DESIGN GUIDE

CISCO TRAFFIC ANOMALY DETECTOR MODULE AND CISCO ANOMALY GUARD MODULE

SUMMARY

This document provides some general design considerations and introductory design guidelines for deploying the Cisco® Traffic Anomaly Detector Module and the Cisco Anomaly Guard Module. These modules work together to detect and mitigate distributed denial of service (DDoS) attacks.

1. CISCO TRAFFIC ANOMALY DETECTOR MODULE

Role of the Cisco Traffic Anomaly Detector Module

The Cisco Traffic Anomaly Detector Module is a passive monitoring device that constantly looks for indications of a DDoS attack against a protected destination (referred to as a "zone"), such as a server, firewall interface, or router interface. The Traffic Anomaly Detector Module analyzes copies of all inbound traffic (via Switched Port Analyzer [SPAN] or passive network tap) destined for the protected zone or zones. This analysis consists of comparing the current traffic to a set of behavioral thresholds (a "zone policy") to detect anomalous traffic behavior. If anomalous behavior is seen and is considered a possible attack, the Traffic Anomaly Detector Module will signal the Anomaly Guard Module (via an out-of-band Ethernet management network) to start mitigating the attack.

Place in the Network

The Traffic Anomaly Detector Module is placed logically downstream from the Anomaly Guard Module, but upstream of any firewall. During non-attack periods, the Traffic Anomaly Detector Module will see all inbound traffic destined for the protected zone. During an attack where an Anomaly Guard has diverted traffic from the targeted zone for mitigation, the Traffic Anomaly Detector will only see the "cleaned" traffic leaving the Anomaly Guard destined for the zone.

Performance and Capacity Limits

Each individual Traffic Anomaly Detector Module is capable of:

- Analyzing up to 1 Gbps of inbound Ethernet traffic
- Containing a maximum of 500 configured zones
- Actively monitoring 90 zones (where each zone is one or more protected destination IP addresses) concurrently

Up to eight Traffic Anomaly Detector Modules can be deployed in a Catalyst 6500 chassis with a minimum Supervisor Engine 2 and SFM.

Redundancy Options

Multiple Traffic Anomaly Detector modules can be configured to analyze traffic for the same set of zones, in an "Active-Active" arrangement. The modules would be configured identically, but would operate independently. When anomalous traffic behavior is detected, triggering a remote activation of an Anomaly Guard for attack mitigation, all of the Traffic Anomaly Detector modules will send activation notices to the Anomaly Guard. While the Anomaly Guard will receive multiple activation notices for the same zone, it will ignore all subsequent notices after processing the first activation notice. In this arrangement, the failure of any single Traffic Anomaly Detector Module will not stop attack detection and Anomaly Guard activation service.

Up to eight Traffic Anomaly Detector modules can be deployed in a single Cisco Catalyst 6500 Series chassis with a minimum Supervisor Engine 2 and SFM in an Active-Active redundancy configuration.

Figure 1 illustrates the deployment of the Traffic Anomaly Detector Module and the Anomaly Guard Module in a single Cisco Catalyst 6500 Series chassis.

Figure 1. Cisco Catalyst 6500 Chassis with Traffic Anomaly Detector and Anomaly Guard

