White Paper

Cisco ASR 9000 Integrated DDoS Mitigation with Arbor Networks Peakflow TMS

Technical Deep Dive

Last Updated: 3/10/2015

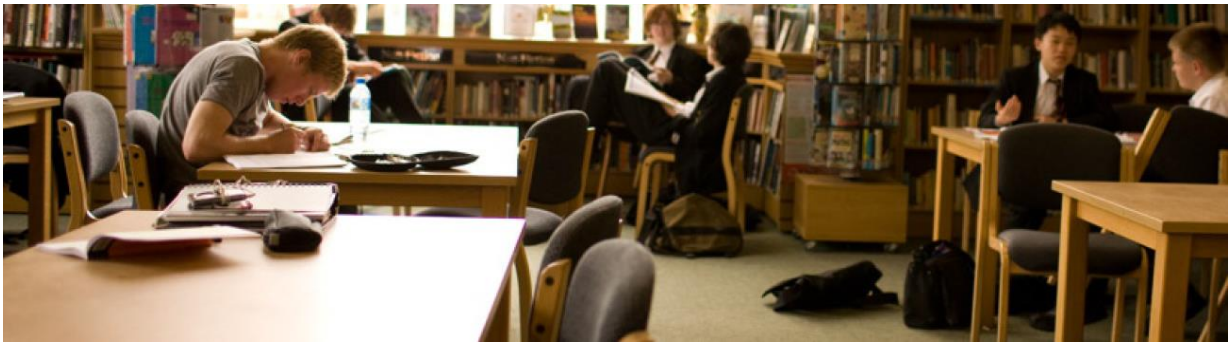

Cisco and Arbor Networks are partnering to deliver industry-leading DDoS mitigation capabilities on the ASR 9000 Series routers. This partnership extends the Peakflow solution by adding the Threat Management System (TMS) DDoS mitigation functionality to the ASR 9000. The TMS will be implemented on the ASR 9000 VSM (Virtualized Services Module) hosted in the ASR 9000 chassis. This solution enables unprecedented scaling of the DDoS mitigation architecture, reduces operational complexity and hardware footprint, and allows network providers to push attack mitigation to the edge of the network which reduces network loads from backhauling attack traffic to centralized scrubbing centers.

Lane Wigley & Nicolas Fevrier – Cisco SP Routing Technical Marketing
w/ Brian Prater and Jorge Escobar - Arbor Networks

Goals

This whitepaper provides an in-depth look into the Peakflow solution with a focus on the new capabilities enabled by integrating the Threat Management System (TMS) DDoS mitigation capability into the ASR 9000.

Executive Summary

Cisco and Arbor have partnered to integrate the Arbor Threat Management System (TMS) into the ASR 9000 via the Virtualized Services Module (VSM). This new solution allows for optimal placement of DDoS mitigation within the network which reduces the need to carry the traffic to scrubbing centers and therefore provides a range of advantages including limiting the scope of attack traffic and simplified routing and operations.

In addition to providing 40 Gbps per VSM of high-touch DDoS scrubbing, the Cisco/Arbor solution will also enable intelligent SDN-driven hardware-accelerated filtering that scales to Tbps rates by distributing knowledge of attacks throughout the network.

Solution Overview

The Arbor Peakflow solution protects customer networks by mitigating undesirable traffic caused by both volumetric and application Distributed Denial of Service attacks. Peakflow accomplishes this by building a baseline for the network based on packet rates per host, high and low packet and bit rate thresholds, and destination-specific automatically-generated profiles. It then recognizes anomalies from this baseline and removes the unwanted traffic while permitting valid traffic. This is implemented via ACL-like or Fingerprint-based filtering, rate limiting, session authentication, and monitoring adherence of traffic to protocol standards.

The high-level operation of the solution is as follows: First, the system will monitor network ingress points via Netflow and BGP to build a baseline for network behavior and traffic patterns. It will then perform ongoing monitoring to detect anomalies and flag them as potential attacks. These potential attacks are presented to network operations via a GUI, email, or SNMP which allows a range of actions to be taken, including initiating a response or marking an event as a false alarm.

The main actions to initiate a response are: a) update the network routing to redirect all traffic for the destination through the TMS-VMS which can remove unwanted traffic and b) clean the traffic as effectively as possible without blocking valid connections.

In partnering with Arbor, the Cisco ASR 9000 will integrate the traffic cleaning functionality into the router and scale its performance operation far beyond what is possible with appliances alone or with competing offerings.

