

Chapter 1

Product Architecture

The JUNOS software provides IP routing protocol software—as well as software for interface, network, and chassis management—specifically designed for the large production networks typically supported by Internet service providers (ISPs). The JUNOS software runs on all Juniper Networks routers. For more detailed information about hardware features, see the hardware installation guide for your router.

This chapter provides an overview of the router hardware and then discusses the relationships between the hardware and the software:

Hardware Overview on page 3

Product Architecture on page 4

Hardware Overview

The JUNOS software runs on two types of Juniper Networks routers: M- and T-series routers. The routers consist of the following major hardware components:

Sheet metal of the chassis.

Power supplies (AC or DC).

Impeller trays.

Fan assemblies.

Routing Engine.

System Control Board (SCB), System and Switch Board (SSB), Switching and Forwarding Module (SFM), Switch Interface Board (SIB), or Forwarding Engine Board (FEB).

Flexible PIC Concentrators (FPCs), each populated by up to four Physical Interface Cards (PICs) for various interface types, including SONET/SDH OC-192, OC-48, OC-12, and OC-3, ATM OC-12, OC-3, DS3 (T3), E3, DS1 (T1), E1, Gigabit Ethernet, Fast Ethernet, and Channelized OC-12. Some PICs do not require an FPC.

Product Architecture

The router is composed of two components (see Figure 1):

Packet Forwarding Engine—Forwards packets through the router. The Packet Forwarding Engine is a high-performance switch that is capable of forwarding 40 million packets of any size per second.

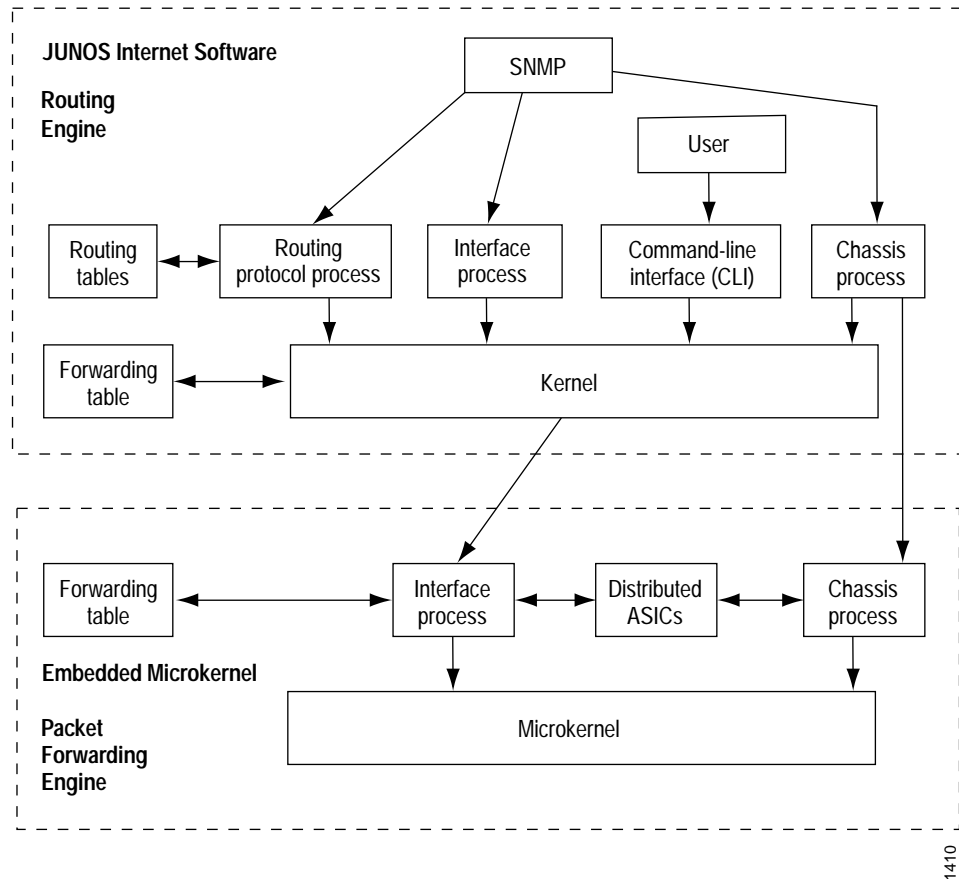
The Packet Forwarding engine uses ASICs to perform Layer 2 and Layer 3 packet switching, route lookups, and packet forwarding. On M-series routers, the Packet Forwarding Engine includes the router (on the midplane, except on the M40 router, where it is on the backplane), FPCs, PICs, and other components, unique to each router, that handle forwarding decisions.

The T-series routers feature multiple Packet Forwarding Engines, up to a maximum of 16 for the T640 Internet routing node and 8 for the T320 Internet router. Each FPC has one or two Packet Forwarding Engines, each with its own memory buffer. Each Packet Forwarding Engine maintains a high-speed link to the Routing Engine. For information about T-series routers, see the *T640 Internet Routing Node Hardware Guide* and the *T320 Internet Router Hardware Guide*.

Routing Engine—Performs routing updates and system management. The Routing Engine consists of routing-protocol software processes running inside a protected memory environment on a general-purpose computer platform. The Routing Engine has a direct 100-Mbps connection to the Packet Forwarding Engine.

