Hardware Overview 2

This chapter provides a general system hardware overview and describes all hardware components.

Hardware Design Overview

The ERX system uses a modular carrier-class design with a passive midplane, active front-insert line modules, and high-reliability rear-insert input/output (I/O) modules. Both the ERX-700 series and the ERX-1400 series use the same line modules and I/O modules; however, the ERX-1440 system uses only the higher-performance line modules. Consequently, service providers can limit their inventories of spare modules, and easily upgrade from the ERX-700 series to the ERX-1400 series as the subscriber base at a POP or central office increases. Both systems support full redundancy and line module hot-swapping to optimize network uptime. Figure 2-1 shows the ERX-1400 series system architecture. The ERX-700 series architecture is similar; however, the modules fit horizontally in the chassis, rather than vertically.

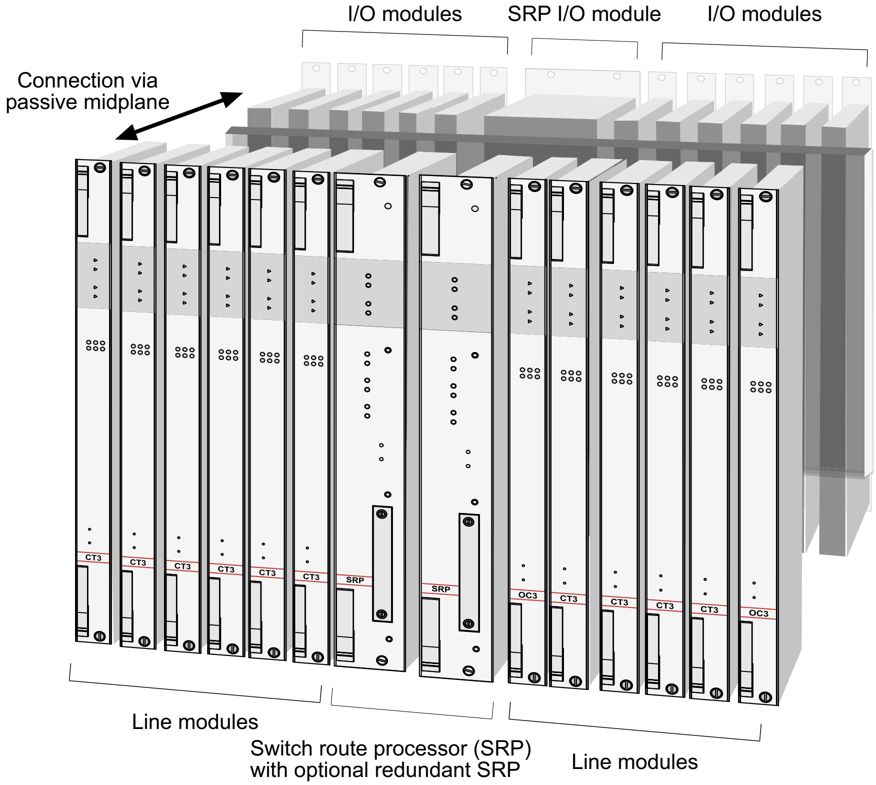

Figure 2-1 Overview of the ERX-1400 series system architecture

ERX System Architecture

The Juniper Networks ERX system architecture is composed of three main elements:

- A shared switch fabric that operates at 5, 10, or 40 Gbps
- Hardware-assisted forwarding engines distributed to each line module
- A high-performance route processor for routing table maintenance and system configuration

The system uses a highly distributed multiprocessor architecture that distributes processing function to each port in the system to speed up decision making and scale system growth. Consequently, the ERX system supports next-generation routing features that provide differentiated services, wire-speed packet filtering, buffer management, and scheduling.

Data Flow

Figure 2-2 shows the data flow through the hardware of the system. Line modules handle packet processing and forwarding. A switch fabric performs high-speed internal packet switching. A route processor gathers the routing table information and sends route tables and updates to the line modules.

