

## Chapter 3 Preferences and Route Selection

### 3.1 Preferences Overview

Preference is the value that GateD uses to select one route over another when more than one route to the same destination is learned from different protocols or peers. Preference can be set in the GateD configuration files in several different configuration statements. Preference can be set based on one network interface over another, one protocol over another, or one remote gateway over another. Preference cannot be used to control the selection of routes within an interior gateway protocol. This control is accomplished automatically by the protocol based on metric. Preference can be used to select routes from the same exterior gateway protocol (such as BGP) learned from different peers or autonomous systems. Each route has only one configurable preference value associated with it, even though preference can be set at many places in the configuration file. Simply, the last or most specific preference value set for a route is the value used.

The **preference** value is an arbitrarily assigned value used to determine the order of routes to the same destination in a single routing database. The active route is chosen by the lowest **preference** value. Some protocols implement a second preference (**preference2**), sometimes referred to as a tie-breaker. BGP and OSPF protocols use **preference2**. **preference2** is for internal use only and is not configurable. Its value is used only when comparing routes with equal values of preference.

### 3.2 Assigning Preferences

A default preference is assigned to each source from which GateD receives routes. Preference values range from 0 to 255, with the lowest number indicating the most preferred route.

The following table summarizes the default preference values for routes learned in various ways. The table lists the statements (some of which are clauses within statements) that set preference and shows the types of routes to which each statement applies. The table lists the preference precedence between protocols and the default preference for each type of route. The more narrow the scope of the statement, the higher the precedence its preference value is given, but the smaller the set of routes it affects.

**Table 2: Preference Selection Precedence**

Preference of	Defined by Statement	Default
Direct connected networks	<code>interface</code>	0
Routes to interface aliases		1
OSPF routes	<code>ospf</code>	10
IS-IS level 1 routes	<code>isis level 1</code>	15
IS-IS level 2 routes	<code>isis level 2</code>	18
Redirects	<code>redirect</code>	30
Routes learned via route socket	<code>kernel</code>	40
Routes installed via SNMP		50
Routes learned via router discovery	<code>router-discovery</code>	55
Static routes from config	<code>static</code>	60
RIP routes	<code>rip</code> or <code>ripng</code>	100
Point-to-point interface		110
Routes to interfaces that are down	<code>interface</code>	120
Aggregate/generate routes	<code>aggregate/generate</code>	130
OSPF AS external routes	<code>ospf</code>	150
BGP routes	<code>bgp</code>	170
Routes in kernel at startup		254

### 3.3 Sample Preference Configurations

```

interfaces {
    interface 138.66.12.2 preference 10 ;
} ;
rip yes {
    preference 90 ;
} ;
import proto rip gateway 138.66.12.1 preference 75 ;

```

In these statements, the preference applicable to routes learned via RIP from gateway 138.66.12.1 is 75. The last preference applicable to routes learned via RIP from gateway 138.66.12.1 is defined in the **accept** statement. The preference applicable to other RIP routes is found in the **rip** statement. The preference set on the **interface** statement applies only to the route to that interface.

