Intrusion Detection System

"Transforming Network Data into Security Knowledge."

White Paper

Distributed Denial of Service and K2-Defender

Author: Arrigo Triulzi
Chief Security Officer, K2 Defender Ltd.
Published: 7th February 2002

Overview of DDoS

In simple words, a DDoS attack is a co-ordinated attempt to starve a site of resources. This need not be a direct attack against a particular site, but may be a 'collateral damage' sustained as a result of an attack on a 'neighbour[1]'.

The simplest form of a DDoS attack involves sending vast quantities of requests, not necessarily malicious, to a specific host, for example, a web site, rendering it overloaded and unresponsive to any other legitimate request. A slightly more sophisticated attack is on the communication links. Although the hosts behind the link may be up and running; by saturating the link, no traffic reaches them. The net effect is again that the site is unreachable to legitimate traffic.

To perpetrate a DDoS attack, it is crucial to have a large number of Internet-connected hosts available – not necessarily on high-bandwidth links, since it is the aggregated bandwidth that matters –, which direct a focused stream of traffic simultaneously to a specific location on the Internet[2]. These hosts are often unmaintained systems in an academic environment; a prominent example has been the Korean academic network. Another source of hosts is the vast number of unsecured home PCs connected to 'always-on' cable networks such as @Home in the USA.

A client program is installed on the host which then becomes what is known as *zombie* ready to take orders from a 'master' or a chain of 'masters'. When the system receives orders from a chain of masters, the process becomes even more complex as it is more difficult to back-trace to the 'owner' of a DDoS network once access to a *zombie* is secured. DDoS software has developed a number of features making it even more difficult to predict attacks. Amongst these are:

- Pre-distribution of victim IP addresses that makes it even more difficult to predict the attack even if the network is being monitored.
- Encrypted control messages between masters and *zombies*, and
- Use of Internet Relay Chat (IRC) channels for control message distribution, all aimed at protecting the 'DDoS network' from being dismantled.

It is important to realise that having multiple redundant sites does not mitigate the effect of a well-planned DDoS attack. Reconnaissance scans will have taken place before the attack to ensure maximum effect and multiple sites would have been noticed. Similarly clusters of web servers only offer a larger target cross-section[3].

The effects of a DDoS attack on the Internet infrastructure clearly depend on the target. The target can be as small as a single server moving to a whole ISP or indeed the core services of the Internet. Some possible scenarios are:

1. *DDoS against a single server*

This is the simplest DDoS attack that one could implement as the resources required are minimal. The target is a single server or it can be one of the web sites hosted on the server in the case of multiple virtual web sites on a single server.

2. *DDoS against a single site*

This is an attempt to isolate a site from the Internet. The closer we get to a particular site the smaller the bandwidth to it becomes compared to the NAPs[4], the transit ISPs and the ISPs themselves. Saturating access to an ISP's customer is relatively trivial compared to the scenarios that follow.

3. *DDoS against a particular ISP*

The smaller ISPs are the more vulnerable they are to this type of attack. Smaller ISPs that obtain their bandwidth from transit ISPs or bandwidth providers normally have a small number of links of a significantly smaller capacity than a back-bone ISP.

1.'Neighbour' is a common networking term used to describe networks of distinct customers which share the same bandwidth provider.
2.Normally this would be a single IP address or a subnet (a contiguous set of IP addresses).
3.It is a well-known saying in network security that one should always assume that "the attacker always has more bandwidth than you do" when trying to evaluate counter-measures.
4.Neutral Access Points (NAPs), also known as IXPs or Internet eXchange Points, are locations where traffic is passed between ISPs. Well-known NAPs include MAE-East and MAE-West, SprintNAP in the USA, LINX in London and AMSIX in Amsterdam.

4. DDoS against well-chosen back-bone ISPs

A small number of ISPs carry the vast majority of Internet traffic – if we exclude academic networks in Europe –, in particular transatlantic traffic. Transatlantic links are of relatively large capacity, but there are not many of them. Saturation of these links can isolate sections of the Internet, in particular as most traffic between Europe and Asia travels via the USA, it would isolate Europe from the rest of the world.

