

CISCO CRS-1 SECURITY

The Cisco® Carrier Routing System and its distributed and modular Cisco IOS® XR Software microkernel architecture supports highly secure, continuous system operation through embedded instrumentation, access control, and process isolation across management, control, and data planes.

THREATS TO SERVICE PROVIDER PROFITS

For service providers, network security is a matter of business survival. Security incidents due to viruses, intrusion, operator error, and software configuration error can involve extensive associated costs and consequences such as service disruption, financial loss, dissatisfied customers, reduced productivity, and even media attention. To protect their revenue and profits, service providers must protect their infrastructures and offer managed services for secure connectivity, threat defense, and endpoint protection.

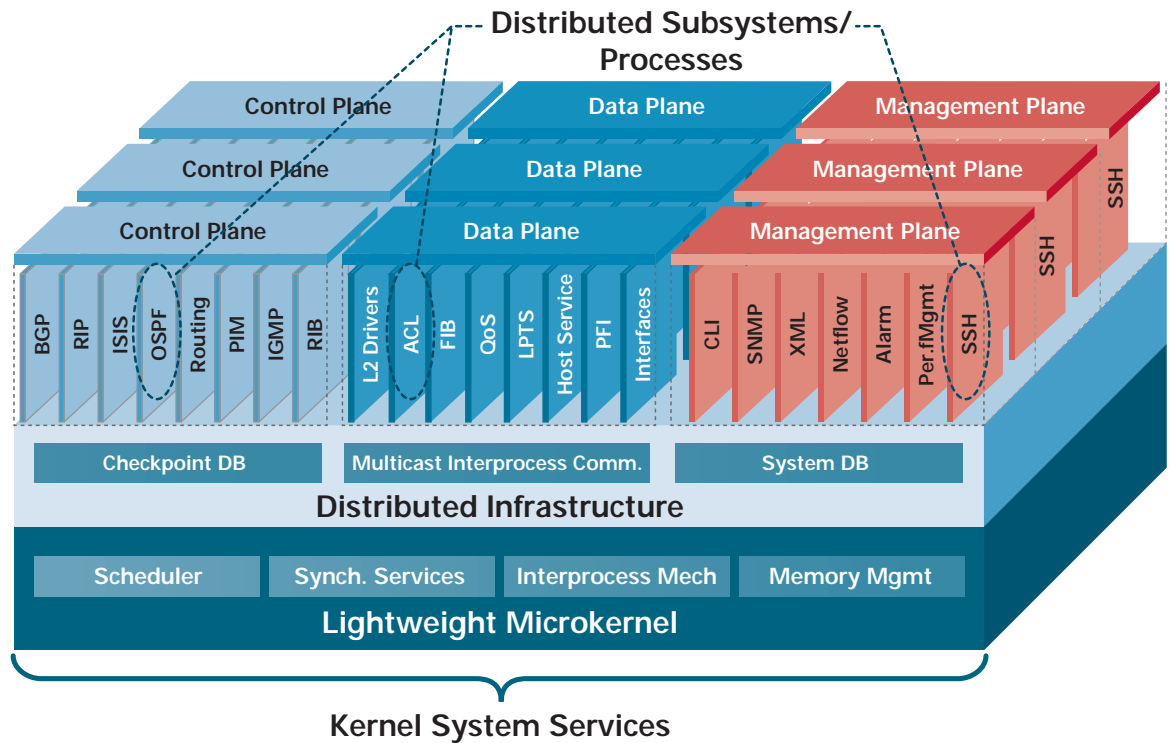
To maintain high availability in an environment of increasing security threat (for example, distributed-denial-of-service [DDoS] attacks) and policy complexity, service providers are looking to new routing and switching solutions—solutions that offer effective and embedded, hardware-based security instrumentation that enables self-defending networks. These new continuous system operation solutions must support:

- Service separation, fault isolation, and memory protection
- Seamless software and hardware recovery
- Configuration and management protection
- Proactive and rapid response

THE CISCO CARRIER ROUTING SYSTEM

The Cisco Carrier Routing System (CRS-1) is a multi-shelf routing platform based on a modular and distributed microkernel operating system, Cisco IOS XR (Figure 1).

Figure 1
Cisco IOS XR Software Architecture



When designing CRS-1, Cisco Systems® developers leveraged years of Cisco IOS Software Internet security experience. The modular, physically and logically distributed architecture of Cisco IOS XR Software (Figure 1) offers tremendous advantages in creating a highly available, secured routing platform and network. Discrete software components (subsystems) are implemented as separate software processes, running in their own protected memory address spaces. This enables true fault isolation and compartmentalization in the event of a security incident, by preventing faults in one subsystem from negatively affecting others.

