NYC Utility

EIGRP Migration

POST Migration documentation

January 2002

Applied Methodologies, Inc.

Table of Contents

NYC Utility EIGRP Migration

POST Migration Cleanup tasks and tuning options

The Migration of the NYC Utility Enterprise from RIP to EIGRP is complete. The new routing protocol, EIGRP, will provide additional scalability and performance benefits to the NYC Utility enterprise. This document is to serve as a road map outlining post migration issues discovered during the migration and enhancements that should be reviewed and implemented to further enhance the stability and scalability of the NYC Utility network.

Post Migration view of the NYC Utility routing Enterprise.

The NYC Utility enterprise routing architecture is composed of EIGRP with the combination of EIGRP and RIP running only in the MECC and the GOLRS. There should be no other deployments of RIP in the entire enterprise. This of course excludes the O&R and Solutions networks, which still run RIP, and were not covered in the migration project.
All routers are running ***IP classless*** and ***Auto summarization***. The only router with no Auto summarization is 708. All routers are set to ***Fair Queuing*** on all interfaces and the bandwidth is set to ***1536***. The SNMP strings and enable passwords have been corrected wherever encountered as well.

Outstanding cleanup tasks:

This section outlines the general cleanup tasks that should be applied to all routers and the more detailed site specific tasks. These tasks should be completed before any tuning or convergence activity begins.

Note: due to the 12.07T CEF bug, all site core routers should be upgraded to 12.1 before any of the tasks outlined in this document are started.

General/global cleanup issues

All Serial interfaces should be unnumbered to an active Ethernet or Loopback interface. The use of a Token-Ring interface for IP unnumbered should be discouraged. If a site's router requires redundant interfaces for its unnumbered interfaces a second Loopback can be used.

Additional global cleanup tasks that are detailed in each site's section are as such:

- Removal of Bridging
- Removal of superfluous protocols such as DEC, IPX, XNS

- Comdisco to VN2 IP address assignment on serial links
- Cleaning up of unneeded static routes in all major site core routers.

Control Center routers:

The Rye, Brooklyn and Staten Island control centers should run HSRP to simplify the control center's UNIX server default route configuration.

The MECC control center routers are still running 11.2. Because the MECC routers are still running RIP, their IOS upgrade should be 12.1 to avoid the Spurious Access RIP bug.

IOS upgrade progress.

For a master list please see the EIGRP migration progress table in appendix A. of this document. This list shows the current IOS version of all routers migrated. A separate site list is provided in each site's section.

Note: The DC1 and 2 RSMs are still running 11.2

Content switches and the relationship to EIGRP.

There are currently two Content switches deployed in the NYC Utility enterprise. These Content switches were deployed during phase one of a dual phase project. The phase one functionality of the Content switches has no bearing or impact in regards to EIGRP operation. The phase one functionality provides only fail over capabilities for services inside NYC Utility via static routes and shared addresses. There is a RIP process active on the Content switch but that function is currently not used.

In Phase two of the Content switch project the use of routing between ports behind the Content switch through the bus to the ports in front of the Content switch that connect to the enterprise will be utilized. The back-end ports will have their own subnet where load balanced servers shall reside. These subnets possibly 158.57.214.0 and 158.57.227.0 can route through the Content switch via RIP to announce these subnets to the NYC Utility enterprise. Since NYC Utility is not running RIP in any of its core site locations this presents a small issue.

There are two approaches. The first one is to add RIP and redistribute into EIGRP on any two routers where a Content switch also resides. This provides a flexible solution but requires pockets of RIP to be maintained in the core. Also, the routers will have to be upgraded to 12.1 to eliminate the Spurious Access bug found when RIP is enabled on 12.07T.

The second approach is simpler and just requires the use of static routes in the Content switches pointing to the enterprise. Static routes in any two routers on the same segment where a Content switch resides can redistribute information about subnets 214 and 227 pointing to the appropriate Content switch. With this approach no pocket of RIP and IOS upgrade is required. Since the Content switch networks are never expected to change, this approach is the simplest and cleanest to implement, it was also recommended by Cisco.

Any EIGRP convergence testing should also include the content switches preferably after phase two is completed to properly determine the impact.

