March 26, 2012

Richard Bejtlich

Chief Security Officer, Mandiant

Testimony before the U.S.-China Economic and Security Review Commission

Hearing on "Developments in China's Cyber and Nuclear Capabilities"

Mr. Chairman, members of the Committee, thank you for the opportunity to contribute to today's hearing. I am Chief Security Officer at Mandiant, a private company that provides software and services to detect and respond to digital intrusions. My testimony draws on our company's experience, as well as four years defending General Electric as Director of Incident Response. I have defended Western interests against serious intruders since 1998 when I worked as an analyst and intelligence officer at the Air Force Computer Emergency Response Team, the Air Force Information Warfare Center, and the Air Intelligence Agency.

Because my most recent experience relies on work done in the private sector and enterprise customers, I am not able to provide first-hand answers to questions concerning China's military, security services, criminal groups, or other parties. Your recently released reported titled "Occupying the Information High Ground" is a better source of information on specific, named organizations within China, such as the People's Liberation Army's Third and Fourth Departments of the General Staff Department.

However, I can comment on the characteristics of the groups that the Mandiant Intelligence Team has identified as Advanced Persistent Threat, or APT, actors. For the most part, our team and I use the strict definition of APT as created by the Air Force in 2006, namely as an unclassified reference to intrusions sets ultimately traced back to actors in China. Members of our team have extensive knowledge of these actors that includes time at Mandiant and other organizations focused on the threat from the Asia-Pacific region. Mandiant's assessment of APT actors is not based on any single aspect of an intrusion, such as an IP address owned by a Chinese registrant, or the presence of Chinese language characters in malicious tools or other code. Rather, Mandiant dynamically tracks, over time and subject to continuous modification and refinement, APT groups using a variety of indicators of compromise.

Our intelligence team currently tracks approximately twenty distinct APT groups. These groups include all of the parties identified by reports publicly released by other security companies, as well as actors that we believe are unknown to many of those other companies. We have seen these groups demonstrate various levels of technical and organization skill, with approximately a quarter having "high" skills, one quarter having "medium" skills, one quarter having "low"

skill, and one quarter too new to make a characterization. Within APT groups we tend to see evidence of "crews," meaning smaller teams who specialize in various stages of a compromise. For example, one crew may be tasked with obtaining access to the victim; a second crew moves laterally through the organization to gather intellectual property or other data; and a third crew steals or exfiltrates the data.

Most of the APT groups we track target the US defense industrial base (DIB). Some of these groups also target US government agencies, think tanks and political organizations, and other commercial or private targets. Our most recent M-Trends research report described our consulting caseload for 2011 in these terms:

- Communications companies: 23%
- Aerospace and defense: 18%
- Computer hardware and software: 14%
- Energy or Oil and Gas: 10%
- Electronics: 10%
- Other, of which the financial sector was the largest component: 25%

The following case studies illustrate the trends we have seen in computer intrusions linked to China. The first case describes APT actors assembling the intellectual property they need to replicate a complete product. The second case describes APT actors present during merger and acquisition activities.

In early 2011, an electronics component manufacturer contacted Mandiant as the result of receiving a notification of compromise from a government agency. After conducting sweeps to obtain forensic evidence, we realized that the attacker had been replacing their malware every six months during the two years they had been resident at the victim organization — and this replacement occurred again during the course of our investigation.

To maintain persistence, the attacker used a variety of backdoors, including some publically available ones. One interesting custom backdoor consisted of a custom miniport driver, which listened for a particular "magic packet" that, when received, would activate the malware. Of the approximately 100 compromised systems at this customer, the intruder installed malware on less than half of them. For access to the other systems, the intruder relied on usernames and passwords stolen from the organization.

Mandiant consultants were able to forensically recover a partial listing of stolen intellectual property. The victim company did not place a high value on the stolen data since it was merely a sub-component of a more advanced technology, and the victim did not even produce the other component parts. While the more advanced product was extremely valuable, it could

only be built by combining the victim's technology with parts from a second company in the supply chain. Within weeks, however, the second company called Mandiant. They had also been the victim of an advanced attack, and they also lost intellectual property for a sub-component. It was only by connecting the dots between the two victims that the attacker's goal was clear: rather than targeting a single company for a particular technology, they had been tasked to acquire the more advanced, broader technology. The attackers had performed reconnaissance to determine what companies produced the component technologies, and then targeted those entities to steal what they needed.

Later in 2011, a large European defense contractor contacted Mandiant just months after acquiring a specialty service provider. The service provider had received information indicating that they had been the victim of a targeted attack, and the parent company was concerned about the extent of the penetration.

The attack began with a phishing email containing a malicious PDF attachment. Prior to sending the email, the attacker had performed enough reconnaissance to uncover the name of an individual at a competing organization with whom the victim user had previously corresponded. The socially engineered email purported to be from that individual. When the victim opened the malicious attachment, an intruder established a foothold in the environment. The attacker leveraged this initial backdoor to move laterally throughout the environment, obtained legitimate credentials, and ultimately stole over 50,000 files.

Based on the lessons learned from this incident, the parent company implemented a process requiring every new acquisition to be vetted by the Mandiant Intelligent Response tool prior to being allowed to join the corporate network. This process paid off in late 2011 when the company discovered an APT group actively operating at another company they were about to acquire. The integration was put on hold until a thorough remediation and damage assessment was completed.

Through these sorts of cases, Mandiant extracted several other statistics which describe trends seen in computer intrusions attributed to APT groups.

