

Cisco 3600 Series Router Architecture

Document ID: 7442

Contents

Introduction

Prerequisites

- Requirements

- Components Used

- Background Information

- Conventions

Hardware Overview

- Block Diagram

- Memory Details

Boot Sequence

Packet Switching

- Buffers Overview

- Switching Paths

Performance Figures

Related Information

Introduction

This document describes the hardware and software architecture of the 3600 Series Routers.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

The information in this document is based on this hardware version:

- Cisco 3600 Series 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.

Background Information

The 3600 Series Routers include these models:

- 3620 equipped with two network module slots
- 3640 equipped with four network module slots
- 3660 equipped with six network module slots

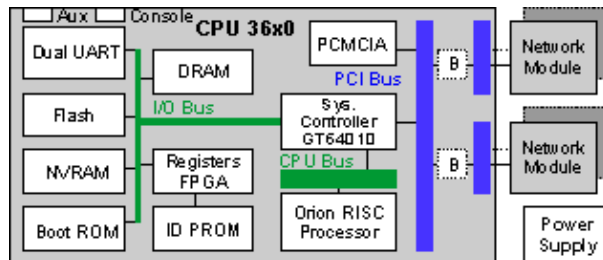
For a product overview, see Cisco 3600 Series – Modular, High-Density Access Routers.

Conventions

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

Hardware Overview

Block Diagram



Main Processor – CPU

- The processor loads instructions defined in Cisco IOS software from the main processor memory and executes them.
- The 3620 uses an **80 MHz** IDT RISC Processor (R4700)
 - ◆ Multiplexed 64-bit address and data bus
 - ◆ Internal 16 KB data cache, 16 KB instruction cache, both 2-way set associative, write back
- The 3640 uses the same processor as the 3620, but with a **100 MHz** internal clock.
- The 3660 uses a **225 MHz** IDT RISC Processor (R5271) with 2 MB of Level 2 cache.

System Controller

The system controller helps the main processor with device control, interrupt handling, counting/timing, data transfer, and communication with slower Input/Output (I/O) devices, and dynamic RAM (DRAM). The system controller contains the cache, Flash, DRAM, PCI, and Interrupt controller.

PCI Bus

The PCI Bus is the medium of communication between the CPU and the network modules (NMs), mainly for packet data transfer.

- The PCI Bus is 32-bit, 20 Mhz for the 3620, and 25 MHz for the 3640 and 3660.
- The **PCI Bridges** (two on the 3640, three on the 3660) provide some isolation and allow inter PCI Bus communication. There is no bridge on the 3620; the NMs are connected directly.
- The **Arbiter** (not shown) controls the communication over the PCI Bus, and resolves master-slave conflicts.

Other Buses

Other buses are used by the CPU to access various components of the system and to transfer instructions and data to and from specified memory addresses.

- The CPU Bus is for high speed operations, with direct Processor access – 64-bit multiplexed, 40 MHz in the 3620, 50 MHz in the 3640, and 75 MHz in the 3660.
- The I/O Bus allows the System Controller to control other devices. It allows a 32-bit DRAM interface in the 3620; 32- or 64-bit DRAM interface in the 3640; and 64-bit SDRAM interface in the

3660.

Dual Universal Asynchronous Receiver/Transmitter (UART)

Dual UARTs provide the necessary user interface.

- It has two RS232 ports: data communications equipment (DCE) for the console port and data terminal equipment (DTE) for the Auxiliary port; both have RJ45 connectors.
- Higher console speeds (up to 115.2 Kbps) are supported.
- Downloading of Cisco IOS software images over the console or auxiliary port is supported through xmodem or ymodem. More information is available at ROMmon Recovery for the Cisco 3600 Series Router.

Network Modules (NMs)

Interfaces on the 3600 Series Routers are configured by means of interchangeable network modules (NMs), some of which can accommodate voice and/or WAN interface cards (voice interface card – VIC, WAN Interface Card – WIC, Voice/WAN Interface Card – VWIC).

In addition to the network modules, the 3660 supports two internal Advanced Integration Module (AIM) slots. The 3660 also supports hot swapping (not supported on 3620/3640 routers) of similar network modules.

Power Supplies

Power supplies provide power to various components of the router.

- The power supply can be AC or DC
- Fans provide the necessary cooling to the router.

