Managing the Targeted Cyber Threat

Maarten Van Horenbeeck
E-mail: maarten@daemon.be

Abstract

This paper illustrates the gap between the handling of cyber threats as generic, versus real-life threats as targeted. This difference, which does not provide sufficient coverage on the notion of the threat agent in cyber attacks, has direct roots in the consideration of the 'class break' as an information security concept. It shows why, despite the importance of this concept, it should not be read to entail a lesser importance of the threat agent. It provides a case study in the form of an information warfare instance between the Falun Gong community and a yet unknown group of attackers. Finally, it looks at the effectiveness of current countermeasures against these types of attacks.

Keywords*: targeted attack; information warfare; trojans; cyber threats; espionage*

Shifting our threat landscape

When making a significant investment, such as the building of a house, or the purchasing of a car, various environmental factors assist in our decision making process. If we live in the city, we may choose to purchase this new home in one of the better areas of the city, or we may elect to park our brand new car in a closed and monitored garage as opposed to leaving it on the streets. The same factors are also taken into consideration when living in a suburban household, but are given much less weight. We're less worried about terrorists blowing up our house than we are of school children accidentally aiming for our front windows during a leisurely game of soccer. This is all common sense, and this type of active risk management allows us to function in an ever challenging world.

When discussing information security, this concept however is oft forgotten. There are good reasons for this. In 2003, Bruce Schneier introduced the concept of the so-called "class break" (Schneier, 2003). This describes that in general, attacks on information security, so called *exploits*, can be reused extensively without significantly increasing the cost of the attack. Stealing a car takes physical effort, and thus cannot be repeated infinitely, so targets are selected carefully. A so-called "drive by exploit", in which a web site infects its own visitors with malicious code, is implemented once, and can infect the machines of hundreds of users within a very short timeframe, while not significantly increasing the work effort. If an exploit is found in a specific piece of software, it can be used against thousands of servers simultaneously without increasing workload.

Schneier's concept has proven very accurate and important: the noise level of attacks on the internet and other shared, public networks, is much higher than the noise level of attacks taking place in our physical reality. As such, most work in information security has been aimed at lowering the noise floor, and stopping these generic, well understood attacks. Many of the security standards that have been developed provide a baseline of security requirements to address specifically this issue. However, they do not evolve in a sufficiently quick manner

to stay updated with all threats, especially those which are tuned to one specific target. Few of them also proactively embed components involved in tracking or responding to these issues.

One standard which does clearly indicate the threat agent as being the root of any attack on information security is ISO 15408, which describes evaluation criteria for information security. This standard describes the overall security context in which threat agents give rise to threats, which exploit vulnerabilities and increase risk to the assets or protection targets (ISO, 1999).

The issue of targeted attacks

A different market economy

Targeted attacks are, as the term itself describes, attacks which are not aimed at compromising a maximum of hosts, but at compromising machines of specific interest. These could be database servers hosting content unique to the organization, but it could also include desktops of people with a specific role in the organization. Both give access to types of data that would normally not be available to an outsider.

These attacks are problematic in a sense, as they disregard the usual economies of internet crime. Most internet based crime is focused on the fact that a vulnerability is relatively easy to exploit, and has a high degree of success in compromising large amounts of systems. These systems are then joined in for example a botnet, a collection of hosts which can be used by the attacker to perform various other nefarious acts, including the transmission of spam. A vulnerability which is hard to exploit in a piece of application software which is not very popular will generally not be exploited due to the lack of a convincing business case.

In the case of a targeted attack, a specific set of information, with a fixed price tag, may be the target. As long as the effort to exploit a series of vulnerabilities is budgeted below that price tag, and still allows for a decent margin, it will still be valuable to the attacker to exploit the vulnerability and attack a machine.

This essentially changes risk management, as for an organization to present its data against these attackers we should be aware of both the cost to exploit known and unknown vulnerabilities in the software we use, as well as understand the exact value of our assets on the black market. There is offcourse no cost-effective approach to achieving this state of information superiority above an attacker, who only needs to find a single vulnerability and a way to exploit it..

Targeting methodology

An attacker who wishes to gain access to specific information, will need to first identify the exact targets. It is highly unlikely that the compromise of a corporate web site would reveal significant data of interest, as these systems are usually segregated from the actual corporate network.

In most cases, targeting data consists of targeting the individuals associated with it:

- *Identifying the actors involved in processing information of value*: this may include personnel that would normally handle sensitive data, such as research & development engineers.

- *Mapping logical communities those actors frequent*: this may include public mailing lists or web site, which could all be used as attack vectors later on;

- *Identifying the actor's reputation within the community*, its interests and contact information;

- *Design of a malicious code specimen* to undertake an action such as searching and stealing for specific data on a system, or granting the attacker remote access;

- *Providing this malicious code* sample to the actor in a way that it would be expected for him to open it.

Attack methodology can vary widely within these constraints: malicious code can be delivered by e-mail, but it could just as well be provisioned by a web site the user knows and trusts. The attack also does not always need to take place directly: instead of attacking an individual directly responsible for the data item, a correspondent of him may be attacked, using his online identity (such as an e-mail address) to provide malicious code to the final target.

Targeted e-mail attacks against Falun Gong

Falun Gong is a by origin Chinese spiritual movement consisting of meditation exercises and spiritual teachings, which was founded by Li Hongzhi in 1992. Both the exercises and the movements principles are described in two works by its founder, 'Zhuan Falun' and 'China Falun gong'.

