NYC Utility

EIGRP Migration White Paper

August 2001

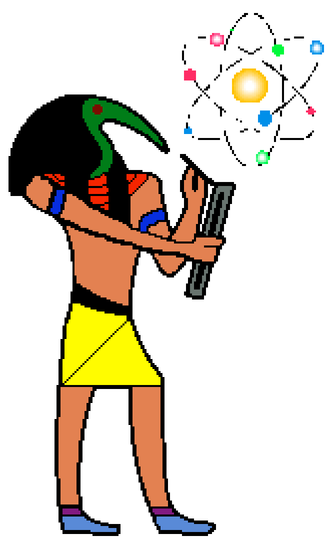

Applied Methodologies, Inc.

Table Of Contents

1.0 Introduction

This document is a working draft and its purpose is to provide NYC Utility a road map and to document the options and issues relating to migrating their current RIP based network to EIGRP. NYC Utility already has some standards in place regarding the use of its SONET Core for enterprise route distribution and path availability between major sites.

As was mentioned in a previous document submitted to NYC Utility titled "Cursory Network Review" in May of 2001, the RIP routing protocol is outliving its capabilities within NYC Utility. The RIP protocol, although very efficient for smaller networks of 20 routers or less, starts to show many of its limitations when it comes to VLSM, summarization and route path convergence. NYC Utility is starting to experience issues such as sub optimal and asymmetrical routing paths, higher hop counts from outages and mis-configurations. Also, fault tolerant solutions based on RIP provide a slower convergence time for redundant paths to be used, thus causing sub optimal fail over results. Since RIP is a classfull protocol NYC Utility cannot summarize or subnet its networks to effectively control routing table sizes and conserve address space.

EIGRP is considered a modern distance vector routing protocol that combines the ease of distance vector type configuration options with convergence and load balancing capabilities that rival some link state protocols. Every large network deploying any interior routing protocol requires careful design, backed up with a solid addressing scheme that allows efficient summarization, and networks implemented with EIGRP are no exception. Although EIGRP is sometimes presented as a routing protocol that requires no network design, that is not true in reality.

It is assumed that the reader understands basic routing protocol theory and knows RIP plus EIGRP and its terminology such as: SUCCESSORS, FEASIABLE SUCCORS, DUAL, ADVERTISED DISTANCE, FEASIABLE DISTANCE and FEASIBILITY CONDITION. A general understanding of Cisco's implementation of routing protocol redistribution, load balancing and path manipulation options is helpful as well. Knowledge of EIGRP's show commands for troubleshooting is also assumed.

NYC Utility's routed network was built using the RIP routing protocol and RIP is present in every router today. NYC Utility has also deployed EIGRP in it's Core and Distribution routers in most major sites. It should be noted that wherever EIGRP is currently active RIP too is active. EIGRP routes are redistributed into RIP with a metric of one. This is one of the major sources of the sub optimal routing issues.

NYC Utility would like to expand their current EIGRP implementation out to the edges of the entire enterprise to improve routing path selection, usage, increase network stability, improve convergence times and offer improved tuning capabilities. NYC Utility would also like to remove RIP completely or leave it where it is absolutely necessary, and provide its routed network a means to select the most optimal routes during normal or outage conditions. NYC Utility's router network has grown to the point where its requirements and needs of a modern routing protocol are necessary to utilize tuning options of vector and composite metrics so the influencing of data paths utilizing a plethora of medium from Ethernet to SONET can be achieved when applicable.

1.1 EIGRP Migration Design Considerations

After careful research the following design considerations were outlined:

1. **Needs to have faster more reliable convergence**
2. **Reduce and/or eliminate current sub optimal routing conditions**
3. **Can be applied/overlapped to NYC Utility's current physical infrastructure**
4. **Low impact and easy to migrate to and from RIP**
5. **Supports full connectivity to the internet and all enterprise networks**
6. **Simple configurations and scalability options used**
7. **Scalable to edges of enterprises**
8. **Scalable for future growth in site backbones with RSM/MSFC or additional sites added**
9. **Easier to troubleshoot and architect with**
10. **Provides summarization and VLSM capabilities**
11. **Provides advanced tuning and route manipulation options**
12. **Move NYC Utility's routed network to one major routing protocol with very limited to no redistribution necessary**
13. **Have all major sites participate in and utilize the SONET Core and provide some hierarchal level of routing within the NYC Utility routed network.**

EIGRP meets the design considerations listed above. However, to achieve design consideration #4 careful planning and execution must be applied.

Two basic migrations paths have been planned and are outlined below, even though many variations of the same approach can be reared, these two are the most prevalent in terms of starting points. The abstract diagram of NYC Utility's current EIGRP and RIP deployment is depicted below. These abstraction diagrams refer to a major site like 4IP, TLC, 30 Flatbush or Rye. Most of the major sites already have EIGRP enabled in Core and Distribution routers and may also be connected by SONET to each other. The same devices in some of these sites perform the Core/Distribution functions, so they were separated here for clarity. The overall goal is to have each major site connected to SONET and just only running EIGRP.

Eventually, EIGRP will be pushed out to the edges of each major site one at a time at a time until all of the sites utilize EIGRP and the SONET properly to access resources from one another. The following abstract diagram depicts a state the NYC Utility network may be in during this transition:

