

# Implementing Access Lists on Cisco 12000 Series Internet Routers

Document ID: 40742

---

## **Introduction**

### **Prerequisites**

Requirements

Components Used

Conventions

### **Overview of ACL Support on the Cisco 12000 Series Internet Router**

ASIC-based ACLs vs. CPU-based ACLs

### **Control and Management Plane Filtering**

Configuring IP Receive Path ACLs

### **IPv4 ACL Support by Line Card Type**

Engine 0 – ACL Processing

Engine 1 – ACL Processing

Engine 2 – ACL Processing

ISE (IP Services Engine) Engine 3 – ACL Processing

Engine 4 (POS) – ACL Processing

Engine 4+ (POS and DPT) – ACL Processing

Engine 4+ (Ethernet) – ACL Processing

### **ACL Logging**

### **IPv4 Output ACL – Line Card Interoperation Matrix**

### **IPv6 ACL Support**

### **Cisco 12000 ACL Command Reference**

### **Glossary**

### **Related Information**

---

## **Introduction**

This document describes support for access control lists (ACLs) on the Cisco 12000 Series Internet Routers.

## **Prerequisites**

### **Requirements**

Cisco recommends that you have knowledge of the basics of how an ACL works on a Cisco router.

Refer to these documents for general information on ACLs and their applications:

- Access Control Lists: Overview and Guidelines
- Configuring IP Services: Filter IP Packets

### **Components Used**

The information in this document is based on Cisco 12000 Series Internet Routers.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure

that you understand the potential impact of any command.

## Conventions

Refer to Cisco Technical Tips Conventions for more information on document conventions.

# Overview of ACL Support on the Cisco 12000 Series Internet Router

On the Cisco 12000 Series Internet Router, ACLs can be processed in hardware (Application-Specific Integrated Circuit – ASIC), software (a line card s CPU), or as a hybrid feature processed in software with hardware assist. Whether an ACL is processed in hardware or software depends on the ACL application, the line card engine type, and the interaction from ACLs in other line cards.

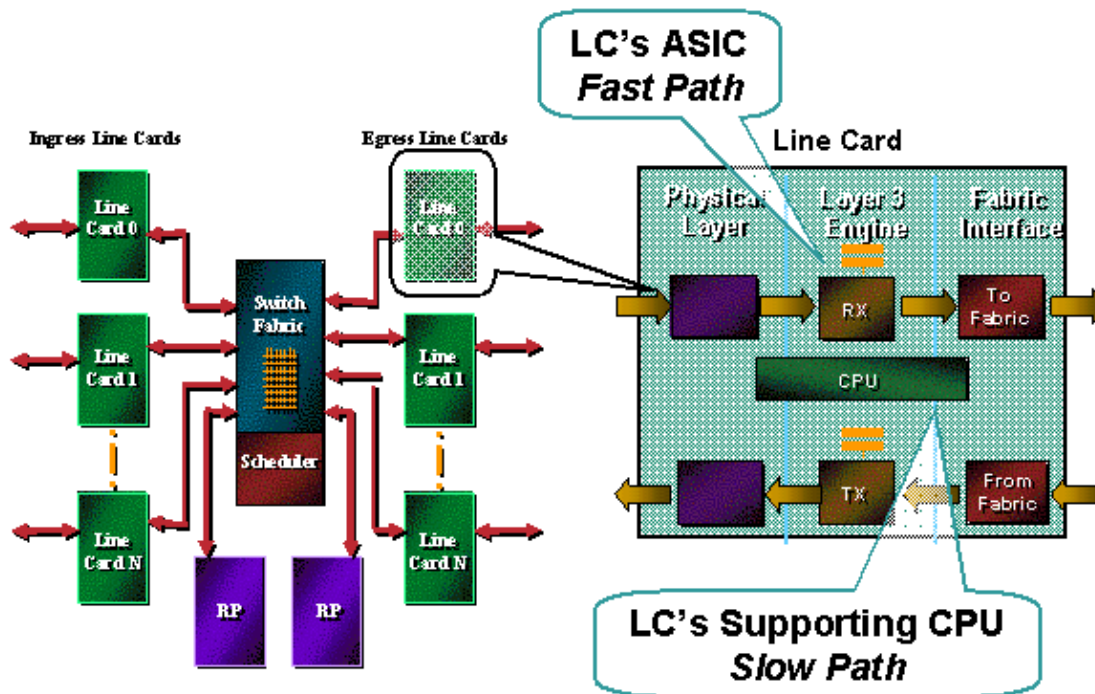
The Cisco 12000 Series line card engines provide different ACL capabilities. For ACL support information for a particular line card engine, go to the corresponding section in this document.

**Note:** IP Multicast ACLs are not supported in Cisco IOS® Software Release 12.0S. The IP Multicast boundary feature can be used where multicast filtering is required. Refer to Fast-Path Multicast Forwarding on Cisco 12000 Series Engine 2 and ISE Line Cards for more information.

