

Digging for Dark IPMI Devices: Advancing BMC Detection and Evaluating Operational Security

Oliver Gasser, Felix Emmert, Georg Carle
Technische Universität München
Chair of Network Architectures and Services
Email: {gasser,emmertf,carle}@net.in.tum.de

Abstract—IPMI is the industry standard for managing devices remotely independent of their operating status. Since there are known vulnerabilities in the protocol, IPMI devices should not be directly reachable on the Internet. Previous studies suggest, however, that this best practice is not always implemented. In this paper we present a new unintrusive technique to find dark IPMI devices through active measurements. These dark devices do not respond to conventional IPMI connection setup requests. Using our technique, we find 21 % more devices than previously known techniques. This adds a significant number of IPMI devices which could be exploited by an attacker using a Man-in-the-Middle attack. We further reveal that IPMI devices are heavily clustered in certain subnets and Autonomous Systems. Moreover, the SSL security of IPMI devices’ web-interface is well below the current state of the art, leaving them vulnerable to attacks. Overall our findings draw a dire picture of the current state of the IPMI deployment in the Internet.

I. INTRODUCTION

Out-of-band network management is the process of managing devices and systems over an auxiliary communication channel, independent of their operating state. Out-of-band management enables administrators to remotely manage servers, routers, switches, and other devices. This management capability is especially important when managing hundreds or thousands of these devices, such as in data centers or colocation centers.

The de facto industry standard protocol for Out-of-band management is IPMI (Intelligent Platform Management Interface). The IPMI protocol also specifies access to IPMI devices over the network via IPMI-over-IP.

The IPMI-over-IP protocol, however, has some inherent weaknesses. Attackers can exploit insufficient authentication checks and other vulnerabilities to compromise the host system as detailed e.g., by HD Moore [1]. These weaknesses allow an attacker to gain access to a powerful interface over the network. This introduces new attack vectors independent of the host system’s security. Once gained access to an IPMI device, an attacker can e.g., power off the device (essentially a Denial-of-Service attack), rebooting into a custom operating system, or installing a rootkit to eavesdrop on the communication. In short, an attacker has full control over the system and can potentially compromise the IPMI device itself.

To better assess these risks it is important to understand the IPMI deployment in the Internet. Therefore we conduct large-

scale scans to find openly accessible IPMI devices and classify their security properties. We use an unintrusive measurement technique which does not attempt any authentication with the IPMI devices. We perform the scans using a modified version of ZMap, which we make available online. As a result of the scans we find significantly more IPMI devices than other current scanning efforts. Moreover, we discover that a large number of IPMI devices are not properly secured.

Outline. The remainder of this paper is structured as follows. In the following Section II we present related work in the area of discovering IPMI devices and assessing their security. Section III provides information about the IPMI protocol and the scanning techniques used during the experiment. In Section IV we detail our scanning approach, first describing *dark IPMI devices* and which IP addresses we target during our measurements. Then we describe the software used for our network measurements and lay out our ethical considerations. In Section V we present the results of our scans and classify them with regards to security of out-of-band management. We conclude in Section VI and give pointers for future work.

II. RELATED WORK

In this section we present previous work related to our research. We start by pointing out other research surveying the security of IPMI devices. Then we detail works in the field of Internet-wide scans for security purposes.

In 2013, Bonkoski *et al.* [2] surveyed the security of IPMI implementations. They analyzed firmwares of IPMI devices from Supermicro and found exploits which can be used to bypass the authentication of the web frontend and gain access to the system. Furthermore, the authors estimated the number of potentially vulnerable IPMI devices by looking at SSL certificates from large-scale scans. They found that more than 40 000 potentially vulnerable IPMI devices and more than 100 000 IPMI devices in total exist in the wild. Similarly to Bonkoski *et al.*, we also use SSL certificates to identify potential IPMI devices as targets for our active measurements. In our Internet-wide measurement we found more than 220 000 IPMI devices.

In the same year Dan Farmer, [3], [4] performed Internet-wide scans for UDP/623 and found more than 230 000 IPMI devices. He analyzed the security of these devices and found

that many were vulnerable to authentication weaknesses and that passwords could be brute-forced. Although we found slightly less devices in 2015 than Dan Farmer in 2013, we found many dark IPMI devices, i.e., devices which do not respond to Farmer’s scanning technique.

Zhang *et al.* [5] in 2014 tried to correlate the maliciousness of networks (e.g., sending spam emails) with mismanagement metrics. One of their mismanagement metrics was the reachability of IPMI devices within Autonomous Systems (ASes). By matching regular expressions on the subject field in SSL certificates, they identified about 100 000 publicly reachable IPMI devices in different ASes. We use a similar technique to match found IPMI devices to vendors. They found a weak correlation between the reachability of IPMI devices in networks and the maliciousness of a network.

In 2014, Costin *et al.* [6] performed a large-scale analysis of embedded firmwares. They gathered firmware images using a web-crawler and a site where users could upload their firmware. Then the authors analyzed more than 32 000 of them and found 38 new vulnerabilities. Moreover, they were able to extract private SSL keys and crack hard-coded password hashes.

