



## EIGRP Wide Metrics

The EIGRP Wide Metrics feature supports 64-bit metric calculations and Routing Information Base (RIB) scaling in Enhanced Interior Gateway Routing Protocol (EIGRP) topologies. The 64-bit calculations work only in EIGRP named mode configurations. EIGRP classic mode configurations use 32-bit calculations. This module provides an overview of the EIGRP Wide Metrics feature.

- [Information About EIGRP Wide Metrics, on page 1](#)
- [Feature Information for Overview of Cisco TrustSec, on page 5](#)

## Information About EIGRP Wide Metrics

### EIGRP Composite Cost Metrics

The Enhanced Interior Gateway Routing Protocol (EIGRP) uses bandwidth, delay, reliability, load, and K values (various constants that can be configured by a user to produce varying routing behaviors) to calculate the composite cost metric for local Routing Information Base (RIB) installation and route selections. The EIGRP composite cost metric is calculated using the following formula:

$$\text{EIGRP composite cost metric} = 256 * ((K1 * \text{Scaled Bw}) + (K2 * \text{Scaled Bw}) / (256 - \text{Load}) + (K3 * \text{Scaled Delay}) * (K5 / (\text{Reliability} + K4)))$$

EIGRP uses one or more vector metrics to calculate the composite cost metric. The table below lists EIGRP vector metrics and their descriptions.

**Table 1: EIGRP Vector Metrics**

Vector Metric	Description
bandwidth	The minimum bandwidth (Bw) of the route, in kilobits per second. It can be 0 or any positive integer. The bandwidth for the formula is scaled and inverted by using the following formula: Scaled Bw = $(10^7 / \text{minimum bandwidth (Bw) in kilobits per second})$
delay	Route delay, in tens of microseconds. Scaled Delay = $(\text{Delay} / 10)$
load	The effective load of the route, expressed as a number from 0 to 255 (255 is 100 percent loading).

Vector Metric	Description
mtu	The minimum maximum transmission unit (MTU) size of the route, in bytes. It can be 0 or any positive integer.
reliability	The likelihood of successful packet transmission, expressed as a number between 0 and 255, where 255 means 100 percent reliability and 0 means no reliability.

EIGRP monitors metric weights, by using K values, on an interface to allow the tuning of EIGRP metric calculations and to indicate the type of service (ToS). K values are integers from 0 to 128; these integers, in conjunction with variables like bandwidth and delay, are used to calculate the overall EIGRP composite cost metric. The table below lists the K values and their defaults.

**Table 2: EIGRP K-Value Defaults**

Setting	Default Value
K1	1
K2	0
K3	1
K4	0
K5	0

Although you can configure K values to produce varying routing behaviors, most configurations use only the delay and bandwidth metrics by default, with bandwidth taking precedence, to produce a single 32-bit metric. Use of the default constants effectively reduces the above-mentioned composite cost metric formula to the following default formula:  $256 * (\text{Scaled Bw} + \text{Scaled Delay})$ .

For example, let us consider a link whose bandwidth to a particular destination is 128 kb/s and the delay is 84,000 microseconds. By using the default formula, you can simplify the EIGRP composite cost metric calculation to  $256 * (\text{Scaled Bw} + \text{Scaled Delay})$ , thus resulting in the following value:

$$\text{Metric} = 256 * (10^7 / 128 + 84000 / 10) = 256 * 86525 = 22150400$$

## EIGRP Wide Metrics

The Enhanced Interior Gateway Routing Protocol (EIGRP) composite cost metric (calculated using the bandwidth, delay, reliability, load, and K values) is not scaled correctly for high-bandwidth interfaces or Ethernet channels, resulting in incorrect or inconsistent routing behavior. The lowest delay that can be configured for an interface is 10 microseconds. As a result, high-speed interfaces, such as 10 Gigabit Ethernet (GE) interfaces, or high-speed interfaces channeled together (GE ether channel) will appear to EIGRP as a single GE interface. This may cause undesirable equal-cost load balancing. To resolve this issue, the EIGRP Wide Metrics feature supports 64-bit metric calculations and Routing Information Base (RIB) scaling that provide the ability to support interfaces (either directly or via channeling techniques like port channels or ether channels) up to approximately 4.2 terabits.




---

**Note** The 64-bit metric calculations work only in EIGRP named mode configurations. EIGRP classic mode uses 32-bit metric calculations.

---

To accommodate interfaces with bandwidths above 1 gigabit and up to 4.2 terabits and to allow EIGRP to perform path selections, the EIGRP composite cost metric formula is modified. The paths are selected based on the computed time. The time that information takes to travel through links is measured in picoseconds. The interfaces can be directly capable of these high speeds, or the interfaces can be bundles of links with an aggregate bandwidth greater than 1 gigabit.

$$\text{Metric} = [(K1 * \text{Minimum Throughput} + \{K2 * \text{Minimum Throughput} / 256 - \text{Load}\} + (K3 * \text{Total Latency}) + (K6 * \text{Extended Attributes})] * [K5 / (K4 + \text{Reliability})]$$

Default K values are as follows:

- K1 = K3 = 1
- K2 = K4 = K5 = 0
- K6 = 0

The EIGRP Wide Metrics feature also introduces K6 as an additional K value for future use.

