



Broadband Provisioning Registrar

Broadband Provisioning Registrar (BPR) automates the configuration and provisioning of network devices supported by a broadband service provider. It is a flexible product that can be scaled to suit virtually any size network. BPR is designed to handle the rapid growth of the service provider. It targets broadband service providers (including multiple service operators), Internet, and voice service providers who want to deploy IP data, voice, and video on hybrid fiber and coaxial cable networks. BPR also provides such critical features as redundancy and failover protection, and can be integrated into new or existing environments through the use of a provisioning application program interface (API) that lets you control how BPR operates. The provisioning API can be used to enable BPR to register devices, device configurations, and configure the entire BPR provisioning system.

Features and Benefits

BPR lets multiple service operators (MSOs) meet the rapidly changing demands for data over cable services. Using BPR, you can realize these features and benefits:

- Distributed architecture
- Highly scalable
- Redundant extensible architecture

Architecture

This section describes the basic BPR architecture including:

- Regional distribution unit (RDU) that provides:
 - Device configuration.
 - Provisioning API support, such as device query and configuration query. For more information about this component, see the *Broadband Provisioning Registrar Developer's Guide*.
- Device provisioning engines (DPEs) that provide:
 - TFTP protocol server
 - Configuration cache
 - DPE redundancy
 - Time of day server

- Cisco Network Registrar servers that provide:
 - Dynamic Host Configuration Protocol (DHCP)
 - Domain Name System (DNS)

Registration Modes

Broadband Provisioning Registrar uses registration modes to help implement the business rules which you use to determine the number of interactions with each device. You must be prepared to process any change to a registered device. Consequently, there is a significant quantitative difference between registering 100 cable modems with unregistered computers behind them, and registering 100 cable modems each of which has a potentially large number of registered computers behind it. To control device changes such as these, BPR includes these registration modes:

- Standard mode—In this mode, the modem and all computers behind it are registered. The standard mode has two sub-options.
 - Roaming—This is the default condition and indicates that a registered computer can receive its class of service behind any registered modem.
 - Fixed—This indicates that a registered computer can receive its class of service only when that computer boots up behind a specific modem. If the computer reboots behind a different modem, the computer becomes unprovisioned.
- Promiscuous mode—In the promiscuous mode only the modem is registered, and the DHCP server maintains lease information about the computers behind the modem. All computers behind a registered modem receive network access.
- Network Address Translation (NAT) mode—In the NAT mode only the modem is registered and no information is maintained about the computers behind the modem, but all computers receive network access. The NAT mode is invoked at the hardware level.

Regional Distribution Unit

The RDU is the primary server in the BPR provisioning system. The RDU provides these functions:

- Manages the generation of device configurations
- Acts as a clearinghouse through which all application programming interface (API) requests must pass
- Manages the BPR system

The RDU supports the addition of new technologies and services as they become available. The RDU also provides:

- Provisioning data access and manipulation by means of the API
- Distributing configurations to the DPEs for scalability
- Supporting external clients, operations support systems (OSSs), and other provisioning related functionality by means of the API

The RDU includes these components:

- Provisioning APIs that allow you to interface BPR with other external systems.
- Configuration generation.

The following sections describe these RDU components:

- [Configuration Generation, page 1-3](#)
- [Regional Distribution Unit Failover, page 1-3](#)

Configuration Generation

Device configurations can include customer required provisioning information, such as: DHCP IP address selection, bandwidth, data rates, flow control, communication speeds, and level of service. A configuration can contain DHCP configuration and TFTP files for any device. When an unprovisioned device is installed, and the boot operation performed, a default configuration for the appropriate technology is pulled from BPR and sent to the device, by means of either DHCP or TFTP. The default configuration can be changed for each supported technology.

Regional Distribution Unit Failover

BPR currently supports one RDU per installation. For failover support, use your own hardware redundant system or obtain a similar system with hot-swap capability.

Device Provisioning Engine

The Cisco Device Provisioning Engine (DPE) communicates with the DHCP server and communicates with the RDU to give new devices their configurations. Each DPE caches information for up to 250 thousand devices. You can use multiple DPEs to ensure BPR scalability.

The DPE handles all configuration requests including providing configuration files for devices. It is integrated with the Cisco Network Registrar DHCP server to control the assignment of IP addresses for each device. Multiple DPEs can communicate with a single DHCP server.

