Technical Note

# Centralized Protection

## Riverhead Long Diversion Method Using MPLS LSP

As distributed denial-of-service (DDoS) attacks escalate, both in terms of frequency and severity, businesses who rely on uninterrupted online availability are demanding their service providers protect them against such assaults.

Delivering such services can be costly for providers, who support everything from large, predictable enterprises to small, low-volume web sites.  While the large enterprises, which concentrate high levels of traffic at single entry points, deserve dedicated DDoS defenses, the smaller, lower-volume customers — who access the provider infrastructure through multiple points of presence — can't justify the expense of dedicated devices at each entry point.

Riverhead Networks overcomes those problems with a unique "centralized protection" feature that allows service providers to consolidate smaller numbers of Riverhead Guards in a single location, creating a centralized, cost-effective solution for protecting lower-volume customers from DDoS attacks.

The key to centralized protection is a concept called "long diversion."  When an ISP customer is under attack in a Riverhead-protected environment, traffic addressed to the victim is long-diverted from all peering routers through the centralized Guard, which removes bad or malicious packets before injecting clean traffic back onto the main path towards the original destination.

This document describes how Riverhead's long diversion mechanism protects ISP networks from DDoS attacks in a pure Multiprotocol Label Switching (MPLS) network with IS-IS running between edge devices.

Riverhead
networks

## Riverhead DDoS Protection

Before describing long diversion, it is first important to understand how the Riverhead solutions protect mission-critical web sites against DDoS attacks.

When a DDoS attack is launched, a Riverhead Detector (or other third-party device) recognizes the threat and immediately alerts the Riverhead Guard to initiate mitigation services. All traffic destined for the targeted device is then diverted off the normal path through the Riverhead Guard, which applies a series of patented technologies and algorithms based on Riverhead's Multi-Verification Process™ (MVP) architecture to identify and remove malicious packets while allowing legitimate transactions to pass.

The diversion process, which is unique to Riverhead, preserves business continuity by rerouting only affected traffic flows and then analyzing those flows to remove attack traffic while allowing "real" requests to complete, reducing the impact on business operations.

## Long Diversion Method

Unlike Riverhead's standard diversion technique, where the Riverhead Guard diverts traffic only from an adjacent router, the long diversion method diverts traffic from remotely located peering routers that reside several hops away from the centralized Guard. This document will examine how diversion is implemented in an ISP backbone that implements MPLS. Figure 1 illustrates the general overview of long diversion, in which R2, R3 and R4 are peering routers, while R1 is the router adjacent to the Guard.

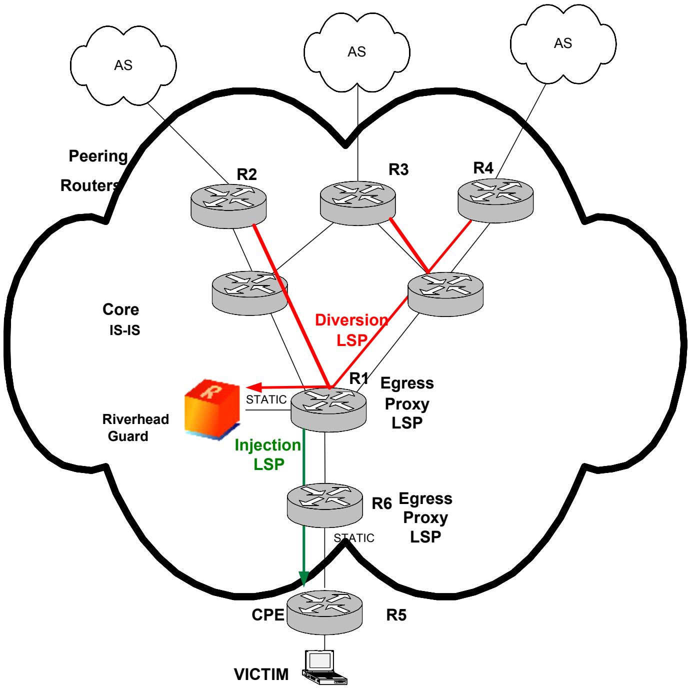

*Figure 1: Long diversion using the Riverhead Guard.*

The long diversion process is divided into three parts:

- **Diversion:** Diverting the victim’s traffic from the peering routers (R2, R3, R4) to the Guard.
- **Cleaning:** The act, by the Guard, of removing malicious packets and forwarding “clean” packets.
- **Injection:** Returning clean traffic back on route to the Victim.