Unlike monolithic kernel architectures such as FreeBSD UNIX, network stacks such as TCP run as separate processes outside the microkernel. As a result, even if the TCP stack is compromised, the system continues to operate. To resume service, the relevant processes are restarted automatically without human intervention. In addition, the modular software architecture and in-service software upgrade (ISSU) support allow, for example, a patch to be installed quickly without a complete system reload.

A key to the deep fault isolation and implementation of security instrumentation within Cisco IOS XR Software is its logical distribution of processes across three planes, each with its own access control for secure network operation:

- Control plane
- Data plane
- Management plane

Control Plane Protection

The control plane is where all routing control information is exchanged, making the control plane and its components a target. Because control plane resiliency depends on CPU processing power and scalability, “out-of-resources” attacks against the CPU are not uncommon.

To support scalability and performance, the CRS-1 control plane is designed with distributed and redundant route processors that use symmetric multiprocessing (SMP) CPUs. Under normal operations, CRS-1 transit traffic is processed by its line cards at wire rate. However, exceptions occur when packets are directed to the router itself. These “punted packets,” which include routing protocol, Internet Control Message Protocol (ICMP), and network management packets, are directed from the line card packet processor to either the line card CPU or route processor CPU.

To safeguard the control plane against DoS attacks in an open environment, multiple, layered security features are distributed to the line card and its packet processors. These features include:

- Dynamic control plane protection (DCPP)
- Automatic control plane congestion filter
- Control plane time-to-live (TTL) sanity check (RFC 3682, Generalized TTL Security Mechanism (GTSM))
- Border Gateway Protocol (BGP) routing protocol filtering and Route Policy Language (RPL)

Dynamic Control Plane Protection

Unauthorized or deliberately malicious routing updates caused by violations such as an intruder diverting or analyzing network traffic can compromise network security. Implementing neighbor router authentication with Message Digest Algorithm 5 (MD5) is a common way to avoid spoofing, and it virtually ensures that the router receives reliable information from a trusted source—but it is only a first step. If spoofed BGP packets start spraying toward the router, receive-path access control lists (ACLs) and modular QoS CLI (MQC) rate limits control exactly where these packets can proceed. However, ACL and MQC controls are not automated. If BGP peers go down or restart, the Layer 4 port number changes with each session reestablishment. As a result, network designers have been asking for an automated, dynamic way to permit configured BGP peering sessions and drop non-configured sessions.

In response, CRS-1 offers a DCP scheme for line card packet processing. With DCP, explicitly configured BGP peering sessions are automatically allocated adequate resources, whereas non-configured sessions are rejected or given minimum treatment. This permit-deny model is based on the association of statically configured IP addresses and dynamic Layer 4 port numbers. Prior to authentication and establishment for maximum admission control, different resource policies exist for initial connections. Control plane packets have to go through multilayer, prescreening schemes until they are authorized through an internal lookup table and allocated adequate resources. This automation frees time spent by network administrators on manual configuration for use on other mission-critical tasks.

Automatic Control Plane Congestion Filter

Under extreme DoS or DDoS attacks that cause line cards to exceed CRS-1 slot capacity, control mechanisms perform at hardware application-specific integrated circuit (ASIC) rate, beyond line card capacity, to drain packets into the Silicon Packet Processor (SPP) on the Layer 3 Modular Services Card (MSC) and assure control plane packet-processing priority. This feature maintains topology while the network administrator uses other security tools to install mitigation schemes to solve the problem.

Control Plane TTL Sanity Check (RFC 3682, GTSM)

Most control protocol peering sessions are established between adjacent or directly connected routers. Prior to GTSM (formerly known as BGP TTL Security Hack [BTSH]), BGP packets directed at the router from non-directed peering points had to be processed by the router CPU. When enough of these packets were generated, it effectively created a massive DDOS attack that exhausted CPU resources. Now, with GTSM, a TTL check on BGP peering packets can effectively block all non-directed BGP spoofing in MSC SPPs.

These techniques may also be applied to many other applications, such as Label Distribution Protocol (LDP) and Resource Reservation Protocol (RSVP), which can take advantage of the features of generalized GTSM. Because of the fully programmable MSC architecture in CRS-1, GTSM support for other application protocols can be easily added to MSCs.

BGP Routing Protocol Filtering and RPL

BGP is one of the most fundamental routing protocols of the Internet. Unfortunately, if BGP is attacked without proper prefix filtering implemented, a flood of traffic “garbage” can be transmitted throughout the Internet. As a result, prefix filtering has been one of the Internet service provider (ISP) community’s best practices for years. (More information is available at <http://www.ispbook.com>.)