## ASIC-based ACLs vs. CPU-based ACLs

The Cisco 12000 supports all generations of ACL processing. An operational understanding of how each of these processing modes work, interact, and support each other is essential to effective ACL use on the Cisco 12000.

Early generations of ACL processing used a programmable CPU to process the ACL. Over time, the packet per second (PPS) processing requirements exceeded the ability of new CPUs to keep up. ASICs were built to achieve higher PPS rates for router forwarding and feature capabilities. ACLs that were loaded on the line card (LC) CPU were then loaded onto the LC ASIC. ASICs continued to be improvised to handle higher PPS rates. These second generation ASICs have been built on the pioneering work of the generation before, and offer more ASIC capabilities. Because the Cisco 12000 is a distributed routing platform, interaction between the various generations of ACL processing can create some operational confusion.



Terms such as ASIC-based ACL, CPU-based ACL, Fast Path, Slow Path, and ASIC Punts are used throughout this document to help explain what occurs with the ACL processing. Here are explanations of these terms:

- ASIC-based ACLs (Fast Path) ACLs are loaded and processed in the ASIC hardware. The performance envelope of the ASIC determines the ACL depth, performance, and capabilities. Fast Path has been used in the path to illustrate the difference between ASIC-based processing and processing done in the LC-supporting CPU. The more generic term, ASIC-based, is used in this document.
- CPU-based ACLs (Slow Path) ACLs are processed in software on the line card CPU. For the early generation cards (Engine 0 and in some cases Engine 1), all processing is done on the LC CPU. ASIC-based LCs perform ACL processing on packets that are punted from the ASIC. Slow Path was used in the past to illustrate how punts to the LC CPU were slower than the ASIC. The more generic term, CPU-based, is used in this document.
- ASIC Punts ASICs have strict design envelopes. When a packet exceeds the designed envelope, it gets punted from the ASIC to be processed on either the LC supporting CPU or sent up to the Route Processor (RP). ASIC-based ACLs punt packets that fall outside the design of the ASIC. An example is an ACL that has an ACE with a log or log-input keyword. The information required to log the packet needs to be processed outside the ASIC, so the packet is automatically punted out of the ASIC, into the LC CPU, and processed like a normal CPU-based ACL.

**Note:** When you configure policy-based routing (PBR) with match statements to match ACLs, the ACLs should not match the source port. The gigabit switch router (GSR) does not support hardware switching for the PBR with ACLs that match the source port. It triggers process switching and GSR performance degrades.

## Control and Management Plane Filtering

The Router Processor provides control and management plane services in the distributed architecture of the Cisco 12000 Series. Receive Path ACLs (rACLs) provide a simple distributed filtering capability for control and management traffic destined for the RP. It can be logically viewed as an additional layer of security that takes advantage of the strengths of a distributed architecture.

## Configuring IP Receive Path ACLs

The rACL was introduced through a special waiver into the maintenance throttle of Cisco IOS® Software Release 12.0(21)S2. It is officially supported in Cisco IOS Software Release 12.0(22)S. Refer to IP Receive ACL for more information.

The Router Processor provides control plane services in the distributed architecture of the Cisco 12000 Series. The Receive ACLs provide filtering capabilities for control traffic destined for the RP, such as routing updates and Simple Network Management Protocol (SNMP) queries.

The rACL is considered Phase 1 of a multi-phase effort to add new protections to the control and management of plane traffic. New rate-limiting enhancements are being added through software updates.

## IPv4 ACL Support by Line Card Type

The 12000 Series line cards provide different ACL capabilities per engine type. This section describes the ACL capabilities of the different line card engines. For ACL support information for a particular line card engine, see the corresponding section of this document.

There are some general characteristics for all ACLs (ASIC and CPU based):

- Only one ACL can be applied to an interface for each direction. For example, interface POS 0/0 can have only one input ACL and one output ACL.
- Testing of the packet against an ACL stops after a match is found. If an ACL that is 300 entries long matches the packet on Access-list Entry (ACE) #45, then the packet is processed and ACL processing is stopped.
- There is an implicit **deny all** entry at the end of every ACL. As a result, if there is no match on the ACL, the packet gets dropped. Cisco ACLs are created with *explicit permit* ACL architecture. This means that there must be an ACE to match the packet for it to be processed and forwarded.
- Newly-added ACEs are always appended to the end of the ACL. Whenever the ACL requires updates, it is a good practice to remove the ACL (use the **no access-list** command) and re-add the new ACL.
- Because non-initial IP fragments do not contain Layer 4 protocol information in the IP header, only standard match criteria are supported for non-initial fragments. Full details on how Cisco ACLs comply with IP fragment filtering can be found in Access Control Lists and IP Fragments.
- Numbered ACLs are processed and applied as soon as they are entered through the command line interface (CLI). With large ACLs, this sometimes results in a CPU spike on the RP or the LC CPU.