By default, the path selection scheme used by EIGRP is a combination of throughput (rate of data transfer) and latency (time taken for data transfer), and the formula for calculating the composite cost metric is as follows:

$$\text{Composite Cost Metric} = (K1 * \text{Minimum Throughput}) + (K3 * \text{Total Latency})$$

$$\text{Minimum Throughput} = (10^7 * 65536) / \text{Bw}, \text{ where } 65536 \text{ is the wide-scale constant.}$$

$$\text{Total Latency for bandwidths below 1 gigabit} = (\text{Delay} * 65536) / 10, \text{ where } 65536 \text{ is the wide-scale constant.}$$

$$\text{Total Latency for bandwidths above 1 gigabit} = (10^7 * 65536 / 10) / \text{Bw}, \text{ where } 65536 \text{ is the wide-scale constant.}$$

With the calculation of larger bandwidths, EIGRP can no longer fit the computed metric into a 4-byte unsigned long value that is needed by the Cisco RIB. To set the RIB scaling factor for EIGRP, use the **metric rib-scale** command. When you configure the **metric rib-scale** command, all EIGRP routes in the RIB are cleared and replaced with the new metric values.

## EIGRP Metric Weights

You can use the **metric weights** command to adjust the default behavior of Enhanced Interior Gateway Routing Protocol (EIGRP) routing and metric computations. EIGRP metric defaults (K values) have been carefully selected to provide optimal performance in most networks.




---

**Note** Adjusting EIGRP metric weights can dramatically affect network performance. Because of the complexity of this task, we recommend that you do not change the default K values without guidance from an experienced network designer.

---

By default, the EIGRP composite cost metric is a 32-bit quantity that is the sum of segment delays and the lowest segment bandwidth (scaled and inverted) for a given route. The formula used to scale and invert the bandwidth value is  $10^7 / \text{minimum bandwidth in kilobits per second}$ . However, with the EIGRP Wide Metrics

feature, the EIGRP composite cost metric is scaled to include 64-bit metric calculations for EIGRP named mode configurations.

For a network of homogeneous media, this metric reduces to a hop count. For a network of mixed media (FDDI, Gigabit Ethernet (GE), and serial lines running from 9600 bits per second to T1 rates), the route with the lowest metric reflects the most desirable path to a destination.

## Mismatched K Values

EIGRP K values are the metrics that EIGRP uses to calculate routes. Mismatched K values can prevent neighbor relationships from being established and can negatively impact network convergence. The example given below explains this behavior between two EIGRP peers (Device-A and Device-B).

The following configuration is applied to Device-A. The K values are changed using the **metric weights** command. A value of 2 is entered for the *k1* argument to adjust the bandwidth calculation. A value of 1 is entered for the *k3* argument to adjust the delay calculation.

```
Device(config)# hostname Device-A
Device-A(config)# interface serial 0
Device-A(config-if)# ip address 10.1.1.1 255.255.255.0
Device-A(config-if)# exit
Device-A(config)# router eigrp name1
Device-A(config-router)# address-family ipv4 autonomous-system 4533
Device-A(config-router-af)# network 10.1.1.0 0.0.0.255
Device-A(config-router-af)# metric weights 0 2 0 1 0 0 1
```

The following configuration is applied to Device-B, and the default K values are used. The default K values are 1, 0, 1, 0, 0, and 0.

```
Device(config)# hostname Device-B
Device-B(config)# interface serial 0
Device-B(config-if)# ip address 10.1.1.2 255.255.255.0
Device-B(config-if)# exit
Device-B(config)# router eigrp name1
Device-B(config-router)# address-family ipv4 autonomous-system 4533
Device-B(config-router-af)# network 10.1.1.0 0.0.0.255
Device-B(config-router-af)# metric weights 0 1 0 1 0 0 0
```

The bandwidth calculation is set to 2 on Device-A and set to 1 (by default) on Device-B. This configuration prevents these peers from forming a neighbor relationship.

The following error message is displayed on the console of Device-B because the K values are mismatched:

```
*Apr 26 13:48:41.811: %DUAL-5-NBRCHANGE: IP-EIGRP(0) 1: Neighbor 10.1.1.1 (Ethernet0/0) is
down: K-value mismatch
```

The following are two scenarios where the above error message can be displayed:

- Two devices are connected on the same link and configured to establish a neighbor relationship. However, each device is configured with different K values.
- One of two peers has transmitted a “peer-termination” message (a message that is broadcast when an EIGRP routing process is shut down), and the receiving device does not support this message. The receiving device will interpret this message as a K-value mismatch.

## Feature Information for Overview of Cisco TrustSec

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to [www.cisco.com/go/cfn](http://www.cisco.com/go/cfn). An account on Cisco.com is not required.

**Table 3: Feature Information for Overview of Cisco TrustSec**

Feature Name	Releases	Feature Information
IPv6 enablement - Inline Tagging	Cisco IOS XE Fuji 16.8.1	The support for IPv6 is introduced.

