



Enhanced Interior Gateway Routing Protocol (EIGRP) Wide Metrics

Last updated: February 2016

Contents

Introduction	3
Prerequisites	3
Issues with EIGRP Composite Metric Formula	3
Problems with Scaling.....	4
Bandwidth Scaling Issues	5
Interface Delay Scaling Issues.....	5
Migrating Your Network to Wide Metric Support	6
Impact on Configuration Commands.....	6
Impact on Routing—Link Utilization	7
Impact on Routing—Metric Scaling.....	7
Impact on Routing—Link Utilization Changes	8
Impact on Routing—Load Balancing.....	9
Impact on VPNv4 PE/CE Support.....	9
Impact on MANET	10
Wide Metric EIGRP Composite Metric Formula	10
Metric Formulation	11
Usage and Definition of K_1 - K_6	13
Example Wide Metric Computation	14
Composite Metric for Connected Interface.....	14
Composite Metric for Paths.....	15
EIGRP Metric vs. RIB Metric.....	16

Introduction

With the advent of higher speed interfaces and terabit routers, the number of high-speed “multi-gigabit” interface deployments has been growing rapidly, as the current trend in the networking industry is consolidation.

It is becoming increasingly common to deploy Fast EtherChannel, Gigabit EtherChannel, 10 Gigabit Ethernet, or even 100 Gigabit EtherChannel throughout the campus core. By default, Enhanced Interior Gateway Routing Protocol (EIGRP) chooses a route based on minimum path bandwidth and accumulative delay, however in the past the formula that computes these values loses resolution on single links or EtherChannel links above two gigabit. This paper will discuss how EIGRP has been changed to handle these higher link speeds.

Prerequisites

Readers of this paper should have knowledge of the following topics:

- [Enhanced Interior Gateway Routing Protocol](#)
- [An Introduction to EIGRP](#)
- Cisco IOS® Software Classic
- Cisco IOS-XE® Software
- EIGRP Release 8.0 (or later)

To determine the EIGRP release being used, use the **show eigrp plugins** command in user EXEC or privileged EXEC mode.

```
Router# show eigrp plugins
EIGRP feature plugins:::
  eigrp-release      :      8.00.00 : Portable EIGRP Release
                    :      1.00.00 : Source Component Release(rel8)
  igrp2              :      3.00.00 : Reliable Transport/Dual Database
```

Issues with EIGRP Composite Metric Formula

EIGRP adds together weighted values of different network link characteristics in order to calculate a metric for evaluating path selection.

These characteristics include:

- Delay (measured in 10s of microseconds)
- Bandwidth (measured in kilobytes per second)
- Reliability (in numbers ranging from 1 to 255; 255 being the most reliable)
- Load (in numbers ranging from 1 to 255; 255 being saturated)

Various constants (K_1 through K_5) are able to be set by a user to produce varying routing behaviors. However by default, only delay and bandwidth are used in the weighted formula to produce a single 32bit metric:

$$\left[\left(K_1 \cdot \frac{K_2 \cdot \text{Bandwidth}}{\text{Bandwidth} + 256 - \text{Load}} + K_3 \cdot \text{Delay} \right) \frac{K_5}{K_4 + \text{Reliability}} \right] \cdot 256$$

Note: Default K values are: $K_1 = K_3 = 1$ and $K_2 = K_4 = K_5 = 0$
 When K_5 is equal to 0 then $[K_5 / (K_4 + \text{reliability})]$ is defined to be 1

Use of the default constants effectively reduces the formula above to:

$$\text{Bandwidth + Delay} \times 256$$

For the final computation, EIGRP scales the interface values by inverting the bandwidth and scaling the delay with following calculations:


$$\text{ScaledBandwidth} = \frac{10^7}{\text{InterfaceBandwidth}}$$

$$\text{ScaledDelay} = \frac{\text{Delay}}{10}$$

Problems with Scaling

Figure 1 below shows an example network, where the links use ten Gigabit Ethernet interfaces for the 10.2.2.0/24 and 10.4.4.0/24 links, and Gigabit Ethernet interfaces on the 10.3.3.0/24 and 10.5.5.0/24 links. Network 10.1.1.0/24 is a Fast Ethernet interface attached on Router 4.

