

NYC Utility

*Cursory Network
Infrastructure review*

May 2001



Applied Methodologies, Inc.

Table Of Contents

| | |
|--|----|
| <u>INTRODUCTION</u> | 4 |
| <u>PROFESSIONAL SUMMARY</u> | 5 |
| <u>OBSERVATIONS DETAILS</u> | 5 |
| <u>2.0 MAJOR & STRATEGIC OBSERVATIONS/RECOMMENDATIONS</u> | 5 |
| <u>1.0 MINOR & HOUSEKEEPING</u> <u>OBSERVATIONS/RECOMMENDATIONS</u> | 6 |
| <u>1.1 Consistency of Configurations Across Routers/Switches</u> | 6 |
| <u>1.2 Syslogging</u> | 7 |
| <u>1.3 Devices in DNS</u> | 7 |
| <u>1.4 Other Housekeeping Issues Observed</u> | 7 |
| <u>2.0 MAJOR & STRATEGIC OBSERVATIONS/RECOMMENDATIONS:</u> | 8 |
| <u>2.1 Router/Switch Password Authentication</u> | 8 |
| <u>2.2 Configuration Archiving</u> | 8 |
| <u>2.3 Redundant RSP</u> | 9 |
| <u>2.4 Unnecessary router hops</u> | 9 |
| <u>2.5 Fox Hills router</u> | 11 |
| <u>2.6 IP Unnumbered</u> | 11 |
| <u>2.7 Site Loop design</u> | 11 |
| <u>2.8 Unnecessary traffic on Core</u> | 13 |
| <u>2.9 Router System Buffers</u> | 14 |
| <u>2.10 IP addressing scheme</u> | 15 |
| <u>2.11 IP Secondary addressing</u> | 15 |
| <u>2.12 Loop-back Interfaces</u> | 16 |
| <u>2.13 RIP migration to EIGRP</u> | 16 |
| <u>2.14 Selection of EIGRP AS number</u> | 16 |
| <u>2.16 Debug Router</u> | 17 |
| <u>2.17 Terminal Server</u> | 17 |
| <u>2.19 Network Management upgrade</u> | 18 |
| <u>2.20 T-1 Circuit Protocol Analysis</u> | 18 |
| <u>2.21 CiscoWorks 2000</u> | 19 |
| <u>2.22 Router Memory Core Dumps</u> | 19 |
| <u>2.23 Network Documentation</u> | 19 |
| <u>2.24 Change Control Approval</u> | 20 |
| <u>2.25 Multicasting Services</u> | 20 |
| <u>2.26 IOS version 12.x for enterprise routers</u> | 20 |
| <u>2.27 Network Time Protocol - NTP</u> | 21 |

| | |
|--|----|
| <u>2.28 Unused Interfaces</u> | 21 |
| <u>2.29 Router Reboot Schedule</u> | 21 |
| <u>2.30 Router Menu System</u> | 21 |
| <u>2.31 Access-Lists</u> | 22 |
| <u>2.32 Network Switching Core Upgrade</u> | 22 |
| <u>2.33 802.1Q/P Integration</u> | 23 |
| <u>2.34 Backbone Switch & RSM Warning Messages</u> | 24 |
| <u>3.0 GENERAL STAFF OBSERVATIONS</u> | 25 |
| <u>4.0 FUTURE STUFF</u> | 26 |
| <u>4.1 Broadband Technologies</u> | 26 |
| <u>4.2 Forensics Protocol Analysis</u> | 26 |

Introduction

This review of NYC Utility's (NYC Utility's) network infrastructure components, which consist of Cisco Routers and Switches, is only cursory. A basic review of the major components was conducted over the last two weeks. There are several items of significant and insignificant importance. Some of these items NYC Utility may already be aware of. Since this review is only cursory, some of the configurations, operating conditions and relationships to ongoing plans and projects relative to the items listed in this document may be incomplete from the view of the outside reviewer. The method to gather the information for this report was the following:

- 1. Physical inspection of Routers/Switches operating conditions and configurations.***
- 2. Interviews and questions asked to the support staff.***
- 3. Experience gleaned from ongoing projects and troubleshooting assistance provided by the reviewer.***

The goal of this review is to provide an outsider's brief interpretation of the current state of affairs regarding the network infrastructure components. This information will also be used for the planning and prioritization of tasks and projects relative to the strategic and tactical direction of NYC Utility's network.

Professional Summary

Based on what was observed during the past two weeks, NYC Utility's network is currently in a precarious state in terms of direction and growth. The original network was designed well and operates as such today. The network runs well for its current requirements despite some outages resulting from configuration oversight, product failures, design issues, scalability, and the many network related changes (floor and remote location moves, et. al.) occurring within NYC Utility. The network is morphing into an extremely complex enterprise class tool that NYC Utility relies increasingly upon each day at a rate that the Network Team (NT) support staff cannot keep up. There are many tactical changes implemented due to project requirements, Moves Adds and Changes (MACs), and outages.

NYC Utility's NT is constantly required to architect this network around an application or a specific need, thus moving away from a symmetrical architecture to an asymmetrical architecture employing unique design considerations for each requirement. The result is a heterogeneous network that does not facilitate management, MAC, flexibility and scalability needs. This network is growing and changing regularly without a definitive strategic mandate.

So, for NYC Utility to scale this network to support new applications, provide additional resiliency, bandwidth, throughput, management and flexibility some major and minor changes shall be required.

The Details section of this report provides a cursory outline of the observations noted and recommendations for such changes necessary to provide NYC Utility a starting point and ideas on scaling their current network to enterprise class status.

