Arbor Threat Mitigation System (TMS)

How It Works and Where It Lives in the Network

TMS or Mitigation Capacity

Arbor Sightline and Arbor Edge Defense (AED) detect attacks and do have a set of actions which could be set up to protect the final customer, for example, BGP FlowSpec and AED filtering can be used if all packets have the same IP dest and Port dest. However, for complex and or multi vector attacks it becomes mandatory to let an intelligent device analyze all the traffic to be able to filter the malicious one and deliver clean traffic to the end user.

TMS Placement

An Arbor Threat Mitigation System (TMS) deployment can scale from 1Gbps to 40Tbps of capacity, meeting the needs of all network operators. In order to protect both end customers and service provider infrastructure, Arbor TMS can operate in a variety of environments and architectures:

- **Boundary protection:** TMS can be deployed on Each Internet Point of Presence (PoP), where external links from Peers or Transit providers aggregate, will require one or more TMS. Thus, mitigation capacity is put close to the source of the attacks (first line of defense). TMS can also be used to prevent any malicious traffic coming from infected customers to quit the network (last line of defense).
- **Distributed all over the network:** to ensure more efficiency, in addition of protecting Internet PoPs, TMS can also be spread over the backbone to be able to stop the internal attacks and limit the impact. TMS can be installed at the aggregation provider edge to control enterprise fiber infrastructure, at the peering point with Data Centers where partners and internal division installed servers and platforms (single sign on, IT services, Customer care systems). Other usual locations for TMS are the interconnection between mobile and fixed backbone and the wholesale trunks. In that configuration, both North/South and East/West attacks can be prevented.
- **Centralized architecture:** Network operators can opt to centralize their scrubbing resources. The assumption here is that there is sufficient backbone capacity to transport attack traffic to a central site or sites, before mitigation occurs. With this approach all kind of attacks can still be monitored and mitigated (mobile, trunk, DataCenter...) however the solution is more scalable since it uses common resources and adding capacity is easier, since there is only one location to upgrade and the new capacity will benefits to all protection types (North/South and East/West).

TMS Redirection Mechanisms

Individual network operator environments and designs are unique and a range of redirection and re-injection solutions have been developed by Arbord. All of them supports inbound mitigation and BGP FlowSpec and MPLS LDP (Centralized TMS) can add outbound mitigation.

TMS deployment can scale from 1Gbps to 40Tbps of capacity, meeting the needs of all network operators.

Inbound Mitigation: BGP FlowSpec

Once a DDoS attack is detected Arbor Sightline will advertise a BGP FlowSpec rule to the Routers impacted by the attack. This FlowSpec rule can specify source and destination traffic parameters, redirected to the next hop IP of the Arbor TMS diversion interface. On the INTERNET ingress interface (INTERNET), the Router will interrogate its FlowSpec rule database, and any matching traffic will be diverted to the Arbor TMS.

The Arbor TMS will process the attack traffic and any clean or legitimate traffic will be re-injected to the network. Since there is no FlowSpec processing on the TMS re-injection interface, the Router will interrogate its routing table and forward the traffic towards the intended destination.

Inbound Mitigation: Clean VRF

With that method, Arbor Sightline will advertise a new BGP route to the Router. This BGP route will specify the destination IP of the system under attack together with the next hop of the Arbor TMS diversion interface. The advertised BGP route will be tagged with a specific BGP community, to mark it as a "poison-route". On the INTERNET ingress interface (INTERNET), the Router will interrogate its GRT database and any traffic destined to the "attacked" IP will be diverted to the Arbor TMS.

The Arbor TMS will process the attack traffic and any clean or legitimate traffic will be sent back to the "clean-VRF". Since the "clean-VRF" is in essence a copy of the GRT, except for the "poison-route", the traffic will be forwarded to the original next-hop. The Router automatically forwards the traffic as per the GRT, without applying any MPLS labels, since the "clean-VRF" is based on the VRF-Lite construct.

Arbor TMS will process the attack traffic and any clean or legitimate traffic will be re-injected to the network.