Similarly, Stefan Viehböck analyzed more than 4000 embedded firmware images in 2015 [7]. Unfortunately, most of the issues previously found still persist: the author was able to extract 580 unique private keys. These keys are used by 9% of all SSL hosts on the IPv4 Internet and 6% of SSH hosts.

Rapid7 performs monthly IPMI scans and publishes the raw response packets as part of *Project Sonar* [8]. Compared to their scanning efforts we use a different scanning technique (see Section III-D) which leads to more detected IPMI devices.

Large-scale network measurements have recently become a valuable tool to assess specific aspects of the Internet’s security. Holz *et al.* [9] conducted active and passive measurements in 2011 to assess the security of the SSL PKI. They identified multiple security issues in the SSL deployment such as incorrect certificate chains and invalid subject names in certificates.

In 2012, Heninger *et al.* [10] evaluated the cryptographic properties of SSL and SSH. They concluded that due to a lack of randomness, many keys were predictable.

In 2014, Gasser *et al.* [11] conducted multiple Internet-wide SSH scans. They were able to confirm many of Heninger *et al.*’s findings and additionally found duplicate yet cryptographically strong keys.

Similar to Heninger *et al.*, we analyze the SSL certificates of web interfaces to evaluate the security of IPMI devices. The certificates found on IPMI devices are not suitable to properly secure connections.

III. BACKGROUND

In this section we will give general information about Out-of-band network management. Then we provide insights into the protocols relevant for this research.

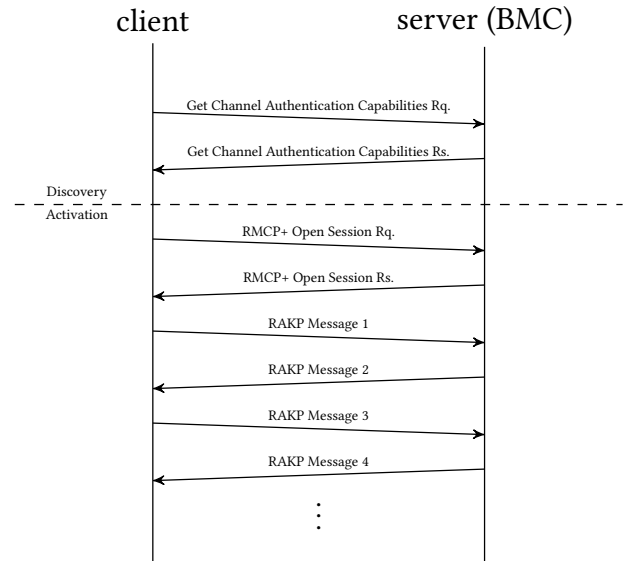


Fig. 1. IPMI-over-IP connection establishment in IPMI 2.0.

A. Out-of-Band Network Management

Out-of-band network management is a term describing different technologies enabling system administrators to remotely manage their network hardware (e.g., switches, routers, servers,...) independently of the system’s operating state. This goal is commonly achieved by independent sub-systems connected to the main system’s network hardware. These sub-systems run their own operating system on dedicated hardware. They are connected to various I/O ports of the main system. Out-of-band network management devices have their own network interface controllers (NICs) or at least access to one of the main system’s NICs via a “side-band” interface. Out-of-band management devices commonly provide some sort of management interfaces for administrators like web interfaces or access via SSH.

B. IPMI Basics

Intelligent Platform Management Interface (IPMI) is the de facto industry standard for Out-of-band network management devices used for server management. The IPMI specification [12] defines the architecture, different functionalities, and user interfaces of Out-of-band network management devices for servers. IPMI devices run on embedded microcontrollers called *Baseboard Management Controller (BMC)*. BMCs are commonly installed via daughter cards or directly integrated in the server’s mainboard. IPMI’s functionality may be extended by BMC manufacturers. A common example for such an extension are web interfaces on TCP ports 80 (HTTP) and 443 (HTTPS).

C. IPMI-over-IP

IPMI defines its own network protocol called *IPMI-over-IP* which uses UDP port 623 [12]. IPMI-over-IP allows for administrators to remotely login to their BMCs and perform

a set of actions like rebooting the server or configuring the BMC. Using IPMI-over-IP, it is possible to take full control over the connected server. If not needed, most devices offer the possibility to deactivate IPMI-over-IP.

IPMI-over-IP has been introduced in version 1.5 of the IPMI specification. It has been updated in the new version 2.0 of the IPMI specification. However, IPMI version 2.0 devices are still required to simultaneously support the old version 1.5 [12].

In the following we describe how an IPMI-over-IP connection is established and authenticated. IPMI-over-IP's connection establishment is divided in two phases, "Discovery" and "Activation" (see Figure 1 for IPMI version 2.0).

