

7 Interconnected Contexts

Distributed Denial of Service Attacks and Other Digital Control Measures in Asia

Hal Roberts, Ethan Zuckerman, and John Palfrey

In early 2008 the Vietnamese government announced plans to mine bauxite, the mineral used to make aluminum, in the Central Highlands of Vietnam in cooperation with a Chinese company. These plans became the subject of increasing protest beginning in 2008 and continuing thereafter. Protesters have expressed environmental concerns about damage to mined areas and toxic by-products of bauxite mining. While some activists involved with the bauxite protests have been connected to banned prodemocracy movements, others have been protesting the Chinese-backed mine on grounds of environmental concern or national pride.¹

In 2009 a group of activists distributed a petition and created a Web site named <http://bauxitevietnam.info> to protest the bauxite mining. According to reports from Vietnamese free-speech advocates, both the bauxitevietnam.info site and the larger bauxite protest movement have been under constant attack since 2009. The government has repeatedly detained and interrogated both the founders of bauxitevietnam.info and many of those who signed the petition. Forged e-mails, purportedly by the founders of the Web site, have been distributed online, falsely claiming that the leaders were quitting the protest. Activists report that the Vietnamese government broke into the site's servers to steal protester information and shut down the site.²

In January 2010 a flood of traffic from compromised computers overwhelmed bauxitevietnam.info, making it inaccessible not only in Vietnam but also throughout the entire Internet.³ Political actors increasingly use this type of attack, known as a distributed denial of service (DDoS) attack, to control content on the Internet. Vietnam has routinely filtered Internet sites the government considers to be controversial, preventing users in Vietnam from accessing them without taking unusual steps. In contrast, a DDoS attack makes a Web site inaccessible to all online audiences by disabling a targeted Web server under a flood of traffic.

This particular DDoS attack used a botnet, an army of "zombie" computers that have been taken over, in the vast majority of cases, without their owners' knowledge. These zombie computers are generally used to commit some sort of fraud on the network. For example, some computers controlled by botnets are used to sign up for thousands of

free e-mail addresses and send spam. In this case, the zombie computers sent an extraordinary number of requests to <http://bauxitevietnam.info>, crashing the site.

Shortly after the DDoS attacks on the site began, Google announced that it would no longer censor its search results in China⁴ because of attacks on its Gmail service, which it found had originated from within China. While investigating the source of those Gmail attacks, Google found evidence that the botnet attacking bauxitevietnam.info—though not involved in the Gmail attacks—consisted largely of computers that had been infected by a malicious program hidden by an attacker within a program called VPSKeys.⁵

Technicians at Google and at the antivirus firm McAfee then unraveled the story of the bauxitevietnam.info DDoS attacks. VPSKeys is the most popular Vietnamese keyboard input program. Distributed by the Vietnamese Professionals Society (VPS), it allows Vietnamese users to enter Vietnamese characters easily using Western keyboards. Some months before the attacks on bauxitevietnam.info, likely in late 2009, the Web site hosting the VPSKeys software had been compromised. The attacker replaced the VPSKeys program with a Trojan version designed to infect the host computer with botnet software. The attackers also alerted thousands of VPSKeys users by e-mail that a new (secretly infected) version of the software was available. Many Vietnamese users updated their software in response. It is likely that the attackers were able to obtain the mailing list used to send this e-mail through a separate attack—possibly intrusions that seized membership databases of popular Vietnamese discussion forum sites in 2009.

Tens of thousands of users downloaded the Trojan software, which infected the host computers and added them to a botnet before the Trojan software was discovered. The makers of VPSKeys replaced the infected software with a clean version, but not before the Trojan software had created the network of compromised computers. This botnet was used to mount the DDoS attack on bauxitevietnam.info and may have been used against additional targets.

Why did the attackers go through the effort of compromising computers and creating their own botnet? There is a thriving underworld business devoted to the sale of lists of infected computers, which in essence allows attackers to rent these computers for the purpose of a one-time attack like the one on bauxitevietnam.info.⁶ A plausible explanation is that a botnet of computers based in Vietnam would be difficult for a site administrator to defeat through geographic filtering. If bauxitevietnam.info were attacked by thousands of computers located in South Korea, an administrator might respond by blocking all requests to the Web site from that country. But blocking requests from Vietnam would defeat the purpose of raising awareness within Vietnam itself. It is also possible that the botnet was an added benefit in a scheme that primarily sought to monitor the activity of Vietnamese-speaking users around the world. The botnet was certainly capable of spying on the owners of the infected computers,

possibly logging keystrokes and capturing passwords to online accounts, even possibly collecting the list of e-mail addresses used to encourage more people to download the Trojan software.

The administrators of bauxitevietnam.info defended the site from the DDoS attack by mirroring the site on multiple hosting providers. They created mirrors at <http://bauxitevietnam.info>, <http://boxitvn.org>, <http://boxitvn.net>, <http://boxitvn.info>, <http://boxitvn.blogspot.com>, and <http://boxitvn.wordpress.com>. The last two of these mirroring environments are especially important, because they are hosted by large blog-hosting services, Blogger (run by Google) and WordPress. These large-scale services offer highly DDoS-resistant services at no direct financial cost to the activists.

