Should We Stop Trusting Trust?

Greg Goth

As Internet disruptions go, the 24 February incident that blocked access to YouTube for two hours wasn't particularly disastrous in economic terms. It wasn't even particularly unique — within weeks, another IP address hijacking blacked out service provider Africa Online Kenya for more than 22 hours.

However, YouTube's global visibility and the context in which it was blocked — by a state-operated telecommunications provider acting as a censor — combined to greatly increase worldwide attention on interdomain routing and the vulnerabilities inherent in routing protocols that can cripple the network.

"If you asked me where the weakest link was in today's Internet infrastructure, there is no doubt this would be my answer," says Danny McPherson, chief research officer for network security firm Arbor Networks. "Subsequent to this would be DNS [Domain Name System] and the name server infrastructure, but that's much more resilient today. There's no doubt [that] I absolutely believe this is the weakest link in the Internet today."

The Symbolic Black Hole

Martin A. Brown, an engineer at network monitoring firm Renesys, compiled a fairly comprehensive timeline of the YouTube hijack and partial blackout on the Renesys blog (www.renesys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml#more). The trouble began after Pakistan Telecom, acting on a government directive to block access to YouTube within Pakistan, began advertising a black hole route to YouTube. However, Pakistan Telecom's upstream provider, Hong Kong-based PCCW, failed to validate the Pakistani advertisement and erroneously supplied the false route to its peering networks elsewhere.

Essentially, the advertised Pakistani address block (208.65.153.0/24) was more specific than YouTube's legitimate block (208.65.152.0/22); that is, the higher the number after the slash, the smaller the address block and the more specific the advertised route. As Border Gateway Protocol (BGP) researcher Iljitsch van Beijnum wrote (http://arstechnica.com/news.ars/post/20080225-insecure-routing-redirects-youtube-to-pakistan.html): "For instance, a router could have both 10.0.0.0/8 and 10.10.0.0/16 in its routing table. Then, if a packet for 10.10.10.10 arrives, how should it be forwarded? The answer is, longest match first. The smallest address block with the largest number after the slash takes precedence."

In this case, the smallest address block advertised was bogus, and any users trying to reach YouTube were sent into a black hole. Although best current practices assume that providers will check the validity of advertised blocks, that obviously doesn't always happen.

"One thing we asked ourselves when we examined it in detail was, how could PCCW have known programatically that YouTube hadn't opened up a data center in Karachi?" Brown says. "Then it would be perfectly legitimate for this advertisement, even though it was a slash 24 and not a slash 22. It would be perfectly legitimate for YouTube to open an office in a new place and advertise the number."

Not quite three weeks after the YouTube hijacking, Africa Online Kenya found one of its prefixes erroneously announced by service provider Abovenet. The resulting route instability lasted for what McPherson termed "an order of magnitude" longer than the YouTube incident, although the problem was noted almost immediately.

Ultimately, Brown says, the fact that bogus routes can be propagated often comes down to the razor-thin margins inherent in core transit. Validating router announcements is still a relatively labor-intensive task, and hijacks thus far have been minor enough in scope and impact that providers haven't focused on preventing them.

"If you're a provider, and you're running a low-margin business, how do you justify spending on two people for a year for this type of monitoring?" asks Brown. "Even if they're not fantastically technical, you're talking [US]$100,000 to $120,000 a year. How do you justify that?"

 1089-7801/08/$25.00

Good Intentions and Bad Possibilities

Currently, consensus in the routing community holds that the old human trust model in the tight-knit global routing community will be sufficient to maintain interdomain routing integrity for the immediate future, as long as each player begins to assiduously strive toward implementing best current practices.

Relying on that model for any length of time has its shortcomings, however. McPherson says that 10 years ago, roughly 500 technicians worldwide had access to BGP-speaking routers; such a small community could quickly communicate network problems and trust expert technicians to promptly address problems. Today, he places that number at 250,000 technicians.

Concurrent with the explosion in the number of technicians charged with running this infrastructure has been what Tomas L. Byrnes, founder and chief technology officer of the network monitoring start-up, ThreatStop, says is a "race to the bottom" by backbone providers caused by economic pressures.

