



The Art of Mitigating DDoS Attacks

For a number of years competitive rivalries, political disputes and organized crime have been spilling over onto the internet in the form of Distributed Denial of Service (DDoS) attacks. These are attempts by hostile parties to disrupt online services by overwhelming infrastructure required for communications. This is done with floods of traffic that IPS and firewalls cannot block. This paper provides some essential background about attacks, explains why traditional security equipment cannot mitigate them, and describes the ONLY method of ensuring uninterrupted service during such attacks.

Service Provision

The Internet is used to provide a wide range of services. "Classic" traffic such as FTP, SMTP and NTP coexists with more recent World Wide Web protocol-based traffic, and even more recent protocols such as Skype or SIP. Common to all is the requirement for reasonably loss-free and low latency communication between service consumers and service providers. For example, disruption of any of the following can prevent customers accessing a business' web site:

- (1) Routers in the IDC that handle the business' traffic
- (2) The business' downlink (to their premises, or their rack co-located at the IDC)
- (3) The business's firewall or IDS/IPS
- (4) The web server's network card
- (5) The web server's OS and TCP/IP stack
- (6) The web server process itself (e.g. Apache)
- (7) Any database backend for dynamically generating the web page

Service Disruption

DDoS attacks disrupt services by overwhelming infrastructure with floods of traffic. The attacks succeed because overwhelmed networking infrastructure tends to drop packets indiscriminately (or crash, causing all packets to be dropped). With the use of botnets attackers can easily generate enough traffic to cause packet loss rates at which provision of satisfactory service is impossible. Moreover, they can overwhelm different parts of the infrastructure with different aspects of traffic. For example:

- The business' downlink (2) can handle a certain bandwidth in bits-per-second. Traffic is dropped if that bandwidth limit is exceeded, regardless of whether the traffic consists of many small packets or few large ones.
- The server's OS and TCP/IP stack (5) is typically limited by the number of packets it can handle each second, meaning some packets will be lost once more than a certain number arrive per second, regardless of packet sizes (and the bits-per-second bandwidth they take up).
- The web server process (6) is typically limited by number of connection attempts. It will drop connection attempts above its' imposed per second limit, regardless of overall traffic levels.

Service Availability

DDoS mitigation is about ensuring legitimate traffic is not crowded out by attack traffic

This is done through a process of selectively dropping traffic, which requires the ability to recognize when and how much traffic needs to be dropped to protect the infrastructure, while avoiding dropping legitimate traffic in the process i.e. it requires knowledge of when to initiate protection, how much traffic to drop, and which traffic to drop.

To protect an entire network the mitigation needs to occur at the network edge. Furthermore, as the first line of defense protecting downstream infrastructure, the mitigation needs to be done at full line-rate to avoid becoming a bottleneck that results in the loss of legitimate traffic.

These principles can perhaps be better appreciated by considering that DDoS mitigation will NOT properly prevent service disruption if:

- Infrastructure is overwhelmed because protection is not activated or too little traffic dropped
- Too much and/or the wrong traffic is dropped i.e. collateral damage and/or false positives
- The mitigation equipment is itself overwhelmed because it cannot process traffic fast enough

Knowing When and How Much Traffic to Drop

For the DDoS mitigation device to know when and how much traffic to drop, rate limits need to be imposed to define acceptable traffic rates and act as thresholds that trigger filtering when they are exceeded. Since each part of the network ('entity') can handle different volumes of traffic, each requires its own limit. And since each entity can be overwhelmed by different traffic aspects, the limits need to be configurable for bits-per-second, packets-per-second, and connections-per-second. Limits of this type are the only way to ensure that traffic is dropped when necessary, and only as much as necessary is dropped. Without limits, attackers can send targeted floods that overwhelm parts of the network needed for a particular service.

Knowing Which Traffic to Drop

With this ability to monitor traffic rates to each entity in the network, a mitigation device has vital information about which traffic to drop. It knows which entity is under attack, and can therefore confine rate-limiting to only traffic destined for this entity. As well as narrowing down the pool of potential attack traffic, this prevents non-targeted services being affected by the attack or the rate-limiting, and in the case of service provider installations prevents an attack against one customer negatively affecting other customers' services.

The remaining task for mitigation is to determine which of the traffic destined for the attacked entity is legitimate. This task is difficult because attackers use networks of bots to send any traffic they like, including traffic that is outwardly indistinguishable from legitimate traffic. They can constantly morph attacks by switching between botnets and altering the composition of attack traffic. This makes futile any effort to predict the type of traffic that attackers will send.

Under these conditions the only advantage a mitigation device can have over attackers is greater knowledge of a specific network's legitimate users, derived from having observed the traffic in that network for an extended period. Having developed an intimate knowledge of the particular network's legitimate traffic, it can identify and prioritize this traffic in the midst of floods of attack or excess traffic. The composition of the attack traffic is unimportant since the mitigation device is only concerned with passing legitimate traffic, meaning the DDoS-mitigation device is impermeable to any conceivable type of attack that the cleverest attacker can conceive.

How Security Appliances Fare

Traditional anomaly detection techniques cannot mitigate DDoS attacks

The required DDoS mitigation functionality described above differs from traditional “anomaly detection”, and is beyond the means of devices that provide such anomaly detection, including firewalls, IPS, and other security appliances that claim to protect against DDoS attacks. This includes the range of currently available dedicated DDoS-mitigation appliances, which draw heavily on intrusion prevention techniques and architectures.

