

US Cities Exposed

A Shodan-Based Security Study of Exposed Assets in the US

Numaan Huq, Stephen Hilt, and Natasha Hellberg
Trend Micro Forward-Looking Threat Research (FTR) Team

TREND MICRO LEGAL DISCLAIMER

The information provided herein is for general information and educational purposes only. It is not intended and should not be construed to constitute legal advice. The information contained herein may not be applicable to all situations and may not reflect the most current situation. Nothing contained herein should be relied on or acted upon without the benefit of legal advice based on the particular facts and circumstances presented and nothing herein should be construed otherwise. Trend Micro reserves the right to modify the contents of this document at any time without prior notice.

Translations of any material into other languages are intended solely as a convenience. Translation accuracy is not guaranteed nor implied. If any questions arise related to the accuracy of a translation, please refer to the original language official version of the document. Any discrepancies or differences created in the translation are not binding and have no legal effect for compliance or enforcement purposes.

Although Trend Micro uses reasonable efforts to include accurate and up-to-date information herein, Trend Micro makes no warranties or representations of any kind as to its accuracy, currency, or completeness. You agree that access to and use of and reliance on this document and the content thereof is at your own risk. Trend Micro disclaims all warranties of any kind, express or implied. Neither Trend Micro nor any party involved in creating, producing, or delivering this document shall be liable for any consequence, loss, or damage, including direct, indirect, special, consequential, loss of business profits, or special damages, whatsoever arising out of access to, use of, or inability to use, or in connection with the use of this document, or any errors or omissions in the content thereof. Use of this information constitutes acceptance for use in an "as is" condition.

Contents

4

Exposed Cyber Assets

6

Exposed Cities

36

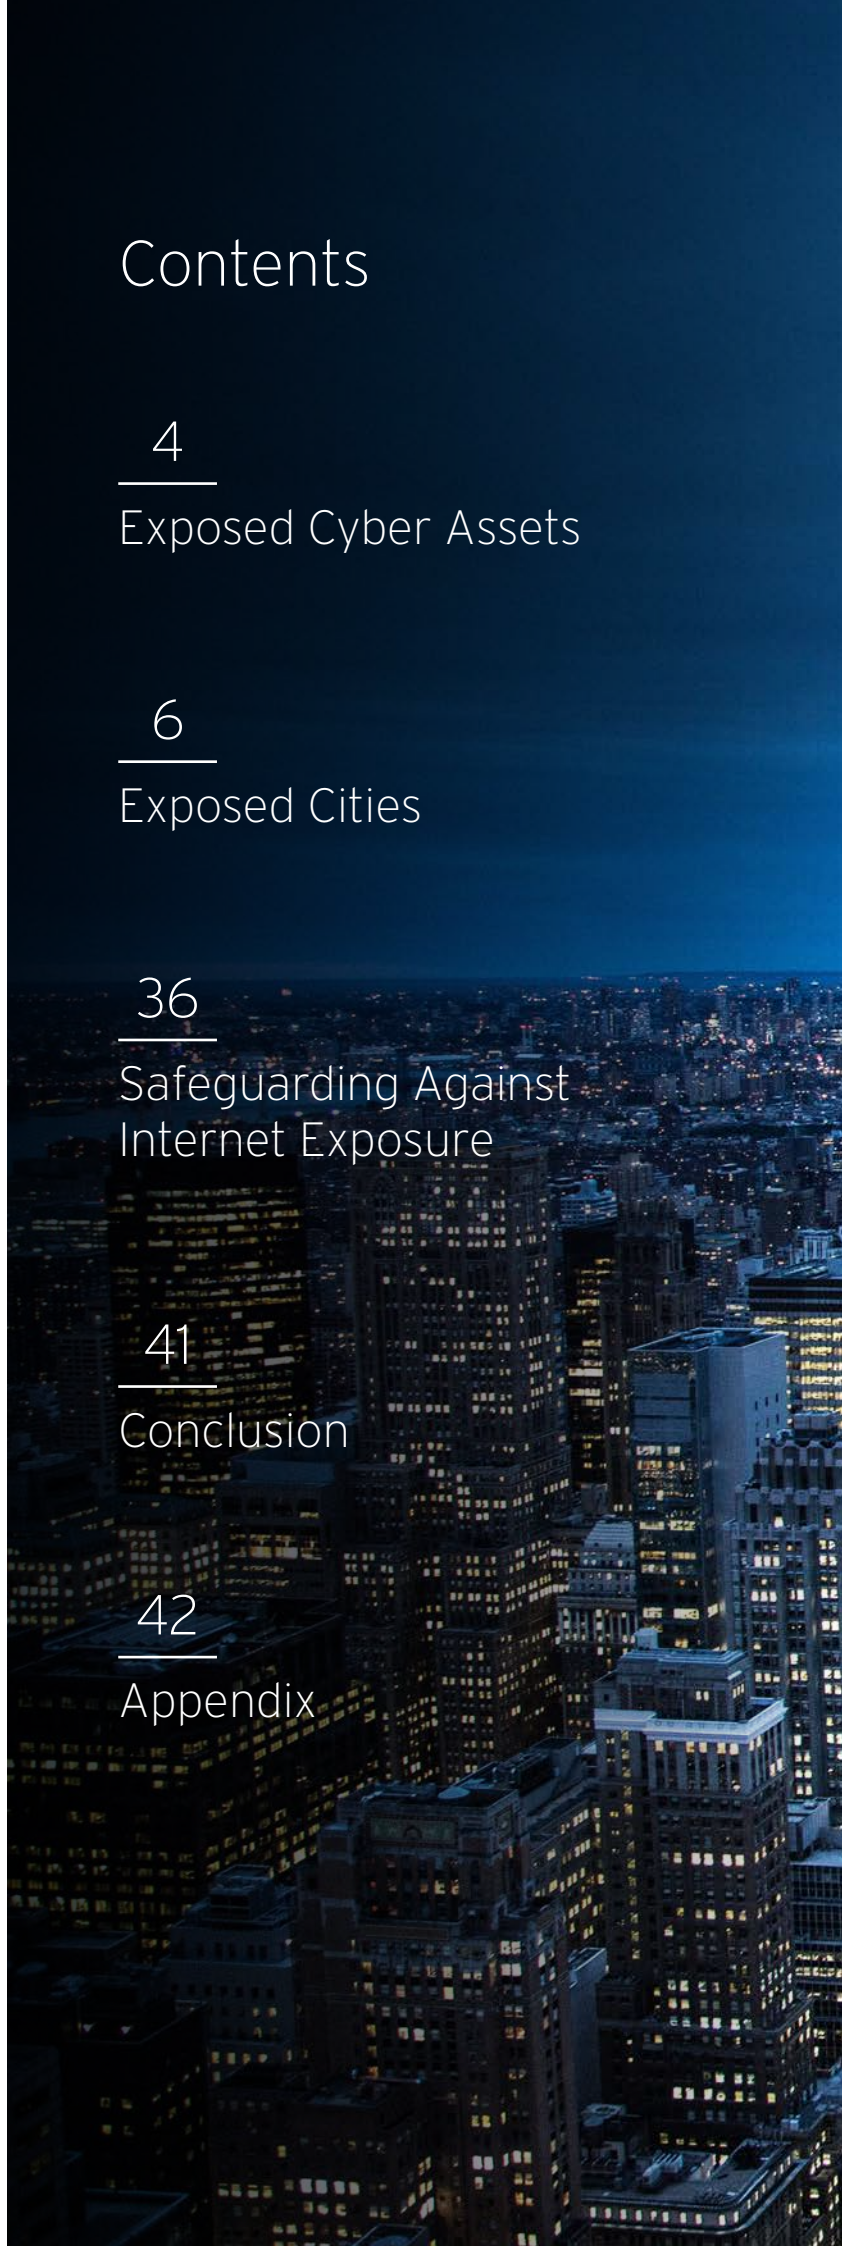
Safeguarding Against
Internet Exposure

41

Conclusion

42

Appendix



DISCLAIMER: At no point during this research did we perform any scanning or attempt to access any of the Internet-connected devices and systems. All published data, including screenshots, were collected via Shodan. Note that any mention of brands in this research does not suggest any issue with the related products but only that they are searchable in Shodan.

The Internet of Things (IoT) is fast becoming the new norm, connecting everything from computers, mobile devices, cars, industrial robots, home appliances, and even smart clothing to the Internet. This interconnected world is very exciting and has created new and unique opportunities to improve our lives. But truth be told, today's society is adopting connected technologies at a faster rate than we are able to secure them. Caution dictates that in addition to exploring new opportunities with IoT, we also examine the implications and repercussions of an all-devices-online world. There is a strong likelihood that some of our Internet-connected devices and systems may be inadvertently exposing information about us and our surroundings online, and that could potentially jeopardize everyone's safety and security.

The main goal of this research paper is to build public awareness about exposed cyber assets and highlight problems and issues associated with them. We define “exposed cyber assets” as Internet-connected devices and systems that are discoverable on Shodan or similar search engines, and can be accessed via the public Internet.

Several research papers and conference talks have been published and presented that explore these problems and issues, but in this paper, we study exposed cyber assets from the macroscopic perspective of cities. The exposed cyber assets profiled refer to all of the popular Internet-connected devices and systems in large US cities, and allows us to do comparative analysis of cities with similar population sizes. In a follow-up research, “**US Cities Exposed: Industries and ICS**,”¹ we profiled exposed cyber assets that are critical to daily city operations (i.e., critical infrastructure and industrial control systems [ICS]).

Research results revealed a significant number of exposed devices such as webcams, network-attached storage (NAS) devices, routers, printers, phones, and media players, many of which are vulnerable to exploitation and compromise. We also found a significant number of Web and email servers, along with databases, including medical databases, which could potentially be compromised by determined threat actors. Finally, we profiled several vulnerabilities that the Shodan crawler scans for, and if these vulnerabilities remain unpatched then attackers can exploit them to compromise underlying systems.

While the connected world struggles with questions about who is responsible for safeguarding and policing exposed cyber assets, how it should be done, and what awareness campaigns must be run to better protect cyber infrastructure, we provide some guidance by outlining a set of security best practices for businesses and home users to follow that will help them secure their Internet-connected devices against potential attacks.

Exposed Cyber Assets

Traditional Web search engines such as Google, Bing®, and Yahoo!® are great if you are looking for information and websites, but not so good if you are searching for device metadata. The solution? Shodan, a publicly available search engine for Internet-connected devices and systems. Shodan finds and lists devices and systems such as webcams, baby monitors, medical equipment, ICS devices, home appliances, databases, and others. In short, Shodan collates and makes searchable both device metadata and banner information (i.e., services running) that Internet-connected devices and systems are freely sharing with anyone who queries them. A majority of these require Internet access to function properly though some such as ICS and medical devices should never be directly connected to the Internet. If not properly configured, then by virtue of being exposed on the Internet, some of these devices and systems may be vulnerable to compromise and exploitation. There is also the elephant in the room—privacy; what, if any, sensitive information is being exposed online?