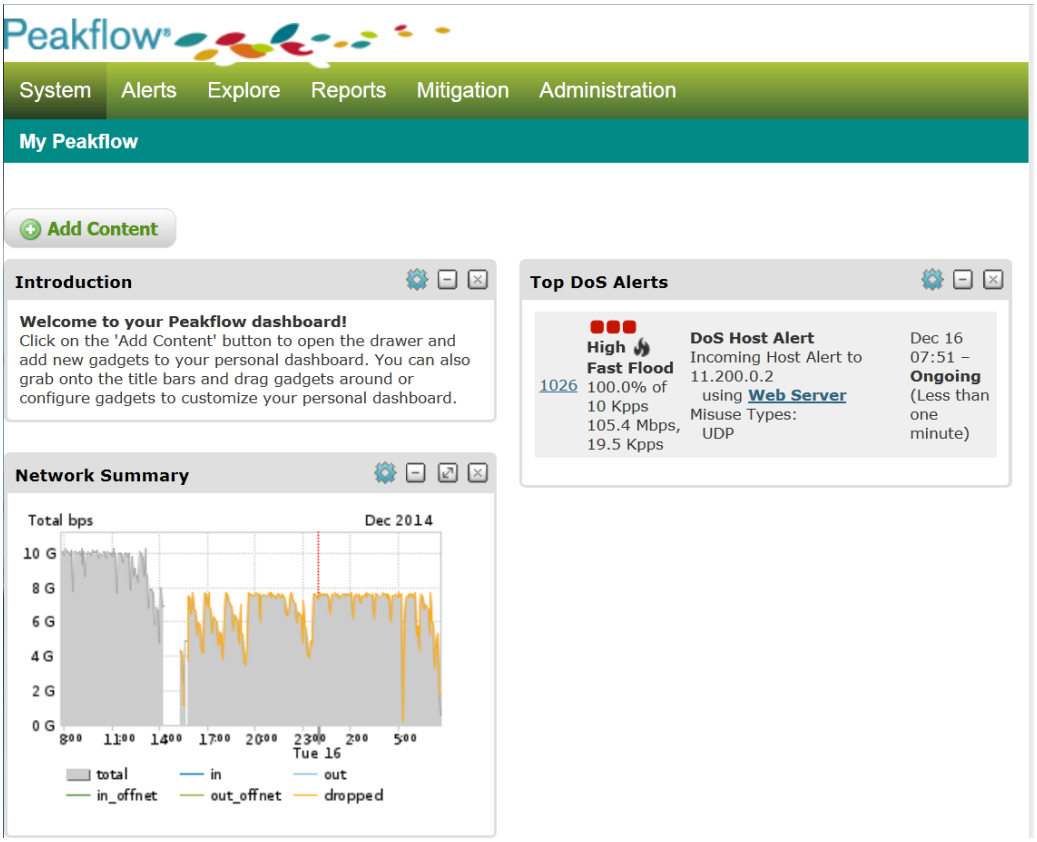

Figure 1 - Peakflow Graphical Interface

Key components

The Peakflow solution comprises a number of functions as well as a set of hardware devices that implement those functions. “Peakflow” is the overall Arbor solution for both network analytics and DDoS detection and mitigation. “Peakflow SP” refers to the control components for Peakflow such as monitoring the network, detecting attacks, and coordinating an attack response. These run on SP appliances (required for leaders) or in VMs. “Peakflow Threat Management System (TMS)” or “Peakflow SP TMS” is the data plane component to remove DDoS attacks. TMS is the component that runs on the ASR 9000 VSM. Peakflow SP runs on a combination of appliances from Arbor (such as the SP 6000) and virtual machines with additional instances allowing scale of the GUI, collector capabilities, and number of Managed Objects to be monitored.

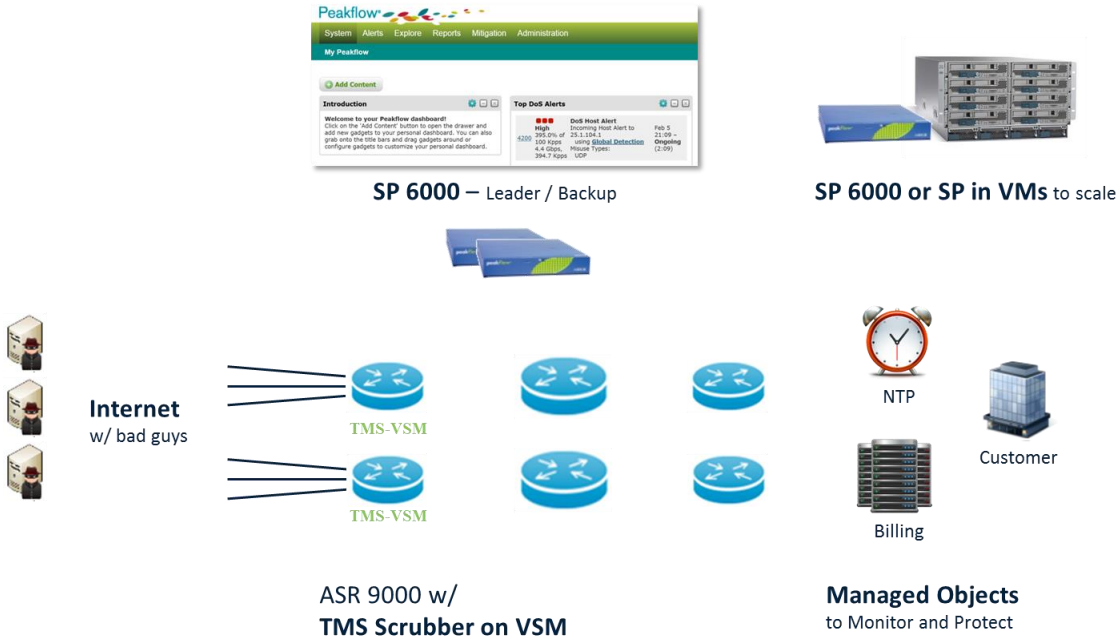

Figure 2 - Key Peakflow Components

Peakflow SP provides the following functions:

- Leads the overall system and manages communication among components
- Presents the GUI to network operations staff
- Receives Netflow and routing information from the routers
- Analyzes the data to detect anomalies and generate alerts
- Creates diversion and reinjection paths via BGP or BGP Flowspec
- Determines appropriate countermeasures (with user input) and program countermeasures into TMS
- Receives statistics and packet samples from the TMS and display via the GUI
- Manages system licenses (from Leader appliances)

Peakflow SP is instantiated on a combination of SP appliances and user-provided virtual machines. A physical appliance (e.g., SP 6000) is required to serve as the leader and (optionally) backup leader. Additional scale occurs via appliances or VMs as well as software licenses. Earlier versions of Peakflow hosted the SP functions on a set of appliances known as Collector Platform (CP), Flow Sensor (FS), Portal Interface (PI), and Business Intelligence (BI) and scaled via adding appliances. These components may still be utilized but aren't presented as part of the reference solution.