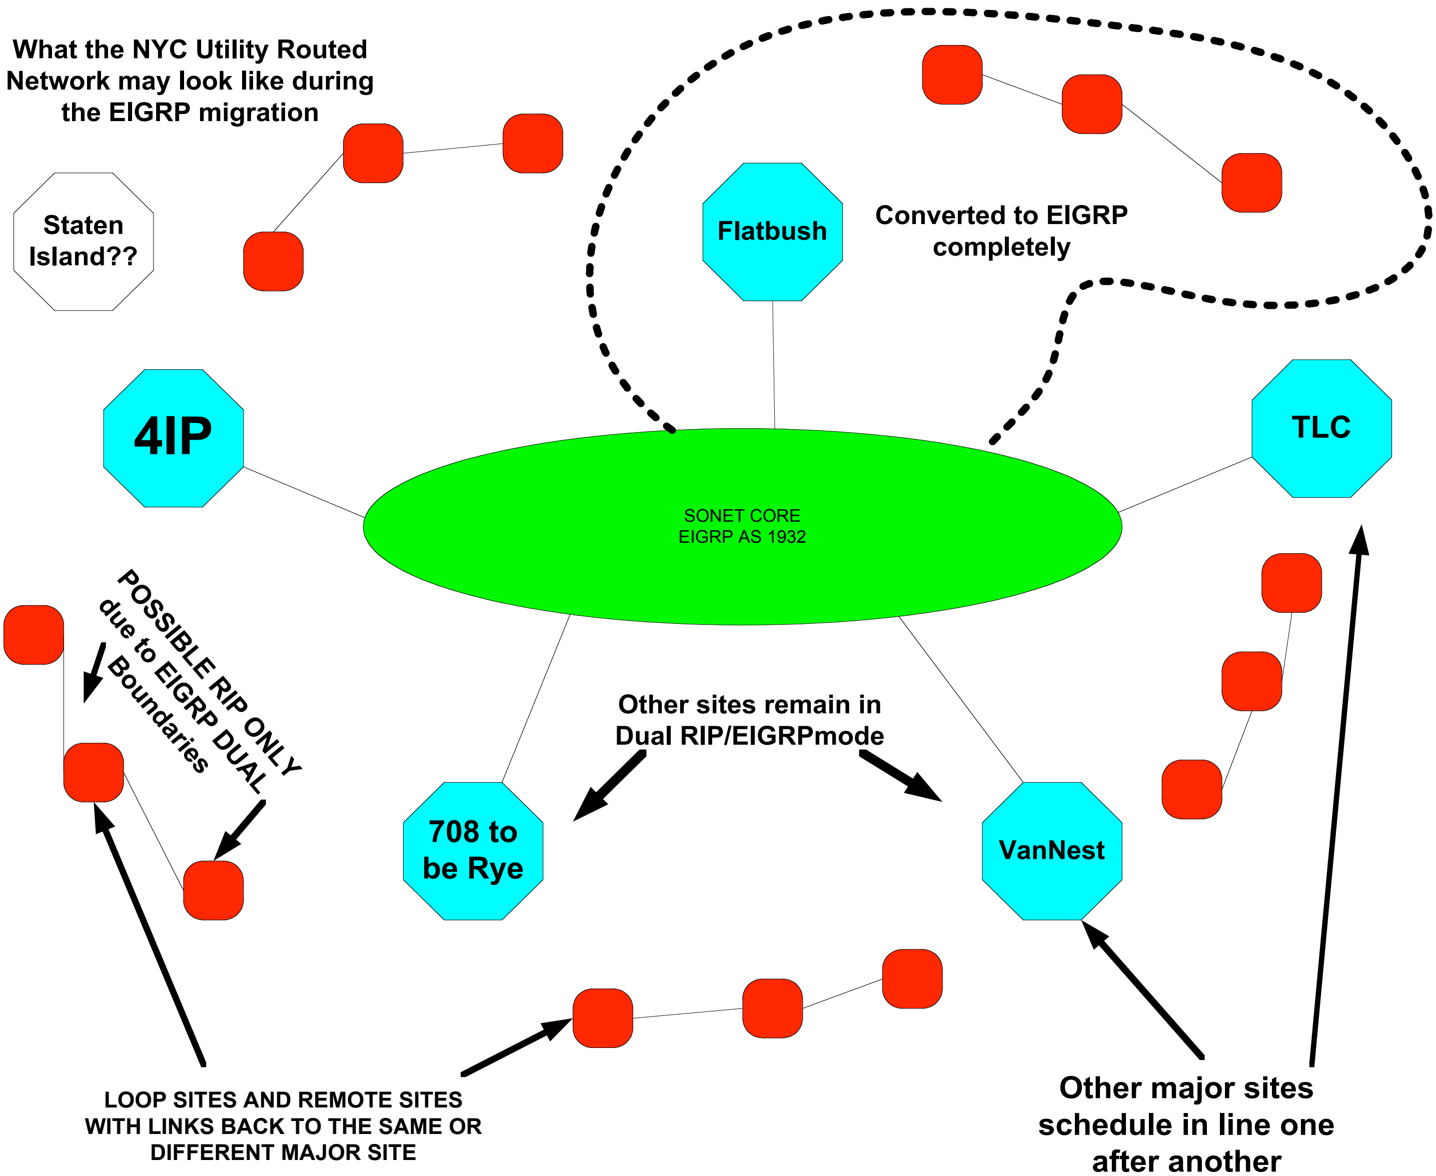

Abstraction of current Routing protocol distribution:

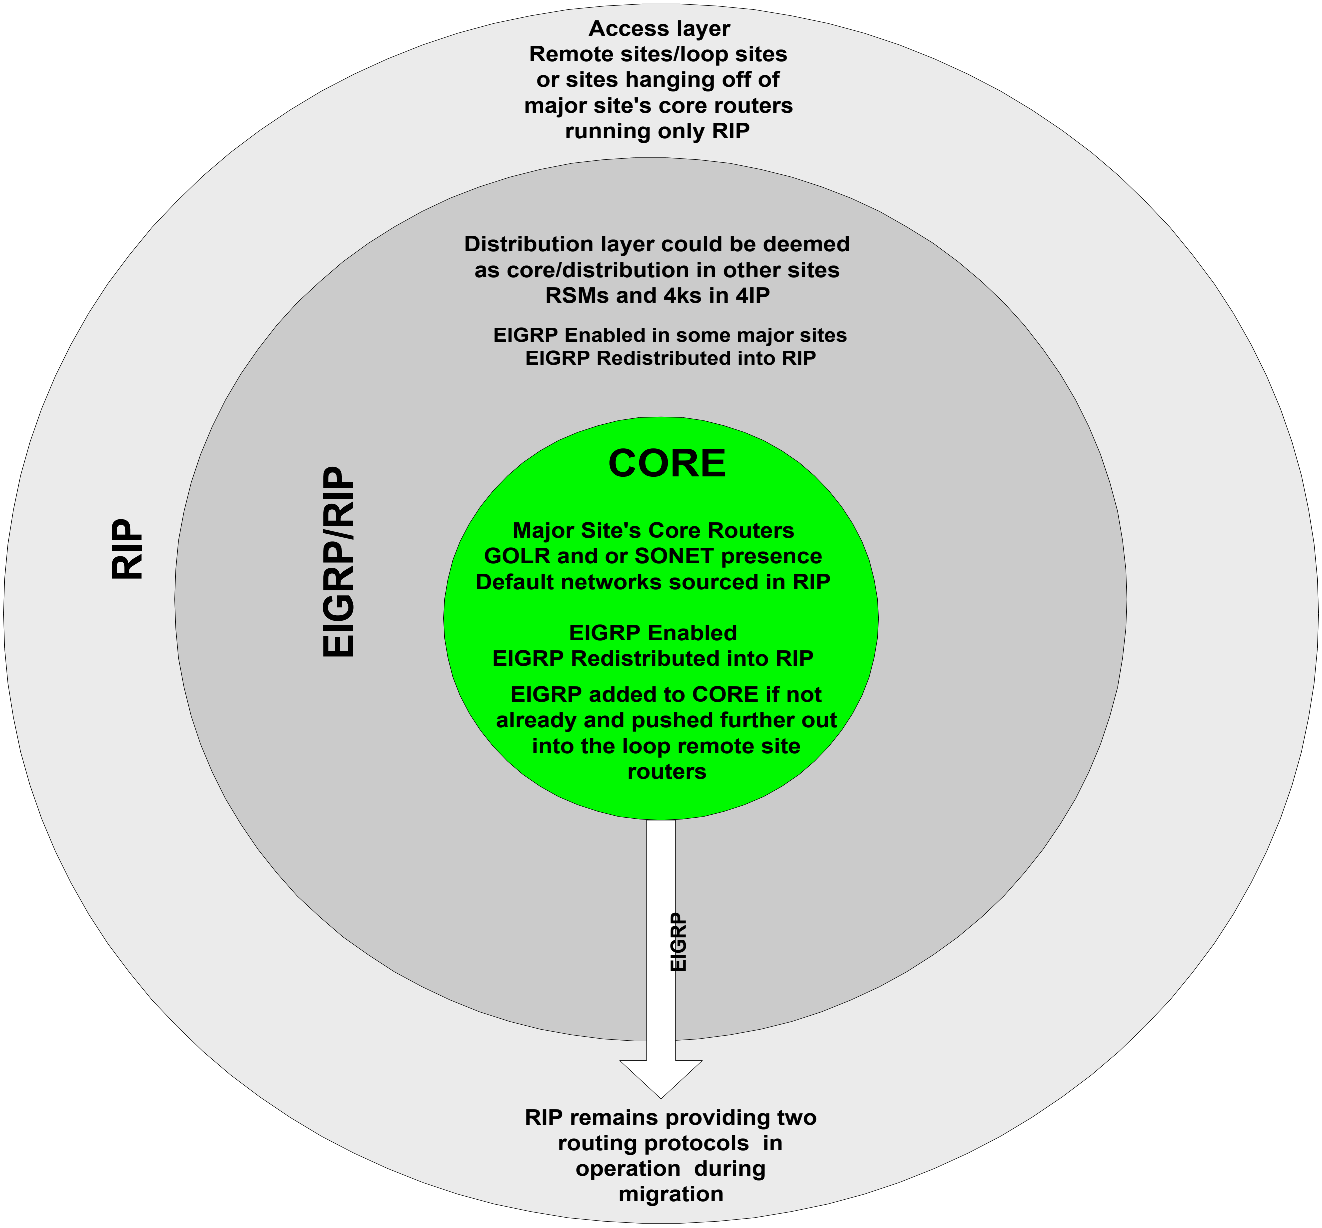

1.2 The dual EIGRP and RIP approach- Migration approach#1

This approach entails pushing EIGRP into all devices(allowing for IOS upgrade or platform capabilities) while leaving RIP activated in same devices. See the following diagram.

Add EIGRP to all Distribution and Loop or edge routing devices incrementally starting from the Core outward while RIP is still active in those same devices. Ships in the Night(SIN) routing can be applied. There is no need to redistribute the RIP routes into EIGRP for the routes are already there via RIP. For EIGRP to obtain it's default route the GOLR must have EIGRP enabled and participate in NYC Utility EIGRP AS of 1932.

