FEATURE

THE WORLD OF BOTNETS

Dr Alan Solomon, Programmer, UK
Gadi Evron, Beyond Security, Israel

With a Trojan horse on one compromised computer, you would be able to do whatever you wanted. That computer would be as good as your own. You would own it. Now imagine that you owned 100,000 such computers, scattered all over the world, each one running and being looked after in someone's home, office, or school. Imagine that with just one command, you could tell all of these computers to do whatever you wanted.

You could tell them to access the same website simultaneously, to search for security weaknesses in nearby computers, or just to send out large amounts of spam. If you sent out one spam message from each computer every second, you could send a million pieces of spam in ten seconds; hundreds of millions per hour. You would effectively be in command of an army of robots, or bots for short. The actions taken by these bots would appear to be the actions of the individuals whose computers were compromised.

Now imagine that you own a million compromised machines. You would be in possession of a force capable, at the very least, of significantly slowing down activity on the Internet or taking down a large Internet-based business. This, of course, would be illegal, but that does not stop some people from doing it. Today, there are 3.5 million bots on unique IP addresses used every day for spam purposes alone.

BOT-HERDING HOW TO

The first step towards becoming a bot-herder is to google the subject, and proceed to recruit an army – a botnet (also called a drone or zombie army). Some of the most commonly used recruitment methods include sending out millions of email messages that contain something tempting to click on, or setting up websites with browser exploits that activate a drive-by (the surfer merely has to access the web page for a trojan to be installed on their computer silently, without any interaction). Another successful method of bot recruitment is scanning networks for vulnerable computers. Not everyone installs patches, and as a bot-herder you don't care whose computer you recruit for your bot army; the name of the game is quantity, not quality.

You'll also need to put some effort into maintaining your army, as there are many rivals, obstacles and enemies on every corner. Users may run one of these pesky anti-virus products that discover and remove your bot, their computer may crash, and they may reinstall *Windows*. Alternatively, another botnet controller or bot could come along and hijack your control of the user's computer. But that's no great problem – there are a lot of potential victims out there, you just need to keep recruiting. There's no need to worry about losing the grunts when there are a lot of fresh troops out there without criminal records (such as having been blacklisted for sending spam).

Your biggest danger as a bot-herder is being discovered, unveiled and arrested. This has happened to other botnet controllers, but it's a rare occurrence – people tend to discount small risks, and the large profits available from bot-herding outweigh the remote possibility of a couple of years in prison. Indeed, if you live in Ruritania, the chances of being collared are near-zero.

Another danger is that your botnet will be discovered by one of the white-hat bot-hunting groups, and your command and control (C&C) will be cut. The C&C is the means by which your bots report back to you and get their marching orders. As a couple of examples, this might be via an IRC channel or nickname (with the advantage of the anonymity they confer – thus reducing the risk of two years eating porridge [*widely believed to be the staple diet in British prisons*] – as well as the ease of moving to a different IRC server at a moment's notice), or it might be via the URL of a website. The C&C channel is the botnet's weakest link, and these white-hat groups hunt bots voluntarily, just because it's the 'Right Thing To Do' (and fun).

MAINTENANCE

So how do you maintain your botnet C&C when an opponent is trying to dissolve it? Some of the methods used include the use of free dynamic DNS services. Alternatively, you can use throw-away domain names and hosts. This allows for a quick response, moving from an at-risk IP address to a new one, while still using the same DNS record for the bot to connect to (only now, it would be pointing at a new IP). In the same way, the DNS records can be discarded and replaced while the IP address remains constant. Alternatively, you can point the host at many different IP addresses, or use many different hosts.

Botnet C&Cs also enjoy the benefits of redundancy, with unlimited fast-changing IP addresses and hosts ('fast-flux'), and failover capabilities with alternate C&C channels in case the first ('the head of the hydra') is compromised or fails. Further, a botnet may operate much like a terrorist cell, where different botnets are used to operate each other in a tree-like structure. For example, botnet A consists of 10 computers, and each bot within that botnet controls 100 computers contained in botnets B–G, and so on, with each of the branches compartmentalized from the others.

