

DDoS Mitigation

Maintaining Business Continuity in the Face of Malicious Attacks

Distributed Denial of Service (DDoS) attacks are a serious threat to businesses everywhere. How can organizations mitigate the effects of these vicious attacks to ensure business continuity?

Riverhead Networks delivers solutions that alleviate the impact of DDoS attacks. Based on a unique Multi-Verification Process (MVP) architecture that employs several patent-pending technologies and algorithms, the Riverhead solutions work together to:

- **Detect** even the most sophisticated DDoS attacks
- **Divert** suspicious traffic from the critical path for further analysis and purification
- **Defeat** the attack by thoroughly scrubbing attack traffic, filtering out bad packets while allowing legitimate transactions to complete

This Technical Note will focus on Riverhead Networks' mitigation capabilities — the ability to defeat DDoS attacks by quickly identifying and eliminating illegitimate traffic without affecting valid transactions, minimizing the impact on business operations. This high-level overview is not intended to provide detailed product information; rather, it is meant to convey the powerful capabilities of the Riverhead solution and how it helps businesses maintain uninterrupted operations.

The Riverhead MVP Architecture

Before describing how the Riverhead solution mitigates the impact of DDoS attacks, it is first important to understand the Riverhead philosophy of maintaining business continuity. That means, in contrast to other DDoS solutions, Riverhead is equally committed to ensuring good packets make their way through the system as they are to blocking malicious traffic.

To achieve that objective, the Riverhead solutions focus on three things: detect, divert and defeat. When a DDoS attack is detected, suspicious traffic is immediately diverted to a Riverhead Guard that resides off the critical path but is connected to key routers upstream from the targeted device. The Guard then subjects the diverted traffic to a rigorous, multi-level evaluation and analysis process designed to remove bad packets while allowing legitimate traffic to pass, defeating the attack with no impact on the business.

That filtering process, based on Riverhead's Multi-Verification Process (MVP) architecture, is what makes the Riverhead solution so unique. All DDoS prevention tools apply some level of filtering to suspicious traffic when an attack is identified. The problem with those solutions is that the thresholds designed to block malicious traffic also block a lot of legitimate traffic. Only Riverhead passes traffic through a detailed, five-step process that thoroughly separates good packets from bad, delivering a granular level of protection that no other solution can provide.

Five-step DDoS Mitigation

The MVP architecture features five separate and distinct enforcement and verification modules that, working together, provide unprecedented protection against DDoS attacks (see Figure 1). Deployed in the Riverhead Guard DDoS Defender, each module applies different technology designed to ferret out malicious traffic while allowing legitimate requests to continue through the system. The modules are tightly integrated and work in a closed-loop fashion, constantly communicating with and updating each other. The result: a highly dynamic solution that is constantly evolving and adapting to the changing landscape, offering the most advanced and impenetrable DDoS protection available.

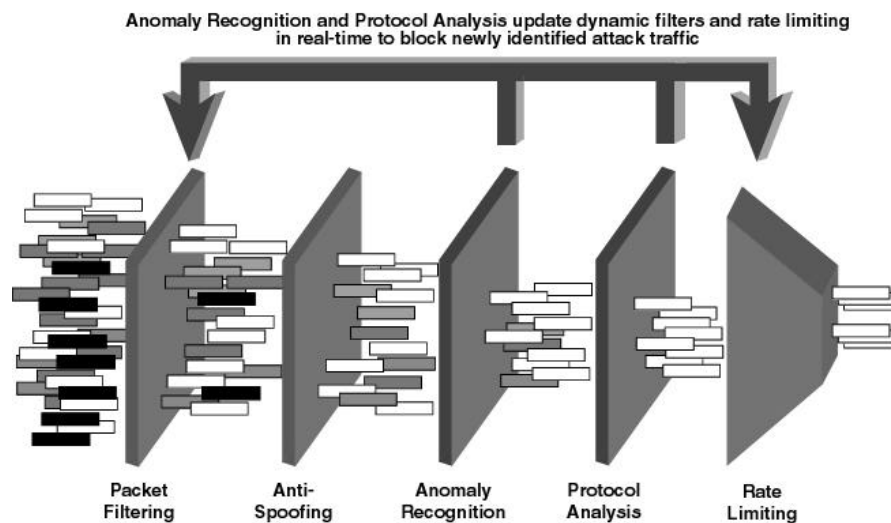


Figure 1: Riverhead Guard's Multi-Verification Process (MVP) Architecture

The five Riverhead MVP architecture modules include:

Packet Filtering: The packet-filtering module performs simple packet inspection, enforcing a basic set of rules on traffic destined for the target device. This set of rules includes two types of filters: static and dynamic. The Static filters, which are configurable, are designed to simply block non-essential traffic from reaching the victim. These are activated only in response to an attack condition.