The DPE manages:

- Last step device configuration file handling
- Communication of the configuration files through an embedded TFTP server
- Integration with Cisco Network Registrar
- Time of day protocol server

This section describes these major DPE components:

- DPE hardware
- DPE server assignments
- TFTP server
- Configuration cache
- Provisioning groups
- DPE redundancy

DPE Server Assignments

BPR supports multiple DPE servers. These servers communicate with devices, a DHCP failover server pair (DHCP failover configuration), and the RDU. During installation, you must make these DPE server assignments:

- Assign a DPE server to be responsible for one or more logical groups of devices.
- The IP address and port number of the RDU.

The DPE's primary objective is to send configurations to customer devices whenever those devices are powered up or rebooted. To do this quickly, the DPE keeps a copy of the configuration for each device in a local cache database. Usually the DPE only has to add in a few minor pieces of information to the configuration, such as a current IP address of the device and one or more security checks, before it provides the configuration to the device.

TFTP Server

The TFTP server receives requests for files, including DOCSIS configuration files, both from device and non-device entities. This server then transmits the configuration files to the requesting entity.

Provisioning Groups

Provisioning groups allow you to scale BPR implementations to support increasing numbers of devices. To ensure that there is adequate support for a reasonable number of devices (per location), the system-wide number of devices must be split between various provisioning groups.



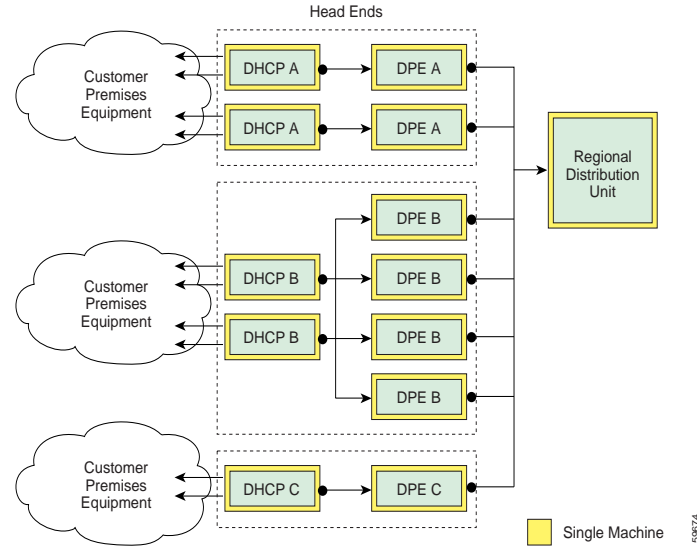
Note

The DHCP server determines which provisioning group a device belongs to. Each device is registered with each DPE within a provisioning group, but not in DPEs belonging to other provisioning groups.

BPR supports multiple DPEs and, during BPR installation, each DPE is assigned to one or more provisioning groups. DPEs are identified, within a provisioning group, as either primary or secondary DPEs and each can support a maximum of 250 thousand devices. The DHCP server assigned to the provisioning group communicates with the DPEs. When handling device requests, the DHCP server first attempts to communicate with one of the primary DPEs in the group and, if none are available, the DHCP server communicates with one of the secondary DPEs. The DHCP server rotates the requests between all DPEs in the assigned provisioning group.

Supporting a group of DPE servers allows the distribution of many device requests across multiple servers. Because DPE servers maintain the configurations locally in their cache, if contact with the RDU is lost, ongoing operations are not affected. New devices cannot be added, but all existing devices continue without interruption.

Figure 1-1 shows an example of provisioning group assignments. In this figure, the provisioning groups to which the DPE and DHCP servers are assigned, are identified as groups A, B, and C. The DHCP servers for each provisioning group service the same customer premises equipment. For example, the pair of DHCP servers for provisioning group A service the same devices.

Figure 1-1 Provisioning Groups and Redundancy

A DPE server assigned to multiple provisioning groups can be a primary or secondary server for each provisioning group. For example, if a DPE is assigned to provisioning groups A, B, and C, it can be a primary server for group A and a secondary server for groups B and C.

The primary DPE servers can service all requests for the provisioning group to which they belong. The secondary servers are backup servers in the event that the primary servers are overloaded or are unavailable. When the DPE starts up, it goes through an automated process of gathering configuration information about all devices present in its assigned provisioning group. In addition, the DPE gathers cached TFTP files and extra data needed to be a productive member of its provisioning group.

