Advanced Data Exfiltration – the way Q would have done it

Going above and beyond DNS and HTTP tunnels…

Iftach Ian Amit – VP Consulting, Security Art, Israel iamit[at]iamit[dot]org.
September 2011

Abstract

Penetration testing and red-team exercises have been running for years using the same methodology and techniques. Nevertheless, modern attacks do not conform to what the industry has been preparing for, and do not utilize the same tools and techniques employed by such tests. This paper discusses the different ways that attacks should be emulated, and focuses mainly on data exfiltration.

The ability to "break into" an organization is fairly well documented and accepted, but the realization that attacks do not stop at the first system compromised did not get to all business owners. As such, this paper describes multiple ways to exfiltrate data that should be used as part of penetration testing in order to better simulate an actual attack, and stress the organization's security measures and detection capabilities of data that leaves it.

Infiltration

Modern attack trees employ multiple realms of information that are not necessarily extensively tested and protected by organizations. Some of these paths involve not only technical factors, but also human, business, and physical ones.

From a technical perspective, even though information security as a practice has been around ever since computer systems have been in use, organizations tend to focus on the issues that are well documented, and have a plethora of products that address them. On the other hand, attackers would do exactly the opposite – try to find infiltration paths that involve some human interaction (as humans are generally less secure than computers), and focus on elements that are less scrutinized by protection and control mechanisms. One example for such an attack path is using formats that are commonly used by organizations, and are known to contain multiple vulnerabilities. Such formats (like WinZIP, Rar, PDF, Flash) and applications (vulnerable Internet Explorer, Firefox, and other commonly used tools) are bound to be found in organizations – even if the installed versions are problematic from a security standpoint. This is due to the fact that a lot of internal applications are still enforcing the use of old versions of such tools (notably Internet Explorer).

The human element kicks in when trying to introduce the malicious code into the organization. A common and still highly useful attack avenue is the use of targeted phishing emails (spear-phishing) that provide an appealing context to the target which would make it more "legitimate" to open and access any content embedded in the email or referred by it (external links to sites that would carry the malicious code).

Another human element that can be easily exploited in getting malicious content into an organization is using physical devices that are dropped or handed to key personnel. This paper is read at a conference, where multiple vendors provide giveaways in the form of electronic media (be it CDs, or USB memory devices). A classic attack vector would use such an opportunity to deliver crafted malicious code over said media (or follow-up emails) to key personnel.

The last element of the infiltration attack surface is the physical one – gaining physical access to the offices or the target (or one of its partners) is easier than perceived, and provides the attacker multiple ways of getting their attack code onto the network. From casually tapping into open ports and sniffing the network via a remote connection (by plugging in a wireless access point), or simply plugging in an infected USB drive into various PCs, such attacks can be carried out by personnel that are not necessarily security professionals (we often use paid actors to carry out such engagements and gain a familiarity bonus for local accent and conversation context).

Data targeting and acquisition

Before launching the actual attack on the organization, the attacker will perform some basic intelligence gathering in order to properly select the target thorough which the attack would be conducted, as well as the data that is being sought after.

Such intelligence gathering would be done through open source intelligence, as well as onsite reconnaissance. From an organizational perspective, the ability to map out the information that is available through public channels on it is crucial. It provides a baseline on which the organization can prepare the threat model, which reflects the attacker's intelligence data, and allows it to properly address any policy/procedure issues in terms of exposing data. Social media channels are one of the more popular means of intelligence gathering, and can easily be used to build an organizational map and identify the right targets to go through.

Finally – in terms of data targeting, the attacker would usually focus on intellectual property, financial and personal information – all of which would be easily sold on the open market and to competitors.

Once the target has been selected, and the data have been identified, the actual payload needs to be created and deployed using the attack tree that was chosen for infiltration. Such payload (often called an "APT" – Advanced Persistent Threat) is usually no more sophisticated than a modern banker Trojan that has been retooled for a singular targeted attack. Such Trojans are available for purchase on the criminal markets for $500 to $2,500, depending on the plugins and capabilities provided. These by themselves offer a fairly low detection rate by security software, but in order to assure a successful attack they are often further obfuscated and packed.

