

Cisco ASA 5585-X Adaptive Security Appliance Architecture

Product Overview

The Cisco ASA 5585-X Adaptive Security Appliance is a modular security services chassis intended primarily for high-performance data center deployments. The chassis can accommodate up to 2 Security Services Processor (SSP) or interface expansion modules in the following combinations:

- Firewall SSP in the bottom slot with IPS or CX SSP application module in the top slot
- Firewall SSP in the bottom slot with one or two half-slot interface expansion cards in the top slot
- Firewall SSP module in both the top and bottom slots

Figure 1 shows an ASA 5585-X configuration with a firewall SSP in the bottom slot and an IPS SSP in the top slot. Each SSP provides a set of 10 Gigabit Ethernet and 1 Gigabit Ethernet interfaces for network attachment as well as out-of-band management. When the top slot houses an interface expansion or an application module, all nonmanagement interfaces are controlled from the firewall SSP module in the bottom slot. When a chassis houses two firewall SSP modules, each firewall operates independently with its own set of data interfaces.

Figure 1. Cisco ASA5585-X Chassis with Firewall and IPS SSP Modules

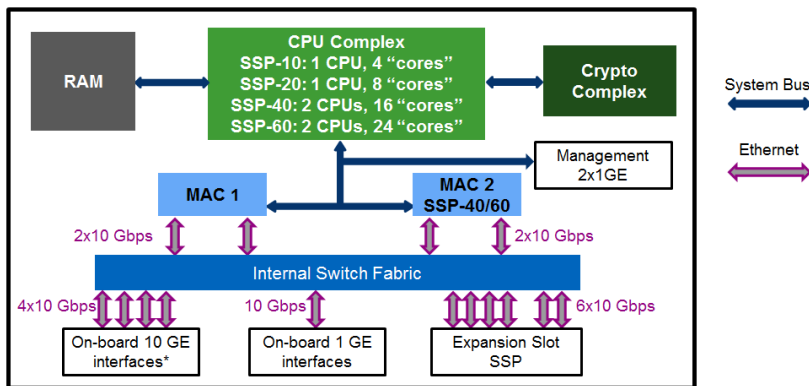


Unlike many other firewall offerings, the Cisco ASA 5585-X provides layered defense-in-depth protection to critical network services with an extendable general-purpose CPU complex and a scalable internal traffic load-balancing architecture. This design avoids having a single external Ethernet interface becoming a processing bottleneck and offers consistently high performance with any traffic profile across the entire feature set.

Internal Architecture

Cisco ASA 5585-X architecture builds on more than 10 years of the award-winning Cisco ASA appliance design. Figure 2 illustrates a basic block diagram of an ASA 5585-X firewall SSP module.

Figure 2. Cisco ASA 5585-X Firewall SSP Block Diagram



The crucial component is the flexible general-purpose CPU complex, which implements the unique software network processor (SoftNP) technology in a parallel processing fashion. Every packet entering the ASA from the network must undergo a complete set of security checks in the SoftNP or be discarded. This approach eliminates the possibility of packet leaks between interfaces that some other security devices exhibit under oversubscription. Unlike many security products that offload the bulk of the processing into static and nonextendable application-specific integrated circuit (ASIC) and field-programmable gate array (FPGA) modules, the ASA SoftNP can easily accommodate new features and protocols while delivering the same predictably high performance. The ASA 5585-X relies on intelligent hardware network interface controller (NIC) modules and specialized accelerator components for certain routine traffic load balancing as well as decryption and encryption tasks.

Unlike most other firewall designs, the physical Ethernet interfaces on the ASA 5585-X do not have direct access to the CPU complex. All incoming traffic is load-balanced across two or four (depending on the SSP model) Media Access Controller (MAC) CPU-complex uplinks through the internal switch fabric. This allows the ASA 5585-X to support high packet rates even when connecting to the network with 1 Gigabit Ethernet physical interfaces. The system can also load-balance traffic from a single physical interface across multiple internal buffering structures and processing queues, thus eliminating single-interface oversubscription scenarios. When a packet flood could easily oversubscribe a single physical interface in an older security device architecture, the ASA 5585-X mitigates such an attack by spreading the traffic across multiple internal 10 Gigabit Ethernet interfaces.