In the optional "Discovery" phase of the IPMI-over-IP protocol in version 2.0 of the IPMI specification, the client sends a *Get Channel Authentication Capabilities Request* packet to the BMC. The BMC answers using a *Get Channel Authentication Capabilities Response* packet. This response packet includes the IPMI version and authentication methods supported by the BMC. If IPMI-over-IP is deactivated, the BMC will not respond.

In version 1.5 of the IPMI specification, IPMI-over-IP also supports discovery using *RMCP Ping* packets. If probed, the BMC responds by sending an *RMCP Pong* packet. This response packet does not include much information other than whether or not IPMI is supported. However, small-scale tests on a *Dell iDRAC 7* show that the BMC replies to *RMCP Ping* packets even if the IPMI-over-IP protocol has been deactivated.

The "Activation" phase of the IPMI-over-IP protocol is only described in version 2.0 of the IPMI specification, the older version 1.5 is out of scope for this paper.

The "Activation" phase of the IPMI-over-IP protocol in version 2.0 of the IPMI specification starts with the client sending an *RMCP+ Open Session Request* packet to the BMC which responds with an *RMCP+ Open Session Response* packet. These packets contain session IDs for further communication between client and BMC. The response packet also contains information about supported cipher suites.

Next, the client sends an *RAKP Message 1* packet answered by the BMC with an *RAKP Message 2* packet. These packets contain nonces for mutual authentication (later signed using the user's password) as well as the client's username and the BMC's GUID (globally unique ID). Since the *RAKP Message 2* packet is already signed using the password of the requested username, it is possible to perform an offline brute-force attack on the password if the requested username is valid. It is also possible to perform an online brute-force attack on the username, since the BMC tells the client whether the username is valid or not.

Finally, the client sends a signed *RAKP Message 3* packet answered by the BMC with a signed *RAKP Message 4* packet. The signature is made using the user's password similar to the *RAKP Message 2* packet.

D. Different Scan Types

It is possible to scan for IPMI devices in various ways. The IPMI-over-IP protocol defines two different discovery methods, both over UDP port 623.

BMCs queried with *Get Channel Authentication Capabilities Request* packets only reply if the IPMI-over-IP protocol has been activated. The response packet contains information about the IPMI version of the BMC (1.5 or 2.0) as well as some information about supported authentication methods.

BMCs scanned with *RMCP Ping* packets reply with *RMCP Pong* packets. The response packets contain little to no information about the BMC other than its presence. However, small-scale tests on a *Dell iDRAC 7* device show that the BMC replies to *RMCP Ping* packets even if the IPMI-over-IP protocol has been deactivated. That is why we presume to find additional dark IPMI devices by scanning with *RMCP Ping* packets.

E. SSL Basics

SSL is a security protocol based on a Public Key Infrastructure. It is used in the WWW, but also for email, chats, and other services. SSL certificates contain identity information about a peer (e.g., a domain name) and a corresponding public key. A certificate therefore creates a binding between an identity and a public key.

We will use SSL certificates in our research in the following two ways: First, to identify potential IPMI devices for scanning (see Section IV-B). Second, to evaluate the security of IPMI devices with regards to cipher security and the potential of Man-in-the-Middle attacks (see Section V-D1).

IV. APPROACH

In this section we describe the rationale and approach of our IPMI measurements. We begin by explaining the concept of dark IPMI devices. Then, we detail the two types of measurements: The first one is limited to a small subset of IP addresses including likely dark IPMI devices. The second type of measurement is a complete scan of the IPv4 space. Finally, we detail the used scanning software and address ethical questions regarding active network measurements.

A. Dark IPMI Devices

With our measurements we want to discover *dark IPMI devices*. These are devices which have the IPMI-over-IP port disabled. Consequently they do not respond to standard IPMI scans such as those executed by Rapid7 [8]. They do, however, respond to *RMCP Ping* requests as required by the IPMI specification [12]. Even though dark IPMI devices do not provide direct IPMI-over-IP access, they are still valuable to attackers. Once identified as an IPMI device, attackers could exploit other attack vectors, e.g., flaws in the web interface implementation or insecure SSH connections [13] to gain access to the BMC or the host system.

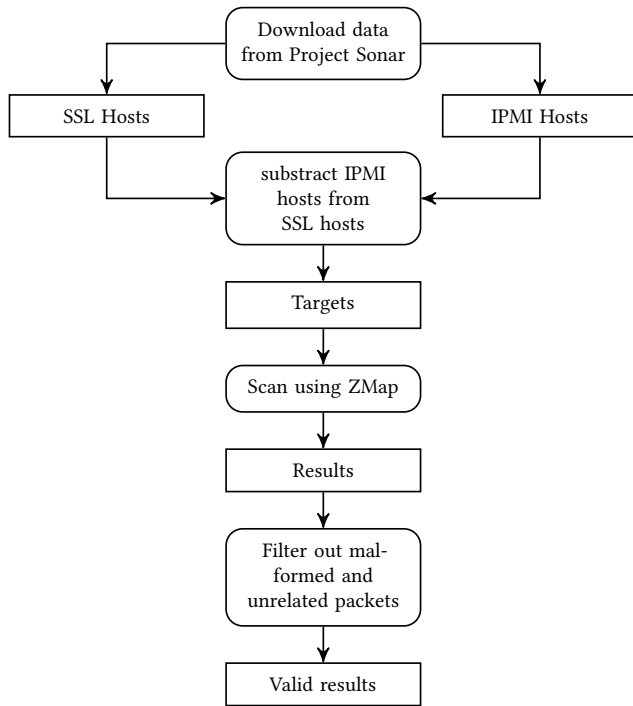


Fig. 2. Flow chart of discovering dark IPMI devices used in the first scan.