Figure 1.



Consider the following interface bandwidth and delay values reported via a 10 gigabit and a 1 gigabit interface:

```
Router1#show interface TenGigabitEthernet 2/0
TenGigabitEthernet2/0 is up, line protocol is up
  Hardware is AmdP2, address is aabb.cc00.0102 (bia aabb.cc00.0102)
  Internet address is 10.4.4.1/24
  MTU 1500 bytes, BW 10000000 Kbit/sec, DLY 10 usec,
    reliability 255/255, txload 1/255, rxload 1/255

Router1#sho interface GigabitEthernet e3/0
GigabitEthernet3/0 is up, line protocol is up
  Hardware is AmdP2, address is aabb.cc00.0103 (bia aabb.cc00.0103)
  Internet address is 10.5.5.1/24
  MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 usec,
    reliability 255/255, txload 1/255, rxload 1/255
```

Bandwidth Scaling Issues

As the bandwidth increases, the effect of inverting the “scaled” value begins to cause issues.

Typical interface metrics might look like the following example:

```
GigabitEthernet:
  Scaled Bandwidth = 10,000,000 / 1000000
  Scaled Delay = 10 / 10
  Composite Metric = 10 + 1 * 256 = 2816

10 GigabitEthernet:
  Scaled Bandwidth = 10,000,000 / 10000000
  Scaled Delay = 10 / 10
  Composite Metric = 1 + 1 * 256 = 512

11 GigabitEthernet:
  Scaled Bandwidth = 10,000,000 / 11000000
  Scaled Delay = 10 / 10
  Composite Metric = 0 + 1 * 256 = 256

20 GigabitEthernet:
  Scaled Bandwidth = 10,000,000 / 20000000
  Scaled Delay = 10 / 10
  Composite Metric = 0 + 1 * 256 = 256
```

Any interface exceeding 10 gig, the composite metric will always be 256, rendering the metrics “equal cost” to EIGRP for the purpose of path selection. This would result in an undesirable routing decision in most scenarios.

Interface Delay Scaling Issues

For end-to-end paths, in which the bandwidth is constant, the use of the scaled bandwidth has sufficient resolution to allow for differentiation up to 10 gigabits, unfortunately no one deploys networks with end-to-end 10 gigabit links. Most paths traverse links considerably slower in the distribution and/or access layers, and effectively remove the scaled bandwidth as a useful path differentiator.

The path through the ten Gigabit Ethernet 10.4.4.0/24 has a minimum bandwidth of 100,000 kbps, and a total delay of 120 microseconds, as does the path through the Gigabit Ethernet 10.5.5.0/24.

For both paths:

- $BW = 10,000,000 / 100,000 = 100$
- $DLY = 120 / 10 = 12$
- $metric = (100 + 12) * 256 = 28,672$

Router 1 shares the load equally over both links. Again, this can be verified with the command “*show eigrp address-family ipv4 topology*”.

```

Router1#show eigrp address-family ipv4 topology 10.1.1.0/24
EIGRP-IPv4(AS 1): Topology entry for 10.1.1.0/24
  State is Passive, Query origin flag is 1, 2 Successor(s), FD is 28672
  Routing Descriptor Blocks:
  10.4.4.2 (TenGigabitEthernet2/0), from 10.4.4.2, Send flag is 0x0
  Composite metric is (28672/28416), Route is Internal
  Vector metric:
    Minimum bandwidth is 100000 Kbit
    Total delay is 120 microseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 2
  10.5.5.3 (GigabitEthernet3/0), from 10.5.5.3, Send flag is 0x0
  Composite metric is (28672/28416), Route is Internal
  Vector metric:
    Minimum bandwidth is 100000 Kbit
    Total delay is 120 microseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 2

```

Router 1 will not prefer the path through the link 10.4.4.0/24 due to EIGRP's use of "minimum bandwidth" along the path, the bandwidth advantage of a single link will hardly be leveraged. If the ten Gigabit Ethernet interface were to report a smaller delay, router 1 would apparently prefer that path.