Dynamic filters are “new” filters that are inserted in real-time by other modules within the architecture to block sources and flows that these other modules have identified as malicious through behavioral analysis of incoming traffic. The Dynamic filters are an excellent example of the closed-loop relationship between the various modules within the MVP architecture.

Anti-spoofing: The anti-spoofing module employs challenge and response mechanisms to verify that each incoming packet has actually been generated by a real application located on the device possessing the packet’s source IP address. Many DDoS attacks use spoofed traffic to evade detection by passive packet inspection and flow monitoring solutions. To thwart these types of attacks, the Riverhead solution employs a variety of patent-pending anti-spoofing mechanisms to distinguish “spoofed” packets from “non-spoofed” ones. It also uses unique “stateless” techniques to avoid becoming a victim of these resource-robbing attacks itself. The anti-spoofing module also employs mechanisms to ensure that no legitimate packets are falsely identified as spoofed, and are processed quickly.

Anomaly Recognition: The anomaly recognition module constantly monitors traffic flow on a per-client and per-protocol basis, looking for any deviation from baseline or normal behavior. The per-flow analysis is superior to global-only trending that cannot distinguish between valid increases in customer traffic and malicious attack sources. If a deviation is found, the module recognizes it as a potential DDoS attack and responds based on default and customized policies. Responses include operator notification and activation of dynamic filters and rate limits in the other enforcement modules. Because it compares current activity to historical traffic behavior as observed by the Riverhead device, the anomaly recognition module can identify DDoS attacks that have never been seen before — a real advantage for proactive DDoS protection.

Protocol Analysis: The protocol analysis module selectively processes flows that the anomaly recognition module finds suspicious in order to identify application-specific attacks such as http-error attacks. Protocol analysis then detects any misbehaving protocol transactions, including incomplete transactions or errors. Although it contains an anomaly detection feature, the Protocol Analysis module focuses on Layer 7 headers looking for suspicious flows.

Weighted Fair Queuing: Once traffic has passed through the filtering, anti-spoofing, anomaly recognition and protocol analysis modules, it is assumed to be legitimate. The weighted fair queuing module serves as a throttle for traffic being returned to the network. Per-protocol and per-flow limits can be set as an alternative enforcement action. Global and per-flow limits are also useful for releasing only as much legitimate traffic as the target can handle to ensure sufficient availability of resources for each flow.

Critical Sequence

While the protocol analysis module plays an important role in the overall traffic scrubbing process, it is really more an extension of the anomaly recognition module, looking for protocol-specific deviant behavior. When broken down to its critical base elements, the Riverhead DDoS mitigation process really consists of four major steps, as shown in Figure 2.




Figure 2: The Riverhead Guard Protection Diagram

The order in which these functions occur is critical to the way the Riverhead solution operates. For instance, it is important that anomaly recognition occurs after spoofed traffic has been removed from the stream by the anti-spoofing module. Anomaly recognition determines if any single source is behaving differently from a normal baseline; if a compromised or hacked source is utilizing IP address spoofing, it is trying to hide this abnormal traffic flow from a single source by making it look as if it is coming from many sources. The anti-spoofing module will defeat this evasion technique. Additionally, if this is not done, a hacker could use spoofed packets to skew the input to the anomaly recognition module and cause it to identify innocent sources as malicious. In other words, anti-spoofing removes the ability of attack sources to evade detection by anomaly recognition, and anomaly recognition removes the ability of any non-spoofed attack source from utilizing a volume attack.

After having passed through the filter, anti-spoofing and anomaly recognition modules, packets reaching the weighted fair queuing stage are assumed to be legitimate. However, because it is impossible to recognize a “bad” traffic flow without monitoring it for a few seconds, the WFQ module may still encounter misbehaving flows, and it is that module’s responsibility to carefully regulate the flow of traffic released back into the network to avoid overwhelming the victim.