Observations Details

The Details section of this report outlines the observations from the reviewer and provides recommendations for strategic and tactical planning.

The observations will be outlined in the following sub sections:

1.0 Minor & Housekeeping Observations/Recommendations

2.0 Major & Strategic Observations/Recommendations

1.0 Minor & Housekeeping Observations/Recommendations

1.1 Consistency of Configurations Across Routers/Switches

There seems to be an inconsistency in terms of IOS features turned ON or OFF on many of the components observed. This lead to questions when troubleshooting, configuring a new services and dubious understanding on how the device is supposed to be operating. To simplify the management, troubleshooting and scalability of all components(Routers/Switches) the basic functions or common denominator across all devices should be the same unless required for a special design purpose.

Some of these inconsistency include:

- Route cache on/off on interfaces
- IP redirects on/off on interfaces
- IP mask reply on/off on interfaces
- Console logging, which leads to unnecessary CPU overhead
- Syslogging, some device are active others are not
- Interface descriptions. Some very good and descriptive others not.
- Unused or out of service interface descriptions
- IP Domain lookup enabled
- IP Classless
- IP subnet zero
- Use of Loop-Back for unnumbered
- Flash setup and partitions
- IP host tables on all routers for Trace Route purposes
- Physical labels on some routers in Data Center but not all
- Banners identifying the device
- ARP timeouts
- EXEC timeout
- Some to no redundancy on remote routers example FKRTR has a T-1 to 4IP but is admin down. Goethals/Fox hills none
- Core Router POS framing set SDH not SONET STS-3(default)
- For EIGRP enabled routers there is no logging of EIGRP neighbor changes
- Logging synchronous on VTY and console lines
- Inconsistent Router Gateway/interface addressing schema
- Current router ports are all over Switch blades and redundant ports should be on different blades/switches but with same mod/port # as primary

Recommendations:

Define a standard common denominator for all components using the list above. Sweeping all the components to validate and update their respective configurations. The use of CiscoWorks 2000 for configuration “snap-in” capabilities should be considered and this activity would be a perfect test of that tool’s capabilities.

1.2 Syslogging

Syslogging must be enabled on all components to facilitate troubleshooting for the Helpdesk and Network Team. Applied Methodologies, Inc. has, as of this writing, submitted a draft strategy for NYC Utility to use.

Recommendation: Implement Syslogging on all components as per the strategy submitted.

1.3 Devices in DNS

The possibility of adding Domain Lookup in all components and adding the component’s addresses into the NYC Utility DNS structure. This may be easier to manager than host tables in each router, but the DNS lookup will be required when doing trace routes and jumping around components.

Recommendation: Look into using DNS for the router/switch names.

1.4 Other Housekeeping Issues Observed

Removal of any legacy commands and configurations not in use on the components.

Again, these legacy configuration options cause unwanted processing, may prevent or impact an IOS upgrade, increase the size of the configuration unnecessarily and again cause the troubleshooter to question its operational validity if not privy to the history of the configuration. Some of these legacy items include:

- **RSRB remote peers and processes**
- **Bridge groups**
- **Access-lists (especially for group 700 SNA)**

Recommendation: Sweeping of all components and removal of unnecessary configurations settings active or inactive. This can also be done at the same time of correcting the configuration inconsistencies mentioned in section 1.1 earlier in this document.

2.0 Major & Strategic Observations/Recommendations:

This section covers a series of major observations that have a strategic and or tactical relevance to the operation of this network.

2.1 Router/Switch Password Authentication

The current Router/Switch password authentication process is extremely basic and vulnerable to inside or outside exploitation. The user level password can be guessed easily as well as the privileged. The privilege level passwords are not encrypted in the configuration. The current approach is a major security violation for NYC Utility. In the event that NYC Utility is audited by an outside governing body, for security compliance, the network infrastructure components will fail such an audit. Also, Network Team, has no control over who can access what component(s), when and what was done to such component(s). Plus, there seems to be too many hands touching these devices from the Help-Desk, LIS and NT. Also, administration of passwords, if ever done, in this environment is decentralized, very tedious and cumbersome.

Recommendation: Implementation of Cisco Secure ACS.

All switches and routers shall authenticate user access ID/passwords from redundant servers. Accounts can be created to stratify levels of access and maintain control of what can be done to each component. Logging of who entered what component, when and what commands were executed can also be reviewed for security or mis-configuration back tracking. Administration and maintenance of the component's access is simplified and centralized.

This will require adding AAA configurations command to all components and installation of Redundant Cisco Secure ACS servers. Applied Methodologies, Inc. will provide a draft Cisco Secure Router and Switch Access policies and procedures document, plans for deployment and operational documentation if requested.

2.2 Configuration Archiving

Component device configurations are not backed up every evening. Due to the number of ongoing MACs occurring against the components, NYC Utility has no effective method to roll back a configuration or know what the previous configuration was after a change, especially if the change had a negative affect. Also, new configurations are not backed up immediately(may be days or weeks) after a change resulting from an upgrade or troubleshooting exercise.

Recommendation: All component configurations should be backed up every evening to a central location for access and use. The use of CiscoWorks 2000 or Perl scripts from a Unix server can facilitate this need.

2.3 Redundant RSP

There are no redundant slave RSP in Core components. It appears that NYC Utility is very fault tolerant aware when it comes to wiring of the components and power. However, there are no redundant RSP in the Core routers for quick recovery in the event of an IOS or RSP/CYBUS failure. A failure to the RSP, regardless of reason, can cause the entire router to fail, thus causing an impact to remote sites and 4IP. Even though there is a redundant Core this is still a major operational fail-over exposure.