It is very rare that an observer can come to identify the owner of any botnet, as Nart Villeneuve and Masashi Crete-Nishihata also find in their fine-grained review of DDoS and defacement attacks in Burma in chapter 8 of this volume. It is the nature of a botnet to be distributed across a broad range of computers infected without the knowledge of their owners. Accordingly, no one (including Google and McAfee, two of a handful of actors most capable of diagnosing this sort of attack) has managed to determine who controlled the botnet during the course of the Bauxitevietnam attacks—or, for that matter, who controls it at the time of this writing. But there are indications that some DDoS attacks against Vietnamese sites have involved more than tacit approval of the Vietnamese government. Viet Tan, a Vietnamese prodemocracy dissident group, reports that their site is routinely subject to DDoS attacks and that many of the attacking computers are based in Vietnam.⁷ Since <http://viettan.org> is generally blocked in Vietnam, these attacks require that authorities lift the blocks on attacked sites to permit attacks from zombie computers in Vietnam. It is difficult to verify this claim without access to Viet Tan's server logs documenting such an attack.

This example—by no means extraordinary, particularly in Asia—shows how DDoS attacks accompany a range of interventions that involve malware and related intrusions into the computers of ordinary Internet users. It demonstrates that governments and other political actors are using a broad array of intertwined methods to contest online (and offline) content that they find offensive. For example, the methods of attack in this case include the following:

- DDoS attacks
- Technical Internet filtering
- Surveillance
- Intrusion by means of malware
- Trojan software
- Online identity forgery
- Offline harassment

The difficulty of diagnosing (and defending against) these attacks is further complicated by the large set of actors, many of whose precise roles are unclear. For example, the attacks on <http://bauxitevietnam.info> may have involved the following:

- *Attackers* A set of attackers, which may or may not have included the Vietnamese and Chinese governments, whose precise identity is unknown and who are responsible for a number of DDoS attacks, intrusions, and forgeries; the hundreds or thousands of compromised computers used to attack <http://bauxitevietnam.info>, most likely without knowledge of their owners.
- *Defenders* The administrators of bauxitevietnam.info, and administrators of the hosting services and Internet service providers (ISPs) they use.
- *Affected third parties* The Vietnamese Professionals Society (which inadvertently distributed malware), McAfee (which detected the attack), and Google (responsible for both investigating the DDoS attack and defending against the attacks via Blogger).

The use of malware particularly complicates this type of analysis. Researchers usually consider malware the province of commercial actors who compromise computers to participate in schemes designed for financial gain. This type of example demonstrates how malware is now playing a major role in how political actors seek to constrain Internet users both within and beyond their borders.⁸ These interconnected controls make diagnosing DDoS attacks an enormous challenge in many cases simply because it is difficult to understand the full array of methods used with the attacks as well as who is executing those methods. In this case, it is likely that some of the computers that attacked bauxitevietnam.info were owned by individuals who supported the goals of the organization. This possibility, in turn, made the attack even harder to block because distinguishing between legitimate and attack traffic was impossible to do on an IP basis.

Finally, this example demonstrates what relatively sophisticated activists can do to defend themselves from DDoS attacks. A common strategy is to diversify hosting and, especially, to flee to large blog hosts, often based in the United States, for cover. While simple to understand and implement, the strategy is extremely effective, allowing the activists behind bauxitevietnam.info to maintain an online presence in the face of a sustained attack without paying for a fee-based DDoS-protection service, the most effective of which start at thousands of dollars per month.

DDoS and Other Next-Generation Control Measures

This chapter is a deep dive into the growing phenomenon of DDoS attacks. We seek to describe the state of DDoS attacks in the context of the interconnected contests to control online content. Our central goal is to situate the phenomenon of DDoS attacks within the theoretical framework developed in OpenNet Initiative (ONI) research. In

particular, this in-depth review of the DDoS phenomenon builds on the observation that the types of control mechanisms that states and others may employ have evolved from the first-generation Internet control process of technical filtering to the second- and third-generation controls that we have observed emerging since the middle part of the 2000s.⁹ As Ronald Deibert and Rafal Rohozinski say of these next-generation controls in the opening chapter of *Access Controlled*:

Although there are several tactics that can be employed within this rubric—deliberate tampering with domain name servers, virus and Trojan horse insertion, and even brute physical attacks—the most common is the use of DDoS attacks. These attacks flood a server with illegitimate requests for information from multiple sources—usually from so-called “zombie” computers that are infected and employed as part of a “botnet.” The ONI has monitored an increasing number of just-in-time blocking incidences using DDoS attacks, going back to our first acquaintance during the Kyrgyzstan parliamentary elections of 2005.¹⁰

A group of ONI researchers have also tracked other early instances of DDoS attacks in the Belarus elections of 2005, the Russia-Estonia dispute in 2007, and the Russia-Georgia conflict of 2008.¹¹ In this chapter, we build upon these previous findings of our ONI partners in this broad-based review of DDoS attacks that were independent of particular sensitive political moments, as well as the detailed research on the 2008 Web defacement attacks in Burma in chapter 8.

Our research method in studying DDoS was multifaceted. We conducted an in-depth analysis of media reports on human-rights and independent media-connected DDoS attacks, surveyed independent media and human rights sites, conducted confidential interviews, and hosted a working meeting with participants from multiple related sectors in Cambridge, Massachusetts, in 2010. We shared our results with knowledgeable peers before disseminating them and discussed possible responses to the rising DDoS threat. Although our research was meant to cover DDoS broadly around the world, Asia proved to be one of two regions we focused on, along with the Commonwealth of Independent States (CIS). Our respondents in Asia came in particular from Burma, China, and Vietnam, and we focus on cases from those countries here.