The following sections will elaborate on each of these tasks.

## Configuration

The example configuration used in this document includes the following network elements:

- Peering routers (R2, R3, R4)
- A Riverhead Guard
- A router to which the Guard is attached (R1)
- An edge router of the victim (R6)

# Diverting Traffic to the Riverhead Guard

When an attack launched against a specific victim has been detected, diversion is achieved by the Riverhead Guard sending out an iBGP announcement stating that, in order to reach the victim, the traffic should be routed to the Label Switching Protocol (LSP) path that ended at the Guard's loopback interface.

To ensure that the BGP announcements will not propagate into all the backbone routers' routing tables, **no-advertise** and **no-export** BGP is applied on the community strings. As a result, only R2, R3 and R4 will receive the BGP announcement about the victim (with longer prefixes), with next hop to the corresponding Guards loopback interface.

## BGP Announcement

The Guard sends out an iBGP announcement (with **no-advertise** and **no-export**) to routers R2, R3 and R4 telling them that the next hop to the victim's destination is the Guard's loopback interface. This process is achieved using route map, which manipulates the next hop attribute in the BGP announcement. The announcement would use a longer prefix then the original victim announcement, and therefore would get priority over the original BGP announcement.

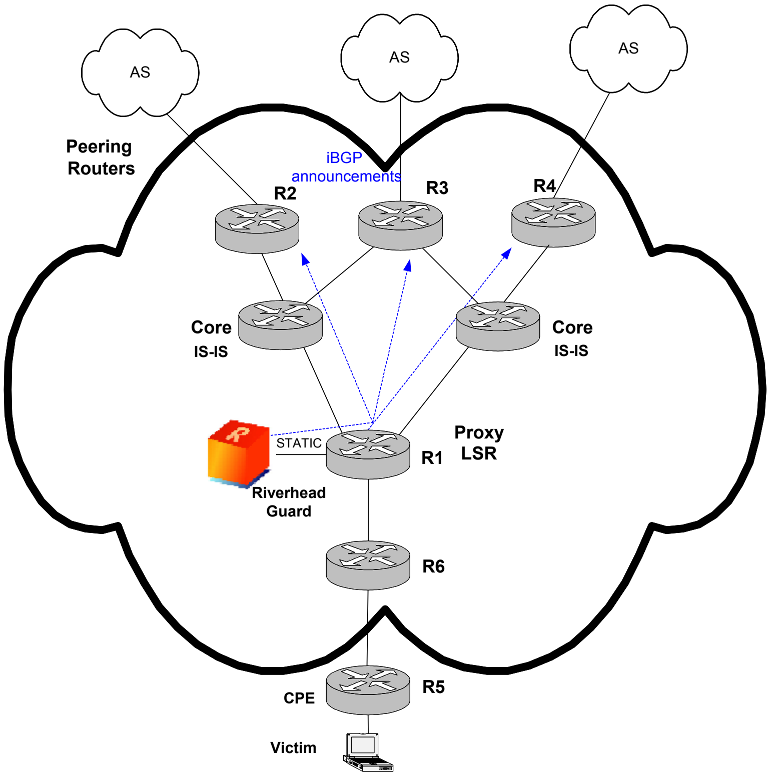

*Figure 2: iBGP announcement to the peering routers.*

## MPLS LSP

After the iBGP announcement reaches the peering routers, the routers reroute the victim's traffic to the LSP path leading from the peering points to the Riverhead Guard's loopback interface.