## Engine 0 – ACL Processing

Engine 0 is the first line card delivered for the Cisco 12000. It is all CPU-based processing and forwarding. Hence, Engine 0 line cards process ACLs in the LC CPU.

These line cards are based on Engine 0:

Line Card Type	Interface Type	Connectivity
12 x DS3	Coaxial	SMB
12 x DS3	Coaxial	SMB
12 x E3	Coaxial	SMB
1xCHOC12->DS3		IR

1xCHOC12/STM4->OC3/STM1	POS	IR
4xOC3c/STM1c	POS	SR
4xOC3c/STM1c	POS	LR
4xOC3c/STM1c	POS	MM
1xOC12c/STM4c	POS	IR
1xOC12c/STM4c	POS	MM
6xCT3->DS1		SMB
2xCHOC3/STM1->DS1/E1		IR
4xOC3c/STM1c	ATM	IR
4xOC3c/STM1c	ATM	MM
1xOC12c/STM4c	ATM	IR
1xOC12c/STM4c	ATM	MM

### Supported Match Criteria

All Cisco IOS Software Release 12.0S Standard, Extended ACL, and Turbo ACLs are supported on Engine 0.

### Number of ACEs Supported

ACL size is limited only by performance requirements and available memory resources.

### Output ACL Processing

Output ACLs are processed in the ingress feature path of the other line cards in the system. A push of the Output ACL to the ingress side of the other LCs protects the backplane from forwarding packets that are going to be dropped. This is an inherited function from the distributed architecture on the Cisco 7500. A detailed explanation, reasons, and operational guidelines are provided in the IPv4 Output ACL – Line Card Interoperation Matrix.

### Line Card Specific Commands

None.

### Operational Guidelines and Line Card Interactions

- If NetFlow is configured on an Engine 0 line card and an output ACL is configured on an egress engine 3 or 4+ line card, the output ACL is processed by both the ingress and egress line cards in order to allow NetFlow to account for packets denied by ACLs as well as forwarded packets.

### Recommendations

Cisco recommends the use of Turbo ACLs on Engine 0 for large ACLs. Small Linear ACLs are more efficient for smaller ACLs because Turbo ACLs require extra memory.

## Engine 1 – ACL Processing

## Overview

The Engine 1 line card is a bridge between the CPU-based processing on the Engine 0 and the first generation forwarding/feature ASIC on the Engine 2. Engine 1 line cards process ACLs in software by default. With Cisco IOS Software Release 12.0(10)S and later, Engine 1 provides hardware ACLs for cards equipped with versions 4 or 5 of the Salsa ASIC (see the Line Card Command Reference below to determine with which version of Salsa a particular card is equipped).

These line cards are based on Engine 1:

Line Card Type	Interface Type	Connectivity
8xFE	(RJ45)	100BaseT
8xFE	(MM)	100BaseF
8xFE	(RJ45)	100BaseT
8xFE	(MM)	100BaseF
1xGE	SX,	GBIC:
1xGE	SX,	GBIC:
2xOC12c/STM4c	DPT	IR
2xOC12c/STM4c	DPT	LR
2xOC12c/STM4 c	DPT	XLR
2xOC12c/STM4c	DPT	MM
2xOC12c/STM4c	DPT	IR
2xOC12c/STM4c	DPT	LR
2cOC12c/STM4c	DPT	XLR
2xOC12c/STM4c	DPT	MM

## Supported Match Criteria

All Cisco IOS Software Release 12.0S supported Standard, Extended, and Turbo ACLs are supported in the LC CPU (Slow Path). In addition, the Engine 1 can process input ACLs in the Salsa ASIC. The Salsa ASIC handles input ACL processing along with route lookup, resulting in increased performance when compared to Traditional Linear ACL processing and Turbo ACL processing. The Salsa ASIC cannot process output ACLs or sub-interface ACLs.

## Number of ACEs Supported

ACL size is limited only by performance requirements and available memory resources.

## Output ACL Processing

Output ACLs are processed in the ingress feature path of the other line cards in the system. See the IPv4 Output ACL – Line Card Interoperation Matrix section for more information.

## Line Card Specific Commands

- **access-list hardware salsa**
- **show controller l3 | include ASIC**

## Operational Guidelines and Line Card Interactions

- The Salsa ASIC and PSA ASIC cannot be operated at the same time. The **access-list hardware** command only accepts either PSA (Engine 2) or Salsa (Engine 1) but not both.
- If NetFlow is configured on an Engine 1 line card and an output ACL is configured on an egress engine 3 or 4+ line card, the output ACL is processed by both the ingress and egress line cards in order to allow NetFlow to account for packets denied by ACLs as well as forwarded packets.