We define “exposed cyber assets” as Internet-connected devices and systems that are discoverable on Shodan or similar search engines, and can be accessed via the public Internet. Important questions that come to mind are:

- What potential risks are associated with exposed cyber assets? Risks include:
 - Exposed cyber assets could get compromised by hackers who steal sensitive data (e.g., personally identifiable information [PII], intellectual property, financial and corporate data, etc.).
 - Exposed cyber assets could be leaking sensitive data online without their owners’ knowledge (e.g., open directories on Web servers, unauthenticated webcam feeds, exposed ICS Human Machine Interfaces [HMIs], etc.).
 - Hackers use lateral movement strategies to gain entry into a corporate or an ICS network by compromising exposed cyber assets then commit espionage, sabotage, or fraud.
 - Compromised cyber assets can be used to run illegal operations such as launch distributed denial-of-service (DDoS) attacks, become part of botnets, host illegal data, be used for fraud, and so on.

- Compromised cyber assets can be held hostage for ransom. This is especially damaging if they are critical to an organization or individual's operations.
- Cyber assets that operate critical infrastructure can jeopardize public safety if compromised.

We also compiled a list of recent notable cyber intrusions in the Appendix, some of which demonstrate the real-world risks that exposed cyber assets pose.

- Why are cyber assets exposed on the Internet? Common reasons for device and system exposure online include:
 - Incorrectly configured network infrastructure that allow direct device or system access
 - Devices and systems need to be Internet connected in order to function properly
 - Remote access is enabled on devices and systems for remote troubleshooting
 - Remote access is enabled on devices and systems for remote operations
- Who is targeting exposed cyber assets? Threats come from a variety of sources, depending on the types of cyber assets targeted. Actors include:
 - Nation-states, both developed and developing, gather intelligence using software espionage tools and customized malware.
 - Criminal syndicates include both criminal gangs who target consumers using different schemes such as ransomware to profit and those contracted by national governments for various political cyber attacks, including cyber espionage and subterfuge.
 - Cyberterrorists launch disruptive or destructive cyber attacks to cause physical destruction of property or potential loss of life and spread fear.
 - Competitors look for information in order to gain strategic advantages over others in the industry.
 - Hacktivists or Internet activists attack cyber assets to draw attention to their causes.
 - Script kiddies represent the vast majority of threat actors who scan the Internet to discover exposed IoT devices either out of curiosity or to cause mischief.

Today's digital warfare is asymmetrical with falling costs for those bent on disruption and fixed or increasing costs for the society disrupted. The cost of finding and exploiting critical infrastructure will continue to fall. The marginal cost of copying vulnerable infrastructure lists or exploits will tend toward zero. The cost of causing disruptions for hackers will continue to fall while that of disruption remediation will remain relatively constant or increase.²

Exposed Cities

Scanning the Internet is important because security flaws can be quickly identified or discovered and fixed before they are exploited. But scanning the Internet is difficult, time-consuming, and poses a set of unique challenges. For our research on exposed cyber assets, we bypassed all of the issues or hurdles and simply used a public data source—Shodan. Technical assumptions and observations about our use of Shodan data in this project can be found in the Appendix that discusses what Shodan is and how we analyzed the Shodan data.

We examined the Shodan US scan data for February 2016. The data set contains a total of 178,032,637 records generated from scanning 45,597,847 unique IPv4 and 256,516 unique IPv6 addresses. The raw scan data was indexed using Elasticsearch and queried using Kibana, which allowed us to search more than 550 fields versus more than 40 fields using Shodan’s Web interface. In this research, we present data on exposed cyber assets in the top 10 largest US cities by population—New York City, Los Angeles, Chicago, Houston, Philadelphia, Phoenix, San Antonio, San Diego, Dallas, and San Jose. The cities were selected using the 2010 US Census data.³ We excluded cloud service providers such as Amazon, Azure, Akamai, CloudFlare, and others from the queries so we can focus on “actual” connected versus online virtual devices. It is also worth noting that not all fields in every scan record were populated (e.g., not every record has the device field populated).

Cyber Asset Exposure Statistics in the Top 10 US Cities by Population

This section provides a general overview of cyber asset exposure numbers and all types of exposed devices, systems, products, operating systems (OSs), and other assets that are visible in the February 2016 Shodan US scan data for the top 10 US cities by population.

Exposed Cyber Assets in the 10 Largest US Cities by Population

It is interesting to note that the volume of exposed cyber assets in large US cities can be disproportionate to their population size. For example, the February 2016 Shodan US scan data shows 3,900,208 exposed cyber assets in Houston, Texas compared with 1,031,325 in New York City, New York. New York City has a far bigger population than Houston, yet it has 3.78 times fewer exposed cyber assets compared with Houston.




Figure 1: Number of exposed cyber assets in the 10 largest US cities by population

How Are Exposed Devices Connected to the Internet?

It is not surprising that most devices are connected to the Internet via modems. Interestingly, we also saw devices connected via virtual private networks (VPNs) and virtual LANs (VLANs) in the Shodan scan data. Should not these devices be private and not respond to queries from the Shodan crawler? Google Fiber™ is slowly being rolled out to many US cities so it is also not surprising to discover Google Fiber network boxes in the Shodan data.




Figure 2: Distribution of means by which devices access the Internet

What OSs Run on Exposed Internet-Connected Devices?

Devices that run Linux® dominated in terms of OS found by the Shodan crawler. These are predominantly IoT devices that run embedded Linux though a fair number of Web servers that run Linux, Apache, MySQL, PHP (LAMP) are also in the mix. The Windows® OS family was also, unsurprisingly, largely prominent. Mac OS X exposure was negligible compared with that of Linux and Windows devices.




Figure 3: Distribution of exposed device OSs

Opportunistic attackers can take this observation as insight into what OS they should focus on finding vulnerabilities for if they want to ensure a broad victim base.

Top 20 Exposed Products

As expected, the list of exposed products (not to be confused with that of device types, which we will cover later) is dominated by Web servers. Shodan also discovered large numbers of Internet-facing MySQL, Simple Mail Transfer Protocol (SMTP), Secure Shell (SSH), and File Transfer Protocol (FTP) servers. Compared with desktops, servers are more vulnerable to zero-day exploits because when compromised, they can be leveraged to attack users that connect to them. On the flip side, a vast majority of daily cyber attacks use weaponized exploits that have been around for a long time instead of zero-day exploits. Administrators should regularly apply security patches to servers in order to prevent hackers from exploiting known patched vulnerabilities.




Figure 4: Top 20 exposed products

Top 20 Exposed Vulnerable Products

The Shodan crawler tests for specific vulnerabilities—CVE-2013-1391 (digital video recorder [DVR] configuration disclosure), CVE-2013-1899 (argument injection in PostgreSQL), CVE-2014-0160 (Heartbleed, OpenSSL), CVE-2015-0204 (Freak, OpenSSL), and CVE-2015-2080 (Jetty remote unauthenticated credential disclosure). It is good to see that aside from a handful, the vast majority of servers scanned by Shodan are patched against these vulnerabilities. Compared with the total number of servers scanned by Shodan, the number of vulnerable servers is negligible. In a targeted attack, threat actors attempt to identify vulnerabilities in the exposed product and use that knowledge to craft social engineering attacks.

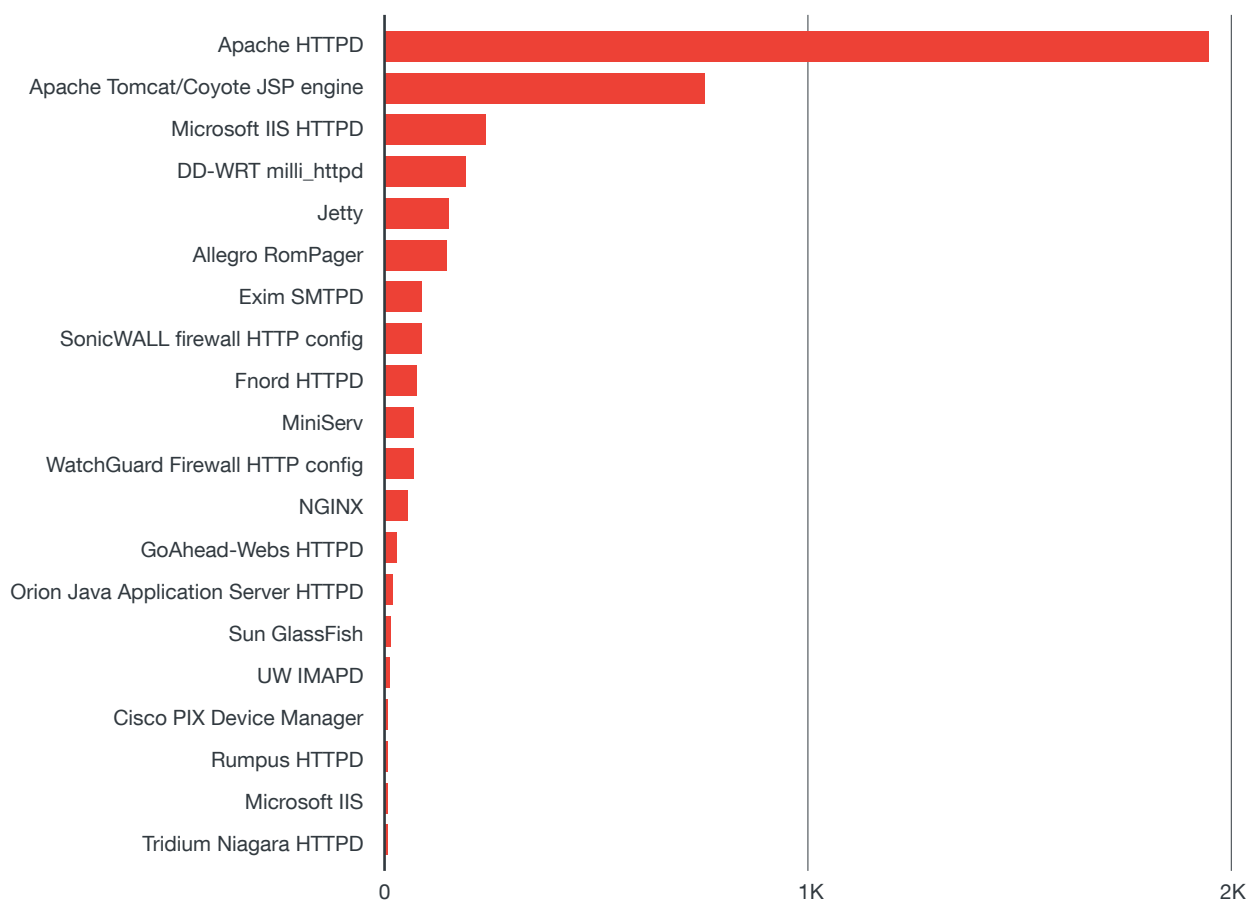


Figure 5: Top 20 exposed vulnerable products

Top 20 Exposed Device Types

Firewalls, webcams, wireless access points (WAPs), printers, routers, and phones dominated the exposed device types seen. The admin interface of the firewall is exposed and this is how Shodan identifies it as such. Attackers can attempt brute-force attacks to gain entry into the firewall’s admin interface and, once inside, change the firewall rules to allow malicious traffic into the network. We also discovered a good number of exposed storage devices, most probably NAS devices. The recent DDoS attack against KrebsOnSecurity.com used compromised routers, webcams, and DVRs to generate a massive volume of network traffic directed at the website.⁴




Figure 6: Top 20 exposed device types

Exposed Devices in Top 10 US Cities by Population

This section digs deeper into exposed devices such as webcams, NAS and media devices, routers, printers, and phones, visible in the February 2016 Shodan scan data for the top 10 US cities by population. Exposed devices are at risk of data theft, lateral movement, forced participation in DDoS attacks, and other threats.