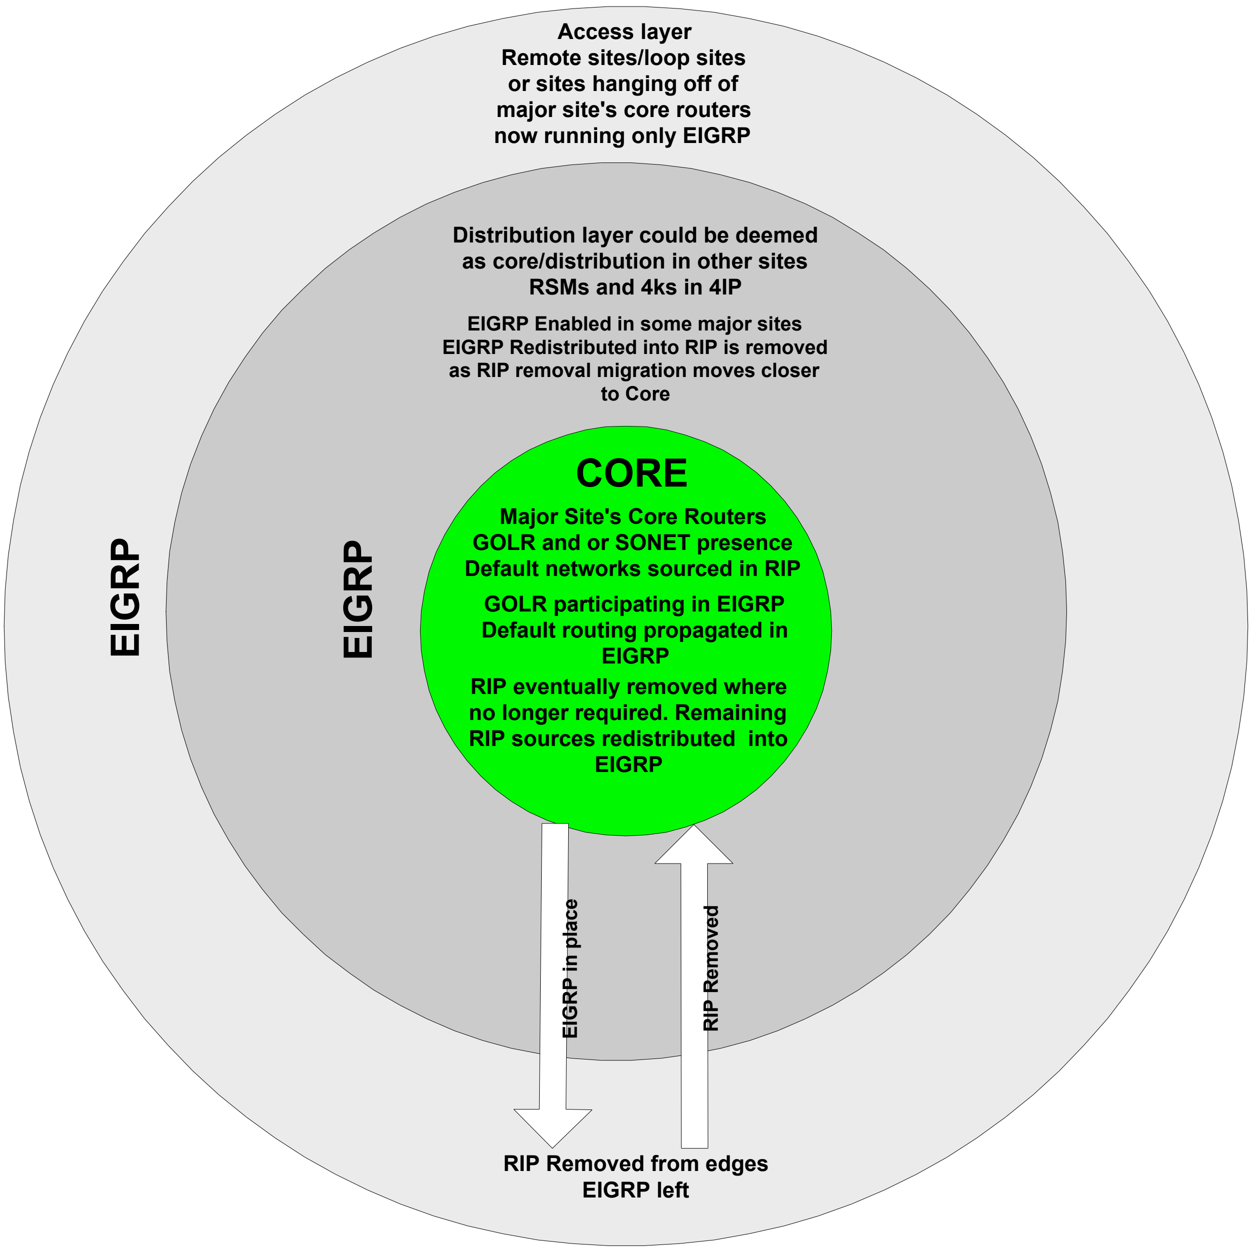

The default routes must be redistributed into EIGRP so each EIGRP speaker will have a candidate default route.

After this is completed on the GOLR, RIP can then be removed from the Loops and edge routers, then the Distribution routers, RSMs et al. at each site until finally removing RIP in the Core routers. The diagram on the previous page (8) depicts an abstract of this approach:

Benefits of this approach:

1. As applied per major site only provides a quick and easy migration path, for both routing protocols are still running.
2. Reduces exposure of an EIGRP issue due to SIN routing.
3. As EIGRP routes are used over RIP routes sub-optimal routing paths will be removed.
4. Easy to rollback just by turning off EIGRP in any device or section of the network causing an issue.

Possible issues with this approach:

1. Increased risk of networking problems due to EIGRP DUAL boundaries such as SIA events originating in Loops and the affects being felt in Core/Distribution routers.
2. Performance of lower scale router platforms ability to handle two routing protocols simultaneously.
3. Segments that have servers or workstations speaking/listening RIP may lose visibility when RIP is removed.
4. Ensuring that Loops and redundant connections are traced so unnecessary RIP/EIGRP borders are not present.
5. Not advisable as per Cisco but mat be unavoidable due to NYC Utility's physical design.
6. Any of the issues listed in section *3.0* ***"Migration Technical Issues"*** of this document can become present.

Summary order of Migration approach #1:

1. Add EIGRP commands to GOLRs
2. Add EIGRP commands to Core and Distribution routers – one at a time and wait a minute between each change.

3. Add EIGRP commands to access layer routers(Loop and remote site routers) no redistribution into RIP required. – one at a time and wait a minute between each change.

At this point EIGRP and RIP are running together in every router and in the loops and remotes sites SIN routing is employed.

4. Remove RIP from each Loop and remote site router towards the major site and verify default route and overall network connectivity.
5. Remove RIP in Distribution and major site Loop edge routers and verify default route and overall network connectivity.
6. Remove RIP in CORE if possible otherwise change redistribution of RIP in Core routers into EIGRP.
7. Control Centers requiring RIP will be handled last and EIGRP will be added but RIP will be redistributed into EIGRP.

Configuration commands required:

Step #1 Pre-change route state output for comparison:

SH IP ROUTE (should only be EIGRP routes only)
SH IP ROUTE RIP(for any rip only routes left)
PING 158.57.x.x (devices in Core)
PING GOLR devices/networks
SH IP ROUTE SUMMMARY
SH IP EIGRP NEIGHBORS
SH IP EIGRP TOP ACTIVE
TRACE to GOLR or other significant end point

Step#2 Add EIGRP commands to GOLRs if not already completed.

router eigrp 1932
  network 158.57.0.0
  network 172.16.0.0 *may not be necessary if staying RIP at other end*
  network 172.28.0.0 *may not be necessary if staying RIP at other end*
  eigrp log-neighbor-changes
  eigrp log-neighbor-warnings
  redistribute static metric 100000 100 255 1 1500
  distribute-list 2 in

RIP commands are untouched or RIP redistributed into EIGRP for the 172.16 and .28 subnets.

