Global Network Pandemic – The Silent Threat

Darren Grabowski, Manager
NTT America Global IP Network Security & Abuse Team

The Internet is in the midst of a global network pandemic. Millions of computers on the Internet are compromised in some fashion. Many of these computers are "zombie" members of a malicious botnet. Most users and operators know a problem exists, but few are in a position to see how big the problem is. Solutions are simple, the right tools, dedicated staff and cooperation. Implementation is the most difficult part. Networks large and small must work together to mitigate this threat.

Quickly for those that may not know how botnets work: they are a group of infected machines across the Internet controlled by one machine. The infected machines log into a central location, like an Internet Relay Chat (IRC) server where one machine, the command and control server (C&C), could then issue commands to them to send spam or launch a DDOS.

The NTT America (NTTA) Security and Abuse Team (aka SAT) are tasked with responding to complaints about security (port scans, malware, DoS, etc.) and abuse (spam) issues. They also monitor the network for security issues using a variety of tools. One of the many tools used is a Darknet. According to the Team Cymru Darknet Project, a Darknet is "a portion of routed, allocated IP space in which no active services or servers reside.[1]" In short, there should be no reason for any traffic to enter this space. In truth there is one server in a Darknet. This server collects packets that enter the Darknet. This data can be used for immediate action or stored for further analysis.

The levels of nefarious traffic from this silent threat are low compared to legitimate traffic. The amount of bandwidth consumed is so low, many network operators may choose to ignore the traffic, or they may not even realize the silent threat hiding amongst their legitimate traffic. This could be due to lack of resources or just plain laziness. Either way this problem cannot be ignored, nor can it be solved by just a handful of providers.

The NTTA SAT's own Darknet data reflects this trend of low bandwidth consumption. The server resides on a 100M Fastethernet connection. The mrtg graph that follows is from June 11th to July 11th, and it shows about 50-70k bps with spikes up to 100k bps of bandwidth consumption, and the packets per second rate hover under 100pps with spikes up to 120pps. The amount of bandwidth being used is just .05% of what is available on this single FastEthernet connection. As you can see it is very easy to overlook or ignore this. The amount of traffic is not going to cause network issues like a Denial of Service attack does. Hence the reason it is a "silent threat."

[1] http://www.team-cymru.org/Services/darknets.html

NTT America Darknet Traffic Graphs

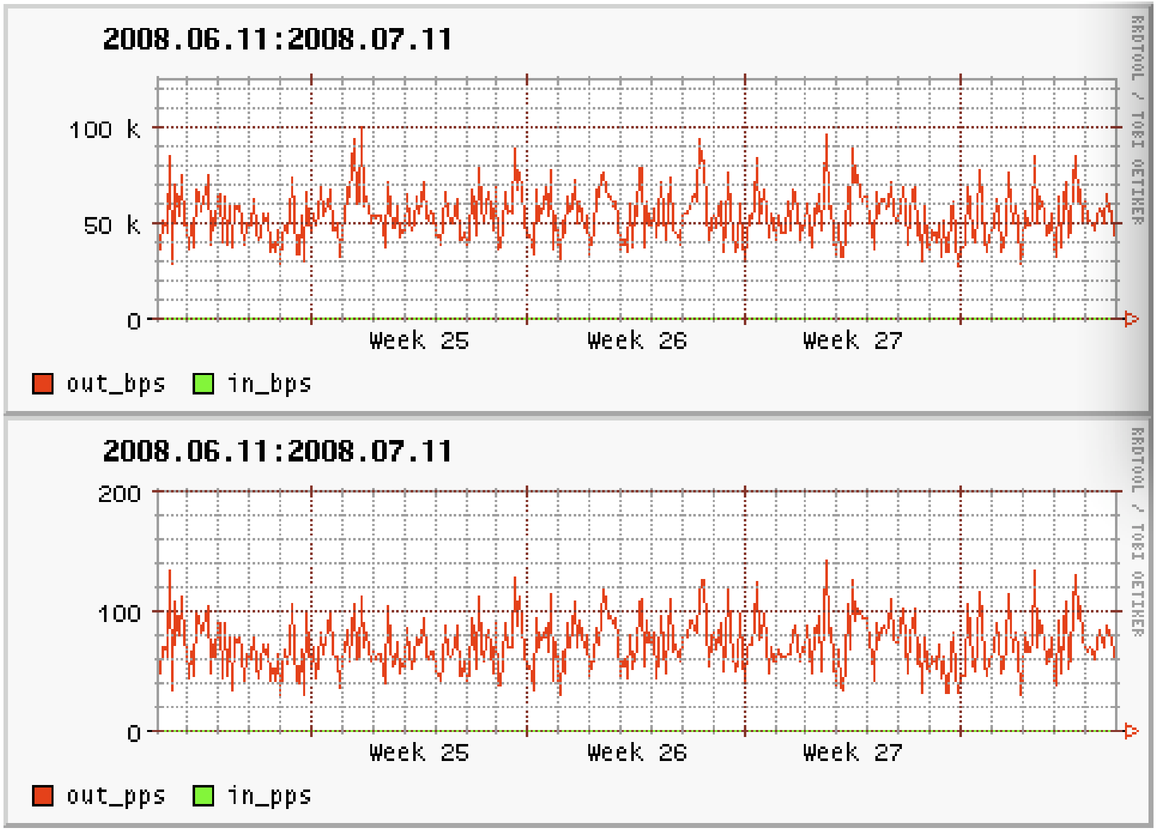

Why is This a Problem?