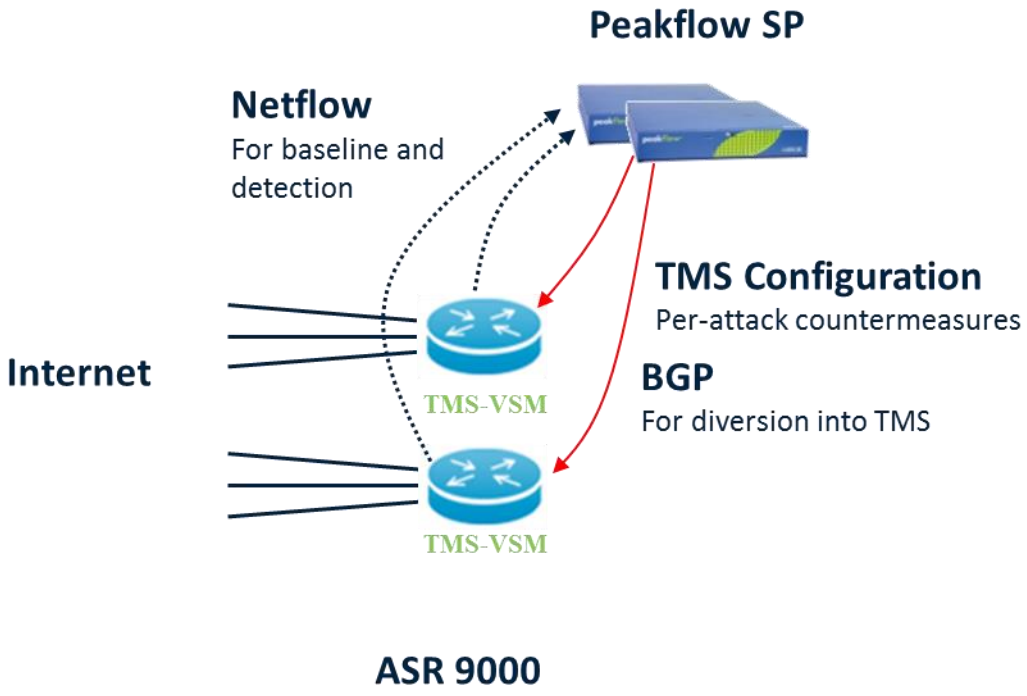

Figure 3 - SP to TMS Communication

Peakflow TMS (or SP TMS) provides the following functions:

- Receives countermeasure programming from Peakflow SP
- Implements countermeasures to remove attack traffic
- Forwards validated traffic to the proper destination
- Sends statistics to Peakflow SP
- Captures packet samples and exports them to Peakflow SP

Peakflow TMS is instantiated in the ASR 9000 Virtualized Services Module. It can co-exist with Arbor TMS appliances which in the same deployment.

As the traffic scrubbing device, the TMS capacity must scale with the volume of mitigation required. Therefore, the hardware must be scoped for the maximum predicted attacks, which will most likely change dramatically over time. Traditionally, the TMS has been implemented on an appliance from Arbor. This solution brings this functionality into the ASR 9000, which allows new scalability, performance, and design options. The following graphic a reference Peakflow topology with TMS integrated into the ASR 9000 and blocking traffic at the peering and customer edge.

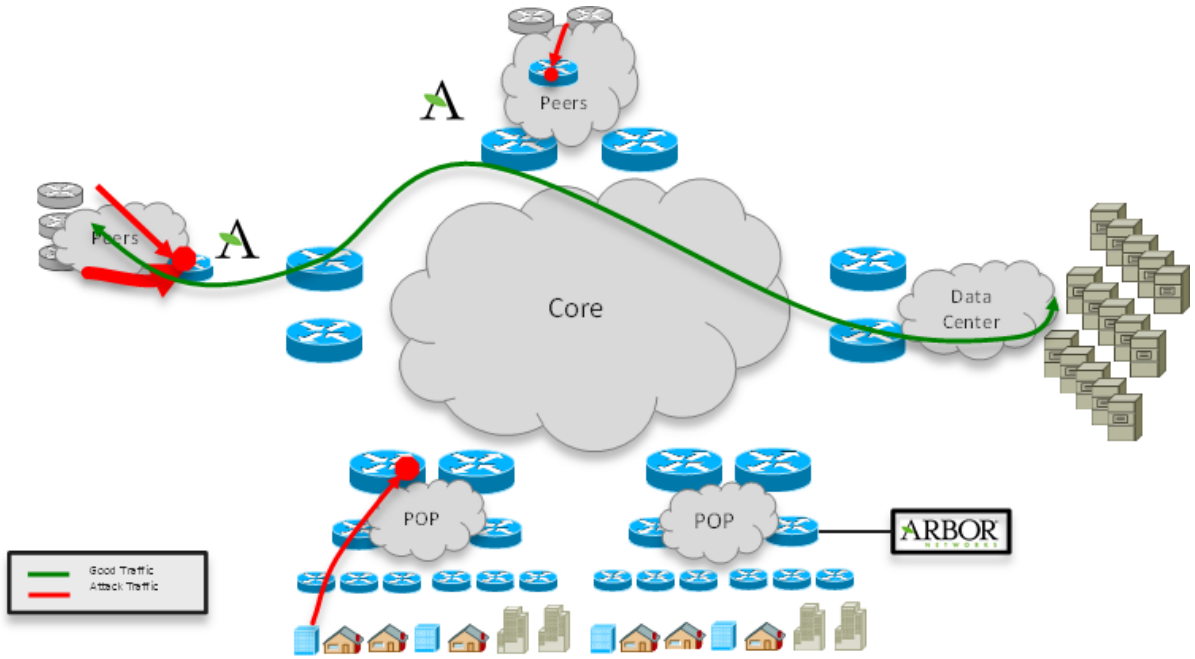

Figure 4 - Traffic flow with distributed TMS

NetFlow for DDoS Attack Detection

Netflow, SNMP, and BGP are the key tools to monitor for DDoS attacks or other anomalies. SNMP can provide interface status and settings. BGP can provide information on network boundaries (based on Managed Object configuration – covered later) and routing status. NetFlow provides the traffic statistics from which the Arbor

algorithms can build baselines and then detect attacks based on the known baseline, traffic rates, and attack fingerprints.

