



CHAPTER 2

Broadband Access Center Architecture

This chapter describes the system architecture implemented in this Cisco Broadband Access Center (BAC) release.

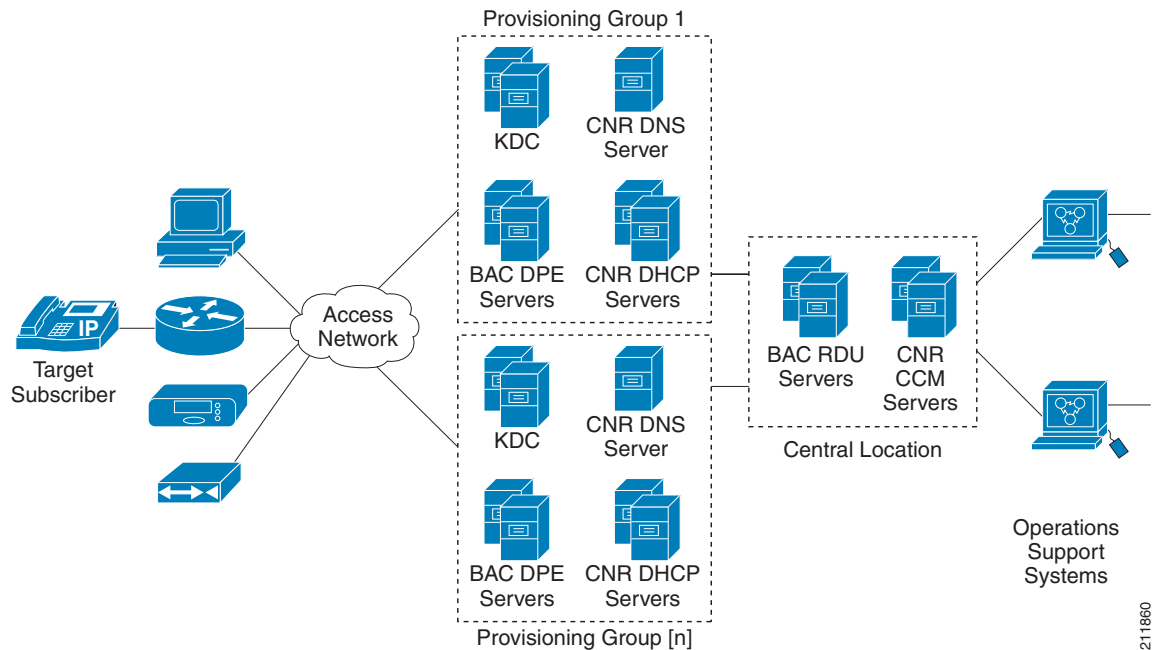
This chapter describes:

- [Deployment, page 2-1](#)
- [Architecture, page 2-2](#)
- [Logging Events, page 2-16](#)

Deployment

Figure 2-1 represents a typical, fully redundant deployment in a BAC network.

Figure 2-1 Deployment Using BAC



211860

Architecture

This section describes the basic BAC architecture, such as:

- Regional Distribution Unit (RDU) that provides:
 - The authoritative data store of the BAC system.
 - Support for processing application programming interface (API) requests.
 - Monitoring of the system's overall status and health.

See [Regional Distribution Unit, page 2-3](#), for additional information.

- Device Provisioning Engines (DPEs) that provide:
 - Interface with customer premises equipment (CPE).
 - Configuration cache.
 - Autonomous operation from the RDU and other DPEs.
 - PacketCable provisioning services.
 - IOS-like command-line interface (CLI) for configuration.

See [Device Provisioning Engines, page 2-5](#), for additional information.

- BAC API that provides total client control over system capabilities.
- Cisco Network Registrar servers that provide:
 - Dynamic Host Configuration Protocol (DHCP).
 - Domain Name System (DNS).

See [Network Registrar, page 2-11](#), for additional information.

