Balkanization from Above

DAN GEER AND HD MOORE

Dan Geer is the CISO for In-Q-Tel and a security researcher with a quantitative bent. He has a long history with the USENIX Association, including officer positions, program committees, etc. dan@geer.org

HD Moore is the Chief Research Officer at Rapid7, responsible for leading Rapid7 research into real-world threats and providing guidance on how to address them. In addition, HD drives technical innovation across Rapid7's products and services, applying technology to the challenge of identifying and defending against current and emerging threats, as well as heading the development of experimental prototypes and free tools. HD is the creator of Metasploit, an open source penetration testing framework, and remains deeply involved in Metasploit's evolution. x@hdm.io

The Internet of 2015 is a different place compared to five years ago. Business models have changed, technology has shifted onward, hundreds of millions of new people have connected to the World Wide Web, and so forth. How they connect, what devices they use, and the threats they face have likewise shifted, and, to our point, the Internet is itself being dragged along.

Where the Internet was transparent and distributed, it is becoming opaque and centralized. The immense, if abstract, value of peer-to-peer communication has been eclipsed by—indeed has become subservient to—consumer demand for downstream content. Nowhere is this more apparent than in the mobile Internet. The IPv4 address space is running out of steam and service providers are compromising bi-directional network communication in favor of scalability. In corporate America, businesses are choosing the economies of scale in cloud offerings and rejecting local datacenters in favor of external on-demand infrastructure.

The end result is an inversion from a peer-to-peer "freedom to connect" model to one consisting of service provider enclaves providing private access to managed offerings. The Internet is increasingly attenuated between broadband on the one end and cloud providers on the other, with decreasing open space in between. Criminals, governments, and curious hackers alike are following this trend and changing their tactics in approximate (if ironic) synchrony. ISP-provided routers are becoming the target of choice for threat actors globally. Vulnerabilities in mobile devices and desktop operating systems are more valuable than ever. Cloud providers are increasingly targeted, and many are failing. The attack surface of the Internet necessarily grows faster than linearly with the count of endpoints, but even that is increasingly difficult to measure.

IPv4 Utilization

The IPv4 Internet has room for approximately 4.3 billion unique addresses, of which 3.7 billion can be used by public networks and hosts. These addresses are a finite resource managed by regional Internet registries, and as of June last year, we ran out. Figure 1 shows the number of /8 network blocks available from 1995 to June 2014.

The Internet relies on DNS to associate a name with an address. Of the 3.7 billion usable addresses, over 1 billion have an associated reverse DNS name. As the IPv4 Internet has run out of free network blocks, growth of named hosts has dropped accordingly. Figure 2 shows the growth of named hosts. (The logistic curve's inflection point was, as shown, November 21, 2008.)

The ITU (International Telecommunication Union) estimates that there are over 3 billion Internet users as of 2015 [1]. This number represents over 2.3 billion mobile broadband subscriptions and another 700+ million fixed broadband subscriptions [2]. Combine these stats with infrastructure equipment such as routers, switches, and all of the servers that actually power the Internet, and it is clear there isn't room for everyone in IPv4. In contrast to the rate of IP allocations and named hosts, growth in total connected devices seems to continue.