5. DDoS against strategically chosen Neutral Access Points (NAPs)

Although the framework of these NAPs is often based on Gigabit Ethernet, the routers of the various ISPs are not necessarily fed by Gigabit connections. Saturating a single NAP would most likely shift traffic to other NAPs, but attacking multiple NAPs could effectively cut the Internet into many 'network islands[5]' unable to communicate with each other.

6. DDoS against the root name-servers

These are thirteen systems[6] deployed strategically around the Internet essential for the Domain Name System to work. Without these systems it eventually becomes impossible to resolve an IP address from a name[7]. The effect is not immediate because some information may be cached at local sites but for a limited time. This would have a devastating effect, but most likely would provoke a very rapid response by authorities, especially in the USA.

The above list shows how the DDoS phenomenon can take different forms and requires a different approach depending on the preferred attack scenario. So far the most publicised episodes were those against Yahoo! and E-bay, both of which fall into the *DDoS against a single site* category. The real issue is of course how to react and, if possible, predict such attacks.

K2-Defender

K2-Defender is a 3rd generation distributed real-time Intrusion Detection System. It supports a large number of high performance multi-threaded sensors and uses sophisticated data-reduction to minimise the number of alerts that are raised. **K2-Defender** enables security analysts to perform alert analysis through historical correlation of the data collected by the sensors and it offers an automated drag & drop facility for report generation.

One potential application of **K2-Defender** is detecting and responding to Distributed Denial of Service (DDoS) attacks.

K2-Defender and DDos attacks

K2-Defender can be used to detect a DDoS attack in progress and can help to track down its perpetrator. Attack detection makes use of real-time statistical tools built into the sensors whereas the forensics would use the historical capabilities of the system. These scenarios will be discussed separately in the following sections.

We shall focus on ISP and *single site* attacks which are the most frequent DDoS attacks.

DDoS detection

The simplest form of DDoS detection, especially in the case of a single site, is that the site stops responding as soon as it gets overloaded and as a result it becomes unable to manage incoming traffic.

Using **K2-Defender** with multiple real-time sensors deployed on the incoming links[8], the DDoS attack would be seen as a surge in bandwidth usage detected at the sensors and displayed through the real-time statistical monitoring tools. This would raise in effect a series of statistical alerts, each of which would be tagged with the sensor recording it. Additional rules linking statistical alerts to specific hosts would also raise an alert on the increased usage of a particular host under attack. In practice this would allow the following:

- Location of the entry point of the DDoS attack because only the relevant sensor would trigger (assuming a multi-homed network).

5.A similar effect can be achieved by propagating bad routing information via BGP.
6.As at the date of wiring.
7.For example, www.mit.edu into 18.181.0.31.
8.If possible one would want to deploy "early-warning" sensors at the upstream provider on a private VLAN.

- Discovery of the protocol being used for the DDoS attack.
- Discovery of the target of the DDoS attack.
- Location of the networks causing the majority of the incoming traffic.

The above would all be visible from the Security Analyst Console.

DDoS reaction

Reacting to DDoS attacks often requires trawling through a mixture of logs, often with non-synchronised timestamps, as well as trial and error procedures. **K2-Defender** can alleviate some of these tasks thanks to the information which is collected by the sensors and can be reviewed by the security analyst.

The collected information enables **K2-Defender**'s users to act pro-actively and contact their upstream ISP to ask for specific filtering and actions. This is the opposite to what currently happens where an ISP would normally notify a client and, in extreme cases, drop all traffic to and from the client.

Another approach, often frowned upon, would be to bring on-line more bandwidth. This often is fruitless as the perpetrator of the attack normally has more bandwidth than can be brought on-line quickly.