- Provisioning Groups that provide:
 - Logical grouping of Network Registrar servers and DPEs in a redundant cluster.
 - Redundancy and scalability.

See [Provisioning Groups, page 2-14](#), for additional information.

- A Kerberos server that authenticates PacketCable Multimedia Terminal Adapters (MTAs). See [Key Distribution Center, page 2-12](#), for additional information.

- The BAC process watchdog that provides:
 - Administrative monitoring of all critical BAC processes.
 - Automated process-restart capability.
 - Ability to start and stop BAC component processes.

See [BAC Process Watchdog, page 9-1](#), for additional information.

- An SNMP agent that provides:
 - Third-party management systems.
 - SNMP version v2.
 - SNMP Notification.See [SNMP Agent, page 9-4](#), for additional information.
- An administrator user interface that supports:
 - Adding, deleting, modifying and searching for devices.
 - Configuring of global defaults and defining of custom properties.See [Administrator User Interface, page 2-13](#), for additional information.

Regional Distribution Unit

The RDU is the primary server in the BAC provisioning system. You must install the RDU on a server running the Solaris operating system.

The functions of the RDU include:

- Managing device configuration generation
- Generating configurations for devices and distributing them to DPEs for caching
- Synchronizing with DPEs to keep device configurations up to date
- Processing API requests for all BAC functions
- Managing the BAC system

The RDU supports the addition of new technologies and services through an extensible architecture.

Currently, BAC supports one RDU per installation. To provide failover support, we recommend using clustering software from Veritas or Sun. We also recommend using RAID (Redundant Array of Independent Disks) shared storage in such a setup.

The following sections describe these RDU concepts:

- [Generating Device Configurations, page 2-3](#)
- [Service-Level Selection, page 2-4](#)

Generating Device Configurations

When a device boots, it requests a configuration from BAC and it is this configuration that determines the level of service for the device. During this process, the DHCP server requests the RDU to build a configuration for the device. The RDU generates a configuration and forwards it to all the DPEs that service the provisioning group that the device belongs to. The DPEs can now provide the device with its configuration without going to the RDU.

Device configurations can include customer-required provisioning information such as:

- DHCP IP address selection
- Bandwidth
- Data rates
- Flow control

- Communication speeds
- Level of service

A configuration can contain DHCP configuration and TFTP files for any device. When you install and boot an unprovisioned device, it is assigned a default technology-specific configuration. You can change the default configuration for each technology that BAC supports.

The RDU regenerates the configuration for a device when:

- Certain provisioning API calls, such as changing the device Class of Service, are made.
- Validation for a configuration fails. This occurs, for example, when certain parameters of a DHCP request from a device change from initial request parameters.
- A DPE is repopulating its cache.

Every time the RDU regenerates a configuration for a device, the updated configuration is forwarded to the appropriate DPEs.

Service-Level Selection

The extension point for service-level selection determines the DHCP Criteria and the Class of Service that the RDU is to use when generating a configuration for a device. The RDU stores this information for each device in its database.

The DHCP Criteria and the Class of Service that the RDU uses to generate a configuration for a device is based on the type of access granted to the device. Device access is of three types:

- **Default**—For devices granted default access, BAC uses the default Class of Service and DHCP Criteria assigned for the device type.
- **Promiscuous**—For devices granted promiscuous access, BAC obtains the Class of Service and DHCP Criteria from the relay agent that the device is behind.
- **Registered**—For devices granted registered access, BAC uses the Class of Service and the DHCP Criteria registered for the device in the RDU database.

There should always be one default extension per device type.

You can enter service-level selection extension points for specific technologies using the default pages at **Configuration > Defaults**. For additional information, see [Configuring Defaults, page 13-6](#). By default, these properties are populated with zero or with one of the built-in extensions.



Caution

Do not modify these extensions unless you are installing your own custom extensions.

Although a device may have been registered as having to receive one set of DHCP Criteria and Class of Service, a second set may actually be selected. The configuration generation extension looks for the selected DHCP Criteria and Class of Service and uses them.