Some debate has taken place on whether or not Falun Gong is to be called a religion. Its founder Li Hongzhi claim it not to be one, and Falun Gong practitioners as such can be a member of any other religion. Practitioners can join out of interest in the physical exercises, only deciding later to consider some of the underlying spiritual ideas.

Repression by the Chinese government

The Falun Gong movement was directly involved in one of the largest public demonstrations in contemporary China. This started with a magazine article, published in the Science and Technology for Youth magazine, which criticized Falun Gong and Qigong in general as being of little benefit to youth – suggesting they should pursue more athletic sports. Practitioners of the movement staged a protest rally at the magazine publisher's office, resulting in several subsequent arrests.

As the Falun Gong practitioners did not feel their needs were adequately addressed by this government approach, a new protest rally took place in Beijing on April 25th, 1999 at the Zhongnanhai, the Chinese government headquarters. While the demonstration was widely reported to have been peaceful in nature, consisting of "standing and meditation"

(Immigration and Refugee Board of Canada, 1999) and resulted in the release of the previously arrested practitioners, it was likely to have been perceived as disruptive by the Chinese leadership. Wide press coverage of the demonstration also increased the movement's visibility outside mainland China.

These events led to the Chinese government officially banning Falun Gong on July 22nd, 1999. In addition to banning the organization itself, any attempts to advertise, defend or promote Falun Gong or spread rumors or 'distorted facts' concerning the movement (BBC, 1999) were prohibited. This in effect, made it impossible to discuss or challenge the decision. Simultaneously, a 'wanted circular' was spread on founder Li Hongzhi, for spreading superstition and fallacies, organizing demonstrations and perhaps dominantly, disturbing public order.

Information war

Ever since its repression by the People's Republic of China, and partially out of necessity due to the ban on public discussion of the group within China, the battle between PRC and Falun Gong has become a very public international one – an example of far-reaching information operations.

Information Operations is defined by the US Department of Defense as 'The integrated employment of the core capabilities of electronic warfare [EW], computer network operations [CNO], psychological operations [PSYOP], military deception, and operations security [OPSEC], with specified supporting and related capabilities to influence, disrupt, corrupt, or usurp adversarial human and automated decision making while protecting our own.' (DoD, 2003). This definition has a very military connotation, but can easily be broken down to its components in order to clarify it from a more public perspective.

Information Operations can be considered the employment of electronic, psychological and operations security to gain competitive advantage over an adversary. In this specific case, the Falun Gong is attempting to bring its message across and allow people to partake in its movement without the risk of prosecution. China, on the other hand, wants to guarantee stability by eliminating the movement from the public picture.

One interesting question to ask is whether China has a similar view of this concept as the US military does. The US Department of Defense is one of the few organizations to have publicly studied the Chinese interpretation of Information Operations. Timothy L. Thomas, from the Foreign Military Studies Office at Fort Leavenworth, one of their subject matter experts, has published a number of books that review Chinese state of the art in information warfare.

One of these books, 'Dragon Bytes', deals with the change the Information Warfare concept has undergone since its inception in China. He discusses how as of the first Gulf war, Chinese strategists have been working on integrating the concept of Information war in Chinese military strategy. This started with a thorough study and discussion of the US methodology and framework, to be followed in 1997-1998 by a specific Chinese approach to the field.

Thomas quotes the well-known Chinese IW strategist Dr Shen Weiguang, who as of 1996 has described information warfare to be linked to control. Controlling the flows of information between all parties involved is considered of prime importance in gaining dominance. It may not be necessary to have direct decision power over a country or province if one can manipulate the information it is receiving and as such have it make the decisions one wants it to make through deception or selective information distribution.

In addition, his book attempts to highlight some of the differences in strategic thinking. One chapter explains the way stratagems influence strategic decision-making. Stratagems are small, practical tools of deception that can be used in an information warfare context. There are in total 36 stratagems, in popular press often referred to as the 36 strategies.

Reviewing these in comparison to US documentation on deception shows a major difference in how these countries formalize their deception techniques: the US has significant guidelines and procedures on how to implement deception with less focus on the actual strategies, while China focuses on integrating ancient strategic thinking, making available less public information on the procedures surrounding them. This collaborates with the description of China as a high context culture assigned to it by many authors (Donghoon, Yigang & Heung, 1998).

Operationalizing the attack stratagems

In his work, Thomas makes an interesting inference. He links a claim by Dr Shen that "people have come up with 36 ways to disrupt the Internet and 36 ways to defend against such disruption" to the actual stratagems, and poses the question whether China may have translated these into the information age (Thomas, 2001).

It does not require much theorization to assess the result of material implementation of the stratagems. The first of the 36 stratagems states that sometimes one should "fool the emperor to cross the sea" (AFPC, 2007). This entails that one should perform ordinary activities and blend in with normal events in order to lower our enemy's guard.

Translated to the e-mail attack vector this could consist of creating accurate context in which the e-mail is interpreted. E-mail content, attachment and sender should be in line with an image that looks trustworthy to the recipient.

The Falun Gong Information war

In 2001, Arquilla and Ronfeldt described the concept of netwar as "an emerging mode of conflict (and crime) at societal levels, short of traditional military warfare, in which the protagonists use network forms of organization and related doctrines, strategies, and technologies attuned to the information age". Netwar as such is not limited to the use of information technology, but deals with the use of information in new forms of organization less typical than standard hierarchies.

In order to correctly view the wide nature of Netwar, we have to be aware of the 'network' existing between different parties in a conflict. Within such network, parties will work

together – however there is little to no hierarchy. Arquilla and Ronfeldt describe in their work *The Advent of Netwar* that such network would simultaneously be "acephalous and polycephalous". Some 'nodes' within the network may have a particular task, while others can be used for generic interventions. Networks may have no heads, or multiple heads, in certain situations.