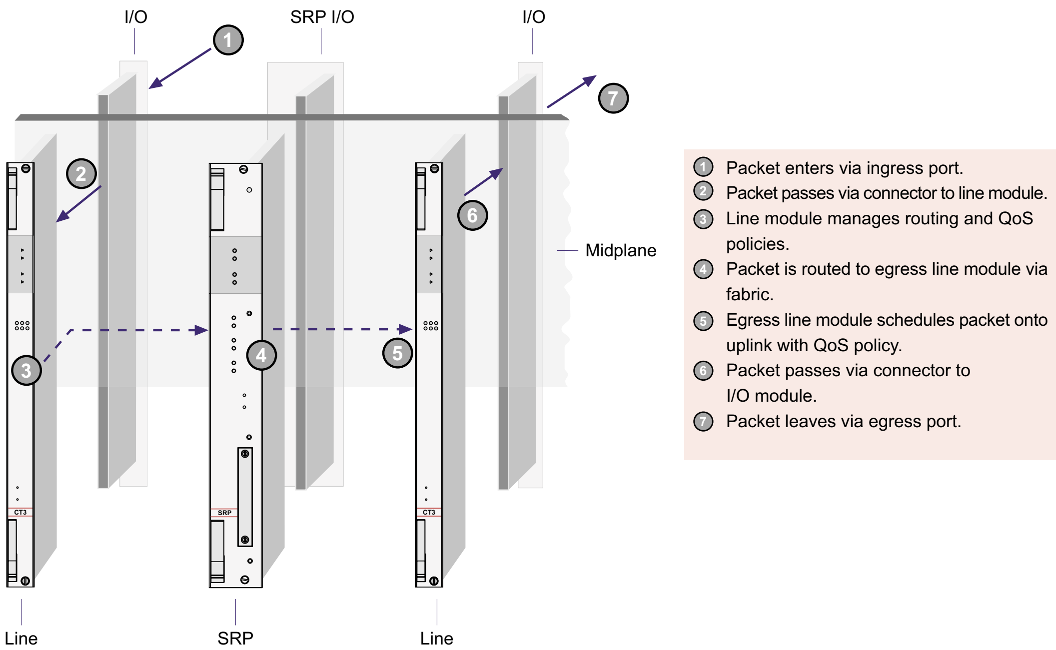

Figure 2-2 Data flow through the ERX system hardware

System Design

The ERX system's innovative design provides a fast path for IP traffic by taking the route processor out of the forwarding path and maintaining an entire routing table on each port. The route processor on the SRP module uses a private multicast channel through the switch fabric to update route tables on forwarding engines and issues updates only when changes occur. For every packet, the system uses its powerful hardware to perform a full table search at wire speed. This design allows the system to avoid the nondeterministic performance of cache-based searches and the associated uncertainty in packet latency.

The ERX system design also supports wire-speed QoS handling; consequently, the system can examine, classify, and queue packets at wire speed. A scheduling function then regulates the packet flows onto the egress link based on the QoS policy assigned to each flow. All QoS policies for all flows are centrally configured and then downloaded for execution to each line module. Because all QoS handling is distributed to the line modules, processing speed is maintained even when you add extra line modules. Regardless of the number of QoS policies or the volume of traffic, the system offers consistently high performance.

ASIC Technology

The ERX system makes use of application-specific integrated circuits (ASICs) in order to speed IP packet processing. The Juniper Networks custom ASICs, called the ESP (edge service processor) family, help meet the forwarding demands of the Internet edge.

Older-generation routers cannot cope with the processor demands of QoS functions, such as classification, queuing, and scheduling while routing packets at wire speed. The ERX system overcomes these limitations with its hardware-assisted architecture.

The system's use of ESP ASICs provides significant benefits to the service provider. ASICs help to:

- Reduce product size
- Reduce product cost, especially packet per second cost
- Increase packet-processing speeds dramatically by applying targeted power to hard-wired tasks

ASIC Chip Sets

The ERX system uses two separate ESP ASICs: an ingress forwarding ASIC and an egress forwarding ASIC. These chip sets are deployed as a pair on each ERX line module. The generic network interface design allows the same chips to be used on all line modules.

