

Chapter 1

Product Architecture

The JUNOS Internet 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 Internet software runs on all Juniper Networks routers. For more detailed information about hardware features, see the hardware installation guide for your router model.

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 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), or Forwarding Engine Board (FEB)

Flexible PIC Concentrators (FPCs), each populated with 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 and OC-3, DS3 (T3), E3, DS1 (T1), E1, Gigabit Ethernet, Fast Ethernet, and Channelized OC-12. Some PICs do not require an FPC.

A fundamental architectural feature is the use of shared memory as the interconnect between slots. Specifically, when a packet arrives on an input interface, it is placed into a buffer where it stays until it is sent out of the output interface. This architecture has several consequences. First, given that the complexity of the system is partly due to the number of buffering stages, the architecture is relatively clean and simple. Second, the centralized buffer can support buffering for each interface that is equal to the bandwidth times the delay and therefore can meet TCP's buffering needs in order to maximize throughput. Finally, the shared memory architecture supports multicast traffic at nearly the theoretical maximum efficiency.

Product Architecture

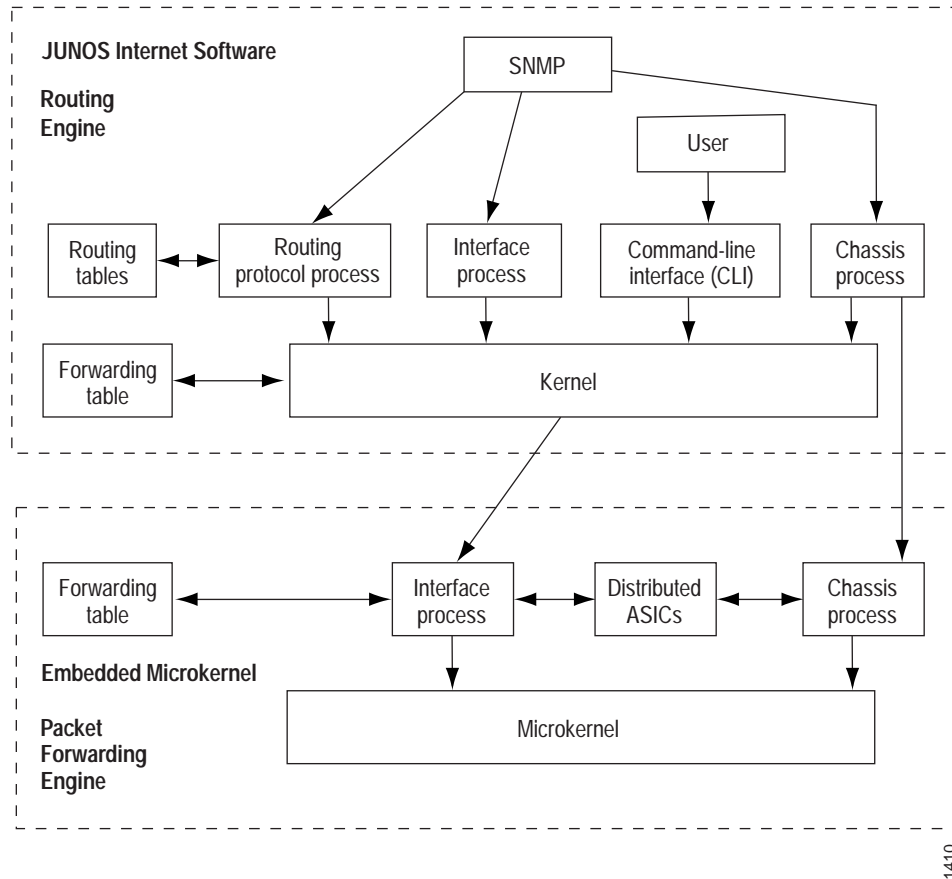
The router is composed of two components (see Figure 1):

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 per second for any packet size.

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.

Packet Flow through the Router

The function of the Packet Forwarding Engine can be understood by following the flow of a packet through the router—first into a PIC on an FPC, 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 details such as framing and checksum verification.

The PIC then passes a serial stream of bits into the FPC, which parses and appropriately decapsulates the packet. 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 then writes them 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, 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 a second Distributed Buffer Manager ASIC, which in turn passes it to the outgoing interface. (Note that there could be multiple outgoing interfaces in the case of multicast.)

It is at this point that a pointer to the packet is queued, not the packet itself. Each output port has four queues, each of which has some 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 RED and WRED algorithms. If the outgoing interface decides to queue the packet for transmission, 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.

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.