This type of conflict also applies to the information war between Falun Gong and the People's Republic of China. While Falun Gong clearly has a leader, Li Hongzhi, many of its proponents are not directly related to that leader in a form of hierarchy. Specific groups within the community have an unofficial messaging role, such as the web sites Clearwisdom.net and Clearharmony.net. The newspaper Epoch Times, with a popular web site and print editions in numerous countries also serves as a significant public relations organ for the movement. In addition, a worldwide network of human rights organizations are loosely affiliated with Falun Gong in the way they fight for its acceptance by the PRC.

China on the other hand has control over a large, more hierarchic public relations and diplomatic affairs organization. However, there is also some indication they may have indirect control over a significant network of citizens outside of the official PRC infrastructure. After the May 8th 1999 bombing of the Chinese embassy in Yugoslavia by NATO, a minor defacement war erupted between Chinese and US based hackers, clearly for patriotic reasons..

Conceptually, this could be interpreted as a modern version of the People's War strategy initially devised by Mao Zedong. By not outstretching the army on the borders and involving the public to join in a form of mobile or guerrilla warfare, information technology and talented young individuals can be used as a force multiplier. While not necessarily purposely incited by the Chinese leadership, strict control over media outlets is likely to have an influence in shaping the inputs of opinion, and any resulting action taken by the population. Combined with the extensive and patriotic Chinese hacker culture, it seems likely these groups are inspired by some of their countries' decisions and may wish to contribute.

One large technical escalation in this information war took place in 2002 and 2003, when Chinese satellite Sinosat-1 was hacked by what were believed to be Falun Gong adherents. The regular China Central Television broadcasts at that time were interrupted by Falun Gong related transmissions (China Embassy, 2002). It is noteworthy however, that this hack was never publicly confirmed by Falun Gong members and could have been part of a propaganda campaign on either site. This would have been valuable due to its broadcast time, interrupting the widely anticipated world cup in mostly rural areas that would be of little benefit strategically to Falun Gong. It only showed relative weakness on part of China, but did aggravate the population and likely decreased acceptance of an organization they would otherwise look upon with relative neutrality.

Attack methodology

As part of the research for this chapter, the author reviewed several attacks aimed at Falun Gong members that took place between April and December of 2007. During this time, it became clear that the bulk of the attacks consisted of the use of malicious e-mail attachments being sent to members of the organization:

In general, members or people familiar with the organization received an e-mail, which was most often spoofed from an address they recognized and contained an attachment to the e-mail. This attachment was maliciously crafted to exploit a vulnerability in application software running on the target's machine. Upon successful exploitation, the file dropped code on the system which then set up a connection to a server abroad, most often located in Taiwan or Mainland China. This server was then able to send commands to the newly compromised machine.

In addition to the samples investigated throughout this year, several other samples were provided to the researcher which indicated these attacks had been ongoing since at least early 2005. Reports on Falun Gong web sites of these attacks dated back even further. Clearwisdom.net, a popular Falun Gong discussion site, published an article in 2003 dealing with 'computer security setup and usage'. The report mentioned recent computer virus infections and hacker attacks.

Interestingly enough, the very same methodology had already been widely reported in a bulletin by a large number of government organizations on critical infrastructure protection. The type of targeted trojan attacks reported in the respective bulletins by the UK Center for the Protection of the National Infrastructure (CPNI), US Computer Emergency Response Team (US-CERT) and Australia's Defence Signals Directorate (DSD) closely matched the real-life experience of Falun Gong members, to the degree that the backdoor Trojans were of the same family..

While this is based on several short-term observations of honeypots, and as such does not constitute conclusive evidence, several data items hosted on the member’s machines appeared to be specifically targeted:

- E-mail contents and contacts;
- Locally stored Word documents.

In at least one case, which was not directed against Falun Gong members, but against a pro-Tibetan non profit which was contacted while investigating a Falun Gong attack, the code specifically copied encryption keys used to encrypt communications between members of the organization (Van Horenbeeck, 2008).

Social Engineering

Social engineering is an oft-used but ill defined term which is as broad as affecting the decision making of an individual by introducing crafted and customized acts to affect the individual's perception. In these specific attacks, several strong techniques were used to make the content submitted by e-mail acceptable to the reader, and encourage him to click on the file and activate the content:

- In some incidents, the writing style of the sender appears to have been investigated prior to sending the message, and very accurately mimicked. This includes the use of hypocoristic names;

- Introduction of cognitive dissonance: an e-mail made a strong statement on an individual, such as the Chinese President, but do not mention his name. The attachment is then named ater the individual. A state of “cognitive dissonance” arises between the reader's pre-existent beliefs and the statement. As such the reader is more likely to click the message to put his beliefs in line with the statement;

- Legitimate, trusted users are often convinced to forward the message to other members of the community, or messages are spoofed as “having been forwarded” by a person they trust. When the target receives the message, he considers it 'previously validated' and is less stringent in assessing its contents;

- In a number of cases, so-called *memes* have been reused. These are messages, presentations or images that have been distributed within the community in a viral fashion. The original document is taken, backdoored, and then forwarded on to members, who are no longer able to distinguish between the original viral meme, and the backdoored version.

Exploit selection

In order to successfully compromise a target machine, code will need to be used which exploits a software vulnerability in an application running on the target machine. In addition, the application must automatically open this specific type of file once it has been clicked from the e-mail client.