Exposed Webcams

In the public’s perception, it seems that exposed cyber assets are synonymous with exposed webcams. This is probably because webcams are easily visible in homes, public places, retail stores, and so on; easy to find online; and extensively used in everyday devices such as phones and laptops. Webcams typically run a light HTTP or HTTP Secure (HTTPS) Web server that allows users to log in and use them. Shodan data shows that three webcam models dominate the results—security camera manufacturers GeoVision and Avtech and home webcam maker D-Link.




Figure 7: Number of exposed webcams




Figure 8: Distribution of exposed webcams by product name

Searching in the National Vulnerability Database (NVD),⁵ we found eight vulnerabilities that directly or indirectly affect D-Link cameras, five that directly or indirectly affect GeoVision cameras, and only three that directly or indirectly affect AVTECH cameras. Just because there is only a small number of known vulnerabilities does not make webcams safe to use. Webcams are rarely patched and most do not have auto-update functionality. This means webcams will remain vulnerable for months or even perpetually after being sold. The Achilles heel of webcams—users do not change their default passwords or use weak passwords that are vulnerable to brute-force or dictionary attacks.

Exposed NAS Devices

NAS devices are popular solutions for sharing files in collaborative work environments, system backups, and data storage. We did not find a lot of exposed NAS devices in the US cities that we profiled probably because either they are not widely used or they have been secured against accidental online exposure. NAS devices are typically used to back up or store important data, making them attractive targets for hackers.




Figure 9: Number of exposed NAS devices

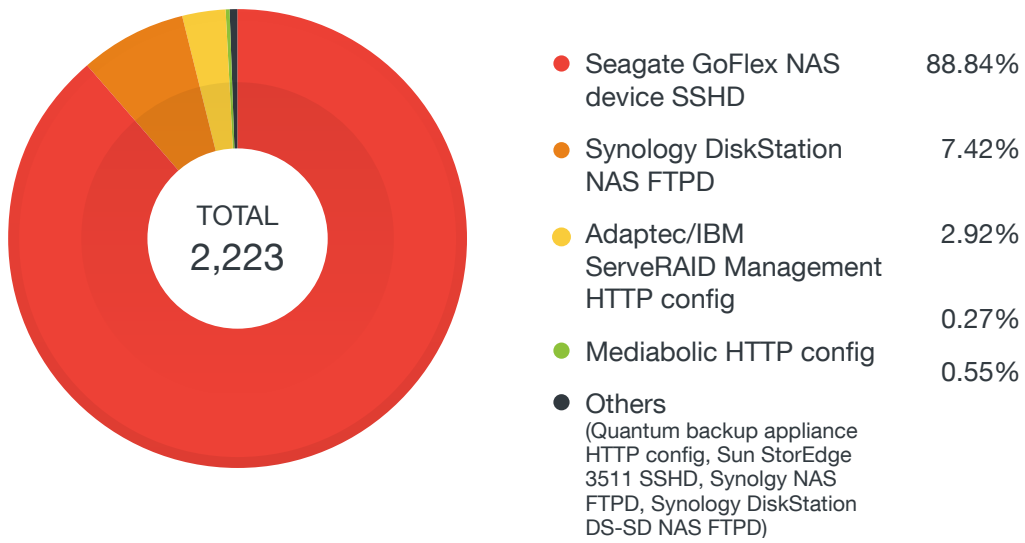


Figure 10: Distribution of exposed NAS devices by product name

Exposed Routers

Routers are present in every home and office that has an Internet connection. Router vulnerabilities are regularly discussed in computer security conferences. Talks disclose new firmware vulnerabilities and how they can be exploited. The security researchers who discovered these vulnerabilities also reach out to router manufacturers who are sometimes slow to respond. After manufacturers release a firmware upgrade or security patch, only a small number of users actually install the fixes.




Figure 11: Number of exposed routers




Figure 12: Distribution of exposed routers by product name

Compromised routers can be used to generate network traffic in DDoS attacks, redirect users to malicious websites that steal credentials, or try to install malware on a user's computer. Cisco routers, which dominate the Shodan results, are typically installed by Internet service providers (ISPs) in customers' homes. Linksys and D-Link are the two most popular home router brands sold in the market. Given that every household with an Internet connection has a router, we were surprised to find so few routers in the Shodan data. We think this could be because the queried routers did not respond to the Shodan crawler or the crawler failed to identify the device as a router.

Exposed Printers

Why would anyone attack and compromise printers? Consider the different types of print jobs an office printer handles on a daily basis—documents containing intellectual property; PII; and financial, customer, and sales data, among others. This scenario is not restricted to only office printers; users print sensitive information on their home printers (e.g., paper copies of e-tax returns, bank statements, travel itineraries, tickets, etc.). At the end of the day, a printer is a computer on a network that stores cached copies of the documents it printed, which makes it a treasure trove of sensitive data that hackers want to steal. Compromised printers can also be used for lateral movement within a target network, to generate network traffic, and to participate in DDoS attacks.

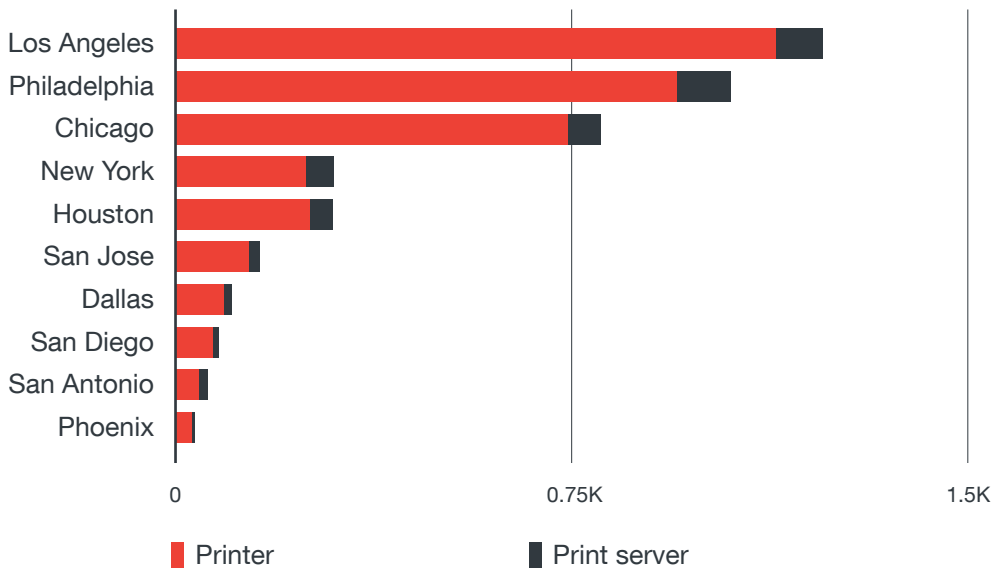


Figure 13: Number of exposed printers




Figure 14: Distribution of exposed printers by product name

Exposed Phones

The world is more connected today than it was a decade ago. Today's companies compete in the global marketplace and instant communication is the key to success. VoIP technology makes making phone calls (both local and overseas) cheap. Thus, many companies are switching to VoIP phones. For this reason, we discovered that Free Private Branch Exchange (FPBX) devices—telephone systems within enterprises—dominate the Shodan results.




Figure 15: Number of exposed phones

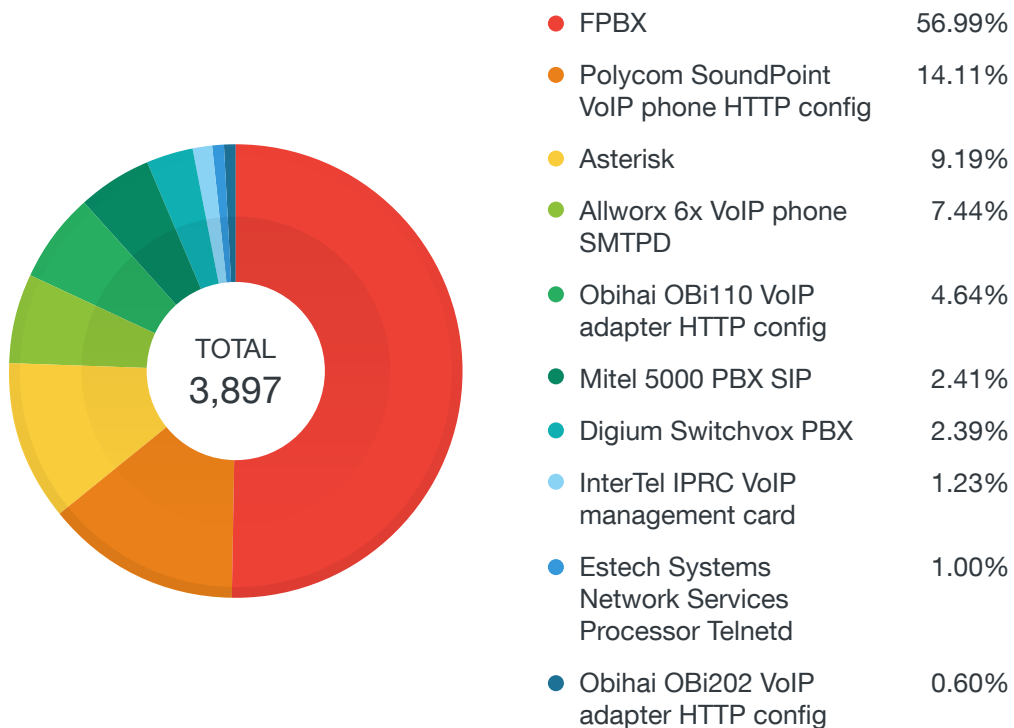


Figure 16: Distribution of exposed phones by product name

Compromising an organization’s telephone system allows hackers to monitor where calls are placed and by whom, eavesdrop on calls, access stored voice mail messages, and, in extreme cases, disrupt voice communications, which may have adverse effects on daily operations.

Exposed Media Devices

TiVo DVRs dominated the Shodan results. A big list of other DVR brands was also exposed online. Exposed DVRs are a security risk for three major reasons—closed-circuit television (CCTV) video feeds are stored in DVRs (for a hacker, these provide valuable surveillance information on targets), compromised DVRs can be leveraged as a point of entry into a corporate network, and compromised DVRs can be used to generate network traffic in DDoS attacks.




Figure 17: Number of exposed media devices




Figure 18: Distribution of exposed media devices by product name

Exposed Servers and Databases in Top 10 US Cities by Population

This section digs deeper into exposed servers and databases such as Web and email servers, and general and medical databases visible in the February 2016 Shodan US scan data for the top 10 cities by population. Server and database exposure puts users at risk of data theft, lateral movement, fraud, and other threats.

Exposed Web Servers

Web servers are Internet facing by design so they can serve browsers the requested data and Web pages. Why do we care about exposed Web servers? Web servers such as Apache, Internet Information Services (IIS), and others are riddled with vulnerabilities that hackers can exploit to compromise them. A quick search in NVD shows 963 vulnerabilities that directly or indirectly affect Apache and 202 vulnerabilities that directly or indirectly affect Microsoft IIS servers. Apache and IIS are the two most popular Web servers in use today. NGINX is another popular Web server. Because of its small memory footprint, NGINX can be run as an embedded Web server. A compromised Web server can be used to redirect visitors to malicious websites, serve malicious content, host illegal data, and so on.