B. Target List

In order to verify that there are indeed dark IPMI devices, we execute active measurements on a specific subset of IP addresses. As a starting point we use all SSL hosts identified by Project Sonar [8] as most IPMI devices provide access via a web interface. Then, we remove the IP addresses which already responded to standard IPMI scans using Get Channel Authentication Capabilities requests. These IP addresses have already been attributed to IPMI devices and are therefore not *dark*. We use the remaining IP addresses in the first type of active measurements. Figure 2 shows the workflow of this type of measurement.

In the second type of active measurements we probe the full IPv4 space to get a complete picture of IPMI deployment.

C. ZMap

ZMap is a network analysis tool designed for scanning different network ports across the IPv4 Internet [14]. ZMap has been optimized for speed, meaning that it is capable of scanning the entire IPv4 space in less than 5 minutes given enough bandwidth [15]. It is possible to load custom modules for packet generation or result processing.

We built a probe module for ZMap which generates *RMCP Ping* packets. Moreover, we extended ZMap to filter out incoming UDP packets that are not addressed to the scanning machine (e.g., multicast packets). The modified version is available on our website.¹

¹<https://www.net.in.tum.de/pub/zmap/zmap-ipmi.tar.gz>

D. Ethical Considerations

Active network measurements have to be conducted in a sensible and ethical way in order not to induce negative consequences. Partridge and Allman [16] propose to evaluate whether the active measurements themselves or the release of the resulting data can harm an individual. Therefore we apply precautionary measures to reduce the impact of our IPMI scans.

First, we try to minimize the load on target networks generated by our research activities. Therefore, we do not scan with maximum speed but rather constrain our scans to 1 Mbit s^{-1} and 10 Mbit s^{-1} respectively. Additionally, ZMap’s randomization feature ensures an even distribution of probes destined to a subnet over time.

Second, we use RMCP Ping requests which are less intrusive compared to Get Channel Authentication Capabilities requests. An RMCP Ping request does not attempt any authentication or login. No sensitive information other than the existence of the device itself is sent in the RMCP Pong response message. Furthermore, it does not show up in the IPMI device’s log file, additionally reducing the number of false alarms.

Third, we set up a web site on the scanning machine with information about our research activities. Moreover, we provide a dedicated email address for information and blacklisting purposes. We offer network administrators the possibility to send their emails encrypted with our PGP key.

Fourth, we exclude IP addresses whose network administrator indicated in the past that they did not want to be scanned. Before our experiments the blacklist contained 148 entries resulting in 2 079 222 IP addresses. During our scans we received only two emails which were automatically sent by intrusion detection systems to the abuse contact listed in the WHOIS database. We answered both emails by providing more information with regards to our activity and offered the network administrators to put their IP ranges on our blacklist. We did, however, not receive a reply.

Fifth, we conduct an internal review at our Chair before starting any network experiment. This ensures that ethical and procedural concerns are addressed in advance and multiple viewpoints are being considered.

This research is conducted under consideration of the two ethical questions raised by Partridge and Allman [16]. We believe that the five precautionary measures ensure that the collection of data in this study does not cause tangible harm to any person’s well-being. Furthermore, since we do not have plans to release the collected data, no private or confidential information is published. The results presented in this paper give a general overview but no one specific individual or host is identified.

V. EVALUATION

In this section we present the results from the two scans conducted during this work. First, we will give an overview of the two scans. Then we will go into detail with regards to the responding hosts by comparing our measurement technique

TABLE I
OVERVIEW OF BOTH SCANS

Scan	Targets	Number of targets	Responding IPs	Valid responses	Hit rate	Scanning rate	Duration
1	HTTPS hosts w/o known IPMI hosts	33 131 598	38 180	37 205	0.11 %	1 Mbit s ⁻¹	7 h 52 min 44 s
2	Complete IPv4 space	3 700 190 401	400 299	225 612	0.01 %	10 Mbit s ⁻¹	2 d 21 h 5 min 11 s

with the state of the art. Following, we evaluate deployment practices and uncover significant clustering in certain parts of the network. Then, we evaluate the security of IPMI device, specifically their SSL certificates and the supported IPMI versions. Finally, we classify the evaluated results and give concrete advice on hardening IPMI deployments in a network.

A. Scan Overview

We conducted two active measurement runs: the first was conducted on likely newly discoverable IPMI hosts, the second was on the complete IPv4 space. Table I shows an overview with statistics about both scans.