"At some point and at some level, we have to do something about what's happened to the backbone providers," Byrnes says. "The consolidation that's been going on there has led to a race to the bottom. They maintain their margins by spending less on people and equipment. The fact is, there's a lot of money to be made by breaking it — but very little money to be made by fixing it — and the people who could fix it, the backbone providers, have no economic incentive whatsoever to do so."

What appears to be the most obvious solution — for legitimate holders of IP address blocks to advertise them as even more granular blocks than the bogus announcements — can be done, and is done, but many BGP routers won't accept such granular blocks, and flooding the Internet with these longer routes can lead to BGP router table congestion.

Several more sophisticated, technical solutions have been proposed over the years, ranging from services offered by individual vendors to standards-level technologies such as Secure BGP (SBGP) and Secure Origin BGP, introduced by BBN Technologies and Cisco, respectively. Thus far, the market has declared both solutions to be computationally and economically too expensive for widespread adoption.

Other technical solutions include monitoring data from the Prefix Hijack Alert System (http://phas.netsec.colostate.edu/index.html) as well as other resources such as Dshield and Whois. Byrnes and his ThreatStop colleagues are following this model and trying to mitigate some of its present limitations, including high false positives for hijacked prefixes. As you might expect, high false positives limit the amount of dynamic filtering, which leads to higher staffing costs for providers adopting such technologies.

Providers might have to absorb higher costs sooner rather than later because the relative ease with which criminals can hijack addresses is beginning to cause more worry among networking experts. Both Byrnes and McPherson outline possible scenarios for IP hijacks with far more sinister motivations than the YouTube or Kenyan incidents.

"The next time, I don't think it's going to be a bunch of medievalists trying to enforce theocracy," Byrnes says. "I think it's highly likely to be a criminal gang trying to hijack a bank."

Byrnes says a possible vector for such an attack might involve a gang hijacking a multinational bank's subsidiary IP address. For bank customers traveling on another continent, a local prefix advertisement could take precedence over the bank's main IP address.

News in Brief

Seven major **mobile telecom equipment** companies — including **Alcatel-Lucent**, **Ericsson**, **NextWave Wireless**, and **Sony Ericsson** — are uniting to establish rules for licensing patents related to **long-term evolution (LTE) technology**. Their effort aims to avoid the bitter battles that erupted during **third-generation mobile equipment development**; as in that case, different companies hold intellectual property rights related to various aspects of LTE technology.

More information is available at http://media.nextwave.com/phoenix.zhtml?c=215860&p=irol-newsArticle&ID=1129196&highlight.

The **International Telecommunications Union (ITU)** has announced a series of **conferences** aimed at encouraging great thinkers to unleash their visions of **next-generation networks (NGNs)**. The first conference, **Innovations in NGN** is cosponsored by the **IEEE Communications Society** and convenes in **Geneva 12–13 May.** According to the ITU, the conference series will highlight **technologies, applications, and services** that capitalize on NGN infrastructure, leading to a **ubiquitous network society** that makes information accessible "anywhere and anytime by anyone and anything."

More information is available at www.itu.int/ITU-T/uni/kaleidoscope.

Hoping to encourage greater interest in Malaysia's December election, **Datuk Rahman Dahlan**, the executive secretary of the youth wing of the **United Malays National Organization (UMNO)**, has announced that all candidates for national youth posts **will launch blogs**. The party, which has been **highly disparaging of bloggers** in the past, issued the blog order in the wake of serious losses for the UMNO-led coalition in the March general elections.

continued on p. 8

News in Brief

continued from p. 7

More information is available at http://thestar.com.my/news/story.asp?file=/2008/4/13/focus/20939569&sec=focus.

As fears and complaints about **Vista** mount, **Microsoft customers** are growing increasingly agitated about the company's plans to pull the widely used **Windows XP** off the shelves in June and gradually erase it from its customer service agenda. In concert with a **"Why Save XP?"** video invitational, **InfoWorld** launched an **online petition** to protest the XP offing and, as of mid-April, had garnered more than **100,000 signatures**.

The online petition, related information, and a countdown to XP's scheduled demise are at http://weblog.infoworld.com/save-xp.

