Chapter 9
Configuring OSPF

This chapter provides details on how to configure OSPF on HP routing switches using the CLI and Web management interface.

A detailed summary of all CLI commands noting syntax and possible values can be found in **Appendix B**.

Overview of OSPF

OSPF is a link-state routing protocol. As a link-state protocol it uses link-state advertisements (LSA) to update neighboring routers of its interfaces and information on those interfaces. These LSAs are flooded to all neighboring routers to update them on each of their neighboring routers. Each router maintains an identical database that describes its area topology to help a router determine the shortest path between it and any neighboring router.

OSPF is built upon a hierarchy of network components. The highest level of the hierarchy is the **Autonomous System (AS)**. An autonomous system is defined as a number of networks, all of which share the same routing and administration characteristics.

An AS can be divided into multiple **areas** as seen in **Figure 9.1**. Each area represents a collection of contiguous networks and hosts. Areas are defined to limit the area to which link-state advertisements are broadcast, thereby limiting the amount of flooding that occurs within the network. An area is represented by either an IP address or a number.

The user can further limit the broadcast area of flooding by defining an **area range**. The area range allows the user to assign a representative value to a range of IP addresses. This representative value becomes the address that is advertised versus all of the addresses it represents. Up to four ranges can be assigned per area.

A router can be a member of multiple areas. These routers with membership in multiple areas are known as **Area Border Routers (ABR)**. Each ABR maintains a separate topological database for each area, which contains all of the LSA databases for each router within a given area. Those routers within the same area will have identical topological databases. The ABR is responsible for forwarding routing information or changes between its border areas.

An **Autonomous System Border Router (ASBR)** is a router that is running multiple protocols and serves as a gateway to routers outside an area and those operating with different protocols. The ASBR is able to import and translate different protocol routes into OSPF via a process known as **redistribution**. For more details on redistribution and configuration examples, please refer to the **Enable Redistribution** section of this chapter.

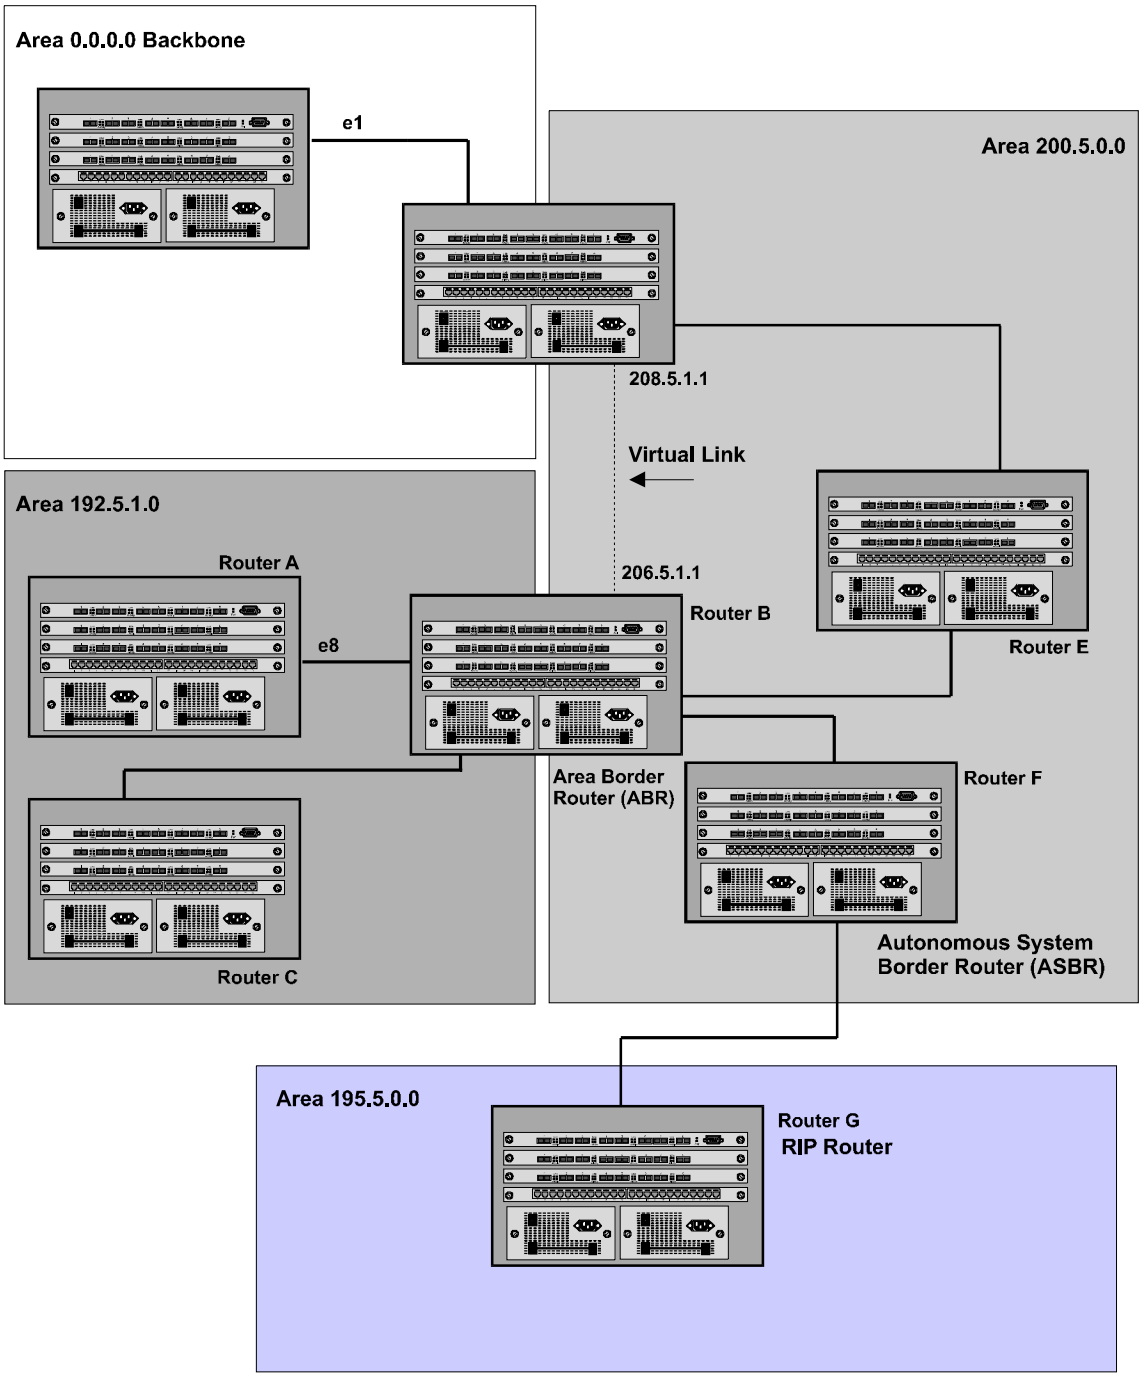

Figure 9.1 OSPF operating in a network

Designated Routers in Multi-Access Networks

In a network that has multiple routers attached, OSPF will elect one router to serve as the designated router (DR) and another router on the segment to act as the backup designated router (BDR). This arrangement will minimize the amount of repetitive information that is forwarded on the network by forwarding all messages to the designated and backup designated routers responsible for forwarding the updates throughout the network.

Designated Router Election

In a network with no designated router and no backup designated router, the neighboring router with the highest priority will be elected as the DR, and the router with the next largest priority will be elected as the BDR, as seen in **Figure 9.2**.

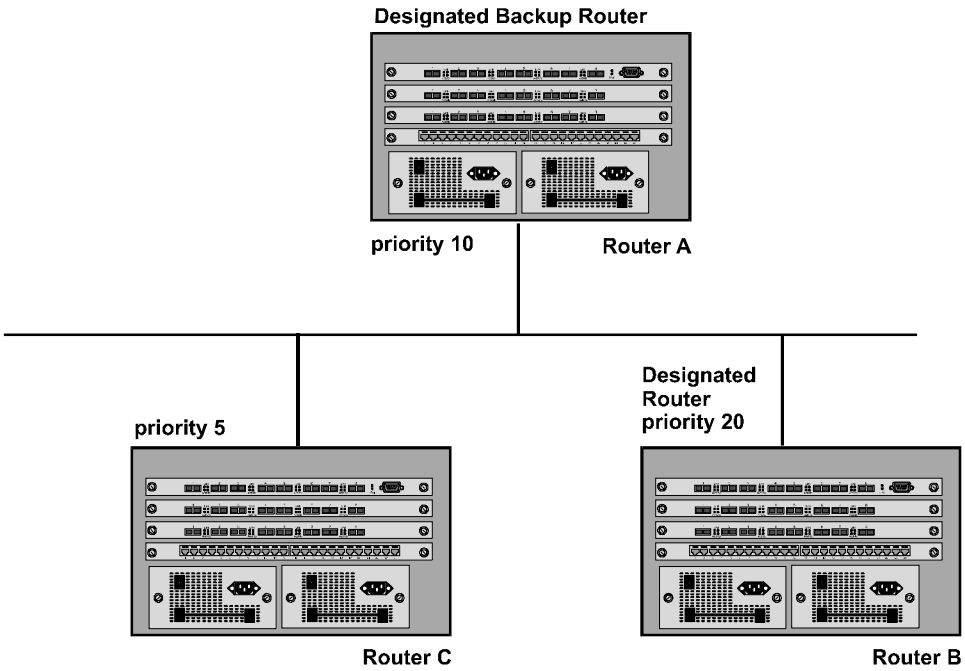

Figure 9.2 Designated and backup router election

Should the DR go off-line, then the BDR will automatically become the DR, and the router with the next highest priority will become the BDR as seen in **Figure 9.3**.