The first scan’s purpose was to gather additional active IPMI hosts compared to other IPMI scanning projects. In contrast to Project Sonar’s regular IPMI scans [8] which use Get Channel Authentication Capabilities packets, we use RMCP Ping packets (see Section III-D). This allows us to find IPMI devices which do not answer to Get Channel Authentication Capabilities requests. These dark devices have IPMI-over-IP deactivated, however other network interfaces (e.g., web interface) might still be accessible. Therefore RMCP Ping requests allow us to estimate the number of accessible IPMI devices more accurately. We decided to scan all hosts with SSL certificates minus the IPs where an IPMI device has already been detected by Get Channel Authentication Capabilities scans. Thus the responding IPs are IPMI devices which could not be detected by Get Channel Authentication Capabilities scans because the IPMI-over-IP interface has been disabled. These devices do not pose a direct security risk as IPMI-over-IP is disabled. However, other access methods such as the web interface still pose a threat as shown by Bonkoski *et al.* [2].

The second scan covers the complete IPv4 space. This allows us to find all publicly accessible IPMI devices in the IPv4 Internet and therefore gives us the complete picture.

For both scans we used the scanning tool ZMap [14] (see Section IV-C for more details). We used a blacklist to exclude hosts and subnets whose network administrators did not want to be scanned. Both scans were run from a physical machine on a mid-range server with a quad core Intel Core i7 CPU and 8 GiB RAM, the operating system was Debian 8 Jessie.

B. Responding Hosts

In our first scan on 33 131 598 target IPs we received replies from 38 180 different IPs. After filtering out malformed and unrelated packets, 37 205 valid responses were left. This corresponds to a hit rate of 0.11%. Since we excluded the




Fig. 3. Number of valid responses from Project Sonar’s IPMI data sets over time.²Note the crunched y axis for increased readability.

results of Project Sonar’s IPMI scans from our target list, these responses come from dark IPMI devices, i.e., devices with IPMI-over-IP disabled.

Our second scan was conducted on the entire IPv4 space minus our blacklist. We got replies from 400 299 different IPs with 225 612 valid responses from IPMI devices. This corresponds to a hit rate of 0.01%.

We compared our results with Project Sonar’s IPMI scan from September 7, 2015. We removed 720 blacklisted IPs from Project Sonar’s results to improve the comparability. Our second scan delivered 21.22% more results than Project Sonar’s IPMI scan. This shows that our scanning approach is able to identify significantly more IPMI devices than other state of the art scanning methods.

When comparing our results to scans conducted by Dan Farmer in 2013 [4], we find slightly less devices than his 230 000. However, this can be explained by the general decrease of reachable IPMI devices over time. Figure 3 shows the number of valid IPMI responses obtained from Project Sonar scans between June 2014 and December 2015. Note that the y axis of the figure is crunched to increase readability. We can clearly see that the number of valid responses is dropping steadily, by a total of 65 090 devices in the observed period. The decreasing number of IPMI host could be a result of previous work pointing out security risks with IPMI [2], [5], [6]. In consequence, network administrators might have isolated IPMI devices from the public Internet. Unfortunately,

²According to remarks by Rapid7 employees the strange valley between February and March 2015 in Figure 3 is most likely a measurement artifact.

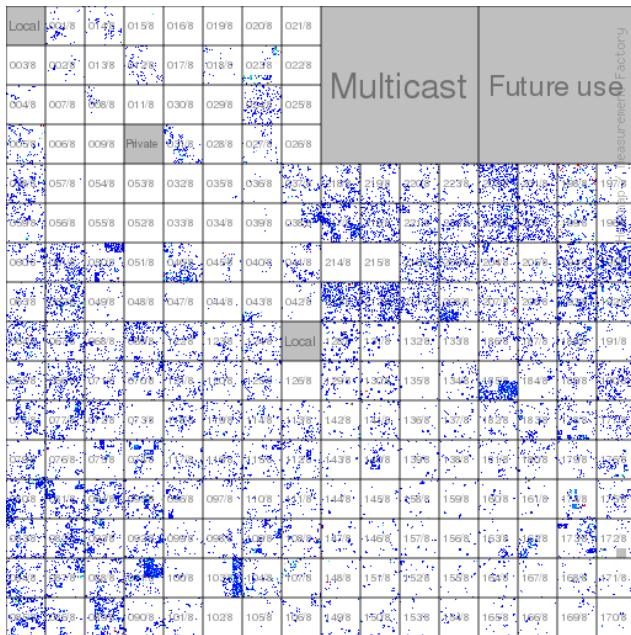


Fig. 4. Hilbert curve of responding IPMI devices in the second scan.

we could not compare our measurement method with Dan Farmer’s as no detailed description was provided. Finally, Dan Farmer did not specify whether a blacklist was used during his scans.

C. Deployment Practices

In this section we evaluate the results of the second scan covering the complete IPv4 space with regards to deployment practices. The question arises if we can find certain subnets with a significantly higher IPMI density.