Netflow collection is performed by Peakflow SP. NetFlow should be enabled inbound on all external interfaces to where potential attacks may enter the network. It may also be enabled in the core to provide additional information to characterize an attack. A range of sampling rates from 1:1000 – 1:10000 will provide effective attack detection. The cumulative size of the sample influences the sampling error (what is predicted from the sample vs. the results of actually seeing every packet) and thus the detection time. Therefore, a higher sampling rate or a longer sampling time can both achieve increased accuracy. That said, attacks will still be detected quickly even with lower sampling rates as the error is very low with even 1000 samples (samples received, not the sampling rate) and only improves about 2% when going beyond 10000 samples.

For a given sampling error, attack detection speed will depend on the type of attack. Volumetric attacks can be recognized more quickly as there are many packets to sample. Application attacks and attacks against lower bandwidth servers represent a small proportion of overall traffic so they may not always be recognized at the edge unless the match a known fingerprint. For this reason, detection and mitigation of application layer attacks should happen after volumetric traffic has been reduced as this results in a higher percentage of application attack packets in the sample near the data center or application host.

Note that while attack detection traditionally takes 1-2 minutes, Peakflow can now detect high volume attacks in a matter of seconds when large attacks occur..This feature is known as Fast Flood Detection.

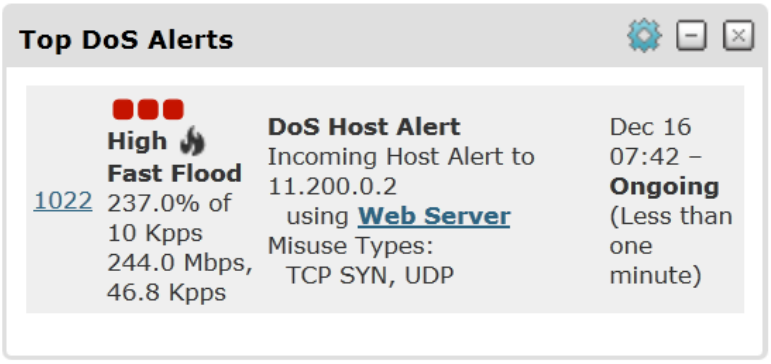

Figure 5 - Sample Alert

Network Baseline for Anomaly Detection

In addition to recognizing attacks via rates and fingerprints, the Peakflow SP system can build a traffic baseline for each Managed Object (protected resource) and recognize increased activity. This may be due to an attack or increased traffic due to a special event so an operator should investigate before initiating mitigation. Three baselines are created: a continuous 30 second baseline, a time-of-day baseline, and a day-of-week baseline. Alerts are generated based on the variance of the current traffic from the baselines. The tolerance before generating an alert is configurable.

Types of Attacks – Volumetric & Application

a. At a high level, there are two main classes of DDoS of attacks. Volumetric attacks overwhelm network or server resources with a high volume of traffic. Application attacks exploit vulnerabilities specific to an operating system or server software. A volumetric attack could be as simple as sending lots of UDP traffic to a single destination in order to congest a network link or server's packet capacity. The NTP monlist attack is a recent example of a volumetric attack capable of significant amplification. A next step in attack complexity would be sending a high volume of TCP traffic with the SYN flag set. This traffic could overwhelm the network capacity, but it will also consume server resources such as the TCP session table. The next level of complexity could include sending a large number of invalid HTTP GET requests to a web server or malformed requests to a DNS server. From these attacks, it is clear that many attacks involve both volumetric and application layer components. Finally, an example of a purely application layer attack is Slowloris (http://en.wikipedia.org/wiki/Slowloris) which builds a number of valid connections to a server and then keeps them open indefinitely in order to reach the server's maximum connection capacity. This attack can enable a single attacking PC to take down a server. It is a good example of the need to detect and mitigate application attacks close to the server under attack or on the server itself. The low volume of traffic characteristic of Slowloris often would not be detected by NetFlow sampling at a peering point. While TMS countermeasures may block application attacks while mitigating a volumetric attack, application layer attacks will often require dedicated mitigation services such as those provided by Arbor Pravail. In addition to individual attacks having both volumetric and application components, multiple different attacks often occur together as part of a larger attack, and Arbor Cloud-Signaling provides an infrastructure that facilitates both local and upstream DDoS mitigation in an automated and real-time manner. Cloud Signaling is an efficient and integrated system coordinating DDoS mitigations from the customer premise to the service provider cloud."

Design Considerations

Protected Resources & Attack Vectors

A key part of designing the TMS deployment is identifying the potential victims and attack sources. The location of the attackers and protected resources (potential victims) will influence the detection and mitigation design. Common sources of attacks are traffic from the Internet via peering or exchanges (from IDCs or from compromised end users) and from other customers of the SP itself. The potential victims may be the link to a customer, a resource inside the customer's network, infrastructure links, or a server within an IDC.

Customer to customer DDoS comprises attacks from one of the SP's customers to another. These attacks require detection at the customer facing network boundary rather than the peering boundary as seen in the more common Internet case. The proximity of the TMS will determine the most effective way to direct traffic for cleaning. One factor that is unique to this scenario is that it is simpler for the SP to take action local to the attacker, which may mean increased reliance on other (non-TMS) features such as ACLs, uRFP, and Black Holes since the attack source is localized.

Protecting Data Centers and infrastructure servers (e.g., DNS, billing, authentication, and Route Reflectors) is critically important to an SP or large Enterprise's ability to minimize the impact of a volumetric attack that will often accompany an application layer attack. Application attacks must be mitigated with L4 through L7 aware countermeasures including TMS as well as other tools in the network and on the servers. Different countermeasures and thresholds are often configured to protect different services. Again, the recommended best practice is to mitigate volumetric attacks as early as possible and to mitigate application attacks closer to the Data Center as those attacks may not be large enough to trigger detection elsewhere.