NOTE: Priority is a configurable option at the Interface Level of the CLI. The user can use this parameter to help bias one router as the DR.

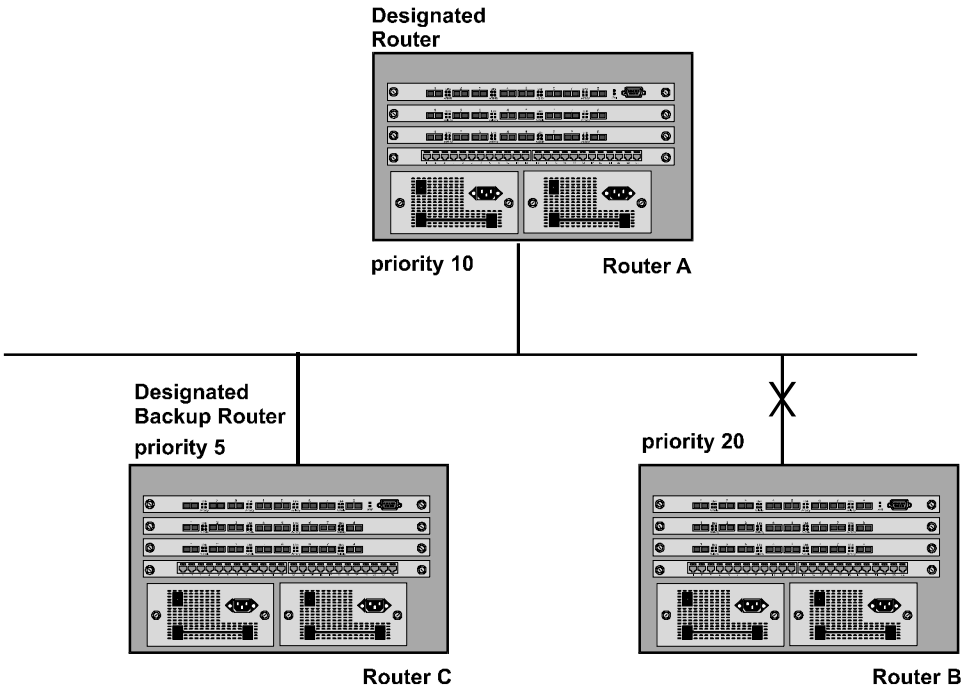

Figure 9.3 Backup designated router becomes designated router

Should two neighbors share the same priority, then the router with the highest router ID will be designated as the DR, and the router with the next highest router ID will be the BDR.

When multiple routers on the same network are declaring themselves as DRs, then both priority and router ID will be used to select the designated and backup designated routers.

When only one router on the network claims the DR role despite neighboring routers with higher priorities or router IDs, it will remain as the DR. This is also true for BDRs.

The election process is performed when one of the following events occurs:

- an interface is in a waiting state and the wait time expires
- an interface is in a waiting state and a hello packet is received that addresses the backup designated router
- a change in the neighbor state occurs, such as:
  - a neighbor state transitions from 2 or higher
  - communication to a neighbor is lost
  - a neighbor declares itself to be the DR or BDR for the first time

OSPF RFC 1583 and 2178 Compliant

HP routing switches are configured, by default, to be compliant with RFC 1583 OSPF V2 specification. HP routing switches can also be configured to operate with the latest OPSF standard, RFC 2178.

NOTE: For details on how to configure the system to operate with the RFC 2178, please refer to the section on **Configuring OSPF**.

Dynamic OSPF Activation and Configuration

OSPF is automatically activated without a system reset.

Additionally, users can configure and save the following OSPF changes without requiring a system reset:

- all OSPF interface-related parameters (e.g. area, hello timer, router dead time cost, priority, re-transmission time, transit delay)
- all area parameters
- all area range parameters
- all virtual-link parameters
- all global parameters
- creation and deletion of an area, interface or virtual link

In addition, the following changes can be done without a system reset by first disabling and then re-enabling OSPF operation:

- changes to address ranges
- changes to global values for redistribution
- addition of new virtual links

NOTE: For more details on OSPF CLI commands, please refer to the examples highlighted in the next section and **Appendix B**.

Configuring OSPF

To begin using OSPF on the routing switch, the user must follow the steps outlined below:

1. Enable the feature on the router.
2. Assign the areas to which the router will be attached.
3. Assign individual interfaces to the OSPF areas.
4. Enable redistribution, if desired.
5. Define redistribution filters if redistribution is enabled.
6. Modify default global and port parameters as required.
7. Modify OSPF standard compliance, if desired.
8. Save the changes to flash.

Configuration Rules

- If a router is to operate as an ASBR, that capability must first be enabled at the system level via the OSPF configuration sheet.
- Redistribution must be enabled on routers configured to operate as ASBRs.
- All router ports must be assigned to one of the defined areas on an OSPF router. When a port is assigned to an area, all corresponding sub-nets on that port are automatically included in the assignment.

OSPF Parameters

The user can modify or set the following global and interface OSPF parameters:

Global Parameters

- Modify OSPF standard compliance setting
- Assign an area
- Define an area range
- Define the area virtual link
- Set global default metric for OSPF
- Define redistribution metric type
- Define deny redistribution
- Define permit redistribution
- Enable redistribution
- Modify database overflow interval
- Modify external LSDB limit
- Modify OSPF Traps generated

Interface Parameters

- Assign interfaces to an area
- Define the authentication key for the interface
- Modify the cost for a link
- Modify the dead interval
- Modify MD5 authentication key parameters
- Modify the priority of the interface
- Modify the retransmit interval for the interface
- Modify the transit delay of the interface

NOTE: When using CLI, the user sets global level parameters are set at the OSPF CONFIG Level of the CLI. The user reaches that level by entering ***router ospf…*** at the global CONFIG Level. Interface parameters for OSPF are set at the interface CONFIG Level using the CLI command, ***ip ospf….***

NOTE: When using the Web management interface, the user sets OSPF global parameters via the OSPF configuration sheet. All other parameters are accessed via links accessed from the OSPF configuration sheet.

Enabling OSPF on the Router

When the user enables OSPF on the router, it is automatically activated on the system.

To enable OSPF on the router, the user would enter the following:

USING THE CLI

```
HP9300(config)# router ospf
```

This command will launch the user into the OSPF router level where the user can assign areas and modify OSPF global parameters.

USING THE WEB MANAGEMENT INTERFACE

OSPF is enabled on the System configuration sheet.

Assigning OSPF Areas

Once OSPF is enabled on the system, the user can assign areas. The user assigns an IP address or number as the **area ID** for each area. The area ID is representative of all IP addresses (sub-nets) on a router port. Each port on a router can support one area.

An area is defined as either **normal** or a **stub**. When an area is defined as **normal,** all external host-routes will be advertised into the area. When an area is defined as **stub**, external routes will not be advertised into the area because only one external route is available for the port.

EXAMPLE: To set up the OSPF areas as seen in **Figure 9.1**, the user would enter the following:

USING THE CLI

```
HP9300(config-ospf-router)# area 192.5.1.0 normal
HP9300(config-ospf-router)# area 200.5.0.0 normal
HP9300(config-ospf-router)# area 195.5.0.0 normal
HP9300(config-ospf-router)# area 0.0.0.0 normal
```

NOTE: An area ID can be defined by either an IP address or a number between 0 and 2,147,483,647.

NOTE: The user can assign one area per port, so if the module has 24 ports, 24 areas are supported on the module.

USING THE WEB MANAGEMENT INTERFACE

1. Select OSPF area from the OSPF configuration sheet. The panel seen in **Figure 9.4** will appear. If areas are already defined for the router, a summary panel will appear and the user will need to select the add area link to reach the OSPF area configuration panel.
2. Enter the IP address for the area in the **area ID** field.

   NOTE: A backbone area of 0.0.0.0 should always be defined for OSPF.