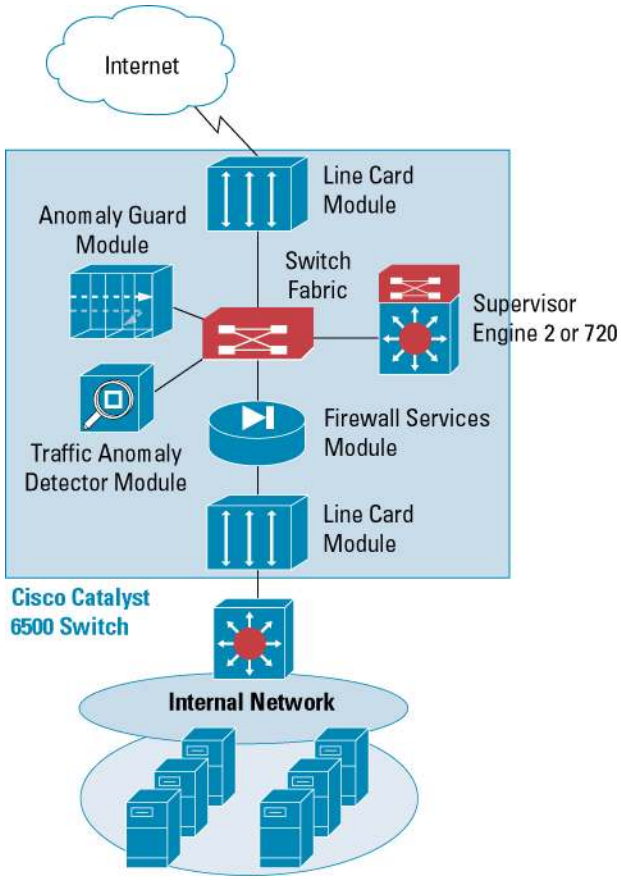

2. CISCO ANOMALY GUARD MODULE

Role of the Cisco Anomaly Guard Module

The Cisco Anomaly Guard Module is an active DDoS attack mitigation device that diverts and processes suspect traffic to drop attack packets and forward legitimate transactions. When a possible attack is detected by the Traffic Anomaly Detector, it notifies the Anomaly Guard to divert all traffic destined to the zone under attack for cleaning. This process is called traffic diversion.

The Anomaly Guard subjects this diverted traffic to analysis and countermeasures to distinguish between attack traffic and legitimate traffic. Attack traffic is dropped by the Anomaly Guard, and legitimate traffic is forwarded to its intended zone. When the Anomaly Guard Module determines that an attack is over, it stops diverting the zone traffic.

The Anomaly Guard is not normally deployed as an inline device; it uses traffic diversion to insert itself into the traffic path when necessary to mitigate an attack. Deploying the Anomaly Guard "inline-on-demand" allows for more efficient provisioning—you only deploy enough Anomaly Guard capacity to process suspect traffic, rather than all traffic. This also allows a more resilient network topology by not having all traffic run through another possible point of failure.

Place in the Network

In general, the Anomaly Guard Module should be placed as far upstream from the protected zones —and as close to the source of the attack traffic— as possible. This allows the Anomaly Guard to protect all downstream resources from DDoS attack traffic. These downstream resources can consist of servers, routers, switches, firewalls, and intrusion detection systems (IDSs).

The Anomaly Guard Module must also be placed upstream of a firewall, to process traffic before any Network Address Translation (NAT) processing occurs, and to protect the firewall itself from becoming a victim of a DDoS attack

Traffic Diversion

In the traffic diversion process, zone traffic is diverted to the Anomaly Guard Module for cleaning during a possible attack against that zone, and is reverted back to its normal path when the Anomaly Guard determines that the attack is over. Traffic diversion allows the Anomaly Guard to operate in an "inline-on-demand" fashion. There are two parts to traffic diversion:

- Traffic hijacking—Diverting affected traffic from the normal path to the Anomaly Guard for cleaning
- Traffic injection—Forwarding cleaned (legitimate) traffic from the Anomaly Guard back into the normal path

Integrated or Dedicated Configuration

There are two network topology choices for traffic diversion—integrated and dedicated.

- Integrated—Zone traffic normally flows through the Cisco Catalyst 6500 Series switch that also contains the Anomaly Guard Module. In the integrated configuration, the Anomaly Guard Module uses Route Health Injection (RHI) to add a static route to the switch's routing table (supervisor engine) to hijack traffic (Figure 2). Zone traffic is then diverted "internally" from the supervisor engine to the Anomaly Guard Module; traffic injection also occurs internally, from the Anomaly Guard Module back to the supervisor engine.