To better visualize connected subnets but not constraining ourselves to a certain prefix length we use a Hilbert space-filling curve. Figure 4 shows the distribution of identified IPMI devices during the second scan. The figure shows a heat map of the IPMI deployment in the complete IPv4 space. We visually highlight /8 networks to make it easier to find specific parts of the Internet. Each pixel represents one /18 network, the color indicates the number of IPMI devices found in this /18, ranging from blue (few IPMI devices) to red (many IPMI devices).

We can see that generally the IPv4 space is sparsely populated with IPMI devices. This is no surprise and corresponds to the hit rate of 0.01%. However, IPMI devices are not uniformly distributed over the IPv4 space. They seem to be concentrated in some subnets whereas other subnets are completely blank indicating that no IPMI device is reachable from the public Internet. We suspect that the former could be stemming from data centers and hosting providers, whereas the latter would include private customers including DSL, cable, and fiber lines. To further analyze this scenario we take a look at parts of the Internet with a high IPMI density in more detail.

Thus we evaluate the density of IPMI devices based on Autonomous Systems (ASes).

TABLE II
TOP 10 ASes WITH MOST IPMI DEVICES

Pos	ASN	AS	# IPMI Devices
1	2914	NTT-COMMUNICATIONS-2914	30308
2	15003	NOBIS-TECH	5447
3	33781	OPQ	4687
4	16596	Univ. de Baja California	4140
5	5461	OKB MEI	3796
6	28227	NOVACIA	3281
7	35662	Redstation Limited	3132
8	60781	LeaseWeb-NL	2836
9	2607	SANET	2830
10	18978	Enzu Inc	2810

We apply CAIDA’s Prefix to AS mapping [17] to match IP addresses to their respective Autonomous Systems.³ We find IPMI devices in a total of 7580 ASes. Table II shows the top 10 ASes with the most IPMI devices. It is astounding that the top AS owned by NTT Communication has more than 30 000 IPMI devices. NTT is a provider of network management solutions. The AS of NTT has more IPMI devices than the next eight ASes combined. Since NTT is one of the largest network services providers announcing numerous IPv4 prefixes (e.g., also for the Akamai CDN) and operating one single global AS this is not surprising.

The rest of the top 10 ASes is with three exceptions made up of hosting providers. Two entries are ASes pertaining to academic institutions from Mexico and Slovakia respectively. Interestingly, one AS is a special Russian agency (“Experimental Design Bureau”) for the purpose of developing aerospace and land-based antenna systems. For research purposes they also operate their own supercomputer.

As can be seen in Table II the number of IPMI devices per AS steeply decreases. In addition, Figure 4 further suggests that there are a few networks with many IPMI devices whereas most have little to none. To further investigate this phenomenon we plot the cumulative distribution function of IPMI devices in Autonomous Systems. Figure 5 shows the percentage of IPMI devices per AS. Note that the x axis’ scale is logarithmic as otherwise the function would almost immediately rise to the top due to the exponential increase. We see that the top 10 ASes are home to almost 30% of the Internet’s IPMI devices. As expected, the number of IPMI devices added per additional AS sharply decreases: The increase from 10 to 100 ASes adds about 30% of IPMI devices, the same percentage as from AS 100 to AS 1000.

To conclude, IPMI devices are not uniformly distributed in the Internet, but rather concentrated in specific subnets (see Figure 4) and Autonomous Systems (see Table II and Figure 5). This hints at a general deployment issue with regards to IPMI: some network administrators and organizations do not seem to deem it necessary to secure their IPMI deployment which leaves them open to exploitation.

³MOAS are counted towards one AS only (according to CAIDA’s deterministic sorting).




Fig. 5. Cumulative Distribution Function of IPMI Devices in Autonomous Systems in the Second Scan. Note: The x axis is log-scaled.

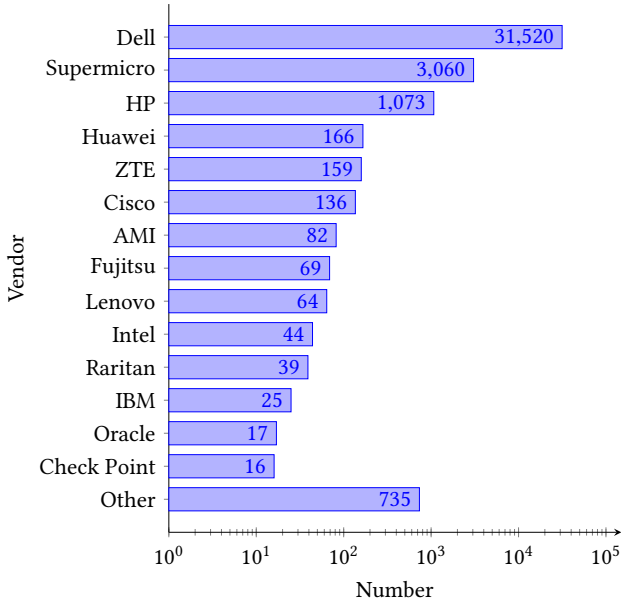


Fig. 6. BMCs of the first scan by vendor. Note: The x axis is log-scaled.