Inbound Mitigation: MPLS LSP

Sightline will advertise a new BGP route to the desired PE Routers. The operator can choose if the traffic should be diverted only on the Internet-PE's (Inbound) or the Customer-PEs (Outbound), or a combination of both (Inbound + Outbound).This BGP route will specify the destination IP of the system under attack with the next hop set to the Diversion IP of the Arbor TMS. On any ingress interface the PE Router will interrogate its GRT database and any traffic destined to the "attacked" IP will be diverted to the Arbor TMS. When the MPLS-PE Router does a recursive lookup for the TMS IP it finds the allocated MPLS LSP. The Router then labels and switches the diverted traffic across the MPLS backbone to reach the central TMS.

The Arbor TMS will process the attack traffic and any clean or legitimate traffic will be sent back to the TMS-PE, and into the Global Routing Table (GRT). Since the TMS-PE has no knowledge of the diverted host (/32) route (it doesn't get any advertisements from Arbor Sightline), the traffic will be forwarded to the original Next-Hop via the allocated LSP.

Countermeasures and User Interface

Arbor TMS is capable of cleaning traffic using predefined countermeasures. The administration is done from Sightline UI where TMS setup can be adapted. On Sightline GUI there is a display associated to each TMS mitigation. Thus, Sightline interacts with TMS to:

- Get info about the traffic (passed/dropped)
- Manage the mitigation
- Get packets sample to get more accurate info about the attack and to build FCAP filters which fits the attack (IPs, Flags, Ports, Protocol, Headers,....).
- Administrate countermeasures and adapt parameters and characteristics.

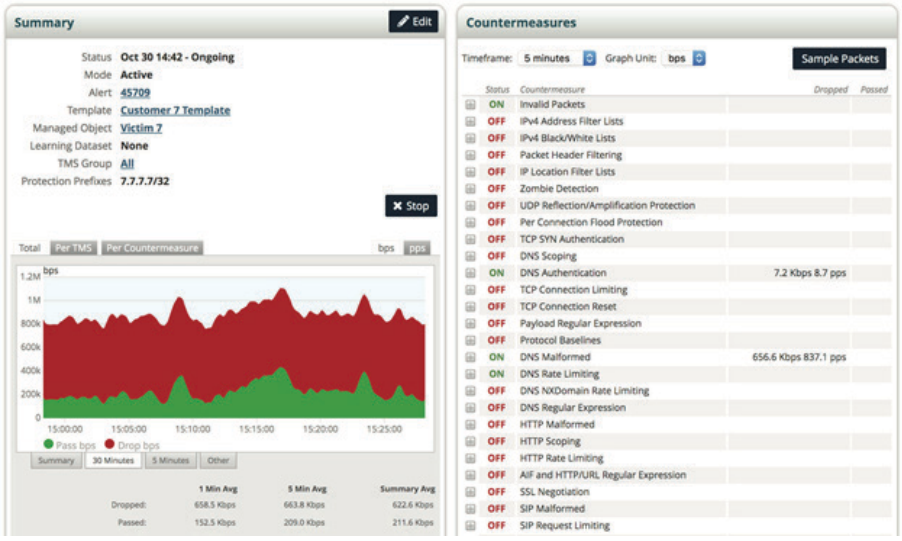

TMS Redundancy

Redundancy and Active/Active Load Balancing can be achieved in a single site by "sharing" a Diversion IP between multiple TMSs, If a single TMS or TMS interface goes offline the routing table entry will be removed automatically by the router and the remaining TMSs will receive the traffic. That's a "Per Site Redundancy".

In a "Network Wide Redundancy" configuration, TMS High-Availability is achieved using standard routing protocols and routing functions within the IGP.

Each TMS "shares" the Diversion IP and the network IGP will ensure that the closest available TMS will receive the attack traffic. The concept of "sharing" IP addresses amongst network devices and services is known as IP Anycast. The "shared" Diversion IP can be assigned to the actual TMS Diversion Interface or deployed via a "loopback".

LEARN MORE

For more information about NETSCOUT Arbor Threat Mitigation System (TMS) visit:

https://www.netscout.com/product/arbor-threat-mitigation-system

NETSCOUT

Corporate Headquarters
NETSCOUT Systems, Inc.
Westford, MA 01886-4105
Phone: +1 978-614-4000
www.netscout.com

Sales Information
Toll Free US: 800-309-4804
(International numbers below)

Product Support
Toll Free US: 888-357-7667
(International numbers below)

NETSCOUT offers sales, support, and services in over 32 countries. Global addresses, and international numbers are listed on the NETSCOUT website at: www.netscout.com/company/contact-us

SECJTF_002_EN-2001 06/2020