Riverhead DDoS Mitigation: How It Works

So far, this technical note has described the overall Riverhead MVP architecture, the five components of the architecture, and what each component does.

This section will describe in greater detail how each of those modules work, providing unique insight into what makes Riverhead DDoS prevention solutions so unique.

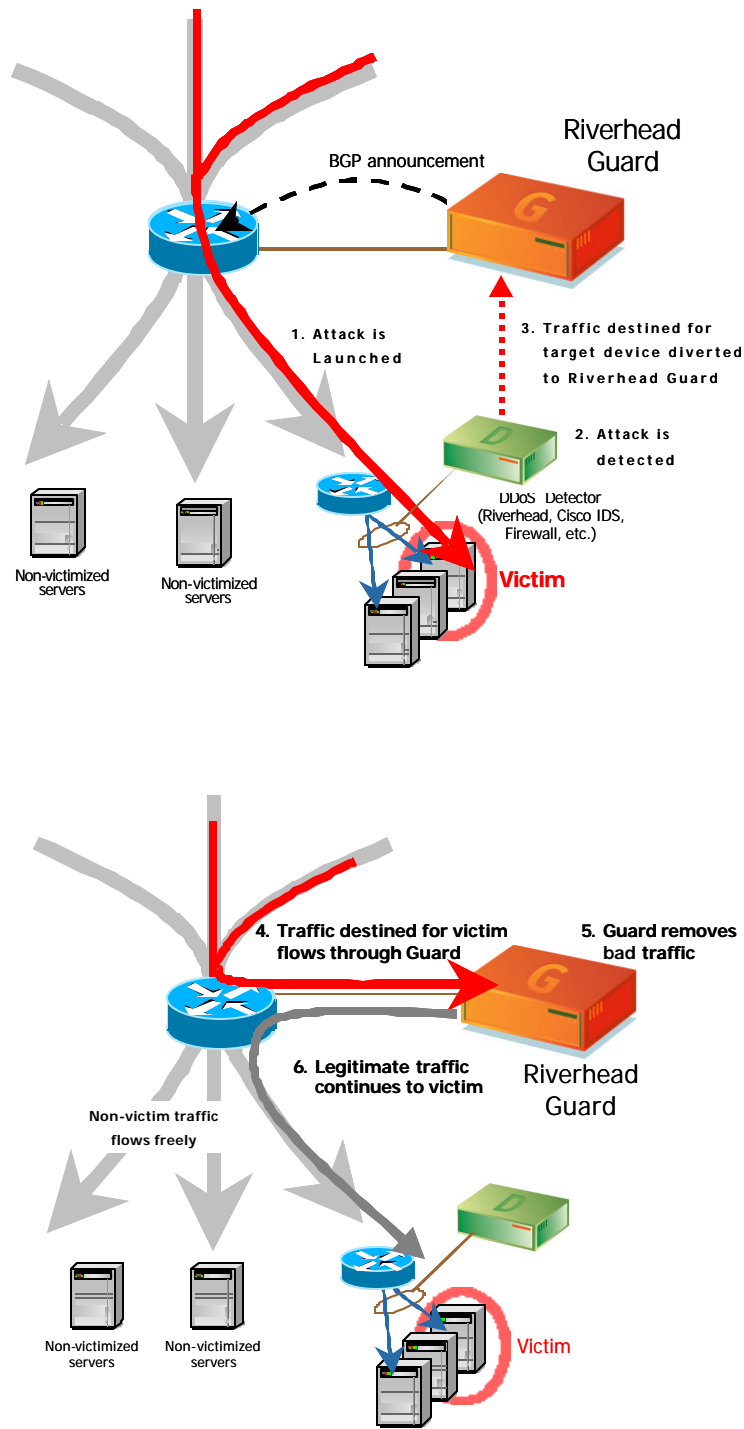


Figure 3: When a DDoS attack is launched (Step 1, top) and detected (Step 2), traffic destined for the targeted device $\frac{3}{4}$ and only that traffic $\frac{3}{4}$ is diverted to the Riverhead Guard (Step 3). Traffic flowing through the Guard (Step 4, bottom) is subjected to a rigorous multi-verification process to remove bad packets (Step 5) while allowing legitimate traffic to pass unimpeded (Step 6).