Figure 19: Number of exposed Web servers

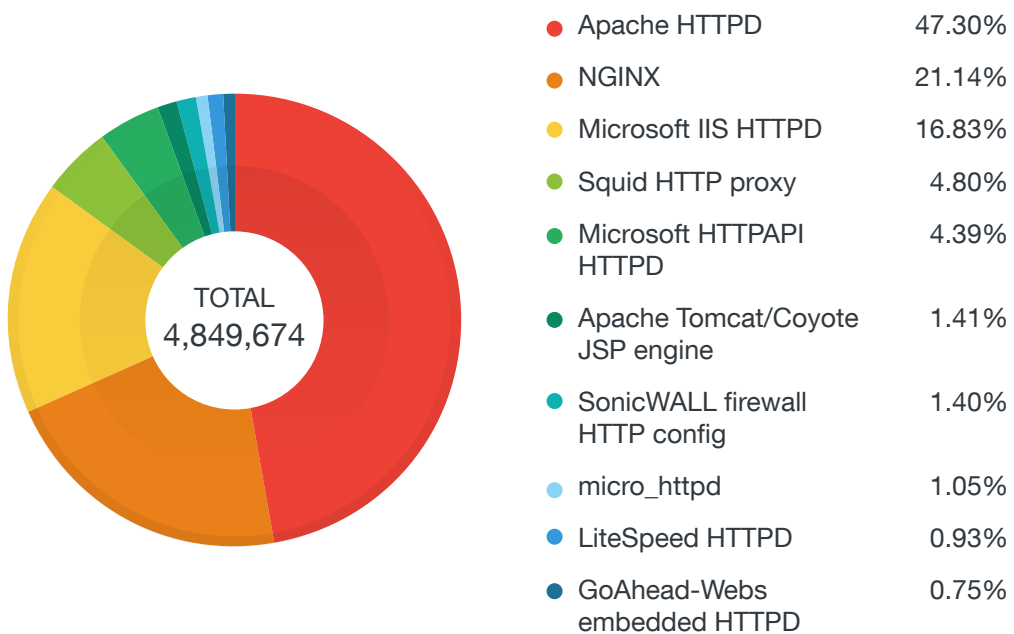


Figure 20: Distribution of exposed Web servers by product name

Exposed Email Servers

Email servers are Internet facing by design so they can send and receive emails. It is interesting to discover that *NIX-based email servers, Exim and Postfix, dominate the Shodan results. We were expecting to find a greater number of Microsoft Exchange servers. Email is one of the main communication tools for modern businesses; a compromised email server means hackers have access to business-critical data (e.g., PII, internal documents, client communication, sales information, etc.). Also, any disruption to email services will severely affect daily business operations. Compromised personal email accounts can lead to the theft of PII, photos, financial information, credentials, and other sensitive information, inflicting damage to the affected individuals.




Figure 21: Number of exposed email servers




Figure 22: Distribution of exposed email servers by product name

Exposed Databases

Databases are the engines of modern business operations. They are used for storing financial, customer, sales, and inventory data; PII; credentials; and other information used by business applications. Databases are treasure troves of critical, sensitive, or important data, which makes them lucrative targets for hackers. For this reason, database theft incidents (where full database dumps are stolen) are regularly mentioned in news about hackers attacking organizations. Recently, we have seen cybercriminals encrypting Internet-exposed MongoDB databases and demanding ransom payment for the decryption key.⁶ From the Shodan data, we found that MySQL was the most popular database exposed on the Internet while MS-SQL and PostgreSQL had comparatively smaller exposure numbers. Banner information returned by MongoDB includes stored table names, which makes it easy to figure out what type of data is stored in the exposed MongoDB databases.




Figure 23: Number of exposed MySQL databases




Figure 24: Number of exposed PostgreSQL databases




Figure 25: Number of exposed MongoDB databases




Figure 26: Number of exposed MS-SQL databases

Exposed Medical Databases

A picture archiving and communication system (PACS) is a medical imaging technology that provides economical storage and convenient access to images from multiple modalities (sources such as Computerized Tomography [CT], X-ray, Magnetic Resonance Imaging [MRI], ultrasound, endoscopy, and other machines).⁷ Digital images and their reports are electronically transmitted via PACSs, giving doctors and physicians instant access to medical imaging results. An electronic health record (EHR) is a digital version of a patient’s paper chart. EHRs are real-time, patient-centric records that make information available instantly and securely to authorized users.⁸ The terms, “EHR” and “electronic medical record (EMR)” refer to the same thing and are often used interchangeably.




Figure 27: Number of exposed PACSs




Figure 28: Distribution of exposed PACSs by server software




Figure 29: Number of exposed EHR and EMR servers




Figure 30: Distribution of exposed EHR and EMR servers by product name

EHRs are shared across different healthcare settings (e.g., laboratories, clinics, hospitals, doctor's offices, etc.) through network-connected enterprise wide information systems or other information networks and exchanges. They may include a wide range of data such as medical history, medications, allergies, immunization status, laboratory test results, medical images, vital statistics, insurance information, and so on.⁹ PACSs and EHR systems are some of the foundation technologies of today's e-healthcare systems. A recent report on healthcare data breaches, prepared for the US Senate by cybersecurity think-tank ICIT, discusses in great detail how stolen patient records are monetized in Deep Web marketplaces.^{10, 11} Given the importance of patient health records, we searched for PACSs and EHR systems in the Shodan US scan data. It is not surprising that we found so many instances of PACSs and EHR systems exposed online. These systems mostly run on Apache, Microsoft IIS, or NGINX. The big question is, "Shouldn't PACSs and EHR systems be operating inside dedicated healthcare or hospital networks, and not exposed online?" Apache and IIS have plenty of known vulnerabilities that hackers can exploit to compromise exposed systems. We believe these systems are being compromised regularly, which is why so many patient health records end up for sale in Deep Web marketplaces.

Additionally, healthcare organizations that unknowingly expose patient data may be liable for civil violations under the Health Insurance Portability and Accountability Act of 1996 (HIPAA).¹²

Exposed Services in the Top 10 US Cities by Population

This section digs deeper into exposed services such as Network Time Protocol (NTP), Universal Plug and Play (UPnP) or Simple Service Discovery Protocol (SSDP), Simple Network Management Protocol (SNMP), SSH, Remote Desktop Protocol (RDP), Telnet, and FTP visible in the February 2016 Shodan US scan data for the top 10 cities by population. Vulnerabilities in the said protocols can be exploited to successfully compromise the devices or systems running them.

Los Angeles		Houston		San Jose		New York		Chicago	
Port	Count	Port	Count	Port	Count	Port	Count	Port	Count
80	1,112,164	443	502,000	80	355,389	80	255,977	80	508,148
3306	264,873	80	456,260	443	116,081	443	130,254	443	298,703
7547	256,912	7547	320,782	3306	85,317	4567	93,256	7547	247,319
443	241,343	53	240,645	7547	83,126	22	73,745	53	187,024
22	205,343	110	185,006	22	52,917	7547	40,261	22	153,046
53	134,473	143	182,756	500	35,189	500	36,566	143	136,078
110	128,142	3306	168,237	23	33,534	4500	33,394	110	133,779
143	125,886	26	138,980	21	28,063	111	33,058	3306	96,103
8080	103,047	8080	128,024	4500	27,648	25	30,528	993	88,680
25	94,884	993	126,746	8080	19,232	8080	21,652	995	87,731

Dallas		San Antonio		San Diego		Philadelphia		Phoenix	
Port	Count	Port	Count	Port	Count	Port	Count	Port	Count
80	526,128	80	190,741	443	145,228	7547	92,330	80	546,959
443	309,742	443	170,719	80	98,815	80	77,444	4567	153,486
53	146,689	7547	134,382	53	85,894	443	47,273	443	126,617
7547	136,427	22	89,318	7547	77,827	4567	46,257	3306	107,821
110	122,829	25	28,719	25	30,526	22	35,530	22	99,809
143	120,149	21	23,092	110	25,741	500	11,838	21	69,155
22	113,839	3306	18,927	143	25,183	23	10,877	8080	39,501
25	85,315	110	16,164	3306	19,929	4500	9,910	110	24,403
3306	82,715	143	15,382	993	16,291	3389	9,702	53	23,588
993	75,921	3389	12,499	995	16,225	8080	7,992	143	23,312

Table 1: Number of exposed services by ports used

This section lists the top 10 exposed ports in each of the profiled US cities. The most popular exposed ports include 22 (SSH), 53 (Domain Name System [DNS]), 80 (HTTP), 110 (Post Office Protocol 3 [POP3]), 143 (Internet Message Access Protocol [IMAP]), 443 (HTTPS), 3306 (MySQL), 4567 (application port used by Verizon Fios® to access routers), and 7547 (CWMP used to access routers). Other interesting exposed ports include 21 (FTP), 23 (Telnet), 26 (unassigned), and 8080 (alternate HTTP).

Exposed NTP

NTP is one of the Internet’s oldest protocols designed to synchronize time between computer systems that communicate over unreliable variable-latency network paths. A recently published paper by Boston University researchers discusses methods of attacking NTP servers. Connections between computers and NTP servers are rarely encrypted, making it possible for hackers to perform man-in-the-middle (MitM) attacks that reset clocks to times that are months or even years in the past. Hackers can wreak havoc on the Internet with these NTP MitM attacks. An attack that prevents sensitive computers and servers from receiving regular time-synchronization updates can cause malfunctions on a massive scale. These attacks can be used to snoop on encrypted traffic or bypass important security measures such as DNS Security Extensions (DNSSEC) specifications, which are designed to prevent DNS record tampering. The most troubling scenario involves bypassing HTTPS encryption by forcing a computer to accept an expired transport layer security certificate.^{13, 14}




Figure 31: Number of exposed cyber assets using NTP

Exposed UPnP and/or SSDP

UPnP is a set of networking protocols that permits networked devices such as computers, printers, Internet gateways, WAPs, and mobile devices to seamlessly discover each other’s presence on the network and establish functional network services for data sharing, communication, and media playback.¹⁵ SSDP, meanwhile, is used to discover UPnP devices. It was first introduced in 1999 and is used by many routers and network devices. According to the NVD, there are 46 vulnerabilities that directly or indirectly affect UPnP while 14 vulnerabilities directly or indirectly affect SSDP. The Metasploit framework includes many UPnP and SSDP modules that can be used to exploit and compromise UPnP- or SSDP-enabled devices.