3. Select either **stub** or **normal** to define the area type.
4. Assign a **stub cost** to assign a priority to the area if stub was selected in the previous step.
5. Select the **add** button after defining each area.
6. Repeat steps 2-5 for each area to be defined. In this example (**Figure 9.1**), the user would also define the areas 0.0.0.0, 200.5.0.0 and 195.5.0.0.

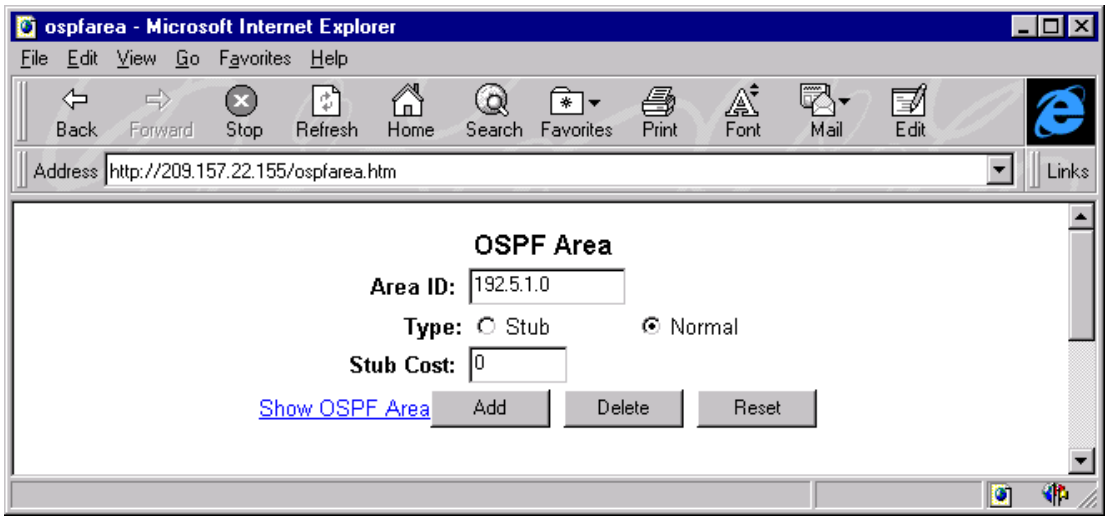

Figure 9.4 Assigning OSPF area 192.5.1.0

Assigning an Area Range (optional)

The user can also assign a **range** for an area, but it is not required. Ranges allow a specific IP address and mask to represent a range of IP addresses within an area, so that only that reference range address is advertised to the network, versus all of the addresses within that range. Up to four range addresses can be assigned per area.

USING THE CLI

EXAMPLE: To define an area range that would be a reference address for sub-nets on 193.45.5.1 and 193.45.6.2 the user would do the following:

```
HP9300(config)# router ospf
HP9300(config-ospf-router)# area 192.45.5.1 range 193.45.0.0 255.255.0.0
HP9300(config-ospf-router)# area 193.45.6.2 range 193.45.0.0 255.255.0.0
```

USING THE WEB MANAGEMENT INTERFACE

1. Select OSPF area range from the OSPF configuration sheet. The panel seen in **Figure 9.5** will appear.
2. Enter the **area ID** to be represented by the area range.
3. Enter the IP address for the range in the **network address** field.
4. Enter the IP mask for the range in the **mask** field.
5. Select the **add** button.

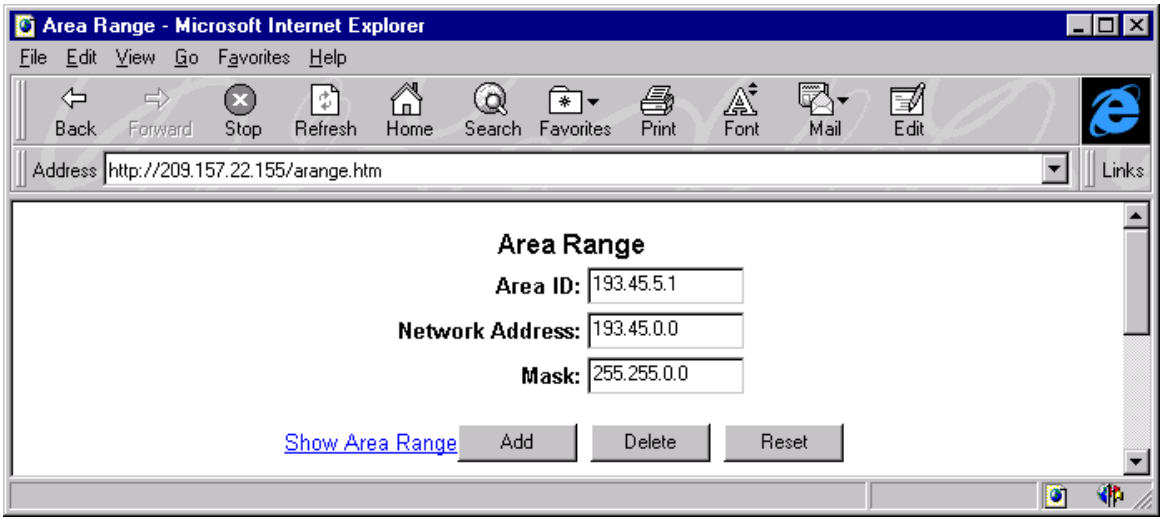

Figure 9.5 Mapping an area to a representative area range

Assigning Interfaces to an Area

Once OSPF areas are defined, interfaces are assigned to the various areas. All router ports must be assigned to one of the defined areas on an OSPF router. When a port is assigned to an area, all corresponding sub-nets on that port are automatically included in the assignment.

To assign interface 8 of Router A to area 192.5.1.0 and then save the changes, the user would enter the following:

USING CLI

To assign interface 8 (slot 2) of Router A to area 192.5.1.0 and then save the changes, the user would enter the following:

```
RouterA(config-ospf-router)# interface e 2/8
RouterA (config-if-2/8)# ip ospf area 192.5.1.0
RouterA (config-if-2/8)# write memory
```

USING WEB MANAGEMENT INTERFACE

All router ports must be assigned to one of the defined areas on an OSPF router. When a port is assigned to an area, all corresponding sub-nets on that port are automatically included in the assignment.

To assign an interface to an area, the user would do the following:

1. Select OSPF interface from the OSPF configuration sheet. The OSPF interface entry panel as seen in **Figure 9.6** will appear.

   NOTE: If OSPF interfaces are already assigned to any OSPF areas on the router, a summary panel of OSPF interfaces assigned and their areas will appear. The user will then need to select the add OSPF interface link.

2. Select the **slot/port** to be assigned to the area from the pull down menu.
3. Select the IP address or router ID of the area to which the interface is to be assigned, from the **area ID** pull down menu.

   NOTE: An area must be defined before assignment of interfaces is supported.

4. Select the **enable** option of the OSPF mode parameter to enable OSPF on the interface.

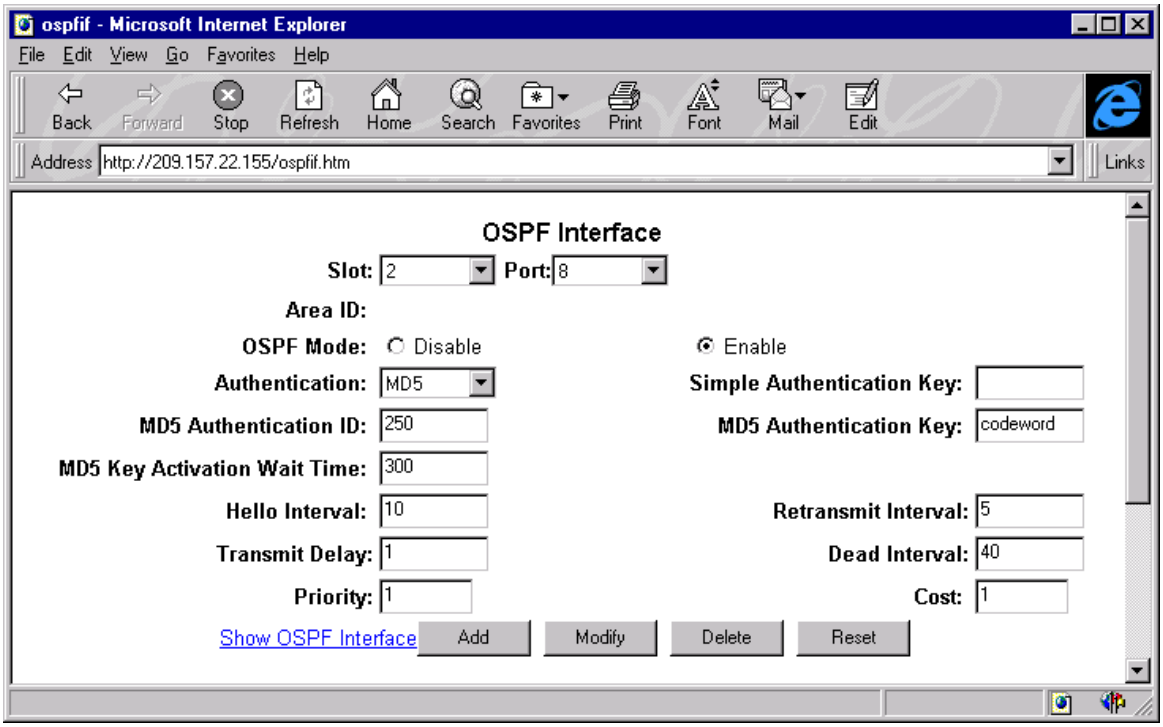

Figure 9.6 Assigning an interface to an area

Modify Port Defaults

OSPF comes standard with a number of port parameters that can be configured. For simplicity each of these parameters comes pre-configured with default values. No change to these default values is required except as needed for specific network configurations.

