

Cisco ASR 9000 vDDoS Protection Solution

Last Updated: 5/8/2015



What You Will Learn

Cisco and Arbor Networks have collaborated to bring industry-leading distributed denial-of-service mitigation capabilities on the Cisco® ASR 9000 Series routers. This collaboration extends the Arbor Peakflow Threat Management System DDoS mitigation functions to the ASR 9000.

The TMS will be implemented on the Cisco ASR 9000 Series Virtualized Services Module hosted in the ASR chassis. This solution enables unprecedented scaling of the DDoS mitigation architecture. It reduces operational complexity and hardware footprint. It allows network providers to push attack mitigation to the edge of the network reducing the need to carry the traffic to scrubbing centers.

Goals

This white paper provides an in-depth look into the Peakflow solution with a focus on the new capabilities enabled by integrating the Arbor Peakflow TMS DDoS mitigation capability into the ASR 9000.

Executive Summary

Cisco and Arbor have collaborated to integrate the Arbor Peakflow Threat Management System (TMS) into the Cisco ASR 9000 Aggregation Services Router through the Virtualized Services Module (VSM). This new solution allows for optimal placement of distributed denial-of-service (DDoS) mitigation within the network, providing 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 ASR 9000 vDDoS Solution solution will also enable intelligent software-defined network (SDN)-directed 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 DDoS attacks. Peakflow accomplishes this mitigation 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 process is implemented through access-control-list (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 follows: First, the system monitors network ingress points with NetFlow and Border Gateway Protocol (BGP) to build a baseline for network behavior and traffic patterns. It then performs ongoing monitoring to detect anomalies and flag them as potential attacks. These potential attacks are presented to network operations through a GUI, email message, or Simple Network Management Protocol (SNMP), which allows a range of actions to be taken, including initiating a response or marking an event as a false alarm.

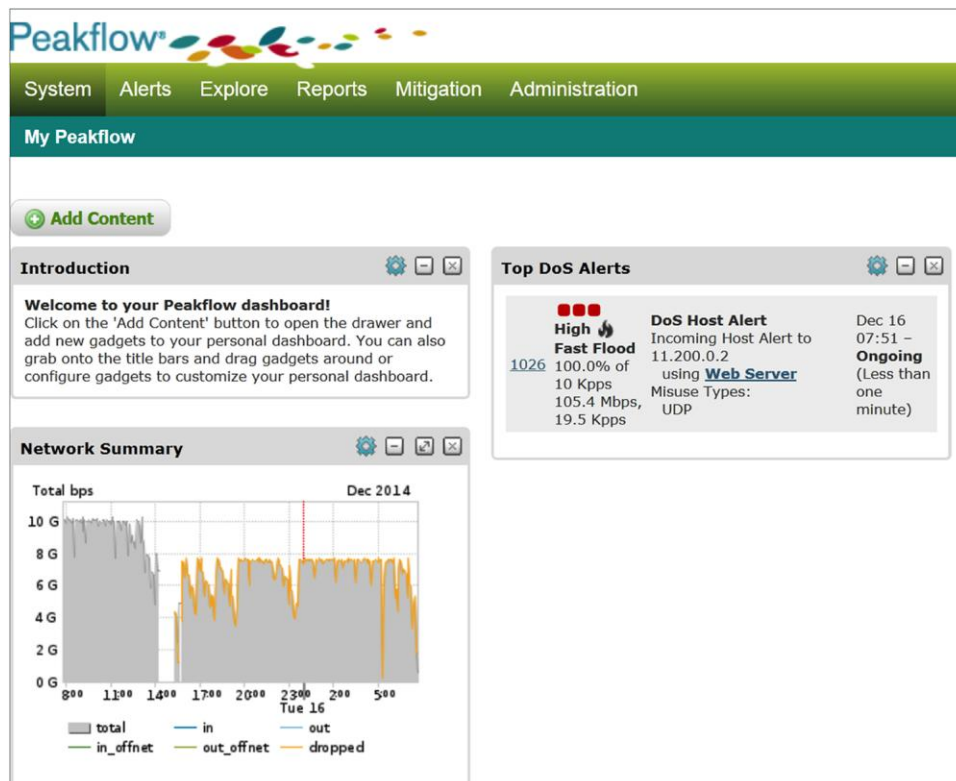
The main actions to initiate a response follow:

- Update the network routing to redirect all traffic for the destination through the TMS-VMS, which can remove unwanted traffic.
- Clean the traffic as effectively as possible without blocking valid connections.