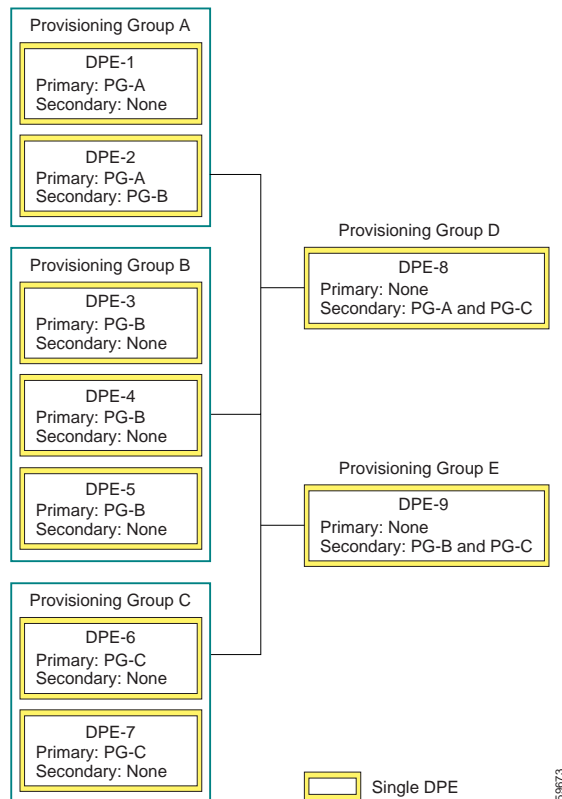
For example, you can configure DPE servers in this way (see [Figure 1-2](#)):

- DPE 1 is a primary server only for provisioning group A.
- DPE 2 is a primary server for group A and a secondary server for group B.
- DPEs 3, 4, and 5 are primary servers only for group B.
- DPE 6 and 7 are primary servers only for group C.
- DPE 8 is a secondary server for groups A and C.

**Note**

The total number of devices supported by a single DPE cannot exceed 250 thousand. Consequently, the total number of devices being supported by all primary and secondary provisioning groups to which a DPE belongs should not exceed 250 thousand.

Figure 1-2 Primary and Secondary DPEs



59673

Device Provisioning Engine Redundancy

DPE redundancy is provided by using multiple primary DPEs, a secondary DPE (both in the same provisioning group), and the DHCP failover protocol.



Note

A provisioning group must contain at least two DPEs to provide DPE redundancy.

Redundancy

A round-robin approach is used to provide DPE redundancy, in the event of a primary DPE failure. This involves using primary and secondary DPEs, within provisioning groups. The servers are configured with the location of the primary and secondary DPEs within the same provisioning group.

A list of primary and secondary DPEs is provided to the RDU and, when DPEs are added to or removed from the provisioning group, the RDU broadcasts an updated DPE list to the Network Registrar servers. In this way the DPEs, and DHCP server, that should receive requests will receive them.

Cisco Network Registrar

Network Registrar provides Dynamic Host Configuration Protocol (DHCP) and Domain Name System (DNS) functionality. It has a complete administrative user interface that, when coupled with customized BPR configuration screens, can be used within a larger enterprise management system.

**Note**

For additional information on Network Registrar, refer to these documents: *Network Registrar User's Guide*, *Network Registrar CLI Reference Guide*, and the *Network Registrar Installation Guide*.

Dynamic Host Configuration Protocol

The DHCP server automates the process of configuring IP addresses and network information devices on TCP/IP networks. This protocol performs many of the functions that a system administrator carries out when connecting a device to a network. DHCP runs a program that automatically manages network policy decisions and eliminates the need for manual configuration. This in itself adds flexibility, mobility, and control to networked device configurations.

DHCP Failover

DHCP failover allows pairs of DHCP servers to act in such a way that one can always be substituted if the other goes down. This is a hot standby arrangement and the server pairs are known as the primary and secondary server. Under normal circumstances the primary server performs all DHCP functions. Should the primary server be unavailable, the secondary server will act as a hot backup for the primary. In this way, DHCP failover prevents loss of access to the DHCP service should the primary fail.

Domain Name System

The Domain Name System (DNS) is a server that contains information on hosts found throughout the network, including IP address hostnames and routing information and DNS uses them primarily to translate between IP addresses and domain names. This conversion of names, such as `www.cisco.com`, to IP addresses simplifies accessing internet-based applications. The DNS directory service consists of:

- DNS data
- DNS servers
- Internet protocols for fetching data from the servers

DNS resides in the Network Registrar component of BPR.

