

# Performance Tuning Basics

Document ID: 12809

---

## **Introduction**

### **Prerequisites**

Requirements

Components Used

Conventions

### **Background Information**

#### **Process-level and Interrupt-level Switching**

#### **Switching Paths**

Process Switching

Fast Switching

Optimum Switching

Cisco Express Forwarding (CEF)

Distributed Fast/Optimum Switching

Distributed CEF

#### **NetFlow Switching**

#### **Distributed Services**

#### **Choosing a Switching Path**

#### **Monitoring the Router**

#### **Related Information**

---

## **Introduction**

This document provides a high-level overview of the issues that affect router performance, and points you to other documents that give more details on these issues.

## **Prerequisites**

### **Requirements**

There are no specific requirements for this document.

### **Components Used**

The information in this document is based on these software and hardware versions:

- Cisco IOS® Software Release 12.1.

### **Conventions**

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

## **Background Information**

The way a router is configured can affect its packet handling performance. For routers that are handling high amounts of traffic, it is worthwhile to know what the device is doing, how it is doing it, and how long it takes to do it in order to optimize its performance. This information is represented in the configuration file. The

configuration reflects the way the packets flow through the router. A sub-optimal configuration can keep the packet inside the router longer than necessary. With a high sustained level of load, you could experience slow response, congestion, and connection timeouts.

In tuning a router's performance, your objective is to minimize the time a packet remains in a router. That is, minimize the amount of time the router forwards a packet from the incoming to the outgoing interface, and avoid buffering and congestion whenever possible. Every feature added to a configuration is one more step an incoming packet must pass through on its way to the destination port.

The two major resources you need to save are the router's CPU time and memory. The router should always have CPU availability to handle spikes and periodic tasks. Whenever the CPU is utilized at 99% for too long, network stability can be seriously impacted. The same concept applies to memory availability: memory must always be available. If the router's memory is almost fully used, no room is left in the system buffer pools. This means that packets that require processor attention (process-switched packets) are dropped as soon as they come in. It's easy to imagine what could happen if the dropped packets contain interface keepalives or important routing updates.

## Process-level and Interrupt-level Switching

In IP networks, packet forwarding decisions in routers are based on the contents of the routing table. When searching the routing table, the router looks for the longest match for the destination IP address prefix. This is done at "process level" (known as process switching), which means that the lookup is considered as just another process queued among other CPU processes. As a result, that lookup time is unpredictable and may take very long. To address this, a number of switching methods based on exact-match-lookup have been introduced in Cisco IOS software.

The main benefit of exact-match-lookup is that the lookup time is deterministic and very short. The time it takes for the router to make a forwarding decision is significantly decreased, making it possible to do this at "interrupt level". Interrupt-level switching means that when a packet arrives, an interrupt is triggered which causes the CPU to postpone other tasks in order to handle that packet. The legacy method for forwarding packets (by looking for a longest match in the routing table) cannot be implemented at interrupt level and must be performed at process level. For a number of reasons, some of which are mentioned below, the longest-match-lookup method cannot be completely abandoned, so these two lookup methods exist in parallel on Cisco routers. This strategy has been generalized and is also applied to IPX and AppleTalk.

In order to perform an exact-match-lookup at interrupt level, the routing table has to be transformed to use a memory structure convenient for this kind of lookup. Different switching paths use different memory structures. The architecture of this so-called structure has a significant impact on the lookup time, making the selection of the most appropriate switching path a very important task. For a router to make a decision on where to forward a packet, the basic information it needs are the next-hop address and the outgoing interface. It also needs information on the encapsulation of the outgoing interface. Depending on its scalability, the latter may be stored in the same or in a separate memory structure.

The following is the procedure for executing interrupt-level switching:

1. Look up the memory structure to determine the next-hop address and outgoing interface.
2. Do an Open Systems Interconnection (OSI) Layer 2 rewrite, also called MAC rewrite, which means changing the encapsulation of the packet to comply with the outgoing interface.
3. Put the packet into the tx ring or output queue of the outgoing interface.
4. Update the appropriate memory structures (reset timers in caches, update counters, and so forth).

The interrupt which is raised when a packet is received from the network interface is called the "RX interrupt". This interrupt is dismissed only when all the above steps are executed. If any of the first three steps

above cannot be performed, the packet is sent to the next switching layer. If the next switching layer is process switching, the packet is put into the input queue of the incoming interface for process switching and the interrupt is dismissed. Since interrupts cannot be interrupted by interrupts of the same level and all interfaces raise interrupts of the same level, no other packet can be handled until the current RX interrupt is dismissed.