Under the EIGRP 1932 process add:

redistribute rip metric 1544 20000 255 1 1500

However, default routing will provide access to those 172.x networks as well. If the redistribution is used then they will appear under EIGRP as external routes.

Step#3 Add EIGRP commands to Core and Distribution routers

router eigrp 1932
network 158.57.0.0
eigrp log-neighbor-changes
eigrp log-neighbor-warnings

***** REDISTRIBUTE STATIC METRIC 100000 100 255 1 1500 *****
To be used if router had static routes configured.

Leave RIP in place in those routers

Step#4 Add EIGRP commands to access layer routers(loop and remote site routers) no redistribution into RIP required
router eigrp 1932
network 158.57.0.0
eigrp log-neighbor-changes
eigrp log-neighbor-warnings

***** REDISTRIBUTE STATIC METRIC 100000 100 255 1 1500 *****
To be used if router had static routes configured.

Leave RIP in place in those routers no redistribution required.

To remove RIP towards the Core:

Step #5 Remove RIP from each loop and remote site router towards the major site and verify default route and overall network connectivity.

no router rip
sh ip route should only be EIGRP routes only
ping 158.57.x.x devices in Core
ping GOLR devices/networks

Step#6 Remove RIP in Distribution and major site loop edge routers and verify default route and overall network connectivity.

no router rip
sh ip route should only be EIGRP routes only
ping 158.57.x.x devices in Core
ping GOLR devices/networks

Step#7 Remove RIP in CORE if possible, otherwise change redistribution of RIP in Core routers into EIGRP.
no router rip – if applicable
no redistribution eigrp 1932 metric 1

router eigrp 1932
redistribute rip metric 1544 20000 255 1 1500 to provide lower metric for those remaining routes -- or can use higher metrics(100000 100 255 1 1500) if sourced on FastEthernet interfaces.

Step# 8 Post migration review of routes

Check and compare information from Step#1

SH IP ROUTE (should only be EIGRP routes only if RIP was totally removed)
SH IP ROUTE RIP(for any RIP only routes left if RIP not removed)
PING 158.57.x.x (devices in Core)
PING GOLR devices/networks
SH IP ROUTE SUMMMARY
SH IP EIGRP NEIGHBORS(compare with Step#1 output for increase or new neighbors)
SH IP EIGRP TOP ACTIVE (for any routes that are marked ACTIVE, should be none)
TRACE to GOLR or other significant end point to see any new or improved paths.

1.3 EIGRP and RIP boundary approach – Migration approach#2

This approach will continue to use RIP and or EIGRP for GOLR default route sourcing but enable EIGRP in all major site networks and keep ONLY RIP in Loop and stub remote sites.

By pushing EIGRP as far into each major site to their respective Loop and redundant link boundaries, for DUAL boundary calculation purposes, removing RIP at those points ensures that that EIGRP DUAL calculation and convergence time dependencies on DUAL queries is mitigated to basically the Core and Distribution of the NYC Utility enterprise. RIP will only be left in the stub and Loop sites. RIP will then be redistributed at stub and Loop boundary router(Access layer to Distribution layer routers) points. This would performed on sites with low bandwidth or platform capabilities issues. The following diagram depicts an abstract of this approach

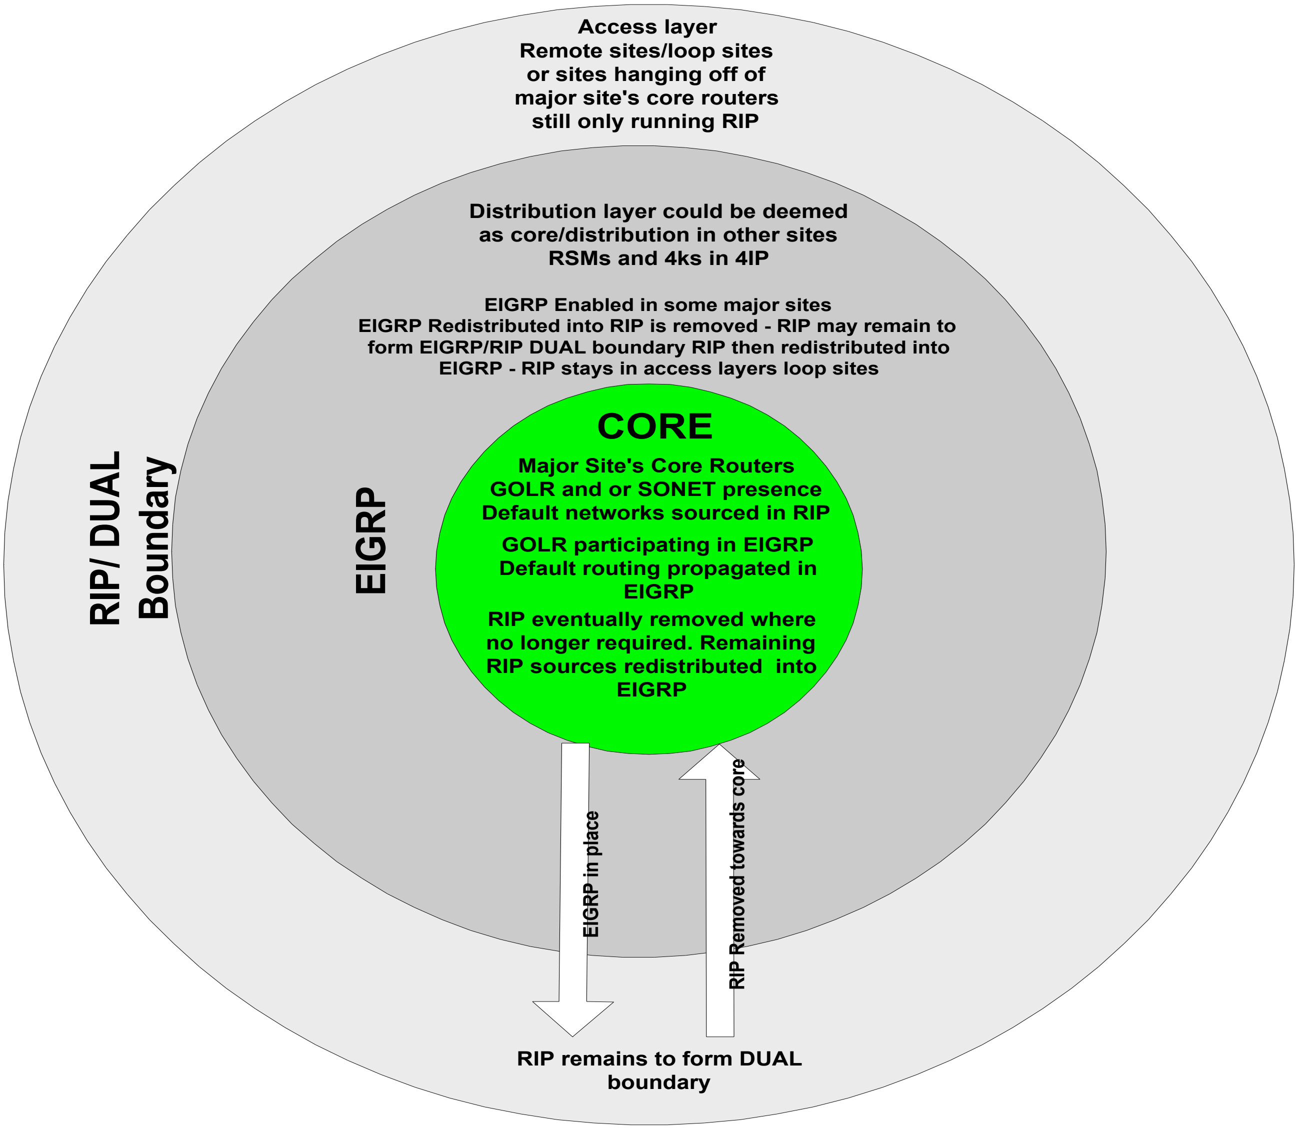

Benefits of this approach:

1. Reduces chances of DUAL SIA activity in Loops causing issues in the Core/Distribution routers.
2. Enable platforms with limited IOS or memory capabilities to remain unchanged or upgraded.
3. RIP listening based devices remain unaffected.

