Future threats to future trust

Herbert Bos[†], Sotiris Ioannidis[*], Erland Jonsson[‡], Engin Kirda[⊕], Chris Kruegel[°]
[†]Vrije Universiteit Amsterdam, [*]FORTH-ICS, [‡]Chalmers, [⊕]Institut Eurecom, [°]TU Wien

I. The FORWARD workshop and threats on the net

Only a few years ago, big worms roamed the planet, spreading within hours, or even minutes, to every nook and cranny of the Internet. The damage caused by them was equally impressive; worms have taken out alarm phone centers, train signalling systems, thousands of cash machines, millions of production PCs and servers, and, oh yes, South Korea[1].

No wonder academics and industry scrambled to counter the threat. Indeed, fast spreading flash worms were all the rage among security experts and millions of euros were spent on projects to counter them. Alliances were formed, research grants applied for, projects started, prototype solutions developed, refined, and discarded. Unfortunately, by the time we developed practical counter measures, flash worms had all but disappeared. Instead, we now worry about stealth attacks, botnets, phishing sites, attacks on mobile phones, and whatever new threats emerged in recent years. The problem is that we tend to work on solutions for today's problems and have no time to worry about the threats of the future. The problem is that we are often caught unawares.

This need not be the case and there are examples of threats that we saw coming before they hit us. A well-known example is RFID. An RFID tag is a small, extremely low-cost chip that can be used for purposes like identification and minimal processing. By adding RFID tags to everything, from pets to products, industry aims to use RFID technology to create the "Internet of Things". However, researchers have shown that tags can be used to propagate malware, which in turn has led a concerned industry to scrutinize security issues in RFID. All of this happened before any real attacks took place.

For this reason the FORWARD initiative intends to bring together experts to discuss future threats and develop realistic threat scenarios. As a first step in that direction, a workshop was organized in Göteborg, Sweden, in April 2008, to discuss future threats [1]. The workshop consisted of broad plenary sessions interspersed with focused experts meetings. This paper summarizes the workshop's findings and their bearings on the future of trust. The remainder of this paper discusses the findings of the targeted expert meetings on critical infrastructure and large scale systems (Section II), fraud (Section III), and malware (Section IV). Concluding remarks are in Section V.

[1]The country virtually dropped off the map as a result of the Slammer worm [2].

II. Trust in critical infrastructure and large scale systems

The systems and networks that constitute critical infrastructure are often taken for granted. Many times people only realize their dependence on these services when there is a disruption. Yet, when such disruptions do happen, they may have serious, even dire consequences. Moreover, as witnessed by the Y2K issue in 1999, even the advertized presence of *potential* problems can be disruptive.

In the past, the systems and networks of the infrastructure were physically and logically independent and separate. They were not connected, and there was little or no interactions between them. With advances in technology, however, this has changed. In each sector, the systems have become automated and interlinked through computers and communication facilities. Furthermore, the trend shows an increase of both automation and linkage, not only within sectors but also between various sectors. Thus, we expect the future will aggravate the interdependencies between systems in general, and systems related to critical infrastructure in particular, leading to a complex "mesh of systems."

While increasing efficiency, interlinked capabilities also render the systems and networks more vulnerable. Not only have the possible vectors for a determined attack or simple harmful influence increased, also the detrimental effects of a service disruption in a single sector have significantly increased. What would have been an isolated incident in the past, can today cause extensive interruptions and/or failures in other sectors as well. In fact, the cascading effects might lead to a more or less

global outage or malfunction, affecting systems and networks in even seemingly unrelated sectors. If such cascading effects cannot be contained, they will directly influence both the economy of society and the physical safety of its citizens. In certain cases, adding to the vulnerability of the system may be unacceptable, and we should question whether interlinking the system should be permitted at all.

As mentioned earlier, serious disruptions have already affected such infrastructures as train signalling systems, cash machines and phone systems. Intelligence services have indicated that targeted cyber attacks have caused power-outages in multiple cities in the past [3].

Future vulnerabilities. Even in the absence of attacks, bugs have shown to have devastating effects on infrastructure. A disturbing example is what is known as the Northeast Blackout on August 14, 2003, which affected some 50 million people and caused approximately 6 billion dollars in financial losses [4]. The outage had a variety of knock-on effects, such as the break-down of much of the public transport. One of the major causes of the black-out was an unlikely race condition that occurred in the system that dealt with failures in the control system.

This brings us to an important conclusion of the experts meeting in Göteborg. By interconnecting more and more systems and adding more parallelism to individual systems (e.g., by multi-core processors) concurrency is entering all aspects of computing. As a result, the future of trust in computing increasingly hinges on our ability to deal with concurrency vulnerabilities that are extremely hard to find and difficult to trigger.

Another crucial factor is the human one. Members of the expert group shared that in their experience the disruption in interconnected networks is often not caused by a deliberate and malicious activity, but simply by human errors (e.g., router misconfiguration.)