Different interrupt switching paths can be organized in a hierarchy, from the one providing the fastest lookup to the one providing the slowest lookup. The last resort used for handling packets is always process switching. Not all interfaces and packet types are supported in every interrupt switching path. Generally, only those that require examination and changes limited to the packet header can be interrupt-switched. If the packet payload needs to be examined before forwarding, interrupt switching is not possible. More specific constraints may exist for some interrupt switching paths. Also, if the Layer 2 connection over the outgoing interface must be reliable (that is, it includes support for retransmission), the packet cannot be handled at interrupt level.

The following are examples of packets that cannot be interrupt-switched:

- Traffic directed to the router (routing protocol traffic, Simple Network Management Protocol (SNMP), Telnet, Trivial File Transfer Protocol (TFTP), ping, and so on). Management traffic can be sourced and directed to the router. They have specific task-related processes.
- OSI Layer 2 connection-oriented encapsulations (for example, X.25). Some tasks are too complex to be coded in the interrupt-switching path because there are too many instructions to run, or timers and windows are required. Some examples are features such as encryption, Local Area Transport (LAT) translation, and Data-Link Switching Plus (DLSW+).

## Switching Paths

The path that a packet follows while inside a router is determined by the active forwarding algorithm. These are also referred to as "switching algorithms" or "switching paths." High-end platforms have typically more powerful forwarding algorithms available than low-end platforms, but often they are not active by default. Some forwarding algorithms are implemented in hardware, some are implemented in software, and some are implemented in both, but the objective is always to send packets out as fast as possible.

The switching algorithms available on Cisco routers are:

Forwarding Algorithm	Command (Issue From config-interface Mode)
Fast switching	<code>ip route-cache</code>
Same-interface switching	<code>ip route-cache same-interface</code>
Autonomous switching (7000 platforms only)	<code>ip route-cache cbus</code>
Silicon switching (7000 platforms with an SSP installed only)	<code>ip route-cache sse</code>
Distributed switching (VIP-capable platforms only)	<code>ip route-cache distributed</code>
Optimum switching (high-end routers only)	<code>ip route-cache optimum</code>
NetFlow switching	<code>ip route-cache flow</code>
Cisco Express Forwarding (CEF)	<code>ip cef</code>

Here is a brief description of each switching paths sorted in order of performance. Autonomous and silicon switching are not discussed since they relate to end of engineering hardware.

## Process Switching

Process switching is the most basic way of handling a packet. The packet is placed in the queue corresponding to the Layer 3 protocol and then the corresponding process is scheduled by the scheduler. The process is one of the processes you can see in the **show processes cpu** command output (that is, "ip input" for an IP packet). At this point, the packet stays in the queue until the scheduler gives the CPU to the corresponding process. The waiting time depends on the number of processes waiting to run and the number of packets waiting to be processed. The routing decision is then made based on the routing table. The encapsulation of the packet is changed to comply with the outgoing interface and the packet is enqueued to the output queue of the appropriate outgoing interface.

## Fast Switching

In fast switching, the CPU makes the forwarding decision at interrupt level. Information derived from the routing table and information about the encapsulation of outgoing interfaces are combined to create a fast-switching cache. Each entry in the cache is comprised of the destination IP address, the outgoing interface identification, and the MAC rewrite information. The fast-switching cache has the structure of a binary tree.

If there is no entry in the fast-switching cache for a certain destination, the current packet must be enqueued for process switching. When the appropriate process makes a forwarding decision for this packet, it creates an entry in the fast-switching cache and all consecutive packets to the same destination can be forwarded at interrupt level.

Since this is a destination-based cache, load sharing is only done per destination. Even if the routing table has two equal cost paths for a destination network, there is only one entry in the fast-switching cache for each host.

## Optimum Switching

Optimum switching is basically the same as fast switching, except that it uses a 256-way multidimensional tree (mtree) instead of a binary tree, resulting in bigger memory needs and faster cache lookup. More details on the tree structures and fast/optimum/Cisco Express Forwarding (CEF) switching can be found in [How to Choose the Best Router Switching Path for Your Network](#).

## Cisco Express Forwarding (CEF)

The main drawbacks of the previous switching algorithms are:

1. The first packet for a particular destination is always process-switched to initialize the fast cache.
2. The fast-cache can become very big. For example, if there are multiple equal cost paths to the same destination network, the fast cache is populated by host entries instead of the network as discussed above.
3. There's no direct relation between the fast-cache and the ARP table. If an entry becomes invalid in the ARP cache, there is no way to invalidate it in the fast cache. To avoid this problem, 1/20th of the cache is randomly invalidated every minute. This invalidation/repopulation of the cache can become CPU-intensive with very large networks.