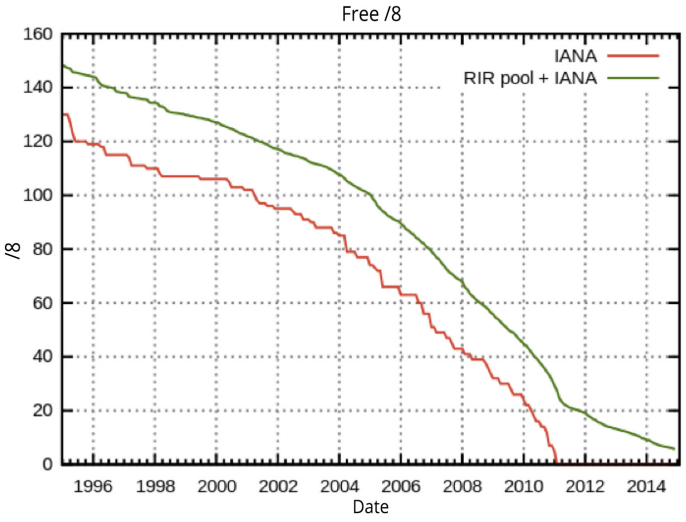

Figure 1: Number of /8 blocks available by date

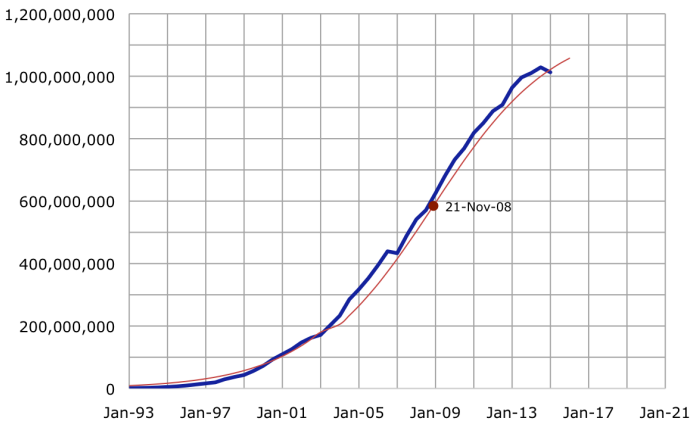

Figure 2: Growth curve and inflection point for number of hosts with PTR records

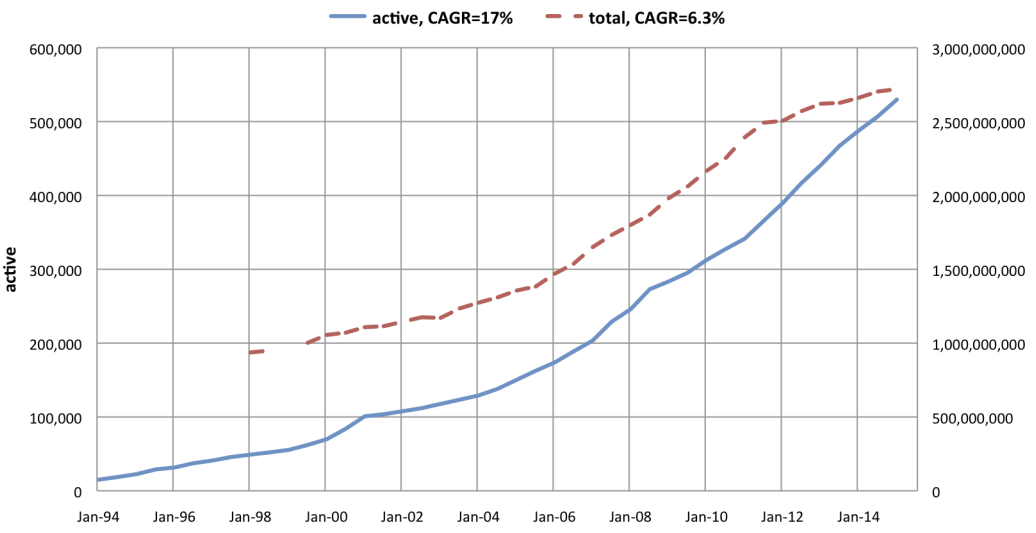

Figure 3: Active space (left vertical axis), total space (right vertical axis)

In a similar vein, growth of the total advertised IP space is slower than growth of subdivision within that space (compound annual growth rate, or CAGR, of 17% versus 6.3% as measured by BGP); see Figure 3.

Note that instead of a lengthy diversion into IPv6 and next-generation addressing, we keep our discussion to the Internet as it stands today. At its most succinct, there are far more users than there are IPv4 addresses, and IPv4 addresses are distributed unequally, sometimes to an absurd degree.