Individual specifications for the power supply are provided in the table below:

Model	Input (W/BTU/KVA)	Output (W)	Output Voltage/Current
3660	385/1320/0.405(AC, DC)	250	5V/40A, 3.3V/15A, 12V/13A, -12V/1.5A
3640	230/800/0.242(AC only)	140	5V/24A, 12V/5A, -12V/3A
	205/705/0.205(DC only)		
3620	95/320/0.10(AC, DC)	60	5V/10A, 12V/1.5A, -12V/0.35A

The 3600 Series can also operate from Redundant Power Supply (RPS). The PWR600-AC-RPS is a Cisco RPS for the 2500, 2600, 3600, and 4000 Series Routers.

Memory Details

The Cisco 3600 Series supports run-from-ram images only. The Cisco 3600 Series Routers support the following types of memory:

- Main processor memory – Used to store the running configuration and routing tables. Cisco IOS software is executed from the main memory.
- Shared (packet) memory – Used for packet buffering by the router's network interfaces.
- Flash memory – Used to store the Cisco IOS software image. Can be stored on either the Flash single in-line memory module (SIMM), or the PCMCIA card.
- Nonvolatile random access memory (NVRAM) – Used to store the system configuration file and the virtual configuration register.
- EPROM-based memory – Stores the ROM monitor that allows booting an operating system software image from Flash or PCMCIA memory when Flash memory does not contain a valid boot helper image.

Important: The DRAM/SDRAM is logically divided into main processor memory and shared I/O memory.

- The Cisco 3620 and 3640 use two types of replaceable or upgradeable memory: DRAM memory and Flash memory. Both types of memory are implemented with SIMMs. Each router has two Flash SIMM sockets and four DRAM SIMM sockets. In the standard Flash memory configuration, the first socket contains an 8 MB Flash SIMM, and the second socket is empty. You can upgrade the Flash memory to 32 MB on both routers. Each Cisco 3620/3640 router ships standard with 32 MB of DRAM. The Cisco 3640 is expandable to 128 MB of DRAM. The Cisco 3620 is expandable to 64 MB of DRAM.
- The Cisco 3660 Series uses two types of replaceable or upgradeable memory: SDRAM memory and Flash memory. Flash memory is implemented with SIMMs whereas SDRAM memory uses dual inline memory modules (DIMMs). The 3660 has two Flash SIMM sockets and two SDRAM DIMM sockets. In the standard Flash memory configuration, the first socket contains an 8 MB Flash SIMM, and the second socket is empty. You can upgrade the Flash memory to 64 MB on a 3660 router. Each Cisco 3660 router ships standard with 32 MB of SDRAM and is upgradeable to 256 MB of SDRAM.
- The Cisco 3600 Series Routers use a DRAM memory configuration that is different from other Cisco routers. Unlike the Cisco 4000 Series Routers, for example, which have separate physical locations for processor memory and packet memory, the Cisco 3600 Series Routers use one pool of DRAM memory. The DRAM is partitioned into processor memory and packet memory areas. For example, the memory of a 16 MB DRAM configuration is split into 12 MB (75% by default) for processor memory and 4 MB (25%) for the packet memory. The processor memory is used for the Cisco IOS software and its subsystems, to store the routing/Cisco Express Forwarding tables, Fast switching cache, the running configuration, and so on. The shared I/O memory known as the packet memory is used to temporarily store the packets in system buffers during process switching and interface buffers during fast and CEF switching.

The logical division (75%/25%) can be changed with the **memory-size** global configuration command. The router will adjust its memory parameters if it does not have enough processor memory to load the Cisco IOS software. The new **memory-size** (I/O memory) command percentage allows DRAM split increases of 30, 40 and 50 percent depending on the type and number of network modules configured.

The Cisco IOS supports an auto-adjusting feature upon startup in which, if the I/O memory-percentage has been configured too high, leaving insufficient processor memory available for bringing up the IOS image, then the IOS automatically reduces the percentage to a lower value. If the default is not sufficient for an image to boot up, then the user does not have enough DRAM/SDRAM for that particular subset image, and the router will complain in the same manner as existing Cisco platforms that encounter insufficient processor memory.