Figure 2. Integrated Configuration

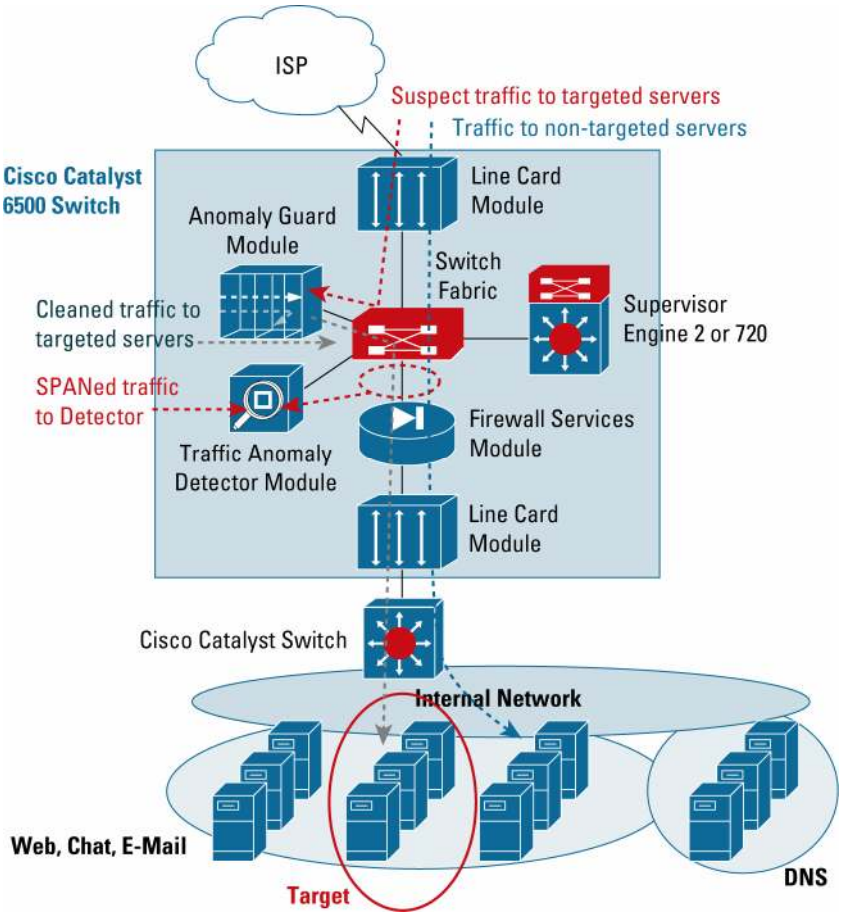

- Dedicated (or "out-of-path")—In the dedicated configuration, the Anomaly Guard Module resides in a Cisco Catalyst 6500 Series switch that is not in the normal traffic path for the zone being protected (Figure 3). Traffic hijacking still starts with the Anomaly Guard Module sending an RHI message to add a static route to the supervisor engine routing table, but in the second step, the supervisor engine (or routing process on the supervisor engine) redistributes this static route to an upstream router using a routing protocol like Border Gateway Protocol (BGP) or Open Shortest Path First (OSPF). Dedicated configuration is intended for building a dedicated appliance with one or more Anomaly Guard modules inside a Cisco Catalyst 6500 Series chassis.