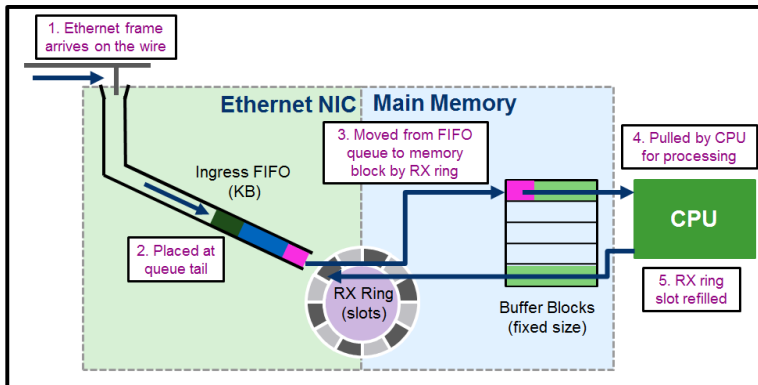
CPU-Complex Uplinks

Internal ASA 5585-X MAC links operate similarly to any external physical NIC on other ASA platforms, but they deliver additional load-balancing capabilities specifically optimized for the ASA 5585-X multiple-core architecture. The main purpose of these links is to deliver incoming packets to the CPU complex for processing as quickly and efficiently as possible. ASA 5585-X MAC uplinks implement the following internal data structures for this purpose:

- **Ingress and egress first-in, first-out (FIFO) queues:** These are simple buffers that hold incoming and outgoing packet data as it is received from or sent into the external network interfaces. These queues are sized in bytes.
- **Receive (RX) and transmit (TX) description rings:** These circular data structures are used as a communication interface between the CPU complex and the MAC links. The CPU complex populates interface descriptor rings with memory locations where incoming packets should be placed (RX) and outgoing packets should be picked up (TX).

Each MAC has two 10 Gigabit Ethernet interfaces, and each of these interfaces has an independent set of FIFO queues and RX/TX rings. Figure 3 shows a simplified data flow when a packet is received from the network into a generic ASA NIC.