CEF addresses these issues by using two tables: the FIB (Forwarding Information Based) table and the adjacency table. The adjacency table is indexed by the Layer 3 (L3) addresses and contains the corresponding Layer 2 (L2) data needed to forward a packet. It is populated when the router discovers adjacent nodes. The FIB table is an mtree indexed by L3 addresses. It is built based on the routing table and points to the adjacency table.

Another advantage of CEF is that the database structure allows load balancing per destination or per packet. The CEF home page provides more information about CEF.

## Distributed Fast/Optimum Switching

Distributed Fast/Optimum Switching seeks to offload the main CPU (Route/Switch Processor [RSP]) by moving the routing decision to the interface processors (IPs). This is possible only on the high end platforms which can have dedicated CPUs per interface (Versatile Interface Processors [VIPs], Line Cards [LCs]). In this case, the fast cache is simply uploaded to the VIP. When a packet is received, the VIP tries to make the routing decision based on that table. If it succeeds, it directly enqueues the packet to the outgoing interface's queue. If it fails, it enqueues the packet for the next configured switching path (optimum switching -> fast switching -> process-switching).

With distributed switching, access lists are copied to the VIPs, which means the VIP can check the packet against the access list without RSP intervention.

## Distributed CEF

Distributed CEF (dCEF) is similar to distributed switching but there are fewer sync issues between the tables. dCEF is the only distributed switching method available from Cisco IOS Software Release 12.0. It's important to know that if distributed switching is enabled on a router, the FIB/adjacency tables are uploaded on all VIPs in the router, regardless of whether their interface has CEF/dCEF configured.

With dCEF, the VIP also processes the access lists, policy based routing data and rate limiting rules, which are all held in the VIP card. Netflow can be enabled together with dCEF to enhance the access list processing by the VIPs.

The table below shows, for each platform, which switching path is supported from which Cisco IOS software version.

Switching Path	Below Low End(1)	Low/Middle End(2)	Cisco AS5850	Cisco 7000 w/RSP	Cisco 72xx/71xx	Cisco 75xx	Cisco GSR 12xxx	Comments
<b>Process Switching</b>	ALL	ALL	ALL	ALL	ALL	ALL	NO	Initializes the switching cache Default for all except IP in high end
<b>Fast</b>	NO	ALL	ALL	ALL	ALL	ALL	NO	
<b>Optimum Switching</b>	NO	NO	NO	ALL	ALL	ALL	NO	Default for high end for IP before 12.0
<b>Netflow Switching (3)</b>	NO	12.0(2), 12.0T & 12.0S	ALL	11.1CA, 11.1CC, 11.2, 11.2P,	11.1CA, 11.1CC, 11.2, 11.2P,	11.1CA, 11.1CC, 11.2, 11.2P,	12.0(6)S	

				11.3, 11.3T, 12.0, 12.0T, 12.0S	11.3, 11.3T, 12.0, 12.0T, 12.0S	11.3, 11.3T, 12.0, 12.0T, 12.0S		
<b>Distributed Optimum Switching</b>	NO	NO	NO	NO	NO	11.1, 11.1CC, 11.1CA, 11.2, 11.2P, 11.3 & 11.3T	NO	Using VIP2–20,40,50 Not available from 12.0.
<b>CEF</b>	NO	12.0(5)T	ALL	11.1CC, 12.0 & 12.0x	11.1CC, 12.0 & 12.0x	11.1CC, 12.0 & 12.0x	NO	Default for high end for IP from 12.0
<b>dCEF</b>	NO	NO	ALL	No	NO	11.1CC, 12.0 & 12.0x	11.1CC, 12.0 & 12.0x	Only on 75xx+VIPs and on GSRs

(1) Includes 801 through 805.

(2) Includes 806 and above, 1000, 1400, 1600, 1700, 2600, 3600, 3700, 4000, AS5300, AS5350, AS5400, and AS5800 series.

(3) Support for NetFlow Export v1, v5, and v8 on 1400, 1600, and 2500 platforms is targeted for Cisco IOS software release 12.0(4)T. NetFlow support for these platforms is not available in the Cisco IOS Software 12.0 mainline release.