However, the increased complexity of routing policy and the number of peers with which each peering router must communicate has created challenges in successfully implementing prefix filtering. In response, Cisco has introduced and incorporated RPL into Cisco IOS XR Software. Developed in an effort to support large-scale routing configurations, RPL has several fundamental capabilities that enhance those found in traditional route-map and ACL or prefix-list-oriented configurations.

The first enhancement is the modular building of policies so that common policy blocks can be defined and maintained independently. These common blocks can then be applied to other policy blocks to create complete policies, thus reducing the amount of configuration information maintained. In addition, parameters can be set for these common policy blocks. This allows for policies that share the same structure but differ in specific values that are set or matched against, to be maintained as independent blocks of policy. For example, three policies that are identical except for the local preference value can be represented as one common policy that takes the varying local preference values as a parameter to the policy.

RPL also introduces the concept of sets, which are containers of similar data that can be used in route attribute matching and setting operations. There are several different set types, such as prefix-sets, community-sets, as-path-sets, and extended-community-sets, which hold corresponding groupings. These sets are analogous to the prefix-lists, community-lists, as-path-lists, and extended-community-lists from traditional Cisco IOS Software configuration, with one significant exception. Sets do not encapsulate the concepts of “accept” and “deny” that are present in their Cisco IOS Software counterparts. Sets are simply containers of data. Most sets also have an in-line variant that allows comparisons against a small number of data values that are fully specified in line, rather than having to refer to a named set that contains just a few values.

Decision making, such as whether to accept or drop routes, is explicitly controlled by policy definitions. RPL allows the user to combine matching operators (which may use set data) with traditional Boolean logic operators (“and”, “or”, and “not”) into complex, conditional expressions. All matching operations return either a true or false result. The execution of these conditional expressions and their associated actions can then be controlled by using simple “if-then, else-if, else” structures specified by the user. This allows the evaluation paths through the policy to be fully user-configurable.

With the introduction of RPL, it is expected that peering policy will become more modular and efficient than current, route-map-based peering statements. RPL helps enable the scalability necessary to make peering with thousands of peers out of a single, multi-shelf routing system such as CRS-1 a reality.

Data Plane Protection

The data plane receives, processes, and transmits network data between network elements, and represents the bulk of network traffic that passes to and through the router. To protect CRS-1 data plane traffic against well-known attacks, many default sanity checks (based on the collective knowledge of the Internet community) have been built into the CRS-1 forwarding engine. In addition, CRS-1 provides features and tools such as ACLs, Unicast Reverse Path Forwarding (uRPF), and NetFlow accounting with dedicated ingress and egress processing on MSCs:

- **ACL**—ACLs, both IPv4 and IPv6, are an important element of many router data plane applications such as packet classification, rate limiting, statistics and accounting, and a permit-deny operator on packets.

Cisco CRS-1 is designed to meet the most stringent performance and scalability requirements so that it can process ACLs at line rate under network load. For example, with 2 million routes and 500 BGP peers, CRS-1 can process thousands of ACLs and their entries without performance penalty.

- **uRPF**—Cisco CRS-1 supports uRPF (strict and loose modes) to help mitigate problems caused by the introduction of malformed or spoofed IP source addresses into a network by discarding IP packets that lack a verifiable IP source address. When uRPF strict mode is enabled on an interface, the router examines all received packets to verify that the source address and interface appear in the routing table and match the interface on which the packet was received.

uRPF loose mode is the foundation of triggered black hole filtering techniques used widely in ISP communities. In loose mode, uRPF can effectively drop DoS and DDoS attack packets based on source IP address, and rapidly push the scheme out to hundreds of routers in a very short period of time.

- **NetFlow**—Accounting is an indispensable part of network management in the areas of traffic engineering, network management, and billing. NetFlow, originally an accounting application, provides a mechanism to look inside individual packet headers, to aggregate packets into traffic classes, and to collect statistics and detailed routing information for each traffic class. Implemented within Cisco IOS XR Software, NetFlow statistics constitute a valuable database that captures network traffic behavior at a microscopic level for use in traffic engineering and security analysis.

- **Static NetFlow and packet sniffing**—Cisco IOS XR Software also supports Static NetFlow, which offers functionality beyond NetFlow. Whereas NetFlow is dynamic, Static NetFlow treats packet flow similarly to how ACLs work with data packets, but with extension fields such as source or destination autonomous-system number and Multiprotocol Label Switching (MPLS) labels. With Static NetFlow, a flow filter may be defined with extended ACLs to keep track of the packet or byte counters of a particular flow. Massive quantities of NetFlow data can be associated with an extended ACL, allowing operators to filter down to and target the exact flow they are interested in, creating another effective tool in defense against DoS and DDoS attacks.