Key Peakflow Design Goals

As a starting point for the design process, it is important to develop a clear view of which resources are being protected and the attack sources that must be managed. Once that is understood, the key design decisions of a successful Peakflow SP design for DDoS mitigation are scalability, redundancy, and routing for traffic diversion and reinjection.

In addition to traffic scrubbing capability, scalability involves the ability to provide sufficient capacity to: collect and analyze data from the network (via syslog, NetFlow, and routing), scale the number of protected destinations (Managed Objects), and providing user access via the GUI. In recent years, the volume of attack traffic is increasing exponentially. For example, the March 2013 attacks against Spamhaus were significantly larger (300 Gbps) than any previous attack. Mitigating such an attack requires a layered approach and significant mitigation capacity. It also involves multiple types of mitigation including ACLs, uRPF, traditional black holes, TMS, and more.

Redundancy involves many aspects of network design. In this paper, the focus is on redundancy of the TMS-VSM devices and how they can be managed via TMS Clusters and Groups.

Traffic Diversion is the process of steering all traffic (good and bad) for the devices under attack to the TMS and then delivering just the good traffic to the destination. The key goals for this aspect of design are minimizing the ability of the attack traffic to congest the network and ensuring that routing loops do not occur. The risk of routing loops comes from the fact that traffic must be sent to two different destinations (first to the TMS and then to the final destination) which implies that they cannot be forwarded based on the same forwarding table. The diversion of traffic to the TMS is sometimes called off-ramp as the packets are taken off their normal path. The reinjection into the network for delivery to the final destination may be called onramp, although diversion and reinjection are the more commonly used terms.

Design Consideration 1: Scalability

Initially, scalability would seem to be simply a function of the amount of scrubbing capacity, and that is a key element. On closer inspection, additional factors become more evident. In the control plane part of the solution, the number of routers to monitor (Netflow collector scale), number of users, as well as the number of Managed Objects are scaled via additional SP devices (appliances or VMs) and software licenses.

For TMS scale, the location of the TMS impacts scalability in two ways. First, if all the capacity is centralized, it can easily be shared among all devices needing protection and therefore provide maximum traffic scrubbing capacity. On the other hand, sending attack traffic to a centralized scrubbing center requires carrying the traffic over the network and potentially causing congestion which may just spread the problem to other network resources.

A distributed model has the opposite strengths and weaknesses. Traffic is cleaned as close to the network ingress as possible, but more total capacity is likely required to handle cases where a large attack has a common ingress point. For example, a central scrubbing center may be able to handle 100Gbps of traffic coming from various ingress points. If all that traffic entered through one part of the network, and the total TMS capacity was split among 4 different ingress points, the TMS for that ingress point would not have the capacity to handle the full attack load on its own. The requirement to add more TMS capacity in this case would likely be offset by reducing the need for extra link capacity to backhaul the attack to a central location.

The optimal design combines both approaches: TMS capacity to handle most attacks can be deployed on the ASR 9000 routers at or near network ingress points (customer and peering edges). If additional capacity is required to mitigate an attack, the traffic can be backhauled to other distributed TMS or to a regional scrubbing center.

Since the mitigation happens per diverted prefix and is often based on finding the nearest TMS via routing metrics, a single victim's mitigation can be spread among multiple TMS via the metric of the diversion route created on the CP GUI. The use of anycast to find the nearest TMS can simplify this process (covered in more detail in the diversion/re-injnection section).

Scalability via multiple VSMs within an ASR 9000 is another likely path to increase capacity without introducing additional complexity. Once traffic reaches the TMS router, the diversion next hop can be statically routed into multiple VSMs which will result in balanced load sharing among the VSMs. This method also provides N:1 redundancy within the chassis similar to a link bundle. Note that multiple VSMs per chassis is a post-FCS feature.

Load sharing within an ASR 9000 occurs when traffic is routed into multiple VSMs and also among the CPU cores within the VSM. When multiple TMS-VSMs are available, the ingress line card picks the VSM, and the VSM picks from one of the 40 cores on the 4 CPUs. The selection of the VSM is based on the default IOS XR 5-tuple load sharing. For IPv4/IPv6 packets, the selection is based on Source IP, Destination IP, Source port (TCP/UDP only), Destination port (TCP/UDP only), and Router ID). MPLS load sharing depends on the size of the label stack and payload type.

Once a VSM is selected, the selection of a CPU core is based on src/dest IP only in order to ensure that traffic from a flow goes through the same core (which contains stated needed to verify the flow).

Groups and clusters are tools to simplify the scale of the TMS installation. A Cluster represents one or more of the TMS-VSMs in a chassis. Clusters are treated as a single entity with regard to capacity – if one VSM goes down the performance is reduced accordingly -- and countermeasure programming. A chassis can contain more than one Cluster.

A TMS Group is a collection of one or more Clusters, which may comprise one or more chassis. Similar to Clusters, Groups enable simplified management as they are all programmed with the same mitigation.

Design Consideration 2: TMS Redundancy

TMS redundancy is effectively an extension of scalability, with the key difference that failure of a TMS must be detected in order to redirect traffic to another TMS or take other measures. Although other methods exist, the primary method of directing traffic to a TMS is via injection of a diversion route pointing to a next hop for the TMS which is on the same IP network as the ASR 9000's diversion interface to the VSM (an internal bundle of TenGE links from the ASR 9000 VSM's NPUs to the VSM's CPUs). This re-direction may occur via an announcement in the global BGP table, via Flowspec, via triggering ABF (Policy Routing) via BGP, or another routing mechanism.