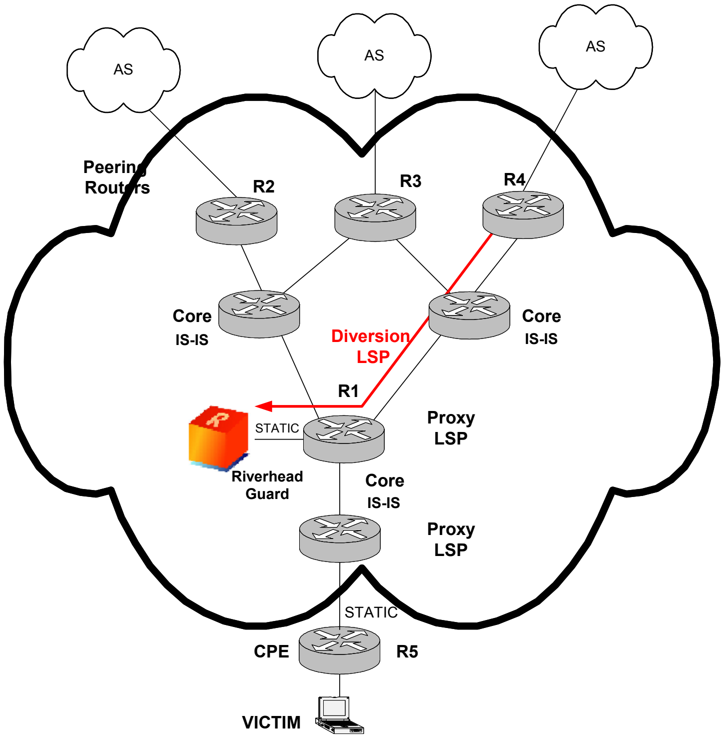

*Figure 3: MPLS LSP from the peering routers (R2, R3, R4) to the Guard.*

While the Guard is at the end of the LSP path, it is not required to run MPLS; it only receives pure IP due to the fact that R1 is the egress proxy LSP path for the Guard. In other words, R1 performs penultimate hop popping on the MPLS packets that arrive at the Guard's loopback and shoots them directly to the Guard over the static route.

The Guard's loopback should be routable via IGP in the entire network. In order to achieve this, R1 is configured with a static route to the loopback of the Guard, which will redistribute this static route with the IGP protocol (IS-IS in the example). Note that the Guard does not run IS-IS.

Each of these detectors monitors the packet rate and compares it to the peace-time rate (profiles of normal or typical behavior built by the Riverhead products during non-attack periods) or to a threshold set manually by the user.

## Injection of Traffic from the Guard to the Victim

Once the Riverhead Guard has cleaned the traffic, it injects good packets back to R1. In this scenario, it is assumed that R1 stores all routes to all possible victims, just as if it were a peering router. Hence, it would forward traffic to the victim using the suitable LSP path.