D. Security Analysis

1) *SSL Evaluation*: Since the target list of our first scan is based on SSL scan data published by Project Sonar, we are able to analyze the SSL certificates of the first scan’s results.

First, we examine the use of default certificates in BMCs by considering the *Common Names* (CNs) and SHA1 checksums of the certificates. We find that 83.31% of BMCs use default certificates, whereas 2.74% of BMCs use custom generated certificates. A special case are 13.95% of BMCs which seem to auto-generate their certificates upon installation. These exhibit the same CN schema but differ in the SHA1 checksum. Default certificates can be misused by extracting private keys from the firmware and therefore compromising the device’s security [6], [7].




Fig. 7. Cumulative distribution function of SSL key lengths in the 1st scan.

Additionally, we are able to determine the manufacturers of most IPMI devices by matching their certificate CN. Figure 6 shows the vendors of the discovered IPMI devices with SSL enabled during the first scan. More than 84% of BMCs have been manufactured by Dell, followed by Supermicro and HP.

We also investigate the number of certificates that are self-signed by building each host’s trust chain. We find that 95.87% of BMCs use self-signed certificates, while the certificates of 4.13% BMCs are signed by issuer certificates. These issuer certificates, however, are mostly not trusted by modern browsers.

In addition we evaluate the key length of every BMC’s client certificate. The results are shown in Figure 7 as a cumulative distribution function. Almost all (89.24%) of keys are 1024 bit or shorter. This is not secure by today’s standard as NIST proposes key lengths of at least 2048 bit [18].

Finally, we analyzed the key types of the BMCs’ client certificates. This reveals that 99.27% of certificates use RSA keys, only 0.73% use DSA keys.

2) *IPMI Versions*: RMCP Pong response packets like those gathered in our scans do not contain information about the supported IPMI version of the BMC. Therefore we use IPMI scan data provided by Project Sonar [8] to analyze supported IPMI versions. After filtering out malformed and unrelated packets, we find that about 63.03% of IPMI devices support IPMI version 2.0 whereas the remaining 36.97% only support IPMI version 1.5. This suggests that IPMI implementations are rarely updated since IPMI 2.0 was specified in 2004.

E. Best Practices

In this section we summarize our findings with regards to IPMI’s operational security and give concrete advice to improve it.

Using the RMCP Ping discovery technique we find 21% more devices than with conventional Get Channel Authentication Capabilities requests. After exploiting other weaknesses,

e.g., in the SSH implementation [13], these additional devices can be used by an evil actor to take over the device and then stage subsequent attacks. This is especially troubling when taking SSL certificate weaknesses into account. Since it is not enough to simply deactivate IPMI-over-IP on the BMC, we strongly recommend to block incoming requests on UDP port 623 using an external firewall.

We found that IPMI devices are heavily clustered in subnets and ASes. Once an IPMI device has been discovered, an attacker can probe for more devices in its vicinity. This makes the scanning approach even stealthier, since not all IP addresses need to be probed to get a certain number of IPMI devices. Again, an outward facing firewall should block foreign IPMI traffic coming to the subnet or AS.

Most IPMI devices also offer access via an HTTPS web interface. This, however, introduces an additional attack vector since secure deployment of SSL is non-trivial. Our results show that most BMCs use default certificates included in the firmware and use very short keys. This makes it possible to perform various attacks preceded by Man-in-the-Middle attacks. We propose to use self-generated SSL certificates with strong keys. These certificates could be signed by a trusted party, e.g., the company's own certificate authority.

All of the above issues can be circumvented from being attacked by the outside if IPMI devices are only accessible within a VPN. This ensures that all interfaces and services (e.g., IPMI-over-IP, HTTPS web interface,...) are contained in a separate network. However, users with access to the VPN can still be a threat to these IPMI devices.

VI. CONCLUSION

IPMI is the de facto protocol for Out-of-band network management. Its ubiquitous use and the potential benefit from compromising an IPMI device makes it a prime attack target.

In this paper we survey the current state of the IPMI deployment in the Internet. We present a method for finding dark IPMI devices. These devices have not been found using conventional methods. We then analyzed the distribution of IPMI devices in the network and found that they are heavily clustered. About one third of all devices are located in only 14 Autonomous Systems. The state of IPMI devices' SSL security is rather troubling as well. Most devices use default certificates included in their firmware and offer weak keys to SSL clients. Finally, we give concrete advice to increase the security of IPMI deployments in the network.

Future Work. To further improve IPMI security more intrusive scans could be conducted. Moreover, we would like to conduct scans with our method and Farmer's in parallel. However, we feel that to adhere to our ethical standards these measurements should only be done with the agreement of the network's owner. Additionally, the work to analyze and extract information from IPMI firmwares could be intensified.