The service-level selection extension selects a second Class of Service and DHCP Criteria based on certain rules that you specify for a device. For example, you may specify that a device must boot in a particular provisioning group for the device to be assigned a specific Class of Service and DHCP Criteria.

The extension returns information on why a specific set of DHCP Criteria and Class of Service is selected to provision a device. You can view these reasons from the administrator user interface on the View Device Details page.

Table 2-1 describes these reasons and the type of access granted in that case.

Table 2-1 Reasons for Device Access as Determined by Service-Level Selection Extension

Reason Code	Description	Type of Device Access Granted		
		Default	Promiscuous	Registered
NOT_BEHIND_REQUIRED_DEVICE	The device is not behind its required relay agent.	✓		
NOT_IN_REQUIRED_PROV_GROUP	The device is not in its required provisioning group.	✓		
NOT_REGISTERED	The device is not registered.	✓		
PROMISCUOUS_ACCESS_ENABLED	Promiscuous access is enabled for the relay agent.		✓	
REGISTERED	The device is registered.			✓
RELAY_NOT_IN_REQUIRED_PROV_GROUP	The relay agent is not in the required provisioning group.	✓		
RELAY_NOT_REGISTERED	The relay agent is not registered.	✓		

Note Most of these reasons indicate violations of requirements for granting registered or promiscuous access, resulting in default access being granted.

Device Provisioning Engines

The Device Provisioning Engine (DPE) communicates with CPE to perform provisioning and management functions.

The RDU generates DHCP instructions and device configuration files, and distributes them to the relevant DPE servers. The DPE caches these DHCP instructions and device configuration files. The DHCP instructions are then used during interactions with the Network Registrar extensions, and configuration files are delivered to the device via the TFTP service.

BAC supports multiple DPEs. You can use multiple DPEs to ensure redundancy and scalability.

The DPE handles all configuration requests, including providing configuration files for devices. It is integrated with the 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 these activities:

- Synchronizes with the RDU to retrieve the latest configurations for caching.
- Generates last-step device configuration (for instance, DOCSIS timestamps).
- Provides the DHCP server with instructions controlling the DHCP message exchange.
- Delivers configuration files via TFTP.
- Integrates with Network Registrar.
- Provisions voice-technology services.

You must install the DPE on a server that runs the Solaris operating system. Configure and manage the DPE from the CLI, which you can access locally or remotely via Telnet. For specific information on the CLI commands that a DPE supports, refer to the *Cisco Broadband Access Center DPE CLI Reference 4.0*.

**Note**

During installation, you must configure each DPE for the:

- Name of the provisioning group to which the DPE belongs. This name determines the logical group of devices that the DPE services.
 - IP address and port number of the RDU.
-

For important information related to DPEs, see:

- [DPE Licensing, page 2-6](#)
- [TACACS+ and DPE Authentication, page 2-7](#)
- [DPE-RDU Synchronization, page 2-7](#)
- [TFTP Server, page 2-9](#)
- [ToD Server, page 2-9](#)

Also, familiarize yourself with the information described in [Provisioning Concepts, page 2-14](#).

DPE Licensing

Licensing controls the number of DPEs (nodes) that you can use. If you attempt to install more DPEs than you are licensed to use, those new DPEs will not be able to register with the RDU, and will be rejected. Existing licensed DPEs remain online.

**Note**

For licensing purposes, a registered DPE is considered to be one node.

When you add a license or extend an evaluation license or when an evaluation license has expired, the changes take effect immediately.

When you delete a registered DPE from the RDU database, a license is freed. Because the DPEs automatically register with the RDU, you must take the DPE offline if the intention is to free up the license. Then, delete the DPE from the RDU database via the administrator user interface or via the API.

Deleted DPEs are removed from all the provisioning groups that they belong to, and all Network Registrar extensions are notified that the DPE is no longer available. Consequently, when a previously deleted DPE is registered again, it is considered to be licensed again and remains so until it is deleted from the RDU again or its license expires.

