



February 4, 2005

To: CTS, CQS, and OPRA Multicast Data Recipients

Subject: New NMS Output Communications Network Source Addresses

SIAC is in the process of upgrading the NMS Output Communications Network, which includes major hardware, software and communications infrastructure changes.

The purpose of this notice is to advise all Data Recipients that there will be new source addresses associated with the NMS network (i.e., CTS, CQS, and OPRA; see Appendix B of the attached IP Multicast Distribution Network Recipient Interface Specification, Version 1.24, dated January 20, 2005). Therefore, Data Recipients who currently employ mechanisms for controlling NMS data flow, such as firewall and/or router filter configurations, may need to make modifications to permit the new source addresses. All Data Recipients will need to verify that they can process data with **both the new and existing** source addresses.

Support of CTS, CQS and OPRA Data Recipient testing with the new network source addresses will begin on February 28, 2005. CTS, CQS and OPRA data is scheduled to begin to be disseminated in production using these new addresses, in June, 2005. If the necessary changes are not made, Data Recipients may not receive data.

Data Recipient test scheduling should be coordinated through Joe Loughran at (212) 383-4908 or Cheryl Boamah at (212) 383-4920.

Please contact Laura Guzzy at (212) 383-8373 or Peggy Liao at (212) 383-2931 if you have any questions regarding these modifications.

Sincerely,

Michael Collazo
Director
NMS Product Planning & Management

cc:

CTA/CQOC Policy Committee	A. Bach
CTA/CQOC Technical Committee	B. Faughnan
OPRA Policy Committee	M. Lewis
OPRA Technical Committee	D. Mrakovic

National Market Systems

Common IP Multicast Distribution Network

Recipient Interface Specification

Prepared by:



Communications Engineering
Planning And Development

Document Number: ml10182000
Date: January 20, 2005
Revision: 1.24

Copyright Notice

Copyright Ó 2005 by the Securities Industry Automation Corporation (SIAC). All Rights Reserved. Except as permitted under the United States Copyright Act of 1976, no part of this document may be reproduced or distributed in any form or by any means, or stored in a data base or retrieval system, without the prior written permission of SIAC.

Brand names and /or Trademarks

Brand names or Products cited in this document may be trade names or trademarks. Where there may be proprietary claims to such trademarks or trade names, the name has been used with a initial capital. Regardless of the capitalization used, all such use has been in a editorial fashion without any intent to convey endorsement what so ever of the product or trademark claimant. SIAC expresses no judgment as to the validity or legal status of any such proprietary claims.

Engineering Services Disclaimer

Information contained in this document is believed to be accurate. However SIAC does not guarantee the completeness or accuracy of any of the published information. This work is published with the understanding that SIAC is supplying information, but not attempting to render engineering or other professional services. If such services are required the assistance of the appropriate professional should be sought.

REVISION LOG

Document Number: ml10182000.doc

**Title: National Market Systems Common IP Multicast Distribution Network
Recipient Interface Specification**

Version	Date	Rev by	Pages affected	Comments	Approval
1.2	3/13/97	ML		Initial Release	
1.3	11/24/97	ML	18	Typo, naming conformance issue	
1.4	12/15/1998	RL	All	Remove references to Bisync and make document present tense with respect to the NMS network; Remove appendix on required bandwidth	
1.5	12/03/99	MC	1	Added references for retransmission and playback data	
1.5	12/03/99	MC	19	Added Retransmissions & Playback IP Group Assignments	
1.9	3/23/00	RL	All	Clean up and Reorganize document. Removed section on logical lines because it served no purpose Added more information on multicast protocols Remove references to Frame Relay support New IP source addresses for RAPs and MPR boxes added as an appendix	
1.10	6/16/00	RL	Appendix C	Added new RAPS IPs for 2 new hosts: RAPSOPRA3 and RAPSOPRA4	
1.11	10/5/00	RL	All (major)	Add time beacon specifications; add new CTS and CQS group numbers	
1.12	10/18/00	RL	All (minor)	Incorporate review comments, fix page numbers	
1.13	7/11/01	RL	Appendix C	Added Appendix C.2.	

Version	Date	Rev by	Pages affected	Comments	Approval
			Appendix C.2	Added text to Appendix C.	
1.14	11/15/01	RL	Entire document	Updates to reflect interface types available on a per service basis. T3 connectivity no longer available to new connections or upgrades.	
1.15	12/04/01	RL	Appendices C and C.2	New OPRA addresses are in production and therefore deleted Appendix C. Changed name of Appendix C.2 to C.	
1.16	08/28/02	RL, CE	All	Removed Legacy Options	
1.17 1.18				Internal draft update, not distributed	
1.19	12/17/02	RL, CE		Include BBO info	
1.20	1/03/03	RL		Correct Typo in Appendix A	
1.21	3/06/03	RL, CE		Updated IP addresses for Multicast playback and retransmissions.	
1.22	8/19/03	RL, CE		Removed non-BBO lines.	
1.23	1/12/04	RL, CE		Transitioned to SFTI interconnection. Deleted Appendix B, renamed appendix C as B, and Appendix D as C.	
1.24	1/20/05	LG	Appendix B; Throughout Document	Update of all source addresses. Removed reference to Site A and Site C. Replaced with Group A and Group B.	

Table of Contents

1	Overview	1
1.1	Data Available via the NMS Network.....	1
1.2	Multicast Primer.....	3
2	NMS Data Types	4
3	Application Considerations	5
3.1	Application Encapsulation	5
3.1.1	Time Beacon Message Encapsulation.....	5
3.2	End to End Data Integrity.....	5
3.3	Line Concept	6
4	Network Layer Connectivity	6
4.1	IP Multicasting – Primer Part II	6
4.1.1	Unicast IP Routing	6
4.1.2	Multicast IP Routing	6
4.2	Multicast Addressing.....	7
4.3	UDP/IP Framing.....	8
4.3.1	IP Header Field Descriptions	8
4.3.2	UDP Header Field Descriptions	9
4.4	Multicast Address Use.....	9
4.4.1	IGMP	9
4.4.2	Subscription Control	9
4.4.3	How Multicast Delivery is Implemented via SFTI.....	10
4.4.4	Multicast Data Retransmission.....	11
4.4.5	Availability of Multicast Services	11
4.4.6	Multicast Transport Protocol	12
4.5	Logical Groups Mappings Versus Physical Access Points.....	12
4.6	Data Entitlement.....	12
4.7	IP Addressing Considerations	12
4.8	Recipient Security	12
5	Physical, Media Layer, and Network Connectivity	12
6	Appendix A - NMS IP Multicast Addresses	13
7	Appendix B - NMS IP Source Addresses	17
8	Appendix C - Time Beacon Message Format	21

Table of Figures

Figure 1 IP Data Block Format.....	5
Figure 2 UDP/IP Datagram Format.....	8

1 Overview

This document provides the interface specifications for customers connecting to the National Market Systems (NMS) distribution network. This includes recipients of the Consolidated Tape System (CTS), Consolidated Quote System (CQS), and the Options Price Reporting Authority (OPRA) real-time production data. Recipients should also use this specification for information related to receiving NMS real-time data retransmission, NMS after-hours playback data, and Time Beacon messages. The NMS distribution network disseminates all market data and Time Beacon information in the form of multicast addressed IP datagrams.