Approximately 370 million IPv4 addresses respond to an ICMP echo request. This represents about 10% of the usable IPv4 space. If we send common TCP and UDP probes as well, this number rises to 466 million IPv4 addresses (13%). The Hilbert graph in Figure 4 represents the density of hosts responsive to ICMP, TCP, and UDP probes. The extreme density in the lower left and center right are in clear contrast to the "empty" blocks in the upper left. The majority of reserved ranges are concentrated in the upper right quadrant and are evenly shaded. Many of the empty blocks are actually in use by government agencies and large corporations, but have been isolated from the rest of the Internet by firewalls (another form of enclave).

This 466 million number is important; it is the number of IPv4 addresses that are remotely discoverable and thus directly targetable by an attacker. The number of directly connected IPv4 systems puts an upper bound on the number of potential targets for any new server-side exploit. At the same time, the number of DNS PTR records at 1013 million is twice as big. What is going on?

3 Billion Users

The number of broadband users, consisting of both fixed-line and mobile, has increased from 500 million in 2007 to over 3 billion in 2014. Figure 5 demonstrates this growth. Contrast the 466 million discoverable IPv4 addresses with 3 billion broadband users and one asks, how are these users connected?

Mobile Broadband

There have been more mobile broadband users than fixed-line broadband users since 2008. In 2014, over 2.3 billion mobile devices were connected through mobile broadband, a mix of feature phones, smartphones, and tablets. If each of these devices required a public IPv4 address, there would be very little room in IPv4 for anything else; see Figure 6.

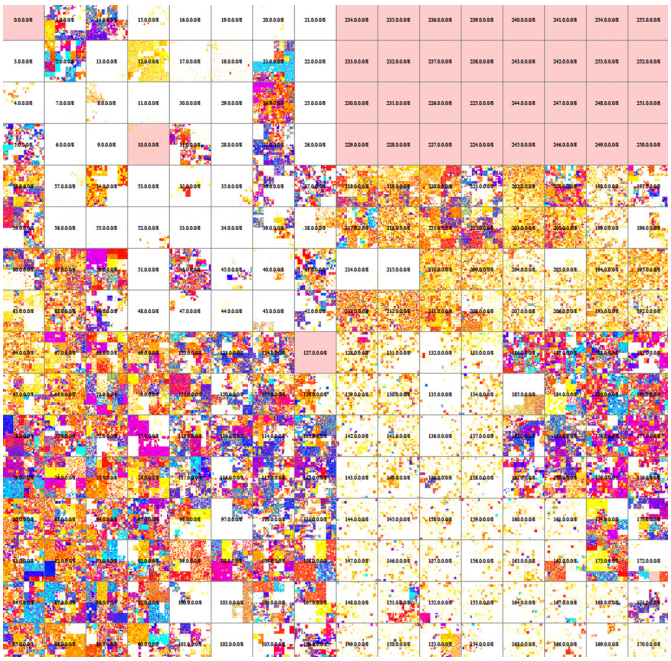

Figure 4: IPv4 Hilbert graph of response to probes as of April 2015

Mobile providers have tackled the IPv4 scarcity problem using so-called "carrier-grade NAT" (CGN). While most Internet-connected devices are routed through some limited private IP space before connecting to an Internet router, the mobile carriers have turned to an altogether industrial version of the same idea, but that industrialization makes for a qualitatively very different Internet. Carrier-grade NAT has created black holes in what was previously a transparent Internet. A single /24 block of IPv4 addresses may handle millions of different customers without discoverability.

CGN networks are essentially private islands on the Internet with a one-way valve for connections to flow outbound. Carriers see commercial benefits of this approach; now, more than ever, mobile providers are looking at "active network management"—a style that only five years ago would have been denounced as both a privacy affront and overt censorship. Not now. Network neutrality lives in a narrow sense, but it is permanently dead for users behind CGN, including essentially all mobile service providers in the US today.

CGN networks do offer an advantage to public IPv4 addressing: devices are not directly discoverable and therefore not directly targetable by Internet-connected attackers. This feature is, however, no panacea—all users within the same CGN network can still reach each other. In other words, governments are not the biggest driver of Balkanization of the public Internet, the mobile providers are. Of course, in countries where the mobile providers are a creature of government, mobile users have never seen a true peer-to-peer, discoverable Internet, and never will.