Although 88 percent of US and UK Internet users recently surveyed said that **personal irresponsibility was the primary cause of identity theft** and fraud, nearly half of those surveyed use the **same password for multiple Web sites**. The April survey results from **Accenture** consultancy also found that only 7 percent of the 800 people queried used any form of password security, such as changing passwords regularly or using password-management software. The impression, says **Robert Dyson**, vice president of Accenture's global security practice, is that "a lot of people don't think there's a problem."

Bigreds.com, an online collectibles retailer, is suing Yahoo! for more than US$1 million, claiming that it did little to prevent **click fraud** that resulted in significant overcharges. The suit alleges that Yahoo's fees were based on ad clicks on Yahoo-affiliated Web sites, but that it was actually site operators doing the clicking to up their own commissions rather than legitimate collectible seekers.

"So, if your bank is headquartered in San Francisco and you're in Prague, you may get redirected and not know it," he says. "A smart cracker will fix it so that as far as you know, your transaction still completes, but in two days when you go to get on a train in Milan, your account is empty. These guys don't have to grab a whole lot of people's information to make this worth their while. That would be the über phish. The only thing you might not be able to make work is the SSL certificate, but how many times do people see [that] it doesn't match the host and click anyway?"

McPherson says he could envision a scenario in which a gang sells a company's stock short, then hijacks its legitimate IP address at a critical juncture in its business cycle, such as at the end of a fiscal quarter. With the address hijacked, orders never reach the company; it misses its sales forecasts, its stock value drops, and the short sellers pocket a tidy profit.

"There are lots of different ways you can monetize something like this," he says.

An Unlikely Rescuer?

Ironically enough, one of the more innovative solutions isn't based on a new technology, but an old one whose obsolescence experts have forewarned of for many years: IPv4.

McPherson contends that the looming exhaustion of IPv4 address space — combined with a slow rollout for IPv6 — will result in those older addresses gaining considerably in value, with commensurate trading mechanisms created for exchanging these resources as needed.

McPherson calls IPv4's newfound value the "egg" answering the chicken-and-egg conundrum of mandating community investment in a definitive hierarchical IP address allocation and authentication mechanism. In this scenario, Regional Internet Registries (RIRs) assume much larger operational roles in allocating and authenticating numbers announcements.

"Enter Resource PKI [public-key infrastructure] and SIDR [Secure Interdomain Routing], with community and specifically RIR work on Resource Certification," McPherson wrote in a blog entry (http://asert.arbornetworks.com/2008/03/ipv4-exhaustion-trading-routing-autonomy-for-security). "In short, this work is aimed at providing an infrastructure that enables certification of 'Right of Use' for IP addresses and AS [autonomous system] numbers with X.509 Resource Certificates. If this infrastructure exists, then it can be used by RIRs to maintain control of IP numbers resources. It could also be used by folks for informational purposes, or to define routing policies based on a verifiable source, or even directly employed by the routing system itself through protocols such as SBGP."

McPherson says the SBGP part of this scenario is "a huge jump" — maybe 15 or 20 years down the road — but the exhaustion of IPv4 addresses, and the need to more closely monitor them, is a near-term incentive to begin incremental deployment of the infrastructure he outlined in his blog entry.

"The registries now have to put an infrastructure in place to control IP resources," he says. "That's been a very big missing link. Another factor is that two, maybe three orders of magnitude more people have access to BGP-speaking global Internet-connected routers. And another thing that's changed is the criticality of availability of e-commerce sites or network elements with reachability provided by BGP."

It's clear that many engineers are indeed concerned about maintaining the integrity of these critical connections. Although the networking experts say the visibility and con-

text of Pakistan hijacking YouTube might give the issue more temporary urgency, they also say they don't know if that urgency will translate to substantive improvements in network architecture.

"It does point to the fact we know there is something wrong with the fabric of the networks, and that this is something the planet has to come to grips with," says Marcus Sachs, Verizon's executive director of government affairs for national security. "It's not a Verizon problem or an AT&T problem. Everybody has to figure out if we are going to put up with the experimental protocols of the Arpanet era as a permanent way to run the networks, or are we going to take the next step and say these were good experimental protocols, but how do we now engineer the 21st century version using everything we've learned so far?"

Greg Goth is a freelance technology writer based in Connecticut.