With respect to physical network connectivity, all data distributed by the NMS systems requires recipients to connect via the Secure Financial Transaction Infrastructure (SFTI). Recipients may connect directly to the physical edge of SFTI, or receive data via a third party value added service provider. For those choosing to connect directly to SFTI, a separate SFTI interface specification, “SFTI Network Interface Specification for directly connected Customers” should be referenced. For more information on the SFTI network and the services available via SFTI, please visit the SFTI website at <http://sfti.siac.com>, contact a SFTI Customer Support representative at SFTI@SIAC.com, or call 1-866-USE SIAC.

1.1 Data Available via the NMS Network

Throughout the remainder of this document there is reference to redundant data streams such that each NMS message is duplicated and available via two separate UDP-based IP Multicast groups, one from Group A and one from Group B. Group A and B refer to SIAC’s data centers. Prior to the installation of SFTI, there were just two geographic points to which users could connect to the NMS network such that users would receive one set of multicast groups from each point, i.e. distribution choices were hardwired to restrict by originating site.

However, with the advent of the SFTI network, there are now several points to which to connect and receive NMS data via SFTI. As a result, now that there are more than two access points for receiving NMS data, directly connected recipients can coordinate with SIAC and choose which multicast groups they wish to receive via each of the SFTI connectivity points.

NMS Real-Time Production Data

Two copies of each NMS real-time production message are available from SIAC’s operational sites, Group A and Group B. These redundant copies are delivered via two distinct multicast data streams. For each unique NMS line (CTS 1-4, CQS 1-5, and OPRA 1-9) there are two redundant multicast data streams. SIAC refers to these streams as the ‘A’ and ‘B’ streams. The ‘A’ stream is available from Group A and the ‘B’ stream from Group B. See Appendix A for the table of multicast group mappings.

NMS Real-Time Data Retransmission Data

The retransmission data streams are available via both sites, but are not delivered via redundant data streams. The recipient may choose to receive the retransmission data from either or both sites. See Appendix A for the table of multicast group mappings.

NMS After-Hours Playback Data

Playback data is available in two ‘flavors’:

- There is a set of Multicast data feeds dedicated for after-hours playback test data. This playback data is made available via a single set of multicast data streams and can be obtained from either site.
See Appendix A for the table of multicast group mappings.
- In addition to the playback test data groups, SIAC will continue to provide dual-sited redundant after-hours playback via the production system expressly for the purposes of redundancy testing.

NMS Network Time Beacon

Each system that sources multicast data within the NMS network generates a single Time Beacon packet once a minute. Five of these systems are located at each site. Each set of 5 will issue Time Beacon packets to the same multicast group, therefore there are two NMS Time Beacon multicast groups.

See Appendix A for the table of multicast group mappings.

Recipients may subscribe to these packets and use the enclosed time stamp for several functions including:

- Verifying the ability to subscribe to and receive multicast data sourced within the NMS network. The Time Beacon is available 24 hours a day, 7 days a week, except during occasional off-hours maintenance periods. These time packets therefore can serve as a “heartbeat” message for indicating that the multicast routing protocols are functioning and that the systems are available.
- Verifying the ability to receive multicast data from all ten NMS multicast source systems, which includes all the primary and backup systems.
- Synchronize to a time source accurate to within 1 second of the Global Positioning System (GPS).

The GPS is a U.S. Department of Defense developed, worldwide, satellite-based radio-navigation system. This system provides time transfer to Coordinated Universal Time (UTC) and is distributed to the NMS systems via redundant Network Time Protocol (NTP) servers. The NTP servers are connected directly to GPS based time clocks located at SIAC. These clocks receive GPS data via directly connected satellite dishes.

Bandwidth Requirements

Bandwidth requirements change with time and recipients are encouraged to contact the SFTI Help Desk and/or NMS planning representatives regarding bandwidth requirements of each of the NMS services. For more information on the SFTI network and the services available via SFTI, please visit the SFTI website at <http://sfti.siac.com>, contact a SFTI Customer Support representative at SFTI@SIAC.com, or call 1-866-USE SIAC.

Data recipients should connect to NMS via SFTI, via at least two access points in order to make full use of the resiliency of SFTI and the redundant data feeds available for each service. Recipients not connecting directly should consult with their value added service provider regarding connectivity options.

Additional considerations for all recipients:

In total, the ten Time Beacon sources contribute a relatively insignificant data rate requirement; (approximately 720Bytes/minute or <100bits/sec).

Message Formats

For details of the message formats utilized by the CTS, CQS, and OPRA systems, please reference the following:

- CTS: CTS, Consolidated Tape System, Output High Speed Line, Interface Specification
- CQS: CQS, Consolidated Quote System, Output High Speed Line, Interface Specification
- Options Price Reporting Authority - Data Recipient Interface Specification

Please note that the message format of retransmission and playback data is also governed by the documents listed above.

The remainder of this specification addresses the communications interfacing requirements for all data types and also includes the message formats for the Time Beacon in Appendix C.

1.2 Multicast Primer

In a nutshell, Multicast is a form of subscription based IP broadcasting. In a traditional broadcasting environment, data is sent out on all links to all LANs (or sub-networks). In contrast, IP Multicasting provides a method for selective delivery of the data via a subscription-based protocol known as the Internet Group Management Protocol (IGMP). The local end-stations (e.g. application hosts) are typically responsible for issuing IGMP requests that are processed by the host's local intermediate-stations (e.g. routers/switches). In response to these IGMP requests a multicast capable network need only deliver the multicast data to those portions of the network that lay in the path between the subscribing host and the original source of the data.

Subscriptions are based on the target multicast group id (which is synonymous with multicast address and multicast host group). The NMS distribution network currently utilizes 36 multicast group ids for the production data, 18 groups for the retransmission data, 18 groups for the after-hours playback data, and 2 groups for the Time Beacon messages.