We sent a survey on DDoS attacks to a sample of 317 independent media and human rights sites. We generated the sample by asking at least three local experts in each of the nine target countries for the most prominent independent media in their countries. We translated the survey into the primary Internet language of each surveyed country and also translated the recruitment e-mail to the primary language of each site. We received full responses from 45 sites, for a response rate of 14 percent.

These survey methods limited our findings in several ways. The sample involved was not large. Despite this limitation, we perceive that the 45 responses amount to a decent response rate for such a survey, given a series of special factors involved. These factors included the difficulty of reaching a key actor at each site, the inherent

sensitivity of the survey subject, and the early stage of research in this field. We used neutral language that did not explicitly refer to DDoS attacks when querying the experts for the list of sites, but some of the experts were familiar with our work and therefore likely to bias their lists of independent media toward sites known to suffer DDoS attacks. It is likely that the 14 percent of responding sites overrepresents sites that have suffered a DDoS attack, since a survey on DDoS attacks may seem more interesting—and worth responding to—to DDoS attack victims. These two factors make the results of the survey less useful for answering questions about overall prevalence of DDoS attacks. We cannot, for these reasons, answer questions about what percent of all independent media sites in our surveyed countries have suffered DDoS attacks. But we believe that the responses are useful for investigating the nature of attacks reported by the surveyed sites and the defenses used by those sites.

We conducted interviews in person, over Skype, and by e-mail with administrators of 12 sites that experienced DDoS attacks. We contacted every survey respondent who reported having been subject to a DDoS attack and requested a more in-depth interview. Six of the interview participants were recruited through this method. We found the rest of the interview participants through media analysis or through referrals from researchers and other contacts in the field. We interviewed administrators of sites based in Australia, Burma, China, Iran, Russia, and Vietnam. The interviews involved a series of questions and answers tailored to each interviewee exploring the technical details of attacks and the experiences of the administrators dealing with them. In a few cases, we obtained and analyzed logs of attacks. We cannot publish the interviews themselves for security reasons, but we include a number of findings, in aggregate form, from the interviews.

Additionally, we studied as many published reports of DDoS attacks as we could find by tracking accounts posted to the Web over the course of six months. Our sample set includes 329 reports of attacks against more than 800 sites going back to 1998. We also had the unexpected opportunity to study a DDoS attack that happened to occur during the course of our research. Our research home, the Berkman Center for Internet and Society at Harvard University, hosts the site of a sister research project, the Citizen Media Law Project, which happened to be attacked by a sustained denial of service attack. We were able to study that attack in progress, as it happened, for which we are grateful to the unknown attackers.

Interconnected Methods of Contesting Information Online

A core finding from our survey and related methods is that DDoS attacks exist within a portfolio of different attacks suffered by these sites. We also found that the same site usually suffers from multiple types of attacks. During the past year, of the surveyed sites,

- 72 percent experienced national network filtering of their sites.
- 62 percent experienced DDoS attacks.
- 39 percent experienced an intrusion.
- 32 percent experienced a defacement.
- Of those experiencing a DDoS attack, 81 percent also experienced at least one of the following other content controls: Internet filtering, intrusion, or defacement.

These numbers provide strong evidence that DDoS attacks are not an isolated problem for independent media sites. Instead, DDoS attacks exist within a larger range of different kinds of attacks against the sites. In addition to the specific range of attacks reported, the surveyed sites reported a high level of unexplained downtime during the past year:

- 61 percent experienced unexplained downtime.
- Of those respondents who experienced unexplained downtime, 48 percent experienced seven or more days of unexplained downtime.

Unexplained downtime can be the result of factors other than attacks. Independent media sites often suffer from a lack of experienced system administrators, leading both to downtime and to the inability to diagnose the reasons for downtime. Still, the very high amount of unexplained downtime experienced by these sites suggests more, and possibly more complex, attacks than described by the answers to the preceding DDoS question.

Our finding that a significant number of sites have experienced 21 days or more of downtime suggests that there is a serious shortage of technical capacity available to respond to threats to independent media and human rights Web sites. Arbor Networks, a leading DDoS mitigation firm, surveys large ISPs annually about their experience with DDoS. Their survey of tier-one and -two ISPs suggests that most administrators of large ISPs respond to a typical DDoS attack within an hour.¹² The administrators we interviewed were unable to bring their sites back online in such a timely fashion.

Our in-depth interviews provided further support for the findings, in both our survey and media research, that DDoS attacks are often accompanied by intrusions, defacements, filtering, and offline attacks. One administrator of more than a dozen independent media sites reported DDoS attacks followed by offline extortion intended to force him to retract a story. (He refused.) That same administrator reported being subject not only to DDoS attacks but also to daily virus-laden e-mails targeting him personally and about topics of confidential interest to him; to weekly intrusion attacks based on guessed passwords; and weekly defacement and complete deletion of at least one of the sites under his control.

Another administrator had been subject to weeks-long, multigigabit DDoS attacks but reported that a greater problem was the harassment of participants in the publication's discussion forums: attackers broke into the discussion forum to steal and publish

the identities of its users and also posted inflammatory content to the forum to trigger governmental prosecution. Yet another administrator reported that intruders had repeatedly accessed internal databases to learn about stories before they were published. And another reported that attackers broke into his site to insert malicious code with the intent of triggering antivirus warnings for the site and thereby scaring users from accessing it. He also reported intrusions to his site that inserted code that slowed the Internet connections of his users by causing them to download large packages of Trojan horse software. In all cases, the DDoS attacks may have been the most visible manifestation that a site was under attack. But the attacks that accompanied the DDoS attacks were often of far more concern and import to the affected administrators.