Figure 3. Dedicated Configuration

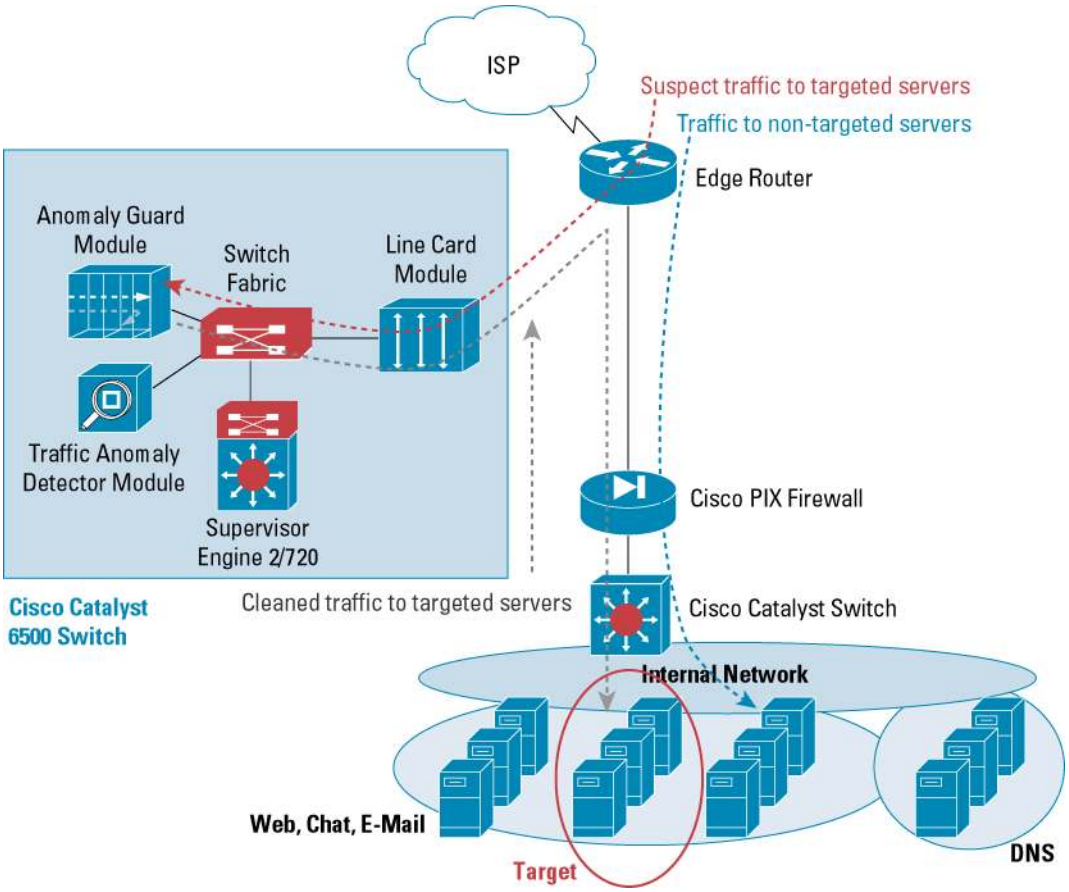

Traffic Hijacking Options

In the integrated configuration, there is only one traffic hijacking option—sending an RHI message from the Anomaly Guard to the routing process of the supervisor engine to add a static route to the supervisor engine routing table that names the Anomaly Guard as the next hop.

In the dedicated configuration, there is only one option for the first step of the process—sending an RHI message from the Anomaly Guard Module to the supervisor engine on the same Cisco Catalyst 6500 Series switch chassis. In the second step, where the supervisor engine routing process has to redistribute the static route created by the RHI advertisement to an upstream router, there are multiple options. Whichever Cisco IOS® Software-supported interior gateway routing protocol is desired for the redistribution from the dedicated guard to the upstream router can be used.

Traffic Injection Options

For both traffic diversion configurations, the traffic injection options are Layer 2 or Layer 3 topologies.

In the Layer 2 topology option, cleaned traffic is forwarded from the Anomaly Guard Module to a statically configured next-hop address residing on a downstream router that is on the same VLAN/subnet as the Anomaly Guard traffic injection interface/VLAN. Layer 2 traffic injection is the simplest to configure—no significant configuration changes are required on the downstream router.

In the Layer 3 topology option, there are two traffic injection choices—VPN Routing and Forwarding (VRF) or tunnel (generic routing encapsulation [GRE] or IP over IP [IPIP]).

VRF traffic injection allows you to configure separate routing and forwarding tables (separate from the global routing and forwarding table) for forwarding injected traffic sent from the Anomaly Guard Module. Without a separate routing table, the injected traffic (clean traffic from the Anomaly Guard) would be sent back to the Anomaly Guard because of the existence of the routing entry in the global routing table that is being used for traffic hijacking.

The VRF traffic injection configuration resides on the upstream/downstream router, from which hijacked traffic is sent to the Anomaly Guard and to which injected traffic is sent. The only configuration on the Anomaly Guard would be a static route pointing to a next hop for the clean traffic destined to the zone.

Whether using the integrated or dedicated configuration, you would configure one tunnel endpoint on the supervisor engine (or routing entity) and the other endpoint at the router adjacent to the zone. A common VLAN would be configured on the Anomaly Guard to carry injected traffic to the supervisor engine for mapping onto the appropriate tunnel.

Performance and Capacity Limits

Bandwidth/Throughput

Each Anomaly Guard Module is capable of receiving up to 1 Gbps of Ethernet traffic. In practical deployment where there is a mix of attack and legitimate traffic the Anomaly Guard Module can process up to 1 Mpps of combined attack and legitimate traffic. This maximum of 1 Mpps can be reduced by the following factors:

- Use of the Strong Anti-Spoofing countermeasure (TCP Proxy).
- The number of active dynamic filters in operation. The maximum number of active dynamic filters available per Anomaly Guard Module is 150,000. A 20-percent performance drop (from 1 Mpps to 800 Kpps) is possible when the active number of dynamic filters exceeds 100,000.

Active Zone Capacity

Each Anomaly Guard Module is capable of cleaning traffic for up to 30 active zones. A total of 500 zones can be configured, but only 30 can be concurrently active. Each zone can contain a minimum of a single host or subnet destination address, and up to a maximum of 100 host or subnet destination addresses.

Clustering Multiple Anomaly Guards for Higher Bandwidth/Throughput

For greater attack mitigation capacity, multiple Anomaly Guard modules can be clustered to operate as one virtual Anomaly Guard providing multigigabit bandwidth/throughput. By using Cisco Express Forwarding load sharing, you can hijack traffic for individual zones to a maximum of eight Anomaly Guard modules for cleaning (Cisco Express Forwarding load sharing allows for a fairly equal distribution of traffic among the multiple Anomaly Guard modules, while maintaining connection persistence through each individual module). This provides a maximum theoretical capacity of 8 Gbps, or a practical capacity of 8 Mpps for cleaning zone traffic.