In the solution with Arbor, the Cisco ASR 9000 integrates the traffic cleaning functions into the router and scales its performance operation far beyond what is possible with appliances alone or with competing offerings.

Figure 1 shows the Peakflow GUI.

Figure 1. Peakflow GUI

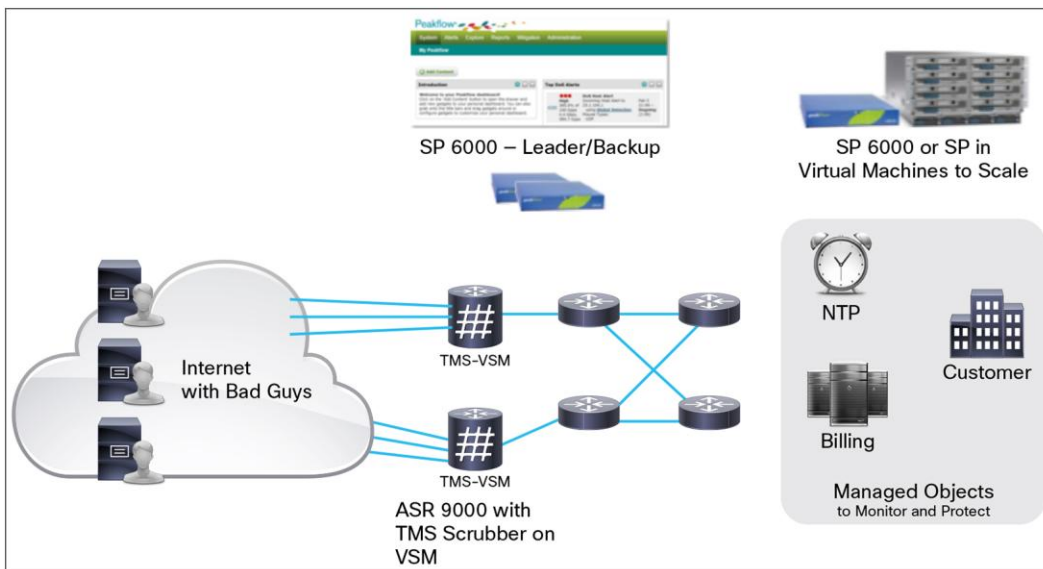


Key Components

Peakflow, the overall Arbor solution for both network analytics and DDoS detection and mitigation, comprises numerous functions as well as a set of hardware devices that implement those functions. “Peakflow SP” refers to the control components for Peakflow such as monitoring the network, detecting attacks, and coordinating an attack response. The 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 Cisco 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. Starting Arbor release 7.0.3 Peakflow SP can run completely as virtualized machines, no hardware appliance necessary. As part of the ASR 9000 vDDoS solution, only the virtualized Peakflow SP will be available on Cisco Price list.

Figure 2 shows the key Peakflow components.

Figure 2. Key Peakflow Components



Peakflow SP provides the following functions:

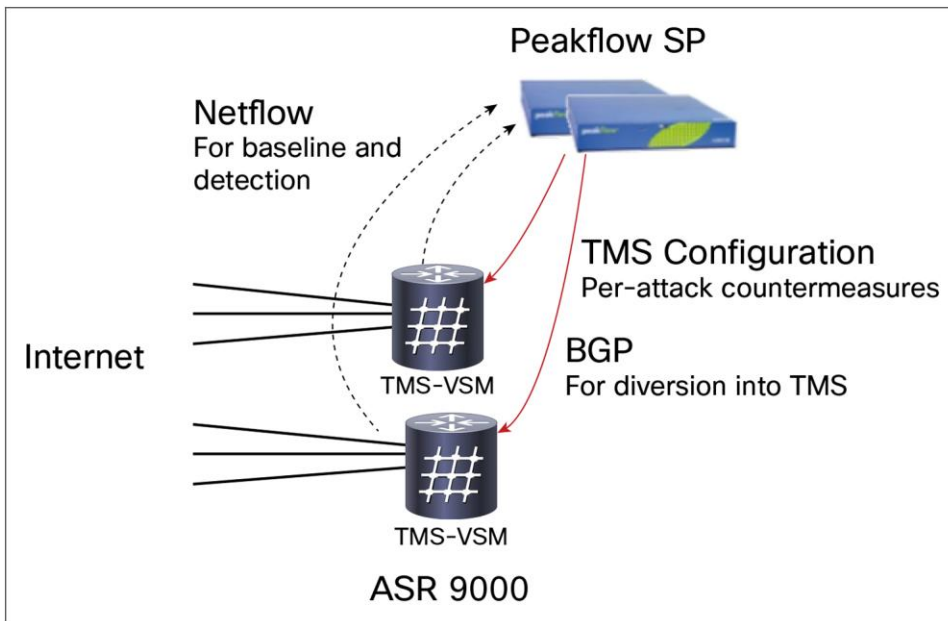
- Leads the overall system and manages communication among components
- Presents the GUI to network operations staff
- Receives NetFlow and routing information from the routers
- Analyzes the data to detect anomalies and generate alerts
- Creates diversion and reinjection paths through BGP or BGP Flowspec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

Peakflow SP can be instantiated on a combination of SP appliances and user-provided virtual machines or starting Arbor release 7.0.3 it can be fully virtualized. Prior to Arbor release 7.0.02 a physical appliance (for example, SP 6000) is required to serve as the leader and (optionally) backup leader but with release 7.0.3 it can be fully virtualized, no need for hardware appliance. Cisco will only be having the virtualized Peakflow SP on its price list.

Additional scale occurs through appliances or virtual machines as well as software licenses. Earlier versions of Peakflow hosted the SP functions on a set of appliances known as the Collector Platform (CP), Flow Sensor (FS), Portal Interface (PI), and Business Intelligence (BI) and scaled by adding appliances. These components may still be used but aren't presented as part of the reference solution.

Figure 3 Illustrates SP-to-TMS Communication.

Figure 3. SP-to-TMS Communication



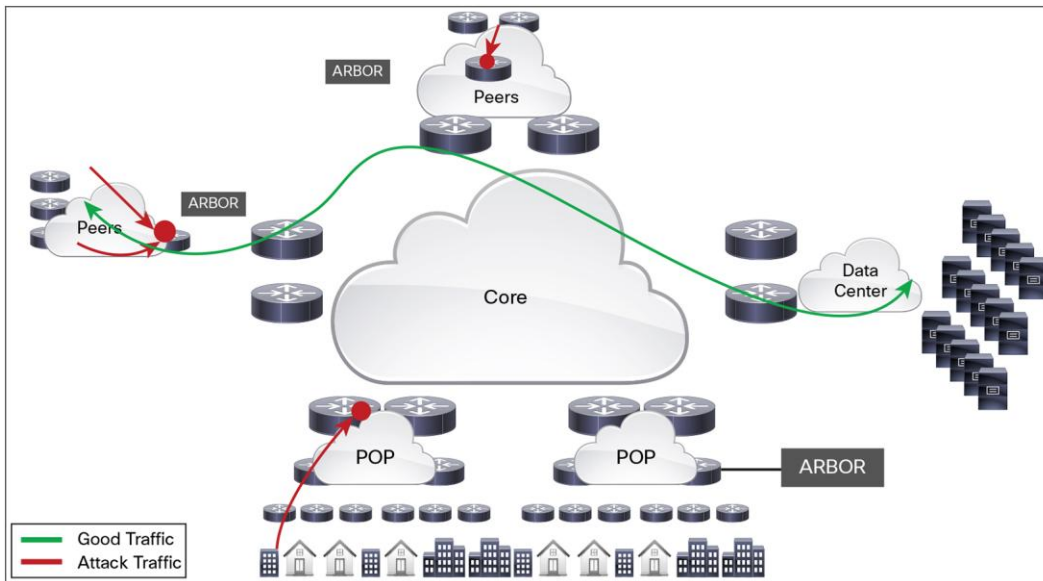
Peakflow TMS (or SP TMS) provides the following functions:

- 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 Cisco ASR 9000 VSM. It can coexist with Arbor TMS appliances that are 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 function into the ASR 9000, allowing new scalability, performance, and design options. Figure 4 shows 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

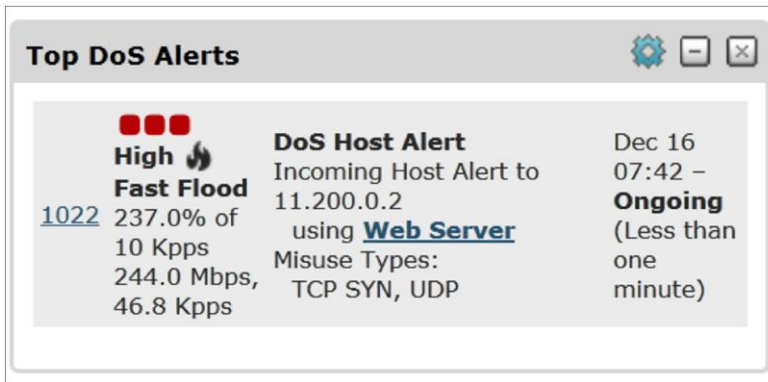
NetFlow, SNMP, and BGP are the key tools to monitor for DDoS attacks or other anomalies. SNMP can provide interface status and settings. BGP can provide information about network boundaries (based on managed object configuration (discussed later in this paper) and routing status. NetFlow provides the 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 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 1:1,000 to 1:10,000 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 because the error is very low with even 1,000 samples (samples received, not the sampling rate) and improves only about 2 percent when going beyond 10,000 samples.

For a given sampling error, attack detection speed depends on the type of attack. Volumetric attacks can be recognized more quickly because 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 they match a known fingerprint. For this reason, detection and mitigation of application layer attacks should happen after volumetric traffic has been reduced because it results in a higher percentage of application attack packets in the sample near the data center or application host.

Note that although attack detection traditionally takes 1 to 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 shows a sample alert.

Figure 5. Sample Alert



Network Baseline for Anomaly Detection

In addition to recognizing attacks through rates and fingerprints, the Peakflow SP system can build a traffic baseline for each managed object (protected resource) and recognize increased activity. This increase in activity may be due to an attack or increased traffic caused by 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 and Application

At a high level, there are two main classes of DDoS 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 User Datagram Protocol (UDP) traffic to a single destination in order to congest a network link or the packet capacity of a server. The Network Time Protocol (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 Domain Name System (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 numerous valid connections to a server and then keeps them open indefinitely in order to reach the maximum connection capacity of the server. 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.

Although 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 or putting the Peakflow to run in an inline mode.

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 Signalling is an efficient and integrated system coordinating DDoS mitigations from the customer premises to the service provider cloud.

Design Considerations

Protected Resources and 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 through peering or exchanges (from IDCs or from compromised end users) and from other customers of the Service Provider 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 attacks comprise attacks from one of the service provider'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 determines the most effective way to direct traffic for cleaning. One factor that is unique to this scenario is that it is simpler for the service provider to take action local to the attacker, which may mean increased reliance on other (non-TMS) features such as ACLs, Unicast Reverse Path Forwarding (uRPF), and black holes because the attack source is localized.