Scale. Most of the challenges in this area are caused by the scale. What we need to deal with is (sets of) large software systems, of huge complexity and sometimes heavily distributed. In addition, we have systems with huge numbers of mobile devices (phones, RFID). Problems in the area of large software systems include concurrency, authorization, and integration. In the area of ”many devices,” the issues revolve around authorization (if people have many devices at home, how do they secure those devices?), sensors that might be fooled, and management of these systems. For both areas, the expert meeting concluded that we must be able to cope with partially compromized systems.

Trusted computing in the form of TPMs is sometimes seen as a panacea for trust in distributed computing. While it is true that TPMs allow one to verify the configuration of remote machines, some inevitable problems come up when applying TPM to large systems. For instance:

- Yes, I can verify that some remote system runs the software that I intend to use, but how do I verify that my own machine is not compromized? The only way to do this is by means of yet another, more trustworthy, device (a mobile phone perhaps?), but how do we make sure that this is not compromized? Yet another, smaller, device? And how do we verify *that*? Where is the root of trust, and what happens if the dog eats it?
- What do we verify? Systems may run a huge amount of complex software and reliably checking the configuration of a large number of devices is exceedingly difficult. Moreover, as end-users cannot be expected to verify each and every systems involved in an interaction with a distributed system, checking has to be delegated in a chain- or tree-like fashion. Any unnoticed compromized system in the hierarchy invalidates the trust in an entire branch of the tree. Worse, this would not be noticed.
- Even if we do notice it, because a remote attestation fails, what does it mean if a chain cannot be verified? Rebooting in a known clean state is often not a solution for mission-critical large scale systems. How can we continue operating when the trust is violated?
- Finally, not all devices are smart enough to be trusted. Phrased differently, they may not have a TPM today, and most likely will not have a TPM in five years time either. Examples include small embedded devices (say the category that sits between between mpg players and RFID).

Trusting the network. A final conclusion is that the most critical infrastructure of all is the communication system to which the critical system as such is connected. In almost all situations, this is the Internet. A disruption in the network that mediates their interrelations might have more devastating effects than a successful attack on one of the connected systems by itself.

One important issue are threats to the Internet routing infrastructure. Internet routing (BGP) is

vulnerable against attacks. In particular, false or spoofed BGP network announcements can be honored by parts of the Internet. This may result in DoS attacks against large parts of the network or hijacking of, for example, well-known web sites during the time the false information is valid. Other problems arise from mistakes caused by (trusted) operators when configuring routers or entering routing information which could have similar effects on the Internet. Yet another of type problem are DDoS attacks against BGP routers, which may have the effect of making parts of the Internet temporarily inaccessible. A single router can also be attacked and its traffic sent via a tunnel (e.g., GRE) to a remote site that can then act as a man-in-the-middle for arbitrary domains and servers.

The problems arise from the fact that the current protocol, BGPv4, is 12 years old, and it was not designed with the current Internet in mind. Furthermore, BGPv4 is here to stay for a very long time, which means that threats are going to follow us in the near and long term future. Even though solutions exist, everyone must start using them at the same time, something which is not likely to happen. Countering future threats would involve (i) motivating vendors to implement solutions, and (ii) somehow extending BGP in a backwards compatible way to make sure the new functionality is used.

More secure routing protocols exist (S-BGP, soBGP), and can be used to verify the origin and correctness of the received information. However, BGP signatures are problematic. The solution may be to move this to out-of-band systems, since all routers are CPU-limited. Also, Moore's law does not help router builders, since density and power remain as issues as more capacity is added. It seems that in the future, there will be no need to propose new routing protocols, unless they offer some really great properties, and as mentioned before, old threats will remain.

III. Fraud and the lack of trust

Online scams are a form of online fraudulent activity in which an attacker aims to steal a victim's sensitive information, such as an online banking password or a credit card number. Victims are tricked into providing such information by a combination of spoofing techniques, social engineering, and sometimes advanced exploitation methods.

According to the participants of the expert meeting on fraud, one of the main reasons why online fraud is increasingly gaining in popularity is because Internet-based attacks are difficult to trace back. Furthermore, fraud on the Internet is easy to perform as a high number of users exist that are technically unsophisticated and are still not highly familiar with the Internet technology. For example, the effort required to launch a physical attack against a bank is very high (e.g., breaking in, armed robbery, etc.) in comparison to hosting a phishing web site and waiting for victims to simply enter their sensitive information.

In addition, the meeting concluded that law enforcement agencies are either slow to react or do not have the necessary technical skills to identify the miscreants. With respect to traditional crime, crime on the Internet is much faster and typically more "international." That is, even if the attack takes place in Europe, the servers participating in the attack (e.g., phishing sites) might not necessarily be located in the same region. Hence, cross-border communications is often necessary, which is a time-consuming and tedious process. Miscreants responsible for the attacks are well-organized and know very well how law enforcement and the targeted organizations operate. For example, many attacks are now launched over the weekend because fewer experts are at work during this time (which in turn results in slower responses).