DPEs that are not licensed through the RDU do not appear in the administrator user interface. You can determine the license state only by examining the DPE and RDU log files (*dpe.log* and *rdu.log*).

**Note**

The functions enabled via a specific license continue to operate even when the corresponding license is deleted from the system.

For detailed information on licensing, see [Managing Licenses, page 13-22](#).

TACACS+ and DPE Authentication

TACACS+ is a TCP-based protocol that supports centralized access for large numbers of network devices and user authentication for the DPE CLI.

Through TACACS+, a DPE can support many users, with each username and login and enable password configured at the TACACS+ server. TACACS+ is used to implement the TACACS+ client/server protocol (ASCII login only).

TACACS+ Privilege Levels

The TACACS+ server uses the TACACS+ protocol to authenticate any user logging in to a DPE. The TACACS+ client specifies a certain service level that is configured for the user.

Table 2-2 identifies the two service levels used to authorize DPE user access.

Table 2-2 TACACS+ Service Levels

Mode	Description
Login	User-level commands at <i>router></i> prompt.
Enable	Enable-level commands at <i>router#</i> prompt.

TACACS+ Client Settings

TACACS+ uses a number of properties that are configured from the CLI. For information on commands related to TACACS+, refer to the *Cisco Broadband Access Center DPE CLI Reference 4.0*.

When TACACS+ is enabled, you must specify either the IP addresses of all TACACS+ servers or their fully qualified domain names (FQDNs) with nondefault values.

You can also specify these settings using their default values, if applicable:

- The shared secret key for each TACACS+ server. Using this key, you can encrypt data between the DPE and the TACACS+ server. If you choose to omit the shared secret for any specific TACACS+ server, TACACS+ message encryption is not used.
- The TACACS+ server timeout. Using this value, you can specify the maximum length of time that the TACACS+ client waits for a TACACS+ server to reply to protocol requests.
- The TACACS+ server number of retries. Using this value, you can specify the number of times that the TACACS+ client attempts a valid protocol exchange with a TACACS+ server.

DPE-RDU Synchronization

The DPE-RDU synchronization is a process of automatically updating the DPE cache to be consistent with the RDU. The DPE cache comprises the configuration cache, with configurations for devices, and the file cache, with files required for devices.

Under normal conditions, the RDU generates events containing configuration updates and sends them to all relevant DPEs to keep them up to date. Synchronization is needed if the DPE is missing some events due to connection loss. Such loss could be because of a network issue, the DPE server going down for administrative purposes, or a failure.

Synchronization also covers the special case when the RDU database is restored from backup. In this case, the DPE cache database must be returned to an older state to be consistent with the RDU.

The RDU and DPE synchronization process is automatic and requires no administrative intervention. Throughout the synchronization process, the DPE is still fully capable of performing provisioning and management operations on the CPE.

Synchronization Process

The DPE triggers the synchronization process every time it establishes a connection with the RDU.

When the DPE first starts up, it establishes the connection to the RDU and registers with the RDU to receive updates of configuration changes. The DPE and RDU then monitor the connection using heartbeat message exchanges. When the DPE determines that it has lost its connection to the RDU, it automatically attempts to re-establish it. It continues its attempts with a backoff-retry interval until it is successful.

The RDU also detects the lost connection and stops sending events to the DPE. Because the DPE may miss the update events from the RDU when the connection is down, the DPE performs synchronization every time it establishes a connection with the RDU.

General DPE States

During the process of synchronization, the DPE is in the following states:

1. **Registering**—During the process of establishing a connection and registering with the RDU, the DPE is in the *Registering* state.
2. **Synchronizing**—The DPE requests groups of configurations that it should have from the RDU. During this process, the DPE determines which configurations in its store are inconsistent (wrong revision number), which ones are missing, and which ones to delete, and, if necessary, updates the configurations in its cache. The DPE also synchronizes deliverable files in its cache for the TFTP server. To ensure that the RDU is not overloaded with configuration requests, the DPE posts only one batch at a time to the central server.
3. **Ready**—The DPE is up to date and fully synchronized with the RDU. This state is the typical state that the DPE is in.