All such devices lack a method of protecting different network segments in sufficient granularity, and are therefore susceptible to all the previously described drawbacks of imprecise filtering and failing to protect against targeted attacks.

They also typically fail to process traffic at line-rate, due to excessive processing overhead incurred from performing unrelated packet-inspection duties and using intensive algorithms, with the result that attackers need only send large numbers of small packets to cause the devices to arbitrarily drop traffic before having the opportunity to determine which packets to pass.

Lastly, current detection techniques are largely ineffective. They assume that traffic from bots is measurably different from legitimate traffic, and can be identified by administering a series of tests designed to detect and discard particular types of anomalous packets. Techniques include:

- Signature detection, to discard traffic that matches known attack patterns. Pre-programmed signatures leave attackers free to tweak attacks to avoid matching the signatures, while signatures generated on the fly take effect after a network is already flooded and their imprecision results in legitimate traffic being discarded.
- Protocol anomaly detection, to discard traffic that does not conform to the traffic models described in RFCs and vendor documents. Attackers can simply avoid detection by ensuring that their traffic conforms. Furthermore, false positives are inevitable because much internet legitimate traffic does not strictly conform to the RFCs.
- TCP flow control, to discard traffic from sources that fail to comply with requests to reduce send-rates. This simply forces attackers to send less traffic from each of a larger number of bots.
- Anti-spoofing, to discard traffic with falsified (‘spoofed’) source IP addresses. In theory this forces attackers to use bots’ real IP addresses, but in practice the defenses can handle too few concurrent connections, allowing attackers to overwhelm the anti-spoofing mechanisms with connection attempts.

In practice these techniques do not so much identify attack traffic as restrict attackers’ options, forcing them to send traffic that appears more “normal”. As attackers craft traffic that eludes detection, the defences need to widen their definition of what constitutes attack traffic. By making the tests more punitive, more collateral damage is inflicted on legitimate traffic (false positives). Where attack traffic closely resembles a network’s regular traffic, such techniques can do no better than arbitrarily discard traffic.

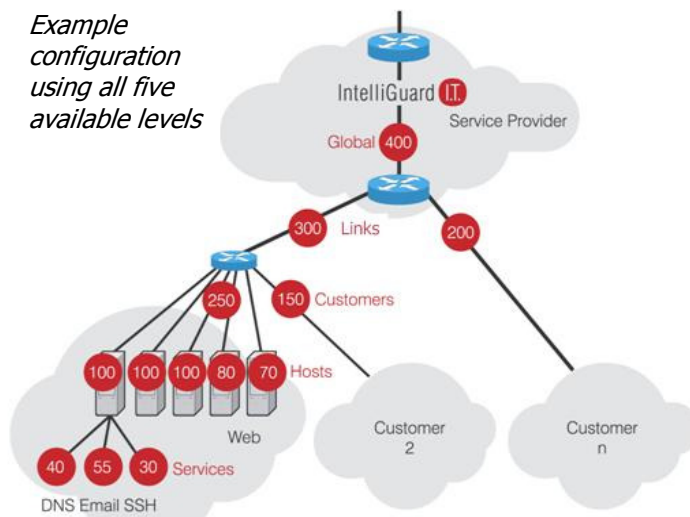
In summary, existing products that claim DDoS defence simply don’t protect against other than the simplest and most naïve attacks and are themselves open to attack.

IntelliGuard's DDoS Protection System (DPS)

IntelliGuard DPS provides a fundamentally different and superior approach.

The first key difference is that IntelliGuard's focus is on letting legitimate traffic through, rather than keeping attack traffic out. This approach maximises the number of legitimate clients able to access protected services at every point in time, thus maximizing the quality of service provided. It was designed from the ground up to handle the worst possible attack scenarios without performance degradation, protect each part of a network, and distinguish legitimate from attack traffic with unprecedented accuracy.

The second key difference is that a single IntelliGuard DPS deployed at the network edge protects every component of a complex network, from large links all the way down to individual services running on a single machine. This is achieved via a system of hierarchical limits and guarantees configured for each entity.



NB: circled numbers represent bandwidth limits applied to each entity

Limits set the traffic thresholds that trigger protection. This protects low bandwidth services from attack, and improves filtering precision by restricting filtering to only the attacked entity's traffic.

Guarantees set the minimum traffic rates available to an entity regardless of network conditions. This provides "fair" protection by allocating network capacity to where it's needed, and protecting non-attacked entities from potential congestion resulting from attacks on a part of the network or server they belong to.

Where traffic to a Protected Entity exceeds a limit, some traffic destined for it is dropped, starting with traffic from senders with the lowest "legitimacy classifications". These classifications are formed by analyzing the behavior of clients that interact with the protected network on a range of attributes that can indicate legitimacy and illegitimacy. Every client's traffic is ranked relative to every other client's, with trust established through long-standing patterns of legitimate interaction reflected in higher rankings.

The third key difference is IntelliGuard's Learn – Rank – Protect technology. With IntelliGuard's *Learn-Rank-Protect* approach, DDoS is treated as a resource allocation issue. When under attack, as much traffic as possible from trusted clients is passed. Unlike other approaches, traffic does not 'pass' or 'fail' any specific tests of legitimacy, and therefore cannot be crafted by attackers to elude detection. As a result, every conceivable type of DDoS attack is mitigated with the highest confidence that all legitimate traffic will be passed.

IntelliGuard builds network security appliances that manage network traffic to ensure business continuity and quality of service for on-line business's and service providers.

For further details, please see www.intelliguardit.net