Each chip has approximately 1,000,000 gates. A number of FPGAs (Field Programmable Gate Arrays) are also used on each line module to add incremental processing assistance. The line modules also use a number of reduced instruction set computing (RISC) processors to maintain software flexibility. This hardware-intensive distributed architecture distributes packet processing to the line modules to achieve maximum wire-speed throughput on all line interfaces while maintaining a highly scalable and flexible architecture.

RISC + ASIC Architecture

The ERX system uses a RISC + ASIC architecture that provides both high-performance packet management and the flexibility to support changes in standards and specialized customer enhancements. The ASIC technology supports wire-speed forwarding of packets on all system interfaces by hard-wiring standard routing functions. The RISC technology provides flexibility, since it can be easily upgraded to handle software enhancements. The combination of the technologies is a powerful and key architectural necessity for high-performance edge routing switches, overcoming the disadvantages of all-ASIC (no flexibility) or all-RISC (too slow) approaches.

Hardware-Assistance Features

As a next-generation edge product, the ERX system uses hardware assistance to maintain wire-speed processing at high bandwidth rates, while supporting the processing-intensive IP service functions that have become a service provider necessity.

The system implements hardware assistance for the following functions:

- Basic router functions – frame cracking, address lookup
- Classification – mapping the IP header to the classification filters
- Rate measurement – packet monitoring and shaping via token buckets
- Buffer management – allocation of frames to buffers and clearing of buffers after packet transport. The system manages buffers from a dynamic pool in contrast to older routers that rely on fixed buffer assignments. This architecture allows faster processing and more flexible distribution of buffer resources per line module.
- Queue management – pointer management
- Scheduling management – scheduling the packets on the output links in relation to their QoS service plans
- Statistics support – incrementing counters in response to packet information

To speed packet processing and reduce packet latency, these functions run in parallel in the system.

Redundancy Features

Both the ERX-700 series and the ERX-1400 series have a passive midplane design. Line modules that contain active components plug into the front of the chassis. I/O modules that contain the I/O ports have mostly passive components and plug into the rear of the chassis. This design provides redundancy by allowing multiple line modules to share cables and by supporting zero-touch sparing. If a line module goes down, the redundant line module uses the same cables, requiring no manual intervention or cable swapping from an on-site technician.

1:1 Redundancy

The SRP module uses a 1:1 redundancy scheme. When two SRPs are installed in the ERX system, one acts as a primary and the second as a standby. Both SRP modules share a single SRP I/O module located on the rear of the chassis. If the primary SRP fails, the secondary SRP assumes control.

1:N Redundancy

Most line modules support a 1:N redundancy scheme to support zero-touch sparing. In this process, an extra line module in a group of identical modules provides redundancy for all those modules. This level of redundancy is critically important because thousands of subscribers connect via single channelized interfaces, and because it is increasingly common for facilities to operate in lights-out mode with no personnel on site to swap cables.

To use this feature, you must install:

- A spare line module in the lowest slot of the redundancy group
- A redundancy midplane that covers the slots in the redundancy group
- A redundancy I/O module in the lowest slot of the redundancy group

The special midplane can cover two through six slots. It provides additional connectivity that enables the spare line module to assume control of the I/O module associated with any failed line module in the redundancy group. The spare I/O module provides connectivity from the spare line module to the redundancy midplane.

The process by which the ERX system switches to the spare line module is called *switchover.* When switchover occurs, the system breaks the connection between the primary I/O module and the primary line module. Next, the system connects the primary I/O module to the spare line module via the redundancy midplane and redundancy I/O module. Protocol processing then takes place on the spare line module.

Port Redundancy

Some channelized and unchannelized OCx/STMx I/O modules support SONET Automatic Protect Switching 1+1 redundancy (APS) and SDH Multiplex Section Protection (MSP). See *Line Modules and I/O Modules,* later in this chapter, for information about which modules support this feature. This feature meets the following specifications:

- Telcordia document GR-253 – Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria, Revision 3 (September 2000).
- ITU-T G.783 – Characteristics Of Synchronous Digital Hierarchy (SDH) Multiplexing Equipment Functional Blocks: Annex A – Multiplex Section Protection (MSP) Protocol, Commands And Operation (1990)

With this type of redundancy, there are two fiber ports on the I/O module. Both ports receive the same SONET/SDH packets. The ERX system examines the SONET/SDH packets to determine whether or not there is a problem with the optical connection to either port. Based on this information, the system chooses which port is active and which port is redundant. If there is a problem with the connection to the active port, the system switches data transfer to the redundant port.