Main and packet DRAM/SDRAM cannot share or borrow from the other as on the 25xx Series Routers. If the main DRAM/SDRAM memory is insufficient, then either a memory increase is required, or, if the memory size has been adjusted through the Cisco IOS software, then a memory decrease is required.

At least 2 MB of free processor memory and 1.2 MB of free I/O memory are required for most average size networks. You can use the **show memory free** command to view the amount of used and available system memory.


The Cisco 3600 allows you to load new system images using a PCMCIA Flash memory card. You can also load images from a local or remote PC through the console or auxiliary ports using the xmodem or ymodem protocols. See ROMmon Recovery for the Cisco 3600 Series Router for details.

Available DRAM/SDRAM Memory Options


The Cisco 3620, Cisco 3640, and Cisco 3660 have different memory options and product numbers due to the difference in memory architectures.

- The Cisco 3640 can support industry standard dual-bank 8 and 32 MB SIMMs, because of its 64-bit memory architecture.
- The Cisco 3620, with a 32-bit architecture, cannot use dual-bank SIMMs, but instead uses a custom Cisco 8 MB single-bank SIMM. It cannot be mixed with the dual-bank 8 MB SIMMs.
- Both the Cisco 3620 and Cisco 3640 use the same 4 MB and 16 MB SIMMs, but have different product numbers.
- By default, the 3640 and 3620 ship with non-parity DRAM SIMMs, but parity SIMMs are supported if all banks contain SIMMs that have parity enabled. If parity SIMMs are mixed in a system with non-parity SIMMs, the parity function is not supported on any of the SIMMs. Only certain combinations of DRAM SIMMs are permitted.


DRAM SIMM Locations in the Cisco 3620 Router



DRAM SIMM Locations in the Cisco 3640 Router



SDRAM DIMM Locations in the Cisco 3660 Router



Below are generic memory combination guidelines:

- On the 3620 and 3640 routers, DRAM is composed of 4 SIMMs (72-pin, 60 ns, all with or without parity) configured as four banks.
 - ◆ The sum of memory in slots 0 and 1 must be a power of 2 in order to stuff slots 2 and 3.
 - ◆ Largest memory sizes must always be first.
 - ◆ No holes are allowed. An unoccupied slot cannot be followed by an occupied one.
 - ◆ The ROM Monitor displays a warning message if illegal SIMM placement or sizing has occurred.
- The 3620 supports 32 bit access, 16 to 64 MB (single bank SIMMs).
 - ◆ In order to use banks 2 and 3, the SIMMs in banks 0 and 1 must be identical.
 - ◆ The SIMMs in banks 0 and 1 must be the same size or greater than the SIMMs located in banks 2 and 3.
 - ◆ The Cisco 3620 does not support double-sided SIMMs or 64-bit addressing.
- The 3640 supports 32 bit (single bank SIMMs) or 64 bit (single or dual bank SIMMs) access, 16 to 128 MB. The following rules are for 64 bit access mode:
 - ◆ The SIMMs in slots 0 and 1 must be the same size in MB, and have the same access time in Ns

- ◆ The SIMMs in slots 2 and 3 must be the same size and also have the same access time in Ns
- ◆ The size of the SIMMs in slots 2 and 3 must be less than or equal to the size of the SIMMs in slots 0 and 1
- On the 3660, there are two SDRAM DIMMs. The following SDRAM rule applies:
 - ◆ The SIMM in socket 0 must be the same size or larger in megabytes, than the SIMM in socket 1.

Note: The valid memory configurations can also be found on the router when in the ROMmon mode. The command **meminfo-1** displays valid memory configurations.

For guidelines on how to order the proper amount of memory to support different network module configurations, refer to Product Bulletin 544 – 3600 Series Memory Options and Configuration Guide.

Network Module Memory Requirements

The amount of packet memory required depends on the number and types of Network Modules configured, and on the amount of memory Cisco IOS software uses for each specific feature. For example, additional processor and packet memory is used when a CT1 or CE1 is configured for Integrated Services Digital Network (ISDN). When the CT1 or CE1 network modules are used in channelized mode, much less memory is required.