Figure 32: Number of exposed cyber assets using UPnP or SSDP

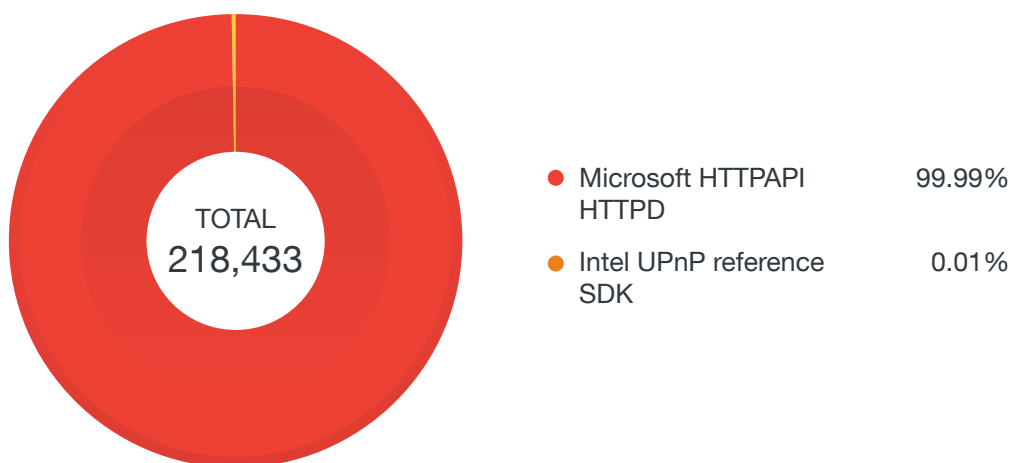


Figure 33: Distribution of exposed cyber assets using UPnP or SSDP by software

Exposed SNMP

SNMP is a popular protocol for network management. It is used to collect information and configure network devices such as servers, printers, hubs, switches, and routers.¹⁶ It is a convenient way for hackers to figure out the network topology, which they can later use for lateral movement within the target network. It can also be used to manage devices (e.g., to shut down a network interface), making it a dangerous tool in the hands of malicious hackers.¹⁷ Another big threat is hackers abusing devices configured to publicly respond to SNMP requests in order to amplify denial-of-service (DoS) attacks. Hackers use the Internet Protocol (IP) address of an individual or an organization they are targeting as the spoofed source of the SNMP request. They can then send bulk requests to devices configured to publicly respond to SNMP requests, which results in a flood of SNMP GetResponse data being sent from the devices to the victims.¹⁸




Figure 34: Number of exposed cyber assets using SNMP




Figure 35: Distribution of exposed cyber assets using SNMP by product

Exposed SSH

SSH is a cryptographic network protocol used to securely operate network services over an unsecured network.¹⁹ It is one of the protocols frequently targeted by hackers usually via brute-force attacks. In an SSH brute-force attack, an automated program tests combinations of usernames and passwords on a server to gain entry. This is effective against weak username/password combinations. To prevent SSH brute-force attacks, an administrator can restrict SSH access by IP address, change SSH to another port, use intrusion prevention tools to dynamically block access, rate-limit SSH sessions, or lock out an account after a defined number of log-in attempts. From the Shodan data, we saw that miscellaneous NAS devices, routers, and firewalls made up the bulk of exposed SSH-enabled devices.




Figure 36: Number of exposed cyber assets using SSH

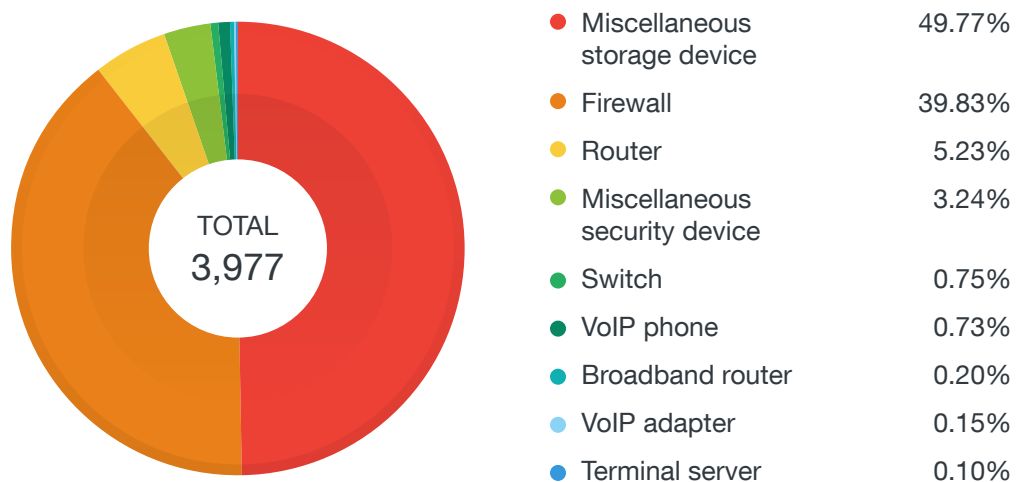


Figure 37: Distribution of exposed cyber assets using SSH by device type

Exposed RDP

RDP is a proprietary protocol developed by Microsoft, which provides users with a graphical interface to connect to another computer over a network connection. Users employ RDP client software for this purpose while a target computer must run RDP server software.²⁰ According to the NVD, 46 vulnerabilities directly or indirectly affect RDP. One of the popularly exploited RDP vulnerabilities is CVE-2012-0002. Proof-of-concept (PoC) code for CVE-2012-0002 was leaked online, leading to widespread exploitation. RDP has traditionally been abused to exfiltrate data as part of a targeted attack, steal information that can be sold in Deep Web marketplaces, and integrate hijacked systems into botnets. Recently, Crysis ransomware were found to be able to brute-force RDP as infection vector.²¹




Figure 38: Number of exposed cyber assets using RDP




Figure 39: Distribution of exposed cyber assets using RDP by software

Exposed Telnet

Telnet is an application layer protocol used on the Internet or a LAN to provide bidirectional interactive text-oriented communication using a virtual terminal connection.²² In a Telnet session, all data is sent and received in clear text; there is no end-to-end content encryption. This makes Telnet highly vulnerable to packet-sniffing attacks. Telnet was first introduced in the early 1970s and, over time, has been replaced by SSH. It is surprising to find so many Telnet-enabled devices in the Shodan data, chief among them being routers and network switches. According to the NVD, 210 vulnerabilities directly or indirectly affect Telnet. A quick Internet search shows numerous tools and instructional websites that discuss methods to exploit Telnet. Administrators are strongly encouraged to disable Telnet if they do not have a use for it.




Figure 40: Number of exposed cyber assets using Telnet




Figure 41: Distribution of exposed cyber assets using Telnet by device type

Exposed FTP

FTP is a standard network protocol used to transfer files between a client and a server over a computer network.²³ It is enabled by default on most Web servers, which makes it a lucrative target for exploitation by hackers. Once FTP is exploited and the server compromised, hackers can access all hosted files and upload new malicious files. Looking at the Shodan data, we found routers, WAP and NAS devices, printers, print servers, and webcams in the list of exposed FTP-enabled devices. We wonder how many users know that the FTP port is open on their WAPs, printers, and webcams? Several tools are available online for exploiting FTP, including a collection of exploits in the Metasploit framework. Device manufacturers have a responsibility to ensure that their products do not have FTP enabled by default in order to prevent device compromise by the exploitation of vulnerabilities.




Figure 42: Number of exposed cyber assets using FTP




Figure 43: Distribution of exposed cyber assets using FTP by device type

Exposed CWMP

The emergence of the Mirai botnet has increased the focus on router security. ISPs use TR-069 or CPE Wide Area Network (WAN) Management Protocol (CWMP) to remotely manage router-modems. Port 7547 has been assigned to this protocol and is left open to outside connections on ISPs' router-modems. Authentication happens via certificates or TR-069 messages encoded with Simple Object Access Protocol (SOAP). TR-069 messages can be used to reboot a device, reset it to factory defaults, and get or set configuration parameters.²⁴




Figure 44: Number of exposed cyber assets using CWMP

There are known vulnerabilities in TR-069 as well as implementation and configuration flaws in many ISPs' Auto Configuration Servers (ACSs) that communicate with router-modems using TR-069. ACSs are a single point of failure and can lead to ISP router-modem fleet takeover, causing customer service disruptions. The number of devices listening on port 7547 is very large (more than 16 million across all municipalities in the US and more than 1.38 million in the top 10 US cities according to the February 2016 Shodan US scan data), but not all of them run vulnerable implementations of TR-069. Some only accept commands from a specific server. It is very difficult to ascertain which router-modems are vulnerable and which are not. It is also reasonable to expect exploit code for TR-069 vulnerabilities will be added to the Mirai and other IoT botnets in the future.^{25, 26, 27}

Safeguarding Against Internet Exposure

Defensive Strategies for Businesses

Exposed cyber assets do not translate to compromise; rather, this means some device, system, or network is poorly configured. On the flip side, by virtue of being exposed on the Internet, this device or system is vulnerable to compromise. Cyber-attack and data breach prevention strategies should be considered an integral part of daily business operations. The key principle of defense is to assume compromise and take countermeasures:

- Quickly identify and respond to ongoing security breaches
- Contain the security breach and stop the loss of sensitive data
- Preemptively prevent attacks by securing all exploitable avenues
- Apply lessons learned to further strengthen defenses and prevent repeat incidents

A strong security checklist includes:

- Securing the network infrastructure by:
 - Segmenting a network according to function, department, geographic location, level of security, or any other logical separation (taking contractors, third-party vendors, and others into account)
 - Implementing log analysis for threat detection and remediation, and building threat intelligence; the data can be fed into Security Information and Event Management (SIEM) software and help the response team understand ongoing attacks
 - Properly configured user access profiles, workstations, and servers, including Internet-connected devices using the least-privilege model

- Protecting sensitive data via:
 - Data classification by determining the sensitivity of data sets and establishing different access and processing guidelines for each category
 - Establishing endpoint-to-cloud protection through identity-based and cloud encryption
 - Building a data protection infrastructure with multitiered access where sensitive tiers are in a disconnected network, others require multifactor authentication, and others can remain on regular file servers
- Building an incident response team consisting of technical, human resources, legal, public relations personnel, and executive management
- Building internal and collecting external threat intelligence, acted upon by knowledgeable human analysts who can determine through identifying patterns in attacker's tools, tactics, and procedures (TTPs), if an attack is ongoing inside the network

Ultimately, no defense is impregnable against determined adversaries. Having effective alert, containment, and mitigation processes is critical. Companies should further look into fulfilling the Critical Security Controls (CSC)²⁸ best practice guidelines published by the Center for Internet Security. The CSC goes through periodic updates to address new risks posed by an evolving threat landscape.

Securing Connected Homes

Today's society is adopting connected technologies at a faster rate than we are able to secure them. Every home is unique and hosts a wide variety of connected devices that serve different functions. Unfortunately, there is no "one size fits all" cybersecurity solution for connected devices. Compared with a business environment, a connected home is unstructured, dynamic, and tends to be function oriented. A vast majority of people are either unaware or unconcerned about the potential security risks that their exposed connected devices pose. The IoT ecosystem is multilayered and risk factors tied to successful compromises increase with each additional layer.




Figure 45: Risk factors increase with the addition of each layer in the IoT ecosystem from “Securing Your Smart Homes”²⁹

To better understand how to secure connected devices at home, we did an inventory of connected devices in one of the paper authors’ homes. The list includes laptop computers, a Wi-Fi-enabled printer, smartphones, a smartwatch, tablets, a wearable health monitor, an Internet-connected television (TV), a router, an IPTV receiver, a VoIP phone, a Nest thermostat, a Wi-Fi-enabled telescope, gaming consoles, NAS, a Wi-Fi bathroom scale, an Apple TV®, and Amazon Kindles. While this list may look long, it is not unusual for the average home to have that many if not more connected devices.

After brainstorming how to secure the list of inventoried connected devices, we came up with a set of general guidelines and best practices that home users should follow. Many of the recommendations are common sense and cybersecurity experts will repeatedly recommend them. When discussing how to secure connected devices at home, we also need to be mindful of three core IoT principles—always online, always available, and easy to use. We also need to remember that the average household does not have a resident IT guru who can secure everything connected, so enabling security features should be made as simple as possible. Our recommendations are as follows:

- Enable password protection on your devices. This is an easy option to enable on most connected devices that support passwords. It should be mandatory for smartphones, tablets, laptops, webcams, and so on.
- Replace default with strong passwords. Users routinely do not change the factory default passwords on their devices and these can be easily discovered using any Internet search engine. The other usual suspect is weak passwords that can be defeated using brute-force or dictionary attacks.