Those unfamiliar with multicast technology are encouraged to reference RFC 1075 -The Protocol Independent Multicast-Sparse Mode (PIM-SM)), and RFC 2117 and RFC 2362 - Host Extensions for IP Multicasting (which includes the Internet Group Management Protocol (IGMP). Also of notable assistance is the text titled “TCP/IP Illustrated, Volume I” by Richard M. Stevens which provides several sections detailing multicast protocols and IGMP.

Recipients are strongly recommended to consult the SFTI interface specification, which provides additional information and considerations for receiving multicast services via SFTI.

2 NMS Data Types

The NMS network distributes data via the multicast addressing and delivery protocols. Each of the three systems (CTS, CQS, and OPRA) has a unique set of multicast addresses assigned to each of its data “lines”. In each case there are redundant data streams provided for daytime production delivery of each line. There are currently 18 different data lines in total for the three systems which include CTS1-4, CQS1-5, and OPRA1-9. Therefore there are 36 unique multicast groups allocated for the redundant delivery of these 18 lines ($18 * 2 = 36$).

Each system line has a single retransmission multicast address assigned to it (18 in total). Entitled recipients have the option of subscribing to any retransmission line as needed.

There are 18 additional multicast group ids allocated for supporting after-hours playback of the NMS data. Note that after-hours playback can also utilize the 36 production multicast groups.

There are 2 additional multicast group ids used for distributing the Time Beacon messages.

In all cases, recipients will only be permitted to receive data to which they are entitled. Note that all recipients are able to receive the Time Beacon messages.

Appendix A provides tables listing of all multicast group ids. The tables also include a listing of the UDP destination port numbers assigned to each data stream. The NMS distribution system utilizes the UDP protocol at the IP transport layer. In order to provide the recipient community with the highest level of flexibility, the NMS systems have assigned a unique UDP destination port number to each multicast data stream. Note that the real-time redundant data streams use unique multicast addresses at the IP layer and unique UDP destination port numbers at the transport layer.

Recipient application software may make use of the UDP port mappings in order to multiplex between each of the datastreams. Typically, applications use a “socket” programming interface which provides the means for requesting data on a per UDP port basis. If a port is not specified the application host’s operating system might pass all IP broadcast data (including all multicast data) to a single process if the application has not specifically requested data on a per port basis. Please consult your application host’s programming and system documentation for information particular to your environment.

3 Application Considerations

This section defines the application data framing and some of the key aspects of the IP distribution environment.

3.1 Application Encapsulation

In the IP environment the NMS application messages are also encapsulated in blocks, which in turn are encapsulated in an Ethernet frame as given in Figure 1 IP Data Block Format.





Figure 1 IP Data Block Format

There are actually several levels of encapsulation that occur within the Ethernet frame. The NMS data block, (which can be a maximum of 1000 bytes), is encapsulated within a UDP datagram, which in turn is encapsulated within an IP datagram which itself is encapsulated within an Ethernet frame. Each number shown in parentheses, e.g. IP (20), refers to the size of the particular header in bytes.

3.1.1 Time Beacon Message Encapsulation

The Time Beacon message encapsulation is illustrated below. Note that the message format for the Time Beacon is included in Appendix C of this document.



3.2 End to End Data Integrity

Integrity checking, on a per packet basis, is available via a checksum value in both the UDP header (Figure 2 UDP/IP Datagram Format) and the Ethernet frame check sequence.

In general, the Ethernet frame checksum validation is performed by the host's interface firmware and the IP checksum validation is performed within the TCP/IP stack and not by the application software.

Unlike TCP/IP based application services, the UDP/IP protocol has no "built-in" automatic retransmission functionality and therefore recipient host applications must examine the sequence numbers embedded within each NMS message on a per line basis in order to determine whether any data has been missed.

3.3 Line Concept

The term “LINE” refers to a specific logical data stream identified by the value pair formed by a unique IP multicast destination address and unique UDP destination port number.

Note that the following terms are all analogous to each other:

- multicast group
- multicast group id
- multicast host group
- multicast host group id
- multicast destination address

The NMS network currently utilizes 72 unique multicast group ids for the purposes of providing NMS data to the recipient community. Each multicast group id also has a UDP destination port number assigned to it, therefore each line of NMS data is uniquely identifiable by the value pair formed by its multicast group and UDP destination port number pair.

Appendix A provides the exact mappings of each line to its identifier pair. The list below summarizes those tables.

- 52 Production data streams for day-time dissemination (2 sets of 26 redundant data streams)
- 26 Production retransmission streams for day-time dissemination (1 set of data streams)
- 26 Playback test data streams for after-hours support (1 set of data streams)

The concept of “lines” does not apply to the Time Beacon. Each system in the NMS network that sources multicast also sources a single Time Beacon message once a minute. Currently there are ten such systems and these messages will be staggered to result in approximately one Time Beacon message every 6 seconds. Five of these messages will be destined to one multicast group id, and five to one other.

4 Network Layer Connectivity

4.1 IP Multicasting – Primer Part II

The Internet Protocol suite, referred to as IP, defines a data encapsulation method that allows data to traverse multiple networks through intermediate network devices known as routers.

4.1.1 Unicast IP Routing

Typically, IP packets are issued from a source host with a single destination host as the target. This type of addressing is usually referred to as “unicast addressing”. Unicast addressed packets are routed by intermediate-stations (i.e. routers) based on the destination network number associated with the destination IP address listed in the IP header portion of the packet. The intermediate-station compares the destination with its local IP routing table and forwards the packet to the appropriate next hop device (router) or to a local host if the router is local to the destination network.

4.1.2 Multicast IP Routing

In contrast, IP multicasting uses a special class of IP addresses that are used to represent a “host group”. These addresses are referred to as Class D and fall in the range of 224.0.0.0 to 239.255.255.255.

The host group id is both an actual number and a concept. It can refer to the actual Class D IP address that is placed in the IP header’s destination address field of the IP multicast packet. It also refers to the

protocol's concept of a host group. A host group represents all end-stations, (or hosts), that have specifically subscribed to the multicast host group id. The subscription functionality and the multicast routing protocols provide the underpinnings that enable a single multicast addressed packet to be delivered to all LANs connected to at least one host that has subscribed to the host group in question.

Each multicast packet sourced by an originating host is forwarded by the local intermediate-stations supporting the multicast routing protocols. Intermediate-stations replicate and forward the multicast packets out each of its interfaces that meet one of the following two criteria.