If the TMS-VSM does down, its diversion and reinjection interfaces will go down so the TMS will no longer be reachable, thus allowing routing to re-converge. Within an ASR 9000, traffic can be load balanced among VSMs via two methods. First, the SP could announce different next hops for different victims to each VSM or Cluster. This method results in load sharing due to the operator spreading the load. Alternatively, the SP could announce a single next hop and use static routes (the advertised address is statically routed to both VSMs) to split the traffic among the VSMs. When combined with a network level anycast design, the second approach greatly simplifies the complexity of selecting a TMS. In anycast, multiple devices advertise the same route, and other routers will pick the

closest TMS based on routing metrics. It also provides seamless redundancy when one TMS goes down, albeit with reduced capacity. This model can be extended to many VSMs in a chassis.

Design Consideration 3: Deployment Routing Options

The mechanisms to create an effective diversion and re-injection path include BGP Flowspec, injecting a more specific route (diverts traffic from target to the TMS – note that only the traffic to the victim goes to the TMS), tunneling traffic to and/or from the TMS, putting the dirty and clean traffic in different VRFs/VPNs, and using ACL Based Forwarding (ABF) to steer traffic. These tools can be used in different combinations (e.g., Tunnel diversion & VRF re-injection, /32 diversion and VPN re-injection, and /32 diversion and GRE tunnel re-injection) to implement a range of routing designs. When making design decisions, make sure to consider the total amount of configuration required (e.g., are tunnels required to every router connected to a protected resource or is traffic just injected into a single clean VRF?) as well as ease of use during an attack. The next section focused on routing in more detail.

Routing Examples

Routing Example 1: Diversion via Flowspec & Default re-injection

Using BGP Flowspec for traffic diversion is particularly well suited to designs that place the TMS functionality on an ASR 9000 at the network edge. The TMS placement ensures that there will be no routing loops since only one router is involved in the traffic diversion. This greatly reduces the operational complexity by removing the need for additional VRFs or tunnels. The Flowspec update is originated from the SP and redirects traffic into the diversion interface. The onramp interface is placed in the default VRF for normal forwarding of clean traffic to the final destination.

Routing Example 2: Routing with VPN re-injection

One of the most flexible routing designs is to use a unique VRF to deliver the clean traffic to the destination. This allows traffic to be diverted via a /32 diversion route in the global routing table of the ingress edge router. After traffic is cleaned, it is placed in a clean VRF which contains routes for customers. Once traffic reaches the router for the protected resource, it is redirected back into “normal” path (PE-CE link or a link within the IDC). Alternatively, dirty traffic could be held in a VRF and clean traffic re-injected into the default table.

Routing Example 3: Diversion via Longest Match & GRE re-injection

The original deployment model with the Arbor TMS appliance was to divert traffic via injecting a more specific route into BGP (/32), and then using a preconfigured GRE tunnel to deliver traffic to a router close to the protected resource. The GRE tunnel termination could be a customer premises router, a Provider customer edge router, or an IDC router. When using this design, verify that all routers support GRE, as this may not always be the case. When using GRE with TMS-VSM, the GRE tunnel should be originated from the router. While GRE is a proven solution, it has the downside of provision tunnels for each protected destination which adds to management overhead.

Peakflow integration with BGP

BGP serves several purposes in Peakflow. First, peering with routers allows SP to detect network boundaries (to identify ingress points) and recognize Managed Objects. Second, abnormal BGP activity can be detected and alerts can be displayed. This use is outside the context of DDoS. Third, SP uses BGP to inject the route to divert traffic to the TMS for cleaning. Finally, BGP Flowspec may be used for traffic redirection as well as filtering or rate limiting.

Countermeasures

TMS implements countermeasures to block attack traffic and pass valid traffic. There are several types of countermeasures that are applied to traffic to protected destinations. They are executed in order so traffic passed by one countermeasure is then subject to the next enabled countermeasure (with the exception of the black/white list which can pass traffic directly). The Black/White list countermeasure permits or denies traffic based on ACL-like parameters. This filter is performed before other countermeasures. If a packet matches the black or white list, no further countermeasures are applied and the packets are dropped or forwarded appropriately. Note that ACLs on routers or filtering via BGP Flowspec provides higher performance and should therefore be seen as the primary option for basic filtering as that allows the TMS to focus on more in-depth countermeasures. Flowspec filters can be generated by the SP and are highly recommended as part of the DDoS mitigation solution.

Other countermeasures operate on L4 protocols. Examples of this type of countermeasure are Zombie Removal, TCP SYN Authentication, and TCP Connection Reset. They limit certain types of traffic or, in the case of TCP authentication, reply on behalf of the destination and await a valid reply before allowing further traffic through.

Another class of countermeasure operates at the application-layer. A common example is the HTTP Countermeasures. One of their functions is to use a mechanism similar to TCP SYN Authentication to validate the client by responding to a GET request and requiring a valid reply before allowing communication with the web server. This prevents zombies from flooding a server with GET requests from spoofed source addresses.

Finally, there are several countermeasures that look all the way into the payload. DNS or Payload Regular Expression fall into this category. This type of filtering can take advantage of the Arbor Threat Feed (ATF) which sends attack signatures from the Arbor Security Engineering and Response Team (ASERT) directly to the Peakflow system. For additional information on specific countermeasures, refer to the Peakflow SP User Guide and TMS Technical Overview. The behavior of individual countermeasures is consistent among VSM and the Arbor appliances.

To optimize performance, many of the countermeasures are capable of generating a dynamic blacklist that can be evaluated as one of the first countermeasures or offloaded into hardware. For example, if the HTTP Countermeasure identifies an invalid source, the source could be pushed into the dynamic blacklist which then can drop future packets more efficiently without running all the other enabled countermeasures. In a post-FCS release, the blacklisting and dynamic blacklisting will be offloaded into the ASR 9000 NPUs for greater performance.

During the design phase, templates for each protected resource (Managed Object or group of MOs) should be constructed as they will later serve as a starting point for building mitigations. During an attack, these templates can be applied and then the countermeasures enabled can modified based on the attack characteristics which may be recognized by the network operator or may be “pushed” from Arbor.