Trust among the good guys. One issue that was discussed in the fraud meeting was whether exchanging data would help mitigate fraud-related attacks. All participants in the expert group thought that this was a good idea and that it could actually help. For example, it is certainly interesting for banks to find out if there are similar attacks happening elsewhere and what solutions other organizations use. Also, organizations are interested in knowing if certain malware specifically targets them before the attack largely seen in the wild. However, it was not clear how such a data exchange should be performed. That is, while many organizations are certainly interested in getting information and data, they are less excited about giving away information as they have privacy as well as security concerns. It is clear that a common basis of trust needs to be created among organizations so that they are willing to share sensitive information. Currently, some organizations (e.g., banks) are not even willing to talk about the problems they face as they are afraid that the information that they give out can be used against them in some way.

The underground economy. One interesting research challenge with respect to online fraud is to be able to understand how the underground Internet economy actually works. For example, if we were to start a botnet business, how would

we actually go about and communicate with our "customers"? How would we sell our services and initiate money transfers? Hence, by understanding the way this new type of illegal economy functions, the participants of the fraud group believe that solutions could be created that actually undermine this economy and significantly increase the effort required by the miscreants.

IV. Malware

Malicious code (or malware) is defined as code that fulfills the harmful intent of an attacker. Typical examples include viruses, worms, and spyware. One reason for the prevalence of malicious code on today's networks is the rising popularity of the Internet and the resulting increase in the number of available vulnerable machines because of security-unaware users. Another reason is the elevated sophistication of the malicious code itself.

Nature and form of malware. One issue raised in the experts meeting was about the behaviour of malicious code and their sources. Surprisingly, perhaps, the basic functionality of malware has not changed much. The samples that are observed today either steal sensitive information (key loggers, password thieves, Bank Trojans), send spam mails, or can be used to launch denial of service attacks. The real development is in the way in which the malicious code is written. In addition to obfuscation to evade traditional, signature-based detection, malicious code increasingly tries to evade analysis. That is, by including code that detects virtual machine environments or debuggers, human or automated analysis is made more difficult. Thus, one finding of the meeting was that we expect a significant increase of novel techniques that stealthy, malicious code uses to resist analysis and thwart detection.

Threat landscape. Another question concerns the change in the threat landscape over the last years. There was agreement that most malware is actually coming from a (relatively) small number of criminal groups that have a well-funded development process and a pool of talented developers. These groups use those venues that can be most easily exploited to inject their code on end-user machines. For this, there is a strong trend towards social-engineering-based attacks (such as email) or browser-based exploits compared to exploiting network services. As a result, novel mechanisms for data collection are needed. For example, a traditional honeypot might not be efficient anymore to capture the current threats. This was confirmed by numbers from VirusTotal, which showed a discrepancy between the malware that they see compared to the samples that are collected via traditional honeypots. Moreover, the adversary might have developed techniques to fingerprint and detect honeypots so that they can avoid detection. Finally, mapping out dark (honeypot) address spaces is an emerging threat. As a result, the expert group saw the need to develop techniques that can accurately capture emerging threats, since a good intelligence is a prerequisite for subsequent mitigation efforts.

Related to the previous issue, the group also discussed emerging targets of malicious code. In particular, the question was raised whether mobile devices (phones, PDAs) might become a target. Everybody agreed that the threat has been hyped in the last few years. However, once there is a business model behind attacking phones (i.e., it turns out to be profitable for the criminals), such attacks can be expected to appear. Also, this development will be supported by the significant growth in the number of mobile devices.

V. Concluding remarks

The expert meetings in Göteborg serve as a starting point for developing a research agenda to deal with future threats. Within the FORWARD project the conclusions of the workshop are used to establish working groups, each of which work towards in-depth analysis of a subdomain. So far, the following working groups have been created: (WG1) Smart Environments, (WG2) Malware and Fraud, and (WG3) Critical Systems. Besides analysis in the broad sense, working groups will develop specific threat scenarios in which future threats are worked out in detail. The threat scenarios will be consolidated and worked into a white book.

Acknowledgements

FORWARD is sponsored by the European Community's 7th Framework Programme (FP7/2007-2013) under grant agreement no 216331.

References

[1] The 1st FORWARD workshop, Göteborg, Sweden, April 2008. `http://www.ict-forward.eu/workshop/`.

[2] P. Boutin. Slammed! an inside view of the worm that crashed the internet in 15 minutes. *Wired!*, 11(07), July 2003.

[3] T. Claburn. CIA admits cyberattacks blacked out cities. *InformationWeek*, January 2008.

[4] NERC. Technical analysis of the august 14, 2003, blackout: What happened, why, and what did we learn? Technical report, North American Electric Reliability Council, Princeton, New Jersey, July 2004.