Redundancy Options

There are two redundancy options for the Anomaly Guard Module:

- Active-Active cluster
- Active-Standby cluster

The Active-Active cluster is operationally identical to a cluster configuration that is deployed to provide multigigabit capacity. All zone traffic that is diverted to the Anomaly Guard module cluster will be load-shared among all available Anomaly Guard modules.

The Active-Standby cluster is created by configuring different relative RHI weights for each Anomaly Guard Module in a cluster. In a simple Active-Standby cluster with two Anomaly Guard modules, one module will insert a static route (via RHI) into the supervisor engine routing table

with a higher weight (lower cost) than the other module, so that this first module acts as the Active Guard while the second module acts as the Standby Guard. In the event of an Active Guard failure, RHI will withdraw the static route entry for the first module, leaving the Standby Guard as the new Active Guard.

3. ENTERPRISE DEPLOYMENT

Anomaly Guard and Traffic Anomaly Detector modules can be deployed in an enterprise network to protect Internet-facing servers and network infrastructures from the debilitating effects of a DDoS attack.

A typical enterprise deployment has an Anomaly Guard Module and Traffic Anomaly Detector Module installed on a Cisco Catalyst 6500 Series switch that is the first point of entry for inbound Internet traffic. During nonattack periods, inbound traffic from the Internet flows through line cards and the supervisor engine without any involvement of the Anomaly Guard Module, while a copy of this inbound traffic (via SPAN or virtual access control lists [VACLs]) is sent to the Traffic Anomaly Detector Module for analysis.

If anomalous traffic behavior is detected by the Traffic Anomaly Detector Module, the module will signal the Anomaly Guard Module to start its mitigation process. This mitigation process consists of:

1. Diverting (“hijacking”) the affected traffic (all traffic destined for a targeted IP address that is under attack) from the normal path to the Anomaly Guard.
2. Subjecting that traffic to multiple layers of analysis and countermeasures to distinguish legitimate sources from attack sources (“cleaning” the traffic).
3. Dropping the attack traffic and forwarding the legitimate traffic (“injection”) back into the normal traffic path for forwarding to the target (zone).

Internet Link Bandwidth

A first point to consider in an enterprise deployment is whether there is enough link bandwidth from the ISP to the enterprise premises to allow an effective deployment of an Anomaly Guard Module downstream from that ISP link. In order to effectively counter high-bandwidth DDoS attacks, the Anomaly Guard mitigation device must be deployed far enough upstream to drop the attack traffic before it can saturate the ISP link.

Given the size of high-bandwidth DDoS attacks that have been observed on the Internet, it is recommended that if an enterprise has Internet bandwidth of greater than 100 Mbps, the Anomaly Guard Module should be deployed downstream of its Internet link in an “on-premises” deployment. If the link bandwidth is less than 500 Mbps, the Anomaly Guard Module should be deployed upstream of the ISP link in a colocation or a managed service arrangement.

On-Premises, Colocation, or Managed Service

The primary consideration for on-premises or off-premises deployment is Internet link bandwidth.

For on-premises deployment, the Anomaly Guard and Traffic Anomaly Detector modules will typically be installed in the same Cisco Catalyst 6500 Series switch, although this is not mandatory (Anomaly Guard and Traffic Anomaly Detector modules can reside in different switches, as long as the Anomaly Guard is logically upstream of the Traffic Anomaly Detector). Figure 4 shows a network diagram of an on-premises enterprise deployment.