Countermeasures	Description	Usage
Black White List	FCAP expression to explicitly drop or pass traffic. Global list.	Pass traffic from critical services, Google crawler; drop spoofed sources, drop invalid ports/protocols.
IP Address Filter Lists	IP hosts or CIDR addresses. Global list.	Pass known partner networks, secure clients, support workers. Block infected subscriber lists, known BOT C&C servers.
Black White List (formerly global exception list in V5.1 and below)	FCAP expression to explicitly drop or pass traffic. Global list.	Pass traffic from critical services, Google crawler; drop spoofed sources, drop invalid ports/protocols.
URL Filter List	Lists of regular expressions that match on the URI fields in HTTP requests. Global list.	Efficient way to block/allow HTTP requests.
DNS Filter List	Lists of regular expressions that match DNS domains in DNS requests. Global list.	Efficient way to block/allow DNS requests.
GeoIP Filter List	GeoIP country lists are an extension of IP address lists. Global list.	Efficient way to drop or pass traffic from specified countries during a mitigation.
GeoIP Policing	Per mitigation list of GeoIP filters (countries) that are passed, dropped or rate-limited.	Effective in blocking/limiting sources that have no legitimate reason to be sending traffic (e.g., traffic from a country to an ecommerce site that does not do business there).
Zombie Removal	Removes sources that exceed defined pps or bps thresholds.	Effective against flood, TCP SYN and protocol attacks. Black lists offending hosts until their behavior falls below thresholds.
TCP SYN Authentication	Event-driven countermeasure designed to block TCP requests from spoofed sources and traffic generators.	Protects servers, firewalls and load balancers from TCP session table exhaustion.
TCP Connection Reset	Event-driven countermeasure that protects servers from excessive idle sessions.	TMS clears idle TCP sessions from back-end servers. Added features ensure graceful recovery for legitimate users.
Payload Regex	Regex countermeasures reassemble incoming packet streams and match on payload data. Event driven countermeasure.	Effective in removing attack traffic with known pattern in the payload.
Baseline Enforcement	Bandwidth and protocol enforcement. Monitors top subnets and protocols and identifies high spikes in traffic from normally low-volume sources or protocols.	Blocks sources and protocols that exceed the norms. Effective in mitigating attacks that use legitimate sources sending legitimate traffic.
DNS Countermeasures	Combination of raw and event-driven countermeasures designed to ensure that only valid DNS requests from valid sources are allowed.	Protects against DNS amplification, cache poisoning and resource exhaustion. Protects recursive and authoritative DNS. Scoping features focus countermeasures for virtualized, shared and NAT'd environments.
HTTP Countermeasures	Combination of raw and event-driven countermeasures designed to ensure that only valid HTTP requests from valid sources are allowed.	Protects Web servers from spoofed sources, traffic generators and bot sources. Scoping features efficiently focus countermeasures for virtualized, shared and NAT'd environments.
SIP Countermeasures	Combination of raw and event-driven countermeasures designed to ensure that only valid HTTP requests from valid sources are allowed.	Protects SIP servers from attack. Rate limiting and malformed detection.
Traffic Shaping/Rate Limiting	Limits traffic to a level that allows protected hosts to continue to function.	Effective in managing flash crowd events and helping control operations gracefully if other countermeasures are not fully mitigating an attack.

Figure 6 - TMS Countermeasures

Managed Objects

Managed Objects are a tool to effectively classify protected resources (usually more than one node or network) that are defined by a network edge. When alerts occur, they are flagged relative to a Managed Object. Managed Objects are treated as a single entity in attack detection and mitigation. For example, a Managed Object for a POP could identify routers by a POP-specific BGP community and then allow monitoring and mitigation specific to that POP. Managed Objects can be identified by a number of other parameters including VPN RD, AS Path regex, CIDR block, and peer ASN.

A key characteristic of Managed Objects is the network boundary that defines the edge of the MO. This boundary is used for defining where traffic is monitored for volumetric or application attacks. Peakflow builds a baseline for each Managed Object and then uses the thresholds configured for the Managed Object to generate alerts. Note that additional licenses may be required for MO scale.

Enterprise Applications of Peakflow

In this whitepaper, the primary focus and examples have been from Service Provider networks. As many large Enterprise networks are effectively of SP scale and design, almost everything covered applies. The unique decision an Enterprise must consider is whether they want to in-source or out-source their DDoS protection. They may also do a combination of both. Out-sourced protection may come from their SPs or from a cloud-based DDoS protection provider. There are many benefits of an Enterprise taking control of its own protection. First among these is that an outside provider's resources may be oversubscribed and unable to fully protect the network when there are multiple large attacks occurring against many of their customers. Another consideration which may favor one approach over the other is CAPEX (in-source) vs. OPEX (out-source) preference for the business. That said, in-sourcing will require an OPEX expenditure as well in order to develop and maintain design and implementation teams.

Footprint

ASR 9000 integration can reduce the facility and interface requirements for TMS. In addition to the reduced physical footprint for the TMS operation, the ASR 9000 does not require physical connections between the routers and the TMS. As the solution scales, the benefits of integration increase, especially once mitigation requirements exceed the capacity of a single appliance chassis (40G). In addition, the ASR 9000 is NEBS complaint, which expands the range of facilities where the TMS can be deployed.