- Change default settings. Many devices have all their supported services enabled by default, many of which are not essential for daily operations (e.g., Telnet on webcams). If possible, users should disable nonessential services. The only caveat is that advanced technical knowledge may be required to decide which services to disable and how to correctly do that. We do not expect the average user to be knowledgeable about this so it is up to device manufacturers to make sure their devices are secure out of the box.
- Do not jailbreak devices. This can disable built-in security features, making it easier for hackers to compromise them. Jailbreaking is popular especially with smartphones, as this allows users with phones locked to a particular service provider to make them work for all service providers or in different countries.
- Do not install apps from unverified third-party marketplaces. Only use verified app marketplaces such as Apple's App Store, Google Play, Amazon Appstore, and others. This is especially a big security risk for jailbroken iOS and Android™ devices. Apps installed from unverified third-party marketplaces can have backdoors built into them that criminals can use to steal personal information or, worse, take control of them. Verified app marketplaces are not immune to hosting malicious apps but the probability of that happening is small.
- Update firmware. This will fix known security vulnerabilities. On the flip side, there are many caveats with firmware updates—some device firmware are not easy to update; the latest firmware may be unstable and introduces new bugs or issues; there are too many devices to update; it is difficult to track firmware updates; why should users update the firmware when the device is functioning properly; and updating the firmware may not even be possible.
- Enable encryption for both disk and communication. Enable disk encryption for smartphones, tablets, laptops, and other devices to secure the data on them even if they are stolen. Encryption is not a bulletproof solution but will secure the data on the disk against theft from the most skilled and resourceful hackers. Enabling HTTPS instead of HTTP for communication secures devices against MitM and packet-sniffing attacks.
- Some router-specific best practices include enabling the firewall, using faster but shorter-range 5GHz Wi-Fi signals to limit access-point-hacking attempts, disabling WPS and enabling the Wi-Fi Protected Access-2 (WPA2) security protocol, and using a strong password for Wi-Fi access.
- Other router security suggestions that unfortunately may limit device usage and functionality include configuring the router to limit device network access to set hours during the day or night, disabling UPnP though this will limit the operations of connected devices such as Wi-Fi-enabled printers, and allowing only a hardcoded list of device media access control (MAC) addresses to access a network (the MAC address list will need to be constantly updated).

- In extreme cases, disconnect the device from the network if Internet access is optional for it to function properly. But this practice goes against one of the core IoT principles—always online. For devices such as the Wi-Fi bathroom scale, Internet access is not required to measure body weight but is required for the bathroom scale to send the measured weight to an online portal that tracks daily changes in weight and provides fitness suggestions.

Connected devices are an integral part of our daily lives. Ideally, device security should not affect availability and should be transparent to the user. As previously stated, there is no “one size fits all” cybersecurity solution for connected devices. In addition to the listed best practices and general guidelines, users must be able to rely on device manufacturers to enable strong security out of the box. Ultimately, we may need to rely on security by obscurity—our connected devices hide among billions of other connected devices online and avoid getting compromised by hackers.

Conclusion

Our analysis of Shodan data for the top 10 US cities reveals that an incredible number of devices are publicly visible over the Internet. For instance, Los Angeles, Houston, Chicago, and Dallas each has more than 2 million exposed cyber assets that make them vulnerable to exploitation and compromise.

Majority of the exposed devices we found were, as expected, connected to the Internet via router-modems. However, we also saw devices connected via VPNs and VLANs, which should be private and not be responding to the Shodan crawler. Devices running Linux dominated in terms of volume, which could be because IoT devices very commonly run embedded Linux.

The top exposed devices discovered by Shodan are firewalls, webcams, WAPs, printers, routers, and phones. Houston has the greatest number of exposed webcams and routers while Los Angeles has the greatest number of exposed printers. San Jose has the greatest number of exposed phones. Chicago has the greatest number of exposed media devices, which includes DVRs like TiVo.

The top exposed product discovered by Shodan is Apache HTTPD. It is important to note that aside from a handful of exposed servers, a vast majority of the servers scanned by Shodan were already patched against Heartbleed and Freak. Alarming, we also found exposed medical databases, which could be a cause for concern for healthcare organizations, as hackers could compromise them and steal sensitive patient information.

Altogether, the sheer volume of exposed cyber assets we discovered in Shodan suggests that even if it is relatively easy to secure a connected device, many device owners still fail to do so. Companies should employ defensive strategies to make sure threat actors will not infiltrate their networks via their Internet-connected devices. Likewise, home users should secure newly purchased and currently installed connected devices to avoid the undesirable consequences of exposed devices.

Appendix

What Is Shodan?

Scanning the Internet is important because security flaws can be quickly identified or discovered and fixed before they are exploited. But scanning the Internet is difficult and time-consuming because of the massive IP address space that needs to be scanned—IPv4 supports a maximum of 2^{32} unique addresses and IPv6 supports a maximum of 2^{128} unique addresses. In addition to this massive address space, carrier and traditional Network Address Translation (NAT) hides millions of connected nodes; IPv6 gateways also support NAT64, which connects IPv6 to IPv4. Other challenges with scanning the Internet include administrators seeing network scans as attacks, some IP ranges are blocked by different countries, legal complaints, dynamic IP addresses, ICS device operations can be affected by active network scanning, powerful hardware required for processing and storage, exclusion lists, agreements with ISPs so they do not block Internet access, and others. For this research, we bypassed all of these issues and hurdles and simply used a public data source—Shodan.

What is Shodan? Shodan is a search engine for Internet-connected devices. The basic unit of data that Shodan gathers is the banner, which contains textual information that describes a service on a device. For Web servers, this would be the headers that are returned; for Telnet, it would be the log-in screen. The banner content greatly varies, depending on service type. In addition to banners, Shodan also grabs metadata about a device such as geographic location, hostname, OS, and more.³⁰ Shodan uses a GeolIP database to map the scanned IP addresses to physical locations.

The Shodan crawler works as follows—first, it generates a random IPv4 address; next, it generates a random port to test from a list of ports that it understands; and finally, it scans the generated IPv4 address on the generated port and grabs any returned banners. This means the Shodan crawlers do not scan incremental network ranges. Completely random crawling is performed to ensure uniform coverage of the Internet and prevent bias in the data at any given time. Scan data is collected from around the world to prevent geographic bias. Shodan crawlers are distributed around the world to ensure that any sort of countrywide blocking will not affect the data gathering.

Shodan provides an easy one-stop solution to conduct open source intelligence (OSINT) gathering for different geographic locations, organizations, devices, services, and others. Software and firmware information collected by Shodan can potentially help identify unpatched vulnerabilities in exposed cyber assets. An adversary can use Shodan to perform detailed surveillance and gather intelligence about a target, which is why it has been called the “world’s most dangerous search engine.”³¹ But, truth be told, insufficiently secured devices are coming online in droves and a vast majority of people are either unaware or unconcerned about the potential security risks. Adversaries do not really need Shodan scan data to

find weaknesses in their targets; they often do their own scanning using open source tools such as nmap and MASCCAN. There are also other publicly available data sources and services similar to Shodan for doing surveillance and intelligence gathering. Shodan was the first search engine to bring awareness to the large variety and massive volume of everyday exposed cyber assets all around us.

Shodan Data Analysis

For this research, we partnered with Shodan who provided us with access to raw scan data in JavaScript Object Notation (JSON) format. Recently, we did a similar but smaller study of exposed cyber assets in Japan,³² which was very well received and that encouraged us to do a more comprehensive study for a different country. In this paper, we examined the US scan data for the month of February 2016 because the US has the largest number of exposed cyber assets among all countries observed by Shodan. *The Shodan crawler roughly takes three weeks to cycle through the entire IPv4 address space; hence a month's worth of Shodan scan data provides a fairly accurate picture of the different online devices and systems in the US.* The February 2016 US scan data contains a total of 178,032,637 records generated by scanning 45,597,847 unique IPv4 and 256,516 unique IPv6 addresses. The raw scan data was indexed using Elasticsearch and queried using Kibana, which allowed us to search more than 550 fields instead of only 40 some fields using Shodan's Web interface. Observations and assumptions include:

- We did not study month-to-month changes in the Shodan US scan data because these tend to be gradual. To observe marked differences, we would need to study changes in the scan data over many months, if not several years, which is outside the scope of this research paper. Realistically, only significant regional or national events will dramatically impact the number of Internet-exposed devices and systems; hence, we assume a month's worth of scan data will give us an accurate snapshot of what devices and systems are exposed online in the US. Profiling exposed cyber assets in different countries as well as tracking long-term trends in Shodan data will make interesting future research.
- IP addresses appear and disappear month to month from the Shodan scan data. In some cases, the devices and systems are offline and the IP address and port scan returns no results. A device or system absent in Shodan does not mean it is not exposed online. On the flip side, Shodan may rescan the same IP address multiple times in the same month (e.g., we found an IP address with 58,143 scan records).
- We found that the volume of exposed cyber assets in large cities can be disproportionate to their population size. For example, the February 2016 Shodan US scan data showed 3,900,208 exposed cyber assets in Houston, Texas compared to 1,031,325 exposed cyber assets in New York City, New York. New York City has a far bigger population than Houston yet it has 3.78 times fewer exposed cyber assets compared with Houston. We think these discrepancies exist because:
 - Shodan's GeoIP database mapping may be returning results for the city only and not the greater metropolitan area to which it belongs.

- ISPs that serve the city are dropping Shodan crawler queries.
- Presence of data centers
- Shodan scans both the IPv4 and IPv6 address spaces. We restricted our research to the IPv4 address space only. IPv6 address scanning accounted for only 0.78% of the total data so examining the IPv4 address space only gives us a fairly accurate snapshot of what devices and systems are exposed online in the US. This will, of course, change over time as the IPv6 address space is better utilized by future connected devices.
- Explosion in the usage of the Internet means the IPv4 address space is fast getting depleted. The IPv4 address space supports a maximum of 2^{32} addresses. IPv6, with its maximum 2^{128} addresses, will more than solve the address space shortage problem but this will still take several years to be fully implemented or adopted. And even then, IPv4 will continue to be used. NAT is an essential tool in conserving global IPv4 address space allocations. NAT allows a single device such as a router to act as an agent between the Internet and a local (or “private”) network. This means that only a single unique IP address is required to represent an entire group of computers and devices.³³ This translates to finding multiple devices and systems visible from the same IP address in the Shodan scan data, most likely sitting behind a router or a firewall.

Recent Notable Cyber Intrusions

As the Internet and the real world increasingly intersect, hackers are infiltrating critical systems and infrastructure. A recent Australian Broadcasting Corporation (ABC) news report lists some of the well-publicized cases of cyber intrusions³⁴ (note that not all incidents listed below are results of cyber asset exposure):

- In 2013, Newsat (one of Australia’s biggest satellite companies), which builds communication satellites for the Australian Defense Force and mining companies, publicly disclosed that its network was infiltrated and compromised by foreign hackers.
- Stuxnet was used to target Iran’s nuclear fuel-enrichment facilities. The attacks were speculated to have originated from nation-states that wanted to deter Iran’s nuclear ambitions. While this was not the first cyber attack against ICS devices, it was the first to infect a programmable logic controller (PLC).
- The first publicly acknowledged successful cyber attack to knock out a power grid happened in the Ukraine in December 2015. Thirty substations were disconnected from the grid, leaving 225,000 customer homes freezing in the Ukrainian winter chill. The BlackEnergy group is believed to have been responsible for this attack.

