many of the hybrid and smaller network types were not as well represented in this edition of the survey as previous editions. There was also a considerable drop in the number of respondents who rely on a customer call to serve as the attack detection mechanism (from 14 percent to 8 percent); this can only be considered an improvement.

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. This should obviously be considered an integral part of the attack incident response, especially given the scale of attacks that have emerged over the past several years.

Seventy percent of respondents indicated that they use flow-based tools to trace attacks back to network ingress interfaces. Another 12 percent said they use SNMP-based tools, 3 percent reported using DPI or other tools, and 15 percent indicated they have no current solutions or tools to trace attacks back to network ingress.

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 decreased from 47 percent last year to 30 percent this year (Figure 9). Source and destination-based BGP real-time blackholing (RTBH) also decreased considerably. Intelligent filtering or "scrubbing"—for example, via the Arbor Peakflow® SP Threat Management System (Peakflow SP TMS) device—increased from 8 percent to 14 percent over the past year, providing more effective mitigation techniques without effectively completing an attack. This increase is likely due to more network operators offering DDoS detection and mitigation services, as discussed in later sections.

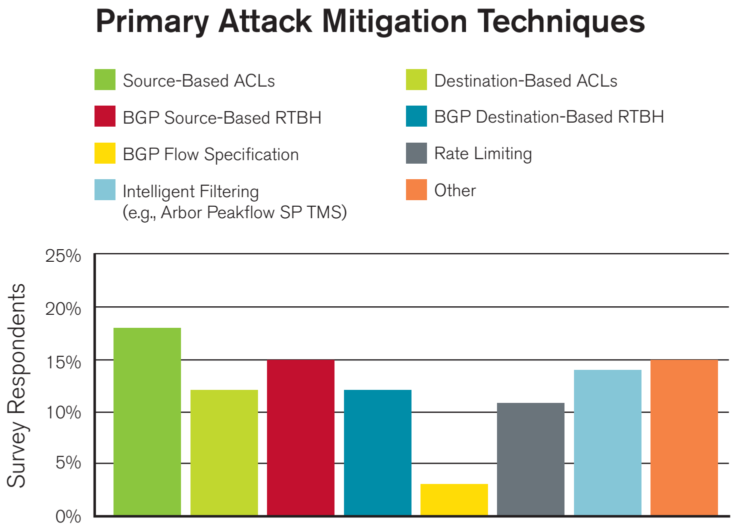

***Figure 9:** Primary Mitigation Techniques*

Source: Arbor Networks, Inc.

For more information on each of the above mitigation techniques, see the *Attack Mitigation* section of the 2007 edition of this survey.[7]

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. Nearly 15 percent of respondents said they typically mitigate attacks in less than 10 minutes once they have been detected, while another 15 percent said less than 20 minutes, and some 14 percent said less than 30 minutes. About 26 percent of respondents indicated that it takes about an hour to mitigate an attack once it is detected. Surprisingly, over 30 percent of respondents reported needing more than an hour to mitigate an attack, even after it has been detected.

[7] www.arbornetworks.com/report

When explaining the primary obstacle in reducing time to mitigation, respondents provided an array of input. Some responses varied widely, although common themes did indeed emerge. Some of the challenges cited in mitigating attacks include:

> "Coordinating with upstreams to filter attack flows, inattentive NOC/SOC [network operations center/security operations center] at peers/upstream."
>
> "Delays in internal escalation to capable staff, senior management authorization."
>
> "Cooperation from internal operations teams for internal policy deployment."
>
> "Internal manpower, resources."
>
> "Manual generation, configuration and deployment of mitigation policies in network."
>
> "Poor quality of data/tools in-house, identification and classification after detection."
>
> "Sometimes difficulty in verifying that it is not a flash crowd or otherwise legitimate customer traffic."
>
> "Ensuring no unintended outage is triggered by mitigation, ability to effectively mitigate without impacting production traffic, esp. with emerging application-layer attacks."
>
> "Taking target offline may deflect attack to another target—dealing with mitigation impact scenarios."
>
> "Budget for infrastructure to surgically mitigate attacks."
>
> "Managing capacity of dedicated mitigation devices."
>
> "Sometime[s] with or without managed services, customer permission to mitigate [an] attack is required before any action can be taken."

The most often referenced obstacles to reducing attack mitigation time included: 1) accurately identifying and separating attack flows from legitimate traffic; 2) communication with upstreams, customers and internal staff; and 3) internal resources and manpower to mitigate attacks. In addition, respondents raised concerns about attack mitigation options and tools to enable best scenario mitigation for a given incident.

Customer-Facing 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 has been born—purely out of necessity. Many organizations generate a majority or oftentimes all of their revenue from Web or other network service transactions—making their Internet presence and availability 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 (page 14) illustrates the number of respondents who currently offer DDoS detection and/or mitigation services to customers. It also shows the number of respondents offering a new service category—traffic visibility across the Internet, Multiprotocol Label Switching virtual private networks (MPLS VPNs) and other types of networks.