RYE

Outstanding items that can be covered in the cleanup:

Transparent bridging still remains on loop routers
West 110 St has a switch using the 192.168.56.240 address but this is not defined in the site's router.
There are still serial interfaces using IP unnumbered to active Token Ring interfaces
W125St Still has open RSRB peers to 100.9(DC2CORE2)

Routers to be upgraded to 12.07T that were already prepped and have been migrated:

EVBR	158.57.220.172
PLBR	192.168.73.254
MWD2BR	192.168.71.254
ELMBR	192.168.77.253

VanNest

Outstanding items that can be covered in the cleanup:

DEC still running on VN2 and Van Nest
Transparent bridging still remains on loop routers

Kingsbridge router still using RSRB and Bridge Group 1
Serial 1/0 and 1/1 using unnumbered to Token-Ring interfaces

East 179St. router still using RSRB and Bridge Group 1
Hell Gate router still using Bridge Group1 and Dec Net

BDG-KNSBDG is just a stub router off of Kingsbridge's Token-Ring with IP routing disabled used for shut down Metrotech link. RSRB was removed no peers.

Bruckner router is still using Bridge Group 1 and Decnet -- RSRB was removed no peers.

CPI router is still using RSRB and Bridge Group 1 -- removed Bridge Group 2 and Decnet
Serial 1/0 using unnumbered to Token-Ring interface 0/0
Serial 1/1 using unnumbered to Token-Ring interface 0/0
Serial 1/2 using unnumbered to a shutdown Ethernet 0/0

JAMABR and QHBR are still running bridging over their serial interfaces between each other. IP routing is disabled

Jamaica router is still using RSRB
Serial interfaces are unnumbered to Token-Ring interfaces.

Routers to be upgraded to 12.07T that were already prepped and have been migrated:

KINGSBR	192.168.230.1
SHERBR	192.168.220.2
RAMBR	192.168.222.2
PLVABR	192.168.223.2
EFISHBR	192.168.224.2
DUNWOODIE	192.168.58.251
SPBRBR	192.168.231.1
EAST179	158.57.162.250
FARRINGTON	158.57.80.250
JAMAICA	158.57.57.1

Brooklyn Staten Island and West End Ave.

Brooklyn

Flat 6 still using RSRB, Bridge Group 1 and Decnet
serial 1/1/1 and 1/1/3 unnumbered to tok4/0
serial 1/1/0 and 1/1/2 unnumbered to fa 1/0/0

30 Flatbush still using RSRB, Bridge Group 1 and Decnet
serial 1/1/0 and 1/1/2 unnumbered to tok4/0 all others using fa1/0/0

BKCC1 still using Decnet and IPX
BKCC2 still using Decnet

FirstSt still using RSRB, Bridge Group 1 and Decnet -- S3/0 using Tok0 for IP unnumbered
Gowanus still using Bridge group 1 and Decnet
Faragut still using Bridge group 1 and Decnet
Gold Bridge Group 1 active -- removed RSRB on 11/27
Atlantic Bridge Group 1 active -- removed RSRB, IPX, XNS on 11/27

Neptune	Bridge Group 1 active -- removed RSRB on 11/27
FKRTR	Bridge Group 1 active
Victory	Bridge Group 1 active
Rockaway	still using RSRB, s1/3 using Tok0/0 for IP unnumbered -- removed Bridge group 1 and Decnet on 11/27

Staten Island

Davis	still using RSRB, Bridge Group 1 and Decnet s 3/0 and 3/1 using Tok2/0 for IP unnumbered
DavisCC	still using RSRB, Bridge Group 1 and Decnet s 3/1 using Tok2/0 for IP unnumbered
SICC1	Decnet, IPX and XNS are active -- removed removed Bridge group 1 on 11/28
SICC2	Decnet, IPX and XNS are active -- removed removed Bridge group 1 on 11/28

West End Ave.