DDoS attacks vary greatly in their nature and magnitude. In our interviews, we heard about a range of attacks, extending from multi-Gbps floods of traffic that overwhelmed the network connectivity of the affected sites to attacks that used as few as a few dozen requests per minute to cripple sites by exploiting holes in Web servers and other applications. Five of the interview participants reported attacks in the range of 500 Mbps to 4 Gbps. One participant, who was the administrator of a large service provider working for an independent media site, reported an attack of greater than 10 Gbps. Some of these attacks may have been bigger, since at greater than 1 Gbps, many local ISPs become saturated and drop any additional traffic. One interview subject, whose site experienced several DDoS attacks in the previous four years, reported an escalation of the size of attacks over time. His site had been successfully disrupted in 2007 with a 1 Gbps DDoS attack, and he moved to more robust, DDoS-resistant hosting provider. His contract with the provider specified that he would be protected from attacks up to 2 Gbps. When an attack in 2010 involved 4 Gbps of traffic, his host took his site offline, offering him the option of either increasing his monthly payments or remaining offline until the attack ended.

Three interview participants reported application attacks at low—even very low—bandwidths that caused significant downtime. One was taken down by fewer than 40,000 requests per day, another by less than ten machines hitting his search page. Two participants reported long-term success using mitigation strategies—caching and Web application optimization—which would be effective against only relatively low bandwidth attacks. We believe these attacks exploited known holes in application software, such as the Slowloris attack against Apache Web servers.¹³

It is likely that most or all network attacks that we encountered in our research involved the use of botnets to generate incoming traffic. Other indicators suggest that some of the attacks involved the use of rented botnets. Two interview subjects reported that attacks began and ended at the top of an hour, suggesting that a botnet had been rented for a specific duration. The DDoS attack against the Berkman Center's Citizen Media Law Project offered further evidence of rented botnet attacks. The DDoS attack was an application attack using HTTP GET requests originating from a shifting set of

exactly 500 IP addresses. The attack was highly effective, rendering the site inaccessible for 12 hours, despite steady work from the Berkman Center's highly experienced technical staff to keep the site online. That the attack came from a round number of attacking IPs and that the IP addresses in use shifted in real time in response to defenses suggests that the application attack came through a rented botnet.

We also saw a strong correlation between DDoS, filtering, defacement, and intrusion attacks in our media analysis. These techniques were often used in conjunction, and may have synergistic effects—making a site more DDoS resistant can make it more difficult to access using a Web proxy, for instance, which makes state-based filtering more effective. Independent media organizations participating in the working meeting repeated the same theme: sites suffer from multiple types of attacks, including DDoS, which in turn have complicated impacts on one another.

A key example of these impacts was the problems that a prominent Burmese independent Web site experienced from a combination of DDoS attacks and national filtering. The Web site has moved to a DDoS-resistant hosting provider to protect itself against high-bandwidth-traffic attacks. The site in question is routinely filtered by the Burmese government, so people within the country must use proxies to access the site. Burmese users gravitate toward a small set of proxies discovered through word of mouth. All the traffic from each of those proxies appears to come from the same IP address. One method the DDoS-resistant hosting provider uses to protect against attacks is to block IP addresses that are submitting too many requests. Since the proxies submit many more requests than other IP addresses, the hosting provider often bans them, to the end effect of blocking Burmese audiences from accessing the site. It is possible to address this problem by providing the hosting provider with a white-list of proxy servers, but that list is difficult to maintain because users in Burma keep seeking new proxies to stay one step ahead of government efforts to block them.

Non-DDoS attacks on a site are often more serious and less tractable than DDoS attacks. A common method for intrusions is to compromise the computer of someone who has administrator-level access to the target server. Access to the server is then used to delete sites; to discover the identities of dissidents, authors, and sources for further on- and offline harassment; to deface the target site; or to implant malware on the target site either to discredit the target site or to execute a DDoS attack on another site or both. Administrators of human-rights-related independent media consistently report being frequently subject to specifically targeted e-mail viruses, often connected to content tailored to be of interest to the administrator in question. These specifically targeted attacks are very difficult to defend against, requiring a high level of training and support for the victims. But many or most of the independent media organizations struggle to maintain even very simple client-side technology infrastructures.

For example, one participant—an administrator of a well-funded and prominent Asia-based nonprofit organization—reported that his organization shared two desktop computers among its staff of several dozen people. Many of these staff members had never touched a computer before working for the publication. Defending client computers that are so widely shared and used by such novice users, and that are specifically and aggressively targeted, is an enormously difficult problem to solve for even one organization, let alone for the field as a whole. We know of at least one organization focused on political rights in Asia that has a policy of reformatting the hard drive of laptop computers that have been removed from the office and used on other networks. Most organizations do not have nearly this level of concern or technical competence, even though most targeted organizations likely need to operate at this level of caution.