Table 2-3 describes some other states that the DPE may be in from time to time.

Table 2-3 Related DPE States

State	Description
Initializing	Is starting up
Shutting Down	Is in the process of stopping
Down	Does not respond to queries from Network Registrar extension points
Ready Overloaded	Is similar to <i>Ready</i> except that there is a heavy load on the system on which the DPE is running



Note

Regardless of the state that the DPE is in, it continues to service device configuration, TFTP, and ToD requests.

You can view the DPE state:

- From the administrator user interface. See [Viewing Device Provisioning Engines, page 12-21](#).
- From the DPE CLI using the **show dpe** command. Refer to the *Cisco Broadband Access Center DPE CLI Reference 4.0*.

TFTP Server

The integrated TFTP server receives requests for files, including DOCSIS configuration files, from device and nondevice entities. This server then transmits the file to the requesting entity.

The TFTP server is located in a home directory that is used for local file-system access. The local files are stored in the *BPR_DATA/dpe/tftp* directory. In this release, all deliverable TFTP files are precached at the DPE; in other words, the DPE is always up to date with all the files in the system.



Note

The TFTP service on the DPE features one instance of the service, which you can configure to suit your requirements.

By default, the TFTP server only looks in its cache for a TFTP read. However, if you run the **service tftp 1..1 allow-read-access** command from the DPE command line, the TFTP server looks in the local file system before looking in the cache. If the file exists in the local file system, it is read from there. If not, the TFTP server looks in the cache. If the file exists in the cache, the server uses it; otherwise, it returns an error.

When you can enable read access from the local file system, directory structure read requests are allowed only from the local file system.



Note

Ensure that you give unique names to all TFTP files instead of differentiating the files by using upper or lowercase. The filename casing is important because the DPE, while looking for a file in its local directory or cache, converts all filenames to lowercase.

You can specify TFTP transfers using IPv4 or IPv6, using the **service tftp 1..1 ipv4 | ipv6 enabled true** command from the DPE command line. You can also specify a block size for these transfers using the **service tftp 1..1 ipv4 | ipv6 blocksize** command. The blocksize option specifies the number of data octets and allows the client and server to negotiate a block size more applicable to the network medium. When you enable blocksize, the TFTP service uses the requested block size for the transfer if it is within the specified lower and upper limits. For detailed information, see the *Cisco Broadband Access Center DPE CLI Reference 4.0*.

The TFTP service maintains statistics for the number of TFTP packets that are processed for TFTPv4 and TFTPv6. You can view these statistics from the administrator user interface on the device details page. For more information, see [Viewing Device Details, page 12-9](#).

ToD Server

The integrated time of day (ToD) server in BAC provides high-performance UDP implementation of RFC 868.

**Note**

The ToD service on the DPE features one instance of the service, which you can configure to suit your requirements.

You can enable the ToD service to support IPv4 or IPv6, from the DPE command line, using the **service tod 1..1 enabled true** command. The ToD service is, by default, disabled on the DPE.

While configuring this protocol on the DPE, remember that the ToD service binds only to those interfaces that you have configured for provisioning. For detailed information on configuring the ToD service, see the *Cisco Broadband Access Center DPE CLI Reference 4.0*.

The ToD service maintains statistics for the number of ToD packets that are processed for ToDv4 and ToDv6. You can view these statistics from the administrator user interface on the device details page. For more information, see [Viewing Device Details, page 12-9](#).

DOCSIS Shared Secret

BAC lets you define a different DOCSIS shared secret (DSS) for each cable modem termination system (CMTS). In this way, a compromised shared secret affects only a limited number of CMTS, instead of every CMTS in the deployment.