EastRiver	still using Bridge Group 1 and Decnet
Hudson	still using Bridge Group 1 and Decnet
W59st.	still using Bridge Group 1 and Decnet
ECC	still using Bridge Group 1 and Decnet
ECC2	still using Bridge Group 1 and Decnet --- removed IPX on 11/28

Routers to be upgraded to 12.07T that were already prepped and have been migrated:

VICTORY	158.57.200.171	11.2
FKRTR	158.57.138.1	11.2
NEPTUNE	158.57.179.1	11.2
ATLANTIC	158.57.198.250	11.2
GOLD	158.57.184.250	11.2
GREENWOOD	158.57.182.250	11.3
FARAGUT	158.57.73.171	11.2
Davis	158.57.137.171	11.1
SICC1	158.57.137.172	11.2
SICC2	158.57.137.173	11.2
GOETHALS	158.57.132.250	11.2
FOX HILLS	158.57.130.1	**10.3**
EAST RIVER	158.57.83.171	11.2
W59TH	158.57.87.1	11.2
ECC	158.57.88.1	11.2
ECC2	158.57.24.172	11.2

ECCG1	158.57.24.173	11.2
ECCG2	158.57.24.174	11.2
BKCC1	158.57.75.174	11.2
BKCC2	158.57.75.176	11.2

Queens/Manhattan

Queens

Rainey	still using Bridge Group 1, change description on Ethernet interface The RAINEYBR router should be removed so the VERNONBR and AST136BR routers be linked to the RAINY router
AST136	still using RSRB, Bridge Group 1 and Decnet -- S3/0 is down but using Tok0 for IP unnumbered
AST137	still using RSRB, Bridge Group 1 and Decnet, S3/0 is using Tok0 for IP unnumbered -- removed IPX on 12/4
VAN_DAM	still using Bridge group 1 and RSRB -- Loopback has no IP address assigned
Corona	still using Bridge Group 1 and Decnet
TLC3NE	Still using DEC, RSRB and Bridge Group 1
TLC2SW	Still using RSRB and Bridge Group 1

Manhattan

W54Th	still using RSRB, s1/3 using Tok0/0 for IP unnumbered - Loopback has no IP addressed assigned
E74Th	Change description on Ethernet interface -- Removed Bridge Group 1 -- 12/4
W49St	Removed IP http server
E36BR1	Removed Bridge Group 1 on 12/4
E36temp	Removed Bridge Group 1 on 12/4
W19th	Removed Bridge Group 1 on 12/4
DC1CORE1	Still using DEC, RSRB and Bridge Group 1
DC1CORE2	Still using DEC, RSRB and Bridge Group 1
DC2CORE1	Still using DEC, RSRB and Bridge Group 1 -- Virtual Access1 int. is unnumbered to Token-Ring
DC2CORE2	Still using DEC, RSRB and Bridge Group 1

Routers to be upgraded to 12.07T that were already prepped and have been migrated:

RAINEYBR	158.57.66.201	**11.2**
VERNONBR	192.168.69.254	**11.2**
QBRBR	192.168.68.254	**11.2**
AST136BR	192.168.67.254	**11.2**

VAN_DAM	158.57.71.1	**12.02**
W54TH	158.57.111.254	**12.2**
60THST	158.57.114.250	**11.3**
E74TH	158.57.89.1	**11.2**
W49ST	158.57.90.1	**12.1**
W49THSS	192.168.40.250	**11.2**
W42NDSS	192.168.41.250	**11.3**
E40BR1	192.168.42.250	**11.3**
E36BR1	192.168.43.250	**11.3**
E36TEMP	192.168.43.252	**11.2**
E29ST	192.168.54.250	**11.2**
W19ST	192.168.44.250	**11.3**
MECC_LAN_R1	158.57.26.250	**11.2**
MECC_LAN_R2	158.57.26.251	**11.2**

Static Route Removal Candidates.

These static routes should all be sourced and redistributed from the GOLRS
There is no need for them to be in all of the CORE routers. Additional review should be undertaken before removal.