Figure 4. On-Premises Enterprise Deployment

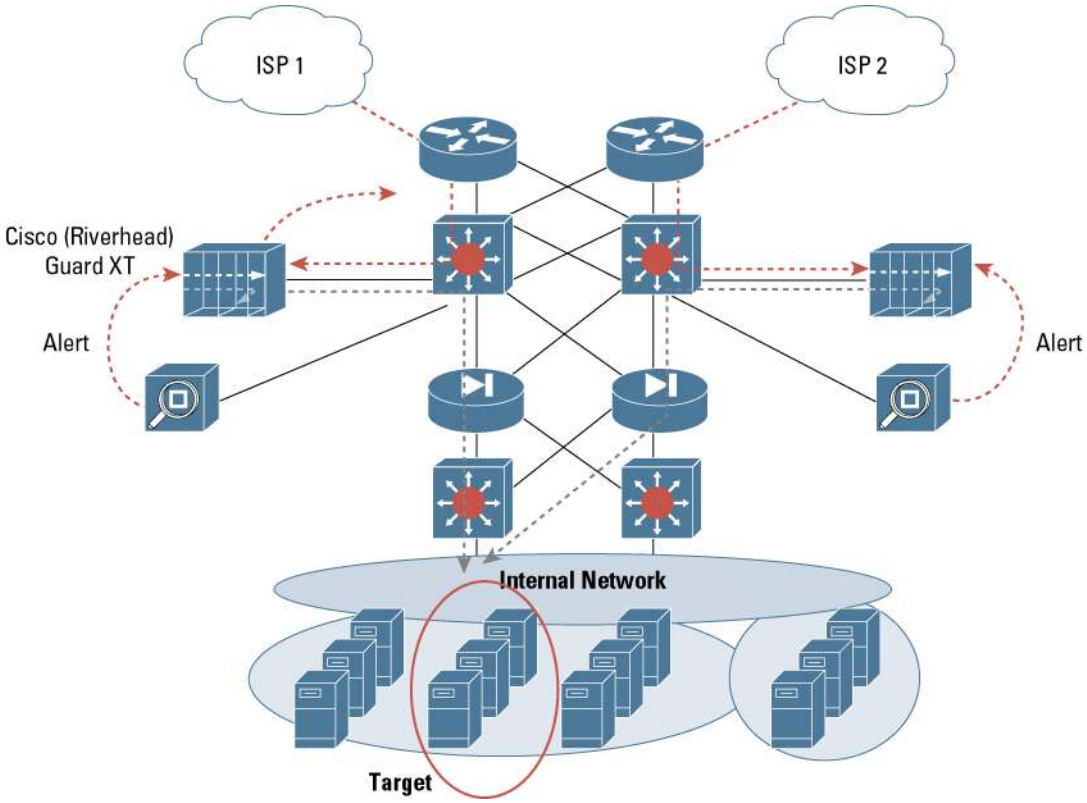

A colocation deployment places the Anomaly Guard Module upstream of the Internet link at a service provider's point of presence (POP), but the Traffic Anomaly Detector Module can reside either at the colocation switch or on an enterprise premises switch.

Managed service arrangements have different deployment options, depending on the service provider. In any managed service option, the Anomaly Guard Module will be deployed inside the service provider network. The Traffic Anomaly Detector Module can be deployed either in the service provider network or on the enterprise premises, or can be replaced by other methods of attack detection using IDS- or NetFlow-based systems.

4. HOSTING PROVIDER MANAGED SERVICE DEPLOYMENT

A hosting provider deployment is one type of managed DDoS protection service. The other type, described below, is a managed service provided by an access service provider. The goal of a hosting provider deployment is to offer a managed security service as a value-added enhancement to the core Web hosting service. Anomaly Guard and Traffic Anomaly Detector modules can be easily incorporated into an existing hosting provider's Cisco Catalyst 6500 Series switching infrastructure to provide DDoS protection for any or all hosted customers.

Deployment Considerations

Shared Use of Anomaly Guard and Traffic Anomaly Detector

- In hosting centers that service small to midsize customers, operators will deploy Anomaly Guard and Traffic Anomaly Detector modules that will be shared among multiple hosted customers. This normally requires the configuration of individual zones for each protected customer. As hosting providers design their DDoS detection and mitigation infrastructures, they need to be aware of the performance capacities—a total of 500 zones can be configured, but only 30 can be concurrently active.

Given that each zone can contain multiple destination IP addresses (each representing a server) that are either host destinations (/32 mask) or subnetted destinations, multiple customers can be configured onto one zone. The most significant configuration constraint is that active zones cannot have duplicate IP address space. Therefore, one can potentially configure DDoS detection (Traffic Anomaly Detector) and mitigation (Anomaly Guard) for thousands of customers on each individual module. Apart from the active zone limit, total module bandwidth capacity, as described above, is the other factor to consider when designing a hosting provider managed service deployment.

5. ISP MANAGED SERVICE DEPLOYMENT

Internet service providers (ISPs) can deploy DDoS protection for their subscribers as a managed service offering to complement their Internet connectivity services. In an ISP-managed DDoS service, Anomaly Guard modules will always be deployed in the service provider network. Traffic Anomaly Detector modules can be deployed on the subscriber premises or within the service provider network, depending on the type of attack detection system being used.

Deployment Considerations

Centralized or Distributed Attack Mitigation

There are two options for deploying Anomaly Guard Modules—centralized or distributed deployment.