Recommendation: Add redundant RSP in Core routers

2.4 Unnecessary router hops

A High number router hops were noted in several RIP announcements from the Core backbone segment: These high hop counts are not a positive occurrence for there may be routing loops, thus causing an artificially larger network diameter for some networks than normal.

These high hop counts could also be the result of a mis-configured or currently failed component. These hop counts were verified days later and are still present. In the event of an outage somewhere in the network this condition could compound the outage for any additional routes added for a divergence path may hit the 15 hop limit. The following networks were briefly noted for this condition and verified several days later.

| Subnet | Hops |
|------------|--------------|
| 158.57.179 | 9 |
| 158.57.132 | 13 |
| 158.57.131 | 12 |
| 158.57.140 | 11 |
| 158.57.200 | 10 |
| 158.57.122 | 8 |
| 158.57.179 | 9 |
| 158.57.179 | 9 |
| 172.150.0 | 16(infinity) |

After careful study of a particular path to subnet 158.57.132.0 via a trace route the results were as follows:

```
1 158.57.34.2 4 msec          30 flat POS 0/0/0
  158.57.100.9 0 msec        dc2core2 Serial5/2
  dc2core2 (158.57.159.174) 0 msec dc2core2 serial4/0 4/1 FA1/1
2 158.57.192.2 0 msec          Flat6 Tok 4/0 s/1-3-4
  158.57.75.175 0 msec 0 msec Flat6 Fa1/0/0 s1/0
3 158.57.206.171 4 msec 0 msec 4 msec First St eth1/0
4 158.57.185.1 8 msec 4 msec 12 msec Gold St. Lo0 S/0
5 158.57.195.1 12 msec 76 msec 12 msec Atlantic Ave. Tok0 S0-1
6 158.57.179.1 12 msec 24 msec 68 msec Neptune Ave. Lo0 S0-1
7 158.57.200.171 56 msec 68 msec 32 msec Victory Eth0 s1/0-1-2
8 158.57.138.1 12 msec 24 msec 96 msec Fresh Kills Lo0 s1/1 and 3
9 158.57.130.1 56 msec 52 msec 20 msec Fox Hills Lo0 s0-1
10 158.57.132.250 32 msec * 128 msec Goethals Eth0 s0-1
```

There appears to be a long path from the Core to the Goethals location. Unattended application processes or users may experience response time issues due to the longer than usual path. Also, in the event of an outage this path may not converge to an alternate. The path was supposed to be shorter from the Core through Davis to Goethals, but this was not the case here.

Further investigation revealed that the reason for this routing activity was because there was no IP address assigned to the other side of the Goethals link to Davis. The Davis router was missing an **IP unnumbered** command. This was corrected and the route from the Core to Davis is now:

```
1 158.57.100.9 0 msec
  dc2core2 (158.57.159.174) 4 msec
  158.57.100.9 0 msec
2 158.57.196.1 4 msec 0 msec 4 msec
3 158.57.132.250 8 msec * 4 msec
```

As you can see a mis-configuration or oversight resulted in an abnormal network condition. Since there were other subnets with higher than usual hop counts the above condition or other similar conditions may still be occurring.

Recommendation: Review and investigate high hop count subnets for RIP routing loops, unnoticed failures, or mis-configuration.

2.5 Fox Hills router

The Fox Hills router is running IOS 10.3 and has a small amount of memory. This was discovered from the activity discussed earlier regarding the HOP count issue. Since this router is in the path of several other routers looped together it is a potential weak link since it is running an “End of Deployment” classification version of the IOS. Since it only has 4096K/2048K of router processor memory performance issues may arise if heavily loaded. Also this configuration prevents the ability to upgrade to a more enhanced and stable IOS version.

Recommendation: Look into upgrading the Fox Hills router.

2.6 IP Unnumbered

The use of IP unnumbered interfaces to save address space is an ideal solution for tactical purposes. However, since the routing updates are RIP based and sourced from another subnet on a sourced interface this can cause issues in hop count and convergence as listed earlier. Plus, some of the sourced interfaces are actual interfaces and others are loop-backs, another issue of inconsistency here. Also, you cannot use the ping EXEC command to determine whether the interface is up, because the interface has no address. Simple Network Management Protocol (SNMP) can be used to remotely monitor interface status. You also cannot netboot a runnable image and you cannot support IP security options over an unnumbered interface.

Recommendation: Review the plausible use of IP unnumbered as a design strategy if it is still only intended to save address space. This issue is also relevant to the addressing issues that will be covered in section 2.10 of this document. This issue is also relevant to the loop design of the remote locations discussed in the next section.

2.7 Site Loop design

The use of the loop(see figure 1 below)or daisy chain concept of chaining routers/sites together and using IP unnumbered and bridging for a single subnet or routing for physical and logical connectivity is a dubious tactical solution. The recent Buchanan outage is a perfect example. This type of configuration lends it’s self to scalability and performance issues. For example, if the end of the loop/chain wants to receive multicast, hoot or holler, video conference or other traffic it must traverse X number of other devices. The previous HOP count issues, as mentioned in section 2.4, can also rear its head if routing is involved and there is an outage or mis-configuration. The use of DEC Spanning-Tree protocol is questionable over IEEE Spanning Tree for DEC Spanning tree has a history of compatibility issues and operation.

The loop design concept works well for AC electrical circuits but does not lend itself well for dynamic packets. An outage on either side can cause one or more devices to be affected. It is understood that some of the sites in the loops are unmanned or only a handful of users reside at them.