In addition, when security vulnerabilities are identified in application, they are generally fixed by the application's vendor, and an upgrade distributed to the users. However, this process is not standardized and depending on its implementation, can be automated or manual.

This means that:

- The likelihood of the application under attack being used by the target increases the likelihood of the cyber attack in being successful;

- Applications that are not automatically patched by the vendor are more useful targets than those applications which are patched either automatically or on a regular basis.

Vulnerabilities in applications are universally identified through a CVE number. CVE stands for Common Vulnerabilities and Exposures, and is a dictionary of significant software vulnerabilities. In total, 7 different CVE listed vulnerabilities were abused in the attacks on Falun Gong dating from April to December 2007.

The attacks on Falun Gong members employed predominantly a common insecure configuration setting in the target operating system, in addition to exploits in a common Office tool, as well as an Archiving application. The Office tool was clearly selected due to its widespread use. While an automated update mechanism does exist it was not in all cases enabled. It should be noted though that Office applications, as each version adds

significant new functionality, are actually fairly likely to be patched, especially in enterprise environments.

The archiver application was in far less widespread use, but did not contain functionality to automatically check whether an updated version was available. In addition, archivers, which have the ability to pack multiple files in a single container file, are low level tools, and unless the file format significantly changes, users will not be inclined to upgrade, despite the availability of a new version.

Deep compromise

One of the unique aspects of these attacks is how the connection to the control infrastructure is maintained. As there is generally no anti virus coverage to detect the threat up-front, the only corporate wide means of detection is to identify the control connection. While the control connection can be set up to an internet address directly, there is an alternative methodology which was observed during these attacks. It entails the use of a host name registered through a dynamic DNS provider.

Dynamic DNS providers are organizations which allow an individual to quickly set up a 'host name' pointing to an IP address of their choice. A user, for example, has the ability to point *name.dynamicdnsprovider.com* to the IP address of his home computer..

The advantage of this is that even if the user's IP address changes, he only needs to adjust the name's configuration, and from another location, users will still be able to connect to the machine without needing to provide them with the IP address. This type of tool was originally developed to allow home users, whom in some cases frequently need to change IP address, to connect to their machines from any other location, such as a cyber cafe. By running an automated tool on their home machine, they could automatically update their host name when the address had changed.

This mechanism is however also appeared to be widely used in a malicious manner as part of these attacks. In this case, an attacker compromises a random machine on the internet and designates it his control server for future attacks. He then anonymously registers a dynamic DNS name, and points it to the control server.

An example from an actual attack is john0604.3322.org. 3322.org, operated by the Chinese company Bendium, is a provider of dynamic DNS services to individuals. In this instance, the attacker had compromised a server at the IP address 63.161.44.241, and pointed john0604.3322.org to this address. He then delivered malicious code to Falun Gong members by means of a social engineered e-mail with a document attachment. The code, once executed set up a connection to john0604.3322.org and fetched commands to be executed on the Falun Gong member's machine.

There are two advantages to using this technique:

- Once the attack is identified, law enforcement may attempt to have the control server shut down. It is much easier to contact the service provider responsible for this IP address (which in this instance was based in the United States) and request them to shut down the host than it is to request the same to a dynamic DNS

provider. Essentially, the host name is not undertaking malicious action, while the server is. However, the attacker can easily re-enable the host name to point to a different physical control server.

- Due to their purpose, allowing people to stay in touch with their home machines once their home IP address changes, dynamic DNS providers apply a very low Time To Live value on their hostnames. This value describes how long someone using the name will remember that it points to the same address. As such, once the server is taken down, an attacker can re-enable the attack in minutes instead of hours, as with usual DNS entries.

In addition, this ability to change a host name entry at short notice has other uses in targeted attacks. Detection of the backdoor connection on the network can take place by identifying the name lookup, or the actual connection which ensues. Due to the massive amount of name lookups emanating from the average network, organizations do not usually log these queries. However, connections visible in firewall logs on the organization's perimeter are usually stored and maintained. As such, a large amount of connections to an IP address based in Taiwan or China could arouse suspicion.

In order to limit traffic to the backdoor to only those timeframes when intelligence collection is required, attackers have previously changed the control server IP address to 127.0.0.1. This IP address is assigned to the loopback interface of the system, and prevents any traffic from leaving the machine.

Below is an example hostname which was tracked during this research and which displays this behavior. Every line preceded by + means that at that time, the hostname started to resolve to the listed IP address. A line preceded by – indicates the hostname stopped resolving to that specific IP address atthat time.

```
+ 2008-01-07 15:28 | ihe1979.3322.org | 127.0.0.1
+ 2008-01-11 14:46 | ihe1979.3322.org | 117.9.209.247
- 2008-01-11 14:46 | ihe1979.3322.org | 127.0.0.1
```

In this case, the attack, which had been deactivated at the end of 2007, has been reactived on January 11th at 14:46 GMT using a new control server IP address.

Based on similar findings, several Intrusion Detection rules have been published which track DNS lookups that resolve to 127.0.0.1. These are immediately suspicious, as there are few reasons why a legitimate host name would resolve to a this loopback IP address.

The technique has however been brought to the next level. Several samples now contain a predefined “parking” address. Once a hostname resolves to this perfectly legitimate parking address, the sample will take no further action and attempt to resolve this same address again after the Time To Live value expires.