- 94% of victims learn of compromise via third parties; only 6% discover intrusions independently. Victim organizations do not possess the tools, processes, staff, or mindset necessary to detect and respond to advanced intruders.
- The median number of days that elapse between compromise of a victim organization and detection or Mandiant involvement is 416 days. Incredibly, this number is an improvement over past intruder "dwell time" measurements of two to three years.
- Advanced intruders installed malware on 54% of systems compromised during an incident. They did not use malware to access the other 46% of systems compromised

during an incident, meaning relying on tools that find malicious software misses about half of all victim computers.

- Mandiant observed intruders using stolen credentials in 100% of the cases it worked in 2011. Intruders seek to use legitimate credentials and access methods as soon as possible, because they can then "blend in" with normal user activity.

APT groups use the level of sophistication required to achieve their objective. For example, in 2011 Mandiant observed an increase in the usage of publicly available malicious tools by APT actors. We assess that the adversary uses publicly available tools for three reasons:

1. They already exist, so the intruder does not need to expend research and development resources to create custom tools.
2. Many organizations allow internal use of the sorts of tools favored by intruders.
3. Publicly available tools rarely stand out against the "noise" created by lower-level intruders pursuing smash-and-grab or "botnet" intrusions.

The use of public tools or leveraging publicly known vulnerabilities is a source of confusion for many security professionals. They assume the "advanced" element of the APT term requires that Chinese actors deploy the most sophisticated digital weapons for all phases of an intrusion. I have personally observed APT actors escalating their technical sophistication to adapt to countermeasures deployed by computer incident response teams, so I know the APT can be as advanced as needed when the target warrants it.

I prefer to emphasize the advanced nature of Chinese intrusion management skills when explaining the sophistication of APT groups. It is significant that the most well-resourced, highly professional, and motivated network defenders on the planet have not yet "solved" the problem of Chinese intrusion activity. At best we can keep them from stealing the bulk of an organization's crown jewels, but only after significant investment in improved technology, business and IT processes, partnerships, and staffing.

Mandiant is not aware of specific attacks against an organization's supply chain or cloud infrastructure in order to steal intellectual property, beyond what has been publicly mentioned in the press. Attacks against the supply chain, when manifested as malicious code in trusted hardware or software, can sometimes be discovered by end user organizations. Local security staff can identify the malicious code by the action it takes on the network, or by the way an adversary interacts with it. It is difficult for end user organizations, and any consultants they hire, to gain visibility and awareness concerning compromise of cloud platforms. In general, do not expect cloud providers to be able to identify adversary activity, because it is difficult for the cloud provider to differentiate between legitimate and illegitimate access and use.

APT groups continue to focus on enterprise Windows computers, although other operating systems have been compromised. Intruders exploit enterprise systems hosted in company-owned data centers, and enterprise systems hosted at third party data centers. For the most part, mobile devices, true "cloud infrastructure," and tablet computers do not yet appear to have been targeted.

Concerning legislative or administrative actions that the U.S. Congress can take, I have two recommendations. First, I believe far too much legislative and regulatory attention is paid to compliance with standards and the question of "are we vulnerable?" In my professional opinion, compliance with standards is, at best, effective at stopping some lower-skilled intruders who opportunistically exploit targets. Compliance regimes tend to devolve into a paperwork exercise based on subjective interpretations and the whims of an auditor.

Regarding the question of "are we vulnerable," the answer for every organization is "yes." Rather than wasting time on this question, organizations should instead ask themselves "are we compromised?" In other words, does the organization suffer an ongoing intrusion by a targeted intruder, whether from China, Russia or a criminal group? It is a waste of time and resources seeking compliance with standards while intruders are actively stealing data from a victim organization. The adversary will adapt to any countermeasures deployed during the compliance exercise; I have seen this pattern repeated regularly during my career.

To this end, I recommend Congress consider the integration of an "are you compromised" assessment into any new requirements levied on specific industries. These assessments should occur no less frequently than once per year, although true continuous assessment on a 30-day cycle is much more effective in my professional judgement and experience. By requiring processes and technology to answer the "are you compromised" question, regulators, Congress, and other appropriate parties will, for the first time, gather ground-truth knowledge on the state of security in selected industries. Without knowing the real "score of the game," it is unreasonable to expect real progress in digital security.

My second recommendation involves sharing threat intelligence. I offer a few principles based on my experience as someone who has created, consumed, and shared threat intelligence in a variety of public and private roles.

1. First, adopt an open standard for exchanging technical data. Mandiant created the Open Indicator of Compromise, or OpenIOC format (http://www.openioc.com) for this very purpose. It allows fine-grained description of threat intelligence for use by analysts and software and is free of charge with an open specification available online.

2. Second, recognize that dozens of effective threat intelligence sharing organizations already exist. These include the Defense Industrial Base Collaborative Information Sharing Environment (DCISE), the Bay Area CISO Council, the Financial Services Information Sharing and Analysis Center (FS-ISAC), as well as other ISACs, and other groups. Understanding and coordinating efforts among these groups is a good precursor to any additional sharing activity.
3. Third, please note that intelligence sharing networks do not necessarily improve as additional members join. Having participated in these networks, I have seen a tendency for participants to guard their contributions as the network adds those for whom trust cannot be established on an interpersonal basis. Intelligence sharing relies on trust in order to succeed, and trust is built on personal relationships.

Thank you again for the opportunity to testify. I welcome your questions and comments.