USING THE CLI

Port default values can be modified using the CLI command ***ip ospf [cost | dead-interval | hello-interval | priority | retransmit-interval | transit-delay]...***at the interface level of the CLI.

For a complete description of these parameters, please refer to the summary of OSPF port parameters in the next section.

USING THE WEB MANAGEMENT INTERFACE

To modify OSPF port parameters, the user would do the following:

1. Select OSPF from the main menu.
2. Select the OSPF interface link from the OSPF configuration sheet. The panel seen in **Figure 9.6** will appear.
3. Select the **slot/port** to be modified from the pull down menus.
4. Select the **authentication** method for the interface from the pull down menu. Options are none, simple or MD5.

NOTE: If MD5 is selected as the authentication method, enter a value for the MD5 authentication ID, key and key activation time in the associated fields. If simple is chosen, enter an authentication key. If no authentication (password) is chosen as the authentication method, there is no need to specify anything in the simple and MD5 fields.

---

5. Modify the default values of the following interface parameters as needed: **hello interval, transit delay, priority, retransmit interval** and **cost**.
6. Select the **add** button to save the changes.

Overview of OSPF Physical and Virtual Interface Parameters

Area: Assigns an interface to a specific area. The user can assign either an IP address or number to represent an OSPF area ID. If a number is assigned, it can be any value between 0 and 2,147,483,647.

Authentication-key:OSPF supports three methods of authentication for each interface-none, simple password and MD5. Only one method of authentication can be active on an interface at a time. The default authentication value is none, meaning no authentication is performed.

The simple password method of authentication requires that the user configure an alphanumeric password on an interface. The password can be up to eight characters.

The simple password setting takes effect immediately. Any OSPF packet transmitted on the interface will contain this password. Any OSPF packet received on the interface will be checked for this password. If the password is not present, then the packet will be dropped. The password can be up to sixteen characters.

The MD5 method of authentication requires the user to configure a key ID and a MD5 Key. The key ID is a number between 1 and 255 and identifies the MD5 key that is being used.

Cost: Indicates the overhead required to send a packet across an interface. This value can be modified to differentiate between 100Mbp and 1000Mbp (1 Gbps) links. The default cost, as defined, is calculated by dividing 100 million by the bandwidth. For 10Mbp links, the cost would be 10. The cost for both 100Mbps and 1000Mbps links is 1, as the speed of 1000Mbps was not in use at the time the OSPF cost formula was devised.

Dead-interval: Indicates the number of seconds that a neighbor router will wait for a hello packet from the current router before declaring the router down. Possible values are 1-65,535 seconds. The default is 40 seconds.

Hello-interval: Represents the length of time between the transmission of hello packets. Possible values are 1-65,535 seconds. The default is 10 seconds.

MD5-authentication: A method of authentication that requires the user to configure a key ID and a MD5 key. The key ID is a number between 1 and 255 and identifies the MD5 key that is being used. The MD5 key consists of up to 16 alphanumeric characters. The MD5 is encrypted and included in each OSPF packet transmitted.

Priority: Allows the user to modify the priority of an OSPF router. This value is used in selecting designated and backup designated routers.Possible values are 1 to 255 The default is 1 second.

Retransmit-interval: The time between retransmits of link-state advertisements to router adjacencies for this interface. Possible values are 0 to 3600 seconds. The default is 5 seconds.

Transit-delay: The time it takes to transmit Link State Update packets on this interface. Possible values are 0 to 3600 seconds. The default is 1 second.

Assign Virtual Links

All area border routers must have either a direct or indirect link to the OSPF area backbone (0.0.0.0). If an **area border router (ABR)** does not have a physical link to the area backbone, it can configure a **virtual link** to another router within the same area, which has a physical connection to the area backbone.

The path for a virtual link is via an **area** that both the **neighbor area border router** (router with a physical backbone connection), and the **ABR** requiring a logical connection to the backbone, share.

Two fields must be defined for all virtual links—**transit area ID** and **neighbor router.**

The **transit area ID** represents the shared area of the two area border routers and serves as the connection point between the two routers. This number should match the area ID value.

The **neighbor router** field is the **router ID (IP address)** of the router that is physically connected to the backbone, when assigned from the router interface requiring a logical connection.

When assigning the parameters from the router with the physical connection, the router ID is the IP address of the router requiring a logical connection to the backbone.

NOTE: When establishing an **area virtual link,** it must be configured on both of the routers, i.e. both ends of the virtual link.