Modes of Operation

First, it is necessary to understand that the Riverhead Guard typically works in one of three modes of operation on a per-flow basis.

- **Protect mode:** In protect mode, the Riverhead Guard actively filters out malicious traffic and forwards legitimate traffic to the victim.
- **Analysis mode:** In analysis mode, the traffic of a potential victim flows through the Riverhead Guard, but the Guard does not do anything to the traffic except monitor its behavior for anomaly recognition purposes. If it notices an abnormal traffic pattern, the Guard will then switch to “protect” mode.
- **Learn mode:** In learn mode, traffic is diverted through the Riverhead Guard so it can learn the normal behavior of different connections and clients to establish a baseline profile. Once the profile is built, the operator interacts with the Guard and may adjust or accept any of the suggested parameters.

Packet Filtering

When a DDoS attack is launched, it is essential that only the relevant applications (protocols) be allowed to pass through the system while other non-essential (or illegitimate) traffic is blocked. The Filtering module serves that purpose, acting as a simple firewall of sorts for DDoS attacks.

Anti-spoofing Mechanisms

Once suspicious traffic is filtered, it is passed to the anti-spoofing module.

As described earlier, in spoofed packets, the source IP address in a packet’s header is bogus. Whether it was randomly generated or calculated in some way to appear legitimate, it is not the real source IP address of the machine from which the packet originated.

Spoofed packets may be of different protocols and types, such as UDP, ping packets, DNS packets, SYN packets from the first phase of the TCP handshake, etc. The Riverhead solution protects against different types of spoof attacks, including:

- Spoofed TCP-based SYN packets
- Spoofed DNS request packets
- Spoofed TCP packets (other than SYN)
- Spoofed UDP packets (other than port 53)

TCP-based SYN Packets

Riverhead uses a dozen different algorithms to ensure that no spoofed packets make it past the anti-spoofing module. Each algorithm is designed for one of the classes of traffic described above and features a specific level of rigor and strictness.

For the purposes of this document, we will use spoofed TCP-based SYN packets to illustrate how this module works. The three-way handshake that SYN uses to establish connections is one of the more popular denial-of-service attack methods. In this type of attack, a huge number of spoofed SYN requests are sent to a targeted server. Since each request must be buffered and kept by the server for a set period of time (typically 30 seconds) until its corresponding SYN-ACK is received, the server’s SYN request buffer quickly fills up. At best, this causes the server to ignore genuine requests to open new connections; at worst, the server fails completely.

To prevent such attacks, the Riverhead device intercepts the SYN request on behalf of the server and performs the three-way handshake with the client. Only after a correct SYN-ACK response is received, confirming that the request is valid, does the Riverhead device open the corresponding connection. If a valid response is not received, the request is considered malicious and is dropped.

Failures are avoided because the Riverhead device is designed to handle and hold open hundreds of thousands of legitimate concurrent connections even while authenticating thousands of new requests per second. Its unique algorithms do not maintain state for malicious connections; therefore, it is not subject to the same failures as the targeted servers and network devices. In a redundant ISP configuration, the effects of an attack are further mitigated because the load is distributed across multiple Riverhead devices, increasing the number of concurrent sessions that can be handled.

DNS Spoof Protection

The vast majority of DNS request-and-reply traffic on the Internet is over UDP, which is a connectionless protocol. Since this protocol does not have built-in message verification sources, it is relatively easy to spoof DNS messages and “choke” a DNS server with a multitude of bogus requests. The Riverhead solution uses proprietary, patent-pending methods to protect DNS servers from such attacks using transparent verification and DNS authentication protocols that protect root servers, top-level domain name servers, and other name servers that are lower in the hierarchy.

Anomaly Recognition

Once all spoofed traffic has been removed, remaining traffic is passed to the anomaly recognition module. The anomaly recognition module monitors the traffic, breaks it down per source and per flow, and compares it to baseline behavior in search of deviations that would indicate whether the source or a flow is behaving maliciously.

This sequence of events is crucial. Spoofed traffic must be removed first to ensure all remaining traffic is coming from valid sources. If the anomaly recognition module detects a pattern that differs from the traffic’s “normal” behavior as observed by the Riverhead device while in Learn mode, it assumes the source is malicious and provides operational rules for blocking the attack without disturbing “innocent” traffic.