## Recommendations

For versions of Engine 1 line cards that do not support hardware ACLs, Cisco recommends the use of Turbo ACLs for large ACLs. Small ACLs (less than 20 lines) can be implemented as linear ACLs to conserve memory.

## Engine 2 – ACL Processing

### Overview

The Engine 2 was the first line card with a forwarding/feature ASIC. With Cisco IOS Software Release 12.0(10)S and later, Engine 2 line cards provide hardware ACL capabilities in the high-performance Packet Switching ASIC (PSA). As with all forwarding/feature ASICs, strict performance envelopes place boundaries on the capability of the ASIC. The key performance envelope on the Engine 2 ACLs are due to memory limitations in the PSA ASIC.

Packet forwarding in Engine 2 is done by the PSA ASIC. PSA has three main external memories:

- PLU (Path-lookup) Used to store trie nodes
- TLU (Table Lookup) Used to store FIB leaves and possibly loadbalance structures. Also used to hold many of the PSA ACL data structures
- SRAM The primary location for loadshare structures

The PSA ACL feature is a microcode-based implementation of ACL checking. A special set of instructions is loaded into the PSA chip that allows for basic ACL checking. There are a number of limitations to this feature that should be carefully understood before deploying. One major drawback to PSA ACLs is the large amount of hardware forwarding memory required.

The PSA ACL feature requires a large block of PLU/TLU memory to be pre-allocated regardless of the number of prefixes, etc. Because this allocation comes primarily from the TLU area, it has a significant impact on the number of routes that can be maintained on these cards when PSA ACLs are configured.

In addition to the initial outlay of PLU/TLU memory, each prefix stored in the TLU memory requires significantly more memory. The amount of memory required for each prefix varies, based on the direction of the ACL applied (ingress vs egress) and the linecard type. In general, egress ACLs require more memory than ingress, and linecards with more physical ports require more memory than those with fewer ports.

In the case where the Engine 2 linecard does not use ACLs, the data structures for ACL are built regardless of actual ACLs configured. In order to change to the smaller non-ACL structures, you must configure **no access-list hardware psa** on the router. This command disables all ACL processing on all Engine2 linecards in all directions. Cisco recommends to use them with extreme caution.

### Overview

In order to provide ACL processing performance that is independent of match depth, Engine 2 ACLs are integrated into the hardware forwarding table. See below for explanations on how this can impact prefix

scalability.

These line cards are based on Engine 2:

Line Card Type	Interface Type	Connectivity
1xOC48c/STM16c	POS	SR
1xOC48c/STM16c	POS	LR
1xOC48c/STM16c	POS	SR
1xOC48c/STM16c	POS	LR
1xOC192c/STM64c	Enabler	SR
16xOC3c/STM1c	POS	IR
16xOC3c/STM1c	POS	MM
4xOC12c/STM4c	POS	IR
4xOC12c/STM4c	POS	MM
4xOC12c/STM4c	POS	IR
4xOC12c/STM4c	POS	MM
4xOC12c/STM4c	ATM	IR
4xOC12c/STM4c	ATM	MM
8xOC3cSTM1c	ATM/TS	IR
8xOC3c/STM1c	ATM/TS	MM
3xGE	SX	GBIC:
3xGE	CWDM	GBIC:
1xOC48c/STM16c	DPT	SR
1xOC48c/STM16c	DPT	LR
1xOC48c/STM16c	DPT	SR
1xOC48c/STM16c	DPT	LR

### Supported Match Criteria

All Cisco IOS Software Release 12.0S supported Standard and Extended ACL match criteria, except Layer 4 source ports. Discontinuous masks, IP precedence fields, and Layer 4 source ports are punted from the PSA ASIC and processed on the LC CPU.

### Number of ACEs Supported

Up to five 448-line input ACLs in the PSA. One ACL can be configured per port. Additional ACLs are administered by the line card CPU. See the Restrictions section below for restrictions on output ACLs.

### Output ACL Processing

An output ACL configured on this line card will be performed in the ingress feature path of the other line cards in the system. See the IPv4 Output ACL – Line Card Interoperation Matrix for details.

## Line Card Specific Commands

- **access-list hardware psa limit 128**
- **no access-list hardware psa**
- **psa bypass**
- **show access-list psa detail**
- **show access-list psa summary**
- **show controller psa feature**