For two or more T1/E1 primary rate interfaces (PRIs), or 12 or more basic rate interfaces (BRIs) configured, Cisco recommends a 60/40 memory split. Currently, the Cisco 3600 Series is shipped with a default memory split of 75/25, so you need to reconfigure this with the Cisco IOS software **memory-size I/O memory 40** command. Otherwise, you may receive a %SYS-2 MALLOCFAIL error message. This error message occurs when less than 1.2 MB of free I/O memory is available for system use. Increasing the I/O memory size eliminates this error message.

See Table 1 in the Cisco 3600 Series Memory Options and Configuration Guide for more details.

Factors that Affect Default Memory Configurations

The size of the network routing tables and access lists has the greatest impact on the amount of DRAM required. Only the main processor memory is affected by the size of the routing tables, not the I/O memory. Increasing the I/O memory only affects the Network Module packet buffers. In an ideal situation, 16 MB of DRAM with the default memory split can support up to 10,000 IP and 9,000 IPX routes in the routing tables.

The use of multiple high density Network Modules, such as T1/E1 or MBRI, requires that additional packets be configured using the Cisco IOS software **memory-size I/O memory** command. A 60/40 split is recommended in most situations where more than two T1/E1 PRIs or three NP-8Bs are configured, but 50/50 also works if the network routing tables allow this. The router error message %SYS-2-MALLOCFAIL is an indication that insufficient I/O memory is available for the network modules configured. This error occurs when less than 1.2 MB of free I/O memory is available to the system.

See Table 1 in the Cisco 3600 Series Memory Options and Configuration Guide for the specific amounts of memory required.

Flash SIMM Memory

The Cisco 3600 Series Routers contain two Flash memory SIMM sockets for storing the Cisco IOS software image. These SIMMs are not interchangeable with the DRAM SIMMs in the Cisco 3640/3620 and the SDRAM DIMMs in the Cisco 3660. The Flash memory can be upgraded by replacing the existing SIMM with an 8 or 16 MB SIMM, or by adding an additional SIMM to the second Flash memory socket. Unequal value

SIMMs are supported by Flash. You must always have at least one Flash SIMM installed in the router for normal operations. Partitioning of the Flash memory (Dual Bank Flash [DBF]) is supported if the router is configured for at least 8MB of memory.

Note that 36x0 platforms have no Boot Flash, and no RxBoot. So, NetBooting is not possible. Therefore, having a good Cisco IOS software image on Flash SIMM or the PCMCIA Flash card is important.

PCMCIA Flash Card

The Cisco IOS software image can also be stored on a PCMCIA Flash card on either of the 2 available PCMCIA slots provided on the Cisco 3660, Cisco 3640 and Cisco 3620. Partitioning is also supported on the PCMCIA cards that support 8 MB or more capacity.

Two PCMCIA Flash cards are currently available: 8 MB and 16 MB.

The format of PCMCIA cards is different on different platforms. PCMCIA cards are not exchangeable without (re)formatting them. The 36x0 Series uses the Filesystem Class B and is compatible with the 100x and the 160x Series in this regard, but you have to reformat them in a 36x0 router to be able to use them.

More information is available at the [PCMCIA Filesystem Compatibility Matrix](#).

Non-Volatile Random-Access Memory (NVRAM)

NVRAM is used for writeable permanent storage of the startup configuration.

- NVRAM is a battery-backed static RAM (SRAM).
- The life of NVRAM is specified by a maximum number of writes and a maximum time limit.
- NVRAM size is 32 KB for the 3620, and 128 KB for the 3640 and 3660 routers.

Boot ROM

Boot ROM is an erasable programmable ROM (EPROM) used for permanently storing startup diagnostic code (ROM Monitor). The boot ROM size is 512 KB.

EEPROM

EEPROM is an electrically-erasable programmable ROM (EEPROM) (128 B) used for permanent storage of hardware revision and identification information, and Media Access Control (MAC) addresses for LAN interfaces.

Registers

Registers are small, fast memory units used for storage of special purpose information, such as "interrupt status" and "currently executing instruction".

- The location of registers depends upon their use. For example, the Main Processor contains the instruction register and other control registers; Dual Universal Asynchronous Receiver Transmitter (DUART) contains its own status register such as other I/O devices, data read/write registers on various components, and so on.
- The CPU contains the field programmable gate array (FPGA) register for various special purpose registers.
- The Main Processor contains general purpose registers for integer and floating point data used in instruction execution.