Migrating Your Network to Wide Metric Support

To address these issues, a new metric formula has been developed to provide the necessary resolution. The EIGRP Wide Metric support is on by default, and backward compatible with existing routes. But consideration should be taken, as there are a few subtle changes to take note of.

Impact on Configuration Commands

EIGRP supports the following existing commands listed below. These commands are used to control the vector metrics for external route sources injecting routes into EIGRP, where the delay component must be able to provide for increased scale. The set of commands affected are:

- delay (interface)
- default-metric
- redistribute metric
- set metric (route-map)
- metric weights
- summary-metric

These commands have not been converted to allow for the configuration of higher precision values. Specifically, the delay values are interpreted in units of 10s of microseconds.

Network Designers Note:

EIGRP supports the ability to match the metric of incoming and outgoing routes learned via the “offset-list” and “match metric” Route Map commands. With the introduction of wide metrics, the computed values will change, requiring the “match metric” configuration to be changed to use the new metric values.

Impact on Routing—Link Utilization

Impact of Scaling Between Old and New Routers

The neighbor discovery/recovery mechanism enables EIGRP to dynamically learn about other routers, and their capabilities, through the use of peering version. With the introduction of wide metrics, the TLV packets were updated to carry additional information.

The change to the TLV format meant the peer version was incremented to 2.0, and due to the additional TLV information, older software versions were not compatible. To address this, EIGRP was modified to send both the older packet formats and the newer ones in a mixed network. Sending twice the data will impact link utilization and performance.

As EIGRP generally uses very little bandwidth on a link, the impact should be minor for typical networks. However, large-scale DMVPN deployments could be impacted more negatively. To address this, EIGRP is able to detect when ALL routers on a link are using the new packet formats, and will send only one version in homogeneous LAN segments.

Impact on Routing—Metric Scaling

Because the wide metrics feature is designed to allow EIGRP to be completely backward compatible with older routers in the network, LANs with mixed version peers experience some minor conversion effects as routes are exchanged between old and new routers.

If a route is received from an older version peer, the scale values will also be “un-scaled” and according to the formula:

$$\text{UnscaledBandwidth} = \left(\frac{\text{EIGRP_BANDWIDTH} \cdot \text{EIGRP_CLASSIC_SCALE}}{\text{ScaledBandwidth}} \right)$$

The conversion of the scaled bandwidth to throughput will result in resolution loss if destinations pass through a mixture of routers. As a result, a destination learned through newer peers will always look better than older peers. Loops are avoided because the metric converted could never be better than the value learned through all peers supporting wide metrics.


Likewise, the scaled delay value sent by the older version peer will be “un-scaled” and according to the formula:

$$\text{UnscaledDelay} = \left(\frac{10 \cdot \text{ScaledDelay}}{\text{EIGRP_CLASSIC_SCALE}} \right)$$

Impact on Routing—Link Utilization Changes

Within a network of high-speed links, if older versions of EIGRP are midstream, destinations that pass through the routes will be adversely penalized resulting in poor metric values and suboptimal routing. As an example, consider the network where all links are 10 gigabit:

Figure 2.



For the routers running the older versions, EIGRP will use the interface reporting the delay of 100 milliseconds, and the bandwidth value of 10,000,000 kilobytes/second.

Looking at R3, normally, if all routers were of the same version of code, R3 would use R2 > R1 to reach 10.1.1.0/24.

Path Selection	Min Throughput	Accumulative Latency
R3 > R2 > R1	10,000,000 KB/sec	300 μ s
R3 > R6 > R5 > R4 > R1	10,000,000 KB/sec	500 μ s

With the router running the newer version, they will calculate the interface delay as 10 μ s, and the bandwidth value of 10,000,000. R2 using older code will experience a loss of resolution, thus delay and bandwidth will be reported as:

Path Selection	Min Throughput	Accumulative Latency
R3 > R2 > R1	10,000,000 KB/sec	210 μ s
R3 > R6 > R5 > R4 > R1	10,000,000 KB/sec	50 μ s

Likewise, from R6's perspective, the cost to 10.1.1.0/24 is no longer equal cost, and instead will prefer the route R5 > R4 > R1.