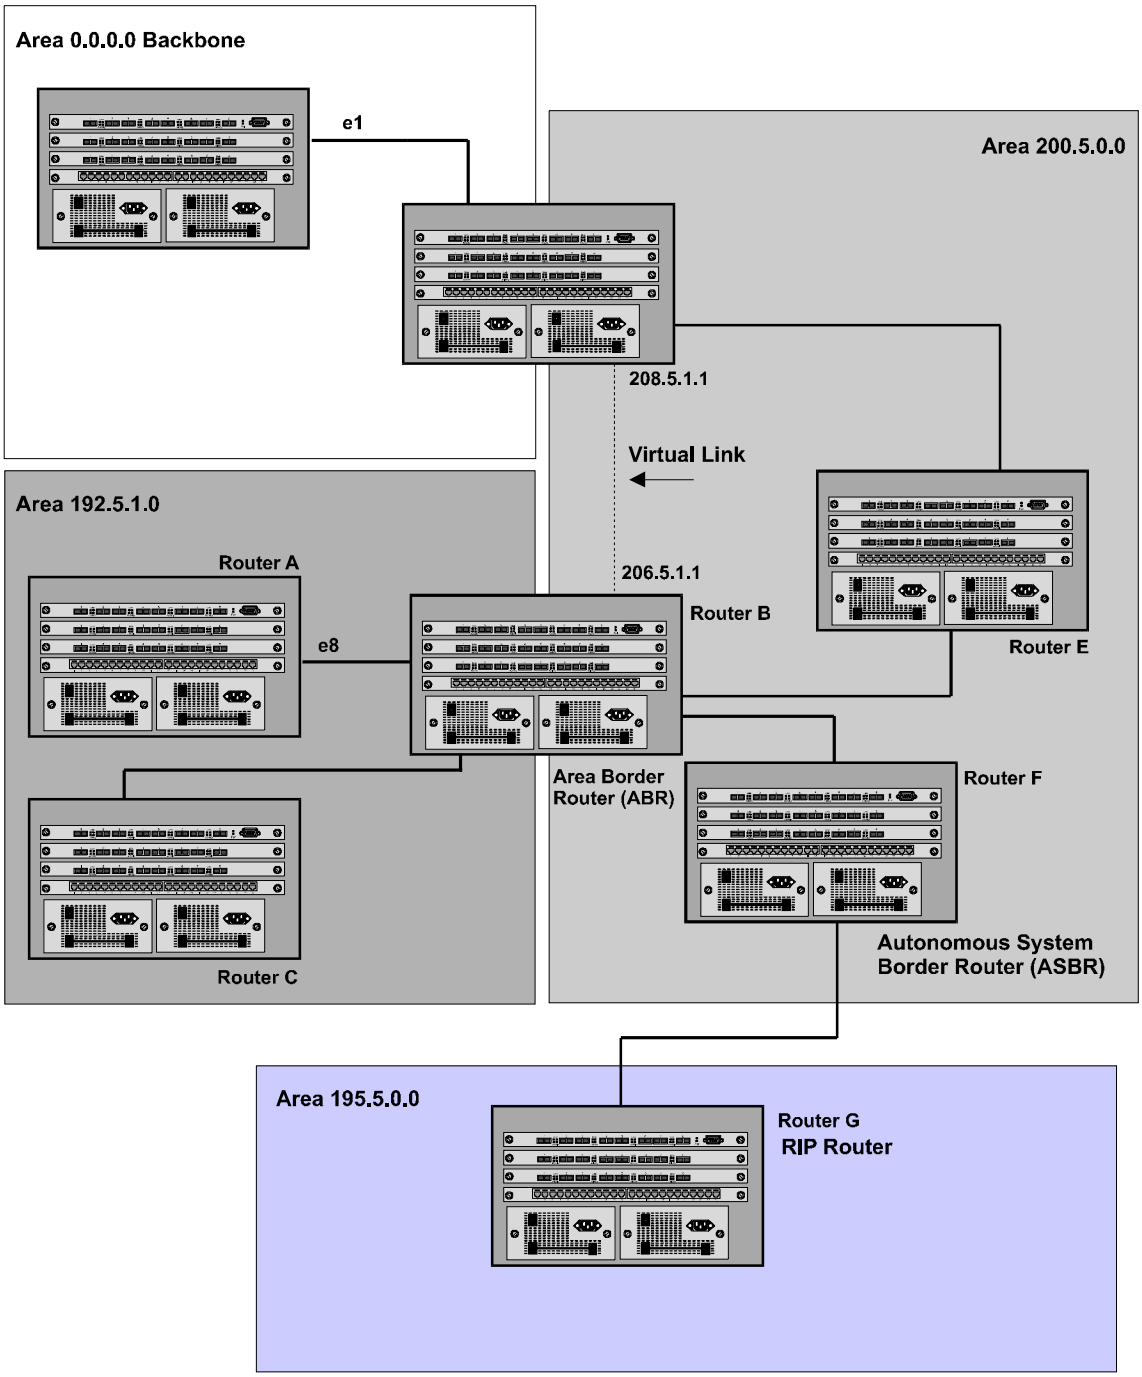

Figure 9.7 Defining OSPF virtual links within a network

USING THE CLI

To define a virtual link between the area border router (Router B) and the area backbone 0.0.0.0 via Router D (**Figure 9.7**) which has a physical connection to the area backbone, the user must first define the physical link to area 0.0.0.0:

```
RouterD(config-ospf-router)# int e 2/1

RouterD(config-if-2/1)# ip ospf area 0.0.0.0

RouterD (config-if-2/1)# write memory
```

To define the virtual link between the ABR, Router B and Router D, the user would define the virtual link on both Router B and Router D and enter the following:

```
RouterB(config-ospf-router)# area 200.5.0.0 virtual 208.5.1.1

RouterB(config-ospf-router)# write mem

RouterB(config-ospf-router)# telnet 208.5.1.1

RouterD(config-ospf-router)# area 200.5.0.0 virtual 206.5.1.1

RouterB(config-ospf-router)# write mem
```

NOTE: The ***area <ip address>*** represents the shared area of the two ABRs—the one with a physical connection to the backbone and the router that needs to define a logical connection to the backbone. The defined area serves as the connection point between the two routers.

When configuring the virtual link on the router requiring a logical connection, **router ID** in the command ***virtual-link <router id>*** represents the lowest IP address of the router that is physically connected to the backbone.

When configuring the virtual link on the physically connected router, **router ID** represents the IP address of the router interface that requires a logical connection to the backbone.

USING THE WEB MANAGEMENT INTERFACE

To define a virtual link between the area border router (Router B) and the area backbone 0.0.0.0 via Router D (**Figure 9.7**) which has a physical connection to the area backbone, the user must first define the physical link to area 0.0.0.0:

1. Select OSPF virtual link from the OSPF configuration sheet.
2. Select the add OSPF virtual link interface from the link summary panel that appears. The panel shown in **Figure 9.8** will then appear.
3. Select the **transit area ID** from the pull down menu. The transit area is the area ID of the area shared by both routers. For the example, to configure the network shown in **Figure 9.7**, the user would enter 200.5.0.0.
4. Select an **authentication** method from the pull down menu. If simple is chosen, then enter the authentication key in the appropriate field. If MD5 is selected, then enter the MD5 authentication ID, key and wait time.

   NOTE: A description of authentication is provided in the **Overview of OSPF Physical and Virtual Parameters** section.

   ---

5. Enter the IP address of the **neighbor router**. This will be the IP address of the interface of the router at the other end of the virtual link. Therefore, if the user is configuring Router B, the neighbor router address would be that of Router D (208.5.1.1). If the user is configuring Router D, the neighbor router address would be that of Router B (206.5.1.1)
6. Modify the default settings of the following parameters if desired: **hello interval**, **transit delay**, **retransmit interval** and **dead interval**.
7. Select the **add** button to apply the changes.
8. Log onto the neighbor router and configure the other end of the virtual link.
9. Repeat steps 2 through 7 for the neighbor router.

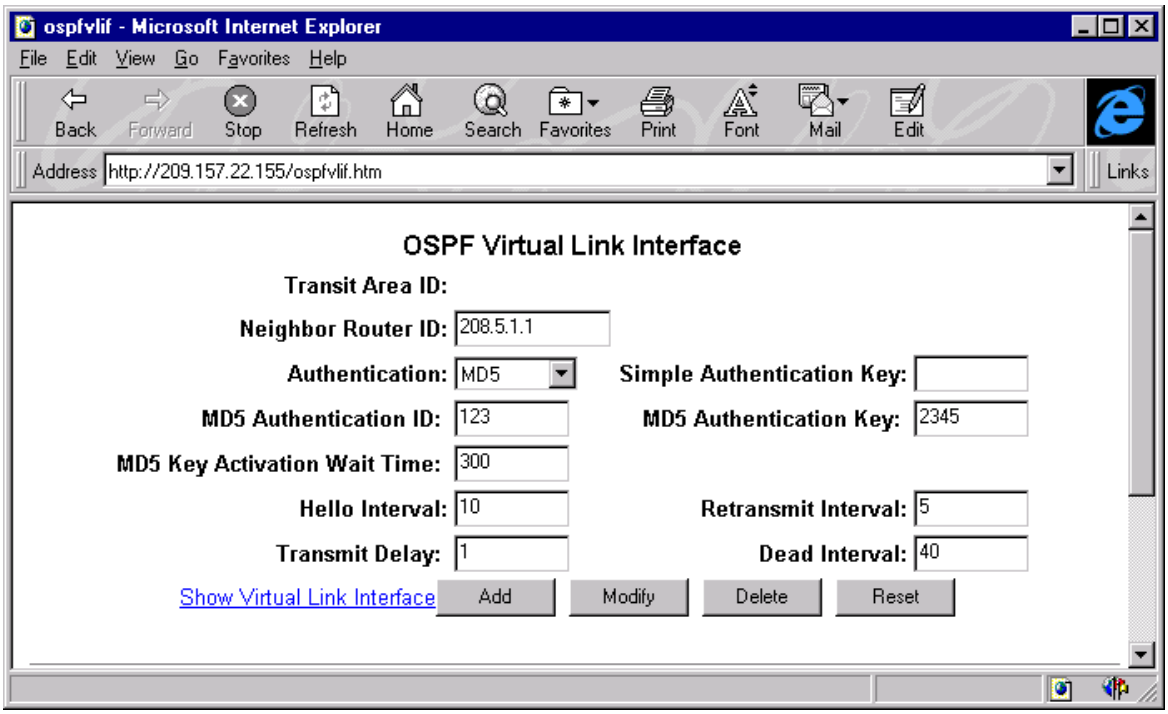

Figure 9.8 Defining a virtual link interface

Modifying Virtual Link Parameters

OSPF comes standard with a number of parameters that can be modified for virtual links. Notice that these are the same parameters that can be modified for physical interfaces.

USING THE CLI

Default values for virtual links can be modified using the CLI command ***area <ip address> virtual-link <ip address> [authentication-key | dead-interval | hello-interval | retransmit-interval | transmit-delay <value>*** at the OSPF router level of the CLI. A description and summary of possible values is summarized in the next section.

For additional command descriptions and value ranges, please refer to **Appendix B**.

USING THE WEB MANAGEMENT INTERFACE

To modify virtual link default values, the user should do the following:

1. Select OSPF virtual link from the OSPF configuration sheet. The panel shown in **Figure 9.8** will appear.
2. Select the **transit area ID** of the shared area for which the virtual link is defined.
3. Enter the area ID of the **neighbor router**.
4. Modify the virtual link parameter defaults as required.

NOTE: A list of all possible virtual link parameters is summarized in the **Overview of OSPF Physical and Virtual Interface Parameters** section.

---

5. Select the **add** button to assign the changes.
6. Log onto the neighbor router and configure parameter changes to match those configured for the local router.

Enabling Redistribution

Redistribution is used to import and translate different protocol routes into a specified protocol type. On HP routing switches, redistribution is supported for OSPF, RIP and static routes. When configuring redistribution for RIP, the user can specify that either OSPF or static routes or both are imported into RIP routes. OSPF redistribution supports the import of static or RIP routes or both into OSPF routes.