The GE I/O module uses a similar scheme with its two ports; one active Gigabit Ethernet connection is supported at a time. In the event of a cut or loss of the fiber cable connected to the primary port, the redundant port becomes active.

High-Availability Features

The ERX system implements two important high-availability features in software to further guarantee subscriber traffic uptime:

- When the redundant SRP module configuration is used, all user traffic continues to be forwarded even in the event of a one-degree SRP module failure. This is handled through fabric connection mirrored on the two-degree backup SRP.
- The ERX system implements IETF drafts to maintain BGP peer information, even in the event that the ERX system reboots. This feature dramatically reduces overall network conveyance frequency and therefore increases uptime.

ERX System Chassis

Both the ERX-700 series and the ERX-1400 series use the same line modules and I/O modules; however, the ERX-1440 system uses only the higher-performance line modules. This design enables service providers to maintain their investment in modules if they upgrade from the ERX-700 series to the ERX-1400 system. In addition, if customers have more than one ERX model, they can minimize inventories of spare modules.

The midplane is built with passive components and merely handles the clock, control, and management traffic for the module-to-module communication, and the passing of subscriber data between modules. Table 2-1 shows ERX system chassis features.

Table 2-1 ERX system chassis features

Chassis Specifications	ERX-700 series	ERX-1400 series
Total number of slots	7	14
Slots reserved for SRP module	2	2
Slots available for line module–I/O module pairs	5	12
Dimensions	10.5 (H) x 19 (W) x 16 (D) inches (26.67 x 48.26 x 40.64 cm)	22.75 (H) x 19 (W) x 16 (D) inches (57.78 x 48.26 x 40.64 cm)
Weight	• Chassis only: 22 lb (9.9 kg) • Chassis fully configured: 46 lb (20.7 kg)	• Chassis only: 42 lb (18.9 kg) • Chassis fully configured: 88 lb (39.6 kg)

Power System

Both the ERX-700 series and ERX-1400 series are designed for DC power input. The Juniper Networks unique distributed power architecture distributes redundant –48 VDC feeds through the system to each line module, SRP module, and fan assembly. Conversion of the required secondary voltages takes place at the line module. This system design allows the addition of faster and more powerful line modules, without the power restrictions of a centralized architecture. If a power component fails, the design ensures that other components will not be affected.

The ERX system offers extremely low power consumption (50 Amps at –48 VDC for the ERX-1400 series and 30 Amps at –48 VDC for the ERX-700 series) and a correspondingly low heat dissipation (less than

2400 W for a fully populated ERX-1400 series and less than 1400 W for a fully populated ERX-700 series). The system's efficient electrical design minimizes the overall power needs in a POP, and satisfies stringent co-location power restrictions.

If AC power is desired, a DC to AC power converter is available as part of the system.

Fans/Air Flow

The system uses a single fan tray assembly with six fans mounted on the tray in the ERX-1400 series and four fans in the ERX-700 series. This assembly provides power- and status-sensing capability. The fan tray consists of two redundant converters that power the fans. If one converter fails, the other takes over. In addition, the system software reports an alarm if any of the fans overrotate or underrotate or if one of the converters fails.

Airflow in the chassis varies according to the model. The ERX-1400 series uses an innovative front-to-back and bottom-to-top airflow design to keep the unit cool even when it is mounted in a rack with multiple other devices. In the ERX-1400 series, the fans are located at the top of the chassis. The ERX-700 series has a conventional airflow, from right-to-left; fans are located on the left side of the chassis.

A separate fan try is available for customers who desire a dual tray option.

System Modules

You can use several different line modules and I/O modules in both the ERX-700 series and ERX-1400 series. A few architectural characteristics are common to certain types of modules:

- The presence of an EEPROM on each module to identify the module and its manufacturing date.
- The location of nonvolatile storage on each line and SRP module to store its own boot image and run-time code.
- A hardware-based lookup, forwarding component, and dedicated RISC processor on each line module. This distributed processing power guarantees wire-speed classification and forwarding of 40-byte packets, even when the system is fully loaded.