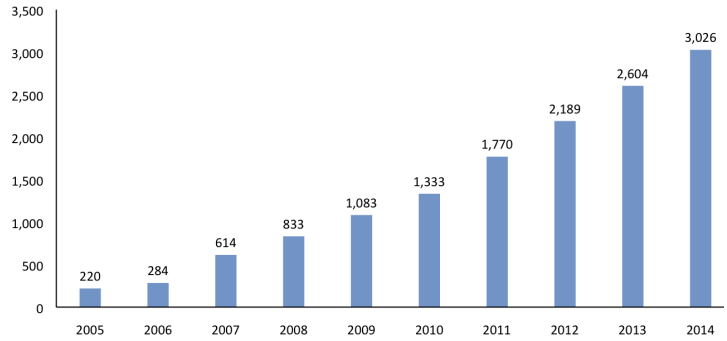

Figure 5: Total broadband users worldwide in millions; CAGR=20.8%

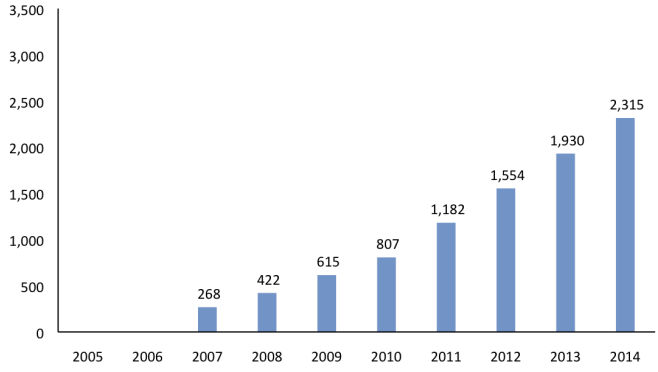

Figure 6: Mobile broadband users worldwide in millions; CAGR=30.8%

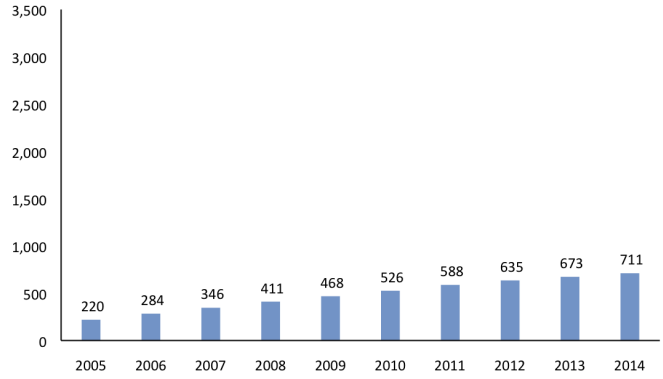

Figure 7: Fixed broadband users worldwide in millions; CAGR=13.1%

Fixed-Line Broadband

Fixed line broadband does continue to increase world-wide, but infrastructure costs have limited its growth to a less aggressive rate than mobile broadband. There are over 700 million fixed-line broadband subscriptions in place as of the end of 2014: the Americas and Europe represent 163 million and 173 million, respectively, while the Asia & Pacific region has skyrocketed to 313 million, as shown in Figures 7 and 8.

US broadband growth is relatively slow compared to Asia but growing consistently all the same. Figure 9 shows the number