Protecting data centers and infrastructure servers (for example, DNS, billing, authentication, and route reflectors) is critically important to a service provider's or large enterprise's ability to minimize the impact of a volumetric attack that will often accompany an application layer attack. Application level attacks must be mitigated with Layer 4-through Layer 7-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 because 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 the resources that are being protected and the attack sources that must be managed. When 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 (through syslog, NetFlow, and routing), scale the number of protected destinations (managed objects), and provide user access through the GUI. In recent years, the volume of attack traffic has been 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, Unicast Reverse Path Forwarding (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 with 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), implying 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 on-ramp, 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 through additional Peakflow SP appliances or virtual machines and software licenses.

For TMS scale, the location of the TMS affects 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 that 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 100 Gbps 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 four 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 TMSs or to a regional scrubbing center.

Because the mitigation happens per diverted prefix and is often based on finding the nearest TMS through routing metrics, a single victim's mitigation can be spread among multiple TMSs through the metric of the diversion route created on the CP GUI. The use of anycast to find the nearest TMS can simplify this process (this topic is covered in more detail in the section "Routing Example 1: Diversion with Flowspec and Default Reinjection").

Scalability with multiple VSMs within an ASR 9000 is another likely path to increase capacity without introducing additional complexity. When traffic reaches the TMS router, the diversion next hop can be statically routed into multiple VSMs, a process that 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-first-customer-shipment (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 four CPUs. The selection of the VSM is based on the default Cisco IOS® XR Software 5-tuple load-sharing feature. For IPv4 and IPv6 packets, the selection is based on source IP, destination IP, source port (TCP/UDP only), destination port (TCP/UDP only), and router ID). Multiprotocol Label Switching (MPLS) load sharing depends on the size of the label stack and payload type.

When 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 the state of the flow required 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 corresponding 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 because 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 by injection of a diversion route pointing to a next hop for the TMS that is on the same IP network as the diversion interface of the ASR 9000 to the VSM (an internal bundle of 10 Gigabit Ethernet links from the ASR 9000 VSM NPUs to the VSM CPUs). This redirection may occur with an announcement in the global BGP table, with Flowspec, by triggering ACL-Based Forwarding (ABF) (Policy Routing) through BGP, or another routing mechanism.

If the TMS-VSM goes down, its diversion and reinjection interfaces will also go down, so the TMS will no longer be reachable, thus allowing routing to reconverge. Within an ASR 9000, traffic can be load balanced among VSMs by two methods. First, the service provider could announce different next hops for different victims to each VSM or cluster. This method results in load sharing because the operator is spreading the load. Alternatively, the service provider 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 transparent 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 reinjection 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 Virtual Route Forwarding (VRFs) or VPNs, and using ABF to steer traffic. These tools can be used in different combinations (for example, tunnel diversion and VRF reinjection, /32 diversion and VPN reinjection, and /32 diversion and generic routing encapsulation (GRE) tunnel reinjection) to implement a range of routing designs.

When making design decisions, be sure to consider the total amount of configuration required (for example, 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 focuses on routing in more detail.

Routing Examples

Routing Example 1: Diversion with Flowspec and Default Reinjection

Using BGP Flowspec for traffic diversion is particularly well suited to designs that place the TMS functions on an ASR 9000 at the network edge. The TMS placement helps ensure that there will be no routing loops because only one router is involved in the traffic diversion. This setup greatly reduces the operational complexity by removing the need for additional VRF or tunnel configuration. The Flowspec update originates from the Peakflow 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 Reinjection

One of the most flexible routing designs is to use a unique VRF to deliver the clean traffic to the destination, allowing traffic to be diverted through 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 that contains routes for customers. When traffic reaches the router for the protected resource, it is redirected back into “normal” path (provider edge-customer edge 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 Through Longest Match and GRE Reinjection

The original deployment model with the Arbor TMS appliance was to divert traffic by 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, because they may not always support GRE. When using GRE with TMS-VSM, the GRE tunnel should originate from the router. Although 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 service providers 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, service providers use BGP to inject the route to divert traffic to the TMS for cleaning. Finally, BGP Flowspec can be used for traffic redirection as well as filtering or rate limiting.