The centralized model is optimal for ubiquitous DDoS managed service deployments, where all subscribers and the service provider's peering ingress points are being protected by Anomaly Guard Modules. The distributed model is easier to deploy for service providers that are introducing the managed service to a subset of customers over a staggered introduction process.

Centralized Model

In the centralized deployment model, Anomaly Guard clusters are installed in a single central location on the service provider backbone network or in multiple regional locations (Figure 5). These Anomaly Guard "scrubbing centers" are engaged to clean traffic from any source to any destination on the service provider backbone. The destinations are addresses (representing servers or network infrastructure that are potential attack targets) of ubscribers who have signed up for the DDoS protection service.

Attack detection in the centralized deployment model can be accomplished using a different detection mechanism than the Traffic Anomaly Detector odule. Given the broad mitigation coverage afforded by a centralized model (attacks destined for any subscriber destination throughout the ervice provider network will be directed to a central location for cleaning), a broadly based attack detection mechanism would be best provided by a NetFlow-based system, such as Arbor Networks' Peakflow SP product.

In a centralized deployment, Anomaly Guard modules can be dedicated for individual subscribers, shared among all subscribers, or a combination of edicated and shared. However, the centralized model is well suited for a ubiquitous managed service deployment, where all subscribers share the entral Anomaly Guard cluster.

A variation of a single centralized cluster is the deployment of multiple regional clusters that would more efficiently process anomalous traffic in a arge backbone network.

Design recommendations for the centralized model include:

- Sufficient backbone network capacity (bandwidth) exists to transport the anomalous traffic (good and bad traffic destined for a target that requires cleaning) from ingress points to the Anomaly Guard cluster.
- Scalable traffic injection methods are configured for forwarding the cleaned (legitimate) traffic from the Anomaly Guard cluster to the intended destinations (using MPLS-VRF rather than GRE tunnel's more efficient router operation, for example).

- The deployment and use of a NetFlow-based, widely distributed traffic anomaly detection system such as Arbor's Peakflow SP.
- Configuring interoperability between the Arbor Peakflow Controller and the Anomaly Guard clusters. Please refer to the document titled "CPOC Case Study on Guard-Arbor Integration" for more details on interoperability between Arbor Peakflow and Cisco Anomaly Guard modules.

Figure 5. A Centralized Service Provider Deployment

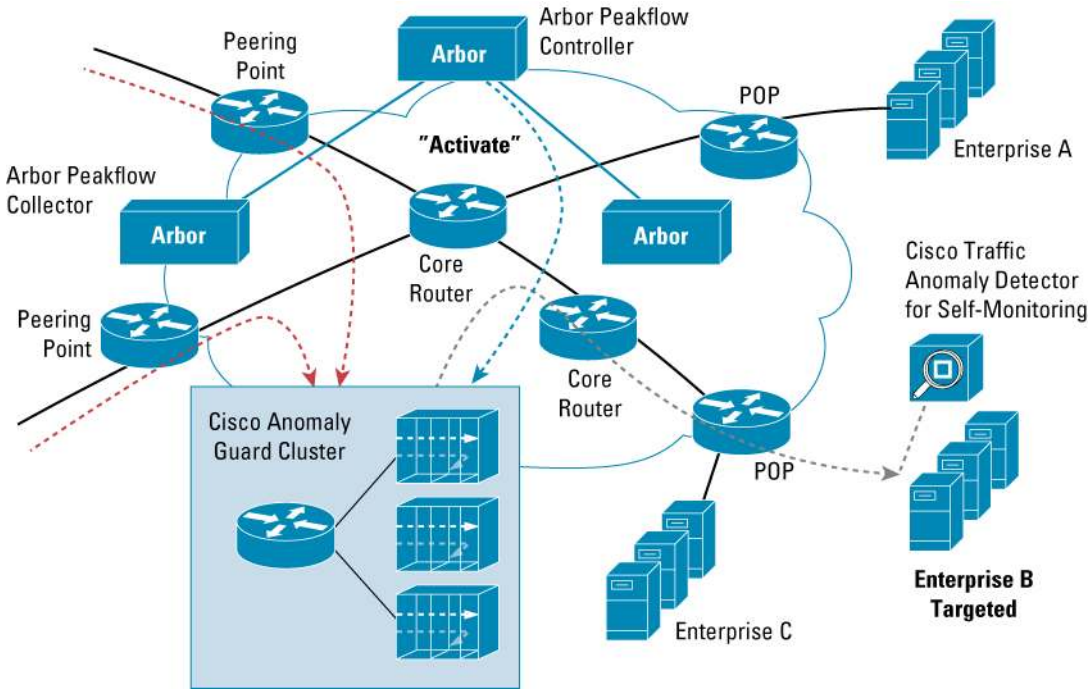

Distributed Model

In a distributed deployment model, Anomaly Guard modules and clusters are distributed at service provider POPs that connect directly to subscribers (Figure 6). In this deployment model, the modules or clusters will be installed at the POPs serving the subscribers who have signed up for the managed service. These Anomaly Guard modules are usually dedicated for those subscribers' exclusive use.