In **Figure 9.9**, the user wants to configure the network so that interface 206.5.1.1 on Router B will serve as the link to the RIP Network and redistribute (translate) all OSPF or static routes into RIP routes. All the other routers in the OSPF network will be configured so that all RIP or static routes are forwarded to Router B.

NOTE: Router B must be running RIP and OSPF protocols to support this activity.

When using the CLI, redistribution is configured at either the RIP or OSPF router level of the CLI.

On the Web management interface, redistribution is enabled on the RIP and OSPF configuration sheets.

USING THE CLI

EXAMPLE 1: To enable Router B (**Figure 9.9**) to redistribute both RIP and OSPF routes, the user would enter the following:

```
RouterB(config)# router rip

RouterB(config-rip-router)# redistribution

RouterB(config-rip-router)# router ospf

RouterB(config-ospf-router)# redistribution
```

EXAMPLE 2: To enable RIP and define the RIP type on interface 1 of Router B, the user would enter the following:

```
RouterB(config-rip-router)# int e 2/1

RouterB(config-if-2/1)# ip rip v1-compatible-v2
```

USING THE WEB MANAGEMENT INTERFACE

To enable redistribution on the router, the user would do the following:

1. Select the OSPF link from the main menu. The panel seen in **Figure 9.12** will be seen,
2. Enable **redistribution**.
3. Select the **apply** button to assign the change.

Defining Redistribution Filters

After enabling redistribution, the user must define the redistribution tables with deny and permit redistribution filters.

NOTE: When configuring with CLI, the ***deny*** and ***permit redistribute*** commands for OSPF are found at the OSPF router level.

When configuring with the Web management interface, the user selects the redistribution filter link from the OSPF configuration sheet.

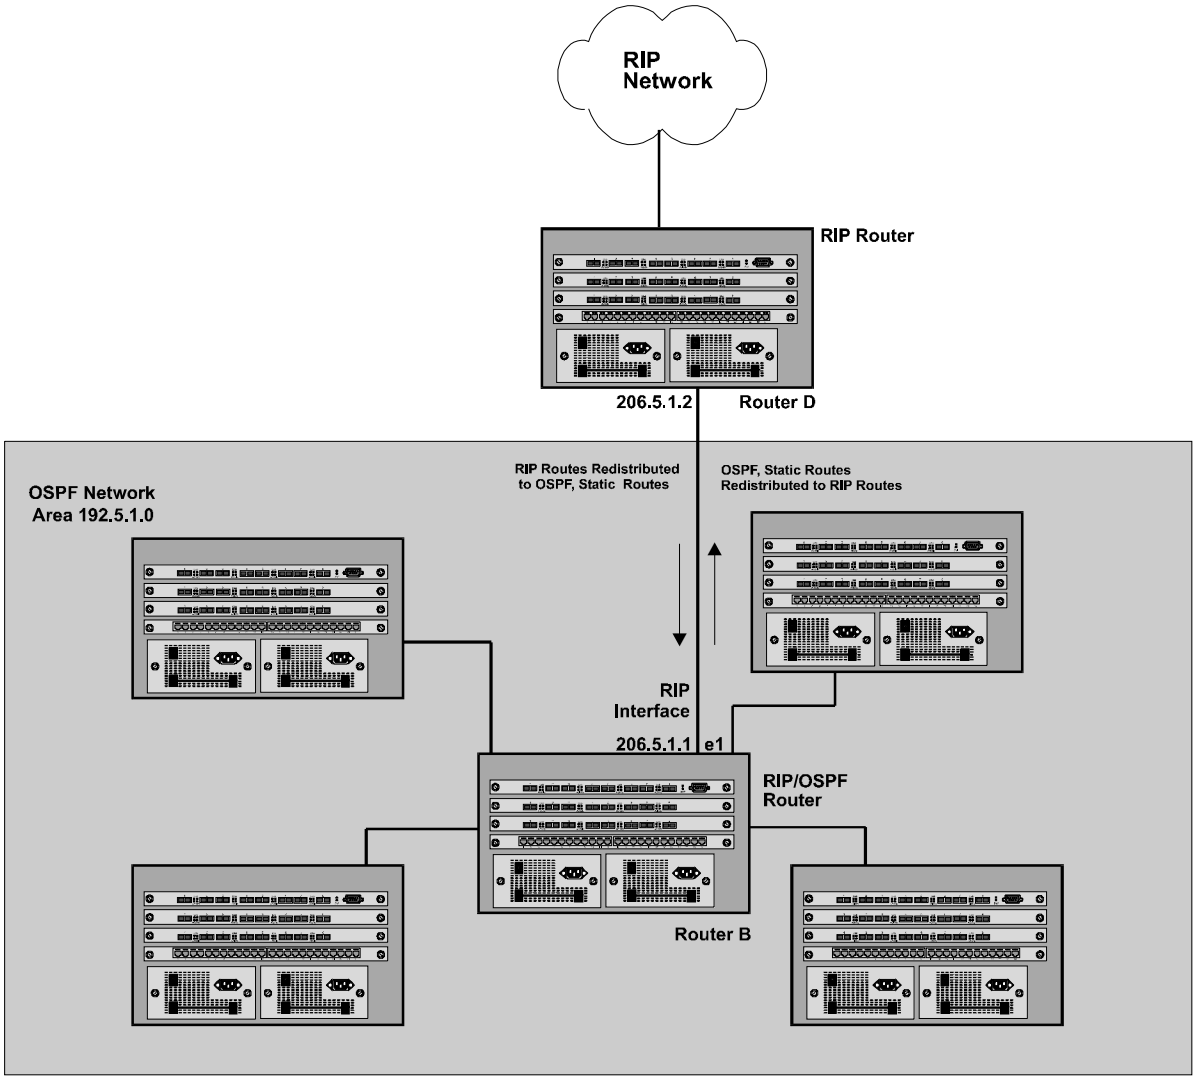

Figure 9.9 Redistributing OSPF and static routes to RIP routes

USING THE CLI

EXAMPLE 1: To allow all OSPF and static routes from the OSPF network to be imported into RIP routes when forwarded to interface 206.5.1.1, the user would enter the following:

```
RouterB (config)# router rip
RouterB (config-rip-router)# permit redistribute 1 all
```

The user also has the option of specifying import of just OSPF or static routes, as well as specifying that only traffic from a specific network or cost (metric) be imported, as noted in the command syntax highlighted below:

syntax: <deny|permit> redistribute <index> <all|ospf|static> address <ip address> <ip address> [match-metric<value>|set-metric <value>]

EXAMPLE 2: To allow all RIP and static routes from the RIP network to be imported into OSPF routes when forwarded to interface 206.5.1.2 on Router D, the user would enter the following commands:

```
RouterB (config)# router ospf
RouterB (config-ospf-router)# permit redistribute 1 all
```

The user also has the option of specifying import of just RIP or static routes as well as specifying that only traffic from a specific network or cost (metric), be imported, as noted in the command syntax highlighted below:

syntax: <deny| permit> redistribute <index> <all|rip|static> address <ip address> <ip address>|match-metric<value>|set-metric <value>

USING THE WEB MANAGEMENT INTERFACE

To define which routes are imported into OSPF, the user can define a redistribution filter. To do so, the user would do the following:

1. Select redistribution filter from the OSPF configuration sheet. The panel shown in **Figure 9.10** will appear.

NOTE: If redistribution filters are already defined on a router, then the summary panel, Show OSPF Redistribution Filter, will appear. The user will then need to enter add OSPF redistribution filter to get to the redistribution entry panel.

---

2. Enter the IP address and mask for routes that are to be permitted or denied. Entering 0.0.0.0 for the IP address and mask will allow routes from all networks to be imported.
3. Enter a **filter ID**. This can be any unused value between 1 and 64.
4. Select either **permit** or **deny**.
5. Select **static**, **RIP routes** or **all** to specify which protocol(s) to allow or deny being imported into OSPF routes.
6. To specify that only those routes that match a specific metric be imported, enable **match RIP metric** and enter a specific value other than zero in the **match metric** field.
7. To apply an OSPF metric, other than that defined at the global level, to all imported routes, enable **set OSPF metric**, then enter a value into the **set metric** field.
8. When all parameters are entered, select **add** to apply the changes.
9. Repeat steps 1 through 8 for each redistribution filter that is to be defined.