The basic principle behind the module’s operation is that the pattern of traffic originating from a “black-hat” daemon residing at a source differs from the pattern generated by such a source during normal operation. In contrast, traffic patterns of innocent sources during an attack resemble their traffic at normal times. This principle is used to identify the attack source and type and provides guidelines for blocking traffic (either by the filter or by other modules in the multi-verification process model) or, conversely, for further analysis of suspected data streams by refined monitoring in the anomaly recognition module.

The parameters monitored by the anomaly recognition module for individual streams of packets include, but are not limited to, the following:

- The volume of traffic from an attacking daemon
- The distribution of packet sizes and port numbers
- The distribution of the packets’ inter-arrival times
- The number of concurrent flows
- Higher-level protocol characteristics
- The ratio of in-band to out-of-band traffic.

For any of these parameters, any departure from normal behavior may indicate that the source (client) is an attacking daemon. Policies define which of these parameters, or triggers, should result in an enforcement action against the attacking source when a threshold is crossed.

Network Flows and Traffic Classification

The basic element studied by the anomaly recognition module is a flow. Each flow is a sequence of packets belonging to the same connection. Generally, a flow is identified by the following parameters:

- Source IP address
- Source port
- Destination port
- Protocol type
- Connection creation — time of day; day of week

The destination IP address is implied, since all the information is collected per destination address. For each such flow the traffic volume is registered.

However, it is not feasible to store all the above information, since it consumes an unacceptable amount of memory. To record these key parameters in a succinct way, Riverhead employs machine-learning methods to identify the basic characteristics of the traffic traveling to each destination.

The learning method studies the typical behavior of groups of users interacting with the destination. For example, a typical web site is accessed either by individual users sitting behind a host (PC); by a group of users sharing one multi-user timesharing host; or by a group of users sitting behind a proxy. Other types of users are possible, such as web crawlers (used in search engines) and monitoring servers such as Keynote.

Traffic Monitoring and Analysis at Attack Time

To detect malicious traffic, the Riverhead solution monitors and classifies the victim's incoming and outgoing traffic and compares it to normal activity. Two major network flow properties are used to identify malicious events:

- Traffic patterns
- Traffic volume

Recognition of Traffic Patterns

The Riverhead solution examines several aspects of traffic patterns, including the following:

Source “IP Geography” Proximity	Sources are classified to resemble the “IP Geography;” i. e., IP addresses residing on neighboring networks (using an IP address prefix) are classified together. A class that generates a relatively large volume of requests is suspected as malicious. Note that such “malicious classes” are likely to form if the attacker planted a collection of daemons in the same network that does not use a proxy.
Periodicity	Sources are examined for the distribution of packet inter-arrival times. It is expected that the inter-arrival time distribution of a malicious daemon differs from the distribution of an innocent source (user or proxy); it is likely that malicious sources act in a relatively consistent manner, while innocent sources exhibit a more irregular pattern.
Packet Properties	Packets of a source are examined for repetitive properties, e. g. distribution or packet size. It is likely that malicious sources generate identical properties (i.e., all packets a similar size) while innocent sources generate packets of a more random nature. Other properties include port number distribution.
Protocol-Specific Properties	Each protocol may have distinct properties to examine; e.g. HTTP (error and refresh rate) and DNS (error rate).

Recognition of Traffic Volume

Traffic volume recognition is used to identify malicious sources that transmit large volumes of data over long periods of time that are significantly different from normal levels. Specifically, Internet data sources are classified as either small or large sources. Small sources relate to individual IP addresses whose traffic volume is normally miniscule.

Large sources relate to Proxy/NAT traffic or Spider (Crawler) traffic whose volume is considerably higher. (Traffic volume resulting from a Spider access is normally higher than that of a single human user, especially on large web sites; a Spider scans the entire site, generating hundreds or even thousands of requests, while a human client requests, on average, a maximum of tens of pages.)

“Zooming”

In addition, the Anomaly Recognition module uses the zooming principle in which it first monitors traffic at an unrefined level and, once an exceptional volume is detected for a certain type of traffic (e.g. UDP), a more refined statistical detection is invoked to identify more specific parameters for the offending traffic.

With TCP, for example, volumes destined for a certain network that are above “normal” levels cause the volumes of different port numbers, and/or different sources, to be examined in order to detect the offending source or the characteristic port number. Or if HTTP is observed deviating from expectations, per-source monitoring is again activated for each source and the error rate and request-packet length are monitored. A source with an exceptional error rate would then be blocked.