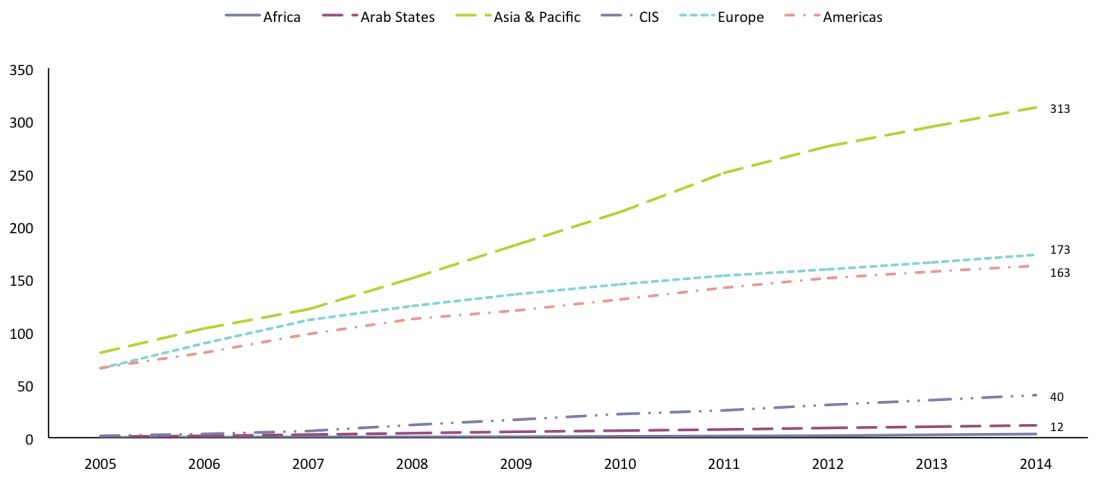

Figure 8: Fixed broadband users by region in millions

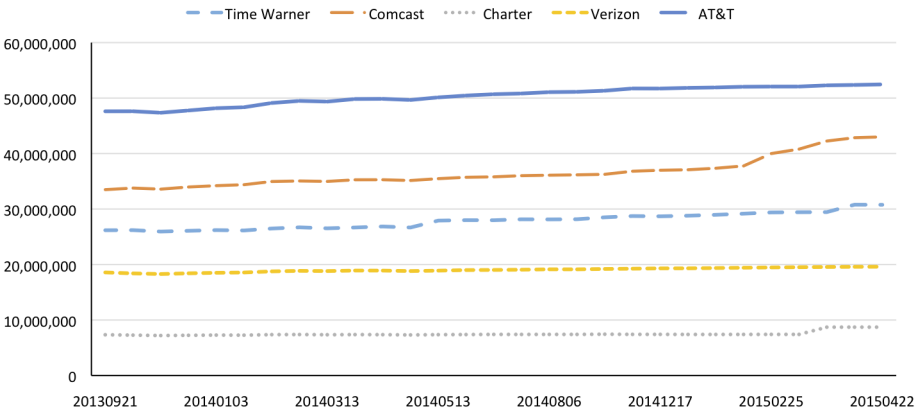

Figure 9: Fixed broadband users by vendor

of IPv4 addresses that correspond to individual US broadband providers between September of 2013 and April of 2015.

In contrast to US mobile carriers, most US fixed-line broadband providers are not using CGN, but instead offer external IP addresses. This provides the (freedom/self-determination) benefit of bi-directional traffic for users at the cost of safety: broadband providers are well known for supplying insecure hardware to their customers, including home routers, TV set-top boxes, and Internet telephony systems. The vast majority of exploitable embedded devices on the IPv4 Internet are ISP-provided systems. Broadband users are rarely given a choice about what equipment they use to connect to the Internet. The end result is that in terms of raw numbers, there are more exploitable broadband devices on the Internet than any other type of system.

Contrary to common belief, populations of vulnerable devices do not always decline with time. In some cases, vulnerabilities can get reintroduced when new hardware is deployed. Figure 10 demonstrates the percentage of devices vulnerable to two stack overflow vulnerabilities in two distinct UPnP software libraries. These libraries are often used in home routers, and both of these vulnerabilities had patches available in 2013. The data shows that the percentage of exploitable devices with UPnP open to the world and exploitable has actually increased; this is the result of broadband ISPs introducing new home gateways that use vulnerable versions of these libraries.

Figure 11 shows another vulnerability that appears to be getting worse over time. In 2014, a configuration weakness was identified in multiple devices regarding the NAT-PMP protocol. This protocol can expose the user's internal network to attack and allow a malicious user to turn vulnerable routers into proxy servers. The continued growth of vulnerable devices can be directly associated with broadband ISP deployments.