But these are today's botnets – it is possible to design a botnet that doesn't have the weak link of a C&C, and maybe that will happen in the future. At the moment, whenever a C&C becomes more complex, it also becomes that much easier to detect. As a result, today's miscreants are focusing on making the take-down of the C&C channel irrelevant.

THE MONEY GAME

Finally, we reach the name of the game – profit! The return on investment from running botnets is significant. The people who run botnets can use them for any purpose. Although within the miscreant community there is a full social structure with many players, most botnets are run by organized crime groups involved with phishing and other types of financial fraud. That said, some botnets are still run by kids with nothing better to do. There is enough profit for everyone.

The uses for botnets are endless, varying from click-fraud and spyware installations to phishing and credit card theft, with identity theft, spam, anonymity, the launching of distributed denial of service attacks (often with the intent of blackmail) and malware-spreading thrown in for good measure. Botnets are weapons, and while they can be used for surgical strikes and intelligence gathering, they can also be utilized as weapons of mass destruction. Once the trojan is installed, it is only a matter of a small change in the source code or new instructions in order for it to be optimized for different uses.

Among the most difficult types of attack to deal with are distributed denial of service (DDoS) attacks. A target site is hit for an hour, by 100,000 different computers all over the world. It is extremely difficult and often impossible to distinguish their access attempts from accesses by genuine, money-spending customers. The site is brought down by the sheer impossibility of servicing all these accesses. When the DDoS starts up again the next day, the owner of the site is wits-ended, and if offered an end to his problems for $5,000, is likely to be very tempted to pay up.

But, of course, once you pay the Danegeld, you never get rid of the Dane. The $5,000 would only be the first instalment. Further, who is to assure you the miscreants won't attack regardless, or that others won't soon follow?

Which brings us to the classic problem that any fishpaster faces (blackmail is such an ugly word): how do you collect the loot, without being collared? The same problem faces phishers. Your botnet won't help you ... or will it? Enter the mule. A mule is someone who carries something of value, and who also carries the can if anything goes wrong. Mules are recruited by sending out spam using botnets. You've probably seen some of these spam messages amid the flurries of advertising for pills, potions and penis enlargement. 'Make money from home.' 'Join our international financial agency.' 'Become an export agent.'

Here's how it works: the mule receives the money and pays it into their bank account, keeping 10% for themselves, and sending the rest onwards using, for example, *Western Union* (which is a good way to receive money because the pickup is anonymous). The only purpose of the transaction is to muddy the money trail, launder the cash. Any problems, such as the FBI turning up at the door, will affect only the mule – and there are plenty more where he or she came from. Of course, the mules aren't told that – 'No risk', the recruitment messages say.

THE PROBLEMS

The bot-herders' C&Cs have become so protected with redundancy and secondary control channels that taking these down no longer has any kind of effect on the problem – other than pushing the miscreants to work harder at developing new technologies, and having them use alternative ISPs.

And indeed, the technology has advanced significantly and steadily over the past years, from simple IRC-based trojans in 1996–7 to today's fully fledged DNS control-based trojans using man-in-the-middle rootkit technology, effectively adding quality where before this game was about quantity alone. These listen in on every HTTPS session to steal credentials, aiming for financial information. Indeed, one of the main businesses of botnets today is stealing credentials and information wherever they manage to infiltrate. On an operational level, the organized groups behind the botnets operate dedicated teams dealing with everything from stolen data and deciding which stolen accounts are worth their time, through money transfer, all the way to real-world operations to support their criminal activity globally.

With a return on investment amounting to tens of thousands of US dollars for a relatively small botnet with click-fraud alone, and with damages from phishing alone likely to reach as much as two billion US dollars in 2006 globally, the bad guys are not likely to change their occupations any time soon.

Much as with every other established threat (which was ignored when it was small enough), it is a never-ending arms race of small victories on both sides as each escalates with a new technique or technology, forcing the other to adapt.

Leaving other concerns aside, millions of identities are being stolen every day across the world, through the use of the man-in-the-middle trojans utilizing rootkit technology. Organizations ranging from a moms&pops shop to Fortune 50 corporations have compromised machines on their networks, which are sending out spam or participating in DDoS attacks. These can easily be found by anonymous third parties and utilized for espionage.