SRP Modules

An ERX system must contain at least one SRP module and its associated SRP I/O module. See Figure 2-3. You must insert the SRP module into the slots reserved for this purpose. In the 14-slot chassis, the two center slots are reserved for SRP modules; in the 7-slot chassis, the two bottom slots are reserved. You can install another SRP module in the second slot to provide redundancy. See *Redundancy Features* earlier in this chapter.

Five models of SRP module are supported:

- SRP-40G+

  This model offers a 40-Gbps switching fabric and is used in the ERX-1440 system. The SRP-40G+ module supersedes the SRP-40G module.

- SRP-40G

  This model offers a 40-Gbps switching fabric and is used in the ERX-1440 system. The software still supports this model, which is superseded by the SRP-40G+ module.

- SRP-10G

  This model offers a 10-Gbps switching fabric and is used in the ERX-1410 system and the ERX-700 system.

- SRP-5G+

  This model offers a 5-Gbps switching fabric and is used in the ERX-705 system. The SRP-5G+ module supersedes the SRP-5G module.

- SRP-5G

  This older model offers a 5-Gbps switching fabric and is used in the ERX-700 system. The software still supports this model, which is superseded by the SRP-5G+ module.

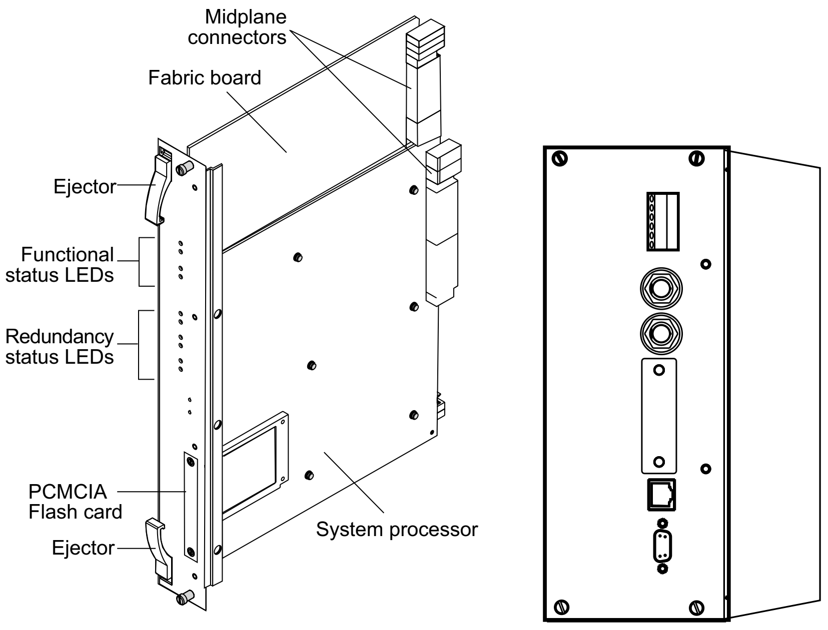

Figure 2-3 SRP and SRP I/O modules

The SRP and SRP I/O functions include:

- Configuration and operation of the 5-, 10-, or 40-Gbps switch fabric and management of the unit interprocessor communication and system control bus master
- Initial system boot, system initialization, and monitoring of fans, power supply, and temperature sensors
- Storage of system software image and system configuration in Type II PCMCIA NVS cards and download of executables and images to the line modules
- Clocking functions, including real-time clock with battery backup, and clock distribution to line modules
- Redundancy control for switch fabric/route server, timing subsystem, and line modules through keepalive messages
- Maintenance of the master route table and link state information via a dedicated RISC processor, and distribution of the route tables to the line modules
- Diagnostic user interface, system alarm contacts
- Watchdog timer for internal management

- System console ports for management access
- Command line interface support

As shown in Figure 2-3, the SRP module comprises two boards:

- Fabric board – 5-, 10-, or 40-Gbps switch fabric
- Processor board – dedicated processor boots the system, manages diagnostics, and supports route protocol processing

The SRP module specifications are provided in Table 2-2.