Two of the independent media site administrators we interviewed reported multiple types of attacks coming from multiple sources, as well as confusion about the source of the attacks. One participant was subject to a DDoS attack when he published a story about a prominent government actor, and then was approached separately both by the government actor with demands to take down the offending story and by the group of cybercriminals who were carrying out the attacks with demands for money. Another participant claimed that his site is sometimes attacked by the government when it is unhappy with a particular story and sometimes attacked by activists in opposition to the government when they are unhappy with a story (and sometimes the activists have taken credit for attacks that the participant thought were certainly coming from the government). Others we surveyed suggest that, in many cases, the effectiveness of DDoS attacks was a matter of gaining press coverage rather than success in taking and keeping a site down. In other words, even when DDoS is the only attack a site faces, the actors and their motivations may be complex and multilayered.

Site Administrators Worry about DDoS, among Other Attacks

Despite their prevalence, DDoS, intrusion, and defacement attacks are not the primary concern for most independent media sites. Asked to rank the impact of various issues, participants placed DDoS, intrusion, and defacement attacks squarely in the middle of the pack among other Internet content-control issues. The issues were ranked in the following order, with the most important issue listed first and with the average rank out of five noted (a higher number implies a lower priority):

- Blocking access to the publication's site by the government (2.47)
- Persecution of authors, publishers, or sources by the government (2.53)
- Intrusions, defacements, and denial of service attacks (2.89)
- Financial support for the publication (3.00)
- Technical issues other than defending against attacks (3.89)

While DDoS attacks are an increasingly prevalent form of Internet control, our respondents listed conventional government filtering as the most serious problem they face. Only 11 percent of respondents chose DDoS, intrusion, and defacement attacks as the most pressing issue, and only 32 percent chose these attacks as one of the two most pressing issues. These are particularly interesting findings given the bias of the study toward respondents facing such attacks. By comparison, 68 percent of respondents chose persecution of authors, publishers, or sources by the government as one of the two most pressing issues. Issues directly related to censorship and control (filtering, persecution, and DDoS and other attacks) all ranked higher than the two issues not directly related to censorship and control (finance and nonattack technical issues).

Effective Responses to DDoS Attacks Are Elusive

A common response to a DDoS attack is to turn to the hosting ISP during the time of the attack. The survey respondents had mixed luck getting their ISPs to defend them against attacks. Of those who experienced a DDoS attack in the past year,

- 55 percent had their site shut down by their ISPs in response to the attack.
- 36 percent report that their ISP successfully defended them against a DDoS attack.

The number shut down by their ISPs is surprisingly high, considering that an ISP will usually shut down an attacked site only when subject to a traffic-based attack (since other type of attacks generally do not directly affect the ISP's network or other customers). The fact that 55 percent of respondents suffering a DDoS attack had been shut down by their ISPs at least once indicates that at least 55 percent, and almost certainly more, of the sites had been subject to a traffic-based attack. This fact, along with the fact that only 36 percent of the respondents subject to DDoS attack had an ISP that defended them against attack, indicates that for many independent media, the local ISP is a weak point rather than a strong ally. We do not know whether the reason for this poor defense of sites by their ISPs is that independent media sites are customers of sites outside the core of ISPs able to respond to an attack in under an hour or whether the reason is that the independent media sites are customers of the core ISPs but are not able to pay for the DDoS protection that those ISPs generally sell as an add-on service.

In our survey, we also asked site administrators about the defenses they had tried when hit by a DDoS attack and how effective those defenses had been. Their responses can be read as a map of how independent media escalate defenses against DDoS attacks:

- 83 percent had fixed problems with their existing Web application software, with 80 percent reporting that this measure was “somewhat effective” or “effective.”

- 75 percent had installed security software or hardware on their existing servers, with 92 percent reporting that this measure was “somewhat effective” or “effective.”
- 62 percent had upgraded their Web server hardware, with 88 percent reporting that this measure was “somewhat effective” or “effective.”
- 43 percent had downgraded the functionality on their existing sites, with 33 percent reporting that this measure was “somewhat effective” or “effective.”
- 40 percent had subscribed to a denial-of-service-protection or other security service, with 100 percent reporting that this measure was “somewhat effective” or “effective.”
- 38 percent had hosted content temporarily on a large hosting provider (Blogger, LiveJournal, etc.), with 67 percent finding that this measure was “somewhat effective” or “effective.”
- 36 percent had changed their hosting providers, with 80 percent reporting that this measure was “somewhat effective” or “effective.”
- 29 percent had changed their Web application software, with 75 percent reporting that the change was “somewhat effective” or “effective.”

The vast majority of sites that experience DDoS attacks try to update the configurations of their local computers by fixing the existing Web application software, installing local security hardware or software, and installing upgraded local Web server hardware, or some combination of these three approaches. These basic strategies can all be taken by individual sites without help from core network providers, though in some cases core technical expertise may be needed to properly apply these upgrades. Each of these approaches rates as at least somewhat effective against DDoS attacks, insofar as these basic changes prove somewhat effective against further attacks.

A much smaller number of sites escalate their responses either by implementing more aggressive (and costly) defenses at the edge—downgrading functionality or changing Web application software—or by moving closer to the core of the network: subscribing to expensive protection services, hosting content on large providers, or changing hosting providers. The success of these defenses is more mixed than the simple edge-based fixes, perhaps because these are the defenses that are valid responses to network attacks, which are much more difficult to fend off than application attacks.

Our results indicate that the number of attacks against each site increased for a slight majority of participating sites:

- 16 percent reported many more attacks in 2010.
- 36 percent reported somewhat more attacks in 2010.
- 48 percent reported no change or fewer attacks in 2010.