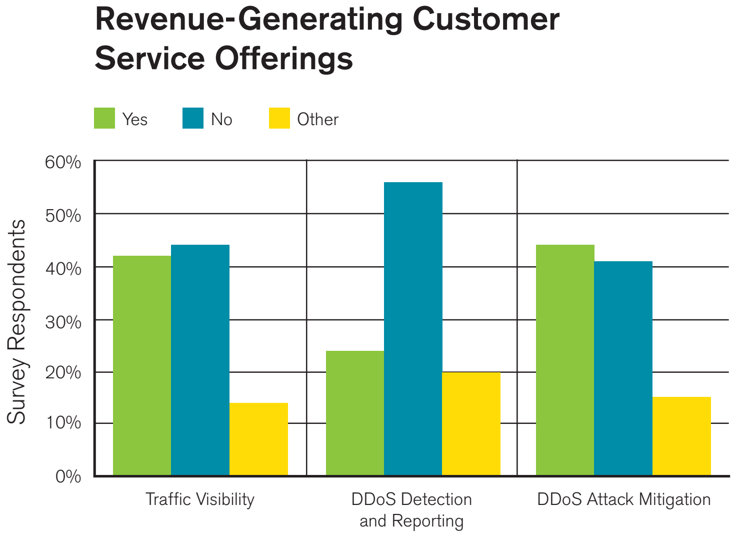

Figure 10: *Revenue-Generating Customer Service Offerings*

Source: Arbor Networks, Inc.

The vast 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 institutions and other organizations that deployed their own scrubbing infrastructure and/or subscribe to MSS offerings from their network services provider(s). Interestingly, many ISPs indicated that they offer traffic visibility services already, seemingly in parallel with DDoS attack mitigation services. This might be largely attributable to the fact that market-leading solutions for DDoS detection and mitigation provide *traffic reporting* and visibility capabilities as well. Therefore, *traffic visibility* provides an easily obtained source of incremental revenue with little or no additional capital expenditures or network infrastructure changes. Only 12 percent of the respondents indicated that they currently provide any type of service level agreement (SLA) for DDoS attack protection services.

Law Enforcement, CERTs and CSIRTS

When respondents were asked how many attacks they have referred to law enforcement over the past years, responses varied widely based on organizational type. Figure 11 (page 15) summarizes how respondents addressed this question this year. Overall, law enforcement referrals dropped for the third year in a row, with 58 percent of ISPs saying they had referred no incidents to law enforcement over the past year (compared with 50 percent last year). Thirty percent said they had only referred five or fewer incidents to law enforcement.

When asked what limited the number of referred attacks, responses were fairly well distributed:

29% – Law enforcement has limited capabilities
26% – Expect customers to report, will not report on their behalf
17% – Little or no utility of reporting
29% – Other

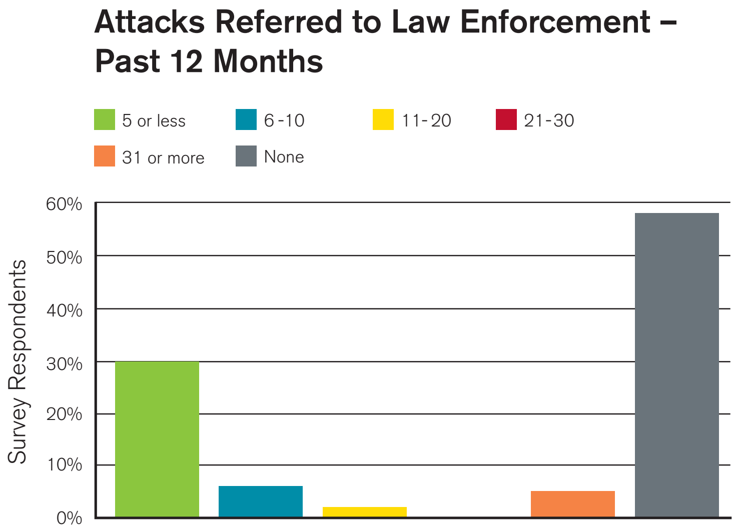

Figure 11: *Attacks Referred to Law Enforcement – Past 12 Months*

Source: Arbor Networks, Inc.

We also asked respondents if they believe law enforcement has the power and/or means to act upon information provided by network operators. Only 21 percent said Yes, while nearly 64 percent said No—indicating that most respondents do not believe law enforcement has the means to act upon information provided about attacks or other security incidents. However, 36 percent of respondents did indicate that they believe law enforcement is becoming more useful to Internet security operations, while only 9 percent said they believe law enforcement was becoming less useful. Of course, 26 percent said they have seen no noticeable change in law enforcement activities, while another 26 percent said “What law enforcement?” (meaning that they see very little law enforcement 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, while 55 percent responded No. Correspondingly (and not surprising), 45 percent of respondents indicated that they work frequently with a government or national CERT or CSIRT. Interestingly, 77 percent indicated that they believe government CERTS/CSIRTS do have a role and responsibility in operational security.

Some 27 percent of respondents indicated that they believe governments have a role in enabling and assisting in infrastructure protection, but that they fail because of a lack of knowledge. Another 18 percent said this failure stems from a lack of cooperation with network operators, while 15 percent said the failure is due to lack of regulation, policy or legislation. Nearly 23 percent said governments fail to enable infrastructure protection because they are slow and far too political, while 11 percent said they seem to be doing a decent job.