Although the DSS can be set for each DPE, you should set it on a provisioning-group basis. Also, ensure that it matches what has been configured for the CMTS in that provisioning group.

**Caution**

Configuring multiple DSS within one provisioning group could, under some conditions, result in degraded CMTS performance. However, this factor has virtually no effect on BAC.

You can enter the shared secret as a clear text string or as an IOS-encrypted string. When entered in clear text, the DSS is encrypted to suit IOS version 12.2BC.

You can also set the DSS from the RDU using the administrator user interface or the API. In this case, the DSS is entered, stored at the RDU, and passed to all DPEs in clear text. Consequently, before a DSS entered this way is stored on the DPE, it is encrypted.

If you set the DSS directly at the DPE using the **dpe docsis shared-secret** command from the CLI, this DSS takes precedence over the one set from the RDU.

Resetting the DOCSIS Shared Secret

You can reset the DSS if the security of the DSS is compromised or to simply change the shared secret for administrative purposes.

To reset the DSS, run the **show running-config** command from the CMTS CLI, then copy and paste the DOCSIS shared secret from the configuration that appears into the DPE configuration. In this way, you can copy the configuration that you enter in a Cisco CMTS into the DPE CLI.

**Note**

To change the shared secret as described, the CMTS must be running a software version later than version 12.2BC.

To change the DSS:

-
- Step 1** Identify the provisioning group on which you need to reset the DOCSIS shared secret.
 - Step 2** Examine the list of DPEs and CMTS associated with the provisioning group.
 - Step 3** Change the primary DSS on the CMTS.
 - Step 4** Change the compromised DSS on the CMTS to the secondary DSS. This change is required to allow cable modems to continue to register until all the DOCSIS configuration files are successfully changed to use the new DSS.
 - Step 5** Determine which DPEs were affected and change the DSS on each accordingly.
 - Step 6** Confirm that the DOCSIS configuration files are using the new DSS and then remove the compromised secondary shared secret from the CMTS configuration.
-

Network Registrar

Network Registrar provides the DHCP and DNS functionality in BAC. The DHCP extension points on Network Registrar integrate BAC with Network Registrar. Using these extensions, BAC examines the content of DHCP requests to detect device type, manipulates the content according to its configuration, and delivers customized configurations for devices that it provisions.

For additional information on Network Registrar, refer to the *User Guide for Cisco Network Registrar 7.0*; *CLIFrame.html* in the */docs* directory; and *Installation Guide for Cisco Network Registrar, 7.0*.

DHCP

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

This BAC release supports DHCP for IPv6, also known as DHCPv6. DHCPv6 enables DHCP servers to deliver configuration parameters, via extensions, to IPv6 hosts. IPv6 hosts by default use stateless autoconfiguration, which enables IPv6 hosts to configure their own addresses using a local IPv6 router. DHCPv6 represents the stateful autoconfiguration option, a technique in which configuration information is provided to a host by a server.

DHCPv6 provides:

- Expanded addressing capabilities via IPv6 addresses
- Easy network management and administration using the stateful autoconfiguration protocol
- Improved support for options and extensions
- Relay agent functionality
- Assignment of multiple addresses to one interface

DHCPv4 versus DHCPv6

Much like DHCPv4, DHCPv6 uses a client-server model. The DHCP server and the DHCP client converse with a series of messages to request, offer, and lease an IP address. Unlike DHCPv4, DHCPv6 uses a combination of unicast and multicast messages for the bulk of the conversation instead of broadcast messages.

Some other differences between DHCPv4 and DHCPv6 are:

- Unlike DHCPv4, IPv6 address allocation in DHCPv6 is handled using a message option.
- Message types, such as DHCP Discover and DHCP Offer supported by DHCPv4 are removed in DHCPv6. Instead, DHCPv6 servers are located by a client Solicit message followed by a server Advertise message.
- Unlike DHCPv4 clients, DHCPv6 clients can request multiple IPv6 addresses.