A derivative of Cisco IOS XR Static NetFlow, in-band packet sniffing, uses the same functions as Static NetFlow such as ACL-like filtering, but can collect samples and forward them to a configured destination.

Management Plane Protection

The management plane is the logical path of all traffic related to the system management of the routing platform. In a distributed and modular environment, the management plane offers new levels of complexity, and hence, increased requirements to maintain secure access. This secure access is best achieved through:

- **Default access denied**—A known and common system vulnerability is that some protocols are enabled by default. These open ports create security loopholes that invite system break in. In response to service provider demand, CRS-1 instrumentation is designed with the default configuration for these services set to off until they are manually enabled by operators.

- **Authentication, authorization, and accounting (AAA) and encryption protocols**— All access to and from a router should be encrypted and access-controlled. Cisco CRS-1 supports AAA authentication and encryption protocols SSH, SSL, IPsec and SNMPv3. Additional control can be implemented through use of ACLs to restrict access to specific source hosts only. Each user can be clearly defined under an AAA domain to reflect user access privileges.
- **Isolation of management ports**—Core routers and switches often come with dedicated management Ethernet ports that can introduce unsecured access to the device. Cisco CRS-1 Ethernet management ports are routable, and thus controlled through AAA access control and encryption, isolating data and control plane traffic so they cannot “hop onto” each other. ACLs may be used to block hopping, and implemented efficiently across multiple ports with a few clicks using the Cisco Craft Works Interface (CWI), a value-added GUI tool specially designed for multi-shelf router management.
- **Role-based privilege model**—Because unauthorized or unskilled network operators also represent a threat to system availability, service providers request flexible ways to assign operator privileges based on user-defined criteria.

To enable a role-based privilege model with a convenient and flexible way of assigning appropriate access levels to specific operators or teams, Cisco IOS XR Software organizes operations as tasks. For example, configuration of BGP is one task and configuration of Open Shortest Path First (OSPF) is another. System reload is also a distinct task. Each task has a uniquely assigned identification number called a task-ID, with defined read or write privileges. Users may be associated with task groups to inherit the appropriate access rights. To enable security, the task-ID works with the AAA server to provide the maximum centralized control for accessing the router.

SUMMARY

DoS and DDoS attacks are part of the reality of today’s Internet, and are among the most critical threats to service provider profitability. To protect service provider profits, the building of a self-defending network is based on next-generation routing systems with embedded security instrumentation, and the success of a system and network-wide best-practices approach.

The Cisco CRS-1 distributed and modular architecture supports highly secure, continuous system operations through memory protection, service separation within logical routers, and process isolation across management, control, and data planes.

In addition to CRS-1 embedded features and recommended best practices, the Cisco Product Security Incident Response Team (PSIRT) is a global panel of experts, available 24 hours per day, who react quickly to resolve customer security incidents that involve Cisco products, as well as handle product vulnerabilities. By leveraging the expertise of the networking market leader, Cisco customers benefit from both proactive and rapid response that is unparalleled in the industry.

For further information about Cisco CRS-1 security features, contact your Cisco account team or visit <http://www.cisco.com>. For up-to-date information about Cisco PSIRT and current advisories, visit <http://www.cisco.com/go/psirt>.

REFERENCES

“CRS-1 System Overview”

<http://www.cisco.com/>

Barry Greene and Philip Smith, “ISP Essentials,”

<http://www.ispbook.com>

Vijay Gill, John Heasley, and David Meyer, “RFC 3682—The Generalized TTL Security Mechanism (GTSM)”

<http://www.ietf.org/rfc/rfc3682.txt>

Vijay Gill, "Lack of Priority Queuing on Route Processors Considered Harmful"

<http://www.nanog.org/mtg-0302/gill.html>

"Improving Security on Cisco Routers"

<http://www.cisco.com/warp/public/707/21.html>

P. Ferguson and D. Senie, "RFC 2827—Network Ingress Filtering: Defeating Denial of Service Attacks which Employ IP Source Address Spoofing,"

<http://www.ietf.org/rfc/rfc2827.txt>



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 Web site 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

All contents are Copyright © 1992–2004 Cisco Systems, Inc. All rights reserved. Cisco, Cisco Systems, the Cisco Systems logo, and Cisco IOS are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website 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. (0403R) 203254_ETMG_CC_05.04