Infrastructure Protection Techniques

In this edition of the survey, we again asked questions regarding several well-known infrastructure security techniques for mitigation of spoofing and protecting 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 garnered nearly half the votes at 42 percent. Both routers and VoIP SBCs received 20 percent each, and load balancers received 8 percent. This might simply be due to the fact that the survey coincided with the disclosure of new DNS cache poisoning vulnerabilities, or it may be a larger indication of the insecurities of the Internet DNS.

When respondents were asked if they have any tools or techniques to monitor for threats against DNS resolvers and/or authoritative name servers, only 39 percent indicated they do have tools in place for detecting such threats. Only 21 percent reported that they have tools in place to detect threats against VoIP infrastructure or services, and of those, all noted that they have solutions in place to mitigate threats against VoIP infrastructure and services.

We had intended to ask respondents about their use of domain name system security extensions (DNSSEC), but those questions did not make it into this version of the survey. We hope to include them in future editions.

Access to Network Elements

Respondents were asked what mechanisms they use to access and configure network devices. Their responses are illustrated in Figure 12.

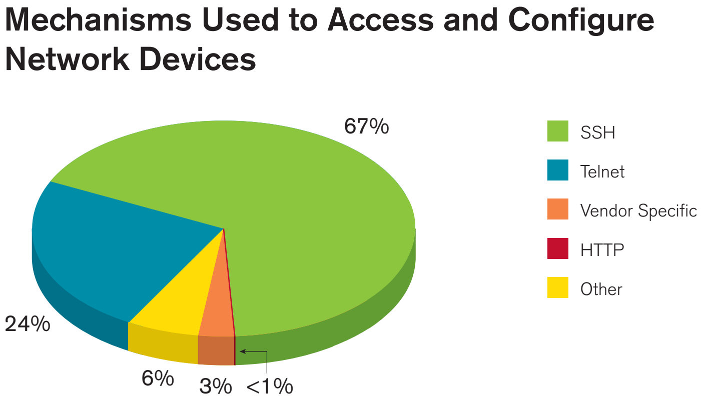

Figure 12: *Mechanisms Used to Access and Configure Network Devices*

Source: Arbor Networks, Inc.

While we suspect that there are some unspecified obstacles for secure shell (SSH) that are supporting such a strong Telnet user base, given its inherent insecurity (e.g., cleartext properties), it is still somewhat surprising that nearly 24 percent of all respondents access their network devices (i.e., routers, switches, etc.) via Telnet, which can easily be snooped or intercepted by other machines on the same local area network. In future editions of the survey, we hope to see a migration to more secure protocols such as SSH as they continue to improve, and will ask for details about any existing obstacles that may be hindering their deployment.

Managing and Monitoring Routers

In addition to asking what tools respondents use for command line access to their network devices, we also asked what protocols they use to manage and monitor their routers. Forty-five percent of respondents indicated that they still use SNMPv1, while only 17 percent have migrated to SNMPv3, which is more secure. Another 3 percent use vendor-proprietary protocols, and 35 percent employ unspecified “Other” mechanisms.

Some 20 percent of respondents indicated that they do enable SNMP write access on network devices, while 70 percent said they do not and 11 percent were unspecified. Unfortunately, given that less than 20 percent of respondents indicated they are using SNMPv3, there are not only SNMP writeable systems out there that use SNMPv3, but there are also SNMP writeable systems out there using SNMPv1. We do hope that operators continue to migrate to SNMPv3 and understand the security considerations with using SNMPv1.

In general, we are rather surprised at the number of operators that: 1) use cleartext Telnet to access routers; 2) employ SNMPv1 for monitoring and configuration; and 3) enable SNMP write capabilities at all. We do trust that they are sufficiently protecting these systems from external attack, and suspect that many of these operators have a considerable number of legacy systems and associated baggage that are keeping them from migrating to more secure network management techniques.

Application of Anti-Spoofing Techniques

IETF BCP 38/RFC 2827[8] provides an overview on anti-spoofing measures that should be employed by network operators. Essentially, it recommends that a network operator should not accept IP packets into the network on a given interface for a specific source IP address unless the packet's destination IP 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.

Unicast RPF (uRPF)[9] provides a mechanism to automate guidelines provided in BCP 38. There are multiple modes in which uRPF may be implemented to provide for 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 within the local routing systems and therefore are 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 and DNS cache poisoning attacks discussed earlier in this report. BCP 84/RFC 3704[10] provides some additional information on uRPF-based anti-spoofing techniques.

Figure 13 illustrates the application of anti-spoofing techniques by survey respondents. Very little has changed since we first asked these questions in last year's survey, essentially illustrating that of the pool of survey respondents, only about half implement any type of source address anti-spoofing measures at various perimeter points in their network.