Below are DNS traces of an attack involving such address:

```
+ 2008-01-26 13:49 | ding.pc-officer.com | 63.64.63.64
- 2008-01-26 13:49 | ding.pc-officer.com | 61.219.152.125
+ 2008-01-27 00:28 | ding.pc-officer.com | 61.219.152.125
- 2008-01-27 00:28 | ding.pc-officer.com | 63.64.63.64
```

```
+ 2008-01-27 02:45 | ding.pc-officer.com | 63.64.63.64
- 2008-01-27 02:45 | ding.pc-officer.com | 61.219.152.125
+ 2008-01-27 08:21 | ding.pc-officer.com | 61.219.152.125
- 2008-01-27 08:21 | ding.pc-officer.com | 63.64.63.64
+ 2008-01-27 08:21 | ding.pc-officer.com | 61.219.152.125
```

In the above example, access to the compromised machines is disabled on January 26th, 2008 from 13:49 GMT, until it is enabled again at 00:28 GMT. Using this mechanism, the attacker has control over when the channel is enabled or disabled. The malicious code samples distributed as part of this specific attack sequence contained code to check whether 63.64.63.64 was returned in the DNS lookup. This is a perfectly valid IP address, and despite its strange constituency would not arouse suspicion on behalf of any intrusion detection tool. However, if the tool sees it, it temporarily stops the connection attempts.

This technique enables *deep compromise*: an attacker can compromise systems on a network, and only activate the visible portion of the backdoor after several weeks or even months. This allows the targeted attack to remain persistent and dormant within an organization to avoid detection. If only a single sample was sent, which bypassed all controls at that time, after-the-fact detection is unlikely.

Threat agent and infrastructure analysis

The control infrastructure to handle an attack spree is often quite complicated. The below graph provides an indication of how a single control server was addressed by multiple malware samples distributed during 2007.

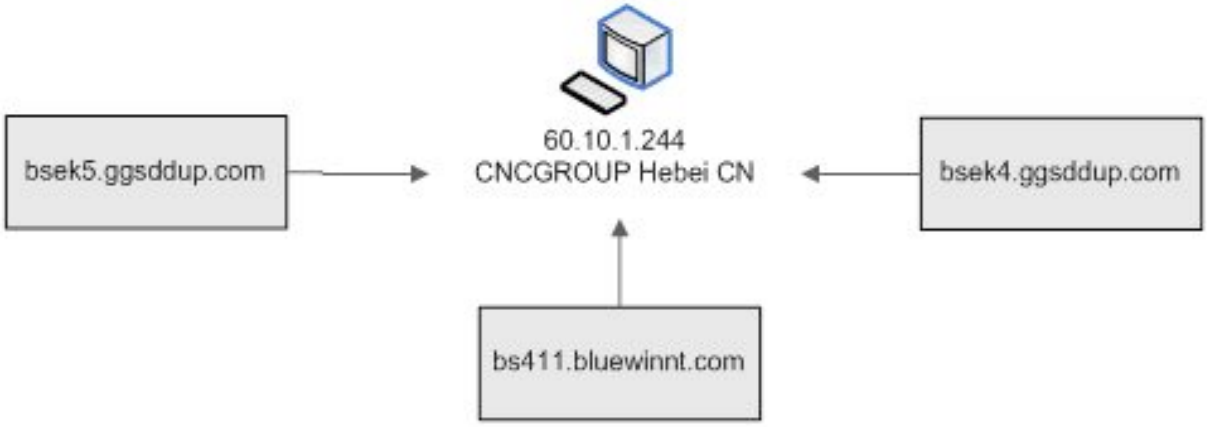

In addition, a unique hostname may point to multiple IP addresses at different stages of the attack spree. This most often interrelates with a machine being reinstalled by the ISP, which causes the attacker to lose his control connection:

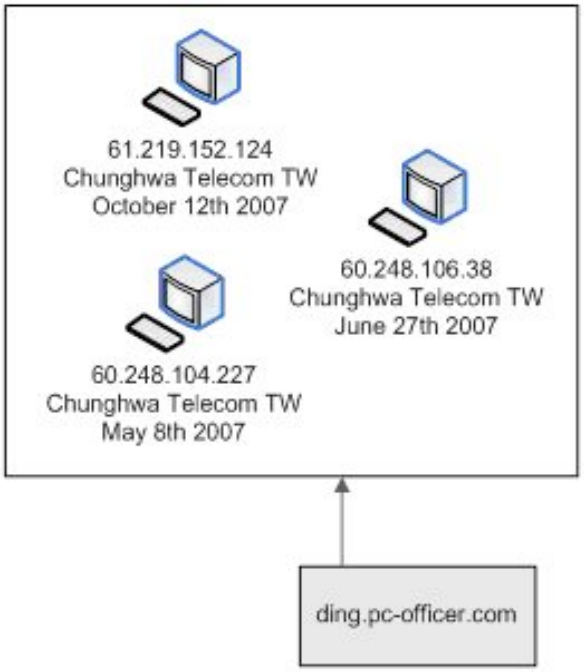

In addition, a single malicious code sample may attempt to connect to multiple hosts to sustain its control connection, should one become unreachable:

The following graph, while not intended to be fully readable within the space confined in this book, provides an appreciation of the complexity of some of the control architecture involved in a relatively simple attack of only 21 malicious code samples.