Boot Sequence

When a 36x0 router is powered on or rebooted, the following events occur:

- The ROM Monitor (in Boot ROM) initializes itself.
- The ROM Monitor checks the boot field (the lowest four bits) in the configuration register.
 - ◆ If the last digit of the boot field is 0 (for example, 0x100), the system does not boot a Cisco IOS software image and waits for user intervention at the ROM Monitor prompt. From the ROM Monitor mode, you can manually boot the system using the **boot** or **b** command.
 - ◆ If the last digit of the boot field is 2 through F (for example, 0x102 through 0x10F), the router boots the first valid image specified in the configuration file or specified by the BOOT environment variable. It goes through each **boot** system command in sequential order until it boots a valid image.

If the router cannot find a valid image, the following events happen:

- If all **boot** commands in the system configuration file fail, the system attempts to boot the first valid file in Flash memory.
- If a fully functional system image is not found, the router does not function and stays in ROM Monitor waiting to be reconfigured through a direct console port connection.

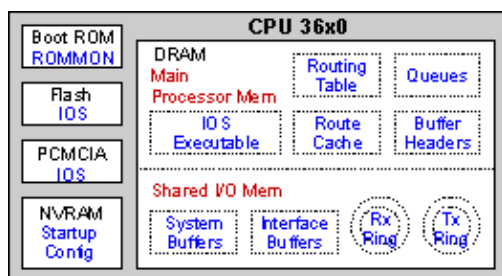
If the router finds a valid image, the following events happen:

- The main Cisco IOS software image is uncompressed into DRAM and loads from there.
- Cisco IOS software makes required data structures such as interface description blocks (IDBs), carves Interface Buffer on DRAM, loads the Startup Configuration, and is ready to go.

If the router is stuck in ROM Monitor mode, you can follow the recovery procedures described in ROMmon Recovery for the Cisco 3600 Series Router.

Packet Switching

Buffers Overview



The 3600 Series Routers are particle-based platforms. The interface buffers are atomic buffers, called particles, into which the packets are split. When a packet cannot be fast-switched, the router has to reassemble it in one system buffer because the process switching code cannot handle particles. Public particle pools cannot be tuned.

First In, First Out (FIFO) Buffer

FIFO buffer memory is used for very temporary storage of a few bytes of data to be transferred between devices (chips).

- FIFO buffer memory is located on network interface chips and other controller chips, as needed.
- There are no queues associated with these buffers.

Rx Ring and Tx Ring

Rx Ring and Tx Ring are buffers used for immediate reception and transmission of packets by the interface driver.

- They are located on Network Modules (NMs), and their size varies.
- Packets reside in Rx Ring until interface buffers are available, or in TX Ring until interface hardware can transmit them.
- They are implemented in a circular FIFO manner, so there are no queues associated with them.

Interface Buffers

Interface buffers (particle-based) store packets for transfer between an interface driver and Fast Switching code.

- The private interface buffer pools are located on shared I/O memory (in DRAM) for all interfaces.
- The number of interface buffers depends on total shared I/O memory, the interface maximum transmission unit (MTU), interface bandwidth, and the total number of interfaces. They are carved upon router startup.
- These buffers are static and do not grow or trim dynamically based on need.
- These buffers have associated "transmit queues" per interface for sending packets.

System Buffers

System buffers store packets for transfer between Fast Switching and Process Switching code.

- They are located on shared I/O memory (in DRAM). Buffer headers for these buffers are of type PAKTYPE, and are located in Main Processor Memory (in DRAM).
- System buffers can be Small (104 bytes), Middle (600 bytes), Big (1524 bytes), Verybig (4520 bytes), Large (5024 bytes), and Huge (18024 bytes).
- The total number of each buffer varies depending on shared I/O memory (in DRAM) size.
- These buffers are dynamic, so they can grow or trim, as needed.
- System buffers are public, so packets from any interface can be placed there, depending on the MTU.
- **Input and Output Hold queues** for each interface are associated with these buffers. The hold queues exist in Main Processor Memory (in DRAM).

Buffer Allocation

Buffer allocation involves several issues such as calculating buffer sizes, the number of buffers, and the physical location in memory.