Countermeasures

TMS implements countermeasures to block attack traffic and pass valid traffic. Several types of countermeasures 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 and white list, which can pass traffic directly). The black and 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 with BGP Flowspec provides higher performance and should therefore be seen as the primary option for basic filtering because that allows the TMS to focus on more in-depth countermeasures. Flowspec filters can be generated by the service provider and are highly recommended as part of the DDoS mitigation solution.

Other countermeasures operate on Layer 4 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 requirement prevents zombies from flooding a server with GET requests from spoofed source addresses.

Finally, several countermeasures 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 about specific countermeasures, refer to the “Peakflow SP User Guide” and “TMS Technical Overview” documents. 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 black list 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 black list, which then can drop future packets more efficiently without running all the other enabled countermeasures. In a post-FCS release, the black listing and dynamic black listing 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 managed objects) should be constructed because 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 be modified based on the attack characteristics that may be recognized by the network operator or may be “pushed” from Arbor.

Table 1. Lists Some TMS Countermeasures.

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.

Countermeasures	Description	Usage
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 SIP 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.

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 point of presence (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 many other parameters including VPN route distinguisher (VPN RD), autonomous-system path regex (AS path regex), CIDR block, and peer Access Service Network (ASN).

A key characteristic of managed objects is the network boundary that defines the edge of the managed object. 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 managed object scale.

Enterprise Applications of Peakflow

In this paper, the primary focus and examples have been from service provider networks. Because many large enterprise networks are effectively of service provider scale and design, almost everything covered applies. The unique decision an enterprise must consider is whether it wants to in-source or out-source its DDoS protection. It may also do a combination of both. Out-sourced protection may come from its service providers 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 multiple large attacks occur against many of its customers.

Another consideration that may favor one approach over the other is capital expenditures (CapEx) (in-source) vs. operating expenses (OpEx) (out-source) preference for the business. That said, in-sourcing requires an OpEx expenditure as well in order to develop and maintain design and implementation teams.

Footprint

Cisco 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 when mitigation requirements exceed the capacity of a single appliance chassis (40 Gbps). In addition, the ASR 9000 is Network Equipment Building Standards (NEBS)-complaint, so the range of facilities where the TMS can be deployed is expanded.

Table 2 compares footprints of the TMS and VSM, and Figure 6 shows the Cisco ASR 9000 Virtualized Services Module.

Table 2. Footprint Comparison

	TMS-4[1/2/3/4]00	VSM
Capacity	10-40 Gbps 80 Gbps for hardware black list	40 Gbps per card 100+ Gbps per VSM with hardware black listing
Dimensions	6 rack units (6RU)	1 line-card slot
Weight	85 lb (38 kg) (TMS-4400)	20 lb (9.1 kg) per VSM
Power	1200W (10 Gbps) to 1650W (40 Gbps)	740W
Interfaces	Eight 10 Gigabit Ethernet	None required