Data stolen from a company's clients often affects the organization too. As an example, stolen credit card information affects the credit company as well as the business from which things are purchased. The problem is: how do you secure something when it is on the remote client side, completely outside your control?

FUTURE

That's today, but what of tomorrow? Here are just a few examples of what we can expect to see in the not very distant future:

- Information already extracted today could be better analysed and used for more advanced purposes. Organizations with bots on their networks could be targeted for industrial espionage. Further aggregation and correlation of data could be introduced, so that financial attacks evolve to intelligence-gathering beyond stealing money or hijacking on-going transactions. The world will slowly be mapped in advanced social networks, seeing who is in business with whom and their level of financial ability.
- Strength of arm, which already rules the Internet, will become even more apparent, with spammers and organized crime groups protecting their business interests, causing damage and hurting businesses and public alike when they are threatened.
- Reputation systems will become even harder to implement when it is very difficult to establish whether a user identifying to a service is a genuine user or a bot. Trojans, i.e. bots, will transform from ill-behaving entities online to advanced critters that simulate regular user behaviour, making their detection extremely difficult.

Perhaps most importantly, there is little in place to change that. Law enforcement organizations cannot deal with the immense number of complaints they receive daily, and very few of them have the expertise to handle these cases. When they do, other parts of the legal system are fearful of computer-related investigations and will avoid them if at all possible. They are heavy-duty, demand a large investment of resources and are very technical, to a level that is often extremely confusing. Furthermore, a murder case is more sexy. Investigations that require wire-tapping, long-term research and global cooperation are even less likely.

The lives of the law enforcement organizations are not easy, though. We work with many skilled, able and smart folks. They want to help. Earlier we mentioned that bot-herding is an illegal activity. That isn't always the case – the law often tends to be one step behind current events, especially where technology is concerned, and in some countries, bot-herding isn't illegal. Furthermore, being illegal does not make it actionable to law enforcement – how do you prove damage from running a command and control server?

SOLUTIONS

On the network side, ISPs are left with the choice of protecting their networks on their own, which may, in the case of law enforcement interest, hinder investigations. The bad guys adapt, change their IP addresses and host names ever faster, and become very difficult to stop. In the past year there has been an on-going migration of C&C servers to China.

Anti-virus products detect samples and in some cases remove infections, yet with an average of 12,000 new bot samples coming out every month, around 10 per cent of which are financial fraud specific (and most of the remaining being multi-purpose trojan samples), the traditional anti-virus solution, as important as it may be to this fight, is no longer up to scratch as a lone solution which is inherently reactive.

Cooperation between ISPs (which in some cases run honey nets, unwillingly host these malware and phishing sites, etc.) and anti-virus vendors (who see what the samples do and what C&C channels they connect to) through operational vetted groups such as DA (Drone Armies) and MWP (Malicious Websites and Phishing) helps to mitigate some of the problems.

Botnets are a serious problem, but this is merely an example of a much larger problem with Internet security today. It is an economic issue, and without an economic solution that changes the miscreants' cost vs. benefit equation by reducing their gains and significantly heightening their risk, not much will change. Every new technology invented will be countered or circumvented by moving to new attack mediums, spam being a good example.

The problem will not go away, but the bad guys behind it can be put under more stress so that it becomes manageable. Law enforcement needs assistance, and not only with more resources. Policy makers should set the pace, allowing the law enforcement bodies to handle these cases to begin with. Further, extradition laws around the world can really come in handy. Some international work to recognize computer crime for what it is and tip the balance would be very helpful.

Every party in this fight is busy, and has their own business to take care of. That said, without better cooperation, intelligence gathering and coordinated response online as well as in the physical world, we believe that the current threats will eventually become unmanageable on the infrastructure level. Get involved with the vetted operational groups mentioned, meet some of those involved and see how you gain information to help your business, get introduced to the latest threats that others outside of your field see first, which will also affect you, and help to turn the tide back, fighting the real, original enemy (which is not the marketing department or your competitor). You can find more information about these groups at http://isotf.org/.