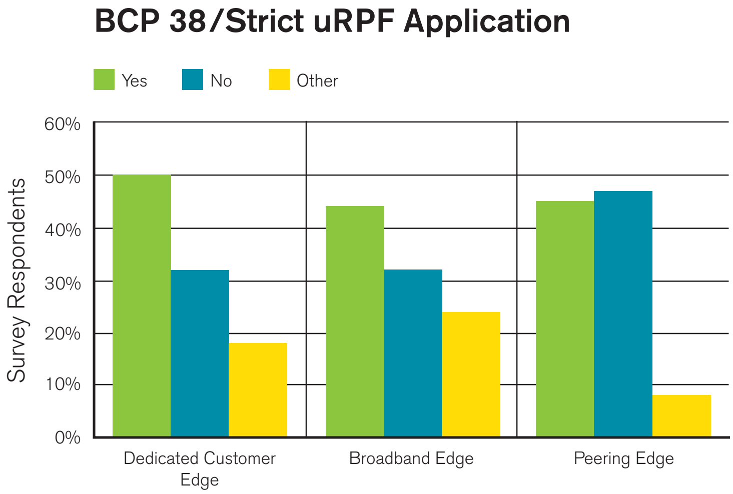

Figure 13: *BCP 38/Strict uRPF Application*

Source: Arbor Networks, Inc.

[8] www.ietf.org/rfc/rfc2827.txt

[9,10] www.ietf.org/rfc/rfc3704.txt

According to the data in Figure 13 (page 17), in general, anti-spoofing techniques are implemented less than 50 percent of the time. Given that the most vehement DDoS attacks today (e.g., reflective amplification attacks) and some of the nastiest targeted attacks (e.g., DNS cache poisoning) are most effective when spoofability exists, this data illustrates little new deployment and is disconcerting. The only good news here is that there is a slight uptick in the application of anti-spoofing when comparing these responses to those received last year, although this may simply be due to our collapsing and rephrasing the questions this year.

The application of anti-spoofing techniques is without question one of the more successful ways to squelch many of today's most effective attacks. ISPs not implementing these policies today should focus very directly on doing so in the near future, and not doing so should be considered irresponsible at best.

Protection of Routing Protocols

Network operators were also asked whether they used the MD5 signature option[11] of the Transmission Control Protocol (TCP MD5) or Internet Protocol Security (IPsec) to protect BGP transport connections in the network, and whether they employ the MD5 protection mechanisms available with their interior gateway routing protocols (IGP) as well. Responses are illustrated in Figure 14.

Quite a large number of network operators use MD5 mechanisms to protect external BGP sessions in particular, and well over half of the respondents also use it to protect internal BGP sessions and their IGPs. A considerable increase was observed over previous editions of the survey for use of TCP MD5 with external peers (eBGP), internal peers (iBGP) and MD5 extensions for IGPs. None of the respondents reported using IPsec for eBGP or iBGP this time around, in contrast with previous editions of the survey. Over 21 percent of the respondents reported using Generalized TTL Security Hack (GTSH).[12]

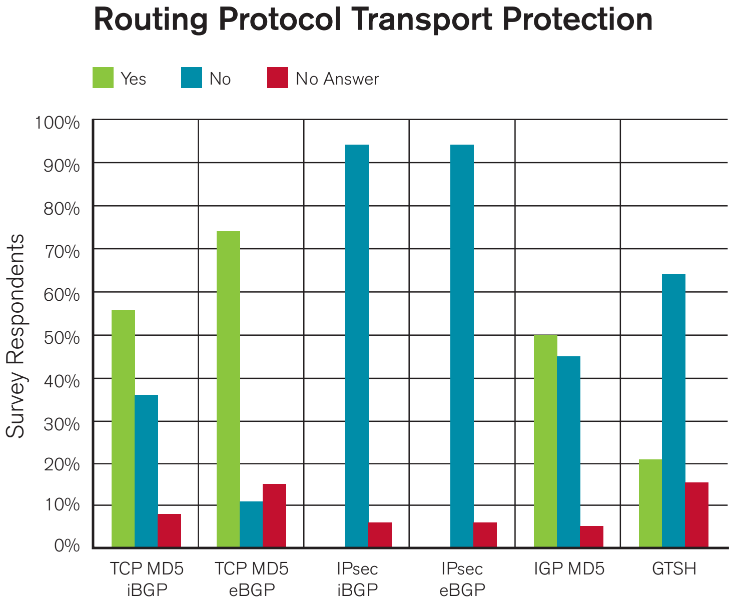

Figure 14: *Routing Protocol Transport Protection*

Source: Arbor Networks, Inc.

[11] www.ietf.org/rfc/rfc2385.txt?number=2385

[12] www.ietf.org/rfc/rfc3682.txt

Use of Internet Routing Registries

Five Regional Internet Registries (RIRs) are responsible for allocating and assigning IP addresses and autonomous system (AS) numbers to ISPs or end sites. Meanwhile, one or more (of many) Internet Routing Registries (IRRs) provide a database for ISPs to register routes and associated routing policies in order to mitigate the risk of malicious or accidental route hijacking and related threats. There currently exists no strict linkage between RIRs and IRRs today, or between RIRs and the actual routing system itself, although some of this work is underway.[13] While many holes exist in the security of the current Internet routing system, the use of 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.