1. The interface is directly connected to a LAN where a member of the host group is attached
2. The interface connects to, either directly or via a shared LAN, to any neighboring routers that lies in the path between a subscribing host and the host that originally sourced the multicast packet

In the NMS network, there may be as many as 56 unique host groups available at any given time during daytime production. This includes the 36 real-time production lines, 18 retransmission lines, and the 2 Time Beacon groups.

4.2 Multicast Addressing

Multicast addresses are known as Class D IP addresses and range from 224.0.0.0 to 239.255.255.255 (using standard IP address notation). The addresses in the range of 224.0.0.0-224.0.0.255 are reserved for local multicast and are non-routable.

The NMS network uses the following ranges, which are presented in further detail in Appendix A. Note that not all these addresses are in use.

224.0.2.192 - 224.0.2.255
224.0.5.128 - 224.0.5.159
224.0.5.176 - 224.0.5.191
224.0.5.240 - 224.0.5.255

4.3 UDP/IP Framing

The application data is encapsulated in an UDP/IP frame as shown in Figure 2 UDP/IP Datagram Format. The IP datagram includes the IP and UDP headers plus the application data. The datagram fields can be read left to right starting at the top and working your way down through the datagram. The size of each field (excluding the UDP data field) is represented in bits across the top and bytes going down. Bits are transmitted across the link starting with bit 0 ,1 ,2 and so forth. This is called the “big endian” representation where the most significant bits are transmitted first.

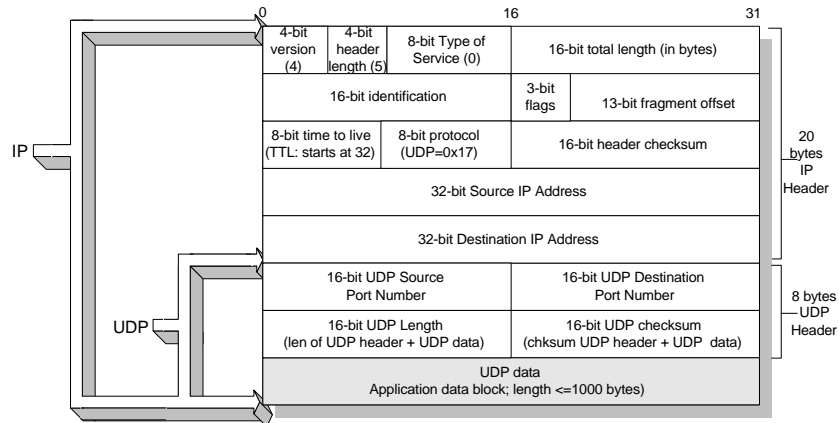


Figure 2 UDP/IP Datagram Format

4.3.1 IP Header Field Descriptions

- **Version** - This is a 4 bit field which defines the current version of the IP protocol. It is currently set to 4.
- **Header Length** - This 4 bit field contains the number of 32 bit words in the IP header portion of the datagram. For all multicast packets being generated by this network the IP header will be 20 bytes long, which means this field will contain the value 5.
- **Type of Service** - The first 3 bits are the precedence sub field and are ignored by most network equipment. The next four bits are flags that define minimize delay, maximize throughput, maximize reliability, and minimize monetary cost respectfully. They are set to zero (0) for this application. The last bit is always set to zero. Based on this description this field will always have the value of zero (0) for all multicast packets.
- **Total Length Field** - This 16 bit field contains the length in bytes of the entire IP datagram. This includes the IP and UDP header plus the application data (UDP data). Since the maximum size of the application data is 1000 bytes, the maximum value for this field is 1028.
- **Identification Field** - This 16 bit field contains a value that is incremented by one for each packet sent by the source system . It only has relevance on the receiving system when packets are either fragmented and/or TCP is used as the transport protocol. IP multicast packets use UDP and will not be fragmented by the multicast distribution network.
- **Flags and Fragment Offset** - The combined 16 bit field is only used when an IP datagram is fragmented. The multicast distribution network will not be fragmenting the data packets.

- **Time to Live (TTL)** - This 8 bit field contains a value that determines the number of routers that this datagram can pass through. Each router that forwards this datagram will decrement this value by one; when it reaches zero the next router throws it away. It is initially set to 32 by the multicast source systems.
- **Protocol** - This 8 bit field contains a value representing the next level encapsulated protocol. In this case it is UDP, which has a value of 0x17, which is 23 decimal.
- **Header Checksum** - This 16 bit field contains a checksum made up of the IP header fields only. The calculation is based on the ones complement sum of the header broken into 16 bit words.
- **IP Source Address** - This 32 bit field contains the IP address of the multicast datagram source system.
- **IP Destination Address** - This 32 bit field contains the IP Multicast Group address designated for this “line” (*see section 3.3*) of data packets. For the mapping of IP multicast group addresses to data lines please consult Appendix A of this document.

4.3.2 UDP Header Field Descriptions

- **UDP Source Port Number** - This 16 bit field identifies the sending process within the multicast source system. It is set by the source system.
- **UDP Destination Port Number** - This 16 bit field identifies the UDP process that should receive this datagram in the recipients receiving system. It will be uniquely set by the multicast source system based on the “line” of data being encapsulated within the packet. For the mapping of UDP port numbers to data lines please consult Appendix A.
- **UDP Length** - This 16 bit field contains the length in bytes of the UDP header plus the application data (UDP data). Its maximum value is 1008.
- **UDP Checksum** - This 16 bit field contains a checksum made up of the UDP header plus the application data (UDP data). In addition it also includes a UDP “pseudo” header, which is made up of selected fields from the IP header (IP Source Address, IP Destination Address, Protocol and UDP Length). The calculation is based on the one’s complement sum of the datagram broken into 16 bit words.

4.4 Multicast Address Use

The multicast group addresses used by SIAC for the dissemination of application data on this network have been registered with the Internet Assigned Numbering Authority (IANA). No recipient will be allowed to connect to the NMS distribution network if it is found that they are using any of these addresses for their own use.

For a list of these addresses please view <http://www.iana.org/assignments/multicast-addresses> .

4.4.1 IGMP

Internet Group Management Protocol (IGMP) is a protocol that end systems use to communicate with multicast compliant routers and is defined in RFC 1112. Recipient host systems that wish to subscribe to multicast groups must be fully compliant with this RFC.

4.4.2 Subscription Control

In order to receive the multicast packets, applications running on recipient end-stations issue IGMP subscription (or “join group”) packets on their locally attached LANs. The local router (which must also be

multicast compliant) adds the multicast group to its registration table and begins to forward all packets destined to that group onto the LAN.