- In July 2015, security researchers Charlie Miller and Chris Valasek demonstrated that they could remotely hack into a 2014 Jeep Cherokee, allowing them to control its transmission system and brakes. They exploited a vulnerability in the multimedia system's Wi-Fi to gain access.
- In 2014, security researcher, Billy Rios, found he could remotely hack into hospital drug-infusion pumps that administer morphine and antibiotics, and change drug dosage levels.
- In 2014, the German government disclosed that hackers attacked an unnamed steel mill in the country and destroyed one of its blast furnaces. Hackers accessed software used to control the plant's operations, which allowed them to stop the blast furnace from shutting down and, in the process, destroyed it.
- In 2013, Billy Rios and Terry McCorkle hacked into the building management system (used to control power management systems, CCTV cameras, security and fire alarms, electronic locks, etc.) of Google's Sydney office. They discovered the exposed building management system using the Shodan search engine.
- Hackers almost gained control of the floodgates at Bowman Avenue Dam near New York City in 2013. Details of the incident remain classified but hackers from a Middle Eastern country are believed to have been behind the attacks.
- French TV station, TV5Monde, fell victim to a sophisticated cyber attack that brought down 12 channels for almost a whole day in April 2015.
- The Australian Bureau of Meteorology suffered a significant cyber intrusion that was first discovered in 2015. The target may have been the Australian Geospatial Intelligence Organization, which provides satellite imagery for sensitive defense operations, and the Royal Australian Air Force's Jindalee Operational Radar Network (JORN), which is designed to detect planes and maritime vessels within a 3,000-kilometer radius of Australia's northern and western shorelines.
- In 2010, Barnaby Jack demonstrated an automated teller machine (ATM) compromise at the Black Hat USA conference. One of the vulnerabilities Jack demonstrated was in the remote-monitoring feature, which is turned on by default in some ATM models.
- In 2014, researchers from the University of Michigan under the supervision of the government road agency demonstrated that they could remotely control a system of 100 intersection traffic lights in an unnamed city in Michigan. The traffic lights use wireless radio to communicate with the central network. The researchers used this radio system to send commands to the traffic lights and could change the lights at will.
- Security researcher, Chris Roberts, is the subject of an ongoing Federal Bureau of Investigation (FBI) case after claiming to have hacked a plane midflight via its entertainment system. He claims to have made the passenger jet fly sideways.

- A recent cyber-attack story that made the headlines was the DDoS attack against the site of security blogger, Brian Krebs—KrebsOnSecurity.com. This was one of the biggest DDoS attacks ever, done in retaliation for Krebs’s exposure of the group who carried out such attacks as a paid service. At its peak, the attack aimed 620GB of data per second at KrebsOnSecurity.com. Security firm, Akamai, said the attack generated such a huge volume of data by exploiting weak or default passwords in widely used Internet-connected cameras, routers, and DVRs. Once in control of these “smart” devices, the hackers used them to swamp the site with data requests. “These new Internet-accessible devices can bring great benefits but they are also increasingly easy and lucrative targets for cybercriminals.”³⁵ On 21 October 2016, a massive DDoS attack was launched to knock down the Dyn DNS server infrastructure, which subsequently caused outages for all sorts of online services and popular websites because one of the major DNS providers became unavailable.³⁶ The Mirai IoT botnet is a malware infrastructure suspected of being used in both DDoS attacks against Dyn³⁷ and KrebsOnSecurity.com.³⁸ People ask why anyone would hack their IoT devices. These attacks are good examples of why IoT devices are compromised and how hackers can abuse them.

Is Your Light Bulb Really Hacking You?

The recent cyber-attack story against KrebsOnSecurity.com clearly demonstrated how compromised IoT devices can be abused to generate traffic in a DDoS attack. With millions of IoT devices connected to the Internet, just how hard is it to get compromised? Max Goncharov and Philippe Lin of the Trend Micro FTR Team set up IoT honeypots to figure the answer to that question. Their research was presented at HITCON 2015 in Taipei. We present a summary of their findings in order to answer the question, “Is your light bulb really hacking you?”³⁹

IoT is omnipresent and will continue to grow especially in the IPv6 era. Max and Philippe created two physical IoT honeypots—the Taipei honeypot ran from 23 March to 23 July 2015 while the Munich honeypot ran from 22 April to 22 June 2015, a combined total of approximately six months of operations.

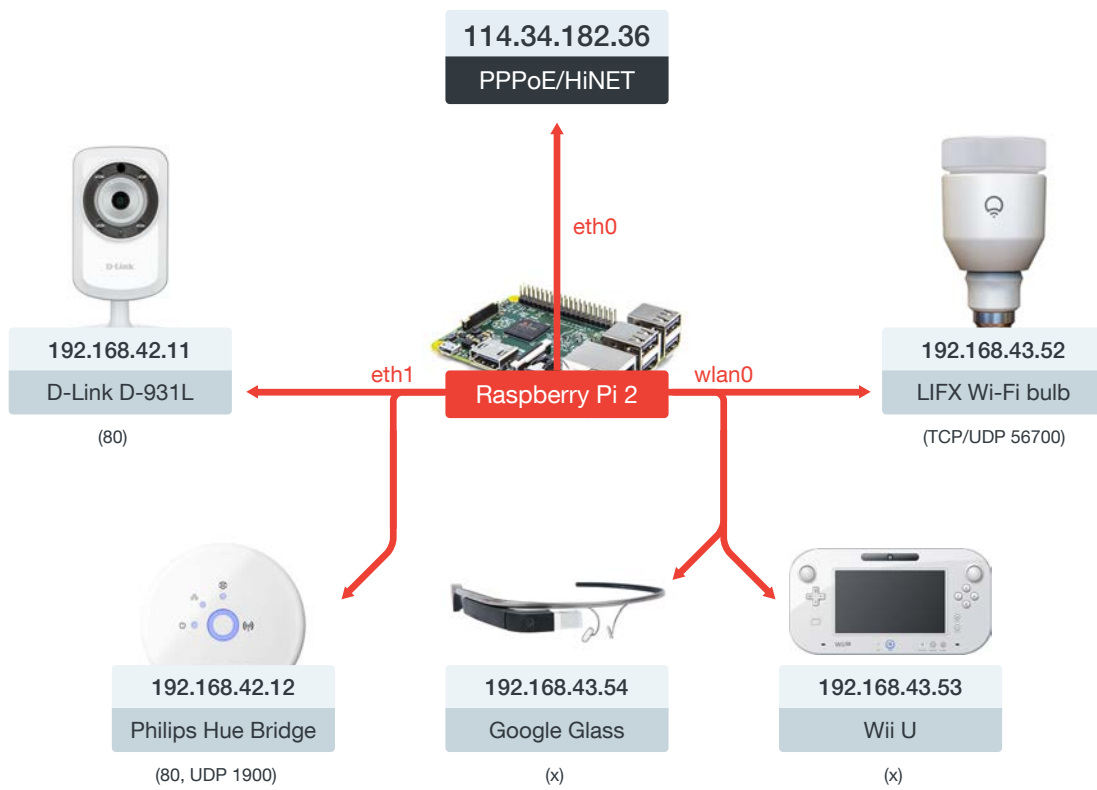


Figure 46: Taipei honeypot block diagram

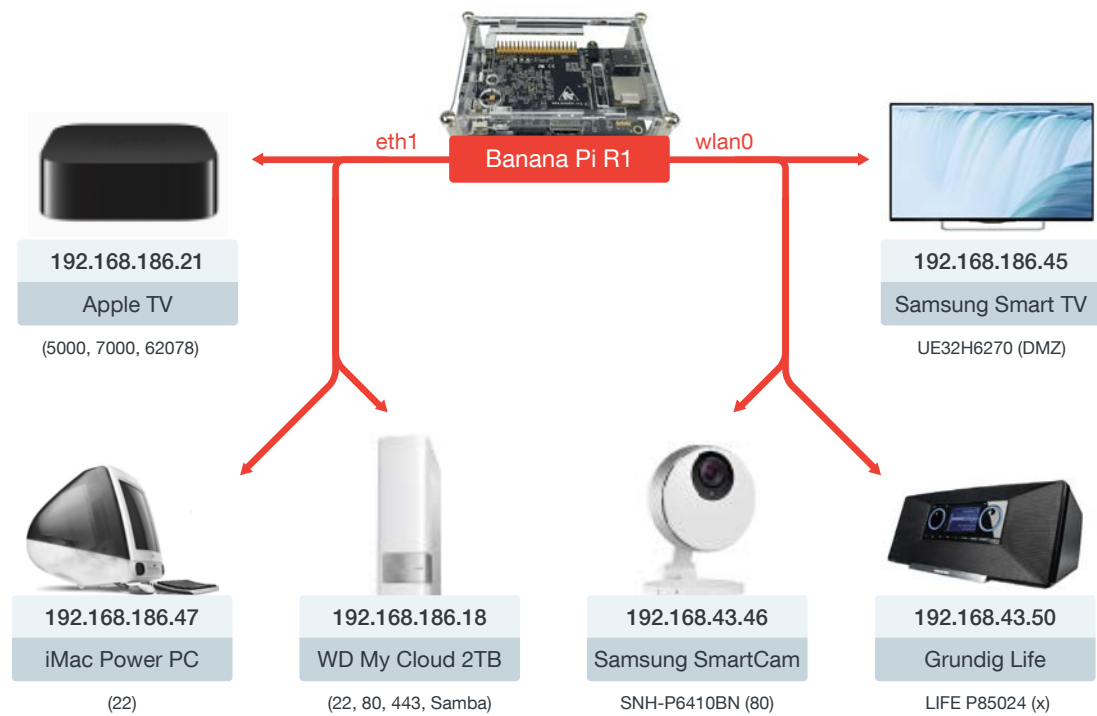


Figure 47: Munich honeypot block diagram

Why use real devices in the honeypot? Shodan can identify honeypots (<https://honeyscore.shodan.io/>); hence, it is safe to assume that hackers can also determine if they are interacting with honeypots or real devices. If the honeypot uses real devices, then we get the correct responses and actions every time. The hackers cannot figure out that in reality they are interacting with a honeypot. The only major caveat with this approach is scalability; adding new devices means purchasing them, which can be expensive. A fake identity and social profile was created for “the person” who owns these IoT devices—fake Facebook, Skype, and documents were uploaded to WDCloud. URLs or credentials were randomly pushed to Shodan and Pastebin. Every effort was made to get hackers to compromise the honeypots.

Analyzing honeypot network access logs, we discovered that automated scanning of connected IoT devices generated the bulk of the network traffic observed. No serious IoT hackers (i.e., people attempting to exploit known vulnerabilities in devices in order to compromise them) were observed. Mostly, we observed queries for popular webcams. In the advent of the recent Mirai attacks, the queries for popular webcams now paints a different and more alarming picture. Mirai looks for popular webcams, among other things, then tries commonly used username/password combinations on them in order to gain access. Once device access is gained, Mirai installs itself and can accept commands to participate in DDoS attacks.