Figure 3. Input Packet Flow on an ASA NIC

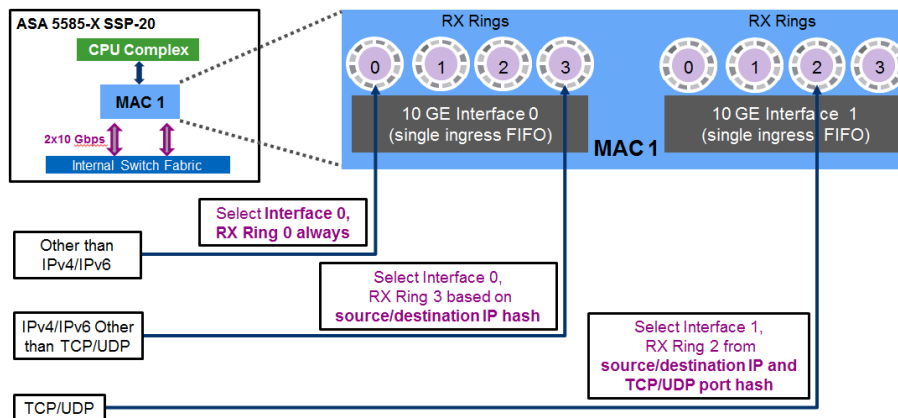


First, the frame arrives at the NIC from the external Ethernet wire. In the case of an ASA 5585-X, external physical interfaces have no FIFO queues; they simply load-balance the traffic across the CPU-complex uplinks using the internal switch fabric. An ASA 5585-X 10 Gigabit Ethernet MAC uplink thus receives the incoming frame from the internal switch fabric. The incoming frame has to be fully stored in the ingress FIFO queue before any decoding is attempted. Once the frame makes it to the head of the FIFO queue, the NIC looks up the next available slot on the local RX ring. The RX ring is prepopulated by the CPU complex with the memory locations (buffer blocks) where the incoming packets should be placed. Once an available memory slot is located, the NIC moves the frame over the system bus into the CPU-complex memory. The CPU complex will recognize the new frame and start processing it. In order to allow a continuous traffic flow from the NIC and maintain high performance, the CPU complex will refill the same RX ring slot with a different memory address.

MAC Traffic Load Balancing

The example above illustrates a single RX ring for simplicity. However, each ASA 5585-X 10 Gigabit Ethernet MAC link has multiple RX rings to properly load-balance incoming traffic. In a multiple-core ASA system, different CPU cores will service different RX rings in parallel. The total number of RX rings across all MAC uplinks in an ASA 5585-X SSP is therefore aligned to match the number of available processing cores in the CPU complex. As mentioned earlier, the external physical interfaces maintain no buffering logic. All incoming traffic is load-balanced by the internal switch fabric to two or four 10 Gigabit Ethernet CPU-complex uplinks. Consider the example in Figure 4. It depicts an ASA 5585-X SSP-20 that has a total of two 10 Gigabit Ethernet CPU-complex uplinks on a single MAC.

Figure 4. Traffic Load Balancing in an ASA 5585-X MAC



The internal switch fabric uses a hash of the source and destination IP addresses and transport ports to select the specific 10 Gigabit Ethernet CPU-complex uplink as well as the RX ring. When the transport ports are not available, only the source and destination IP addresses are used to compute the hash. For all non-IP traffic, the first RX ring of the first 10 Gigabit Ethernet MAC interface is selected. Certain network control traffic is automatically prioritized into a separate dedicated RX ring in order to eliminate any contention with the data connections. By nature of the hash, all packets in a single direction of a connection always land on the same 10 Gigabit Ethernet CPU-complex uplink and the RX ring. This load-balancing approach allows the ASA to effectively and fairly direct processing resources to all of the transit flows and head-drop frames at the RX ring level if necessary. In the worst possible scenario, the oversubscription impact from a single offending flow is limited to a single RX ring. This allows the ASA to handle the majority of transit connections even when subjected to a packet-flood attack.

CPU Complex

The Cisco ASA 5585-X general-purpose CPU complex uses multiple threads to process transit traffic flows in parallel. All but one core run data path processes, which continuously scan the memory for new packets, carry out the entire set of the SoftNP security checks, and release the permitted packets back into the network. One of the cores always runs a dedicated control plane process that handles management and network control traffic as well as more complex application inspection functions. All CPU-complex cores take turns in running the control plane process in order to achieve the best resource use. Since the control plane process typically inspects a very small portion of the transit flows, data path processes are the primary consumers of the CPU-complex resources.

Each data path process works on packets received from one interface RX ring at a time. Since the ASA 5585-X platform aligns the number of RX rings across all 10 Gigabit Ethernet CPU-complex uplinks to the available cores, the CPU complex never ends up in a situation where some data path processes are starved for new work. Different cores periodically take turns attaching to different RX rings, which increases the overall capacity-use efficiency. To preserve the packet order and help ensure accurate state checking, each stateful flow can be processed by only one CPU core at any given time. To make the resource distribution across connections even fairer, the data path processes load-balance the incoming packets across 32,000 CPU work-dispatch queues. The same source and destination IP address and transport port hash is used as when load-balancing traffic across the MAC uplinks in the NIC subsystem. All packets for a single flow always select the same work-dispatch queue, and this mechanism is used to further contain the damage from packet-flood attacks. If a particular stateful flow is generating packets at an unreasonably high rate, the ASA will limit the impact to the particular associated work-dispatch queue. Once the

