For Good Measure
Nameless Dread

DAN GEER AND PAUL VIXIE

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, includi ng officer positions, program committees, etc. dan@geer.org

Dr. Paul Vixie is an Internet pioneer. Currently, he is the Chairman, CEO, and cofounder of Farsight Security, Inc. He was inducted into the Internet Hall of Fame in 2014 for his work related to DNS. Dr. Vixie began his career as a programmer (Cron, RTTY, BIND) before co-authoring *Sendmail: Theory and Practice* and more than a dozen RFCs, and contributing a chapter to *Open Sources: Voices from the Open Source Revolution*. More recently, he has become a serial entrepreneur (ISC, MAPS, PAIX, MIBH, DNS-OARC, Farsight). He was a member of the ARIN Board from 2004–2013. He completed his PhD in 2010 at Keio University. vixie@fsi.io

> What's in a name? That which we call a rose
> By any other word would smell as sweet;
>
> *Romeo and Juliet*, Act 2, Scene 2

Each generation of global commerce and culture has to decide for itself what the Internet "means" to them. Some of that meaning will depend on how large the Internet is at that time. Delightfully, the unit of measure of that largeness will also change with every era.

There was a time when to be "on the Internet" meant that your host's name was published in a central registry called HOSTS.TXT—and then the wheels came off. The original text-only terminals were replaced by graphical workstations, later by personal computers, then by virtual servers, followed by smartphones, and, eventually, smart devices. But whereas the time-shared minicomputers that once serviced text-only terminals had names, the workstations and personal computers that came later were given names mostly out of habit: we wanted to know where connections to our time-shared computers and servers were coming from, but we would rarely have any reason to try to connect back to those origins.

There was also a time when to be "on the Internet" meant that your IP address block was present in the global routing table. Those wheels also came off pretty early: network address translation (NAT), whether deployed as a security measure or due to a real or perceived shortage of address space, meant that only a small island of a university or enterprise network would use so-called "global addresses," and these would act as gateways to private networks that serviced a much larger population of possible endpoints hidden behind such gateways.

In 2018 ("now") the fashion is to measure the number of connected people and not the number of connected devices. We round this number to the nearest billion, as if we neither know nor require any further accuracy.

Because it's hard to secure something we don't understand, it's necessary that we *fathom* the Internet in some way, so that we can account for and predict the risks it poses and the risks it experiences, and ultimately make some plan as to how to manage some risks and how to cope with others we cannot manage.

Scale

The Domain Survey, operated since 1987 by Network Wizards, Internet Systems Consortium, and 3Way Labs, gives us a general baseline of one measure of Internet size: the population of endpoints that have names. Notably, not all of these names are actually used—many are assigned by network operators from a pool of machine-generated and meaningless names with no expectation that any of these names will ever be used to initiate a connection. This is due to ancient prejudices whereby a service might reject as "low value" any connection from an endpoint lacking a name. Even though this prejudice is wrong, the optics generated by its adherents have helped chart the growth of the Internet to a population size just over

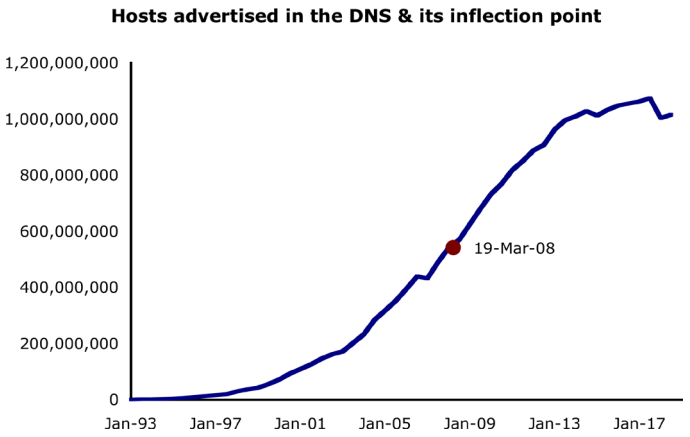

Figure 1: Hosts advertised in the DNS, and its inflection point [1]

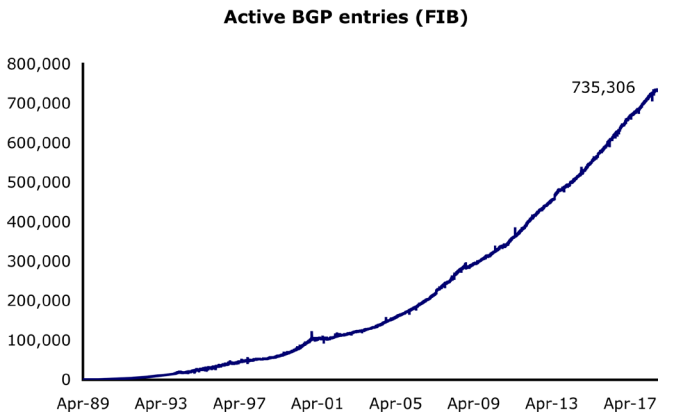

Figure 2: Active BGP entries [2]

one billion "endpoints having names," as seen in Figure 1. We highlight the logistic inflection point, 19-Mar-08, the point in time at which the rate of growth in advertised names changed from accelerating to decelerating. As of today, 66% of the total IPv4 space is advertised as compared to 0.0026% of the total IPv6 space.