Trend Micro puts the number of malware infections in November 2007 to over 7 million[2]. Message Labs estimates that over 70% of all email are spam[3]. Security analysis firm Marshal claims that 85% of spam comes from just six botnets[4]. In a recent report, Commtouch made the news with claims that in the second quarter of 2008, there were an average of 10 million active botnet members on any given day, and botnets are winning the spam war[5]. Companies are already spending thousands, if not millions of dollars, to filter incoming spam but they ignore their own outbound e-mail. Infected machines sending spam via corporate servers can lead to those servers being blocked by not only public blocklists such as the Spamhaus SBL, but also private blocklists. This can be avoided!

The large number of compromised machines, if directed by a malicious botnet, has the combined ability take down key Internet infrastructure. The compromised machines can also be used for other harmful activities that could cause a severe financial

[2] http://us.trendmicro.com/imperia/md/content/us/pdf/threats/securitylibrary/tre_threat_report.pdf
[3] http://www.messagelabs.com/mlireport/MLI_Report_April_2008.pdf
[4] http://www.marshal.com/trace/traceitem.asp?article=567
[5] http://www.networkworld.com/news/2008/070908-botnets-winning-spam-wars-says.html

impact (i.e., phishing). According to a Gartner Survey, in 2007, 3.6 million adults lost money in phishing schemes resulting in an estimated loss of $3.2 billion[6]. Phishing is only one part of the problem. Attacks have already caused issues for countries like Estonia[7] and infrastructure like DNS[8]. Thankfully, the attacks against DNS went mostly un-noticed; but sadly, that was not the case for Estonia.

How Did This Pandemic Happen?

The problem has swelled largely due to a number of contributing factors. First and foremost is the lack of proper security. Security is often an afterthought. Perhaps the cost is too high, or it is just not convenient. People also feel safe in their homes and offices. They trust that their ISP or IT department are protecting them using the "latest and greatest" technologies. Meanwhile, reality shows that budget cuts did not allow for a firewall, or the latest security patches were not installed due to cuts in staffing. There is also the issue of the unknown. Anti-virus companies cannot block that which they do not know. A recently released virus can create havoc on a network, clogging e-mail servers and bringing networks to a crawl. The push to make a deadline for the latest software release may mean lack of proper security audits.

Phishing attacks became so sophisticated that users were simply fooled. This problem continues to rise and it is a global problem. The UK's fraud prevention service, CFIAS, reports a 182% increase in phishing fraud in the second quarter of 2008 when compared to the same time period in 2007[9]. The virus/worm authors also play on emotions and desires. Sensational news stories or promises of adult content drive people to click on attached malware or visit links where, instead of finding their favorite celebrity in a compromising photo, they were treated with a nice bit of malware installed without their knowledge, or in some cases, with their permission!

There is also a segment of the community that blames software developers, such as Microsoft. In reality, we have all contributed in some way to this problem. The lack of educating end users contributes to this. Users are given flashy marketing materials that boast the impressive nature of their newly installed service, but there is no real warning, short of some wordy legalese, in those materials.

In all fairness, Microsoft has done their part in decreasing some botnet activity. In a presentation during NANOG 41 on Storm Worm, Josh Ballard's research showed a decrease in the number of storm-infected machines[10]. This decrease coincided with software updates and a release of the Malicious Software Removal Tool. Solutions like the MSRT and free virus scanning/removal tools are extremely helpful in mitigating this threat, but how do we prevent it?

[6] http://www.gartner.com/it/page.jsp?id=565125

[7] http://asert.arbornetworks.com/2007/05/estonian-ddos-attacks-a-summary-to-date/

[8] http://www.icann.org/announcements/factsheet-dns-attack-08mar07.pdf

[9] http://www.telegraph.co.uk/money/main.jhtml?xml=/money/2008/07/14/ymfraud114.xml

[10] http://www.nanog.org/mtg-0710/presentations/jballard-ispbof.pdf

What Can Be Done to Mitigate This Threat?

We are not going to rid the Internet of compromised machines. That does not mean the problem should be ignored, or that we can't mitigate it. What we need to do is reduce the capability of botnets, which means reducing the number of infected machines. Networks of all sizes can assist by properly monitoring their networks and removing infected machines. Tools exist to monitor traffic at relatively low costs. A Darknet, or any other similar monitoring device, allows networks to find potential compromised machines by watching their IP space. Some monitoring devices can be deployed at a relatively low cost using existing hardware, or using data from existing intrusion detection systems. This has proven to be a valuable tool for the NTT America Security Team.