- Buffers consist of a Packet Data portion, and a Header which has pointers to the Data portion.
- The size of interface buffers depends on the MTU, and is grouped for similar MTUs. System Buffer pools are grouped as Small, Medium, Large, Huge, and so on.
- The number of interface buffers to be allocated depends upon available shared I/O memory (in DRAM) and the total number of interfaces. The number of System Buffers is configurable with some available defaults (also in shared I/O memory).
- Once the size and number of buffers in a pool is decided, Buffer Headers are created and assigned to a Free-List for each pool.
- Calls to a function **malloc()** determine the availability and location of physical memory. If sufficient contiguous memory is not available during a call to malloc(), **malloc failures** may be reported.

Memory is returned after use by calling a function **free()**.

- When a buffer is needed, a function **getbuffer()** is called, which grabs a Buffer Header from the Free-List. If Buffer Header is not available in the Free-List, a new buffer has to be created. This is done immediately at the Process level, but not at the Interrupt level (in which case the Net background process is requested that creates needed buffers later on). Buffers are returned to the Free-List when done, using a **Return queue**.
- Processes that allocate memory but fail to eventually release it after use cause **memory leaks**. Also, allocating various sizes of buffers at random times and random locations causes **memory fragmentation**, and contiguous memory becomes unavailable. Note that buffer allocation is not the only portion of Cisco IOS software that uses `malloc()` and `free()`. See Troubleshooting Memory Problems for more information on troubleshooting memory issues.

Switching Paths

Packet Receive Stage

Incoming packets on a wire are handled by the Network Interface logic chip and Interface Driver. The process is as follows:

1. The chip places the incoming bits in the limited FIFO buffers. Physical layer functionality and signaling is handled by the chip.
2. The chip flags an **Rx Interrupt** for the CPU which runs Interface Driver code to write the packet to an Interface Buffer in Shared I/O Memory (in DRAM using the DMA by the System Controller) based upon available Buffer Headers in the Free-List. If Interface Buffers are not available in the private Interface Buffer pool, then the Public Buffer pools are checked. The packet may be copied to a network module's local memory (if any) before transfer to an Interface Buffer. If the System Controller is not fast enough in pulling the data, the chip overflows the FIFO buffer and an **overrun** is reported.
3. Basic data integrity and Data Link layer checks are performed by the Interface Driver, and appropriate counters are incremented (such as **cyclic redundancy check (CRC)**, **giants**, **frame error**, and so on).
4. If buffers are not available, the Interface Driver ignores the incoming frame (and increments the no buffer counter).
5. The CPU classifies the packet and tries to Fast switch (see details below). If Fast Switching fails from Interface Buffers, the packet is placed in System Buffer (by moving buffer pointers in Shared I/O Memory), and another Interface Buffer is returned to its Free-List.
6. The packet is now enqueued for Process Switching code, and the interface's Input Hold queue counter is incremented. If the Input Hold queue is full, an input queue drop is reported.
7. At this time, the CPU's Rx Interrupt has been handled, and the Scheduler gives control of the Main Processor back to the previously running process.

Packet Switching Stage

Process switching involves the forwarding of packets after creation of the outgoing Frame Header. The outgoing Frame Header is based upon the destination address lookup in the routing table of the related protocol for the packet's next-hop address and outgoing interface.

This operation is performed by the router's Main Processor at the process level (under Scheduler control).

1. Packets are stored in System Buffers in Shared I/O Memory (in DRAM) while they wait for the processor to forward them.
2. All packets destined for the router (such as pings, routing updates, and so on) are enqueued for related application code at the process level, but stay in System Buffers.
3. After a forwarding decision has been made, the outgoing Frame Header is copied on top of the old header. The Processor attempts immediate transmission. If the Interface is not busy and the Transmit

Queue Limit has not been reached, the frame is transferred, and System Buffer is returned to its pool's Free-List. Otherwise, the frame is enqueued for the outgoing interface (in Output Hold queue). If the Output Hold queue is full, an output queue drop is reported.

4. Outgoing Frame Headers from successful forwarding decisions are placed in the Fast Switching cache for future packets.

Fast switching involves the forwarding of packets based upon the destination address lookup in a special route cache that has been built previously by the process switching of the first packet.

This operation is performed by the router's main processor (CPU) at the Rx Interrupt level (under Interrupt handler).