Recipients have the option of subscribing to any combination of multicast groups but as mentioned previously, SIAC will allow recipients to receive only those groups to which they have been entitled.

4.4.3 How Multicast Delivery is Implemented via SFTI

As explained in detail by the SFTI interface specification, the SFTI architecture includes providing access to the NMS services via the use of the 802.1Q protocol, which provides for the definition of logically separate virtual LANs, or VLANs. In SFTI there is a single VLAN configured for transporting the aggregate multicast traffic.

In order to facilitate the delivery of Multicast data, SFTI must employ the use of a multicast routing protocol. SFTI uses Protocol Independent Multicast (PIM) to accomplish this task.

As the SFTI specification describes, customers will have two methods for receiving multicast data from SFTI. That specification refers specifically to the configuration of the customer router port connected to SFTI. Customers can implement any network solution they wish beyond that interface. Beyond the SFTI demarcation point, SIAC places no restrictions on the manner in which a customer designs its networks to support multicast reception. This is true from both from a protocol and physical topology perspectives. Customers are responsible for implementing a working design that best suits their environments.

The following applies to customers connecting directly to SFTI, and though it may also apply to customers connecting via a third party value added service provider, customers must consult with that entity with respect to specifications for receiving multicast data because their service offerings may deviate from the following.

Method I: For Customer routers supporting PIM Sparse-Dense Mode

- Configure PIM Sparse-Dense Mode on the router that connects to SFTI.
- Use “auto-RP” to learn the SFTI RP addresses and multicast group mappings.
- Configure RIP2 in listen mode to learn the routing information for the multicast source networks and the routes to the PIM RPs.

Method II: For Customer routers unable to support PIM Sparse-Dense Mode

- Customers can use PIM Sparse or Dense mode
- SIAC, upon the request of the customer, will define IGMP static joins on the SFTI edge router connected to the customer. This will result in statically forward all entitled multicast groups to the customer edge router;
- Customer routers learn multicast source routes by listening to RIP2.
- Customers can implement whatever solutions they require on their edge router in order to correctly forward the multicast data into their networks. Typically, router vendors provide the option of importing the multicast data at the edge into their routing trees using the routing information learned via RIP2. Some customers might implement “multicast proxies”, which presumably would translate the header information of the multicast datagrams into unicast UDP destined to one or more end-stations within the customer network. As is the case with everything described within this specification, customers must check with their chosen vendor for protocol support and recommended solutions.

4.4.3.1 Multicast Entitlement Control

Multicast entitlement will be enforced at the SFTI Edge Routers by application of PIM join filters on the logical interface (and VLAN) connected to each individual Customer. The use of filters allows for the control of transmission/reception of multicast groups. Different customers will have different definitions based on their service entitlements. For those customers where SIAC has defined static IGMP joins on the SFTI edge, SIAC will by definition use the static joins to control entitlement.

Ingress traffic filters on the Edge Router logical interfaces (VLAN) supporting multicast will silently discard any incoming packets except those used by the multicast (PIM Sparse-dense mode) or unicast routing protocols. These filters will also be used to protect SFTI from any customer-originated multicast traffic.

SIAC can reconfigure these filters dynamically to allow for timely re-provisioning of entitlements.

4.4.4 Multicast Data Retransmission

Some of the multicast services offered via the various SIAC Financial Services Networks (FSNs) provide an inband retransmission request mechanism via unicast UDP based applications. These types of transmissions will not be supported via the same logical interfaces on which the Customer is receiving the multicast data. Unicast based retransmission requests will be routed handled by the unicast VLAN logical interface for the particular FSN involved. For example, CAP retransmission requests for multicast services will be handled by the CAP unicast VLAN, not by the multicast VLAN. This traffic will be transported through SFTI in the same manner as other unicast traffic to the particular destination FSN.

Inband retransmissions are not currently offered via the NMS Distribution Network, but plans are in place to provide this service in the future.

4.4.5 Availability of Multicast Services

Customers will receive a list of the multicast source networks, multicast destination group addresses, and all other relevant information from SFTI Customer Service once the customer becomes a licensed subscriber.

The multicast group addresses used by SIAC for the dissemination of application data on this network have been registered with the Internet Assigned Numbering Authority (IANA).

4.4.6 Multicast Transport Protocol

SFTI IP multicast datagrams will use the connectionless UDP protocol at the transport layer.

4.5 Logical Groups Mappings Versus Physical Access Points

In order to provide a resilient/redundant distribution environment for the recipient, the recipient is provided with the ability to connect to SFTI at several geographically diverse access centers. As of January 2004, there were eight operation access centers, including four in the New York Metro area, two in Chicago, IL, and two in the Boston, MA metro area.

As mentioned previously, each NMS message is provided via redundant data streams for the purpose of allowing recipients to leverage the redundancy of SIAC's data centers. Each multicast group is available via any and all of the SFTI access centers.

4.6 Data Entitlement

For a recipient host system to receive a particular data stream it must subscribe to the data stream's corresponding multicast group id via IGMP. Appendix A lists all multicast group id assignments.

In order to restrict a recipient from subscribing to data streams that they are not entitled to, outbound packet filters are employed on SIAC's distribution routers interfaces connecting to the recipients. These filters block data from being sent to non-entitled recipients on a per service basis (CTS, CQS, and OPRA).

4.7 IP Addressing Considerations

Please consult the SFTI interface specification for details.

4.8 Recipient Security

SIAC protects its network and hosts using several methods. Traffic filters and routing policies prevent sharing of information and data between entities connected to the SFTI network. Additional measures are in place as well, however these security measures maintain the integrity of SIAC's distribution environment by protecting SIAC's network and hosts from intentional or accidental access from within a recipient network.

These measures are in no way intended to provide the same level of security to the recipients themselves. If a recipient believes that additional security is required to protect their network they are encouraged to take action to implement additional security measures.

For the purposes of aiding in the implementation of security measures (e.g. traffic filters), the source IP addresses associated with the NMS systems have been provided in Appendix B.

5 Physical, Media Layer, and Network Connectivity

Please consult the SFTI interface specification.

6 Appendix A - NMS IP Multicast Addresses

This appendix contains the mapping of IP multicast group ID's (addresses) to the currently available data lines. To receive a particular data stream the recipient host system would typically subscribe to that particular multicast group ID. Two multicast group ID's are available for each real-time production data line. The data originating from Group A is generally referred to as the 'A' streams and the data from Group B as the 'B' streams. Also provided in the table are the UDP destination ports associated with each logical line.