It is also understood that this design may be the result of prudent economic factors such as fiber that is owned by NYC Utility and is a significant cost savings as opposed to using the Telco exclusively to provide connectivity.

Recommendations: Economic and user population factors aside. If NYC Utility wants to provide best of class services to all loop locations and maintain a highly redundant, scalable network and be able to offer the same application and services to all users regardless of a HQ or loop location then a redesign to a Hierarchical Mesh should be considered. A hierarchical mesh will provide increased fault isolation, enhanced redundancy, increased bandwidth scalability and more flexibility in terms of offering new services over the network.

Figure 1.

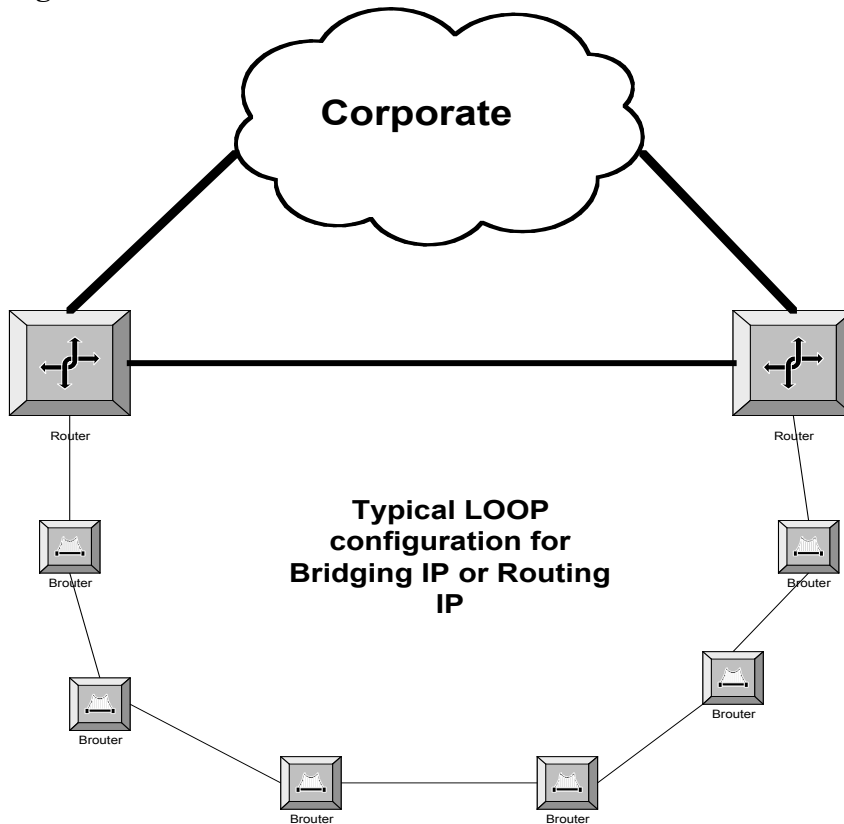



Figure 2.



2.8 Unnecessary traffic on Core

Unwanted and unnecessary broadcast traffic was observed on the Core backbone VLAN1. A traffic and protocol analysis conducted on the switching Core for one day showed that there is a high level of unnecessary broadcast activity leaking onto the Core. This unnecessary broadcast traffic is mostly IPX SAPs from Synoptics hubs and Jet Direct printers. Also, AppleTalk ZIP and DEC traffic was observed on the Core. VLAN1 has been noted to spike up to 50% utilization levels and this is mainly from broadcast activity.

The Apple and DEC traffic is of little concern due to the volume but should be contained at the distribution layer. If the Core was not designed to be a pure IP Core then this activity may be acceptable. The traffic leaking onto the Core should be turned off or filtered.

See attached protocol analysis reports and charts in Appendix A.

Recommendation: Remove these protocols entirely from any components still utilizing them such as routers, Jet-direct printers and Synoptic hubs. If this activity is too much of a logistical issue to handle all at once, then implement filters on the Core and/or distribution routers to prevent these packets from leaking onto the Core segment.

2.9 Router System Buffers

Higher buffer allocation and failures on Route Switch Modules(RSM).

All of the Data Center Core routers exhibit low CPU utilization in the area of under 10% average. The RSMs have an average in the range of 15-20% utilization due to the distribution layer of traffic flows. The Core routers all show adequate system memory buffer allocation except for the Very Big buffers which are experiencing some failures. However on the RSM the buffer allocation activity is more severe as the example below shows:

```
DC2RSM1#sh buff
```

```
Buffer elements:
```

```
  499 in free list (500 max allowed)
 948747129 hits, 0 misses, 0 created
```

```
Public buffer pools:
```

```
Small buffers, 104 bytes (total 120, permanent 120):
```

```
  118 in free list (20 min, 250 max allowed)
 1174042461 hits, 319 misses, 503 trims, 503 created
25 failures (0 no memory)
```

```
Middle buffers, 600 bytes (total 129, permanent 90):
```

```
  127 in free list (10 min, 200 max allowed)
 762223760 hits, 14948 misses, 11408 trims, 11447 cr
6453 failures (0 no memory)
```

```
Big buffers, 1524 bytes (total 90, permanent 90):
```

```
  89 in free list (5 min, 300 max allowed)
 8506212 hits, 34 misses, 28 trims, 28 created
16 failures (0 no memory)
```

```
VeryBig buffers, 4520 bytes (total 10, permanent 10):
```

```
  10 in free list (0 min, 300 max allowed)
 17817128 hits, 1 misses, 0 trims, 0 created
 1 failures (0 no memory)
```

```
Large buffers, 5024 bytes (total 10, permanent 10):
```

```
  10 in free list (0 min, 30 max allowed)
 1 hits, 0 misses, 0 trims, 0 created
 0 failures (0 no memory)
```

```
Huge buffers, 18024 bytes (total 0, permanent 0):
```

```
  0 in free list (0 min, 13 max allowed)
 0 hits, 0 misses, 0 trims, 0 created
 0 failures (0 no memory)
```