In conclusion, none of the IoT devices were hacked or compromised. Only one person “looked” through the D-Link webcam in the four months the Taipei honeypot was operational. The honeypots may have potentially seen more attacks if they were run in multiple countries, popular device models that automated scanners look for were used, they were run for a longer duration, and the devices were connected to multiple IP addresses. At the end of the day, with billions of IoT devices online, the probability of successfully getting compromised looks small.

References

1. Numaan Huq, Stephen Hilt, and Natasha Hellberg. (15 February 2017). *Trend Micro Security News*. “US Cities Exposed.” Last accessed on 15 February 2017, <http://www.trendmicro.com/vinfo/us/security/news/internet-of-things/us-cities-exposed>.
2. Éireann Leverett and Marie Moe. (March 2016). RSA Conference 2016. “From Ukraine to Pacemakers! The Real-World Consequences of Logical Attacks.” Last accessed on 21 September 2016, https://www.rsaconference.com/writable/presentations/file_upload/hta-f03-from_ukraine_to_pacemakers_the_real-world_consequences_of_logical_attacks.pdf.
3. Infoplease. (2016). *Infoplease*. “Top 50 Cities in the US by Population and Rank.” Last accessed on 30 September 2016, <http://www.infoplease.com/ipa/a0763098.html>.
4. Yoshizawa Torushi. (28 June 2016). *ScanNetSecurity*. “外部にさらされているインフラは東京が突出～SHODANで見る5都府県（トレンドマイクロ）.” Last accessed on 18 October 2016, <http://scan.netsecurity.ne.jp/article/2016/06/28/38644.html>.
5. National Institute of Standards and Technology (NIST). (2017). *NVD*. “Search CVE and CCE Vulnerability Database.” Last accessed on 1 October 2016, <https://web.nvd.nist.gov/view/vuln/search>.
6. Steve Ragan. (3 January 2017). *CSO Online*. “Exposed MongoDB Installs Being Erased, Held for Ransom.” Last accessed on 9 January 2017, <http://www.csoonline.com/article/3154190/security/exposed-mongodb-installs-being-erased-held-for-ransom.html>.
7. Wikimedia Foundation Inc. (27 May 2016). *Wikipedia*. “PACS.” Last accessed on 3 October 2016, https://en.wikipedia.org/wiki/Picture_archiving_and_communication_system.
8. Health IT. (16 March 2013). *HealthIT.gov*. “What Is an EHR?” Last accessed on 3 October 2016, <https://www.healthit.gov/providers-professionals/faqs/what-electronic-health-record-ehr>.
9. Wikimedia Foundation Inc. (3 October 2016). *Wikipedia*. “EHR.” Last accessed on 3 October 2016, https://en.wikipedia.org/wiki/Electronic_health_record.
10. James Scott and Drew Spaniel. (September 2016). *Institute for Critical Infrastructure Technology*. “Your Life, Repackaged and Resold: The Deep Web Exploitation of Health Sector Breach Victims.” Last accessed on 3 October 2016, <http://icitech.org/wp-content/uploads/2016/09/ICIT-Brief-Deep-Web-Exploitation-of-Health-Sector-Breach-Victims2.pdf>.
11. Marianne Kolbasuk McGee. (27 September 2016). *Data Breach Today*. “Research Reveals Why Hacked Patient Records Are So Valuable.” Last accessed on 3 October 2016, <http://www.databreachtoday.com/interviews/research-reveals-hacked-patient-records-are-so-valuable-i-3341>.
12. HIPAA Journal. (24 June 2015). *HIPAA Journal*. “What Are the Penalties for HIPAA Violations?” Last accessed on 4 January, 2017, <http://www.hipaajournal.com/what-are-the-penalties-for-hipaa-violations-7096/>.
13. Dan Goodin. (21 October 2015). *ArsTechnica*. “New Attacks on NTP Can Defeat HTTPS and Create Chaos.” Last accessed on 2 October 2016, <http://arstechnica.com/security/2015/10/new-attacks-on-network-time-protocol-can-defeat-https-and-create-chaos/>.
14. Aanchal Malhotra, Isaac E. Cohen, Erik Brakke, and Sharon Goldberg. (October 2015). *Boston University*. “Attacking the NTP.” Last accessed on 2 October 2016, <http://www.cs.bu.edu/~goldbe/papers/NTPattack.pdf>.
15. Wikimedia Foundation Inc. (31 August 2016). *Wikipedia*. “UPnP.” Last accessed on 2 October 2016, https://en.wikipedia.org/wiki/Universal_Plug_and_Play.
16. Microsoft. (28 March 2003). *Microsoft TechNet*. “What Is SNMP?” Last accessed on 2 October 2016, <https://technet.microsoft.com/en-us/library/cc776379%28v=ws.10%29.aspx>.

17. John McCormick. (11 April 2001). *TechRepublic*. "Lock IT Down: Don't Allow SNMP to Compromise Network Security." Last accessed on 2 October 2016, <http://www.techrepublic.com/article/lock-it-down-dont-allow-snmp-to-compromise-network-security/>.
18. Kelly Jackson Higgins. (22 May 2014). *Dark Reading*. "SNMP DDoS Attacks Spike." Last accessed on 2 October 2016, <http://www.darkreading.com/attacks-breaches/snmp-ddos-attacks-spike/d/d-id/1269149>.
19. Wikimedia Foundation Inc. (29 September 2016). *Wikipedia*. "SSH." Last accessed on 2 October 2016, https://en.wikipedia.org/wiki/Secure_Shell.
20. Wikimedia Foundation Inc. (8 September 2016). *Wikipedia*. "RDP." Last accessed on 2 October 2016, https://en.wikipedia.org/wiki/Remote_Desktop_Protocol.
21. Jon Oliver. (19 September 2016). *Trend Micro Security Intelligence Blog*. "A Show of (Brute) Force: Crysis Ransomware Found Targeting Australian and New Zealand Businesses." Last accessed on 2 October 2016, <http://blog.trendmicro.com/trendlabs-security-intelligence/crysis-targeting-businesses-in-australia-new-zealand-via-brute-forced-rdps/>.
22. Wikimedia Foundation Inc. (19 September 2016). *Wikipedia*. "Telnet." Last accessed on 2 October 2016, <https://en.wikipedia.org/wiki/Telnet>.
23. Wikimedia Foundation Inc. (19 September 2016). *Wikipedia*. "FTP." Last accessed on 2 October 2016, https://en.wikipedia.org/wiki/File_Transfer_Protocol.
24. Shahar Tal and Lior Oppenheim. (April, 2015). *RSAC 2015*. "The Internet of TR-069 Things: One Exploit to Rule Them All." Last accessed on 18 December 2016, https://www.rsaconference.com/writable/presentations/file_upload/hta-r04-the-internet-of-tr-069-things-one-exploit-to-rule-them-all_final_copy1.pdf.
25. Johannes B. Ulrich (29 November 2016). Sans ICS Infosec Forums. "TR-069 NewNTPServer Exploits: What We Know So Far" Last accessed on 18 December 2016, <https://isc.sans.edu/forums/diary/TR069+NewNTPServer+Exploits+What+we+know+so+far/21763/>.
26. Dan Goodin. (28 November 2016). *Ars Technica*. "Newly Discovered Router Flaw Being Hammered by in-the-Wild Attacks." Last accessed on 18 December 2016, <http://arstechnica.com/security/2016/11/notorious-iot-botnets-weaponize-new-flaw-found-in-millions-of-home-routers/>.
27. John Matherly. (August 2016). *Leanpub*. "Complete Guide to Shodan." <https://leanpub.com/shodan>.
28. CIS. (2016). *CIS*. "CIS Controls for Effective Cyberdefense." Last accessed on 7 January 2017, <https://www.cisecurity.org/critical-controls/>.
29. TrendLabs. (3 November 2016). *Trend Micro Security News*. "Securing Smart Homes." Last accessed on 5 January 2017, <http://www.trendmicro.com/vinfo/us/security/news/internet-of-things/securing-smart-homes>.
30. Paul Didier, Fernando Macias, James Harstad, Rick Antholine, Scott A. Johnston, Sabina Piyevesky, Mark Schillace, Gregory Wilcox, Dan Zaniewski, and Steve Zuponcic. (9 September 2011). *Cisco and Rockwell Automation*. "Converged Plantwide Ethernet (CPwE) Design and Implementation Guide. Last accessed on 3 May 2016, http://www.cisco.com/c/en/us/td/docs/solutions/Verticals/CPwE/CPwE_DIG/CPwE_chapter2.html.
31. Sam Clements. (26 April 2013). *Vice*. "Is Shodan Really the World's Most Dangerous Search Engine?" Last accessed on 21 September 2016, http://www.vice.com/en_uk/read/shodan-exposes-the-dark-side-of-the-net.
32. Yoshizawa Torushi. (28 June 2016). *ScanNetSecurity*. "外部にさらされているインフラは東京が突出～SHODANで見る5都府県（トレンドマイクロ）." Last accessed on 18 October 2016, <http://scan.netsecurity.ne.jp/article/2016/06/28/38644.html>.
33. Jeff Tyson. (2 February 2001). *HowStuffWorks.com*. "How Network Address Translation Works." Last accessed on 16 September 2016, <http://computer.howstuffworks.com/nat.htm>.

34. ABC News. (29 August 2016). *ABC News*. "The Internet of Hacked Things." Last accessed on 24 September 2016, <http://mobile.abc.net.au/news/2015-10-07/four-corners-internet-of-hacked-things/7778954>.
35. BBC Technology. (22 September 2016). *BBC*. "Massive Web Attack Hits Security Blogger." Last accessed on 24 September 2016, <http://www.bbc.com/news/technology-37439513>.
36. BBC Technology. (22 October 2016). *BBC*. "Smart Home Devices Used as Weapons in Website Attack." Last accessed on 23 October 2016, <http://www.bbc.com/news/technology-37738823>.
37. Kyle York. (22 October 2016). *Dyn*. "Dyn Statement on 10/21/2016 DDoS Attack." Last accessed on 24 October 2016, <http://dyn.com/blog/dyn-statement-on-10212016-ddos-attack/>.
38. Brian Krebs. (16 October 2016). *KrebsonSecurity.com*. "Source Code for IoT Botnet 'Mirai' Released." Last accessed on 24 October 2016, <https://krebsonsecurity.com/2016/10/source-code-for-iot-botnet-mirai-released/>.
39. Max Goncharov and Philippe Lin. (August 2015). *HITCON 2015*. "Your Light Bulb Is Not Hacking You—Observation from a Honeypot Backed by Real Devices." Last accessed on 26 September 2016, <https://hitcon.org/2015/CMT/download/day2-f-r1.pdf>.

Created by:

TrendLabs

The Global Technical Support and R&D Center of TREND MICRO

TREND MICRO™

Trend Micro Incorporated, a global cloud security leader, creates a world safe for exchanging digital information with its Internet content security and threat management solutions for businesses and consumers. A pioneer in server security with over 20 years experience, we deliver top-ranked client, server, and cloud-based security that fits our customers' and partners' needs; stops new threats faster; and protects data in physical, virtualized, and cloud environments. Powered by the Trend Micro™ Smart Protection Network™ infrastructure, our industry-leading cloud-computing security technology, products and services stop threats where they emerge, on the Internet, and are supported by 1,000+ threat intelligence experts around the globe. For additional information, visit www.trendmicro.com.



Securing Your Journey
to the Cloud

www.trendmicro.com