Because this architecture separates control operations such as routing updates and system management from packet forwarding, the router can deliver superior performance and highly reliable Internet operation.

Figure 1: Product Architecture



Packet Forwarding Engine

The Packet Forwarding Engine forwards packets between input and output interfaces. The M-series and T-series routers use a different architecture for packet forwarding, and T-series routers can have multiple Packet Forwarding Engines.

Packet Flow through an M-series Router

You can understand the function of the Packet Forwarding Engine by following the flow of a packet through the router—first into a PIC, then through the switching fabric, and finally out another PIC for transmission on a network link.

When a packet arrives on an input interface, a media-specific PIC performs all the media-specific actions, such as framing and checksum verification.

The PIC then passes a serial stream of bits into the FPC, which parses and appropriately deencapsulates the packet. For M-series routers, the FPC also breaks the packet into 64-byte memory blocks and passes each memory block to the Distributed Buffer Manager ASIC. The Distributed Buffer Manager ASIC or the Queuing and Memory Interface ASIC (on the T640 routing node), then writes the blocks into packet buffer memory, which is distributed evenly across all the FPCs installed in the router.

In parallel with the buffering, the Distributed Buffer Manager ASIC extracts the information from the packet needed for route lookup and passes that information to the Internet Processor ASIC or Internet Processor II ASIC, which performs a lookup in its full forwarding table and finds the outgoing interface and the specific next hop. The forwarding table can forward all unicast packets that do not have options and multicast packets that have been previously cached. Unicast packets with options and noncached multicast packets are sent to the Routing Engine for resolution.

After the Internet Processor ASIC has determined the next hop, it passes the results of the lookup to the second Distributed Buffer Manager ASIC or Queuing and Memory Interface ASIC, which in turn passes the results to the outgoing interface. (Note that there can be multiple outgoing interfaces if you are using multicast routing.)

At this stage, a pointer to the packet is queued, not the packet itself. Each output port has four queues, each of which has a configured share of the link bandwidth. Several factors can account for queuing order, including the value of the precedence bits, utilization of the input interface, destination address, and Random Early Detection (RED) and Weighted Random Early Detection (WRED) algorithms. If the outgoing interface decides to queue the packet to be sent, when the packet reaches the front of the queue and is ready for transmission, the memory blocks are read from packet buffer memory. Then the packet is reassembled and passed to the media-specific PIC for transmission on the line.

Packet Flow through a T-series Router

When a packet arrives on an input interface, the PIC performs media-specific operations such as framing and checksum verification and passes the packet to the FPC housing it. On the FPC, the Layer 2/Layer 3 Packet Processing ASIC parses the packet and divides it into data cells. The cells are sent to the Switch Interface ASIC, which extracts the notification and passes it to the T-series Internet Processor. The data cells are sent to the Queuing and Memory Interface ASIC, which manages data buffering on the FPC.

The Queuing and Memory ASIC sends the notification to the Switch Interface ASIC facing the switch fabric, which sends bandwidth requests to the destination FPC and issues read requests back to the Queuing and Memory ASIC to begin reading the data cells out of memory. The Switch Fabric ASIC on the destination FPC sends bandwidth grants to the originating Switch Fabric ASIC, which sends a cell to the destination FPC for each grant received.

On the destination FPC, the Switch Interface ASIC receives cells from the switch fabric. It extracts the route lookup key, places it in a notification, and forwards the notification to the T-series Internet Processor. The T-series Internet Processor performs the route lookup and forwards the notification to the Queuing and Memory ASIC, which passes it to the Switch Interface ASIC facing the network. The Switch Interface ASIC passes the cells to the Layer 2/Layer 3 Packet Processing ASIC, which reassembles them into packets, performs encapsulation, and sends the packets to the outgoing PIC. The PIC sends the packets out into the network.

Routing Engine

The Routing Engine handles all the routing protocol processes and other software processes that control the router's interfaces, a few of the chassis components, system management, and user access to the router. These routing and software processes run on top of a kernel that interacts with the Packet Forwarding Engine.

The Routing Engine has these features:

Process routing protocol packets—All routing protocol packets from the network are directed to the Routing Engine, and therefore do not delay the Packet Forwarding Engine unnecessarily.

Software modularity—By dividing software functions into separate processes, a failure of one process has little or no effect on the other software processes.

In-depth Internet functionality—Each routing protocol is implemented with a complete set of Internet features and provides full flexibility for advertising, filtering, and modifying routes. Routing policies are set according to route parameters, such as prefix, prefix lengths, and BGP attributes.

Scalability—The JUNOS routing tables are designed to hold all the routes in current and near-future networks. Additionally, the JUNOS software can efficiently support large numbers of interfaces and virtual circuits.

Management interfaces—System management is possible with a command-line interface (CLI), a craft interface, and SNMP.

Storage and change management—Configuration files, system images, and microcode can be held and maintained in one primary and two secondary storage systems, permitting local or remote upgrades.

Monitoring efficiency and flexibility—Alarms can be generated and packets can be counted without adversely affecting packet forwarding performance.

The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the *forwarding table*, which is then copied into the Packet Forwarding Engine. The forwarding table in the Packet Forwarding Engine can be updated without interrupting the router's forwarding.