A buffer miss occurs when there are no free buffers on the free list. A miss triggers the system to attempt to allocate more buffers so that the next time an attempt is made to get a buffer, one will be available. If the attempt to allocate more buffers is unsuccessful a **failure** is generated. A large number of buffer misses for a particular buffer size indicates that the minimum number of buffers and the number of permanent buffers could be increased.

Recommendation: Research into increasing the system buffer allocations for the buffers experiencing repeated failures and misses. The buffers should be monitored to see the frequency of the changes. The routers memory allocation must be considered when tuning such buffers so as to not over or under allocate memory from other processes. A Cisco TAC engineer should review the current buffer conditions to determine the optimal tuning parameters. Also, the resolution of the unwanted traffic leaking onto the Core segment as discussed in section 2.8 may provide some additional relief.

2.10 IP addressing scheme

The IP addressing scheme needs to be reviewed. NYC Utility is utilizing a Class B registered address of 158.57.0.0 and they are running out of address space since they are using a Class C mask. The networks are as contiguous as possible for future summarization. However, since the RIP protocol does not support Variable Length Subnet Masking(VLSM) many addresses can be wasted. NYC Utility implemented “unnumbered interfaces” to resolve this issue but there is still an address shortage. The advent of a classless VLSM based routing protocol such as EIGRP should help but the schema needs to be reviewed for future growth.

Recommendation: Review the address schema and plan for a classless routing protocol such as EIGRP. A possible restack of addresses may be in order to recoup any wasted addresses. The use of VLSM addressing for building/floor location and hierarchies for easier summarization and management should also be considered and planned for. Classless Interdomain Routing(CIDR) techniques should also be considered as well.

2.11 IP Secondary addressing

Removal of IP secondary addresses and the use of VLAN1 as the main network. This issue is relative to the IP addressing issues and NYC Utility’s current shared infrastructure migration. Router interfaces should not have secondary networks for this causes discontinuous addressing, routing protocol considerations and adds to the overall complexity and administration of the network. The use of VLAN 1 should only be used for VTP and/or SNMP management only. The Cisco Switches Supervisor modules are needlessly processing broadcast traffic on VLAN 1 as a result of all of the secondary addressing. Cisco recommends one subnet per VLAN and this could be accomplished with limited trunking at the distribution layer.

Recommendation: Continued migration of shared components to switches and planning the network to properly reflect one VLAN per subnet as per Cisco design standards. Also, keep VLAN 1 only for VTP traffic and as a component management subnet. This type of design facilities improved traffic management, flexibility, scalability, simplifies the network and improves administration.

2.12 Loop-back Interfaces

The use of Loop-back interfaces for management and utility purposes. The use of the virtual Loop-back interface could be used as the interface for sourcing the VLAN 1 subnet since it is always up. All components could source their SNMP traffic and other protocol operations that require an interface to always be active.

Recommendation: Research into the use of Loop-back interfaces for SNMP and VLAN 1 VTP traffic. A separate network address schema will need to be developed throughout the enterprise. The use of private address space 10.0.0.0 can be utilized and injected into the routing protocol. However this solution would work best if EIGRP is deployed throughout the enterprise.

2.13 RIP migration to EIGRP

The RIP routing protocol is outliving its capabilities within NYC Utility. The RIP protocol, although very efficient for smaller network, 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 these issues with 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.

Recommendation: Research and planning into migrating from RIP to EIGRP. This can be accomplished through SIN(Ships in the Night) routing with one EIGRP process. NYC Utility currently has EIGRP enabled on its Core routers and some redistribution is occurring. Pushing the protocol out to the edges will provide a more stable and scalable routing environment with increased convergence times during an outage. Also, VLSM and proper summarization can be implemented to reduce routing table sizes, simplify routing paths and conserve address space.

2.14 Selection of EIGRP AS number.

NYC Utility utilizes the AS number of 1932. NYC Utility stated that this number “was assigned” to them. Is this the same AS number used as their BGP AS? The EIGRP AS has only local significance for a process definition across all internal routers.

Recommendation: Review if the BGP and EIGRP AS numbers are the same and why.

2.15 EIGRP Logging

The use of logging EIGRP neighbor changes on the core routers running EIGRP should be active.

EIGRP relies on neighbors to provide a hybrid link state relationship to other routers. When neighbors change, due to connectivity, traffic, IOS or medium issues, the normal router logs will not show these events. It is critical that these events are logged not only for fault prevention but for troubleshooting as well.

Recommendation: Implement logging of EIGRP neighbor changes on all EIGRP enabled routers.

2.16 Debug Router

Implementation of a DEBUG ROUTER onto the Core and any other important segments for increased monitoring capabilities.

A Debug Router is a router with one interface into the main backbone Core segment and another interface through a back door segment. Its only use is to run debugging commands when issues are happening on the segment so as to not run such processor intensive commands on the production routers. Also, a Debug Router will show what is on the wire in terms of routing protocol traffic as seen/interpreted by a Cisco router, not a Sniffer.