The following graph shows the number of IP addresses seen scanning the NTTA Darknet from June 11th to July 11th. The total number of individual addresses scanning the Darknet was as low as 22,257 on June 11th, and the highest number we saw during this time period was 47,027 on June 23rd. The total number of NTTA individual addresses scanning the Darknet was as low as one on July 7th and as high as 18 on July 4th. These low numbers from NTTA IP space make it easy to identify the infected user and notify them. This process could be easily automated, so there would be minimal impact on staff that may already be stretched thin.

NTTA Darknet Individual IP's Per Day

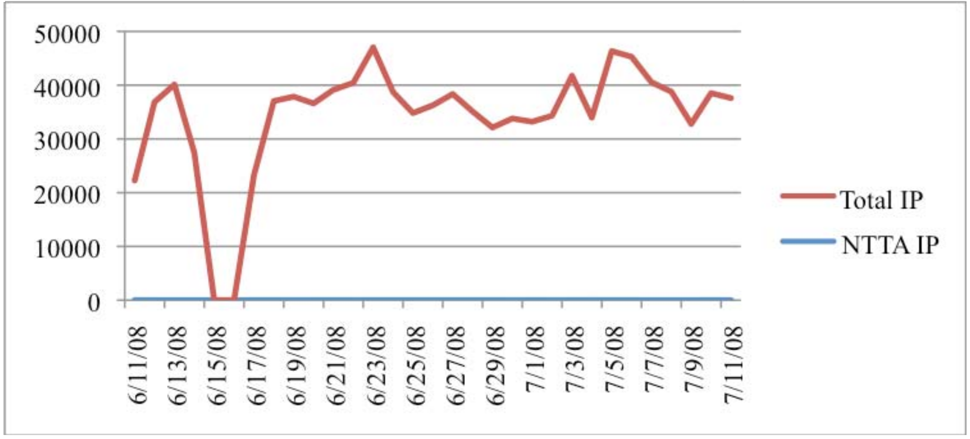

Using some scripts and netflow data, you can monitor your network for activities like DoS attacks. IP addresses participating in a DoS can be investigated a bit further. By combining data from a DoS attack, or a Darknet, and other sources such as grey-listing or spam traps, you can potentially find a botnet member. Once suspicious hosts are located, you can check to see if these hosts are communicating with a common host, which could be a command and control server. Taking down a command and control server can disrupt a botnet, even for a short while. If the compromised host's owner can be contacted, there may be a chance that a list of bots can be obtained and further notifications can be sent out. Clever use of existing data and equipment is one way to keep costs down while helping secure your network. U.S. based providers may have already purchased equipment for CALEA compliance. The Communications Assistance for Law Enforcement Act, or CALEA, is a U.S. statute that covers lawful intercepts on digital transmissions, including data and VOIP. Companies, such as Packet Forensics[11], sell surveillance platforms. These devices are capable of doing deep packet inspection, stealth packet filtering, transparent redirection, as well as a host of other services. A network operator could leverage the pattern matching capabilities of these machines in their hunt for compromised hosts on their network. Even if CALEA is not a concern, these devices could be useful to a network operator who wants to monitor their network for harmful activities.

Most Internet providers block port 25 from their dynamic IP space, and in some cases, from their static IP space. This is great in helping stop the flow of spam and other nefarious activity using email, but it does not stop infected machines from launching attacks, nor does it fix the underlying problem of a compromised host. There is now a trend to move toward a walled garden approach. These allow providers to restrict the activity of a user until their machine is clean. This also allows another method of communicating the issue to the user. Users may ignore email notifications sent to them. With a walled garden, those users can be notified via a redirect to a web site on their browser, and access to the Internet can be severely restricted, or cut off completely. There are those who argue that providers should call these infected customers, but, depending on the size of the provider and the number of infections, that may not be practical. Providers should also be willing to suspend infected user accounts if the problem persists. A walled garden does not have to be limited to an ISP. Networks of any size could benefit from this approach!

The hunt for compromised machines is not limited to network providers. Anyone hooked up to the Internet can watch their traffic and report their findings. Instead of ignoring warnings from an intrusion detection system, automated reports could be sent out. Tools exist to locate the source network, Team Cymru's IP to ASN mapping project[12] is a good example of that. Tools like the abuse.net whois[13] or DNS based lookup services can be used to find out the correct reporting address. Most intrusion detection systems have some sort of reporting process, and hopefully include enough automation so that it does not become like a 2nd job. Automation means people might be willing to spend a little bit of time reporting intrusions.

These are only a few suggested solutions to this problem. The cost of tools for monitoring this threat can be very low if budget is a concern. If you take stock of what is already on your network, chances are you may already have the tools needed. It just takes a little bit of time and effort to use them to your advantage.

[11] http://www.packetforensics.com/products.safe

[12] http://www.team-cymru.org/Services/ip-to-asn.html

[13] http://www.abuse.net