Impact on Routing—Load Balancing

Load balancing is the capability of a router to distribute traffic over all the router network ports that are the same distance from the destination address. EIGRP supports two types of load balancing.

Equal cost path—Applicable when different paths to a destination network report the same routing metric value.

Unequal cost path—Applicable when different paths to a destination network report different routing metric values. The router determines how much traffic to forward along each path, dividing the metric through each path into the largest metric, rounding down to the nearest integer, and using this number as the traffic share count.

With the classic metrics, EIGRP could not distinguish between 1 gigabit and 10 gigabit lines, resulting in traffic being load-shared across both links. With wide metrics, this is no longer the case, and EIGRP will favor the 10 gigabit link.

Network Designers Note:

Upgrading from classic to wide metric could change data flows on routers with links over 1 gigabit. While this is desirable behavior in most cases, if the utilization of the slower links is desirable, the delay values would need to be manually configured in the EIGRP af-interface submode.

Impact on VPNv4 PE/CE Support

EIGRP currently requires redistribution between BGP and EIGRP on the PE routers. To communicate the EIGRP metric information, EIGRP injects extended community information into BGP to enable the remote sites to recreate the original routes.

With the adoption of a wider bandwidth, delay, and variable length fields for multidimensional routing, this approach is no longer viable. Therefore, the existing EIGRP support for PE/CE **will not support** wide metrics, and routes learned through PE/CE will have the “scaled metric” extended attributed attached to the advertisement.

SNMP

The managed objects of all EIGRP processes will be implemented as five object groups or tables on a per-AS and per-VPN basis:

- EIGRP VPN table
- EIGRP traffic statistics
- EIGRP topology data
- EIGRP neighbor data
- EIGRP interface data

The topology group is the only object that describes EIGRP routes within an AS. While it does not report the vector metrics, it reports back an **eigrpDistance** object and **eigrpReportDistance** object. Only the existing 32 bit scaled numbers will be available through the SNMP MIBS.

Impact on MANET

With the VMI interface, the quality of the connection to a neighbor will vary based on various characteristics computed dynamically based on the L2's feedback to L3. These values will map to the basic EIGRP interface parameters according to the mapping:

VMI	EIGRP	Conversion
current data rate	bandwidth	
relative link quality resources	reliability	
latency	delay	Converted to picoseconds
load	load	

Wide Metric EIGRP Composite Metric Formula

EIGRP historically has used five vector metrics: minimum throughput, latency, load, reliability, and Maximum Transmission Unit (MTU).

These values are accumulated from destination to source as follows:

- Throughput—minimum value
- Latency—accumulative
- Load—maximum
- Reliability—minimum
- MTU—minimum
- Hop count—accumulative

To this, there are two additional values being added: jitter and energy. These two new values are accumulated from destination to source:

- Jitter—accumulative
- Energy—accumulative

These extended attributes, as well as any future ones, will be controlled through K6. If K6 is non-zero, these will be an additive to the path's composite metric. Higher jitter or higher energy usage will result in paths which are worse than those paths that either do not monitor these attributes, or which have lower values. EIGRP will not send these attributes if the router does not provide them. If the attributes are received, then EIGRP will use them in the metric calculation (based on K6) and will forward them with that router's values, assumed to be "zero", and the accumulative values will be forwarded unchanged.

Of these vector metric components, by default, only minimum throughput and latency are traditionally used to compute the best path. Unlike most metrics, minimum throughput is set to the minimum value of the **entire path**, and it does not reflect how many hops or low throughput links are in the path, nor does it reflect the availability of parallel links. Latency is calculated based on one-way delays, and is a cumulative value, which increases with each segment in the path.