The distributed model allows more options for using the traffic anomaly detection mechanism, including:

- Using Traffic Anomaly Detector modules at the POP for each subscriber
- Using Traffic Anomaly Detector modules at the subscriber premises
- Using a combination of NetFlow-based traffic anomaly detection on the service provider network and Traffic Anomaly Detector modules on the subscriber premises

In general, the NetFlow-based mechanism scales better over a large network (analyzing NetFlow data from all relevant router interfaces) but provides less-thorough detection (sampled NetFlow), whereas the Traffic Anomaly Detector Module provides more-thorough detection (looks at all packets) but relies on SPAN ports or VACL capture for receiving traffic for analysis.

When Traffic Anomaly Detector modules are deployed at the service provider POP, it is possible to configure them for shared use such that there is not a strict one-to-one relationship between a Traffic Anomaly Detector Module used for attack detection and an Anomaly Guard Module used for attack mitigation. If designed within the capacity limits described above, a single Traffic Anomaly Detector Module can provide attack detection for multiple subscribers that are using dedicated Anomaly Guard modules for attack mitigation.

Design considerations for the distributed model include:

- When using a Traffic Anomaly Detector Module on subscriber premises for attack detection and automated alerting of appropriate Anomaly Guard modules, an out-of-band management network must be configured between the subscriber Traffic Anomaly Detector and the service provider's Anomaly Guard.
- In a combination deployment where both a NetFlow-based system and Traffic Anomaly Detector modules are used for attack detection, the NetFlow-based system can be used for course-level detection further from the target and Traffic Anomaly Detector modules can be used for finer, more thorough detection closer to the target.

Figure 6. A Distributed Service Provider Deployment

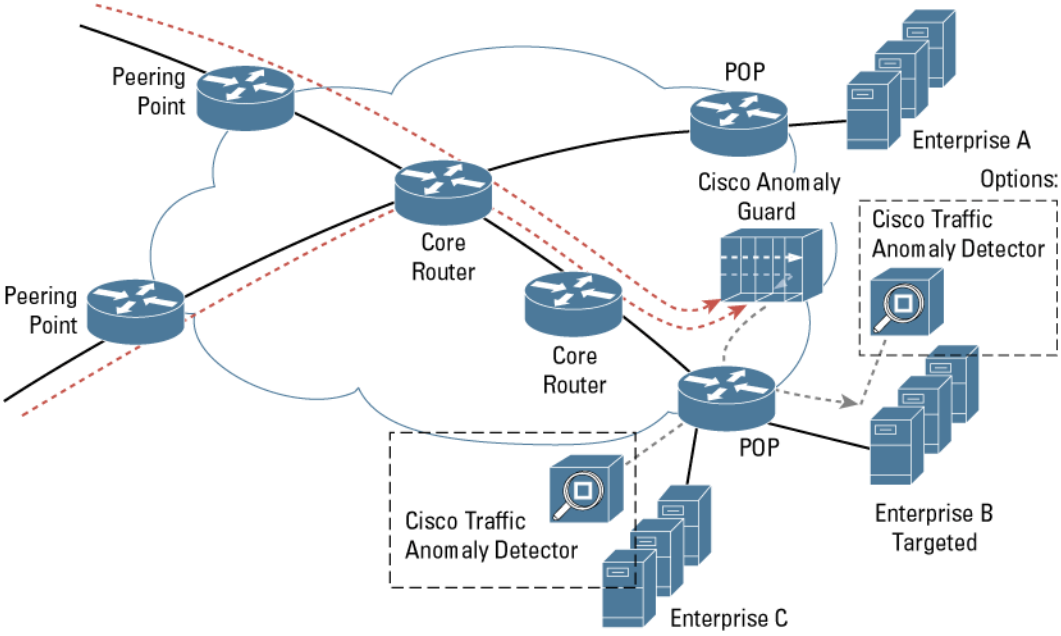

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

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

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

204177.t_ETMG_MH_2.05

Printed in the USA