1. Packets are stored in Interface Buffers and get immediate attention from the processor to forward them.
2. A packet may not be Fast switched if:

- ◆ Fast Switching is not supported or enabled for that protocol
- ◆ An entry for this packet's destination doesn't exist in the Route Cache
- ◆ The packet is destined for the router

If a packet cannot be Fast switched, it is placed in a System Buffer (by moving buffer pointers in Shared I/O Memory) and enqueued for Process Switching code (as mentioned earlier). The Free-List pool of the Incoming Interface Buffer is adjusted.

3. After a forwarding decision has been made, a new outgoing Frame Header is copied on top of the old Frame Header, and the packet stays at the current location (but counted towards the outgoing interface by modifying pointers). The Incoming Interface Buffer pool's Free-List is adjusted.
4. If everything has gone well, the packet is now enqueued for Interface Driver code in Transmit queue. If that is full, the packet is dropped. At this time, the CPU's Rx Interrupt has been handled.

NetFlow switching involves the forwarding of packets after classifying the traffic per flow. A flow is defined as a unidirectional sequence of packets between the given source and the 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.
- The minimum supported Cisco IOS software version for NetFlow switching is 12.0(2), 12.0T, and 12.0S.

Cisco Express Forwarding

The main drawbacks of the previous switching algorithms are:

- The first packet for a particular destination is always process-switched to initialize the fast cache.
- 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.
- There is no direct relation between the fast-cache and the Address Resolution Protocol (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 and repopulation of the cache can become CPU-intensive with very large networks.

Cisco Express Forwarding addresses these issues by using two tables: the Forwarding Information Based (FIB) 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 Cisco Express Forwarding is that the database structure allows load balancing per destination or per packet. The Cisco Express Forwarding home page provide more information about Cisco Express Forwarding.

The minimum supported Cisco IOS software version for Cisco Express Forwarding is 12.0(5)T.

Packet Transmit Stage

Outgoing packets are handled by the Interface Driver and Network Interface logic chip.

- In case of Process Switching, the packet is in a System Buffer in Shared I/O Memory (in DRAM) while it waits in the Output Hold queue. Once Scheduler gives time to the output process, the packet is enqueued for the Interface Driver in the Transmit queue. After packet transfer, the System Buffer is returned to its pool's Free-List.
- In case of Fast Switching, the packet is in Interface Buffers in Shared I/O Memory (in DRAM), and is enqueued in the outgoing interface's Transmit queue.
- The packet is taken from the Interface Buffers or System Buffer, and placed in the outgoing interface's FIFO buffer (or temporarily in local memory, if present). The network interface logic chip pulls bits out of the FIFO buffer and places them on the wire as electrical signals with appropriate timing.
- If the Interface Driver is not fast enough in providing the data to the chip in FIFO buffer, the chip will run out of bits to transmit, and an **underrun** will be reported. All other output-related problems, such as **collisions**, are reported at this stage.
- Once the packet has been transmitted, a **TX Interrupt** is flagged for the CPU's Main Processor. Control of the Processor is given to the Interface Driver code, which returns the Interface Buffers or System Buffer to their pool by placing them in the Return queue. At this time, TX Interrupt has been handled, and the Processor handles other operations.

Performance Figures

The table below indicates the performances for the different switching paths in packets processed per second (pps):

Platform	Process	Fast
3620	2,000 pps	20,000–40,000 pps
3640	4,000 pps	50,000–70,000 pps
3660	10,000–12,000 pps	100,000–120,000 pps

Related Information

- **Password Recovery Procedure for the Cisco 3600 Series Routers**
- **Installing Power Supplies in Cisco 3600 Series Routers**
- **Software Installation and Upgrade Procedure**
- **ROMmon Recovery for the Cisco 3600 Series Router**
- **Data Compression Advanced Integration Modules**
- **Cisco 3600 Series Memory Options and Configuration Guide**

- **Maximum Number of Interfaces and Subinterfaces for Cisco IOS Software Platforms: IDB Limits**
 - **Software Advisor Tool (registered customers only)**
 - **Cisco 3600 Series Multifunction Platforms (3620 and 3640/3640A)**
 - **Cisco 3660 Multiservice Platform**
 - **Technical Support & Documentation – Cisco Systems**
-

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

© 2010 – 2011 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: May 15, 2006

Document ID: 7442