ISPs—who are best positioned to defend sites against many types of DDoS—are often unable or unwilling to defend their customers. This finding leads us to speculate that many of the sites we surveyed are (or were, as many have been dropped by those

providers) tier-three providers, who may lack a fiscal incentive to protect their customers. Tier-three Web-hosting providers sell their services for a small margin over costs—the hours worth of system administration time necessary to fend off a DDoS attack is more costly than the annual profit for the average account. These providers evidently do not see a reputation risk in failing to fend off a DDoS, and they find it more profitable to end relationships with “troublesome” customers than to provide protection to them.

The apparent efficacy of upgrading servers and fixing Web server software strongly suggests that attacks are not all based on clogging network connectivity (where these defenses would be ineffective) and point to application-level vulnerabilities. These sorts of fixes are only really helpful for either very small traffic attacks or application attacks, both of which can be reasonably dealt with by individual publishers at the edge of the network.

Best Practices for Human Rights and Independent Media Sites Are Emerging for DDoS Response

According to experts with whom we consulted, the responses that a site might take to a DDoS attack include the following:

- Blackholing the IP address of the attacked site (i.e., taking the attacked site offline).
- Deploying additional network and server infrastructure for the attacked site.
- Downgrading the content and/or functionality of the attacked site to reduce resource consumption.
- Filtering out attack traffic.
- Using a service with a distributed architecture to scale and absorb attacks on demand.

These responses range from the simplest to implement (taking the site offline, which is essentially giving up in the face of an attack) to complicated and difficult to implement.

Blackholing the IP address of the attacked site fulfills the aims of the attacker by making the site unavailable. But this response also makes the attack traffic disappear entirely from the Internet. In so doing, it protects the network hosting the site. This is the approach taken by many ISPs that are faced with a large traffic-based attack that is either too big or too expensive for them to defend against.

An attacked site may deploy additional servers and bandwidth to protect itself. Our survey results show that this is indeed the most popular method of protection. But for all but the biggest sites, deploying additional infrastructure for a single site is cost effective for small, application-based attacks only, because the peak traffic of a large, traffic-based DDoS attack will be orders of magnitude larger than the peak legitimate traffic of a site.

An alternative to increasing the server resources is to reduce the resource consumption of each page, allowing the server to handle more traffic with the existing server and network. There are some methods for reducing resource consumption that are effective and have little cost, such as caching dynamic content to reduce database queries. As attack size increases, though, an attacked site has to make changes that have costly side effects, like disabling site functions that require expensive database queries, reducing or eliminating images and streaming media, or creating an entirely separate failover site with simpler and less-interactive content.

Another way to reduce resource consumption is to distinguish attacking traffic from legitimate user traffic and filter out the attacking IP address. This approach is frequently used, and several of our meeting and interview participants reported success with this method, but only when the number of attacking machines is small and relatively static. It is simple for a competent system administrator to find and block a hundred static IP addresses that are flooding a site with requests for a single page, but that job becomes much, much more difficult when there are tens of thousands of IP addresses that are rotating every couple of hours and actively trying to make their traffic look legitimate. In these cases, it is sometimes possible to filter attacking traffic based on a signature for the particular traffic, but this approach can be very difficult against a moderately skilled attacker even for a highly skilled defender. It is possible to defend against a range of common attacks by using ModSecurity, an open-source attack-filtering system. But this sort of filtering helps against generic attacks only, and it uses up machine resources for the process of filtering and can therefore make the site more vulnerable to traffic-based attacks.

Finally, a site can protect itself by paying for a hosting or DDoS protection service to serve the content of the Web site. There are many services capable of handling all but the biggest attacks, and a few capable of handling the biggest observed attacks, simply because they have sufficient bandwidth and server resources to accept and process the attack traffic. The advantage of using such a service is that these services have economies of scale both in learning how to defend against particular attacks and in the necessary bandwidth and servers. When using such a service, the attacked site needs to pay for the peak attack traffic only while the attack is happening, rather than paying for the entirety of the resources needed to handle peak attack traffic.

These services, however, can command a very high markup on those resources. Even without the high markup, simply paying for the bandwidth to handle the peak attack traffic can be prohibitively expensive, especially for an independent media site. An attacked site may be able to hire a provider capable of handling millions of requests per second but not be able to afford the resulting bandwidth charges. The economies of scale work best for these sites if a large proportion of the site is not likely to be attacked at the same time, which is important to keep in mind given the model we found in interviews of a single local expert managing many sites from a given area

(meaning that all or many of those are likely to be attacked at critical times for the country). As we noted previously, sometimes an attack will outstrip an administrator's ability to pay the associated bandwidth charges, and the site will be forced to go dark until the attack ceases.

Given the trade-offs of the various defense mechanisms, it is critical for sites that know they are likely to be attacked to weigh the various options before they are affected by DDoS. For instance, site administrators will need to know whether to pay the startup costs to hire a protection service, how much to pay a service to withstand a traffic-based attack, and at what point to accept that the cost of defending against a given attack is too high.