Recommendation: Utilize an old 4000 router with 16-32mbs or RAM and two interfaces. Place the router in a strategic location on the Core and enable debugging when necessary to troubleshoot routing protocol issues et. al.

2.17 Terminal Server

Implementing a Terminal Server for all Core routers and switches for the use of local and remote reverse telnet access of these components through their console ports. This facilitates guaranteed access to the component, even if all interfaces are down on a component, through remote access at the console level.

Recommendation: Look into obtaining a Cisco based 2509 or 11 Terminal Server and define a strategic location for access. The terminal server should also have modem dial-up access as well.

2.18 Dial-In Access

Core routers and switches do not have direct modem dial-in access for support purposes. Without this capability on the main components a support exposure issue is present. For example, if Cisco TAC needed to access the routers directly to validate and troubleshoot an issue remotely, NYC Utility would currently have to set up and test modem connectivity during an outage. Also, if NYC Utility needed to remotely dial into a router directly then this facility can be achieved, especially if the Terminal Server solution, discussed in the previous section 2.17, is not implemented. The modems can be turned off during MTBF for security reasons.

Recommendation: Implement dial up access through the AUX ports on core components.

2.19 Network Management upgrade

The network management system needs to be upgraded to a more dynamic system to facilitate MAC and network map administration. The current system, Sun Net Manager, although very good, needs to be manually updated every evening and may no longer be supported by the vendor.

Recommendation: Look into utilizing either HP OpenView or Cabletron's Spectrum management platform as the base tool for NYC Utility's enterprise management requirements. CiscoWorks 2000 and Distributed Sniffer platforms can reside atop of the base platform thus providing a full view and access to all components for troubleshooting, monitoring and MAC activity.

2.20 T-1 Circuit Protocol Analysis

The WAN circuits currently do not have concurrent protocol analysis or Sniffer capabilities. Currently, when NYC Utility needs to sniff a WAN circuit, it has to be intrusive, thus breaking the link to install a remote pod. This activity limits the ability of support personnel to troubleshoot application issues, model transactions in real time and look into physical or protocol events before getting their carrier involved. Also several legs of a transmission path cannot be traced concurrently for application or device based issues.

Recommendation: Research into deploying distributed Sniffers for the major WAN circuits including the SONET ring, to increase visibility, transaction modeling and troubleshooting capabilities. If that approach is too economically prohibitive then obtain one or two WAN/LAN protocol analyzers with non-intrusive monitoring port access using Bantam cables. Applied Methodologies, has demonstrated such a tool and NYC Utility's CSU/DSUs by DataComm, Inc. support such non-intrusive operation. These units can be shared between LIS and NT.

2.21 CiscoWorks 2000

The use of CiscoWorks 2000 should be a priority in tandem to an upgraded management platform mentioned earlier in section 2.19. CiscoWorks 2000 will provide added capabilities for management and administration of all components. Some noted are as follows:

- Configuration backups every evening
- Configuration commands destined for many components can be snapped in and sent.
- IOS repository and distribution for upgrades/migration
- Device administration
- Traffic Director and VLAN director

The use of CWSI for the switched components also provides the ability to create a dynamic map of the switched network for management and administrative purposes. Also, the CWSI interface provides easier access and administration of VLANS for the Help Desk, LIS and NT.

Recommendation: Implement CiscoWorks 2000 for the enterprise, either initially as a standalone solution with the current Sun Net Manager platform or in tandem with an enterprise based solution. A complete dedication to time and effort is required for CiscoWorks 2000 to be implemented and utilized properly. This requires one or two person(s) to solely focus on this product for at least four months for this size of enterprise.

2.22 Router Memory Core Dumps

In the Core components, the use of core dump configurations commands should be considered. These commands enable the component to capture and send to another device the last state of the component before a failure. This information can be very helpful for troubleshooting device or IOS issues. Also, the support personnel can use a **SH STACK** command and send the information to Cisco or use the Cisco Stack Decoder on CCO after an outage to provide some clues as to the component's failure.

Recommendation: Implement **EXCEPTION** commands on core components.

2.23 Network Documentation

There are no hardcopies of an overall network documentation map that outlines subnets to components and location. A singular view map can help facilitate troubleshooting. This can also be accomplished with enterprise level maps and possibly using CiscoWorks 2000 and CWSI.

Recommendation: Creation of an overall network documentation map or utilize the newer network management tools to accomplish.

2.24 Change Control Approval

No confirmation of approval or disapproval of submitted changes to the helpdesk was observed during a Change Control exercise when updating a component's configuration. For example, when a changed is submitted to change a component's configuration no approval/disapproval is sent back to the personnel performing the change. A disruptive change can go through that should not have and cause an outage since changes are not reviewed.

Recommendation: Look into workflow of Change Control process for critical disruptive changes.

2.25 Multicasting Services

The use of Dense mode Multicasting in all components is not efficient if NYC Utility expects to increase its utilization of multicast based services in the future. This method of multicasting is based on the "push and prune" approach and thus creates a leave latency of over three minutes on each segment with unneeded multicast traffic. Also, Dense mode is not as flexible for shared sourcing of streamed media. The use of CGMP helps a little but the leave latencies are still high. Also, with the high number of components and the loop design in place causes Dense mode based traffic to travel further across lesser capable components.

Recommendation: The use and migration of Sparse mode for all multicasting traffic. This approach ensures that multicast traffic will be sent only when called "Pull and Prune". The use of shared multicasting trees facilitates the ease of distribution and administration of the streams. Also, with the use of CGMP and IGMP v.2 the leave latency is almost instant which results in less traffic on a segments and a quicker pruning of the multicast tree. Cisco supports a dual Sparse/Dense mode for migration purposes.