Possible issues with this approach:

1. Increase complexity of configurations due to RIP/EIGRP border demarcation determination and a reversal of redistribution that NYC Utility performs today must be implemented.
2. Loop and stub sites that have two links will not recover as quickly as pure EIGRP sites
3. Not as scalable and limited advanced path tuning options are available.
4. More difficult to troubleshoot with issues between EIGRP/RIP boundaries
5. Routing loops can become present and chances of errors increase with changes regarding RIP sections in Loops or EIGRP/RIP based boundary routers during or after migration.
6. Any of the issues listed in section *3.0* ***"Migration Technical Issues"*** of this document can become present.
7. Two way redistribution or default routing must be applied at the EIGRP/RIP boundaries on both sides of the loops.

Summary order of Migration approach#2

1. Add EIGRP commands to GOLRs
2. Add EIGRP commands to Core and Distribution routers – one at a time and wait a minute between each change.
3. Add EIGRP commands to access or Distribution layer routers(Loop and remote site routers) boarder routers only, with RIP redistribution into EIGRP required. Execute changes one at a time and wait a minute between each change.

At this point EIGRP and RIP are running together in every router and in the loops and remotes sites SIN routing is employed.

4. Leave RIP from each Loop and remote site router towards the major site and verify default route and overall network connectivity.
5. Remove RIP in Distribution and major site loop edge routers and verify default route and overall network connectivity.
6. Remove RIP in CORE otherwise, change redistribution of RIP in Core routers into EIGRP.
7. Control Centers requiring RIP will be handled last and EIGRP will be added but RIP will be redistributed into EIGRP.

Configuration commands required:

Migration approach #2 has some unique requirements and may not be implemented exactly the same per site. Differences are noted in BLUE.

Step #1 Same as in Migration approach #1 Step #1
Step #2Same as in Migration approach #1 Step #2
Step #3 Same as in Migration approach #1 Step #3

Step#4 Add EIGRP commands to access or Distribution layer routers(Loop and remote site routers) boarder routers only, with RIP redistribution into EIGRP required. (FORM DUAL boundary)

router eigrp 1932
network 158.57.0.0
eigrp log-neighbor-changes
eigrp log-neighbor-warnings
redistribute rip metric 1544 20000 255 1 1500 to provide lower metric for those remaing routes -- or can use higher metrics(100000 100 255 1 1500) if sourced on FastEthernet interfaces.

router rip
redistribute eigrp 1932 metric 1 (leave for two way redistribution)

or

no redistribute eigrp 1932 metric 1

or use default routing via a static route which is the preferable method to avoid using two way redistribution

If the LOOP(RIP only) side and the EIGRP/RIP boarder side utilizes the same supernet 158.57.0.0 then the following commands are required

```
access-list 10 deny   158.57.0.0 255.255.0.0
access-list 10 permit 158.57.0.0 255.255.255.0

router rip
 redistribute static
 network 158.57.0.0
 distribute-list 10 out
```

This is to ensure that the RIP only side network of 158.57.0.0/16 does not show up in the EIGRP side of the network as and external EIGRP route for 158.57.0.0/16 since this supernet originates in the EIGRP domain in the first place

For LOOPs that have the 192.168.x.x subnets defined for the site and serial links this is not an issue. All that is need on the boarder EIGRP/RIP routers are:

```
router eigrp
 redistribute rip metric 1544 20000 255 1 1500
 network 158.57.0.0
```

```
eigrp log-neighbor-changes
!
router rip
 redistribute static
 network 192.168.x.x
 network 158.57.0.0
!
ip route 0.0.0.0 0.0.0.0 towards GOLR
```

Step #5 Leave RIP from each loop and remote site router towards the major site and verify default route and overall network connectivity.

No changes at this step
Do not remove RIP

Step #6 Same as in Migration approach #1 Step #6

Step #7 Same as in Migration approach #1 Step #7

Step#8 Same as in Migration approach #1 Step #8

1.4 Roll back procedures for Migration approaches #1 and 2

Depending on the location of migration or the point of migration progress, rollback options shall vary.

1. If problems are experienced when adding EIGRP to a Loop or remote site router then turn it off if the problem does not resolve itself after a few minutes. Remember, that when activating EIGRP out towards the edges of a site, updates will be seen in the Core and Distribution routers as well.

2. If there are problems pushing EIGRP into a new site's Core and Distribution immediately remove EIGRP and investigate.

3. If EIGRP has been successfully pushed all the way out to the edges for a site and all of the site's routers are running in dual RIP and EIGRP mode, depending on where the problem arises, either remove RIP or EIGRP, so this will require some investigation and EIGRP knowledge.

4. If there is a problem after removing RIP from a site's router once EIGRP was activated, then check the EIGRP routes and if the problem cannot be resolved within an acceptable amount of time, re-activate RIP.

5. If there are problems anywhere in the Core or Distribution routers at the site being migrated or another site running EIGRP connected to the site being migrated try to isolate the problem and turn off EIGRP if in dual mode.

6. If SIA events appear in any of the routers either during a migration or after, follow the neighbors by checking all router logs, look for SIA messages and try to isolate the router or interface causing this issue.

7. If default network connectivity or internet access does not work during or after the migration then look for candidate default route information and check the GOLR distribution method at the GOLR.

2.0 Overall Migration approach:

NYC Utility must try to select one approach over the other and apply it to all major sites. For example, if the migration site used is VanNest, the approach used should be considered in all of the other sites selected to reduce the complexity of the network and support issues involved.

A surgical approach per site should be performed to reduce or prevent any issues originating from the design considerations affecting other sites. A site by site approach should be taken and careful selection and approach to each site must be analyzed to determine whether one site should not be started until the previous site is completely EIGRP based. Also, there are numerous relationships to other router based issues such as: IP subnet addressing, IP unnumbered, current redistribution et al. that are covered in section *3.0* ***"Migration Technical Issues"*** of this document.

Migration candidate starting points: Flatbush upgrade, VanNest or Westchester(Rye HQ and Eastview)

These starting points are viable, for major changes are about to occur to them. Flatbush is the best candidate for a backbone upgrade is forthcoming which involves introducing RSMs at Faltbush's Core/Distribution layer. This is an exceptional opportunity to just enable EIGRP on the RSM for Flatbush and turn the Flatbush local site into a true EIGRP based site. Remember that the GOLRs must be moved into the EIGRP AS also so that the Flatbush RSMs receive their candidate default routes.

VanNest is also a prime candidate due to the fact that the Multilink is utilizing EIGRP on one side(4IP) and RIP on the other. Multilink or other related network failures on either side currently requires 3+ minutes for RIP to converge on the VanNest side. Moving VanNest to EIGRP will help reduce this convergence time greatly.

The Rye headquarters is also a third candidate since there is currently no high bandwidth link to this site(all T-1 based) and a SONET connection is planned for this site. So, moving to EIGRP also prepares the migration for RYE to participate in NYC Utility's OC3 SONET based backbone. Also, by moving to EIGRP before such participation will help RYE eliminate some of the sub optimal routing conditions currently experienced. This would also apply to the Eastview site in Westchester as well. However, since the RYE HQ routers have many redundant links back to 4IP, Flatbush, Loops and remote sites, very careful planning should be considered in relation to all of these link dependencies so as not to cause additional routing issues during the migration. This would apply to either migration approach considered.

3.0 Migration Technical Issues

This section is used to list all of the potential issues with migrating to EIGRP and should be considered and referenced when utilizing either approach outlined earlier in this document.

3.1 NYC Utility Routing Protocol Design Rules:

Below are two rules that should be adhered to as much as possible when deploying EIGRP regardless of the migration approached utilized.

NYC Utility Routing Protocol Rule#1

Low speed links in Core and Distribution routers to similar destinations should not be used for unequal cost load balancing and only for redundant fail-over purposes. For example: any T-1 links in Flatbush or 4IP routers used as a backup for SONET should be considered a Feasible successor if SONET is not available. Very little traffic should be flowing through these links when SONET is operational. Traffic should not be load balanced or split between SONET and T1 links to the same sites. Deviations to such a rule should be avoided as much as possible. The use of the EIGRP ***VARIANCE*** command should be avoided due to NYC Utility's non-hierarchical architecture. If absolutely required then either the VARIANCE or TRAFFIC share options could be used only in each sites main routers. Their use in the Loops may disrupt application traffic flows inadvertently.