Command and Control

One of the main differences between targeted attacks and the more common malware is that targeted attacks need to take into consideration the fact that the connectivity between the payload and the attacker cannot be assured for long periods of time (and sometimes are consistently nonexistent). As such, the C&C (command & Control) scheme for such an attack needs to take that into considerations and the payload should be well equipped to operate fairly independently.

Nevertheless, some form of control communication is needed, and usually utilizes a hierarchical control structure, where multiple payloads are deployed into the organization at different locations, and are able to communicate with each other to form a "grid" that enables the more restricted locations to communicate through other layers to the attacker outside the organization.

Exfiltration

So far we have reviewed how an attacker would infiltrate an organization, target the data it is after, and find a way to somehow control the payload deployed. Nevertheless, getting the actual data out is still a challenge, as more often than not, it is located so deep inside the network that traditional means of communications (DNS/HTTP tunneling for example) are not applicable.
However, the way that organizations build their infrastructure these days basically call for other means of getting the data out. Following are a few concepts that should be used for testing exfiltration capabilities as part of penetration testing – which have proved to be useful on multiple major corporations, as well as government/defense organizations.

First off – the obvious: use "permitted" protocols to carry the information out. These are usually DNS traffic and HTTP traffic. The data itself may be sensitive and filtered by DLP mechanisms, and as such should be encrypted using a method that would not allow a filtering/detection device to parse through it. After encryption, the data can be sent out through services such as Dropbox, Facebook, Twitter, blog comments and posts, wikis, etc… These are often not filtered by the corporate control mechanisms, and are easy to set up if needed to (a Wordpress blog for example, where the payload can post encrypted data using the attacker's public key).

The next exfiltration method that usually works is simply printing documents. Obviously, we won't print out the original data as it would be easily detected and shredded. Instead – the encrypted information can be sent to shared printers (which are easy to map in the network), and made to look like print errors (i.e. remove any header/footer from the encryption method we utilize). Printouts like that are more likely to end up in the paper bin rather than the shredder, and later extracted as part old-school dumpster-diving. Such documents just need to be OCR'd after their retrieval and decrypted to reveal the sensitive data that has been stolen. This method is usually more efficient where proper mapping of the paper disposal process of the target has been performed.

An alternative to printing the encrypted data uses the same means of exfiltration – the shared printers. When a shared printer is found to be a multi-function device with faxing capabilities, the payload can utilize it to fax out the encrypted documents.
In this situation the payload would still need to keep "operational awareness" as some DLP products would actually look at the fax queue for information that is not supposed to leave the organization, hence using the encrypted text is better form.

Exfiltrartion through VoIP

This is the main concept that is being displayed here. As VoIP networks are usually accessible to the local network (usually to accommodate soft-phones, and just a simpler network topology), crafted payloads are able to utilize this channel to exfiltrate data. The method proposed here is to initially sniff the network and observe recurring patterns of calls, and user identifications (to be later used when initiating the SIP call). After some initial pattern can be mapped out, the payload encodes the data to be exfiltrated form its binary format to audio.
A proposed encoding maps out the half-byte values of the data stream to a corresponding scale of audio tones using 16 distinct octaves on the human audible frequency range (20Hz to 20,000Hz). Therefore, a byte value is split into it's high, and low values, and then the value is used to select an octave (out of the 16 available ones).

```
For each byte do:
    hb_low = byte & 0xF
    hb_high = byte >> 4
    voice_msg += octave[hb_low]
    voice_msg += octave[hb_high]
sip_call.play(voice_msg)
```

Sample 1: pseudo-code for voice encoding of data to 16-octave representation

The octave is then played for a short period of time (for example ½ second) on the final output voice channel. The final output is then played back on an opened SIP call that can be made to almost any number outside the organization (for example a Google voice account's voicemail) for later decoding back to the original binary data.

For the decoding itself, an approximation analysis is performed on the input sound file in order to identify the maximum frequency detected for each "time slice" which carries the generated tone, and then comparing the frequency to the octaves used in generating the original sound. As sounds get distorted and downsampled as they go through older telephony systems, and cannot be guaranteed the same quality as used on pure VoIP circuits, the spacing between the frequencies used should be enough to create a distinction between tones.

```
For each sample_pair do:
    max_f = getMaxFreq(sample0)
    bh = getByteFromFreq(max_f)
    max_f = getMaxFreq(sample1)
    bl = getByteFromFreq(max_f)
    byte = bl | bh << 4
    out_stream += byte
file.write(out_stream)
```

Sample 2: pseudo-code for voice encoding of data to 16-octave representation

This method can obviously be optimized in several ways – first, using more octaves (as long as they are distant enough in their frequencies and non-harmonic) to represent more data in each tone being played, and again in the time each tone is played to essentially compress the data over less time.