## Operational Guidelines and Line Card Interactions

- Fast path ACL processing requires these conditions to be met:
  - ◆ The applied ACL is within the 128- or 448- ACE limit.
    - ◇ The length must be less than 128 ACEs if the **access-list hardware psa limit 128** command is configured.
    - ◇ The length must be less than 448 ACEs when the 448-line ACL microcode bundle is required.
  - ◆ Input and output ACLs are not configured together per card.
  - ◆ Up to five output ACLs may be configured on the *router*.
- Only 128-line ACLs are supported on 8- and 16-port OC-3/STM-1 POS line cards. 448 line ACLs are supported on the 4-port OC-12/STM-4 POS, 1-port OC-48/STM-16 POS, and 3-port Gigabit Ethernet line cards.
- Input ACLs take priority in the fast path over output ACLs when both are configured concurrently on the same card (the output ACL is processed in the slow path).
- If an output ACL is configured on an Engine 2 card, and the ingress line card is Engine 0/1/2/4, an output ACL will be processed in the ingress card. For other engine types, the output ACL will be processed in the Engine 2 egress slow path.
- Output ACLs are not supported for IP-to-MPLS traffic (first MPLS label being Pushed onto an IP packet).
- ACL processing information is integrated into the hardware FIB and can impact prefix scalability. Prefix memory exhaustion is reported by memory allocation failures with the `exmem=1` signature in the accompanying log message.

## Recommendations

- ACL processing information is integrated into the CEF forwarding table, which reduces prefix scalability. Applications that do not use ACLs can disable ACL support in the CEF table and thereby increase available prefix memory by issuing the **no access-list hardware psa** command.
- The configuration of the **no access-list hardware psa** command disables all ACL processing by Engine 2 cards in addition to disabling PSA support for ACLs. It does not force software execution of ACLs. This condition also applies if the egress line card has an output ACL configured.
- The configuration of the **access-list compiled** command after the **access-list hardware psa** command converts ACEs that exceed the capacity of the PSA into a Turbo ACL. This provides optimal ACL performance for ACLs over 448 ACEs in length.
  - ◆ The default ACL microcode is 128 (as from Cisco IOS Software Release 12.0(14)S/ST). If smaller ACLs are in use and the 448-line capability is not required, configuring the **access-list hardware psa limit 128** command conserves forwarding (TLU) memory, which improves prefix scalability).
  - ◆ Turbo ACL processing should be enabled with the **access-list compiled** command for ACLs longer than 129 lines along with the **access-list hardware psa limit 128** command. This combination processes the first 128 lines in the PSA ASIC and the remaining lines with

Turbo ACLs, which optimizes performance while conserving forwarding memory.

- The 4–port OC12 ATM line card does not support input ACLs, but provides output ACL detection in microcode, which allows the process of output ACLs in the slow path.
- The 8xOC3 ATM line card supports per–vc 128 line ACLs with Cisco IOS Software Release 12.0(23)S and later. A maximum of 16 distinct input ACLs can be configured in fast path. 448–input ACL is supported on a per–VC basis in slow path only. Output ACLs are not supported.

## ISE (IP Services Engine) Engine 3 – ACL Processing

### Overview

The Engine 3 is the first Dual stage forwarding line card. The Engine 3 has forwarding/feature ASICs on the ingress and egress path. This allows ACLs to be placed in ASIC on both the ingress and egress paths. In addition, the Engine 3 ASIC structure is a hybrid pipeline/parallel array. The ASIC structure implements ACL processing in parallel high–speed ternary content addressable memory (TCAM), which provides line–rate processing of up to 20K ACEs per ingress, and 20K ACEs per egress.

These line cards are based on Engine 3:

Line Card Type	Interface Type	Connectivity
4xOC12c/STM4c	POS	IR
4xOC12c/STM4c	POS	MM
4xCHOC12/STM4–>OC3/STM1–>DS3/E3	POS	IR
16xOC3c/STM1c	POS	IR
16xOC3c/STM1c	POS	MM
8xOC3/STM1c	POS	IR
8xOC3c/STM1c	POS	MM
4xOC3c/STM1c	POS	IR
4xOC3c/STM1c	POS	MM
4xOC3c/STM1c	POS	LR
1xOC48c/STM16c	POS	SR
1xOC48c/STM16c	POS	LR
1xCHOC48/STM16–>STM4–>OC3/STM1–>DS3/E3	POS	SR
4xOC12c/STM4c	ATM/IP	IR
4xOC12c/STM4c	ATM/IP	MM
4xGE	GE	
4xOC12c/STM4c	DPT	IR
4xOC12c/STM4c	DPT	XLR

### Supported Match Criteria

All Cisco IOS Software Release 12.0S Standard and Extended match criteria are supported in the fast path except for log ACEs which are processed by the line card CPU.

## Number of ACEs Supported

- Line rate processing in both ingress and egress direction per port, per VLAN, per Frame Relay subinterface, and per ATM subinterface. Up to 20,000 extended ACEs per direction and per card are supported.
- Match criteria for TCP/UDP source/destination port range, lt, and gt are all handled in hardware using L4 operator resources.
- The number of distinct L4 operands is limited to 32 for the whole line card. Source port operators are limited to a maximum of six.

## Output ACL Processing

Native fast-path support for line-rate output ACL processing in the transmit-path Packet Processing ASIC. See the IPv4 Output ACL – Line Card Interoperation Matrix for details.