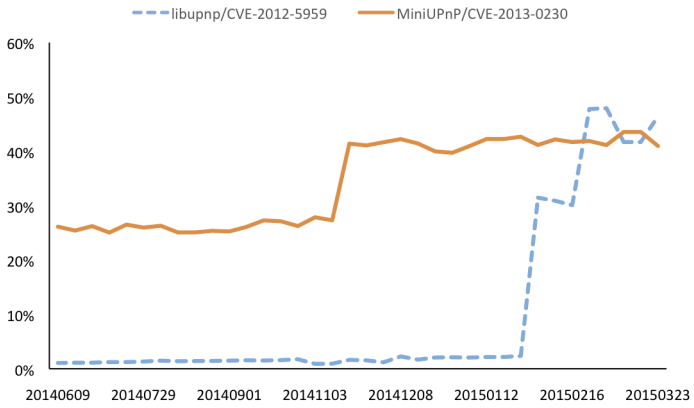

Figure 10: Percentage of devices vulnerable to SSDP over time

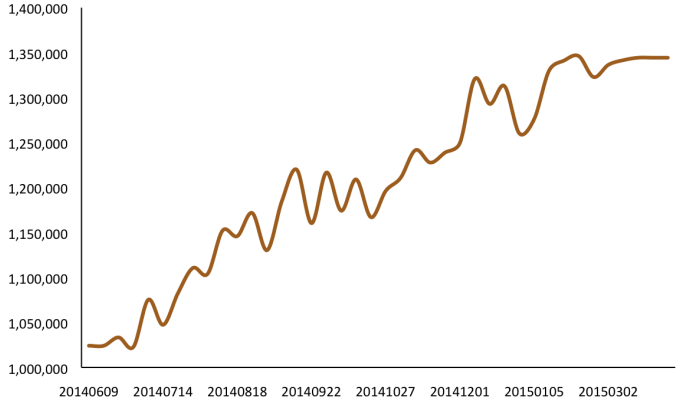

Figure 11: Number of devices vulnerable to NAT-PMP over time

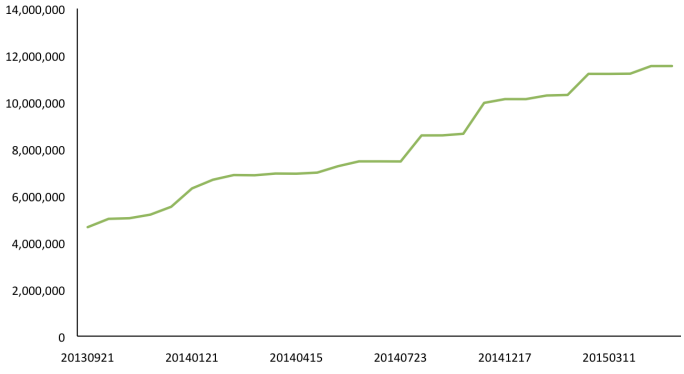

Figure 12: Amazon AWS PTR record allocations over time

These are just two examples. The authors are aware of others, but these two demonstrate how security practices by broadband providers contribute to the overall vulnerability of the Internet. Globally, broadband providers either need to significantly improve the security management of their deployed hardware or provide their users with more control over the devices used. We assume that readers of this column can take care of themselves if given a choice. Those who cannot do so are more numerous, and whose responsibility is that, exactly?

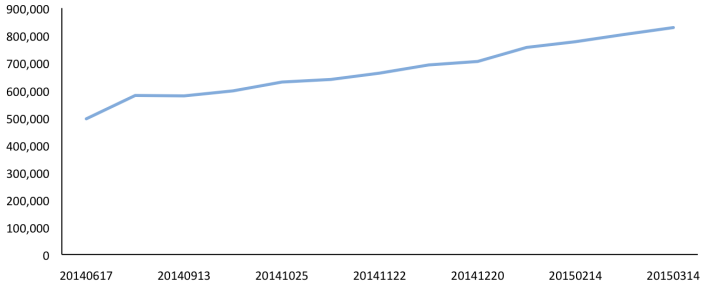

Figure 13: Growth of .com domains using Outlook.com hosted email

Cloud Providers

Businesses have voted with their feet—choosing cloud providers for nearly every aspect of operations. Everything from email to data analytics has been pushed outside of the corporate firewall. In some cases, this is great for security; not every organization has the bandwidth to handle a direct DDoS attack, and external hosting is one way to build a resilient environment. On the other hand, the siren song of on-demand resources fragments an already complex security process. Cloud service providers excel at on-demand scalability, but how they achieve this can be frightening to any CISO.