Acknowledgments. We thank the Leibniz Supercomputing Centre (LRZ) of the Bavarian Academy of Sciences and

Humanities (BAW) for the provisioning and support of Cloud computing infrastructure essential to this publication. This work has been supported by the German Federal Ministry of Education and Research, project Peeroskop, grant 01BY1203C, and project SURF, grant 16KIS0145, and by the European Commission, project SafeCloud, grant 653884.

REFERENCES

- [1] HD Moore, "A Penetration Tester's Guide to IPMI and BMCs," Jan. 2013, accessed: 2015-12-14. [Online]. Available: <https://community.rapid7.com/community/metasploit/blog/2013/07/02/a-penetration-testers-guide-to-ipmi/>
- [2] A. Bonkoski, R. Bielawski, and J. A. Halderman, "Illuminating the security issues surrounding lights-out server management," in *Proceedings of the 7th USENIX Workshop on Offensive Technologies*. Berkeley, CA: USENIX, 2013.
- [3] D. Farmer, "IPMI: Freight Train to Hell," 2013. [Online]. Available: <http://fish2.com/ipmi/train.pdf>
- [4] —, "Sold Down the River," 2013. [Online]. Available: <http://fish2.com/ipmi/river.pdf>
- [5] J. Zhang, Z. Durumeric, M. Bailey, M. Liu, and M. Karir, "On the mismanagement and maliciousness of networks," in *Proceedings of the Symposium on Network and Distributed System Security (NDSS)*, San Diego, CA, 2014.
- [6] A. Costin, J. Zaddach, A. Francillon, and D. Balzarotti, "A large-scale analysis of the security of embedded firmwares," in *Proceedings of the 23rd USENIX Security Symposium*. San Diego, CA: USENIX Association, Aug. 2014, pp. 95–110.
- [7] S. Viehböck, House of Keys: Industry-Wide HTTPS Certificate and SSH Key Reuse Endangers Millions of Devices Worldwide. SEC Consult. Accessed: 2015-12-01. [Online]. Available: <http://blog.sec-consult.com/2015/11/house-of-keys-industry-wide-https.html>
- [8] Project Sonar. Rapid7. Accessed: 2015-10-28. [Online]. Available: <https://sonar.labs.rapid7.com/>
- [9] R. Holz, L. Braun, N. Kammenhuber, and G. Carle, "The SSL landscape: a thorough analysis of the X.509 PKI using active and passive measurements," in *Proceedings of the 2011 ACM SIGCOMM Internet Measurement Conference (IMC)*. Berlin, Germany: ACM, 2011, pp. 427–444.
- [10] N. Heninger, Z. Durumeric, E. Wustrow, and J. A. Halderman, "Mining Your Ps and Qs: Detection of Widespread Weak Keys in Network Devices," in *Proceedings of the 21st USENIX Security Symposium*, Bellevue, WA, 2012, pp. 205–220.
- [11] O. Gasser, R. Holz, and G. Carle, "A deeper understanding of SSH: Results from Internet-wide scans," in *Proceedings of the Network Operations and Management Symposium (NOMS)*. Krakow, Poland: IEEE, 2014, pp. 1–9.
- [12] *Intelligent Platform Management Interface Specification Second Generation*, Intel, Hewlett-Packard, NEC, and Dell, Feb. 2014, revision 1.1. [Online]. Available: <http://www.intel.com/content/www/us/en/servers/ipmi/ipmi-v2-rev1-1-spec-errata-6-markup.html>
- [13] Juniper. Out of Cycle Security Bulletin: ScreenOS: Multiple Security issues with ScreenOS (CVE-2015-7755, CVE-2015-7756). Juniper. Accessed: 2015-12-22. [Online]. Available: <http://kb.juniper.net/InfoCenter/index?page=content&id=JSA10713>
- [14] Z. Durumeric, E. Wustrow, and J. A. Halderman, "Zmap: Fast internet-wide scanning and its security applications," in *Proceedings of the 22nd USENIX Security Symposium*, Washington, D.C., 2013, pp. 605–620.
- [15] D. Adrian, Z. Durumeric, G. Singh, and J. A. Halderman, "Zipper zmap: internet-wide scanning at 10 gbps," in *Proceedings of the 8th USENIX Workshop on Offensive Technologies*, San Diego, CA, 2014.
- [16] C. Partridge and M. Allman, "Addressing ethical considerations in network measurement papers: Abstract," in *Proceedings of the 2015 ACM SIGCOMM Workshop on Ethics in Networked Systems Research*, ser. NS Ethics '15. New York, NY: ACM, 2015, pp. 33–33.
- [17] Routeviews Prefix to AS mapping. CAIDA. Accessed: 2016-03-17. [Online]. Available: <http://data.caida.org/datasets/routing/routeviews-prefix2as/2015/09/routeviews-rv2-20150926-1000.pfx2as.gz>
- [18] E. Barker, W. Barker, W. Burr, W. Polk, and M. Smid, "NIST Special Publication 800-57," *NIST Special Publication*, vol. 800, no. 57, pp. 1–147, 2012.