(4) The performance impact of the use of UHP on these platforms: RSP720–3C/MSFC4, RSP720–3CXL/MSFC4, 7600–ES20–GE3CXL/7600–ES20–D3CXL, SUP720–3BXL/MSFC3 is Explicit null that causes a recirculation and decrease the performance in PE. The throughput is reduced to 12 Mpps from 20 Mpps on RSP720–3C/MSFC4, RSP720–3CXL/MSFC4, and SUP720–3BXL/MSFC3, and the 7600–ES20–GE3CXL/7600–ES20–D3CXL has a reduced throughput to 25 Mpps from 48 Mpps.

## NetFlow Switching

NetFlow switching is a misnomer, aggravated by the fact that it is configured in the same way as a switching path. Actually, NetFlow switching is not a switching path because the NetFlow cache does not contain or point to information needed for Layer 2 rewrite. The switching decision has to be made by the active switching path.

With NetFlow switching, the router classifies the traffic per flow. A flow is defined as a unidirectional sequence of packets between given source and destination endpoints. The router uses the source and destination addresses, the transport layer port numbers, the IP protocol type, the Type of Service (ToS), and the source interface to define a flow. This way of classifying the traffic allows the router to process only the first packet of a flow against CPU–demanding features such as large access lists, queuing, accounting policies, and powerful accounting/billing. The NetFlow home page provides more information.

## Distributed Services

On high-end platforms several CPU intensive tasks (not just the packet switching algorithms) can be moved from the main processor to distributed processors like the ones on the VIP cards (7500). Some of these tasks can be exported from a general-purpose processor to specific port-adapters or network modules that implement the feature on dedicated hardware.

It is common to offload tasks from the main processor to the VIP's processors whenever possible. This frees resources and increases the router performance. Some processes that might be offloaded are packet compression, packet encryption, and weighted fair queuing. See the following table for more tasks that can be offloaded. A complete description of the services available can be found in Distributed Services on the Cisco 7500.

Service	Features
Basic Switching	Cisco Express Forwarding IP fragmentation Fast
VPN	EtherChannel ACLs--- extended and turbo Cisco encryption Generic route encapsulation (GRE) tunnels IP Security (IPSec) Layer 2 Tunneling Protocol tunnels (L2TP)
QoS	NBAR Traffic shaping (dTS) Policing (CAR) Congestion avoidance (dWRED) Guaranteed minimum bandwidth (dCBWFQ) Policy propagation via BGP Policy routing
Multiservice	Low latency queuing FRF 11/12 RTP header compression Multilink PPP with link fragmentation and interleaving
Accounting	Output accounting NetFlow export Precedence and MAC accounting
Load Balancing	CEF load balancing Multilink PPP
Caching	WCCP V1 WCCP V2
Compression	L2 SW and HW compression L3 SW and HW compression
Multicast	Multicast distributed switching

## Choosing a Switching Path

The basic rule is choose the best switching path available (from fastest to slowest): dCEF, CEF, optimum, and fast. Enabling CEF or dCEF gives the best performances. Enabling NetFlow switching can enhance or decrease performance depending on your configuration. If you have very big access lists, or if you need to do some accounting, or both, NetFlow switching is recommended. Usually NetFlow is enabled on edge routers having a lot of CPU power and using many features. If you configure multiple switching paths such as fast-switching and CEF on the same interface, the router will try all of them from best to worst (starting from CEF and ending with process-switching).

## Monitoring the Router

Use the following commands to see if the switching path is used effectively and how loaded the router is.

**show ip interfaces:** This command gives an overview of the switching path applied to a particular interface.

```
Router#show ip interfaces
Ethernet0/0 is up, line protocol is up
Internet address is 10.200.40.23/22
Broadcast address is 255.255.255.255
Address determined by setup command
MTU is 1500 bytes
Helper address is not set
Directed broadcast forwarding is disabled
Outgoing access list is not set
Inbound access list is not set
Proxy ARP is enabled
Security level is default
Split horizon is enabled
ICMP redirects are always sent
ICMP unreachable are always sent
ICMP mask replies are never sent
IP fast switching is enabled
IP fast switching on the same interface is disabled
IP Flow switching is disabled
IP CEF switching is enabled
IP Fast switching turbo vector
IP Normal CEF switching turbo vector
IP multicast fast switching is enabled
IP multicast distributed fast switching is disabled
IP route-cache flags are Fast, CEF
Router Discovery is disabled
IP output packet accounting is disabled
IP access violation accounting is disabled
TCP/IP header compression is disabled
RTP/IP header compression is disabled
Probe proxy name replies are disabled
Policy routing is disabled
Network address translation is disabled
WCCP Redirect outbound is disabled
WCCP Redirect inbound is disabled
WCCP Redirect exclude is disabled
BGP Policy Mapping is disabled
```

From this output we can see that fast switching is enabled, NetFlow switching is disabled, and CEF switching is enabled.