## Line Card Specific Commands

- **hw-module <slot #> tcam compile no-merge**  
*!---12.0(21)S3*
- **show-access-list hardware interface <interface name>**
- **show cef int pos[x/y] | inc if\_number**

## Operational Guidelines and Line Card Interactions

- Packets matching logging ACEs are processed in the slow path.
- Packets matching deny ACEs (throttled to ensure against system interruption) are processed in the slow path.
- When an ACL includes a range of addresses, the hardware uses special ACEs called Range ACEs which require up to three ACEs.
- ACL merging can conserve TCAM resources by sharing common ACEs across individual ACLs. To determine whether an ACL is merged, use the **show-access-list hardware interface** command.
- ACL counters are not supported for merged ACLs. With Cisco IOS Software Release 12.0(21)S3 and later, ACL merging can be disabled with the **hw-module <slot #> tcam compile no-merge** command. In order to determine whether an ACL is merged, use the **show-access-list hardware interface** command.
- If NetFlow is configured on an Engine 0/1 line card and an output ACL is configured on an egress Engine 3 or 4+ line card, the output ACL will be processed by both the ingress and egress line cards in order to allow NetFlow to account for packets denied by ACLs as well as forwarded packets.

## ACL Counter Support

	Per-ACE	Per-ACE (hardware counters)	Aggregate
21S3/ST3		X	
22S		X	X
23S	X	X	X

## Definitions:

- Per-ACE Normal Cisco IOS software support, the **show access-list <number>** command on the RP/LC displays the ACL and counter associated with each ACE. It is available only when **merge** is disabled before you configure any ACLs. This can be done by using this configuration command:

```
Router(config)#hw-module slot <number> tcam compile acl no-merge
```

This option when enabled turns off some TCAM merge optimizations and affects scalability. The exact effect depends upon individual ACLs.

Also note that the counters will not be correct if policy-based routing is applied on that interface. In that case, aggregate counter should be used.

- Per-ACE (TCAM) Hardware counters associated with each TCAM entry. No configuration is necessary and there is no impact on performance/scalability. Available only on the line card using this CLI. These counters cannot be cleared by software.

```
LC-Slot4#show contr tofab alpha acl <if-number> vmr2ace
```

A new generic CLI for this command will be available in Cisco IOS Software Release 22S:

```
LC-Slot4#show access-list hardware interface p0:1 in
```

As with the per-ACE counter, the TCAM counters are valid only when PBR is not used on that interface with ACL.

- Aggregate Each ACL shows a summary permit/deny counter. This is the sum of all individual ACE counters. No configuration is necessary and there is no impact on performance or scalability.

## Recommendations

None at this time.

## Engine 4 (POS) – ACL Processing

### Overview

Engine 4 provides this ACL support with Cisco IOS Software Release 12.0(18)S and later:

- Output ACLs are supported on E0/1/2 line cards if an Engine 4 line card is the ingress card. In this configuration, the output ACL is processed by the egress line card CPU.

These line cards are based on Engine 4:

Line Card Type	Interface Type	Engine Type	Connectivity
4xOC48c/STM16c	POS	E4	
4xOC48c/STM16c	POS	E4	LR
1xOC192c/STM64c	POS	E4	IR
1xOC192c/STM64c	POS	E4	SR
1xOC192c/STM64c	POS	E4	VSR-1
10xGE	SFP	E4	

## Engine 4+ (POS and DPT) – ACL Processing

### Overview

Engine 4+ introduces ACL functionality to the Cisco 12000 Series 10-Gigabit portfolio.

Up to 1024 ACEs are supported in each of the ingress and egress paths. Both Input and Output ACLs are processed at line rate for up to 96 ACEs. Performance for longer matches varies with match depth.

These POS line cards are based on Engine 4+:

Line Card Type	Interface Type	Connectivity
4xOC48c/STM16c	POS	SR
4xOC48c/STM16c	POS	LR
1xOC192c/STM64c	POS	IR
1xOC192c/STM64c	POS	SR
1xOC192c/STM64c	POS	VSR-1
1xOC192/STM64c	POS	LR
4xOC48c/STM16c	DPT	SEP;
1xOC192c/STM64c	DPT	IR
1xOC192c/STM64c	DPT	SR
1xOC192c/STM64c	DPT	VSR-1
1xOC192c/STM64c	DPT	LR

### Supported Match Criteria

All Cisco IOS Software Release 12.0S supported Standard and Extended ACL criteria are supported in the fast path except for log or fragment ACEs.

### Number of ACEs Supported

Up to 1024 ACEs are supported per-direction in the fast path.

**Note:** 1021 of the ACEs are configurable. Three entries are reserved for the ACEs implicit **permit ip any any**, **deny ip any any**, and **send to CPU** commands.

There is no upper limit to the number of ACEs supported. Any ACEs beyond the 1021 limit are performed in the line card slow path.

### Output ACL Processing

Output ACLs are processed in the transmit-side fast path. See the IPv4 Output ACL – Line Card Interoperation Matrix for details.

