

Determining How Information Is Exchanged with EIGRP



Sean Wilkins

NETWORK ENGINEER AND AUTHOR

@Sean_R_Wilkins www.infodispersion.com



Module Overview



Module Overview



**EIGRP Route States, Split Horizon,
and Poison Reverse**



Module Overview



**EIGRP Route States, Split Horizon,
and Poison Reverse**

EIGRP Message/Packet Types



Module Overview



**EIGRP Route States, Split Horizon,
and Poison Reverse**

EIGRP Message/Packet Types


Common EIGRP TLVs



Let's move deeper into
EIGRP communications



EIGRP Route States



Passive



Active



Passive State

Successor exists



Passive State

**Loop free path
is available**



Passive State

**Indication of a
stable network**





Active State





Active State

**Previous successor path lost
without feasible successor**





Active State

**Previous successor path lost
without feasible successor**

Indicates active search for replacement



Split Horizon & Poison Reverse



Split Horizon & Poison Reverse

Rules used for all situation types



Split Horizon & Poison Reverse

Rules used for all situation types

Let's go into more detail



Split Horizon

Implemented in DV protocols



Split Horizon

"Never advertise a route out of the interface through which it was learned"



Split Horizon

“If you have a successor route to a destination, never advertise the route out the interface on which it was learned”



Poison Reverse

Often implemented
with split horizon



Poison Reverse

**"Poisons" routes
out learned
interface**



Poison Reverse

EIGRP uses infinite delay metric



Also used when device
wants to advertise
route unreachability





Let's tie back to the
previous EIGRP packet section



Let's tie back to the
previous EIGRP packet section

Opcode used to indicate
message type inside packet




Let's tie back to the
previous EIGRP packet section

Opcode used to indicate
message type inside packet


Message type usually
used for packet type





Opcodes






Opcodes

- Update (1)






Opcodes

- Update (1)
- Request (2)






Opcodes

- Update (1)
- Request (2)
- Query (3)






Opcodes

- Update (1)
- Request (2)
- Query (3)
- Reply (4)






Opcodes

- Update (1)
- Request (2)
- Query (3)
- Reply (4)
- Hello (5)






Opcodes

- Update (1)
- Request (2)
- Query (3)
- Reply (4)
- Hello (5)
- SIA-query (10)





Opcodes

- Update (1)
- Request (2)
- Query (3)
- Reply (4)
- Hello (5)
- SIA-query (10)
- SIA-reply (11)



Hello Packet



Opcodes will be covered in common order



Hello Packet



Opcodes will be covered in common order




Used primarily for neighbor discovery and keepalive



Hello Packet



Opcodes will be covered in common order



Used primarily for neighbor discovery and keepalive




Usually sent out every 5 seconds



Hello Packet



Opcodes will be covered in common order



Used primarily for neighbor discovery and keepalive



Usually sent out every 5 seconds



Includes k-values and hold timer



Hello Packet

Often sent using
multicast



Hello Packet

**Unicast also
supported**




Hello Packet

**Don't require
acknowledgment**




Hello Packet



Hello Packet


Also used for acknowledgment



Hello Packet

Also used for acknowledgment

Acknowledgment packet is empty with matching acknowledgment number





Update Packet





Update Packet

Relays new or updated information





Update Packet

Relays new or updated information

Use unicast and multicast





Update Packet

Relays new or updated information

Use unicast and multicast

Require acknowledgment



Update Packet - Flag Use

INIT

Conditionally received (CR)

Restart (RS)

End of table (EOT)



INIT Flag

Used in initial neighbor communications



INIT Flag

Used in initial neighbor communications




Ensure bidirectional
communication



INIT Flag

Used in initial neighbor communications




Ensure bidirectional
communication



Instruct neighbor to
advertise known
routes




Conditionally Received (CR)



Conditionally Received (CR)


Used when packet
acknowledgment problems exist



Conditionally Received (CR)

Used when packet acknowledgment problems exist

Commonly used when slow or congested links exist





Conditionally Received (CR)

Used when packet acknowledgment problems exist

Commonly used when slow or congested links exist

Often results in EIGRP recovery using unicast





Restart (RS) Flag





Restart (RS) Flag

Often used in two scenarios:

- Soft resets





Restart (RS) Flag

Often used in two scenarios:

- Soft resets
- Graceful restarts



Restart (RS) Flag

**Adjacency
remains in both
scenarios**



Restart (RS) Flag

Typically sent
along with INIT
flag




Restart (RS) Flag

**Reduces packet
exchange**




End of Table (EOT) Flag




Indicates all routes have been sent



End of Table (EOT) Flag



Indicates all routes have been sent



If used with RS flag,
old entries are removed





Query Packet





Query Packet

Used when a route goes active





Query Packet

Used when a route goes active

Sent to all available other neighbors





Query Packet

Used when a route goes active

Sent to all available other neighbors

Two functions:

- Requests alternate path





Query Packet

Used when a route goes active

Sent to all available other neighbors

Two functions:

- Requests alternate path
- Poisons former path



The use of multiple query packets is possible for a single event



Query Packets



Query Packets

Require acknowledgment



Query Packets

Require acknowledgment

**Common acknowledgment using
empty hello**



Query packets can use
unicast or multicast



Reply Packet

Used with query
packet



Reply Packet

**Always sent as
response to query**



Reply Packet


**Unknown routes
are returned with
infinite delay
metric**




TLVs are common to
update, query and reply
packets




Query Process




Query Process




Query Process



Query Process



Query Process



Query Process

Problematic on large
networks



Query Process

Multiple options exist to limit scope



Query Process


Route
summarization

Stub routing

Route filtering




Query Process



Sometimes
replies are
delayed or lost



Query Process




Sometimes
replies are
delayed or lost



Can cause stuck
in active
condition



Query Process



Sometimes replies are delayed or lost




Can cause stuck in active condition



Caused by lack of reply packet



Query Process



Sometimes replies are delayed or lost



Can cause stuck in active condition





Caused by lack of reply packet



Can result in neighborhood reset








Latest Cisco versions implement modified behavior






Latest Cisco versions implement modified behavior

Two additional messages added:

- SIA-query
- SIA-reply




SIA-query sent at $\frac{1}{2}$ query timer



SIA-query sent at $\frac{1}{2}$ query timer


SIA-reply sent if neighbor still waiting




SIA-query sent at $\frac{1}{2}$ query timer

SIA-reply sent if neighbor still waiting


Neighbor will send SIA-reply or reply
when complete





SIA-query transmission resets query timer
to $\frac{1}{2}$ initial value






SIA-query transmission resets query timer
to $\frac{1}{2}$ initial value

Three SIA-queries can be sent






SIA-query transmission resets query timer to $\frac{1}{2}$ initial value

Three SIA-queries can be sent

Route considered stuck afterwards





SIA-query transmission resets query timer to $\frac{1}{2}$ initial value

Three SIA-queries can be sent

Route considered stuck afterwards

Will typically result in adjacency reset



Request Packet



Rarely used



Request Packet



Rarely used



Used to obtain addition information



Finishing Touches

Process is
completed per
route



Finishing Touches

Single packet can
relay multiple
route entries



Finishing Touches

Update, query,
reply, SIA-query,
SIA-reply are
reliable



Now let's cover common TLVs





Hello Packet





Hello Packet

Two TLVs used typically





Hello Packet

Two TLVs used typically

Including:

- Parameter (0x0001)





Hello Packet


Two TLVs used typically

Including:

- Parameter (0x0001)
- Software version (0x0004)




Parameter TLV



Parameter TLV


Includes:
Local K-values



Parameter TLV

Includes:

Local K-values
Hold timer value




Parameter TLV

Includes:

Local K-values
Hold timer value

Used to determine compatibility



Software Version TLV




Used to ensure
supportable EIGRP
version



Software Version TLV



Used to ensure
supportable EIGRP
version




TLV formats are
different between
versions




Software Version TLV



Used to ensure
supportable EIGRP
version



TLV formats are
different between
versions



Version 2.0 devices
are backwards
compatible



Authentication TLV


EIGRP supports:

MD5

SHA-256




Authentication TLV



Authentication TLV


Type 0x0002



Authentication TLV

Type 0x0002

Used to relay
authentication information



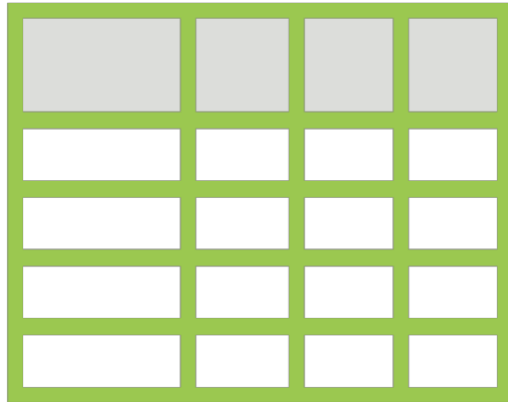
When configured, used in
all packet types except
acknowledgment





IPv4 & IPv6 TLVs





IPv4 & IPv6 TLVs

Included in:

- Update
- Query
- Reply
- SIA-query
- SIA-reply



IPv4 TLVs

Multiple TLVs can be used



IPv4 TLVs

Originally IPv4 advertised in
two types



IPv4 TLVs

0x0102 (Internal)

0x0103 (External)



IPv6 TLVs




Same is true with IPv6



IPv6 TLVs




Same is true with IPv6



0x0402

(Internal)




0x0403

External




Cisco used these TLVs in
older software versions






Older TLVs deprecated
in favor of multi-protocol TLVs



Older TLVs deprecated
in favor of multi-protocol TLVs


0x0602 - Used for
IPv4 & IPv6 internal routes



Older TLVs deprecated
in favor of multi-protocol TLVs

0x0602 - Used for
IPv4 & IPv6 internal routes

0x0603 - Used for
IPv4 & IPv6 external routes



Let's now move into the lab



Summary



Summary



**EIGRP Route States, Split Horizon
and Poison Reverse**



Summary



**EIGRP Route States, Split Horizon
and Poison Reverse**

EIGRP Message/Packet Types



Summary



**EIGRP Route States, Split Horizon
and Poison Reverse**

EIGRP Message/Packet Types

Common EIGRP TLVs