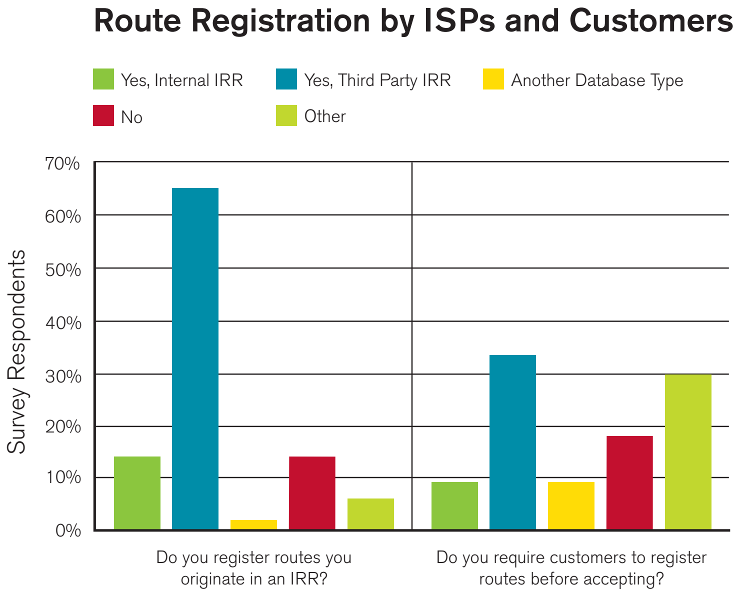

Figure 15: *Route Registration by ISPs and Customers*

Source: Arbor Networks, Inc.

As illustrated in Figure 15, about 80 percent of ISP respondents register the routes they originate in any IRR, while only 42 percent require customers to register their routes before the ISPs accept them. However, 71 percent of ISPs said they explicitly filter routes announced to them by customers, and 44 percent explicitly filter routes announced to them by peers. We would certainly be interested in better understanding how these filters are derived, given that 21 percent of ISPs do not even register routes they originate in an IRR.

The responses above seem to highlight the fact that very little or no inter-domain route filtering is applied on the Internet today, even less than that employed ten years ago. As a result, the state of routing security on the Internet has seemingly deteriorated over the past decade. The primary reasons for this certainly include the insecurities of the IRRs themselves, as well as the staleness of the data they contain. However, this diminished security should also be attributed to the fact that no authoritative database exists today for who has been allocated what IP address space, and what AS is authorized to advertise that address space. Fortunately, work is underway by several RIRs to develop a resource public key infrastructure (RPKI) that will provide this allocation and route announcement authorization database. The RPKI can be used either to populate IRRs or other filtering and policy generation tools and databases, or directly by routing protocols themselves in the future—all of which will enable more secure routing on the Internet.

[13] http://asert.arbornetworks.com/2008/05/using-rpki-to-construct-validated-irr-data

Along the same lines, we asked respondents if they currently have tools that monitor for the hijacking of their routes or the routes that belong to their customers. Fifty-two percent of respondents said they do have tools to monitor for route hijacking, while 48 percent indicated that they currently have no tools or services to monitor for such route hijacks. Given the frequency and exposure of route hijacks along with the insecurities of the routing system over the past year, we expect to see an increase in tools and services for monitoring route hijacking over the coming years.

Another question we asked in the survey, by request, was whether respondents believe that ISPs should be allowed to announce prefixes allocated by an RIR outside of the geographic region for which the RIR is responsible. About 55 percent of respondents said Yes (i.e., you should be able to use that prefix outside of the region where it was allocated), 33 percent said No, and 12 percent said something along the lines of "It's the Internet—who cares!" The reason this was asked has to do with whether network operators should be able to engage with RIRs in other regions in order to exploit address allocation policies that are less stringent than those of the RIRs where the operators reside and operate, particularly given the near-term exhaustion of the available IPv4 address pool.[14]

Finally, we asked respondents if they had experienced any unintentional configuration changes that produced the same effect as a DDoS; 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. Fifty-five percent of respondents indicated that they had indeed observed outages or adverse effects from such incidents over the past 12 months.

Tools for Event Correlation

When respondents were asked if they have tools to enable event correlation, nearly all respondents indicated they do have such tools or systems in place. However, the types of tools or systems in use vary widely and include:

- In-house developed tools
- Commercial tools (e.g., ArcSight, NetCool Micromuse, netForensics, HP Openview)
- In-house tools built upon FLOSS (freely available community-maintained software)
- Mix of commercial and custom tools
- Basic tools, but nothing automated
- Highly modified version of open source security information management (OSSIM)

Several respondents mentioned that they are in the process of developing or evaluating tools for event correlation at this time.

ACL Revision Control

We also asked respondents if they have tools to generate, maintain and deploy ACLs to network devices and firewalls. Obviously, these policies could be used to install persistent security policies, respond to DDoS attacks or other security incidents, or enable routing policy application. Nearly all respondents, some 97 percent, indicated they have some tool to perform these functions. These tools include:

- In-house developed tools
- Concurrent versions system (CVS)
- Open source tools (e.g., Rancid)
- Commercial tools (e.g., Arbor Peakflow SP for DDoS response)

Most respondents indicated that they use some mixture of in-house developed tools, CVS and open source tools.