2.26 IOS version 12.x for enterprise routers

Migration to IOS 12.x GD version on Core and eventually on all components should be planned. The IOS 12.x release provides enhanced capabilities for performance, stability and QOS features. For NYC Utility to provide an enterprise class level network that supports QOS, streaming voice/video conferencing and additional traffic and administration options too numerous to list in this document, then this IOS must be deployed for the advanced features and stability.

Recommendation: Research into introducing IOS 12.x into the enterprise.

2.27 Network Time Protocol - NTP

The use of Network Time Protocol on all components so every component has a consistent synchronous clocking for logging of activities and configuration change management. This approach will ensure that all components are running against the same date/time clock for logging of activities and facilities troubleshooting and administration activities. Also manual clock setting in components is no longer necessary.

Recommendation: Look into deploying NTP on all components in the enterprise.

2.28 Unused Interfaces

The use of unused interfaces for a maintenance or backdoor segment should be considered if these interfaces are to never be used. These interfaces can provide a path into the component for monitoring and administration activities. This item can also be related to the Terminal Server recommendation discussed in section 2.17. The creation of a Console Segment should be considered as well.

Recommendation: Look into using unused interfaces on core components for utility purposes.

2.29 Router Reboot Schedule

All core components should be on a reboot schedule for six month windows to reallocate memory and flush all memory pools. Also, this activity can test the operational integrity of the devices and provide two already scheduled windows for major upgrades. This activity may also provide some relief to the System Buffer issues noted in section 2.9.

Recommendation: Devise a reboot schedule of Core components for maintenance.

2.30 Router Menu System

The implementation of a menu system in all routers for basic support personnel to use for simple and efficient trouble identification and isolation should be considered. If the support personnel cannot be trained for the CLI effectively then a menu system can be supplied below the network management of CiscoView to further facilitate their daily operation.

Recommendation: Research into the design and implementation of a router menu system for Help-Desk personnel. Cisco has a menu system already built into the IOS.

2.31 Access-Lists

The limited use of access lists for traffic leakages and security results in wasted bandwidth and more open control of all components. If NYC Utility prefers to keep the network simple and let traffic flow where it may not belong and just plan for over capacity then this may be acceptable. However, if NYC Utility wants to restrict different traffic types and limit certain broadcast types then a standard access list schema should be considered. Also for security reason on egress devices towards the Internet, access-lists on the routers should be considered. An example security based ACL is as follows:

```
access-list 106 deny tcp 220.11.97.0 0.0.0.255 any eq telnet
access-list 106 deny tcp 220.11.97.0 0.0.0.255 any eq chargen
access-list 106 deny tcp 220.11.97.0 0.0.0.255 any eq finger
access-list 106 deny tcp 220.11.97.0 0.0.0.255 any eq exec
access-list 106 deny tcp 220.11.97.0 0.0.0.255 any eq login
access-list 106 deny tcp 220.11.97.0 0.0.0.255 any eq cmd
access-list 106 deny tcp 220.11.97.0 0.0.0.255 any eq 2000
access-list 106 deny tcp 220.11.97.0 0.0.0.255 any eq 6000
access-list 106 deny tcp 220.11.97.0 0.0.0.255 any eq 87
access-list 106 deny tcp 220.11.97.0 0.0.0.255 any eq uucp
access-list 106 deny tcp 220.11.97.0 0.0.0.255 any eq domain
access-list 106 deny udp 220.11.97.0 0.0.0.255 any eq tftp
access-list 106 deny udp 220.11.97.0 0.0.0.255 any eq snmp
access-list 106 deny udp 220.11.97.0 0.0.0.255 any eq tacacs
access-list 106 deny udp 220.11.97.0 0.0.0.255 any eq sunrpc
access-list 106 deny udp 220.11.97.0 0.0.0.255 any eq 2000
access-list 106 deny udp 220.11.97.0 0.0.0.255 any eq domain
access-list 106 deny udp 220.11.97.0 0.0.0.255 any eq 2049
access-list 106 deny udp 220.11.97.0 0.0.0.255 any eq 6000
access-list 106 permit tcp 220.11.97.0 0.0.0.255 any eq ftp
access-list 106 permit ip any any
```

Recommendation: Look into a standard access-list for restriction of unnecessary traffic.

2.32 Network Switching Core Upgrade

The network switching Core and major locations(Brooklyn) comprising of the Cisco Catalysts 5500 platform is currently not heavily utilized but has the potential to be oversubscribed in regards to port density. Also, the underlying 5500 switching fabric does not scale very well beyond 3.2Gbs. The 5500 series also does not provide the level of QOS and advanced OSI layer switching as other enterprise level switches posses. These current limitations are adequate for NYC Utility's current needs but to achieve a robust and extremely efficient, flexible and tolerant switching core NYC Utility must consider a forklift upgrade of these devices.


Recommendation: Upgrade of core Catalyst 5500 switches to the 650x platform. In light of recent developments of a server farm re-architecture the Core should also be redesigned with enhanced flexibility and scalability in mind. This issue is also relevant to the IP addressing, VLAN distribution and EIGRP recommendations mentioned earlier in this document. Careful planning should be conducted for the Core and Server farm networks so that NYC Utility can provide a robust enterprise switching platform that supports all OSI layers and future needs such, MPLS, QOS, management, streaming audio/video and advanced applications. A set of core design consideration must be drafted for the Core and Server farm and several designs should be created according to such considerations. Also, Cisco should provide some additional designs and a complete review should be conducted to pick the best solution.