Hiding Their Tracks? Ample Suspicion, but No Hard Evidence, That States Are Involved in DDoS

Most sites participating in the interviews expressed a strong belief that the national government of the country their site reported on was ultimately responsible for the attacks. None, however, had clear evidence of state responsibility. One participant had reported a large, ongoing attack to the state's security service but got no help since "it is very difficult to look into this because it is very difficult to catch yourself." He asserted that the security service shut down its own attack only when other publications better connected to the government complained. One Vietnamese site pointed to a press report of a Vietnamese military official claiming responsibility for the attacks.¹⁴ As mentioned previously, a Viet Tan administrator noted that his site was normally filtered from within Vietnam but that the filtering was taken down at precisely the time that a botnet from within Vietnam attacked the site. Most interview participants asserted the opinion that the national government was responsible for the attacks but did not claim any direct evidence for the responsibility. This inability to attribute direct responsibility for DDoS attacks is typical for the attacks. The distributed nature of the attacks makes it difficult to assign responsibility—it is certainly possible that either a government or progovernment individuals could attack a site critical of a specific regime, and our inability to trace the attack would not be an unusual circumstance. Our findings in these respects are consistent with the findings of Villeneuve and Crete-Nishihata in chapter 8.

As a related matter, we also found no obvious connection between the particular ideology of an attacker and the choice of DDoS as an attack method. We saw attacks from ostensibly right- and left-wing groups, attacks that targeted governments, and attacks that suggest government involvement. Neither is there an apparent geographic pattern to the DDoS attacks we saw in our media analysis. We found attacks reported in widely disparate corners of the world. Asian states were a common site for DDoS attacks, but certainly not the only region where they appear. While there is

speculation that some attacks are traceable to governments—for instance, the example of <http://bauxitevietnam.info>—it is unclear that this is an assumption with any merit. DDoS is a technique used by individuals, groups, and, perhaps, states. The accessibility of easy-to-use tools and the apparent success of single-user attacks on small Web sites, as well as the technique’s visibility in the media, suggest that aggrieved individuals may look to DDoS as an easy way of making a political point or settling a score. We note, too, that the widely reported DDoS attacks in the context of the release of U.S. State Department cables by Wikileaks in the fall of 2010 involved attacks both on Wikileaks itself and on major banks and others in apparent retaliation. In an ironic and perhaps inevitable twist, 4chan—an online community that claimed responsibility for many retaliatory attacks—was taken down by a DDoS on December 28, 2010. As with other Internet control mechanisms, DDoS is an approach used by a variety of actors to accomplish a variety of ends.

Conclusion: Situating DDoS in the Context of “Next-Generation Controls” and Other Online Contests

In response to the growing usage of next-generation Internet controls, citizens may be banding together to fend off DDoS and related attacks, at least on a modest scale. In three of our interviews, we heard of local technical experts acting as hubs of technical expertise for their countries (in Vietnam, China, and Iran, specifically). The most productive and satisfied of these local experts was far along in the process of moving sites in his country to a common infrastructure well supported by a hosting provider that was well connected to the core of the Internet (in all senses of “core”: community, expertise, and resources). He was able to exert a great deal of control over the structure of the moved sites, including imposing onerous security and posting restrictions on the sites’ administrators. The most concerned and least content of these local experts was struggling daily with many poorly written sites on broken, incompatible code bases, often reinstalling a site from scratch following an intrusion and manually fighting off the simpler of the constant DDoS attacks. He told us that he had the desire, but not the resources, to fix the underlying problems with the supported sites, as well as gratitude for the help he has received from other individuals, but he was frustrated by his inability to fend off high-bandwidth traffic attacks.

The threat of DDoS attacks is inextricable from other security considerations, including human resources concerns, technical resources, and community connections. Ultimately, what human rights and independent media organizations face, in Asia and elsewhere around the world, is a combination of a shortage of skilled site-administration skills, the bandwidth needed to fend off large network attacks, and the community connections needed to ask core network operators for help to fend off attacks. The difficulty of responding effectively to DDoS attacks is a symptom of a

larger problem: most small, independent organizations simply do not have the talent, bandwidth, or connections to administer independent Web sites in the face of potential attack. The online environment not only offers new ways to reach a broad audience, inside a state and beyond, but also poses new challenges in keeping that online accessible in the face of the many types of attacks described in this book.

There is a final twist to the story. Citizens who wish to publish independent media sites but who do not have significant technical savvy are most likely to be able to resist DDoS and related attacks by signing up with a large, free hosting service. These services, such as Google's Blogger or WordPress, are often run by large, for-profit companies that are not based locally where the activists are situated. This approach was the strategy used by <http://bauxitevietnam.info> in the attacks described at the beginning of this chapter, which Google ultimately diagnosed and then defended the site by providing resistant hosting. The interconnected nature of these attacks, along with the possible responses, puts citizens in Asia and elsewhere in common cause with multinational companies based elsewhere, pitted together against an elusive opponent that may or may not include their own state. Rebecca MacKinnon takes up this topic in greater detail in chapter 10 of this volume.

Though increasingly unavoidable, this allegiance between human rights organizations and large corporations can be a tenuous and complicated one. In the fall of 2010, when Wikileaks was subject to a DDoS attack after releasing U.S. State Department cables, they turned to Amazon.com to serve their Web site.¹⁵ A few days later, Amazon.com decided to stop hosting Wikileaks, which continued to be subject to DDoS attacks, just as the perceived allies of Wikileaks launched DDoS attacks against large banks and others perceived to have turned against Wikileaks.¹⁶ While these independent media sites may have interests aligned with large corporate players to some extent, their allegiance may break down in the context of pressure from states or other powerful interests. It is important to note, however, that any ISP providing services to Wikileaks would likely have come under political pressure from the U.S. government. It is possible that other providers would have acquiesced under similar pressure.