DHCPv4 failover allows pairs of DHCP servers to act in such a way that one can take over if the other stops functioning. The server pairs are known as the main and backup server. Under normal circumstances, the main server performs all DHCP functions. If the main server becomes unavailable, the backup server takes over. In this way, DHCP failover prevents loss of access to the DHCP service if the main server fails.

DNS

The DNS server contains information on hosts throughout the network, such as IP address hostnames. DNS uses this information primarily to translate between IP addresses and domain names. The conversion of names such as `www.cisco.com` to IP addresses simplifies accessing Internet-based applications.

Lease Query

The lease query feature allows you to request current IP address information directly from the Network Registrar DHCP servers in a provisioning group. To find a device's IP address, the RDU sends DHCP lease query messages only to the DHCP servers in the device's provisioning group, which prevents querying all DHCP servers in the network. Among all the responses, the response from the server that last communicated with the devices is taken as the authoritative answer.

In earlier BAC versions, the lease query feature relied on the operating system to select the source interface and the source port for sending lease query requests. In this release, you can configure the RDU to use a specific interface and source port.

For detailed information on lease query support in this BAC release, see [Lease Query, page 6-19](#).

Key Distribution Center

The Key Distribution Center (KDC) authenticates PacketCable MTAs and also grants service tickets to MTAs. As such, it must check the MTA certificate, and provide its own certificates so that the MTA can authenticate the KDC. It also communicates with the DPE (the provisioning server) to validate that the MTA is provisioned on the network.

You must install the KDC on a server that runs the Solaris operating system.

The certificates used to authenticate the KDC are not shipped with BAC. You must obtain the required certificates from Cable Television Laboratories, Inc. (CableLabs), and the content of these certificates must match those that are installed in the MTA. For additional information, see [Using the PKCert.sh Tool, page 14-2](#).

**Caution**

The KDC does not function if the certificates are not installed.

The KDC also requires a license to function. Obtain a KDC license from your Cisco representative and install it in the correct directory. For details on how to install the license, see [KDC Licenses, page 7-9](#).

The KDC has several default properties that are populated during a BAC installation into the *BPR_HOME/kdc/solaris/kdc.ini* properties file. You can edit this file to change values as operational requirements dictate. For detailed information, see [Default KDC Properties, page 7-7](#).

The KDC also supports the management of multiple realms. For details on configuring additional realms, see [Multiple Realm Support, page 7-10](#).

BAC Process Watchdog

The BAC process watchdog is an administrative agent that monitors the runtime health of all BAC processes. This watchdog process ensures that if a process stops unexpectedly, it is automatically restarted. One instance of the BAC process watchdog runs on every system which runs BAC components.

You can use the BAC process watchdog as a command-line tool to start, stop, restart, and determine the status of any monitored processes.

See [BAC Process Watchdog, page 9-1](#), for additional information on how to manage the monitored processes.

SNMP Agent

BAC provides basic SNMP v2-based monitoring of the RDU and DPE servers. The BAC SNMP agents support SNMP informs and traps, collectively called notifications.

You can configure the SNMP agent:

- On the RDU, using the SNMP configuration command-line tool (see [Monitoring Servers Using SNMP, page 10-9](#)) or via the API.
- On the DPE, using the **snmp-server** CLI commands. Refer to the *Cisco Broadband Access Center DPE CLI Reference 4.0*.

Administrator User Interface

The BAC administrator user interface is a web-based application for central management of the BAC system. You can use this system to:

- Configure global defaults
- Define custom properties
- Add, modify, and delete Class of Service

- Add, modify, and delete DHCP Criteria
- Add, modify, and delete devices
- Group devices
- View server status and server logs
- Manage users

Refer to these chapters for specific instructions on how to use this interface:

- [Chapter 11, “Understanding the Administrator User Interface,”](#) describes how to access and configure the BAC administrator user interface.
- [Chapter 12, “Using the Administrator User Interface,”](#) provides instructions for performing administrative activities involving the monitoring of various BAC components.
- [Chapter 13, “Configuring Broadband Access Center,”](#) describes tasks that you perform to configure BAC.