NYC Utility Routing Protocol Rule#2

All IP unnumbered subnets must be on same major network for EIGRP auto summarization to work properly. Otherwise, DUAL boundaries resulting where they do not belong may occur and IP unnumbered connectivity issues may arise such as adjacencies not forming and a link's line protocol not becoming active.

3.2 IP Unnumbered for serial interfaces

The use of IP unnumbered interfaces on active local(Ethernet or Token-Ring)interfaces may cause unnecessary EIGRP DUAL processing if a local problem is experienced on a segment. Plus, if the interface is down from failure or manually, the neighbor will remain active thus showing a false positive to the troubleshooter and network operations support personnel. This was tested in NYC Utility's lab. Also, refer to the CCO excerpts on issues with IP unnumbered listed in Appendix A of this document.

3.3 Redistribution of routes from one routing protocol into another.

NYC Utility currently redistributes EIGRP routes into RIP. This is strongly discouraged according to several Cisco published sources. The use of Multipoint point one way redistribution should be carefully looked at if migration approach #2 is considered. Multipoint two way redistribution should not be used in either migration approach and is actually not necessary if migration approach #1 is utilized. **Multipoin**t **Two-Way** redistribution can cause routing loops and count to infinity problems due to metric incompatibilities. This was tested in NYC Utility's lab. Route redistributions is a complex tool that must be carefully designed and implemented. Improper implementation of route redistribution can result in sub optimal routing, routing loops or overall routing instability as it appears today in NYC Utility's network. **One-Way** redistribution of routing information is easier to implement than two-way redistribution and is therefore the preferred design choice for migration approach #2. The use of Multipoint two-way redistribution is best avoided if possible.

One-Way redistribution into EIGRP usually works as expected. The default administrative distances of EIGRP routes have proper values to ensure optimum routing. The amount of information inserted into the EIGRP topology database must be carefully evaluated or the whole redistribution design might not yield any increase in network stability. The information inserted into the EIGRP process must be summarized and filtered before redistributed, not after it has already appeared in the EIGRP topology database.

NYC Utility's use of natural masks for the 158.57.x.x, 192.168.x and 10.x networks makes this easier, however if migration approach #2 is used, then all loop dependencies must be traced so proper redistribution into EIGRP is handled when RIP is turned off in the Core and Distribution boarders. Also, any discontinuous routes must be tracked in each site so the same route is not sourced in two different places by two different routing protocols. All static routes redistributed must be considered.

3.4 EIGRP Update throughput pacing and WAN links

Utilizing NYC Utility's Routing Protocol Rule#1 and a little planning can prevent slower or busy T-1 links from being overloaded due to a change in the EIGRP network or additional routers being added or upgraded in the enterprise. EIGRP can easily overload a WAN link with routing updates in scenarios where, for example, several neighbors are connected to a router through a single low speed WAN link, such as RYE today. The SONET sites have more flexibility because of the bandwidth available. A central router with high speed links could also congest neighbor routers with low speed WAN links(T-1) thus resulting in dropped packets on the lower speed links.

EIGRP pacing was designed to mitigate this activity and give the designer some control over how much bandwidth a neighbor can consume when sending a full update. Correct operation of EIGRP pacing is extremely important in scenarios where a router has several neighbors reachable through the same physical interface, or where the link speed mismatch between the endpoints of a WAN connection is large. Another example, in NYC Utility's case, is the SONET and any Token-Ring links that also tie the routers together. Also, any of the serial links between major sites can be affected by update pacing. EIGRP uses a sliding window of 1 for its reliable transport protocol to help prevent congestion as well. Pacing intervals are computed independently for each physical and logical interfaces(including Loopbacks).

Actual EIGRP pacing also works independently on each interface. It is therefore very important to set the bandwidth properly on each interface. The bandwidth should match the interface's actual line speed. Otherwise the combined EIGRP traffic sent over a physical interface might overload the WAN link. The bandwidth available to EIGRP over an interface is shared between all the EIGRP neighbors reachable through that interface. A round robin algorithm is used to ensure fairness

The use of IP bandwidth-percent EIGRP parameter may be required in some locations and could be planned for accordingly. There are calculations to validate its need and should be used as well.

3.5 DUAL time allocation and its diameter based on the size of the EIGRP AS

EIGRP specific network design should focus on query boundaries that limit the diameter of diffusing computations. Several scalability tools are available for EIGRP which include summarization, route filters, default routes and redistribution to assist in limiting the querying diameter.

DUAL processing can be felt across all enterprise routers participating in EIGRP for even simple changes. For example, flapping links, IOS bug issues or a failed router, especially in one of the Loops can cause DUAL to execute when there is no feasible successor present. The computation is no longer local and must be diffused across all other neighbors. When failures are present, routes are lost and neighbors time out, the up or downstream EIGRP neighbors will send out queries for the route if it has no feasible successor. These queries will be sent to the edges of the EIGRP AS. In most circumstances this happens very quickly(under 5 seconds). However, depending on the size of the network(by the number of routers) this can take longer. SIA events can be present thus causing more issues and possibly an EIGRP meltdown where one or more routers are stuck querying or waiting for query replies for routes marked active.

In NYC Utility's case migration path#1 may still be suitable since the number of routers NYC Utility has total is under 120(latest Cisco Works count). However, if migration path #1 was chosen and problems were present from the result of DUAL processing and SIA events then DUAL boundaries may need to be considered, hence the network will look like what would be performed in Migration path#2.

There are two ways to successfully reduce the EIGRP query diameter resulting in reduced convergence time and prevention of SIA events:

1. You can reduce overall EIGRP AS size(resulting in hard boundaries that EIGRP cannot cross) by introducing additional routing protocols like RIP, which is already present, in the access layers or BGP in the Core(not applicable to NYC Utility). This approach is usually employed in large networks with a good multi-layer structure. Migration path#2 lends itself to this approach. NYC Utility currently has a form of DUAL query boundaries in existence in the Core and Distribution layer in some sites, however due the to reverse redistribution, sub optimal routing exists.

2. You can establish query boundaries by using tools, such as route summarization, route filters and default routes. This does cause the configurations to become more complex and requires additional administrative planning when changes are made to such routers.

Query Processing and Range

When a router processes a query from a neighbor, the following rules apply:

Query from	Route state	Action
neighbor (not the current successor)	passive	reply with current successor information
successor	passive	attempt to find new successor; if successful, reply with new information; if not successful, mark destination unreachable and query all neighbors except the previous successor
any neighbor	no path through this neighbor before query	reply with best path currently known
any neighbor	not known before query	reply that the destination is unreachable
neighbor (not the current successor)	Active	if there is no current successor to this destinations (normally this would be true), reply with an unreachable
		if there is a good successor, reply with the current path information
successor	Active	attempt to find new successor; if successful, reply with new information; if not successful, mark destination unreachable and query all neighbors except the previous successor

The actions in the table above impact the range of the query in the network by determining how many routers receive and reply to the query before the network converges on the new topology.

DUAL cause and effect behavioral processing was tested in NYC Utility's lab under both migration paths. Migration path#1 does show that a link or router failure in the Loop, Core or Distribution routers, even if a feasible successor is present, that the processing and recovery from several failures was very quick(under 5 seconds) but this was a lab environment so production conditions will vary.

3.6 IOS version compatibility

The IOS versions of the routers must also be considered before deploying EIGRP regardless of migration path utilized. This is a serious issue for an older IOS that has EIGRP related bugs may be present and cause intermittent problems such as unneeded DUAL processing. Also, the handling of default routes and the uses of connected and static routes play a role in the IOS used. The IOS on EIGRP candidate routers should have an IOS equal or above that of the IOS on the Core and Distribution platforms(11.2(14-16)) for these versions are tested for and have been running EIGRP for some time already.

3.7 Static routes treated as connected EIGRP networks even though RIP is redistributed

Static routes pointing to an interface were considered to be ***static*** in old IOS versions; then the IOS was changed to consider them ***connected***(recent IOS up to and including 11.2). The latest IOS versions again treat the static routes pointing to an interface as static(IOS 11.3 and 12.0). Configurations relying on static routes pointing toward physical interfaces could break when you upgrade your router from IOS 11.2 or lower to 11.3 or 12.x. This should be considered especially for the GOLR, Core and Distribution routers before upgrading to 12.x.