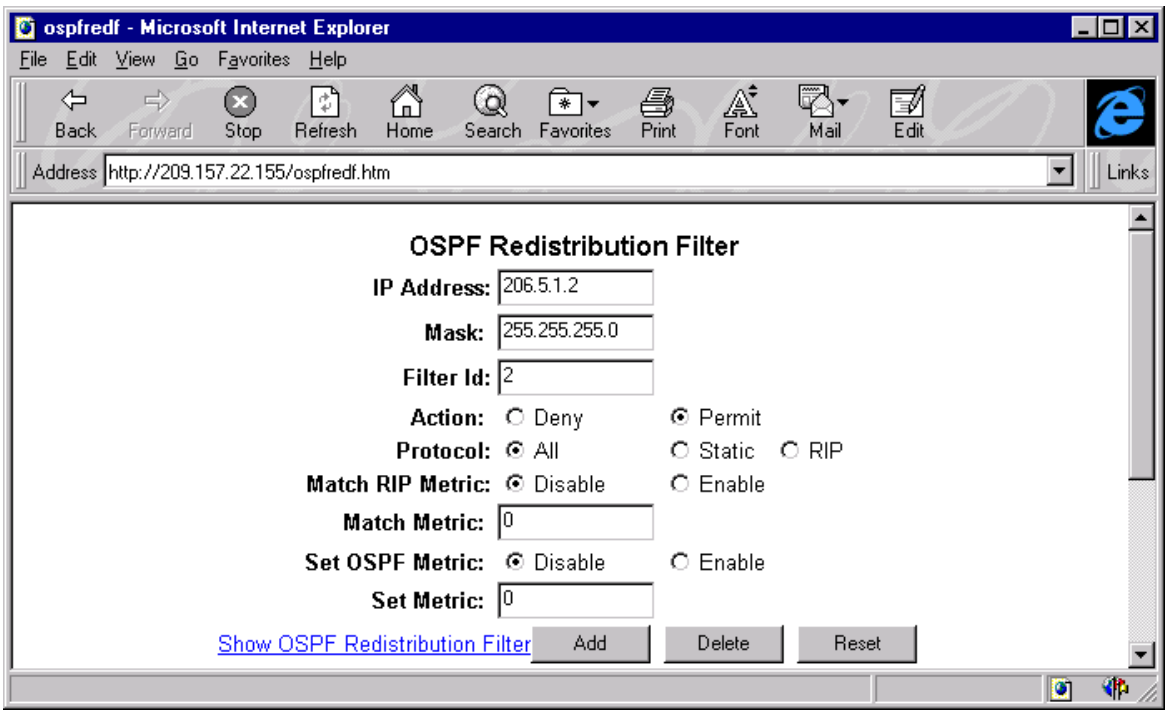

Figure 9.10 Importing of RIP and static routes into OSPF routes for interface 208.5.1.1 (Router B)

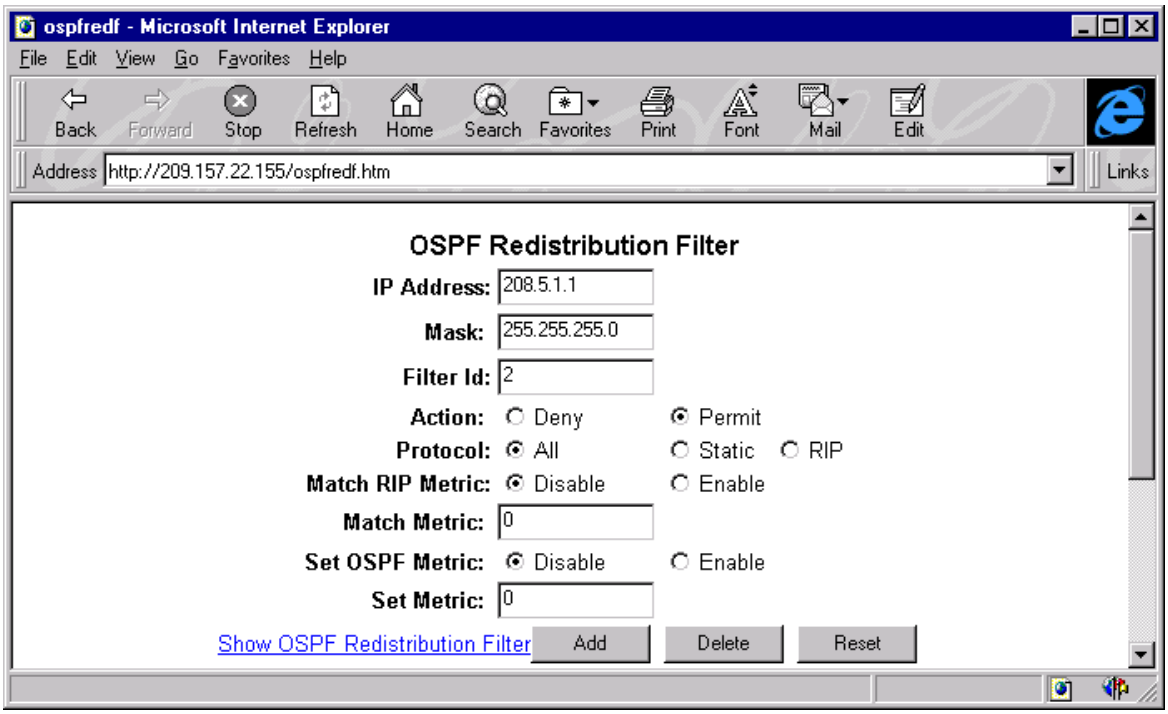

Figure 9.11 Importing of RIP and static routes into OSPF routes for interface 206.5.1.2 (Router D)

Modify Default Metric for Redistribution

Default metric is a global parameter that specifies the cost that will be applied to all OSPF routes by default unless defined at the interface level. The default value is 1 and it can be assigned a value of between 1 and 15.

USING THE CLI

To assign a default metric of 4 to all routes (RIP and Static) imported as OSPF, the user would enter the following:

```
HP9300(config)# router ospf
HP9300(config-ospf-router)# default-metric 4
```

USING THE WEB MANAGEMENT INTERFACE

To modify the cost that is assigned to redistributed routes:

1. Select the OSPF link from the main menu and the panel seen in **Figure 9.12** will appear.
2. Enter a value between 1 and 15 in the **default metric** field.
3. Select the **apply** button to assign the changes.

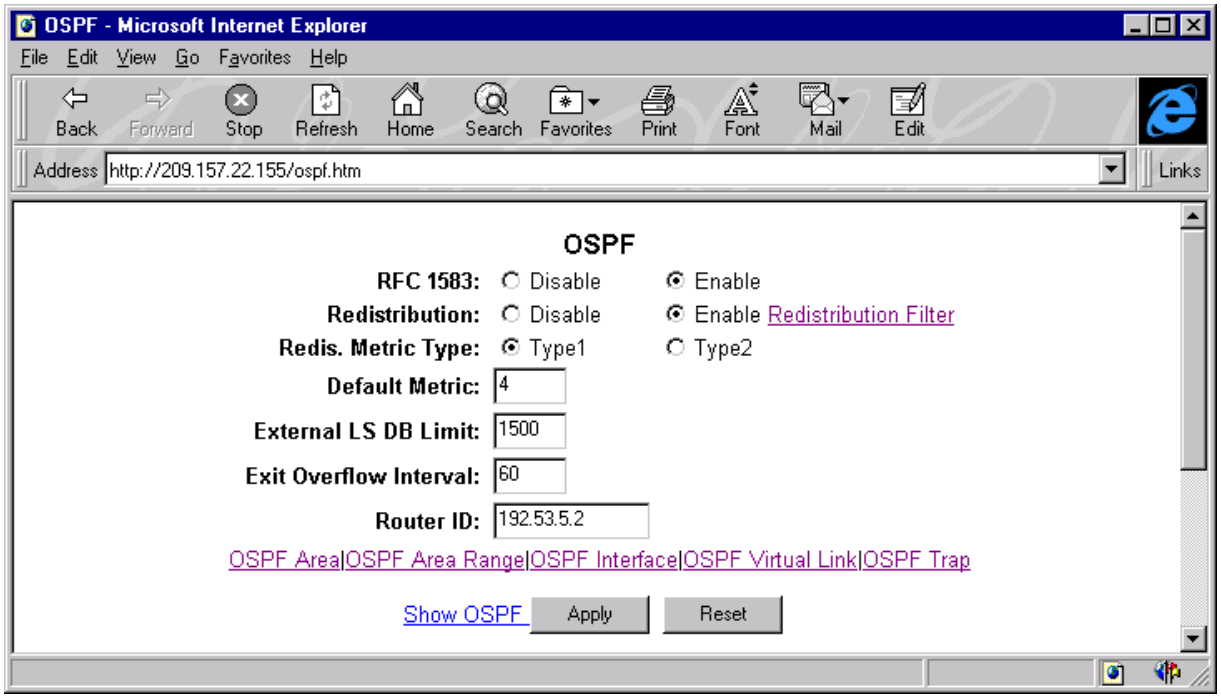

Figure 9.12 OSPF configuration sheet used to define global parameter values and provide access points to other configuration panels

Modify Redistribution Metric Type

This metric type will be used by default for all routes imported into OSPF unless specified differently via the redistribution filter. Type 2 specifies a big metric (3 bytes). Type 1 specifies a small metric (2 bytes). The default value is type 2.