Open Recursive DNS Resolvers

Only 56 percent of respondents indicated that they restrict access to their recursive DNS resolvers to customer IP space only, while 21 percent indicated that they do not. Another 23 percent provided no answer to this question (perhaps because they may not operate recursive resolvers for "customers").

[14] http://asert.arbornetworks.com/2008/08/the-end-is-near-but-is-ipv6

On Disclosure

Given the activities surrounding the DNS cache poisoning vulnerability and the three-phased disclosure method (pre, partial, full) employed, we asked what impact network operators believe such techniques have on scanning and attempted exploits. Seventy-one percent of respondents indicated that they believe that partial disclosures such as those with DNS cache poisoning simply result in more scanning, reverse engineering and exploit activities.

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. By request, we added several new questions to this section in 2008, as provided below.

Composition of Security Teams

Respondents were asked where their network security team resides within their organization. Most respondents indicated that it is either part of Network Engineering (53 percent) or that it falls within Network Operations (27 percent). While some 11 percent of 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. Only 5 percent 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 as displayed in Figure 16. Nearly 23 percent of respondents indicated that their organization has no staff dedicated explicitly to network security, with another 8 percent reporting that just one individual ("me") is responsible. However, 41 percent of respondents indicated that their network security team consists of two to four team members, 20 percent have five to eight team members, and nearly 9 percent have nine 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.

It is worth noting that the content networks have the largest staffs for dedicated security. This is likely due to the fact that their businesses provide data as a service; therefore, they are more aware of security issues that must be addressed to keep the services active. Tier 1 and Tier 2 ISPs also saw a growth in the number of security team members. This is likely due to additional managed service offerings for their customers.

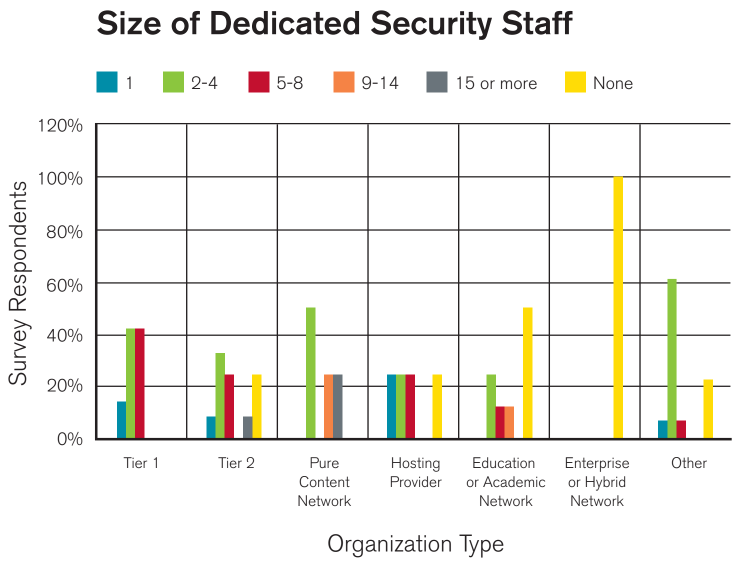

Figure 16: Size of Dedicated Security Staff

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 9 percent of respondents indicated they have a dedicated SOC as part of an MSS offering, while 26 percent indicated that there is a dedicated SOC in their organization focused on network infrastructure and security operations of network services. Nearly 65 percent of respondents indicated that there is no dedicated SOC within their organization.

Management and Executive Support

When respondents were asked if they believe their team receives adequate management-level support for security initiatives and projects, only 61 percent of 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 58 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.

ISPs: Bots, Botnets, AV and Malware

We asked respondents an array of questions ranging from botnet sizing, to distribution of antivirus (AV) software 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.

Bot Sizing

When respondents were asked about the size of the largest botnet they have observed over the past year, responses varied widely. Many respondents indicated they were not sure, did not measure, or did not care. Some cited that solely spoofed-based attacks (such as activities targeting DNS resolvers) made gauging actual bot sizes extremely difficult. Others indicated that their "Botnet observation posts have gone obsolete [or at least the command and control (C&C) techniques used]," or that "Now a criminal will run 1,500 compromised Web servers or 100 Unix-based bots instead of massive collections of IRC [Internet Relay Chat]-based Windows bots." Some respondents indicated that they "Still do not track C&C," while "the Storm botnet" came up in more than a few responses. For those who did provide numbers, they ranged from 100 to 200 hosts per botnet or botnet partition, to 1000 bots from a given botnet "within our constituency," to 10k, 30k, 50k, 175k, 200k and even 700k bots that were members of a single botnet.

As usual, some respondents look much deeper into bot and C&C activities, while others are most concerned only with what is occurring "within their constituency." Apparently, Storm made lots of noise over the past year, and there seems to be a continued migration away from vanilla IRC-based Windows bots to more resilient C&C mechanisms. Several respondents indicated that bot sizes are continuing to increase, while further partitioning, C&C separation and obfuscation, and flux techniques are being employed more commonly. Two respondents also suggested that there is a more decisive separation of bot activities; noisy bots used for things like DDoS are being separated from the herd, while the bulk of bots are still being employed for things like spam cannons or proxies/Haxtor nodes with sole economic aspirations.