Measurement of the Border Gateway Protocol (BGP) global routing table is another proxy for some measurement of the Internet's size. A single entry in this table can contain as few as 256 potential endpoint addresses or as many as 16 million. We can constrain our estimate of the average number of potential endpoint addresses in a routing table entry by noting that about three billion endpoint addresses are globally reachable, there is no new IP version 4 address space remaining in the free pool, and the global routing table contains about 750,000 entries. So a routing table entry represents, on average, perhaps 4000 potential endpoints. In CIDR terms that's a "/20." Notably, many of the smallest routing table entries are just NAT gateways, and so each might represent a vast population of endpoints that could reach outward or accept inbound transactions (see Figure 2).

In 2018, mobile Internet devices such as smartphones began to reach a saturation point—most humans who want or need and can afford a mobile Internet device already have several of them, which means device sales are now principally for upgrades and replacements (see Figure 3). The market is still strong with vigorous competition between handset and platform makers, but the decade of Internet growth driven by new mobile Internet devices may be reaching a plateau. Notably, the vast majority of mobile Internet devices do not have resolvable names since they are only outbound traffic sources and not also inbound traffic sinks. Most do not have fixed addresses and will make their outbound requests from a new address every few minutes due to mobility, roaming, or virtual network grooming.

The fastest source of Internet growth since 2015 is the Internet of Things (IoT), and this is expected to continue, more or less forever (see Figure 4). A "thing" in this context can be a home appliance, an embedded device, or a component in some system like home audio. These devices are cheap to build and cheap to buy, such that very little thought goes into life-cycle management either by producers or consumers of these tiny and plentiful devices. Many such devices are shipped with known or discoverable security vulnerabilities, and many will never be patched whether because of supply chain churn or because the resulting software engineering economics would drive unprofitability. Most importantly, precisely none of these devices have names.

Internet Protocol version 6 (IPv6) has come a long way in a short time, and now represents about a quarter of all inbound traffic seen by Google's network [5]. This fraction is characteristic of other cloud, search, Application as a Service, and social network providers. There is no confident estimate of the relative size of the IPv6 vs. IPv4 endpoint populations due to technical differences in the format and allocation of endpoint addresses between the v6 and v4 systems. Generally speaking, it's easier to enable IPv6 in a new device or online system than to add

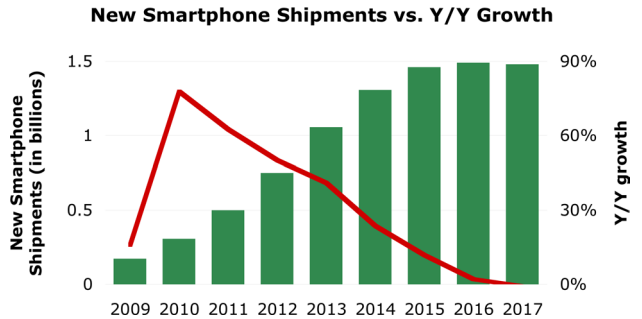

Figure 3: New smartphone shipments vs. year-to-year growth [3]

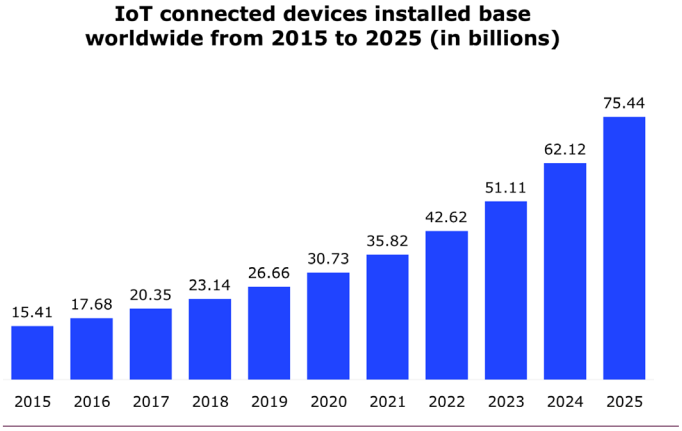

Figure 4: IoT-connected devices installed base worldwide from 2015 to 2025 [4]

IPv6 to an existing IPv4-only system, which dovetails with the trend toward seamless automation without end-user configuration or awareness. While many endpoints from the two largest populations (mobile Internet and IoT) are now using IPv6, most of their traffic is outbound-only, and these devices rarely have or require names. One hopeful difference between the IPv4 and IPv6 systems is that IPv4 addresses are dense enough to permit brute force automatic network scanning by an attacker, so even an endpoint that never advertises its presence and has no name might still be attacked. The sparse addressing of IPv6 makes this kind of attack far more expensive in terms of brute force than for dense IPv4. Of course, security regimes that walk the corporate address space to discover what addresses exist on "their" network are similarly disabled by IPv6's sparseness.