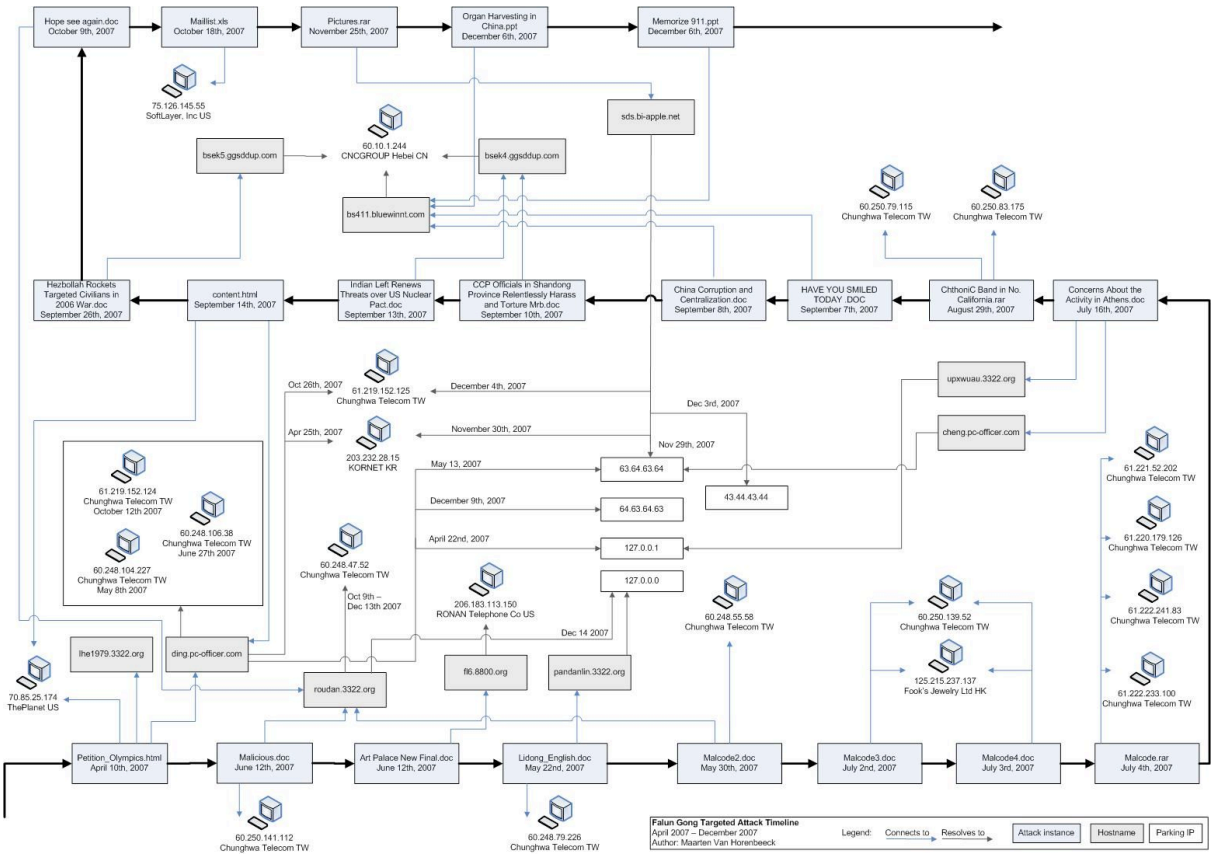

It should be noted that technical evidence does not provide sufficient data to identify the specific attackers involved in this attack spree. As part of this research, several owners of control servers were contacted. From their feedback it became possible to deduce two significant types of control servers:

- Servers which have been compromised and are running an application which terminates the connections by the attacked clients. The hacker logs into the server directly to manage the clients under his control;

- Servers which have been compromised and are running a “port forwarder”, being a tool which merely forwards all incoming connections from attacked clients to another machine, often located in yet another jurisdiction.

It should be noted that technical evidence only allows the data to be traced back to an IP address. Identifying the owner of that address in another country, or even further, the person behind the keyboard, is challenging at best. Without comprehensive law

enforcement cooperation, criminal incidents related to targeted attacks are notoriously difficult to solve.

This becomes increasingly more difficult when there is some uncertainty whether the attacks might be considered acceptable by the government in whose jurisdiction the control server is located. In all, sharing evidence with a foreign government reveals significant data on detection capabilities (Henderson, 2007).

Countermeasures and controls

Technical controls

Targeted attacks are only restricted by the cost an attacker wishes to expend, given the incentive he has to successfully compromise the target. However, in most cases they involve the introduction of malicious code into the target network. This can take many forms: it could consist of the execution of so-called *shellcode* on the target system, which executes a malicious action, such as granting a remote attacker full access to a system. On the other hand, it could simply download more advanced malicious software such as a backdoor trojan. The majority of preventative technical controls implemented to deal with these attacks are located in the *anti-malware* product area.

Within this area, three types of solutions are recognized and widely deployed: black listing, white listing and behavioural detection.

White listing consists of generating a list of all objects that are permitted within the context of a system. This could include a list of permitted web sites, binaries or e-mail senders. Essentially, several components within the network or system enforce that only those white listed objects are accessed. Any violations are blocked and can be logged for the corporate security team to investigate.

In general, white listing has a very good track record in preventing targeted attacks. The attack vector does not need to be known for it to be blocked – as the administrator does not consider it “secure”, it will not be permitted.

However, the devil is in the detail: white listing solutions have a very specific scope, and as such may not always be fully effective. In the case of a web white list, the list may approve a commonly used research portal, but if that web site itself is hacked and contains compromised code to exploit the user's web browser, the white list would not be effective in preventing a compromise.

In addition, some white list functionality may be bypassed depending on the level at which code is introduced into the targeted system. As an example, an executable white list may be implemented which hashes each of the executables on a “clean” system, and only allows these to execute. This may include a network-based service that listens on a specific internet port, but has a yet unknown vulnerability. An attacker may exploit this vulnerability, and if the service runs with sufficient privileges, introduce sufficient code to disable the white listing solution. Much of this is a generic concept, and any such attack would depend very much on the actual implementation of the operating system, the component enforcing the white list, and the actual vulnerability.