Botnet Activities

Respondents were asked what activities they have personally observed bots performing over the past year. Not surprisingly, spam took the lead, followed closely by DDoS (Figure 17, page 23).

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

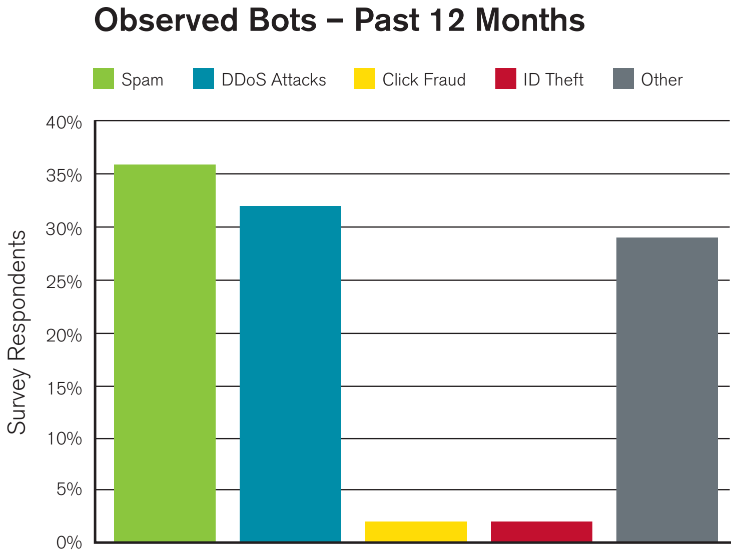

Figure 17: *Observed Bots – Past 12 Months*

Source: Arbor Networks, Inc.

Tracking Botnet Activities

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

> "End system IP address logging of suspicious hosts."
>
> "Flow data analysis."
>
> "We don't monitor directly, we just filter the cr— — they spew."
>
> "Rely on CSIRTs, security groups and data sharing."
>
> "Honeypots and darknet monitoring."
>
> "Snort with bleeding-edge rules."
>
> "Collaboration."
>
> "Internally developed magic."
>
> "Monitor DNS queries and unauthorized traffic."
>
> "It all depends, really on what the botnet is doing…If it's a spam cannon type botnet (e.g., Srizbi), simple extrapolation of originating IP vs. the message run in question will give you a general estimate. HTTP and IRC-based DDoS can be guessed from measuring the impact of the DDoS and/or infiltrating the C&C server."

Perhaps most humorous of all the responses was "Team Cymru and Jose Nazario/Arbor, not that I think they're tools." Arbor's Active Threat Level Analysis System (ATLAS®) was mentioned several times as well. Information sharing and collaboration were common responses; their value cannot be overemphasized. Many of the respondents reiterated this point.

We also asked if respondents believe that ISPs should be responsible for detecting and monitoring botnets. Sixty-one percent said Yes, while 23 percent disagreed, and another 17 percent responded Yes, with some criteria.

When respondents were asked how successful they believe anti-botnet tools and techniques have been over the past year, 20 percent indicated such tools and techniques were sufficient, while 68 percent said insufficient.

Quarantine, Walled Gardens and Cleanup

Some 36 percent of respondents surveyed said they do employ automated techniques for quarantine or walled-garden isolation of infected or malicious subscribers, while 52 percent said they do not. Twenty-nine 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 it has been quarantined.

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

42%	– No
26%	– Yes, via internal resources, for free
11%	– Yes, via internal resources, for a fee
3%	– Yes, via third party, for free
2%	– Yes, via third party, for a fee
17%	– 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.

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

Paying the Bill

When respondents were asked who should be responsible for covering losses associated with banking and related financial credential theft, 42 percent of respondents said banks, 15 percent said consumers or individuals and 35 percent did not respond.

Over 61 percent of respondents said they believe that AV and host security protection applications are not keeping up with Internet security threats.

Given the media coverage surrounding the Russian Business Network (RBN) and alleged sponsorship of illegal activities, we asked respondents to indicate whether they believe that hosting providers should be held accountable for selling services to known "shady" organizations. Twenty-two percent said No, while 50 percent said Yes and 21 percent provided no answer. Furthermore, 64 percent of respondents indicated that there are specific ASNs or networks on the Internet that they consider "evil" and that they would like to see go away.

Trends and the Future

When respondents were asked whether they see the scale and frequency of security threats increasing or decreasing as IPv6 becomes more widely deployed,[15] 55 percent believe threats will increase, while only 8 percent believe threats will decrease.

Regarding respondents who have deployed data retention methods, 26 percent said their methods are based on logs only (i.e., smtp, pop/imap, login/logout, etc.), 8 percent said they are based on flow data, 53 percent said they are based on both logs and flow data and 14 percent said they are based on other techniques.