Network Designers Note:

When manually influencing EIGRP path selection through interface configuration, it is recommended to use delay. The modification of bandwidth is **discouraged** for following reasons:

1. The change will only effect the path selection if the configured value is the lowest bandwidth over the entire path.
2. Changing the bandwidth can have impact beyond affecting the EIGRP metrics. For example, Quality of Service (QoS) also looks at the bandwidth on an interface.
3. EIGRP throttles to use 50 percent of the configured bandwidth. Lowering the bandwidth can cause problems like starving EIGRP neighbors from getting packets due to throttling.

Changing the delay does not impact other protocols, nor does it cause EIGRP to throttle. And because it is the sum of all delays, delay has a direct effect on path selection.

Metric Formulation

To accommodate high bandwidth interfaces, and to allow EIGRP to perform the path selections listed above, the EIGRP packet and composite metric formula will be modified to choose paths based on the computed time information it takes to travel though the links. It is measured in picoseconds.

EIGRP uses a number of defined constants for conversion and calculation of metric values. These numbers are provided here for reference:

EIGRP_BANDWIDTH	10,000,000
EIGRP_DELAY_PICO	1,000,000
EIGRP_INACCESSIBLE	-1
EIGRP_MAX_HOPS	100

EIGRP_CLASSIC_SCALE	256
EIGRP_WIDE_SCALE	65536
EIGRP_RIB_SCALE	128

When computing the metric using the units above, all capacity information will be normalized to kilobytes and picoseconds before being used. For example, delay is expressed in microseconds per kilobyte, and would be converted to kilobytes per second; likewise energy would be expressed in power per kilobytes per second of usage.

Bandwidth Formulation

Interfaces currently report bandwidth in units of kilobytes/second. This value will be used as is, and no conversion of this value will be performed when encoding in packets being sent between peers.

Delay Formulation

For IOS interfaces that do not exceed 1 gigabit, this value will be derived from the reported interface delay, converted to picoseconds. This includes delay values present in configuration-based commands (i.e. interface delay, redistribute, default-metric, route-map, etc.)

$$\text{Delay} = \text{InterfaceDelay} \cdot \text{EIGRP_DELAY_PICO}$$

Beyond 1 gigabit, IOS does not report delays properly, therefore a computed delay value will be used.

$$\text{Delay} = \left(\frac{\text{EIGRP_BANDWIDTH} \cdot \text{EIGRP_DELAY_PICO}}{\text{InterfaceBandwidth}} \right)$$

This formula also applies to the loopback interfaces that are reported as 8G by IOS.

Note that if the BW is manually set we use the classic formula, not the scaled formula. When the bandwidth is configured by the user, it is expected that the delay should also be configured by the user and does not derive/calculate delay from the user defined interface bandwidth. Instead it falls back to the classic way of calculating delay which converts the delay picked up from the interface into picoseconds.

Network Designers Note:

1. In all cases, if a user configures a delay value for an interface, the configured user value will be used, and not calculated from the interface's bandwidth.
2. The interface delay value must be multiplied by a factor of 10 to convert from the IDB value to microseconds. This conversion is not shown in this paper for simplicity.

Wide Metric Formulation

$$\left[\left(K_1 \cdot \text{Min Throughput} + \left\{ \frac{K_2 \cdot \text{Min Throughput}}{256 - \text{Load}} \right\} \right) + K_3 \cdot \text{Total Latency} \right] + K_6 \cdot \text{Ext Attrib} \cdot \left[\frac{K_5}{K_4 + \text{Reliability}} \right]$$

By default, the path selection scheme used by EIGRP is a combination of throughput and latency, where the selection is a product of **total** latency and **minimum** throughput of all links along the path:

$$\text{metric} = K_1 \cdot \text{min Throughput} + K_3 \cdot \sum \text{Latency}$$

IOS RIB Metric and Scaling

With the calculation of larger bandwidths, EIGRP can no longer fit the value needed by the IOS RIB. In order to scale the values to fit into the limited 32bit value, a scale constant will be applied according to the conversion formula:

$$\text{RIBMetric} = \left(\frac{\text{Wide Metric}}{\text{EIGRP_RIB_SCALE}} \right)$$

The default conversion value for **EIGRP_RIB_SCALE** is defined to allow the full range of IOS interfaces from 9K though 4.295 terabit to fit into the RIB.