TLC3NE
ip route 0.0.0.0 0.0.0.0 158.57.48.251

TLC2SW
ip route 0.0.0.0 0.0.0.0 158.57.48.251 85
ip route 0.0.0.0 0.0.0.0 158.57.12.175 130

DC1DIST1
ip route 0.0.0.0 0.0.0.0 158.57.12.175

DC2DIST1
ip route 0.0.0.0 0.0.0.0 158.57.12.175
ip route 0.0.0.0 0.0.0.0 158.57.48.251 2

DC2DIST2
ip route 0.0.0.0 0.0.0.0 158.57.12.175
ip route 0.0.0.0 0.0.0.0 158.57.48.251 2

DC1RSM1
ip route 0.0.0.0 0.0.0.0 158.57.12.175
ip route 0.0.0.0 0.0.0.0 158.57.48.251 2
ip route 192.168.16.0 255.255.255.0 158.57.10.252

DC2RSM1
ip route 0.0.0.0 0.0.0.0 158.57.12.175

ip route 0.0.0.0 0.0.0.0 158.57.48.251 2

DC2GATE1
ip route 158.57.0.0 255.255.0.0 158.57.12.171
ip route 158.57.0.0 255.255.0.0 158.57.48.251

TLCGATE1
ip route 158.57.0.0 255.255.0.0 158.57.48.252
ip route 158.57.0.0 255.255.0.0 158.57.48.253
ip route 158.57.0.0 255.255.0.0 Serial3/2 130

On all 4IP/TLC Core, dist, and RSM routers
Remove 158.57.45.0 and .150 static routes for they are already sourced at the GOLRs.

EIGRP Tuning options

The tuning of EIGRP to determine optimal use of all available bandwidth and paths should be considered. If tuning is employed then failure and convergence testing can commence afterwards. If NYC Utility chooses to perform its failure testing first then tune its EIGRP network a second round of failure testing may be required to determine any issues or improvements that may have resulted from the tuning.

EIGRP can currently load balance up to 4 equal cost paths. If NYC Utility has more than 4 paths to any one or multiple destinations then this can be accommodated by changing to a maximum of 6 paths with the **MAX PATH** command under the EIGRP process. This tuning parameter also applies if unequal cost load balance is employed as discussed further. The MAX PATH should be the same on all routers with relationships to each other regardless if all paths are used.

Every routing protocol supports equal cost path load balancing. IGRP and EIGRP also support unequal cost path load balancing, which is known as variance. The **variance** command instructs the router to include routes with a metric less than or equal to ***n*** times the minimum metric route for that destination, where ***n*** is the number specified by the variance command. For example, **variance <*n*>**. Traffic is also distributed among the links with respect to the metric.

Note: If a path isn't a feasible successor, then it isn't used in load balancing.

The variance should be considered and possibly combined with the Traffic Share options if the flow of traffic is extremely ill proportioned. Also, if these commands are applied in the RYE core routers for example, then all of the other core routers connecting to RYE should have the same options.

The use of Traffic share and its options should also be considered as described below.

Multi-interface load splitting allows you to efficiently control traffic that travels across multiple interfaces to the same destination. The **traffic-share min** router configuration command specifies that if multiple paths are available to the same destination, only paths with the minimum metric will be installed in the routing table.

The number of paths allowed is never more than six. For dynamic routing protocols, the number of paths is controlled by the **maximum-paths** router configuration command. The static route source can always install six paths. If more paths are available, the extra paths are discarded. If some installed paths are removed from the routing table, pending routes are added automatically.

When the **traffic-share min** command is used with the **across-interfaces** keyword, an attempt is made to use as many different interfaces as possible to forward traffic to the same destination. When the maximum path limit has been reached and a new path is installed, the router compares the installed paths. For example, if path X references the same interface as path Y and the new path uses a different interface, path X is removed and the new path is installed.

Failure testing should then be conducted to determine proper path selection and that bandwidth utilization levels are consistent around the failed component and its relationships.

The use of setting the Eigrp neighbor **Hold** and **Hello** timers to match **HSRP** for instant convergence should be considered and tested.

EIGRP Convergence and failure testing.

Now that the entire NYC Utility enterprises is running EIGRP as their routing protocol some tests should be conducted to determine the behavior various parts of the network may exhibit upon a failure of a router or a link. Below is a general approach to testing the convergence capabilities of the new routing protocol. NYC Utility personnel can add to this approach to whatever level of detail is necessary.