### Line Card Specific Commands

- **show tcam appl** [*acl-in* / *acl-out*] **tcam** <*label-no*>
- **show tcam appl** [*acl-in* / *acl-out*] **memory** <*port*> <*number of entries*>

### Operational Guidelines and Line Card Interactions

- Sub-interface ACLs are not supported.
- Performance varies with match depth.
- Range entries use two ACL rules (three if the two entries cross a boundary).

- One ACL is supported per physical interface.
- Up to 1024 ACEs (per direction) are supported in the fast path.
- Any of the 1024 fast path ACEs can be shared across ports.
- ACEs that use the fragment keyword are filtered in the slow path.
- Denied packets are not counted for ACEs being processed in the slow path.
- If NetFlow is configured on an Engine 0 line card and an output ACL is configured on an egress engine 3 or 4+ line card, the output ACL will be processed by both the ingress and egress line cards to allow NetFlow to account for packets denied by ACLs as well as forwarded packets.

## Recommendations

None at this time.

## Engine 4+ (Ethernet) – ACL Processing

### Overview

Engine 4+ Ethernet line cards introduce per-vlan input ACL functionality in hardware to the Cisco 12000 10-Gigabit Ethernet portfolio. These are some of the characteristics:

- Input and output ACLs can be applied simultaneously on a single port without a performance impact.
- ACLs can be applied per VLAN or per port.
- Input ACL performance up to 15K ACEs does not degrade with match depth.
- Output ACLs are processed at line rate for up to 96 ACEs. Performance for longer matches varies with match depth.

These Ethernet line cards are based on Engine 4+:

Line Card Type	Interface Type	Engine Type
10xGE Rev B ("X-B")	SFP:	E4+
Modular	SFP:	E4+
1x10GE	10G	E4+
1x10GE	10G	E4+

### Supported Match Criteria

All Cisco IOS Software Release 12.0S supported Standard and Extended ACL criteria are supported in the fast path except for log or fragment ACEs.

### Number of ACEs Supported

- Up to 15,000 input ACLs which can be configured per port or per VLAN.
- 1024 output ACEs per card which can be applied on a per port basis.

**Note:** 1021 of the ACEs are configurable. Three entries are reserved for the ACEs implicit **permit ip any any**, **deny ip any any**, and **send to CPU** commands.

## Output ACL Processing

Output ACLs are processed natively in the transmit-side fast path. See the IPv4 Output ACL – Line Card Interoperation Matrix for more information.

### Line Card Specific Commands

- **hw-module slot <number> ip acl merge**

### Operational Guidelines and Line Card Interactions

- ACEs that contain the fragment keyword are processed in the slow path.
- ACL counters are not supported for ACLs combined with other features.
- ACL counters are not supported for merged ACLs. Merged ACLs are configurable with the **hw-module slot <slot number> ip acl merge** command.
- Up to 168 L4 operations are supported per line card. Once this is exceeded, the ACL is run in the slow path.
- If an Engine 1 line card has sampled NetFlow enabled and an output ACL is enabled on an egress Engine 3 or 4+ line card, the output ACL is processed by both the ingress and egress line cards in order to allow NetFlow to account for packets denied by ACLs as well as forwarded packets.

### Recommendations

None at this time.

## ACL Logging

Before Cisco IOS Software Release 12.0(21)S, ACL logging information was sent to the RP exclusively over the Maintenance Bus (MBUS). During high levels of ACL logging activity, it was possible to exceed the capacity of the MBUS. Cisco IOS Software Release 12.0(21)S introduces several optimizations that prevent this scenario.

MBUS overload situations are reported by Cisco IOS software with these error messages:

```
LCLOG-3-INVSTATE
```

```
MBUS_SYS-3-SEQUENCE
```

With Cisco IOS Software Release 12.0(21)S and later, high severity (severity 0–4) logging messages are delivered to the RP through the MBUS while lower severity (severity 5–7) log messages are delivered to the RP through the higher-capacity switching fabric. ACL log messages are high severity, thus are now delivered to the RP through the switching fabric.

This added logging functionality is configurable using these commands:

- **logging method mbus [severity]** Determines which messages, by severity, will be sent to the RP using the MBUS. Higher severity messages will be sent through the switch fabric.
- **show logging method** Displays the current logging method for all message severity levels.
- **logging sequence-nums** This command enables the sending line card to sequence number log messages so that messages can be properly re-ordered by the RP. Without this command, log messages can be delivered to the RP in non-sequential order.

# IPv4 Output ACL – Line Card Interoperation Matrix

Before the introduction of egress ACL processing with the release of Engine 3 and Engine 4+, output ACLs were processed by the ingress line card. Output ACLs have been updated to take advantage of the high performance Engine 3 and Engine 4+ output ACL processing capabilities.