Network Designers Note:

The **EIGRP_RIB_SCALE** value has only **LOCAL significance**, as EIGRP does not advertise the composite value. This allows routers with "higher speed links" to use a smaller scaling, favoring the resolution of faster links at the sacrifice of slower links. The use of differing scales on different routers at redistribution points would result in the "external data" showing different values, but these values are not used for routing decisions and are deemed cosmetic.

Usage and Definition of K₁-K₆

Throughput: Usage and Definition of K₁ and K₂

K₁ is used to allow **throughput-based path selection**: EIGRP can use one of two variations of throughput based path selection:

- **Maximum Theoretical Bandwidth**: paths chosen based on the highest reported bandwidth
- **Network Throughput**: paths chosen based on the highest **'available'** bandwidth adjusted by congestion-based effects (interface reported load)

By default, EIGRP computes throughput using the maximum theoretical throughput expressed in picoseconds per kilobyte of data sent. **This inversion results in a larger number (more time), ultimately generating a worse metric.** The formula for the conversion for max-throughput value directly from the interface without consideration of congestion-based effects is as follows:

$$\text{Max-Throughput} = \left(K_1 \cdot \frac{\text{EIGRP_BANDWIDTH} \cdot \text{EIGRP_WIDE_SCALE}}{\text{Bandwidth}} \right)$$

If **K₂** is used, the effect of congestion, as a measure of load reported by the interface, will be used to simulate the available throughput, by adjusting the maximum throughput according to the formula:

$$\text{Net Throughput} = \left[\text{Max-Throughput} + \left(\frac{K_2 \cdot \text{Max-Throughput}}{256 - \text{Load}} \right) \right]$$

Latency: Usage and Definition of K₃

K₃ is used to allow latency-based path selection. Latency and delay are similar terms that refer to the amount of time it takes a bit to be transmitted to an adjacent peer. EIGRP uses one-way based latency values provided either by IOS interfaces or computed as a factor of the links bandwidth.

$$\text{Latency} = \left(K_3 \cdot \frac{\text{Delay} \cdot \text{EIGRP_WIDE_SCALE}}{\text{EIGRP_DELAY_PICO}} \right)$$

Reliability: Usage and Definition of K₄ and K₅

K₄ and **K₅** are used to allow for **packet loss-based path selection** for networks where link quality and packet loss is critical. Packet loss caused by network problems result in highly noticeable performance issues or jitter with streaming technologies, voice over IP, online gaming and videoconferencing, and will affect all other network applications to one degree or another.

Critical services should pass with less than 1% packet loss. Lower priority packet types might pass with less than 5% and then 10% for the lowest of priority of services. The final metric can be weighted based on the reported link quality according to the adjustment:

$$\text{Reliability Quotient} = \left(\frac{K_5}{\text{Reliability} + K_4} \right)$$


Setting **K₅** values up to 255 scales the impact of reliability quotient on the final metric.

Extended Attributes: Usage and Definition of K₆

K₆ allows for **extended attributes**, which can be used for higher aggregate metrics than those having lower energy usage. Currently there are two extended attributes, **jitter** and **energy**, which are not considered for this paper.

Example Wide Metric Computation

To find the metric when using EIGRP and the new wide metric formula, let's look at this earlier diagram again:



Composite Metric for Connected Interface

Using the bandwidth and delay values reported from the interfaces, the metrics can be calculated for a connected route:

Gigabit Ethernet	10 Gigabit Ethernet
Delay = InterfaceDelay • EIGRP_DELAY_PICO	Delay = $\frac{\text{EIGRP_BANDWIDTH} \cdot \text{EIGRP_DELAY_PICO}}{\text{InterfaceBandwidth}}$
Delay = 10 • 1,000,000	Delay = $\frac{10,000,000 \cdot 1,000,000}{10,000,000 \text{ Kbit/sec}}$
Delay = 10,000,000 picoseconds	Delay = 1,000,000 picoseconds