Approach:

Various routers in each major site in the NYC Utility network should be shut down to determine how EIGRP will react and understand the impact of the failure. This approach enables the NYC Utility engineers to determine what the network may look like in the event of such an outage(even if any tuning or change is required from a configuration or physical perspective) to ensure that even under a failure atmosphere the network is behaving in a predictable manner and still providing access to all necessary corporate resources. Failures should comprise of a complete router failure where redundant routers are employed and a link failure where a single router is employed(i.e. loop sites). Before any router is brought down a list of before failure topology information should be complied and compared after the affected router is brought down with the router remaining active.

It is presumed that the engineer performing this exercise is familiar with what to look for from the command output. Some of the information to be compared before and after the failure is:

SH IP EIGRP NEIGHBORS
SH IP EIGRP TOPOLOGY
SH IP ROUTE SUMMARY
SH IP ROUTE 0.0.0.0
SH IP ROUTE
SH IP CEF
SH IP CEF NON-RECURSIVE
SH IP CEF UNRESOLVED

The engineer performing this exercise can add any additional commands that are deemed necessary to determine routing path behavior.

From a major site router or a loop site router perform a ***ping, trace route and a show IP route*** towards the subnet from the failed router to ensure that even after the failure the edges of the enterprise see the convergence and can reach the redundant path around the failed device.

GOLRS
The GOLRS are still running 12.07T. Because the GOLR routers are still running RIP their IOS upgrade should be 12.1 to avoid the Spurious Access RIP bug.

The GOLRS should be tested first to ensure that default route propagation is maintained if any of the GOLRs fail. This should be done before any static route GOLRE clean up is performed and again after the clean-up to ensure that redundancy was not broken from the removal of a static route.

Failure of a GOLR on each location and testing from workstations on the failed GOLR side should consist of checking internet access, access to standard applications and access to any O&R or Solutions resources.

The GOLR remaining up should be looked at and routes to important subnets verified as available.

The routers around the failed GOLR and the running GOLR should be checked to ensure that a default route is still available and routes to O&R and Solutions are available.
A loop router at the edge of the enterprise should also be checked.

This test should be reversed so both GOLRS are tested.

Control center router convergence testing.

At each of the control centers one router should be shut down during a controlled test and application access tested. The remaining active router should be checked for the number of routes and EIGRP neighbors as mentioned earlier. HSRP and EIGRP convergence should be verified.

Major Site Core routers

At each of the major business locations one POS based router should be shut down during a controlled test and application access tested. The remaining active routers should be checked for the number of routes and EIGRP neighbors as mentioned earlier. A review of the EIGRP topology for Successor and Feasible Successor changes noted. HSRP and EIGRP convergence should be verified. A second test at the site should be conducted where this time the POS router stays active and a non-POS router is shut down. If both routers at the site have POS then a test of just the POS interfaces shut down should be conducted to see how convergence operates and traffic flows over the remaining T-1 links.

Loop Routers

At some selected loop sites, preferably one in each borough, should have a serial interface shut down on one side and applications access tested. The remaining active loops routers should be checked for the number of routes and EIGRP neighbors as mentioned earlier. Default route and EIGRP convergence should be verified

Core Routers with bridging enabled interface table

Router	Interface bridging	Other End	TX/RX Traffic PRESENT
DC1CORE1	FA1/1		
DC1CORE2	S1/1	VAN DAM	
	S/2	CORONA	
	S6/1	RYECC	TX/RX
	S6/3	ASTORIA137	
DC2CORE1	FA1/1		
DC2CORE2	S5/4	ECC	
	S5/6	FLAT6	TX/RX
	S5/7	TLC2SW	TX/RX
	S5/1	FARAGUT	
	S6/4	EASTVIEW	
	FA1/1		
RYE_HQ	FA1/1		
RYECC	FA1/0/0		
	S4/1/3	RYESV	
	S4/1/0	DC1CORE2	TX/RX
VN2	FA1/0/0		
VANNEST	FA1/0/0		
	S4/1/2	BRUCKNER	
TLC3NE	FA4/0/0		
TLC2SW	FA4/0/0		
	S5/3	DC2CORE2	TX/RX
FLAT6	FA1/0/0		
	S1/1/3	FIRST AVE	TX/RX
	S/1/1/4	DC2CORE2	TX/RX
30FLATBUSH	S1/1/4	DC1CORE2	

Appendix A. EIGRP Migration Progress Table

Appendix B. Project Plan Gantt Chart