**show processes cpu :** This command displays useful information on the CPU load. For more information, see Troubleshooting High CPU Utilization on Cisco Routers.

```
Router#show processes cpu
CPU utilization for five seconds: 0%/0%; one minute: 0%; five minutes: 0%
PID  Runtime(ms)  Invoked  uSecs   5Sec   1Min   5Min  TTY  Process
  1         28      396653     0    0.00%  0.00%  0.00%  0  Load Meter
  2        661      33040     20    0.00%  0.00%  0.00%  0  CEF Scanner
  3       63574     707194     89    0.00%  0.00%  0.00%  0  Exec
  4    1343928    234720    5725    0.32%  0.08%  0.06%  0  Check heaps
  5         0         1         0    0.00%  0.00%  0.00%  0  Chunk Manager
  6         20         5    4000    0.00%  0.00%  0.00%  0  Pool Manager
  7         0         2         0    0.00%  0.00%  0.00%  0  Timers
  8     100729     69524    1448    0.00%  0.00%  0.00%  0  Serial Backgroun
  9        236     66080         3    0.00%  0.00%  0.00%  0  Environmental mo
 10     94597    245505     385    0.00%  0.00%  0.00%  0  ARP Input
 11         0         2         0    0.00%  0.00%  0.00%  0  DDR Timers
 12         0         2         0    0.00%  0.00%  0.00%  0  Dialer event
 13         8         2    4000    0.00%  0.00%  0.00%  0  Entity MIB API
 14         0         1         0    0.00%  0.00%  0.00%  0  SERIAL A'detect
 15         0         1         0    0.00%  0.00%  0.00%  0  Critical Bkgnd
```

```

16      130108    473809    274    0.00%  0.00%  0.00%   0 Net Background
17         8       327     24    0.00%  0.00%  0.00%   0 Logger
18       573    1980044     0    0.00%  0.00%  0.00%   0 TTY Background
[...]
```

**show memory summary** : The first lines of this command give useful information on the memory usage of the router and on the memory/buffer.

```

Router#show memory summary
      Head    Total(b)    Used(b)    Free(b)    Lowest(b)    Largest(b)
Processor  8165B63C    6965700    4060804    2904896    2811188    2884112
      I/O    1D00000    3145728    1770488    1375240    1333264    1375196
[...]
```

**show interfaces stat** and **show interfaces switching**: These two commands show which path the router uses and how the traffic is switched.

```

Router#show interfaces stat
Ethernet0
      Switching path    Pkts In    Chars In    Pkts Out    Chars Out
      Processor          52077    12245489    24646    3170041
      Route cache        0         0           0         0
      Distributed cache  0         0           0         0
      Total              52077    12245489    24646    3170041
```

```

Router#show interfaces switching
Ethernet0
      Throttle count    0
      Drops             RP          0           SP          0
      SPD Flushes       Fast        0           SSE         0
      SPD Aggress       Fast        0
      SPD Priority       Inputs      0           Drops       0

      Protocol          Path        Pkts In    Chars In    Pkts Out    Chars Out
      Other             Process     0           0           595         35700
      Cache misses
      Fast              0           0           0           0
      Auton/SSE         0           0           0           0
      IP                Process     4           456         4           456
      Cache misses
      Fast              0           0           0           0
      Auton/SSE         0           0           0           0
      IPX               Process     0           0           2           120
      Cache misses
      Fast              0           0           0           0
      Auton/SSE         0           0           0           0
      Trans. Bridge     Process     0           0           0           0
      Cache misses
      Fast              11          660         0           0
      Auton/SSE         0           0           0           0
      DEC MOP           Process     0           0           10          770
      Cache misses
      Fast              0           0           0           0
      Auton/SSE         0           0           0           0
      ARP              Process     1           60          2           120
      Cache misses
      Fast              0           0           0           0
      Auton/SSE         0           0           0           0
      CDP               Process     200         63700       100         31183
      Cache misses
      Fast              0           0           0           0
      Auton/SSE         0           0           0           0
```

## Related Information

- [Troubleshooting High CPU Utilization on Cisco Routers](#)
  - [The show processes Command](#)
  - [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: Jul 11, 2007

Document ID: 12809

---