The NMS data messages are encapsulated in an identical manner in both streams. For example, a datagram issued Group A on OPRA Line 2 destined to multicast group 224.0.2.227 will have a corresponding datagram (containing the identical UDP data payload, i.e. same NMS messages and same sequence number range) sourced from Group B destined to multicast group 224.0.2.243.

Multicast Address Ranges:

NMS Production IP Multicast Feeds Group A:
224.0.2.192-224.0.2.207
224.0.2.224-224.0.2.239

NMS Production IP Multicast Feeds Group B:
224.0.2.208-224.0.2.223
224.0.2.240-224.0.2.255

The full table of address mappings is shown on the next page.

Production, Real-Time IP Multicast Feeds, Dual Sets

Group A Originated Data Lines	Multicast Group ID	Destination UDP Port Number	Group B Originated Data Lines	Multicast Group ID	Destination UDP Port Number
OPRA 1 (FCO)	224.0.2.192	53540	OPRA 1 (FCO)	224.0.2.208	53541
<i>SPARE</i>	224.0.2.193	53542	<i>SPARE</i>	224.0.2.209	53543
<i>SPARE</i>	224.0.2.194	53544	<i>SPARE</i>	224.0.2.210	53545
<i>SPARE</i>	224.0.2.195	53546	<i>SPARE</i>	224.0.2.211	53547
<i>SPARE</i>	224.0.2.196	53548	<i>SPARE</i>	224.0.2.212	53549
<i>SPARE</i>	224.0.2.197	53550	<i>SPARE</i>	224.0.2.213	53551
<i>SPARE</i>	224.0.2.198	53552	<i>SPARE</i>	224.0.2.214	53553
<i>SPARE</i>	224.0.2.199	53554	<i>SPARE</i>	224.0.2.215	53555
<i>SPARE</i>	224.0.2.200	53556	<i>SPARE</i>	224.0.2.216	53557
TIME BEACON	224.0.2.201	53558	TIME BEACON	224.0.2.217	53559
CQS 1	224.0.2.202	53560	CQS 1	224.0.2.218	53561
CQS 2	224.0.2.203	53562	CQS 2	224.0.2.219	53563
CQS 3	224.0.2.204	53564	CQS 3	224.0.2.220	53565
CQS 4	224.0.2.205	53566	CQS 4	224.0.2.221	53567
CTS 1	224.0.2.206	53568	CTS 1	224.0.2.222	53569
CTS 2	224.0.2.207	53570	CTS 2	224.0.2.223	53571
CTS 3	224.0.2.224	53572	CTS 3	224.0.2.240	53573
CTS 4	224.0.2.225	53574	CTS 4	224.0.2.241	53575
CQS 5	224.0.2.226	53576	CQS 5	224.0.2.242	53577
OPRA 2	224.0.2.227	53578	OPRA 2	224.0.2.243	53579
OPRA 3	224.0.2.228	53580	OPRA 3	224.0.2.244	53581
OPRA 4	224.0.2.229	53582	OPRA 4	224.0.2.245	53583
OPRA 5	224.0.2.230	53584	OPRA 5	224.0.2.246	53585
OPRA 6	224.0.2.231	53586	OPRA 6	224.0.2.247	53587
OPRA 7	224.0.2.232	53588	OPRA 7	224.0.2.248	53589
OPRA 8	224.0.2.233	53590	OPRA 8	224.0.2.249	53591
OPRA 9	224.0.2.234	53592	OPRA 9	224.0.2.250	53593
<i>SPARE</i>	224.0.2.235	53594	<i>SPARE</i>	224.0.2.251	53595
<i>SPARE</i>	224.0.2.236	53596	<i>SPARE</i>	224.0.2.252	53597
<i>SPARE</i>	224.0.2.237	53598	<i>SPARE</i>	224.0.2.253	53599
<i>SPARE</i>	224.0.2.238	53600	<i>SPARE</i>	224.0.2.254	53601
<i>SPARE</i>	224.0.2.239	53602	<i>SPARE</i>	224.0.2.255	53603

Retransmission and Playback Test Data, Single Sets

Unlike the production real-time feeds, the day-time production retransmission data and the after-hours playback test data are provided via a single stream only, i.e. redundant ‘A’ and ‘B’ streams are not available. Playback data is only available after-hours.

In addition to the playback test data groups as shown below, SIAC will continue to provide after-hours playback via the production system expressly for the purposes of redundancy testing. The multicast groups will be identical to those listed above for the real-time production system.

Recipients wishing to receive retransmission and/or playback feeds must subscribe to the multicast feeds based on the addressing information shown in the following table.

Multicast Address Ranges:

NMS Retransmission Multicast Group ID Ranges:

224.0.5.128 – 224.0.5.136

224.0.5.138 – 224.0.5.143

224.0.5.176 – 224.0.5.191

The full table of address mappings is shown on the next page.