Implications

Security risk is a function of defects and vulnerabilities, exposure, opportunity, and motivation. Factors like the relative motivations and skills of defenders vs. attackers can often be more decisive than the number of defects or the overall reachability of a victim endpoint. However, when other things are equal, as they tend to become in a maturing market with established equilibriums, the best predictors of risk are exposure and reachability. A device that never receives inbound messages from any other device can contribute very little risk. Of course, outbound-only means that it is infeasible to push messages to that device—an auto-update process has to be initiated by the remote device asking to be updated, for example, or a reserve channel has to be secretly designed-in.

We have placed special focus on names because, for security analysis, a name makes a device more reachable, thus increasing its exposure. If successfully attacked, a device will often become a beachhead by which other more private and less reachable devices can be probed and perhaps also successfully attacked, thus increasing the risk posed by the exposed device. Having a name is a risk factor, just as being reachable from outside the local network due to firewall weaknesses or misconfiguration is a risk factor.

Mobile devices can and do join botnets. But the initial vector for a successful attack on such devices will invariably be that it was induced to make an outbound transaction whose results were damaging in some way and against which the device had no working defense. The same will be true of IoT devices for the most part, although in this class of victim, inbound transactions either from the local network or from selected parts of the outside world are part of the product design, and in that case a name, either in the domain name system (DNS) or some other less public naming scheme, will contribute to reachability and therefore to overall risk.

Defenders should consider a mostly closed reachability policy. Nothing should be externally reachable unless there is a hard requirement. This includes both giving an endpoint a globally resolvable DNS name and giving it any kind of reachability in the firewall configuration. But more than this, internal firewalls have to be deployed so that a successful attack on one part of the network does not necessarily create a beachhead for attacks on the rest of the network. This kind of internal segmentation is costly, but at least it's an up-front cost that defenders can budget for—much cheaper than answering questions from the press, customers, shareholders, or regulators after a successful attack—plus whatever damage was actually caused.

There are far-reaching design questions here. One involves the resurrection of a 20-year-old debate: assuming that myriad, nameless devices will need to be able to cryptographically protect their messages, where is the key for that looked up? Will each device have one of its own? Will internal firewalls include a key-centric, rather than a name-centric, PKI of sorts [6]? Does a MAC address or UUID-in-ROM distinguish keys in a nameless world and thus imply an identity-based PKI? Either way, is the key-management job going to be harder or easier absent names? Will we not bother with keys at all and trust that the internal firewalls are resilient to lax operation? Perhaps especially interesting, what would a name mean when the end user has a half-dozen devices that mutually self-synchronize?

Evolution

Because small nameless devices tend to be cloud-associated but typically do not accept inbound transactions or connections, they will (by design) make long-running outbound connections to their maker's command and control infrastructure and simply wait to be told over that connection what action or report they should make next. The identity of the device might be encoded as a client-side TLS certificate or some hardware serial number.

The command and control service will associate the device's identity with a subscriber, and when the subscriber also connects in, this elbow-shaped pair of connections will allow the subscriber to apparently but indirectly control their device. This synchronization-design language is both the result and supporter of the trend toward namelessness in modern Internet-connected devices. Even where direct LAN-based connectivity is used to connect a subscriber to a device, it will as often be negotiated through the maker's command and control network, as discovered locally by some broadcast or multicast protocol along the lines of mDNS or UPnP. Whatever the motive or method, the universal consensus among system designers is that using names to reach Internet-connected devices is considered a legacy. Services need names; servers who provide those services need names; devices which are not servers, will be reached in other ways.

We are at a fork in the road. The choices to be made will be expensive to later reverse in either dollars or clock-ticks. Momentum says that, soon, the majority of Internet endpoints will not be describable by name or discoverable by scanning. Another layer of indirection will, as ever, solve some problems and create others. Provenance and forensics will all but surely be affected. The CAP theorem [7] is licking at our heels.

References

[1] Internet Domain Survey, July 2018: ftp.isc.org/www/survey/reports/2018/07/index.html.

[2] Active BGP entries: http://bgp.potaroo.net/as2.0/bgp-active.txt.

[3] New smartphone shipments vs. growth: https://www.slideshare.net/kleinerperkins/internet-trends-report-2018-99574140, slide 6.

[4] IoT-connected devices: https://www.statista.com/statistics/471264/iot-number-of-connected-devices-worldwide/.

[5] Proportion of IPv6 traffic: https://www.internetsociety.org/wp-content/uploads/2018/06/IPv6-infographic.pdf; https://www.google.com/intl/en/ipv6/statistics.html.

[6] Name-Centric PKI (Ellison & Metzger) vs. Key-Centric PKI (Ford & Kent): http://static.usenix.org/publications/library/proceedings/ec98/pki.html.

[7] Simon S. Y. Shim, "The CAP Theorem's Growing Impact," *IEEE Computer*, vol. 45, no. 2, February 2012, pp. 21–22: https://www.computer.org/csdl/mags/co/2012/02/mco2012020021.pdf.