3.8 Major Server locations

During deployments at each site major server locations and subnets must be identified and monitored once the migration has been executed to ensure that full connectivity to server subnets is still available and if any RIP speaking/listening servers are affected.

3.9 EIGRP interface characteristics

When you initially configure EIGRP, remember these two basic rules if you're attempting to influence EIGRP metrics:

- The bandwidth should always be set to the real bandwidth of the interface; multipoint serial links and other mismatched media speed situations are the exceptions to this rule.
- The delay should always be used to influence EIGRP's routing decisions.

Because EIGRP uses the interface bandwidth to determine the rate at which to send packets, it is important that these be set correctly. If it is necessary to influence the path EIGRP chooses, always use delay to do so.

At lower bandwidths, the bandwidth has more influence over the total metric; at higher bandwidths, the delay has more influence over the total metric.

Loopback interface have special properties and when used for IP unnumbered and sourcing of EIGRP adjacencies will contain the actual unnumbered interfaces metrics.

This must be considered and reviewed to ensure there are no IOS issues with using Loopbacks for unnumbered and there are metric discrepancies between the Loopback and any serials interfaces.

The use of the Delay metric should be considered for influencing routes relating to shared Token-Ring segments in major sites. If faster routes persist such as Fast Ethernet then those interfaces should be chosen and used over the Token-Ring. Care must be taken to ensure that traffic is not needlessly passing shared Token-Ring segments that does not belong there. This can be accomplished by increasing the delay metric on Token-Ring tinterfaces to ensure that the Fast Ethernet, Ethernet and possibly serial lines are chosen. The use of Token-Ring segments as Feasible Successors should also be reviewed for its viability.

3.10 Default Route/Network Origination and Distribution

The GOLRs in NYC Utility provide the default route and network of 0.0.0.0 into the NYC Utility enterprise. RIP is currently propagating this route into all routers some Core and Distribution routers also have static routes pointing towards the GOLR. Even where there is currently EIGRP and RIP, if RIP were turned off connectivity to the internet would be lost. EIGRP supports the IP default route of 0.0.0.0/0 as well as candidate default routes(default candidates). EIGRP is the only classless routing protocol that supports default candidates. The default route is redistributed into the EIGRP topology database, the default candidate marker is set automatically on the entry into the topology database. This can be seen with a **sh ip route** command. The last line with the **asterisk*** is the default candidate. EIGRP automatically redistributes connected networks(or subnets) marked as **IP default-network** into the EIGRP process. No other classless protocol performs this redistribution behind the scenes.

There are two possible ways of configuring the GOLR, however both methods require EIGRP to be running in the GOLRs.

The first method is to declare the external subnet connecting to the Firewall and ISP as the default network. The external subnet is automatically redistributed into EIGRP with a vector metric of the interface connecting the public router to the ISP. Or, in NYC Utility's case, the GOLR to the FW. It is also flagged as the default candidate, making all the other routers aware that they should use the next-hop router toward the GOLR as the gateway of last resort.

The second method is to configure the static default route pointing to the external subnet or to the physical interface itself and manually redistribute the default route into EIGRP. (This is already done in the GOLR but only RIP is present.) The redistributed route would normally inherit the interface parameters, but you could also overwrite the interface metrics by specifying metrics directly in the **redistribute** command.

Both alternatives are almost identical, with a few minor differences

The EIGRP vector metric of the default route can be better controlled in the second setup because you can control the redistribution of the default route into the EIGRP process. (In the first method, the redistribution is automatic and you cannot configure or tune it.)

The second method works even when the IP subnet on the link between the public internet router and the ISP belongs to the customer's address space.

The second method should be considered and used in NYC Utility's case. The second method was tested in the lab environment and worked as planned thus providing a default route to the EIGRP only speaking routers.

This method should be implemented first in the GOLRs so that RIP and EIGRP will be able to send out default network information. When sites are converted and RIP is removed connectivity to the internet and other outside services will not be broken.

3.11 Handling of Discontinuous subnets and major networks

How to handle starting points for the 192.168.x., 172.x.x.x and 10.x.x.x major supernets?

A review of the major networks such as the 192.168.x.x, 172.x.x.x and 10.x.x.x to ensure that discontiguous subnets are not present is highly recommended. Even though NYC Utility uses the natural masks of the major networks and subnets, EIGRP will auto summarize by default. For example, if a 10.x.x.x network is defined in the Core but it is also used somewhere else in a Loop or another site, intermittent routing loops for those networks may occur within EIGRP depending on where any destination packet is sourced. EIGRP may be confused in regards to which 10.x.x.x network is the real originating supernet and which direction to send the packet. The same issue applies to the 158.57.x.x and the 192.168.x.x networks used to remove the bridge Loops.

Additional planning must be conducted to ensure this issue will not cause any problems. This was not tested in the NYC Utility Lab and should be.

3.12 Ability to keep auto Summarization Classless/Classfull requirements

Auto summarization was initially introduced in EIGRP to facilitate smooth migration from IGRP to EIGRP. Not all networks have a classful addressing scheme where the Core would be one major IP network and each site/region would have a separate major IP network assigned, thus making a natural fit for EIGRP auto summarization. NYC Utility's classful scheme starts in the Core and extend out to the remote sites.

Auto summarization is definitely beneficial to networks being migrated from RIP to EIGRP because it retains all the routing properties the network had before. This guarantees that routing loops for dynamic or static routes would not occur or changed traffic flows after the migration would not occur. In NYC Utility's case the previous statement is half correct. We want to preserve the network properties but want the traffic flows to change for the flows that are currently utilizing sub optimal routing paths.

Auto summarization also introduces query boundaries in the network that use many different major IP networks. Not in NYC Utility's case though, unless all loop sites were reconfigured with a 192.168.x.x to reclaim 158.57.x.x and provide natural DUAL query boundaries thus ensuring a true EIGRP based network with little DUAL performance requirements.

Conversely, auto summarization hurts all of those network designs that deploy discontiguous subnets of major networks(which could be present in NYC Utility and must be reviewed before any migration). Some cases are as follows:

- **Using one major IP network for the Core network and subnets of another major IP network in various sites/regions. Not applicable to NYC Utility.**

- **Using subnets of a major public IP network on the LANs throughout the network and private IP addresses on the WAN links. (Somewhat applicable to NYC Utility)**

In these cases, it is best to turn off auto summarization and replace it with manual summarization where needed or desired. An alternative design consideration might propose turning off auto summarization only in those points in the network where discontiguous subnets appear. Such a design increases the complexity in relation to configuration troubleshooting and operation plus requires careful evaluation to verify that networks are properly summarized.

The use of the **IP classless** command could and should be turned on in all routers migrating to EIGRP, even though it will not be fully used until some form of VLSM is utilized. Since NYC Utility uses natural masks the use of the **ip classless** command may not be neededin some routers but it should be present just in case VLSM is ever used especially for the serial links, loopback addressing , 10.x.x.x 172.x.x.x and 192.168.x.x network deployment.

Where the **ip classless** configuration command falls within the routing and forwarding processes is often confusing. In reality, IP classless only affects the operation of the forwarding processes in IOS; it doesn't affect the way the routing table is built. If IP classless isn't configured (using the **no ip classless** command), the router won't forward packets to supernets. Since we are using auto summarization in EIGRP we will use IP Classless

The auto summarization was tested in the LAB for Migration approach #1 where it would most likely be an issue. There were no lost networks or any issues with interfaces having secondary addresses. However discontiguous subnet testing was not performed and should be tested. Since NYC Utility uses natural masks auto summarization should be left on regardless of either migration approach.