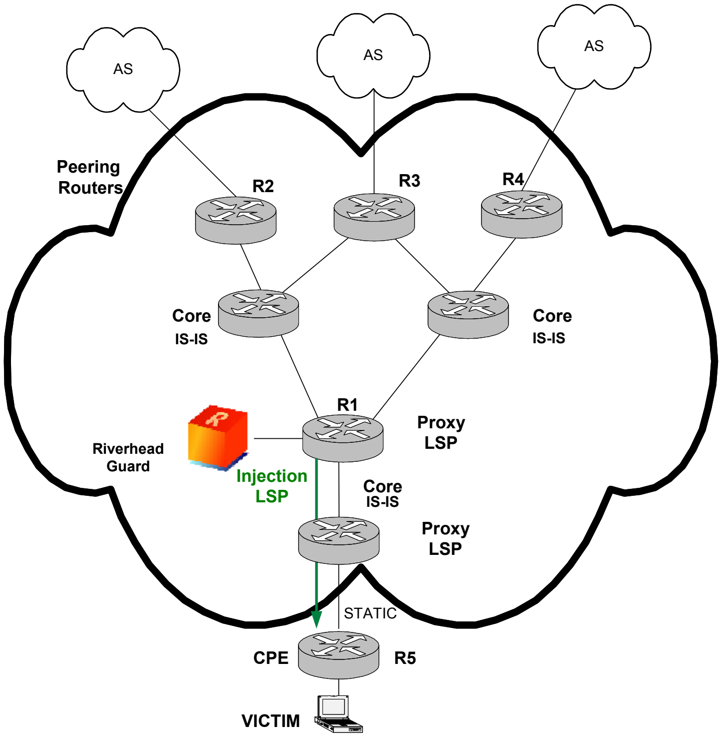

*Figure 4: Injecting the traffic to the Victim.*

## Caveats and Limitations

In order for the proposed solution to work correctly, there are few points that must be taken into consideration.

### Router (R1) Connected to the Guard

When forwarding traffic back to the victim, the traffic is injected to R1, which then performs IP lookup. Router R1 should have a route to the victim (in fact, the router should have routes to all possible victims). Note that for a regular core router, it is enough to store routes to all the loopbacks of the routers in the network.

Router R1 cannot be a peering router.

## Backbone Capacity

The ISP's backbone infrastructure must be capable of handling the volume of the attacked traffic.

## MPLS Enabled

MPLS needs to be implemented on the backbone infrastructure. Note that there are several other tunnel techniques that can also be implemented (such as GRE).

## Topology Assumption

In a situation where the LSP from R1 to the victim ends at an edge router (for instance, R6), and R6 does not implement egress proxy LSP, then the Guard cannot divert traffic from that router (in other words, R6 cannot be a peering router). If, however, the LSP from R1 to the victim ends in the customer premises equipment (CPE), then the Guard can divert traffic from R6 (i.e., R6 can be a peering router). Note that an LSP may end in CPE, even if the CPE doesn't support MPLS, by using R6 as an egress proxy LSP.

# Appendix A: Example of Packet Flow

If the victim is not currently under attack, traffic flows in its normal route to the victim's IP addresses (based on the loopback address that holds the LSP).

Once an attack has been launched, the network operator activates the Guard to start protecting the attacked Victim. The following steps will automatically take place:

1. The Guard informs the peering routers (R2, R3, R4) about the new route to the victim, with next-hop of the Guard loopback interface.
2. The victim's traffic, which is received by the peering routers (e.g., R2), is routed over the RED LSP and is received by the victim.
3. The Guard forwards the clean traffic to R1.
4. R1 performs IP lookup and routes the packets to the victim using the appropriate LSP.