2.33 802.1Q/P Integration

The use of 801.Q/P should be considered over ISL for the standard trunking protocol. The DotQ encapsulation is an industry standard and requires less overhead via smaller encapsulation sizes, 4 bytes for DotQ as opposed to 30 bytes for ISL. Plus, DotQ interoperates with non-Cisco devices if required. Cisco switches supports both trunking encapsulations natively. This type of configuration positions NYC Utility’s switching fabric to also utilize the DotP Data Link Layer priority Class of Service(COS) capabilities. This level of COS at the Data Link layer can provide optimum throughput for selected applications. The Use GMRP for increased leave latency of multicast traffic is another available option for consideration.

See figure 3 below on priority packet features.

Figure 3.



Recommendation: Research into migrating from ISL to DotQ for trunking encapsulation and implementation of DotP for low level COS.

2.34 Backbone Switch & RSM Warning Messages

While looking through both of the backbone switches and RSMs the following message appeared from both DC BB switches and RSMs.

```
DC2_BB1_5500> 2001 Apr 18 11:08:49 %MLS-4-MOVEOVERFLOW:Too many
moves, stop MLS
for 5 sec(20000000)
2001 Apr 18 11:08:54 %MLS-4-RESUMESC:Resume MLS after detecting too
many moves
```

Here is the Cisco explanation.

```
MLS-4-MOVEOVERFLOW: Too many moves, stop MLS for [dec] sec ([hex])
```

Explanation This message indicates that there have been too many Layer 2 source address changes during a short period of time. This message might be caused by topology changes or spanning-tree loops. MLS is stopped for [dec] seconds and all entries are purged; [dec] is the number of seconds, and [hex] is the event code for moves.

Action Check the topology for any loops. Call your technical support representative

```
MLS-4-RESUMESC: Resume MLS after detecting too many moves
```

NYC Utility is aware of this activity but it should be investigated to prevent an outage. This may be indicative of a larger issue or the RSMs are not keeping up. Also, this activity could be the result of a software revision upgrade.

Recommendation: Investigate this warning message and resolve.

3.0 General Staff Observations

A weekly issues and design meeting to list new issues and design considerations should be conducted between members of the Help-Desk, LIS and NT, if one is not already conducted. This type of meeting enables all parties to be informed of issues and each other's activities.

Network Team(NT) needs the bandwidth to set the direction for where the network is going in terms of growth and functionality. NT cannot currently do this for it is in constant tactical mode while the network morphs into many different shapes. Network Team needs to create a corporate wide manifesto for the network and start to plan strategically to achieve NYC Utility's ongoing requirements. A weekly networking design meeting with all members will help identify issues, foster ideas, and provide all members Jr/Sr. of the group an opportunity to learn and hone new skills by thinking critically. Plus, a feeling of teamwork and being part of the overall corporate direction prevails among all team members.

To accomplish this either LIS or the Help Desk must assume more of the daily tactical issues or additional staff in NT may be required to balance the load among all members of the team. It is critical that this be done soon for all of the issues listed in this report need to be addressed otherwise NYC Utility will have a very difficult time migrating their network to enterprise class status if NT continues to spend a majority of their time performing tactical day to day activities.

The NT staff is very talented and has the capacity to grasp and learn the issues listed in this report easily. The team also shows a strong willingness to move forward and apply creative solutions to complex problems.

An endless cycle is upon NYC Utility that is a major trend in many corporations. The NT group is currently too busy to handle the house cleaning and strategic movement of the network to alleviate these issues, thus the strategic movement of the network never commences. Tactical changes are always applied with the assumption that they apply to some informal direction or meet a current need. However, if resources are pulled to provide such strategic planning, daily fires continue to crop up and not enough resources are available to handle the daily issues. Thus, resulting in the strategic initiatives being put on hold again.

This cycle will continue until personnel is added and a clear demarcation of roles is established. Also, a process of network design requests must be in place and enforced so a buffer is available for NT to focus not only on the tactical issues but, more importantly on strategic goals.

A network design planning manifesto and a strategic planning group within NT to focus solely on the network's direction from the standpoints of high level, operational, configuration standards and overall standards should be created.

4.0 Future stuff

4.1 Broadband Technologies

Since NYC Utility owns many fiber runs between it's various offices in the metropolitan or Tri-state area NYC Utility should consider deploying a DWDM MAN for increased bandwidth, reduced costs and provide enterprise class services easily across the area.

A DWDM based MAN can provide over 10Gbs. of bandwidth in basic Ethernet encapsulation, so NYC Utility can use just what they need and add more bandwidth when necessary. A DWDM MAN between 4IP, Flatbush, TLC, RYE and VAN Nest could provide significant savings over the current T-1 and SONET implementation, plus reducing SONET overhead signaling tax and simplifying the network. Also, the switching fabric can be expanded to all of these locations thus creating an efficient Ethernet switching matrix across the enterprise (not just in 4IP and Flatbush) for advanced IP based applications and services.

Recommendation: Research into NYC Utility's next generation MAN utilizing existing fiber and DWDM broadband technologies. An economic and functional ROI report must be drafted to understand the validity of using this resource to save NYC Utility money and enhance their network's overall capabilities.

4.2 Forensics Protocol Analysis

Training on Agilent Advisor, T-1 troubleshooting, Sniffer and forensic protocol analysis for Network Team personnel to increase their understanding of protocols, operation, theory and how to utilize a Sniffer for troubleshooting, application impact and benchmarking analysis.

Recommendation: Conduct protocol analysis labs to enhance support personnel skills and knowledge.