[15] www.arbornetworks.com/IPv6research

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.”
>
> “Zero-sum game.”
>
> “More customer awareness on the needs to monitor their networks.”
>
> “Seeing more behavioral patterns in attacks (e.g., students always attacking on national holidays).”
>
> “Less resource, more work, less board-level interest/understanding.”
>
> “More time on revenue-generating services/products.”
>
> “Attackers are getting smarter. Tools keep up, but users don’t.”
>
> “Web 2.0 attacks on the increase.”
>
> “Much more awareness regarding routing security.”
>
> “Large Web mail operators like Google don’t give a sh— — about spam originating from their networks because they know they are too large to be blacklisted. This causes significant pain.”
>
> “New services are not prepared or developed to deal with Internet security threats.”
>
> “Attackers are smarter and more professional.”
>
> “More law enforcement as they appear to be getting more resources.”
>
> “More attacks are unintentionally impacting other services (e.g., our VoIP).”
>
> “Attackers are ‘definitely’ getting smarter, developing better techniques to hide exploits in Web sites, better rootkits to avoid AV detection, and better 2nd+ stage download techniques to avoid honeypot grabbing.”
>
> “Customers see ISP as ‘the Internet’ and blame for things outside their control. On one hand protect, on the other, don’t even think about looking at our packets or invading our privacy.”

Respondents were also asked if they have deployed DPI equipment, and if so, for what purpose. Fifty-two percent said No, they have not deployed such equipment. Of those answering Yes, 11 percent said it is for lawful intercept, 20 percent for service enforcement and protection, and 3 percent for both lawful intercept and service enforcement and protection. Fifteen percent responded with “Other,” with no details provided.

Conclusions

As this year's survey makes clear, the ISP security landscape continues to change rapidly. ISP optimism about security issues that was reported in last year's survey 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.

Aggressively focused on new revenue, ISPs say they are increasingly deploying more complex distributed VoIP, video and IP services. But surveyed ISP security engineers also say these new services are often poorly prepared to deal with the new Internet security threats. Overall, more than half of the surveyed ISPs believe serious security threats will increase in the next year while their security groups make do with "fewer resources, less management support and increased workload."

ISPs were also unhappy with their vendors and the security community. Most believe that the DNS cache poisoning flaw disclosed earlier this year[16] was poorly handled and increased the danger of the threat. The surveyed ISPs also said their vendor infrastructure equipment continues to lack key security features (like capacity for large ACL lists) and suffers from poor configuration management and a near complete absence of IPv6 security features.

While most ISPs now have the infrastructure to detect bandwidth flood attacks, many still lack the ability to rapidly mitigate these attacks. Only a fraction of surveyed ISPs said they have the capability to mitigate DDoS attacks in 10 minutes or less. Even fewer providers have the infrastructure to defend against service-level attacks or this year's reported peak of a 40 gigabit flood attack. Over the last several years, survey trends also suggest that attack size may be on pace to approach 100 gigabits by this time next year, or perhaps more sophisticated attack vectors will be employed and bits per second won't be as significant.

[16] www.us-cert.gov/cas/techalerts/TA08-190B.html

About the Editors

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

danny@arbor.net

With over 15 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. Previously he held network operations and architecture positions for nearly a decade at internetMCI, Genuity (acquired by GTE Internetworking), Qwest Communications and the U.S. Army.

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.

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.

Contributing Editors

Michael Hollyman, Manager of Consulting Engineering, Arbor Networks

mhollyman@arbor.net

With more than 12 years in the network, security and telecommunications industries, Mike 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. Jose Nazario, Senior Security Researcher, 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.

Dr. G. Robert Malan, Founder and Chief Technology Officer

rmalan@arbor.net

Rob Malan brings over ten years of research experience in computer networking and security to Arbor Networks. Dr. Malan, whose thesis work at the University of Michigan formed the basis for Arbor Networks' technology, is the author of the company's patents. Dr. Malan has successfully transitioned technology from research prototype to product during his tenure in industry, which includes work at the IBM T.J. Watson Research Laboratory and Hewlett-Packard. Dr. Malan began his networking career working as a researcher on the Mach operating system project at Carnegie Mellon. He has authored 18 papers published in top-tier computer security and networking journals and conference proceedings. Dr. Malan holds a Ph.D. and MSE in Computer Science from the University of Michigan and a B.S. in Computer Engineering from Carnegie Mellon.

ARBOR®
NETWORKS

Corporate Headquarters

430 Bedford Street
Lexington, Massachusetts 02420
Toll Free (USA) +1 866 212 7267
T +1 781 684 0900
F +1 781 768 3299

Europe

T +44 208 622 3108

Asia Pacific

T +65 6327 7152

www.arbornetworks.com

About Arbor Networks

Arbor Networks is a leading provider of secure service control solutions for global business networks, including more than 70 percent of the world's Internet service providers and many of the largest enterprise networks in use today. Arbor is addressing the most strategic issues for service providers—security and service control; delivering best-in-class network protection and the means for delivering revenue-generating, differentiated secure services and service plans. Arbor allows service providers to employ both flow-based and DPI-based technologies to enable measurement and protection of the entire network, from the 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.