3.13 Handling of route flapping conditions and recovery

A single constantly flapping interface can introduce a constant stream of diffusing computations in the network; every time the interface goes down, all of the routers in the network have to agree that there is no alternate route to the lost subnet. Over a period of time the number of outstanding queries can grow to the extent that one of the diffusing computations exceeds the SIA timeout. This was tested in the lab and since the DUAL processing is performed in the background no performance issues were noted. However on a larger network this could be an issue and must be planned for. Other conditions can cause the same behavior listed above:

- Loss links or dirty links thus causing retransmissions in packets that could also extend a convergence period.
- Heavy loaded links can cause packet drops that may cause EIGRP retransmissions finally resulting in SIA routes
- Queuing mismatches, priority queuing on serial interfaces may starve out EIGRP packets during convergence times or when the link utilization is heavy. Either a priority query for EIGRP is created or the use of Weighted fair Queuing(WFQ) is implemented on all low speed serial links. WFQ will give a higher priority to EIGRP since EIGRP will use the precedent bit of 7. Packets marked with a precedent bit of 7 are processed above all other in Cisco routers.

Because adjacency resets following a configuration change are a fact of life, Network Systems should follow a few rules to make their network more stable:

- Make EIGRP related changes only during maintenance periods of your network. Any router configuration change that is linked to EIGRP operation might bring your network down for 5-60 seconds(depending on the hello timer values)because EIGRP neighbors only reestablish an adjacency after hello packets are received by neighboring routers.

- Any EIGRP related change causes a massive flurry of local and diffusing computations following the route loss caused by the adjacency resets throughout the network, more so if changes are done on the Core routers with many neighbors. It is therefore strongly advisable to plan enough time for the network to recover from a potential meltdown situation within the maintenance period.

- All EIGRP related changes on the Core routers should be done in a batch to prevent repetitive adjacency resets. CiscoWorks can be of assistance here.

- The Networked Desk should be trained on how to identify such a condition and Network Systems must be trained on how to identify and resolve such conditions.

3.14 EIGRP Load balancing

Since most of NYC Utility's routers are using fast/optimal switching when EIGRP is enabled, per destination load balancing is activated by default for paths that have equal cost. Unequal cost load balancing should be avoided and not used as per
NYC Utility Routing Protocol Rule#1. The **max paths** command can be implemented to reduce or increase the overall number of equal cost paths available to load balance traffic especially in the Core and possible major site locations. This should be reviewed as a tuning option and to provide a further control of traffic paths.

Electric and Gas Control Centers have special considerations to continue running RIP.

For each site being migrated their respective control centers should also be reviewed. For if there are Raptor based or like firewalls in between the NYC Utility network and the Control Center network, those firewalls do not speak EIGRP. It is in those cases that both RIP and EIGRP will be required, however one way redistribution into RIP/EIGRP will be needed. The MECC is one such example.

4.0 Next Steps

Below is an initial set of next steps that should be considered when creating a project plan or formalizing this endeavor as a project.

1. Review this document and determine migration prerequisites steps such as IOS upgrades, platform upgrades et al. based on information provided in the Migration Technical Issues section.
2. Select a candidate site and review all of the link/route dependencies of such site.
3. Select order of other major sites following first site migrated.
4. Create migration schedule and basic project plan (this also could be the first step as well).
5. Allocate resources to perform migration, team members involved, vendor(Cisco) onsite coverage if applicable.
6. Create change control submittals.
7. Execute migration.
8. Create or update any documentation.

5.0 TUNING EIGRP

EIGRP Tuning options available to use and may be required during a migration. Below are some of the features. A full explanation of such features is currently beyond the scope of this document.

- **Bandwidth metric**
- **Delay metric**
- **Distribution lists**
- **Route tags**
- **Load balancing**
- **Max Paths and Variance**
- **Summarization options**
- **Hello and SIA Timers**
- **Redistribution options**
- **Policy routing**
- **Bandwidth Percentage parameter**
- **Traffic share**
- **Prefix lists**

Appendix A.
Sample IP unnumbered and routing protocol issues

Bug Id : CSCdi63289

Headline	major network summary not sent by rip over unnumbered link		
Product	all	**Model**	
Component	rip	**Duplicate of**	
Severity	3	**Status**	R
Version Found	11.1(4)	**Fixed-in Version**	11.1(6)CA 11.2(0.20) 11.1(5.000.003)

Release Notes

```
Rip doesn't send a major network summary over an unnumbered link.
This will only be a problem in situations where subnets of
different major nets are connected via an unnumbered link.
```

```
Work around is to use a different routing protocol or to use
static routes.
```

Bug Id : CSCdj54973

Headline Address change of unnumbered interface confuses FIB table
Product rsp2 **Model**
Component fib **Duplicate of**
Severity 3 **Status** R
Version Found 11.1CC 12.0(2.5)S0113 **Fixed-in Version** 12.0(3.1)S 11.1(24.2)CC 11.1(25)CC 11.1(24.2)CT 11.1(25.1)CC 11.1(24.2)CT01 11.1(25.1)CT 12.0(4.2)S 11.1(25)CT 12.0(4.00.04)S 12.0(4.1) 12.0(4.00.04)PI05 12.0(4.3)T 12.0(4.2)PI05 11.1(25)CT01 11.2(17)GS0.2 12.0(4.6)W05(11.00.09) 11.2(18)GS0.1 11.1(25)CT02
Release Notes

```
When you change the ip address on an interface
that has unnumbered interfaces pointing to it, the
ip address on the unnumbered interfaces will also
change.  This will then allow the receive entry
for the numbered interface to be removed if any
of the unnumbered interfaces are shutdown.
```

Bug Id : CSCdi59425

Headline Incorrect entry of Major Net Summary RIP/IGRP w/ ip unnumbered PPP

Product all **Model**
Component ip **Duplicate of**
Severity 2 **Status** R
Version Found 11.1(1) 11.0(8) 11.1(3) 11.0(4) **Fixed-in Version** 11.0(10) 11.2(1) 11.1(5)
Release Notes

```
With IGRP and RIP, IP unnumbered interfaces using PPP encapsulation
receive but do not process routing updates of Major Network Summaries
correctly.  The major networks show up as host routes instead of
network routes.  This affects all ip unnumbered interfaces using PPP
encapsulation which are pointing to different major networks.  This
affects dedicated links as well as DDR links using RIP or IGRP.
The workaround for this problem is a floating static route for the
majornet matching the PPP created host route using a majornet mask
pointing to the PPP created host route.  i.e., if the host route is
192.1.1.1, then using 'ip route 192.1.1.0 255.255.255.0 192.1.1.1 250'
should cure the problem.
```

EIGRP or 12.0 migration considerations

Bug Id : CSCdm29563

Headline EIGRP does not summarize over ip unnumbered serial links
Product all **Model**
Component eigrp **Duplicate of**
Severity 3 **Status** V
Version Found 12.0(4) **Fixed-in Version** 11.3(11) 11.3(11)T 12.0(6) 12.0(6)S 11.3(11)NA 12.0(6)PI06 12.0(6)S 12.0(6)T 11.3(11)AA 12.0(6)W05(13.00.12)
Release Notes

```
When a router first starts up or eigrp first configured, auto-summary should
be on by default.  EIGRP should summarize routes to their major net if
advertising
them over a link that is a different major net than the advertised route.
This
is not happening when the link is an unnumbered link.  EIGRP advertises all
components of that network rather than just a summary.  The work-around is
to
go into eigrp and configure "no auto-summary" then "auto-sumary".  EIGRP
will
then behave as expected.  This was observed in late 11.3 and 12.0 code but
not
present in 11.2.

rtr#conf t
rtr(config)#router eigrp 1
rtr(config-router)#no auto-summary
rtr(config-router)#auto-summary
```

Bug Id : CSCdi92226

Headline removing secondary address on LAN kills EIGRP on unnumbered serial
Product all **Model**
Component eigrp **Duplicate of**
Severity 3 **Status** R
Version Found 11.2(3) **Fixed-in Version** 11.1(11) 11.2(5) 11.2(5)P 11.2(5)F 11.1(11)AA 11.1(11)CA01
Release Notes

Removing secondary address on an interface may cause EIGRP stops sending hellos/updates to an unnumbered link which uses that interface''s primary address.
Reload is the only way to recover.

What Are the Disadvantages of IP Unnumbered?

Question: What are the disadvantages of IP unnumbered?

Answer:
Network management is the biggest thing. Normally you would want to monitor all

interfaces on the routers, and an IP address on every interface is the only way

to do so. Also, if the interface you are pointing to on the unnumbered interface

goes down, then you lose connectivity to the router even though the serial interface

is up (no Telnet capabilities).