The difference between a security-conscious provider and an amateur can be hard to distinguish without a deep dive into the provider's operations. For every service provider doing a great job of segmenting customer data and producing secure software, there are dozens that are not. CISOs who resort to questionnaires and live testing when choosing a provider also know that the questionnaire and the testing valid today are obsolete tomorrow.

Traffic to Amazon's EC2 platform now exceeds that reaching Amazon's own storefronts [3]. Hundreds of new SaaS providers are building their infrastructure on top of existing cloud providers. Figure 12 shows the growth of PTR record allocations within Amazon's compute cloud. This figure covers September 2013 to April 2015 and doesn't take into account resources without a public address, such as those hosted within VPCs and exposed through load balancers.

On the email front, thousands of organizations have pushed email outside of their firewall and now depend on services provided by the likes of Google and Microsoft. Figure 13 shows the growth of .com domains that use Microsoft's Outlook.com hosted service. This figure covers June 2014 to March 2015 and shows consistent growth.

Precise, and painfully derived, threat models become irrelevant the minute organizations outsource their core IT functions to the cloud. Visibility is the first casualty; most service providers offer some form of logging or audit function, but the customer is

at the mercy of this implementation, and their hands are often tied if they need to respond to a novel attack. The bigger these service providers grow, the more complicated their support model becomes. As numerous high-level defacements have shown (Twitter, *New York Times*, etc.), one mistake by a low-level support technician undermines the security of the entire platform. An Internet built this way is one vulnerable to cascade failure, and that vulnerability is by design. This is not hardening in the sense of toughening but hardening in the sense of embrittlement. Cloud platform failures have a disproportionate effect on the businesses that depend on them. These failures are infrequent, but have resulted in the economic loss of hundreds of millions of dollars [4].

Summary

A shortage of IPv4 addresses leads to carrier-grade NAT. CGN leads to Balkanization of the public Internet. Consumer demand for downstream content leads to a service-oriented Internet, not a communications-oriented one. The divergence between discoverable assets and overall growth places further blinders on defenders who are already struggling with complexity. Consistently insufficient security management by broadband providers has increased the portion of the Internet that is vulnerable to trivial compromise. Mobile providers offer less targetable enclaves, but at the cost of freedom to connect. Corporate consolidation into cloud providers places ever more eggs into ever fewer baskets. Attackers have adapted—mobile devices are targeted through malicious applications, desktop PCs are at risk from embedded network devices, and cloud providers are the richest hunting ground for corporate secrets. Freedom to connect, the Internet principle of record, led to preferential attachment. Preferential attachment led to innovation and resiliency to random faults. In 2015, carriers and governments alike clearly want non-preferential attachment for end-users: carriers in their desire for economic hegemony, free-world governments in their desire for safety built on attribution, and unfree-world governments in their desire to manipulate information flow.

References

[1] Number of Internet users: http://www.internetlivestats.com/internet-users/.

[2] Users by connection type: http://www.itu.int/en/ITU-D/Statistics/Documents/statistics/2014/ITU_Key_2005-2014_ICT_data.xls; http://www.pewinternet.org/2015/04/01/us-smartphone-use-in-2015/.

[3] Network traffic to Amazon's EC2: http://news.netcraft.com/archives/2013/05/20/amazon-web-services-growth-unrelenting.html.

[4] Downtime due to cloud failures: http://iwgcr.org/wp-content/uploads/2013/06/IWGCR-Paris.Ranking-003.2-en.pdf.