Use of the internet or networked technology however provides a great asset to an organization: suddenly it has the ability to interact with and take advantage of an almost unrestricted source of information and potential partners. Limiting the amount of resources available may be acceptable to some roles, such as the organization's reception or help desk, but may not be acceptable for R&D groups that need to evaluate new technologies yet unconsidered by the white list.

A second, more common technology is *black listing*. This consists of the opposite of a white list: generating a list of those objects which should not be permitted, and are known as “unacceptable” or “bad”.

This market space includes the popular signature-based anti virus solutions we are all familiar with. Once a piece of malicious code is identified by such vendor, it is classified, a detection signature is written, and this signature is implemented in the tool's signature database. Machines running the tool regularly update their database, scan the system or any new objects being opened or installed, and block known malicious code samples. Naturally, this type of approach is much less effective than white listing. Samples need to be known before they can be blocked. On the other hand, it does still allow the organization to access any resource – but those that are known to be malicious.

There are a number of issues that make this approach particularly problematic for targeted attacks:

- Targeted attacks will generally use malicious code which is not reused outside a single target, or is at least very limited in distribution. Attack samples will generally not, or in a belated manner, make their way to an anti virus company. As such, specific detection signatures will not be added, causing the malicious code to remain undetected;

- Detection of the sample and generation of a signature always take place after the actual attack. It is essentially a reactive response to the threat. Once the sample has however installed itself on a system, this system is no longer in a “clean” state, and there is no guarantee the object scanner will still function as expected.

A third technology is *behavioural monitoring* for malicious code. This consists of monitoring the behaviour of a sample in action, and triggering an alert when it undertakes specific action known to be malicious. There are however many types of malicious behaviour, and some may only be considered malicious when seen in a specific context. As such, these tools are not generally tuned to trigger on “any” suspicious activity, such as installing a file to automatically run when starting up the system, but apply profiling of executing processes. They maintain a list of all actions undertaken by it, and assign a 'score' or 'risk rating' to each.

Once the sum of this score exceeds a certain threshold, an alert is flagged and the code is considered suspicious. This strongly limits the number of false positives which are inherently involved with flagging behaviour.

However, this type of solution also has specific disadvantages:

- Scores are assigned to types of behaviour because they are known to be related with these types of attacks. A yet unconsidered vector may not have a score attached;

- Despite the intelligent use of “profiling” and “scoring”, this type of solution is more likely to generate false positives. To what degree these can be controlled depends highly on the target environment;

- In order for this type of solution to detect a compromise, the malicious code would first need to execute. During its execution, and before being stopped, it may already make changes to the local system which are not expected or desired by its administrators.

In addition, some solutions are not focused on prevention, but limiting the scope of the attack once it has taken place, or ensuring detection.

A well known approach is the use of automated integrity recovery tools. These essentially make all changes to the system after its boot ephemeral, and ensure they are not stored to disk. Once the system reboots, it is in a clean, uncontaminated state. While this does not prevent the system from becoming compromised, it does prevent the compromise from having an active impact on the organization for much longer than one business day. These tools can prove valuable in those situations where machines are shared among various users, and users are not expected to store content locally.

Human awareness

While technical controls have the ability to shield against specific types of technical attacks, in many cases, they only need to become active because users are not aware of the threat posed. It is important for organizations to assess the risk of a targeted attack on a specific asset, and to ensure users remain vigilant. As one parameter, in some organizations it could be appropriate to involve information security even when an application throws a mere error message upon being received through e-mail.

Targeted attacks generally are highly dependant on social engineering efforts to be successful. In several instances that have been published, mails were distributed to a target, spoofed as originating from a known contact of his, containing valid content which looked very normal and did not give rise to any suspicion on behalf of the reader. However, these attacks are generally preceded by attacks which apply a lower quality of social engineering, and may even contain blatant typographical errors.

Such messages should not merely be ignored by the reader, but should be brought to the attention of the security team for investigation. The outcome of this review will enable the organization to provide further detail to the corporate users on detection of new incidents.

Security Intelligence

The inherent issue with targeted attacks is the lack of a proactive strategy to deal with every possible security incident. As the creativity of the attacker is only limited by his

budget available to obtain access to the information assets, the issue area truly constitutes a situation in which the attacker only has to find “one way in”, while the defender needs to protect against every single attack vector.

While many of the technical controls offer a specific degree of certainty, none provides full coverage against all possible attack vectors. As such, part of the threat requires mitigation through consideration in the security intelligence process.

In it s most basic definition, Intelligence is “information that reduces uncertainty in decision making”. However, at the same time the term is also used for the process used to acquire this information.

The intelligence process is most often described as a cycle of the following four major components:

- *Direction and Planning*: Based on the business requirements of the organization, a decision is made on the questions to be answered by the intelligence process. These may include the questions most of interest to the organization, or a more generic request to follow up on a trend, in which case the question translates into “What risk does and will ... pose”. A collection plan is developed which defines how the questions can be most accurately answered;

- *Collection:* Based on these questions, the collection plan is executed. Sources are contacted or reviewed to gain an understanding of the issue and provide answers to complete the organization's view of a situation;

- *Processing:* The data gathered during the collection phase is normalized and verified. This often entails triangulation: information obtained from a source is checked against several other pieces of information to ensure it is valid, and the source trustworthy. Finally, the current view of the situation, or a pre-developed “model” that explains it, is completed with the acquired information. As a final result, the questions which served as input to the intelligence process are answered;

- *Dissemination:* The output of the intelligence process is rendered readable by the target audience and disseminated accordingly.