USING THE CLI

To modify the default value to type 1, the user would enter the following:

```
HP9300(config-ospf-router)# metric-type type1
```

USING THE WEB MANAGEMENT INTERFACE

To modify the default metric type, the user would enter the following:

1. Select the OSPF link from the main menu and the panel seen in **Figure 9.12** will appear.
2. Select either type 1 or type 2 next to the **redistribution metric type**.
3. Select the **apply** button to assign the changes.

Modify External LSDB Limit

This system (global) level parameter allows the user to modify the number of IP OSPF external link state advertisements that the routing switch will allow before a database overflow condition is declared on the system. It is in compliance with RFC 1765.

By default, HP routing switches support up to 2000 IP OSPF external link state advertisements before a database overflow condition is reported. The minimum configurable value is 1.

USING THE CLI

To modify the default value to 1500 advertisements, the user would enter the following:

```
HP9300(config-ospf-router)# external-lsdb-limit 1500
```

USING THE WEB MANAGEMENT INTERFACE

To modify the number of IP OSPF external link state advertisements, the user would:

1. Select the OSPF link from the main menu and the panel seen in **Figure 9.12** will appear.
2. Enter a value between 1 and 2,000 in the **external LSDB limit** field.
3. Select the **apply** button to assign the changes.

Modify Exit Overflow Interval

Once a database overflow condition exists on a routing switch, it will seek to eliminate that condition by removing entries that originated on the routing switch. The exit overflow interval allows the user to set how often a routing switch will check to see if the overflow condition has been eliminated. If the configured value of the database overflow interval is zero, then the routing switch will never leave the database overflow condition. The default value is 0, with a possible configurable range of 0 to 86,400 seconds (24 hours).

USING THE CLI

To modify the exit overflow interval to 60 seconds, the user would enter the following:

```
HP9300(config-ospf-router)# data-base-overflow-interval 60
```

USING THE WEB MANAGEMENT INTERFACE

To modify the exit overflow interval, the user would:

1. Select the OSPF link from the main menu and the panel seen in **Figure 9.12** will appear.
2. Enter a value between 0 to 86,400 in the **exit overflow interval** field.
3. Select the **apply** button to assign the changes.

Modify OSPF Traps Generated

OSPF traps as defined by RFC 1850 are supported on HP routing switches. OSPF trap generation is enabled on the router, by default.

USING THE CLI

When using the CLI, the user can disable all or specific OPSF trap generation, by entering the CLI command: ***no snmp-server trap ospf***. To later re-enable the trap feature, the user would enter ***snmp-server trap ospf.***

To disable a specific OSPF trap, the user can enter that command ***no snmp-server trap ospf <ospf trap>***.

Both of these commands are found OSPF router Level of the CLI.

A summary of OSPF traps supported on HP routing switches, their corresponding CLI command and their associated MIB objects from RFC 1850, is provided below:

- interface-state-change-trap[MIB object: OspfIfstateChange]
- virtual-interface-state-change-trap[MIB object: OspfVirtIfStateChange
- neighbor-state-change-trap [MIB object:ospfNbrStateChange]
- virtual-neighbor-state-change-trap [MIB object: ospfVirtNbrStateChange]
- interface-config-error-trap [MIB object: ospfIfConfigError]
- virtual-interface-config-error-trap[MIB object: ospfIfConfigError]
- interface-authentication-failure-trap[MIB object: ospfIfAuthFailure]
- virtual-interface-authentication-failure-trap[MIB object: ospfVirtIfAuthFailure]

- interface-receive-bad-packet-trap[MIB object: ospfIfrxBadPacket]
- virtual-interface-receive-bad-packet-trap[MIB object: ospfVirtIfRxBadPacket]
- interface-retransmit-packet-trap[MIB object: ospfTxRetransmit]
- virtual-interface-retransmit-packet-trap[MIB object: ospfVirtIfTxRetransmit]
- originate-lsa-trap [MIB object: ospfOriginateLsa]
- originate-maxage-lsa-trap [MIB object: ospfMaxAgeLsa]
- link-state-database-overflow-trap [MIB object: ospfLsdbOverflow]
- link-state-database-approaching-overflow-trap [MIB object: ospfLsdbApproachingOverflow

EXAMPLE 1: To stop an OSPF trap from being collected, the user would use the CLI command: ***no trap <trap>,*** found in the Router OSPF level of the CLI. To disable reporting of the neighbor-state-change-trap, the user would enter the following:

```
HP9300 (config-ospf-router)# no trap neighbor-state-change-trap
```

EXAMPLE 2: To reinstate the command, the user would just enter this command:

```
HP9300 (config-ospf-router)# trap neighbor-state-change-trap
```

syntax: [no] snmp-server trap ospf <ospf trap>

USING THE WEB MANAGEMENT INTERFACE

To disable a specific OSPF trap or traps, the user would:

1. Select the OSPF Trap link from the OSPF configuration sheet and the panel shown in **Figure 9.13** will appear.
2. Select the **disable** option beside the OSPF trap to disable it.
3. Select the **apply** button to assign the changes.

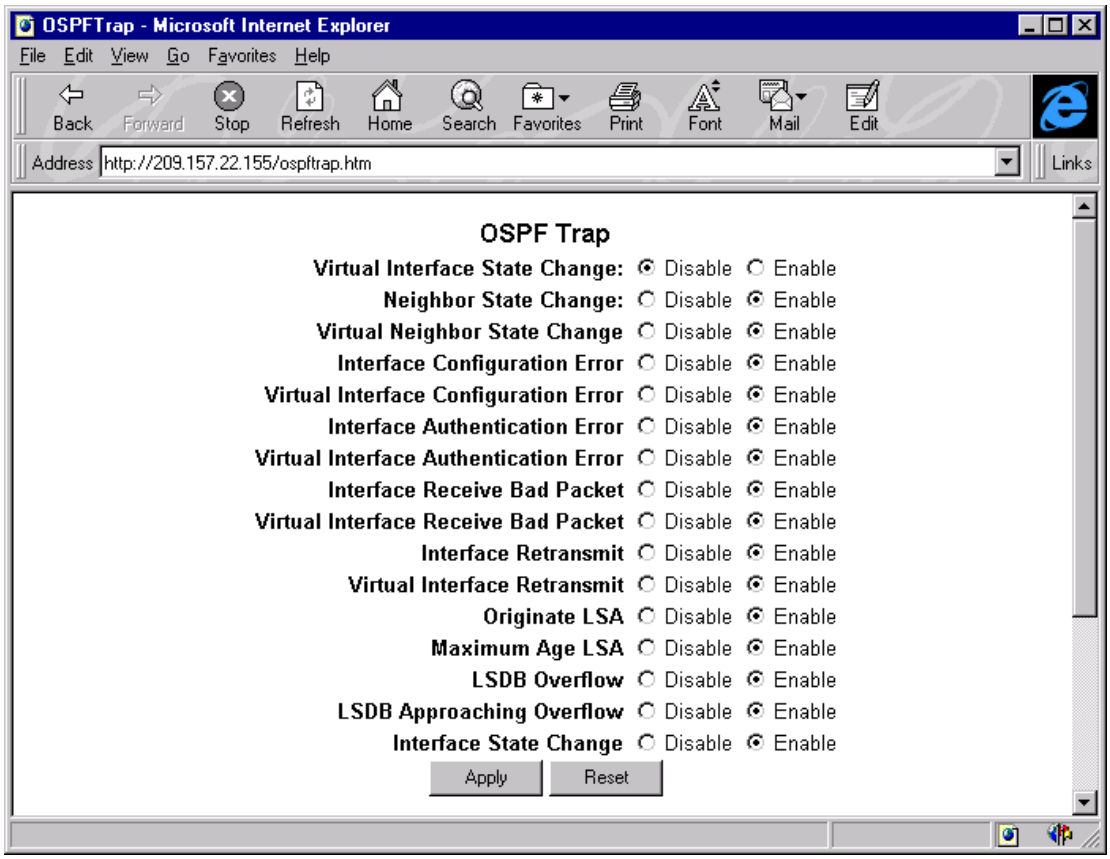

Figure 9.13 OSPF trap summary noting state of traps

Modify OSPF standard compliance setting

HP routing switches are configured, by default, to be compliant with RFC 1583 OSPF V2 specification.

USING THE CLI

To configure a router to operate with the latest OSPF standard, RFC 2178, the user would enter the following command:

```
HP9300(config)# router ospf
HP9300(config-ospf-router)# no rfc1583-compatibility
```

USING THE WEB MANAGEMENT INTERFACE

To configure a router to operate with the latest OPSF standard, RFC 2178, the user would

1. Select the OSPF link from the main menu and the panel shown in **Figure 9.12** will appear.
2. **Disable** RFC 1583.
3. Select the **apply** button to assign the change.