One could also link **K2-Defender** 'reactions' to the BGP routing system and systematically drop route announcements from specific networks which contain the attacking machines. If we consider large cable modem providers in the USA then often these consist of a single network prefix covering the equivalent of an old 'Class A' network, that is to say up to 16 million hosts. Because of the size and relative predominance of home systems on these networks they are often the preferred candidates for providing zombies to a DDoS attack. Isolating that prefix may isolate the attack.

The key to using **K2-Defender** is that sensors can be re-focused at run-time to provide additional information which might be essential, for example in as in the case of dropping routes asking for statistics on packets from cable network prefixes. Having this information available without having to deploy specific hardware or software beyond the intrusion detection sensors can make a real difference in the reaction phase.Once the immediate problem has been dealt with, often with the result of having to isolate the target site, a forensic process can begin.

DDoS tracking

Tracking DDoS attacks is a very complicated process which has not really been codified. A number of techniques are available, but almost invariably they require cooperation from a number of ISPs in different countries. Most tracking is done while the DDoS attack is actually taking place, following flows of traffic and trying to isolate them while at the same time making sure that networks are not isolated excessively[9].

One of the major issues is the lack of data. Very few targets run Intrusion Detection Systems and even fewer collect statistical information correlated with the packets. It then becomes a painstaking process of back-tracking packets one-by-one to try to discover a system on which the *zombie* process has not been removed and from there, perhaps, finding the 'master'.

The data stored in **K2-Defender** can be used while tracking DDoS. Very often, the perpetrator of the attack would have performed some reconnaissance. The **K2-Defender** historical search feature can assist in the process of gathering evidence pointing to a particular address or set of addresses[10]. It is possible that a slow port-by-port scan has taken place during reconnaissance, often over an extended period of time to avoid detection. These scans can be discovered with the aid of the **K2-Defender** surveyors which trawl the database searching for precisely these patterns.

This brings us back to the reconnaissance part of a DDoS attack which is often overlooked. The motives behind a DDoS attack can range from plain 'fun' to disgruntled employees, and even to industrial sabotage in some extreme cases. **K2-Defender**'s ability to trawl through large archives of data and to search through it efficiently, allows a thorough investigation back in time to determine the possible origins of a DDoS event. It is not rare for a DDoS attack to be the

9. This dilemma is not simple to solve. ISPs are left with the choice of balancing their core network and a customer under attack with the reachability of their other customers. Dropping a network the size of @Home can cause many repercussions.
10. It is very rare for the perpetrator of a DDoS attack to use their own system as part of the attack.

outcome of the inability to break into a web site in which case this activity would have been recorded as a breaches of the 'whitelists[11]' defined for the web site.

Earning the ISP's trust and prosecuting

The relationship between a client and the ISP can often be difficult, in particular with respect to security. Very often an ISP ignores security alerts from customers because of the high false-positive rate which they observe. Using **K2-Defender** to enforce a security policy by aggregating both network-based and host-based intrusion detection at an enterprise-wide level, can help improve the relationship with the ISP, which in a situation like a DDoS attack might require strong actions and difficult decisions. Being able to support security alerts with hard data and statistical information is essential and this is one of the roles that K2-Defender can undertake.

Prosecuting in DDoS cases is often difficult and it can also be impossible especially when prosecuting in the absence of data. K2-Defender has been designed with specific support for prosecution, including GPS-locked time-stamps on data from the moment it is taken from the wire to trusted 3rd-party off-site public-key data encryption. For all the data stored on the database it is always possible to reconstruct a time sequence of events and prove that the data has not been tampered with, after being collected from the network.

11.'Whitelisting' is a process by which instead of attempting to wire rules for all malicious traffic, explicit rules for all *known good* traffic are recorded and everything else logged, K2-Defender has supported this view of rule-based Intrusion Detection since its initial design.

Intrusion Detection System

For further information contact:
K2 Defender Ltd.
Garden House
Cloisters Business Centre
8, Battersea Park Road
London SW8 4BG, UK
http://www.k2defender.com

P/N: K2DWHPA001

Part Number: K2-Def-WHPA-DDoS-001