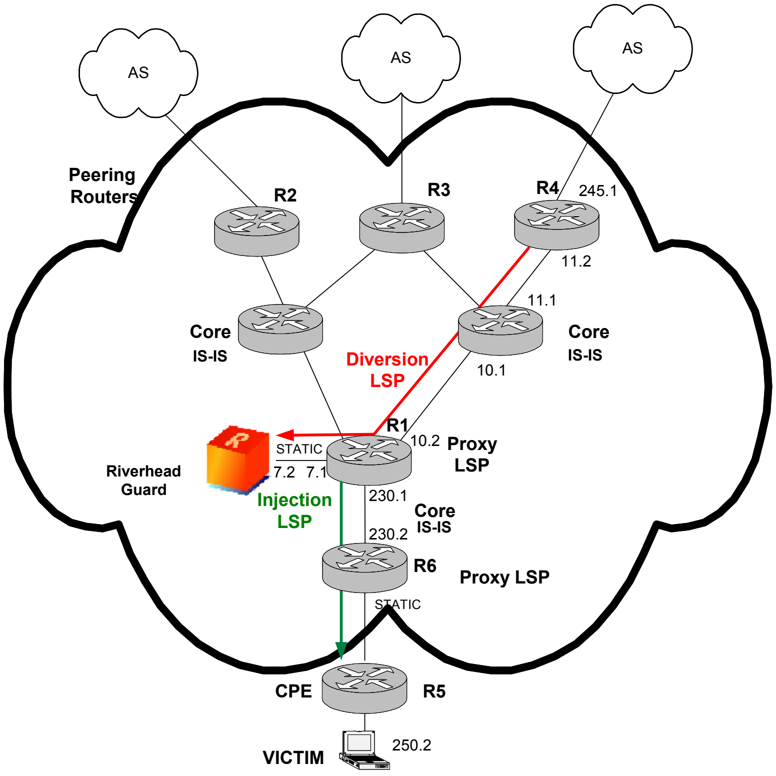

## Peering Router Configuration (Routers R2, R3, R4)

**MPLS Global Configuration:**

```
mpls ip
ip cef
```

**Interface Loopback 0:**

```
interface Loopback0
ip address 3.3.3.3 255.255.255.255 (used to build the LSP via IS-IS)
no ip directed-broadcast
load-interval 30
```

**Connectivity Interface:**

```
interface FastEthernet5/0
ip address 192.168.11.2 255.255.255.0
no ip directed-broadcast
load-interval 30
full-duplex
tag-switching ip (MPLS enabled)
no cdp enable
ip rsvp bandwidth 512 512  (RSVP enabled)
```

**IS-IS Configuration:**

```
router isis
redistribute static ip
net 49.0001.0000.0000.0003.00
```

**BGP Configuration:**

```
router bgp 1680
no synchronization
bgp log-neighbor-changes
redistribute connected
neighbor 192.168.8.16 remote-as 1680 (iBGP session)
neighbor 192.168.8.16 description ---<<< iBGP to the Guard >>>---
neighbor 192.168.8.16 soft-reconfiguration inbound
```

**Router Configuration (R1):**

(Only relevant commands)

**Loopback Interfaces:**

```
interface loopback 0
ip address 2.2.2.2 255.255.255.255  (used to build the LSP via IS-IS)
no ip directed-broadcast
```

**Connectivity Interface (to the network):**

```
interface FastEthernet1/0
ip address 192.168.10.2 255.255.255.0
no ip directed-broadcast
load-interval 30
full-duplex
tag-switching ip (MPLS enabled)
no cdp enable
ip rsvp bandwidth 512 512 (RSVP enabled)
```

**Connectivity Interface (to the Guard-no MPLS):**

```
interface FastEthernet1/0
ip address 192.168.7.1 255.255.255.0
no ip directed-broadcast
```

### Connectivity Interface (to the Victim-MPLS):

```
interface FastEthernet0/1/1
ip address 192.168.230.1 255.255.255.0
tag-switching ip (MPLS enabled)
no cdp enable
ip rsvp bandwidth 512 512 (RSVP enabled)
```

### IS-IS Configuration:

```
router isis
redistribute static ip
net 49.0001.0000.0000.0002.00
```

### Static Route Configuration:

```
ip classless
ip route 1.1.1.1 255.255.255.255 192.168.7.2
end
```

## Guard Configuration

### BGP:

```
router bgp 1680
redistribute Guard
bgp router-id 192.168.8.16
neighbor 192.168.8.1 remote-as 1680
neighbor 192.168.8.1 description ---<<< iBGP session to peering Router >>>---
neighbor 192.168.8.1 soft-reconfiguration inbound
neighbor 192.168.8.1 route-map _new_next-hop out
!
route-map _new_next-hop permit 10 (Define the new next-hop)
set ip next-hop 1.1.1.1
!
line vty
!
end
-
```

## For More Information

For more information about the Riverhead DDoS prevention solutions, visit our web site at www.riverhead.com, call us directly at 408.253.5700, or e-mail us at info@riverhead.com.

**Riverhead Networks, Inc.**
20195 Stevens Creek Blvd.
Cupertino, CA 95014
Tel: (408) 253-5700
info@riverhead.com