Retransmission and Playback Test Data, Single Sets

Retransmission Group Assignments			Playback Test Group Assignments		
NMS Line Name available from both Group A and Group B	Multicast Group ID	Destination UDP Port Number	NMS Line Name available from both Group A and Group B	Multicast Group ID	Destination UDP Port Number
OPRA 1 (FCO)	224.0.5.128	54540	OPRA 1 (FCO)	224.0.5.144	55540
<i>SPARE</i>	224.0.5.129	54541	<i>SPARE</i>	224.0.5.145	55541
<i>SPARE</i>	224.0.5.130	54542	<i>SPARE</i>	224.0.5.146	55542
<i>SPARE</i>	224.0.5.131	54543	<i>SPARE</i>	224.0.5.147	55543
<i>SPARE</i>	224.0.5.132	54544	<i>SPARE</i>	224.0.5.148	55544
<i>SPARE</i>	224.0.5.133	54545	<i>SPARE</i>	224.0.5.149	55545
<i>SPARE</i>	224.0.5.134	54546	<i>SPARE</i>	224.0.5.150	55546
<i>SPARE</i>	224.0.5.135	54547	<i>SPARE</i>	224.0.5.151	55547
<i>SPARE</i>	224.0.5.136	54548	<i>SPARE</i>	224.0.5.152	55548
OPEN	224.0.5.137	54549	OPEN	224.0.5.153	55549
CQS IP line 1	224.0.5.138	54550	CQS IP line 1	224.0.5.154	55550
CQS IP line 2	224.0.5.139	54551	CQS IP line 2	224.0.5.155	55551
CQS IP line 3	224.0.5.140	54552	CQS IP line 3	224.0.5.156	55552
CQS IP line 4	224.0.5.141	54553	CQS IP line 4	224.0.5.157	55553
CTS IP line 1	224.0.5.142	54554	CTS IP line 1	224.0.5.158	55554
CTS IP line 2	224.0.5.143	54555	CTS IP line 2	224.0.5.159	55555
CTS IP line 3	224.0.5.176	54556	CTS IP line 3	224.0.5.240	55556
CTS IP line 4	224.0.5.177	54557	CTS IP line 4	224.0.5.241	55557
CQS IP line 5	224.0.5.178	54558	CQS IP line 5	224.0.5.242	55558
OPRA IP line 2	224.0.5.179	54559	OPRA IP line 2	224.0.5.243	55559
OPRA IP line 3	224.0.5.180	54560	OPRA IP line 3	224.0.5.244	55560
OPRA IP line 4	224.0.5.181	54561	OPRA IP line 4	224.0.5.245	55561
OPRA IP line 5	224.0.5.182	54562	OPRA IP line 5	224.0.5.246	55562
OPRA IP line 6	224.0.5.183	54563	OPRA IP line 6	224.0.5.247	55563
OPRA IP line 7	224.0.5.184	54564	OPRA IP line 7	224.0.5.248	55564
OPRA IP line 8	224.0.5.185	54565	OPRA IP line 8	224.0.5.249	55565
OPRA IP line 9	224.0.5.186	54566	OPRA IP line 9	224.0.5.250	55566
<i>SPARE</i>	224.0.5.187	54567	<i>SPARE</i>	224.0.5.251	55567
<i>SPARE</i>	224.0.5.188	54568	<i>SPARE</i>	224.0.5.252	55568
<i>SPARE</i>	224.0.5.189	54569	<i>SPARE</i>	224.0.5.253	55569
<i>SPARE</i>	224.0.5.190	54570	<i>SPARE</i>	224.0.5.254	55570
<i>SPARE</i>	224.0.5.191	54571	<i>SPARE</i>	224.0.5.255	55571

7 Appendix B - NMS IP Source Addresses

The following table lists all the possible source IP addresses associated with the each of the NMS Multicast Data services.

Service	Data Function Type	IP addresses listed as the source address in the Multicast packets (network number and mask in parenthesis)
CTS/CQS	Production	198.140.61.161 (198.140.61.160 / 29) 198.140.61.162 (198.140.61.160 / 29) 198.140.61.169 (198.140.61.168 / 29) 198.140.61.170 (198.140.61.168 / 29) 198.140.62.161 (198.140.62.160 / 29) 198.140.62.169 (198.140.62.168 / 29) 198.140.41.161 (198.140.41.160 / 29) 198.140.41.162 (198.140.41.160 / 29) 198.140.41.169 (198.140.41.168 / 29) 198.140.41.170 (198.140.41.168 / 29) 198.140.42.161 (198.140.42.160 / 29) 198.140.42.169 (198.140.42.168 / 29)
CTS/CQS	After-hours Playback Test	198.140.61.65 (198.140.61.64 / 29) 198.140.61.66 (198.140.61.64 / 29) 198.140.61.89 (198.140.61.88/ 29) 198.140.61.90 (198.140.61.88/ 29) 198.140.62.65 (198.140.62.64 / 29) 198.140.62.66 (198.140.62.64 / 29) 198.140.62.67 (198.140.62.64 / 29) 198.140.62.89 (198.140.62.88/ 29) 198.140.62.90 (198.140.62.88/ 29) 198.140.62.91 (198.140.62.88/ 29) 198.140.41.65 (198.140.41.64 / 29) 198.140.41.66 (198.140.41.64 / 29) 198.140.41.89 (198.140.41.88/ 29) 198.140.41.90 (198.140.41.88/ 29) 198.140.42.65 (198.140.42.64 / 29) 198.140.42.66 (198.140.42.64 / 29) 198.140.42.67 (198.140.42.64 / 29) 198.140.42.89 (198.140.42.88/ 29) 198.140.42.90 (198.140.42.88/ 29) 198.140.42.91 (198.140.42.88/ 29)
CTS/CQS	Production Retransmission	198.140.61.97 (198.140.61.96/ 29) 198.140.61.105 (198.140.61.104/ 29) 198.140.62.97 (198.140.62.96/ 29) 198.140.62.105 (198.140.62.104/ 29) 198.140.62.106 (198.140.62.104/ 29) 198.140.41.97 (198.140.41.96/ 29) 198.140.41.105 (198.140.41.104/ 29) 198.140.42.97 (198.140.42.96/ 29) 198.140.42.105 (198.140.42.104/ 29) 198.140.42.106 (198.140.42.104/ 29)

NMS IP Source Addresses, cont'd

Service	Data Function Type	IP addresses listed as the source address in the Multicast packets (network number and mask in parenthesis)
OPRA	Production	198.140.61.49 (198.140.61.48 / 29) 198.140.61.50 (198.140.61.48 / 29) 198.140.61.51 (198.140.61.48 / 29) 198.140.61.57 (198.140.61.56 / 29) 198.140.61.58 (198.140.61.56 / 29) 198.140.61.73 (198.140.61.72 / 29) 198.140.61.74 (198.140.61.72 / 29) 198.140.61.75 (198.140.61.72 / 29) 198.140.61.81 (198.140.61.80 / 29) 198.140.61.82 (198.140.61.80 / 29) 198.140.62.49 (198.140.62.48 / 29) 198.140.62.50 (198.140.62.48 / 29) 198.140.62.51 (198.140.62.48 / 29) 198.140.62.57 (198.140.62.56 / 29) 198.140.62.58 (198.140.62.56 / 29) 198.140.62.73 (198.140.62.72 / 29) 198.140.62.74 (198.140.62.72 / 29) 198.140.62.75 (198.140.62.72 / 29) 198.140.62.81 (198.140.62.80 / 29) 198.140.62.82 (198.140.62.80 / 29) 198.140.41.49 (198.140.41.48 / 29) 198.140.41.50 (198.140.41.48 / 29) 198.140.41.51 (198.140.41.48 / 29) 198.140.41.57 (198.140.41.56 / 29) 198.140.41.58 (198.140.41.56 / 29) 198.140.41.73 (198.140.41.72 / 29) 198.140.41.74 (198.140.41.72 / 29) 198.140.41.75 (198.140.41.72 / 29) 198.140.41.81 (198.140.41.80 / 29) 198.140.41.82 (198.140.41.80 / 29) 198.140.42.49 (198.140.42.48 / 29) 198.140.42.50 (198.140.42.48 / 29) 198.140.42.51 (198.140.42.48 / 29) 198.140.42.57 (198.140.42.56 / 29) 198.140.42.58 (198.140.42.56 / 29) 198.140.42.73 (198.140.42.72 / 29) 198.140.42.74 (198.140.42.72 / 29) 198.140.42.75 (198.140.42.72 / 29) 198.140.42.81 (198.140.42.80 / 29) 198.140.42.82 (198.140.42.80 / 29)