Figure 6. Cisco ASR 9000 Virtualized Services Module



Innovative Aspects of the Cisco ASR 9000 Implementation

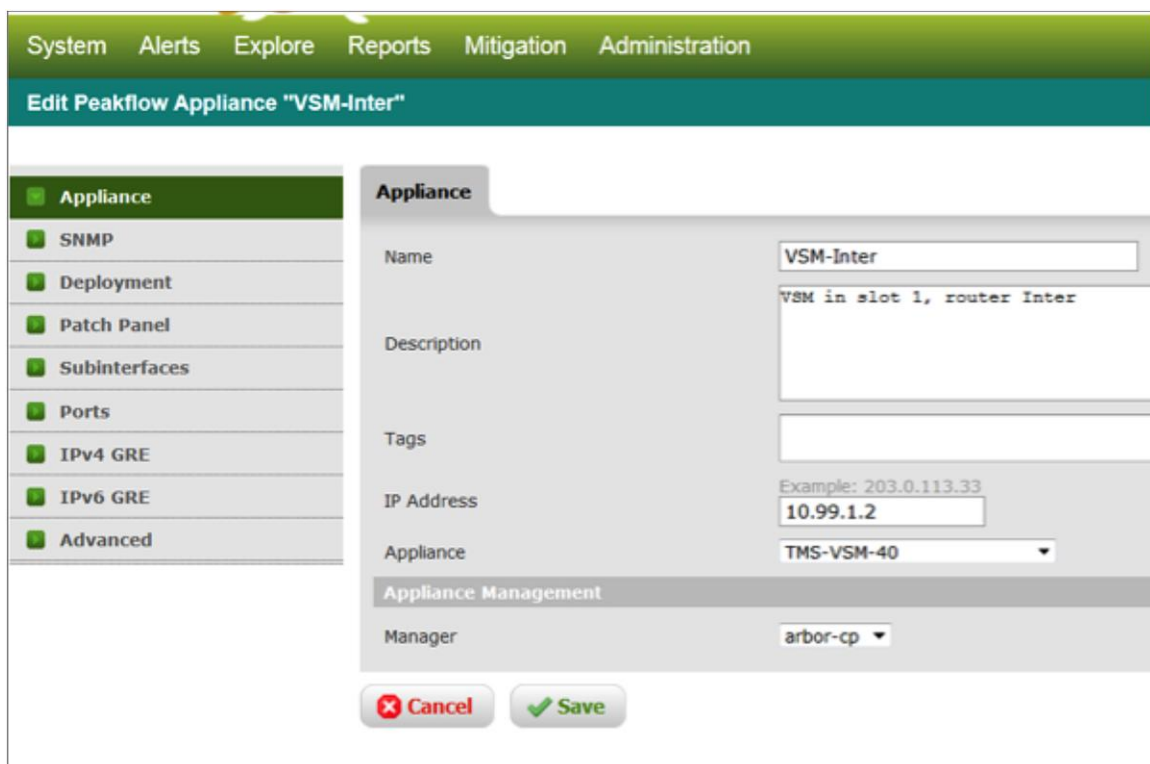
Although the software that implements the TMS is the same on the appliance and the VSM, there are minor differences due to the hardware. The VSM uses four 10-core Intel CPUs to implement the TMS countermeasures. Packets for mitigation are load-shared among the cores, and all packets between a pair of addresses will go to the same CPU core, thus enabling stateful countermeasures such as TCP and HTTP authentication. Note that this setup 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 the Cisco ASR 9000 uses the Cisco IOS XR virtual machine hosting capability, which provides an application-programming interface (API) for third-party applications. There is, therefore, no TMS-specific configuration present in the Cisco IOS XR Software configuration. The following configuration is required (additional details are in the configuration guide):

- Step 1. Configure the 0/<slot#>/0/5-6 10 Gigabit Ethernet interfaces as management links. The 0/<slot#>/1//6 interface is unused at this time, and 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 through the GUI and must be on the same address as the TMS management as configured in Figure 7 under Administration->Peakflow Appliances (the router is 10.99.1.1).

Figure 7. TMS Management



- Step 2. Place the remaining 10 Gigabit Ethernet links on the VSM into a link bundle for the data path.

- Step 3. Configure diversion and reinjection sub-interfaces on the bundle. The interfaces and addresses for the diversion and reinjection 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 vsml
  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

```

Step 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 through the GUI on the SP. For initial setup, the TMS CLI must be accessed through the Cisco IOS XR Software CLI. After setup, the TMS CLI may also be accessed with Secure Shell (SSH) Protocol to the management address as well. The TMS CLI is not required for normal operation, but it can be helpful in troubleshooting.

To attach to the TMS application from the Cisco IOS XR Software CLI, use the following command: `virtual-service connect name TMS1 console node 0/<slot#>/cpu0`. Return to the Cisco IOS XR Software CLI through 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

Arbor Peakflow is the industry leader in DDoS detection and mitigation. The Cisco ASR 9000 is the industry leader in high-performance edge and peering routing. By collaborating 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 now offer a full suite of tools for service providers to provide high-performance, cost-effective, value-added services to their customers. The solution also enables enterprises to control 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.

 Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)