	TMS-4[1/2/3/4]00	VSM
Capacity	10–40 Gbps 80 Gbps for hardware blacklist	40 Gbps / card 100+ Gbps per VSM with hardware blacklisting
Dimensions	6 RU	1 LC slot
Weight	85 lbs (38 kg) (TMS-4400)	20 lbs (9.1 kg) per VSM
Power	1200 W (10G) – 1650 W (40G)	740 W
Interfaces	8 x TenGE	None required

Table 1 - Footprint Comparison

Figure 7 - ASR 9000 Virtualized Services Module

Unique Aspects of the ASR 9000 Implementation

While the software that implements the TMS is the same on the appliance and the VSM, there are minor differences due to the hardware. The VSM utilizes 4 10-core Intel CPUs to implement the TMS countermeasures. Packets for mitigation are load shared among the cores, and all packets between a pair of addresses will go to the same CPU core, thus enabling stateful countermeasures like TCP and HTTP authentication.. Note that this does limit the connection between two IP addresses to the processing capacity of that core.

Another aspect of the software running on different hardware is the performance of the various countermeasures. Overall, the performance of the VSM is roughly equal to the same capacity on the TMS-4400 (40 Gbps).

The TMS implementation on ASR 9000 uses the Cisco IOS XR VM hosting capability which provides an API for third-party applications. There is therefore no TMS-specific configuration present in the IOS XR configuration. The following configuration is required (additional details in the configuration guide):

1) Configure the 0/<slot#>/0/5-6 TenGE interfaces as a management links. 0/<slot#>/1//6 is unused at this time but should not be included in the data path link bundle. Assign an IP address to 0/<slot#>/1/5. This address will be used in the TMS setup via the GUI and must be on the same address as the TMS management as configured below under Administration->Peakflow Appliances (the router is 10.99.1.1)

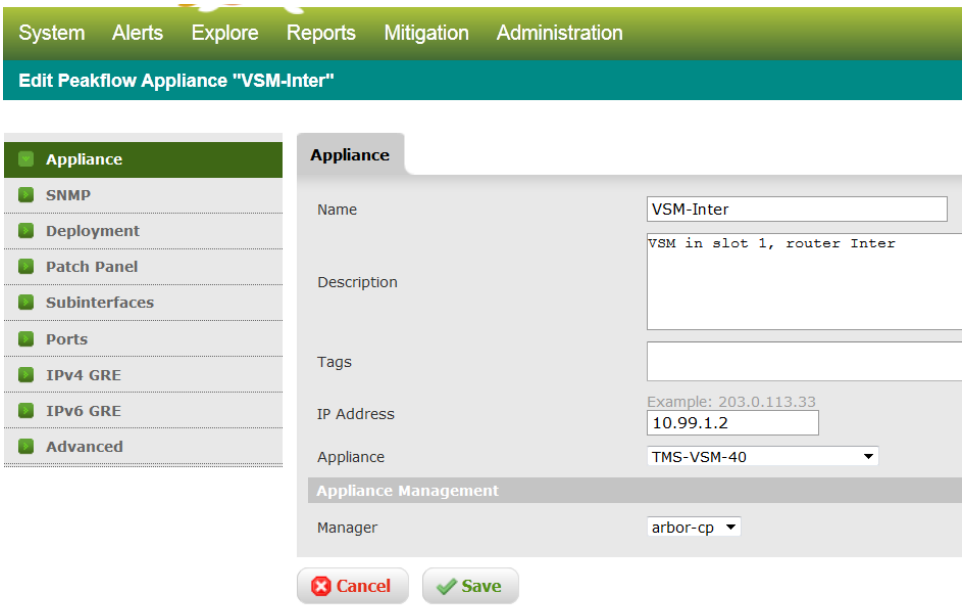

2) Place the remaining TenGigE links on the VSM into a link bundle for the data path.

3) Configure diversion and reinjection sub-interfaces on the bundle. The interfaces and addresses for the diversion and re-injection interfaces will be configured in the GUI to communicate the values to the SP.

```
interface TenGigE0/1/1/0
 bundle id 2 mode on
 load-interval 30
```

```
!
interface TenGigE0/1/1/1
 bundle id 2 mode on
 load-interval 30
!
…
!
interface TenGigE0/1/1/5
 description mgt0 on TMS1
 ipv4 address 10.99.1.1 255.255.255.0
 load-interval 30
!
interface Bundle-Ether2
 description bundle to-from vsm1
 load-interval 30
!
interface Bundle-Ether2.100
 description diversion subinterface
 ipv4 address 198.51.100.1 255.255.255.252
 load-interval 30
 encapsulation dot1q 100
!
interface Bundle-Ether2.101
 description reinjection subinterface
 ipv4 address 198.51.100.5 255.255.255.252
 load-interval 30
 encapsulation dot1q 101
```

4) Configure the SP with these addresses in the Administration->Peakflow Appliances->Patch Panel per the configuration guide.

Hosted Arbor TMS CLI

After initial setup, operation of the TMS will be executed via the GUI on the SP. For initial setup, the TMS CLI must be accessed via the IOS XR CLI. After setup, the TMS CLI may also be accessed via ssh to the m anagement address as well. The TMS CLI is not required for normal operation but can be helpful in troubleshooting.

To attach to the TMS application from the XR CLI, use the following command: **virtual-service connect name TMS1 console node 0/<slot#>/cpu0**. Return to the IOS XR CLI via the escape sequence: **"CTRL+SHIFT+6 e".**

```
RP/0/RSP0/CPU0:ASR9000#virtual-service connect name TMS1 console node 0/1/cpu0
Mon Dec 15 23:46:14.983 UTC
Trying <address>
Connected to <address>
Escape sequence is '^^e'.
admin@arbos:/# ?
Subcommands:
  ip/              IP and network configuration
  services/        System services
  system/          System configuration
```

Summary

The Arbor Peakflow system is the proven industry leader in DDoS detection and mitigation. The Cisco ASR 9000 is the industry leader in high performance edge and peering routing. By partnering with Arbor, Cisco has expanded the reach, performance, scalability, and places-in-network of the DDoS solution. When combined with Cisco's existing security and DDoS feature sets, Cisco and Arbor are now offering a full suite of tools for Service Providers to offer high performance cost effective value added services to their customers as well as enabling Enterprises to take control over their DDoS protection.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA CXX-XXXXXX-XX 10/11