The days of simple filtering of offensive Web sites, in the manner pioneered by Saudi Arabia and a few other states roughly a decade ago, are long past. The interplay of this range of public and private actors and next-generation mechanisms in cyberspace is becoming increasingly complicated and unpredictable. Independent media and human rights operations, especially in Asia, have a much harder job than ever before to keep their Web sites accessible in times of conflict.

Notes

1. BauxiteVietnam Blog, <http://boxitvn.wordpress.com/>; Bauxite Vietnam Web site, <http://bauxitevietnam.info>; John Ruwitch, "Web Attacks Hit Vietnam Bauxite Activists: Google," Reuters, April 2, 2010, <http://www.reuters.com/article/idUSTRE62U0TM20100402>.

2. Ben Stocking, "2 Popular Web Sites Blocked in Vietnam," *Sydney Morning Herald*, February 27, 2011, <http://news.smh.com.au/breaking-news-technology/2-popular-web-sites-blocked-in-vietnam-20100211-nupw.html>.
3. Viet Tan, "Denial of Service: Cyberattacks by the Vietnamese Government," April 27, 2010, <http://www.viettan.org/spip.php?article9749>.
4. In January 2006, Google launched google.cn, a search engine hosted in China and censored to comply with Chinese law. In January 2010, Google shut down google.cn and redirected traffic to their unfiltered Google.com.hk site.
5. Neel Mehta, "The Chilling Effects of Malware," Google Online Security Blog, March 30, 2010, <http://googleonlinesecurity.blogspot.com/2010/03/chilling-effects-of-malware.html>.
6. See the Shadowserver Foundation's resources on botnets and related activity at <http://www.shadowserver.org/wiki/pmwiki.php/Information/Botnets> for some of the most detailed analysis about the formation, growth, and application of botnets. One of the best of the white papers posted on the Shadowserver site, which addresses this question of rented botnets, is Krogoth, *Botnet Control, Construction, and Concealment: Looking into Current Technology and Analyzing Future Trends*, March, 2008 (special version for Shadowserver Web site), http://www.shadowserver.org/wiki/uploads/Information/thesis_botnet_krogoth_2008_final.pdf (see especially section 2.5, "Motivation and Usage"). See also *PC Magazine* Encyclopedia's entry on botnets: "Also called a 'zombie army,' a botnet is a large number of compromised computers that are used to create and send spam or viruses or flood a network with messages as a denial of service attack. The computer is compromised via a Trojan that often works by opening an Internet Relay Chat (IRC) channel that waits for commands from the person in control of the botnet. There is a thriving botnet business selling lists of compromised computers to hackers and spammers." Available at http://www.pcmag.com/encyclopedia_term/0,2542,t=botnet&i=38866,00.asp.
7. Viet Tan, "Denial of Service."
8. There have been an increasing number of reports of malware-related attacks on human rights organizations that may or may not have involved states. See, for instance, the path-breaking reports by our colleagues at the Information Warfare Monitor, available at <http://www.infowar-monitor.net/research/>.
9. For a description of first-generation Internet controls, see Robert Faris and Nart Villeneuve, "Measuring Global Internet Filtering," in *Access Denied: The Practice and Policy of Global Internet Filtering*, ed. Ronald Deibert, John Palfrey, Rafal Rohozinski, and Jonathan Zittrain (Cambridge, MA: MIT Press, 2008). For discussion of second- and third-generation controls, see Ronald Deibert and Rafal Rohozinski, "Beyond Denial: Introducing Next-Generation Internet Controls," in *Access Controlled: The Shaping of Power Rights and Rule in Cyberspace*, ed. Ronald Deibert, John Palfrey, Rafal Rohozinski, and Jonathan Zittrain (Cambridge, MA: MIT Press, 2010).
10. Deibert and Rohozinski, "Beyond Denial," 8.
11. Ibid.; OpenNet Initiative, "The Internet and Elections: The 2006 Belarus Presidential Election (and Its Implications)," OpenNet Initiative Internet Watch Report 001, April 2006, <http://opennet.net/sites/opennet.net/files/2006%20Internet%20Watch%20Report%20Belarus.pdf>.

12. Danny McPherson, Roland Dobbins, Michael Hollyman, et al., "Worldwide Infrastructure Security Report: Volume V, 2009 Report," Arbor Networks, January 19, 2010, http://staging.arbornetworks.com/dmdocuments/ISR2009_EN.pdf.
13. Christian Folini, "Apache Attacked by a 'Slow Loris,'" LWN.net, June 24, 2009, <http://lwn.net/Articles/338407/>.
14. Radio Free Asia, "Hackers in Vietnam: Do Not Confess Also That," May 19, 2010, <http://www.rfa.org/vietnamese/programs/ReadingBlogs/%20VN-police-general-confesses-state-s-hacking-cb-s-web-and-blog-NHien-05192010161301.html>.
15. Andrew R. Hickey, "Wikileaks Turns to Amazon Cloud to Dodge DDoS Onslaught," CRN, November 30, 2010, <http://www.crn.com/news/cloud/228400232/wikileaks-turns-to-amazon-cloud-to-dodge-ddos-onslaught.htm>.
16. Chloe Albanesius, "DDoS Attacks Continue to Plague Human Rights Sites," *PC Magazine*, December 22, 2010, <http://www.pcmag.com/article2/0,2817,2374654,00.asp>.