Now with the computed delay, the EIGRP latency and throughput values can be calculated:

Gigabit Ethernet	10 Gigabit Ethernet
Latency = $\frac{\text{Delay} \cdot \text{EIGRP_WIDE_SCALE}}{\text{EIGRP_DELAY_PICO}}$	Latency = $\frac{\text{Delay} \cdot \text{EIGRP_WIDE_SCALE}}{\text{EIGRP_DELAY_PICO}}$
Latency = $\frac{10,000,000 \cdot 65,536}{1,000,000}$	Latency = $\frac{1,000,000 \cdot 65,536}{1,000,000}$
Latency = 655,360	Latency = 65,536

Gigabit Ethernet	10 Gigabit Ethernet
Throughput = $\frac{\text{EIGRP_BANDWIDTH} \cdot \text{EIGRP_WIDE_SCALE}}{\text{InterfaceBandwidth}}$	Throughput = $\frac{\text{EIGRP_BANDWIDTH} \cdot \text{EIGRP_WIDE_SCALE}}{\text{InterfaceBandwidth}}$
Throughput = $\frac{10,000,000 \cdot 65,536}{1,000,000 \text{ Kbit/sec}}$	Throughput = $\frac{10,000,000 \cdot 65,536}{10,000,000 \text{ Kbit/sec}}$
Throughput = 655,360	Throughput = 65,536

If we want to see the new composite metric for each path, and default coefficients ($k1=k3=1, k2=k4=k5=k6$), the formula reduces to:

$$\text{metric} = \min(\text{Throughput}) + \sum \text{Latency}$$

Gigabit Ethernet	10 Gigabit Ethernet
metric = 65,536 + 655,360	metric = 65,536 + 65,536
metric = 720,896	metric = 131,072

Composite Metric for Paths

Looking at the path metrics, the accumulative delay is used to compute latency and the minimum bandwidth for throughput.

Router1 sees the cost through Router2, Router4 as:

Minimum bandwidth is 10,000,000 Kbit
 Total delay is 3,000,000 picoseconds

The metric is computed as:

$$\text{Metric} = \frac{3,000,000 \cdot \text{EIGRP_WIDE_SCALE}}{\text{EIGRP_DELAY_PICO}} + \frac{\text{EIGRP_BANDWIDTH} \cdot \text{EIGRP_WIDE_SCALE}}{10,000,000}$$

Metric = 262,144

Router1 sees the cost through Router3, Router4 as:

Minimum bandwidth is 1,000,000 Kbit
 Total delay is 21,000,000 picoseconds

$$\text{Metric} = \frac{21,000,000 \cdot \text{EIGRP_WIDE_SCALE}}{\text{EIGRP_DELAY_PICO}} + \frac{\text{EIGRP_BANDWIDTH} \cdot \text{EIGRP_WIDE_SCALE}}{10,000,000}$$
$$\text{Metric} = 2,031,616$$

EIGRP Metric vs. RIB Metric

The IOS RIB can only support 32 bit metric values. Because of this, EIGRP and RIB will have different numbers based approximately on the **EIGRP_RIB_SCALE** value. The output of the show command will now display the RIB value, as well as the EIGRP value:

```
Router# show eigrp address-family ipv4 topology
EIGRP-IPv4 VR(WideMetric) Topology Entry for AS(4453)/ID(3.3.3.3) for
100.1.0.0/16
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 262144, RIB
  is 2048
  Descriptor Blocks:
  2.0.0.2 (Ethernet0/2), from 2.0.0.2, Send flag is 0x0
    Composite metric is (262144/196608), route is Internal
  Vector metric:
    Minimum bandwidth is 10000000 Kbit
    Total delay is 3000000 picoseconds
    Reliability is 255/255
    Load is 1/255
    Minimum MTU is 1500
    Hop count is 2
    Originating router is 100.1.1.1
```



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)