NMS IP Source Addresses, cont'd

Service	Data Function Type	IP addresses listed as the source address in the Multicast packets (network number and mask in parenthesis)
OPRA	After-hours Playback Test	198.140.61.65 (198.140.61.64 / 29) 198.140.61.66 (198.140.61.64 / 29) 198.140.61.89 (198.140.61.88/ 29) 198.140.61.90 (198.140.61.88/ 29) 198.140.62.65 (198.140.62.64 / 29) 198.140.62.66 (198.140.62.64 / 29) 198.140.62.67 (198.140.62.64 / 29) 198.140.62.89 (198.140.62.88/ 29) 198.140.62.90 (198.140.62.88/ 29) 198.140.62.91 (198.140.62.88/ 29) 198.140.41.65 (198.140.41.64 / 29) 198.140.41.66 (198.140.41.64 / 29) 198.140.41.89 (198.140.41.88/ 29) 198.140.41.90 (198.140.41.88/ 29) 198.140.42.65 (198.140.42.64 / 29) 198.140.42.66 (198.140.42.64 / 29) 198.140.42.67 (198.140.42.64 / 29) 198.140.42.89 (198.140.42.88/ 29) 198.140.42.90 (198.140.42.88/ 29) 198.140.42.91 (198.140.42.88/ 29)
OPRA	Production Retransmission	198.140.61.97 (198.140.61.96/ 29) 198.140.61.105 (198.140.61.104/ 29) 198.140.62.97 (198.140.62.96/ 29) 198.140.62.105 (198.140.62.104/ 29) 198.140.62.106 (198.140.62.104/ 29) 198.140.41.97 (198.140.41.96/ 29) 198.140.41.105 (198.140.41.104/ 29) 198.140.42.97 (198.140.42.96/ 29) 198.140.42.105 (198.140.42.104/ 29) 198.140.42.106 (198.140.42.104/ 29)

NMS IP Source Addresses, cont'd

Time Beacon	198.140.61.49	(198.140.61.48 / 29)
	198.140.61.50	(198.140.61.48 / 29)
	198.140.61.51	(198.140.61.48 / 29)
	198.140.61.57	(198.140.61.56 / 29)
	198.140.61.58	(198.140.61.56 / 29)
	198.140.61.73	(198.140.61.72 / 29)
	198.140.61.74	(198.140.61.72 / 29)
	198.140.61.75	(198.140.61.72 / 29)
	198.140.61.81	(198.140.61.80 / 29)
	198.140.61.82	(198.140.61.80 / 29)
	198.140.61.161	(198.140.61.160 / 29)
	198.140.61.162	(198.140.61.160 / 29)
	198.140.61.169	(198.140.61.168 / 29)
	198.140.61.170	(198.140.61.168 / 29)
	198.140.62.49	(198.140.62.48 / 29)
	198.140.62.50	(198.140.62.48 / 29)
	198.140.62.51	(198.140.62.48 / 29)
	198.140.62.57	(198.140.62.56 / 29)
	198.140.62.58	(198.140.62.56 / 29)
	198.140.62.73	(198.140.62.72 / 29)
	198.140.62.74	(198.140.62.72 / 29)
	198.140.62.75	(198.140.62.72 / 29)
	198.140.62.81	(198.140.62.80 / 29)
	198.140.62.82	(198.140.62.80 / 29)
	198.140.62.161	(198.140.62.160 / 29)
	198.140.62.169	(198.140.62.168 / 29)
	198.140.41.49	(198.140.41.48 / 29)
	198.140.41.50	(198.140.41.48 / 29)
	198.140.41.51	(198.140.41.48 / 29)
	198.140.41.57	(198.140.41.56 / 29)
	198.140.41.58	(198.140.41.56 / 29)
	198.140.41.73	(198.140.41.72 / 29)
	198.140.41.74	(198.140.41.72 / 29)
	198.140.41.75	(198.140.41.72 / 29)
	198.140.41.81	(198.140.41.80 / 29)
	198.140.41.161	(198.140.41.160 / 29)
	198.140.41.162	(198.140.41.160 / 29)
	198.140.41.169	(198.140.41.168 / 29)
	198.140.41.170	(198.140.41.168 / 29)
	198.140.41.82	(198.140.41.80 / 29)
	198.140.42.49	(198.140.42.48 / 29)
	198.140.42.50	(198.140.42.48 / 29)
	198.140.42.51	(198.140.42.48 / 29)
	198.140.42.57	(198.140.42.56 / 29)
	198.140.42.58	(198.140.42.56 / 29)
	198.140.42.73	(198.140.42.72 / 29)
	198.140.42.74	(198.140.42.72 / 29)
	198.140.42.75	(198.140.42.72 / 29)
	198.140.42.81	(198.140.42.80 / 29)
	198.140.42.82	(198.140.42.80 / 29)
198.140.42.161	(198.140.42.160 / 29)	
198.140.42.169	(198.140.42.168 / 29)	

8 Appendix C - Time Beacon Message Format

The Time Beacon message is delivered as the data portion of a UDP/IP packet.

Each packet will contain a single message.

Each message is 14 bytes in length and consists of two fields, the MPR Identifier and the Time Stamp and is formatted as shown below.

MPR Identifier	Time Stamp
<p style="text-align: center;">NN [2 Bytes]</p>	<p style="text-align: center;">MMDDYYHHMMSS [12 Bytes]</p>

Description of each field:

Multicast Packet Replicator (MPR) Identifier:

2 Bytes, Numeric - Identifies the MPR that initiated the message.

As of 10/1/00 this number ranges from 1-10 but can change without notice.

Time Stamp:

12 Bytes, Alphanumeric/Special Character - Format is MMDDYYHHMMSS where,

M=Month

D=Day

Y=Year

H=Hours (specified as '0' through '23', in the same manner as military time)

M=Minutes

S=Seconds

The time stamp will reflect the current local time of the United States' Eastern Time zone.

The time stamp will reflect daylight savings time when in effect.