The proof-of-concept that is being released along with this paper is intentionally designed to act as an example (although it can be easily tooled to carry out a significant amount of data, and has been used in several occasions to do so in penetration testing engagements to exfiltrate highly sensitive data).

In terms of protection against such exfiltration techniques, the recommended strategy is to basically extend the same kind of monitoring and scrutiny that is being applied to traditional electronic communication channels (web, email) to the VoIP channels. Although voice communication monitoring has been traditionally associated with more government type organizations, the move to VoIP enables small companies and businesses to extend the same security controls to the voice channel as well. DLP systems for example could easily be used to transcribe (voice to text) phone conversations, and apply the same filtering to them, while alerting on calls that contain non-verbal communications as suspicious.

Future Directions

The concepts presented here in relation to advanced techniques in data exfiltration are not only theoretical. We have been observing progress in the way that advanced threats are addressing this issue, and adding more capabilities and techniques to the arsenal od data exfiltration beyond simply staging data in archives and pushing it out through FTP connections. The proliferation of VoIP networks that are being configured mainly for convenience with not much security concern into them have allowed us to observe a few cases where similar methods of utilizing such channels were used in the transmission of data outside of the targeted organization.

Additionally, VoIP networks also allow simple bridges between networks with different classifications that may not have a direct data connection in the "classic" sense of a TCP/IP network infrastructure.

The other techniques mentioned in this paper (namely the use of covert channels in legitimate services such as blogs, social networks, Wikis, and DNS) are already in full use and should be a reality that corporate security should already address.

Mitigation Strategies

When attempting to address data exfiltration the first important thing to realize is that infiltration is almost taken for granted. With so many attack surfaces encompassing different facets of the organization (outside technical scopes), security managers need to realize that detection and mitigation of data exfiltration is an integral part of the strategic approach to security.

Identifying data in transit and in-motion is the basic element that allows detection of exfiltration paths, and many tools already allow organizations to address this issue. The missing components usually lie in the realms that traditional products and approaches neglect such as encrypted communications, and "new" channels such as VoIP. Addressing these channels is not an unresolved problem, and in our experience simple solutions can provide insight into the data carried in them.

For encrypted channels a policy (both legal as well as technical) of terminating such encryption at the organizational perimeter before it is being continued to the outside should be applied. This approach, coupled with an integration of existing DLP products to identify and detect misplaced data, will provide the required insight into such communications. An unknown data type carried over legitimate channels should be flagged and blocked until proven "innocent" (for example custom encryption used inside a clear-text channel that cannot be correctly identified and validated as legitimate).

For VoIP channels, the same approach that is being applied to the more traditional web and email channels should be used as well. Full interception and monitoring of such channels can be applied, even when not in real-time – such as recording all conversations, processing them using speech recognition software, and feeding the results back to the DLP. This approach yield the same results as a DLP installed on the email and web channels does. Additionally, the investment in terms of time, human resources, and materials is negligible when compared with the added security in terms of detection and mitigation of such threats, and complements the layered security approach that should have covered these aspects in the first place.

Summary

This paper covered both the more advanced infiltration techniques utilized by targeted attack on organizations (which should have been covered by the security industry to a point, although organizations are still struggling with this aspect), as well as raises the awareness to the more problematic issue of detecting data in transit outside of the organization as it is being exfiltrated. Several methods of exfiltration have been discussed, with the more evasive one being the use of voice channels on VoIP infrastructure. We believe that the current practices do a disservice to the layered security approach that is being preached in the security industry by leaving gaping holes in the exfiltration paths monitoring and mitigation. While there may be claims of privacy issues, such gaps are similar in the way that data is being processed and inspected to existing channels, and should adhere to the same standards of privacy protection and abuse as traditional solutions that address data leakage do.