Provisioning Concepts

This section describes those concepts that are key to provisioning and include:

- [Provisioning Groups, page 2-14](#)
- [Static versus Dynamic Provisioning, page 2-15](#)
- [Provisioning Group Capabilities, page 2-15](#)

Provisioning Groups

A provisioning group is designed to be a logical (typically geographic) grouping of servers that usually consists of one or more DPEs and a failover pair of DHCP servers. Each DPE in a given provisioning group caches identical sets of configurations from the RDU, thus enabling redundancy and load balancing. As the number of devices grows, you can add additional provisioning groups to the deployment.



Note

The servers for a provisioning group are not required to reside at a regional location. They can just as easily be deployed in the central network operations center.

Provisioning groups enhance the scalability of the BAC deployment by making each provisioning group responsible for only a subset of devices. This partitioning of devices can be along regional groupings or any other policy that the service provider defines.

To scale a deployment, the service provider can:

- Upgrade existing DPE server hardware
- Add DPE servers to a provisioning group
- Add provisioning groups

To support redundancy and load sharing, each provisioning group can support any number of DPEs. As the requests come in from the DHCP servers, they are distributed between the DPEs in the provisioning group and an affinity is established between the devices and a specific DPE. This affinity is retained as long as the DPE state within the provisioning group remains stable.

Static versus Dynamic Provisioning

BAC provisions devices in the network using device configurations, which is provisioning data for a specific device based on its technology type. You can provision devices using BAC in two ways: static provisioning and dynamic provisioning.

During static provisioning, you enter static configuration files into the BAC system. These configuration files are then delivered via TFTP to the specific device to generate its configuration. BAC treats static configuration files like any other binary file.

During dynamic provisioning, you use templates, which are text files containing DOCSIS, PacketCable, or CableHome options and values that, when used with a particular Class of Service, provide dynamic file generation. A dynamic configuration file provides more flexibility and security during the provisioning process.

[Table 2-4](#) describes the impact of static and dynamic provisioning using the corresponding files.

Table 2-4 *Static Provisioning versus Dynamic Provisioning*

Static Provisioning Using Static Files	Dynamic Provisioning Using Template Files
Used when fewer service offerings are available	Used when many service offerings are available
Offers limited flexibility	Offers more flexibility, especially when devices require unique configurations
Is relatively less secure	Is more secure
Offers higher performance	Offers slower performance, because every time you update a template assigned to a device, configurations for all devices associated with that template are updated.
Is simpler to use	Is more complex

Provisioning Group Capabilities

To provision a subset of devices in a deployment, provisioning groups must be capable of as well as enabled to provision those devices. For example, a provisioning group cannot provision a PacketCable MTA in Secure mode if its DPEs are not configured to support this functionality.

In previous BAC releases, each DPE in a provisioning group registered what it was capable of supporting with the RDU at startup. This information was combined with that of other DPEs in the provisioning group to determine the device types that the group could support. The servers registered their low-level capabilities and if those capabilities were enabled or disabled. After server registration, the provisioning group was automatically enabled to support the device types it was capable of supporting. However, in this BAC release, you must enable device support manually:

- From the administrator user interface on the Provisioning Group Details Page. See [Viewing Provisioning Groups, page 12-27](#).
- From the API using the *ProvGroupCapabilitiesKeys* constants. For details, refer to the API Javadoc.

Logging Events

Logging of events is performed at the RDU and the DPE, and in some unique situations, DPE events are additionally logged at the RDU to give them higher visibility. Log files are stored in their own log directories and can be examined by using any text processor. You can compress the files for easier e-mailing to the Cisco Technical Assistance Center or system integrators for troubleshooting and fault resolution. You can also access the RDU and the DPE logs from the administrator user interface.

For detailed information on log levels and structures, and how log files are numbered and rotated, see [Log Levels and Structures, page 10-1](#).