Table 2-2 SRP module and SRP I/O module specifications

Feature	Specification
Management ports	• One RJ45 10/100 Base-T Ethernet management port • One DB9 RS232 connector for system management
Alarm ports	• One terminal block for external alarm contacts
Clock ports	• Two 3-pin wire-wrap posts for US external clock sources • Two BNC connectors for E1 clock sources • Two connectors to support primary and secondary clock sources
Line module front panel	• Indicates board, power, fan, link, and redundancy status • PCMCIA slot access
Processor support	• PowerPC 750 with 1-MB level 2 cache
Switch fabric	• 5-, 10-, or 40-Gbps switch fabric
Memory support	• Four banks of SDRAM; max of 128-MB each • Upgradable PCMCIA Type II NVS; 220-MB card provided, 440-MB card available

Switching Fabric Functions

The ERX system uses a high-performance, 5-, 10-, or 40-Gbps, packet-aware switching fabric to connect all internal data paths. The fabric handles per flow queuing for large numbers of packet flows.

The fabric’s functions include buffer management, queuing, scheduling, and packet discard. The fabric manages the internal connections between ingress and egress ports and can support point-to-point and multicast connections.

Fabric features include:

- 5-, 10-, or 40-Gbps performance – supports wire-speed throughput for fully loaded configurations
- Per flow queuing – each queue can be assigned QoS parameters and is independently managed and scheduled, so that if congestion conditions occur on one queue, it does not affect other queues. This allows for fairness in QoS policy management and control.
- Dedicated CPU – provides management and control to ensure fast response times to management requests
- Early packet discard (EPD) support – reduces retransmissions by ensuring that only complete packets are sent

The fabric also distributes operational software images and route table updates to the line modules.

System Start-Up

When the system is powered on, the SRP module executes boot code from NVS. The SRP module downloads a software image from NVS into SDRAM, then executes code from SDRAM. Once it is initialized, the SRP module downloads an executable image to each line module in the chassis.

When each line module is active, the SRP module communicates with the line modules through a dedicated, in-band, 150-Mbps Utopia bus connection through the switch fabric.

NVS Cards

If you have two SRP modules installed in a system, you can use NVS cards of different capacities on the SRP modules. The effective capacity of the higher-capacity NVS card will equal that of the lower-capacity NVS card.

The ERX system includes a number of software and hardware features that optimize the way a system restores its configuration if the system is shut down improperly.

Line Modules and I/O Modules

A range of line modules and Input/Output (I/O) modules is available for the ERX system. Line modules process data traffic; I/O modules provide the connections from the ERX system to the network. All line modules, with the exception of the IPSec Service and Tunnel Service modules, combine with I/O modules to provide particular capabilities and connections. For information about available line modules and I/O modules, see *Appendix B, Module Specifications*.

You can use any line module for *access* or *uplink*. Commonly, access line modules receive traffic from low-speed circuits (such as T1 or E1), and the system routes the traffic onto higher-speed uplink line modules (such as OC12/STM4) and then to the core of the Internet.

Most I/O modules have fixed connectors. However, Gigabit Ethernet (GE) modules use small form-factor pluggable transceivers (SFPs). A range of SFPs that support different optical modes (multimode and single mode) and cabling distances is available. You can replace the SFPs without disabling the interfaces or removing the module from the system.

Line Modules for Tunnel Support

The Tunnel Service line module (TSM) supports IP tunnels, such as DVMRP and GRE tunnels, and LNS termination for L2TP. The IPSec Service module (ISM) encrypts and/or decrypts packets encoded with the IPSec protocol.

Unlike other line modules, TSMs and ISMs do not pair with I/O modules that contain ingress and egress ports. When the system receives a packet that requires processing on a TSM, the system forwards the packet via the switch fabric to the TSM or ISM. The TSM or ISM processes the packet and returns it to the switch fabric. Finally, the processed packet leaves the system via an egress module.

Future Modules

Juniper Networks will continue to develop new modules for the ERX system to increase density and functionality. Additional modules will be developed in response to customer demand as network bandwidth needs increase or as interface requirements change.