These are just two simple examples of how the Riverhead solution uses zooming. While there are countless other examples of how the product utilizes this tremendously rich feature, they are far too detailed to cover in this document.

Policies and Enforcement Options

To discover and defeat malicious sources, traffic characteristics are compared with baseline “profiles” of normal network operation. Policies specify thresholds for any of the various triggers, or analyzed properties, and the resulting action taken if the threshold is crossed. Potential actions include notify, insert dynamic filter to block, upgrade level of protection, and apply rate limiting, among others. Therefore the Guard supports either automated responses or notification with information and recommendation with operator authorizations before proceeding. Policies are designed to ensure legitimate traffic is not blocked, potentially at the expense of passing some malicious

traffic. However, since DDoS attacks depend on a high volume of traffic to achieve their goal, a small amount of malicious traffic is relatively harmless. The result of this step is a sequence of filters that block subsequent malicious traffic.

Each filter has two parameters:

- **Network flow:** Identified by a combination of source IP address (which can be prefixed), destination IP address (may be one or more), destination port number, and protocol type. (One may consider blockage that disregards port numbers, i.e. all the traffic originating from a compromised IP address, be it a proxy or a host.)
- **Duration:** Identifying the duration for which that class of traffic is blocked.

Termination Detection

Another function of the dynamic filtering module is to monitor the continued matching of traffic to the filters set to block attack traffic. When all malicious flows have ceased for a predetermined period of time, then traffic diversion is turned off and the Riverhead Guard reverts to standby/learning mode.

Weighted Fair Queuing

After an incoming packet has passed through the various filters and processing phases of the Riverhead multi-verification process filter, it is sent to the weighted fair queuing (WFQ) module — the final module in the process — before being forwarded to its destination. The WFQ module acts as a throttle by carefully regulating and restricting the amount of outgoing traffic, releasing it in a round-robin fashion to ensure sufficient resources are available without overwhelming the target device. Traffic flows are based on the victim's receiving capability.

The true value of the WFQ module, however, is that it is specifically tailored to perform sophisticated, per-flow DDoS traffic shaping to prevent traffic from overwhelming victim devices. Normal traffic is subject to occasional bursts, and unless traffic is monitored for a fair length of time, it is difficult to determine whether a “bursty” condition is normal or a sustained traffic spike attributable to a DDoS attack. The WFQ module releases traffic in an orderly, measured fashion, ensuring all pending transactions are released back into the network in a fair manner.

If the queues fill up, or are nearing capacity, an additional function provides a packet-dropping policy. The security administrator's role is to apply priorities to the various protected servers. For example, if a specific protected server (victim) has a maximum capacity of 60Kpps of a certain type, the security manager can configure the WFQ accordingly, while adding prioritization based on different protocols — for instance, assigning a higher priority to the VoIP traffic over data traffic such as ICMP or SMTP.

Summary

The Riverhead MVP architecture is designed to thoroughly examine and analyze suspicious traffic and remove malicious packets while returning legitimate requests to the network. Unlike other DDoS prevention tools, Riverhead MVP technology establishes multiple lines of defense, subjecting suspect traffic to a rigorous five-step purification process to ensure that malicious traffic is removed from the network.

At the same time, the Riverhead solution ensures that valid or “innocent” packets continue unimpeded to their final destination, delivering a new level of business continuity even in the face of large-scale attacks. No other solution provides the same level of scrutiny — or accuracy — when it comes to separating the good traffic from the bad.

Key Differentiators For Riverhead DDoS Mitigation Solutions:

- Multi-Verification Process (MVP) architecture ensures business continuity, even in the face of severe DDoS attacks

- Recognition technology, based on advanced machine-learning research, detects any deviation from “normal” behavior that would indicate a DDoS attack is underway
- Only traffic destined for the target device is diverted and inspected, minimizing the impact on business operations
- MVP technology subjects suspicious traffic to rigorous, five-step process to separate good traffic from bad traffic more thoroughly and accurately than any other solution
- Anti-spoofing technique employs a dozen algorithms to eliminate spoofed traffic

For More Information

For more information about the Riverhead DDoS prevention solution, 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, Suite 110
Cupertino, CA 95014
Tel: (408) 253-5700
info@riverhead.com