Reduced to its essence, security intelligence focuses on the security threats posed to our organization. Its concept closely resonates with Sun Zhu's popular saying: *know thyself, know thy enemy, and you will never be defeated*. The goal is to understand the organization's security weaknesses, and to attempt to understand the strengths of the attacking threat agents.

While security intelligence is not often formally implemented in business, every organization does use components of it to a certain degree. For example, a basic patch management process in many organizations consists of maintaining an asset register of the assets under their control, and contracting vulnerability tracking services from third party service providers, to obtain notification on new vulnerabilities which need to be patched on these assets.

In the case of targeted attacks, the process can add additional significant value. For example, it can be used to co-operate and share information with other organizations

affected by the same threat agent. In many cases, a targeted attack is executed not merely to a single organization, but a small number of organizations active in the same market or issue area. Several defence contractors can be attacked by the same threat agent as they work on similar issues. While the social engineering component of the attack may be different in both cases, the control infrastructure and the code samples, which take time and effort to set up, is often similar or identical.

As an example, attacks on two targets may use a specific type of trojan software, which generates very similar requests to a malicious web site. While the server hosting the web site may be different in both attacks, the type of traffic generated can be recognized.

Within the United States, as well as several other countries, so-called "Information Sharing and Analysis Centers" have been set up to improve information sharing between organizations within the same industry. The FS-ISAC, for one, provides information sharing services between financial services organizations.

In this case, if one organization is attacked, it has the ability to anonymously share information on the attack with the other organizations. This type of information can include specific data such as the methodology the trojan employs to tunnel data back to its control infrastructure. Such data can be put to use by the other organizations to identify similar compromises on their network

Future Concerns

Outside of the relative lack of technical options to deal with these attacks, another matter of concern is that knowledge on a specific topic is no longer contained within a single enterprise, but often reflects a complete "issue area" of organizations working on the same type of content.

As the attackers implementing this methodology have a tremendous opportunity to reuse attack components, even social engineering, they are likely to attack multiple, similar organizations. As such, targeted theft of information is likely not to upset only a single company, but an entire industry. While the effects may not be clearly visible immediately after the compromise, this type of strategic espionage is likely of having a significant impact on the competitive advantage of existing industry players.

From within another non-profit community, indications have also arisen that attackers are targeting "shared resources" of the issue area. In March of 2008, a web site belonging to a UK pro-Tibetan group was compromised, serving malicious code to users visiting the site. This researcher has reviewed the code and found similarities in process to the e-mail borne attacks discussed in this paper.

Interesting is that such sites are less intended to provide information to the constituency of the organization than they are intended to provide information to other actors within the issue area, such as the press and peer organizations. In 2008, a major press agency reported having received an attack very similar in nature when covering themselves cyber attacks against pro-Tibetan groups (AFP, 2008).

Conclusion

This paper showed how targeted attacks differ from more generic internet crime, especially in methodology but also from an economic perspective. It reviews an overall targeted attack methodology and then provides an in-depth view at a targeted attack on a Chinese minority community, which identifies the various mechanisms “in action” and shows how they are tuned to the specific target to ensure the success of the attack.

Defensive concepts on the technical layer only are not adequate to protect against these attacks, mainly as they need to defend against all possible means of attack, while the attacker only needs to find one viable path of entry into the organization. They need to be supported by clear intelligence on novel attack methodologies, and sufficient and continuously engaged awareness training for those people who are at increased risk of targeted attacks.

An interesting element, which is visible from these attacks is the interest in ‘information’ and not merely a specific company. As combining information from various organizations provide a significant information advantage on the market, it is highly likely that these attacks may not purely focus on a single organization but are targeted in fact at members within a certain “issue area” or market.

References

AFPC (2007). Thirty-Six Strategies. Paris, France: Association Francaise des Professeurs de Chinois.

Arquilla, J. and Ronfeldt, D. (2001). The advent of Netwar (Revisited). In Arquilla, J. and Ronfeldt, D. (Ed.) Networks and Netwars, (pp. 1-25). Santa Monica, CA: RAND Institute.

BBC (1999). Text of notice banning Falun Gong. Available at: http://news.bbc.co.uk/2/hi/world/monitoring/400943.stm. London, UK: BBC News.

China Embassy (2002). Press release: Chinese satellite TV hijacked by Falun Gong cult. Washington DC: China Embassy.

DoD (2003). Information Operations Roadmap. Washington DC: Department of Defense.

Henderson, S. (2007). The Dark Visitor. Privately published.

ISO, (1999). ISO 15408: Evaluation Criteria for IT Security. Geneva, Switzerland: International Standards Organization.

Immigration and Refugee Board of Canada (1999). CHN33180.EX Country of Origin Research. Available at: http://www.irb-cisr.gc.ca. Ottawa, Canada: IRB.

Newey, G. (2008) Pro-Tibet groups bombarded with abusive calls, viruses (Press article). Beijing, China: AFP.

Schneier, B. (2003). Beyond fear: Thinking Sensibly About Security in an Uncertain World. New York, NY: Springer.

Thomas, T. L. (2001). The Internet in China: Civilian and Military Uses. Fort Leavenworth: Foreign Military Studies Office.

Van Horenbeeck, M (2008). Guarding the Guardians: a story of PGP key ring theft. Bethesda, MD: Internet Storm Center.