This chart provides a summary of where output ACLs are processed for different line card combinations:

Ingress line card (output ACL applied to member interface)	Egress line card					
	E0	E1	E2	E3	E4	E4+
E0	Ingress	Ingress	Ingress	Egress	n/a	Egress
E1	Ingress	Ingress	Ingress	Egress	n/a	Egress
E2	Ingress	Ingress	Ingress	Egress	n/a	Egress
E3	Egress	Egress	Egress	Egress	n/a	Egress
E4	Egress	Egress	Egress	Egress	n/a	Egress
E4+	Egress	Egress	Egress	Egress	n/a	Egress

## IPv6 ACL Support

IPv6 extended ACLs are supported in the slow path (Ingress and Egress) on E0, E1, E2, E3, and E4+ in Cisco IOS Software Release 12.0(23)S.

In Engine 3, IPv6 ACL functionality is supported in hardware in Cisco IOS Software Release 12.0(25)S. ACLs are applied to a specific interface, with an implicit deny statement at the end of each access list. IPv6 ACLs are configured using the **ipv6 access-list** command with the deny and permit keywords in the global configuration mode. Engine 3-based cards support filtering of traffic-based IPv6 option headers, flow labels, and optionally, upper-layer protocol type information.

## Cisco 12000 ACL Command Reference

### Engine 1 Commands

- **access-list hardware salsa**
- **show controller I3 | include ASIC**

### Engine 2 Commands

- **access-list hardware psa limit 128**
- **no access-list hardware psa**
- **psa bypass**
- **show access-list psa detail**
- **show access-list psa summary**
- **show controller psa feature**

### Engine 3 Commands

- **hw-module** <slot #> **tcam compile no-merge**

*!--- as of Cisco IOS Software Release 12.0(21)S3*

- **show-access-list hardware interface** <interface name>
- **show contr** [tofab|frfab] **alpha acl** <int> **vmr2ace**

## Engine 4+ Commands

- **show access-list gen7 label**
- **show tcam appl** [acl-in / acl-out] **tcam** <label-no>
- **show tcam appl** [acl-in / acl-out] **memory** <port><number of entries>

## Engine 4+ Ethernet Commands

- **hw-module slot** <number> **ip acl merge**

# Glossary

This section provides standard definitions of relevant terms:

- **Planes of Processing** A network device can be logically divided into three planes of processing:
  - ◆ Data Plane Processing on the packets flowing through the network device.
  - ◆ Control Plane Processing on the packets used to glue network devices together. This includes line protocols (such as Point-to-Point Protocol – PPP and High-Level Data Link Control – HDLC), routing protocols (Border Gateway Protocol – BGP, Routing Information Protocol version 2 – RIPv2, Open Shortest Path First – OSPF, and so on), and timing protocols (such as Network Time Protocol – NTP).
  - ◆ Management Plane Processing on packets that are used to manage the network devices. This includes telnet, Secure Shell (SSH), File Transfer Protocol (FTP), Trivial File Transfer Protocol (TFTP), SNMP, and other management protocols.
- **Standard ACLs** Standard ACLs filter exclusively at Layer 3.
- **Extended ACLs** Extended IP access lists use source and destination addresses for matching operations as well as optional protocol type information for finer granularity of control.
- **Linear Processed ACLs** Processed linearly in software. Performance varies with match depth (the number of entries that must be checked before a match is determined).
- **Turbo ACLs (Compiled)** Turbo ACLs optimize software ACL processing by compiling an ACL into a highly-optimized series of lookup tables which speed software processing. Performance of Turbo ACLs does not vary with match depth.
- **Input ACLs** An ACL applied to traffic entering the port to which it is applied.
- **Output ACLs** An ACL applied to traffic exiting the port on which it is applied. With some exceptions, output ACLs are processed by the input line card.
- **Receive Path ACLs** Receive Path ACLs provide filtering for control traffic destined for the router itself, such as routing updates and SNMP queries.
- **Dual Stage Forwarding Line Card** Line cards that have forwarding/feature ASICs on both the ingress and egress path. This allows the line card to perform features on both the ingress packet flow and egress packet flow without punting packets to the LC CPU. It also allows for new waves of dual stage forwarding algorithms to be used within the Cisco 12000. The Engine 3 Line Card is an example of a Dual Stage Forwarding Line Card.
- **Single Stage Forwarding Line Card** Line cards that have forwarding/feature ASICs on just the ingress path. These line cards only perform ASIC-based processing on the packets that flow on the ingress path. Egress traffic is either not processed (just forwarded), handled by the ingress ASICs of other LCs, or managed by the LC CPU. Engine 2, Engine 4, and Engine 4+ are examples of Single

## Related Information

- **Cisco 12000 Series Internet Routers**
  - **Technical Support & Documentation – Cisco Systems**
- 

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2008 – 2009 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

---

Updated: Feb 19, 2007

Document ID: 40742

---