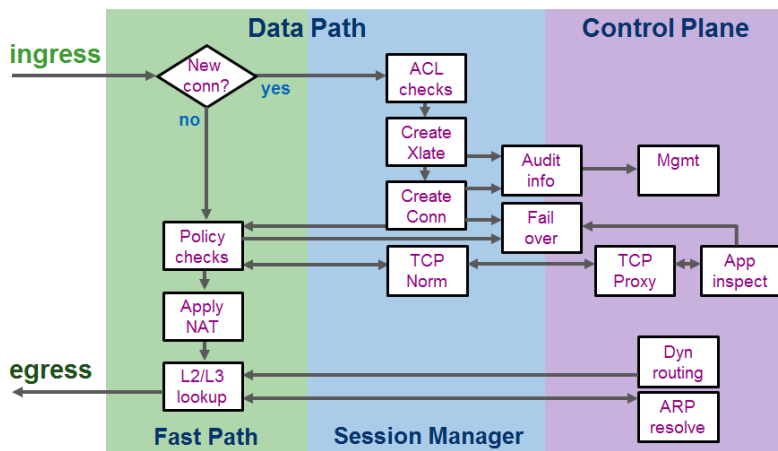
queue is full, the ASA drops any subsequent packets that hash to the same queue. As the result, the oversubscription impact from a single flow is contained to just one work queue out of 32,000; this translates into about a 0.003 percent chance of one offending flow affecting other legitimate transit connections. This process is yet another example of how the ASA 5585-X architecture is explicitly designed to contain and self-mitigate common packet-flood attacks.

As mentioned earlier, the ASA performs all packet-processing tasks in the flexible SoftNP. On the multiple-core Cisco ASA platforms, such as the ASA 5585-X, the SoftNP components are spread across the data path and control plane processes. Most of the connection-processing functions are implemented directly in the data path with the following logical components:

- **Fast path:** As the name implies, this component allows to forward packets that match already established stateful connections at a very high rate. It uses the previously evaluated security policy for the given flow to perform the full scope of stateful checks with extremely low latency.
- **Session manager:** This component evaluates the complete security policy when attempting to create the stateful connection entry. If the connection is permitted by the policy used by the first packet, the complete inspection action set is programmed into the fast path for future packet processing.

Packets that match certain connections may be escalated from any data path process to the control plane. Figure 5 provides a brief view of the functional separation between the fast path and session manager components within the data path as well as the control plane modules of the SoftNP.

Figure 5. The ASA SoftNP Logical Diagram



In this hierarchical ASA architecture, a defense-in-depth approach can be effectively implemented, where every connection is permitted or denied after the minimum necessary set of security checks. While ASA can effectively manage most security threats at the basic Layer 3 and 4 levels, advanced application inspection engines as well as IPS and CX modules can examine the permitted traffic all the way up to Layer 7 in order to stop the most complex attacks. At every step, the ASA 5585-X architecture aims at optimizing the processing resources toward potentially malicious traffic.

Conclusion

Cisco ASA 5585-X offers a modular and redundant security services architecture for high-performance data center deployments. The effective combination of the general-purpose CPU complex and hardware-accelerated components achieves consistently high performance with the current feature set as well as enabling easy future extendibility and portability. The intelligent internal load-balancing design with 10 Gigabit Ethernet CPU-complex uplinks and multiple processing queues limits the impact of common packet-flood attacks and provides uninterrupted forwarding of legitimate traffic even under system oversubscription. The defense-in-depth approach of the SoftNP architecture mitigates the most complex network attacks from Layer 3 to 7 with the minimum necessary security checks.

For More Information

https://www.ciscolive.com/online/connect/sessionDetail.wv?SESSION_ID=76600&tclass=popup

<http://www.cisco.com/c/en/us/products/security/asa-5585-x-adaptive-security-appliance/index.html>



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)