

TCP/IP for NMS Participant Input

Participant Interface Specification

Prepared by:



**Communications Engineering
Planning and Development**

Date: 12/27/19
Revision: 4.3

Copyright Notice

Copyright © 2019 by 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 an initial capital. Regardless of the capitalization used, all such use has been in an 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:****Title:****TCP/IP for NMS Participant Input, Participant Interface Specification**

Version	Date	Rev by	Pages affected	Comments	Approval
0.2	7/18/97			Initial Draft	
0.3	7/21/97	ML	2,3	Material added	
0.4	7/23/97	ML	All	New sections	
1.0	8/7/97	ML	All	First release	
1.1	6/15/98	ML	All	Issues resolved	
1.2	6/20/98	ML	All	Additional chugs	
1.3	7/7/98	ML	All	Additional chugs	
1.4	12/31/98	RL	All	Change to be an official specification for distribution to Participants	
1.5	2/2/99	RL		Text added to section 6.2.1 regarding 4 byte header	
1.6	2/8/99	RL		- Expanded detail to section 6.2.1 regarding 4 byte header - TCP/IP connections chart added to Appendix A	
1.9	7/11/02	CE, RL	All	Removed references to legacy connections. Treatment of TCP/IP as the only option and not as the “new” option. Remove reference to data center locations. Remove references to non-Ethernet connectivity options	
1.91	7/22/02	RL	Appendix A	minor change to table	
2.05	10/7/03	RL, CE	All	Added ITS and compliance with SFTI	
2.06	10/16/03	RL	Appendix A	Table updated for ITS	
2.07	12/26/03	LG	Page 8	Addition of Pad character for ITS in message block	
3.0	10/3/07	ML, PH	All	- Removed references to ITS - Changed references to NMS mainframe to NMS system	
4.0	5/22/12	MCP	Various	-Performed document cleanup -Revised document where applicable	
4.1	11/22/13	MC	Various	Performed further document cleanup	
4.2	5/9/16	WG	Various	Updated SFTI contact information	
4.3	12/27/19	WG	Various	Modified all SFTI references to IGN Updated contact information	

Table of Contents

1.	Introduction	1
1.1	Requesting ICE Global Network (IGN) Services	1
2.	Overview TCP/IP.....	1
3.	Existing NMS Multicast Environment.....	1
4.	Application Overview.....	2
4.1	Architecture.....	2
4.2	Logical Connection Allocation and Fairness	4
5.	NMS Input Network	5
5.1	TCP/IP Support.....	5
5.2	Redundancy	5
5.2.1	Redundancy Example	5
5.3	Security	5
5.4	Data Pacing.....	7
6.	Participant Requirements.....	7
6.1	Applications	7
6.1.1	General Message Format	7
6.1.2	Safe-Store Exposure	7
7.	Network Layer Connectivity	8
7.1	IP.....	8
7.2	IP Addressing.....	8
7.3	TCP/IP Segment.....	8
7.3.1	IP Header Field.....	9
7.3.2	TCP Header Field Description.....	9
7.4	Dual Access Network.....	10
7.5	IP Network Considerations	10
8.	Physical, Media Layer, and Network Connectivity via IGN.....	11
9.	Appendix A: References.....	12

1. Introduction

SIAC has deployed the NMS network to support the dissemination of the NMS market data to the Recipient community. Since the Participants are also Recipients with a connection to this network, it was determined that this network would also support the input from the Participants. Participants connect to the NMS network via the ICE Global Network (IGN) (formerly Secure Financial Transaction Infrastructure; SFTI – pronounced ‘safety’). With respect to connecting directly to IGN, customers are urged to review the IGN Interface Specification for Directly Connected Customers for more information, call US (770) 661-0010 option 1 | EU +44 203 808 6638 option 1, or visit the [Data Services Dashboard](#) portal.

1.1 Requesting ICE Global Network (IGN) Services

In order to more efficiently and accurately handle your IGN service requests, IGN has added functionality to the Data Services Dashboard portal that will allow you to submit your service requests online. This secure portal can be used for requesting service changes to your IGN connection(s), which do not require a contract modification such as provisioning for new core services, new subnets, network maintenance and disconnects.

To submit your request, please log into the [Data Services Dashboard](#) portal, fill out all appropriate information. Should you need an account set-up for access to the Data Services Dashboard portal you may create one by clicking on the “Register” link on the dashboard login page. Alternatively, you may contact an IGN sales representative (clientnetworks@theice.com) and they will facilitate the request for you.

2. Overview TCP/IP

TCP/IP is a network layer connection oriented protocol that will provide a reliable network based point to point connection between the Participants input system and NMS. Reliability in transmission is guaranteed using a handshaking mechanism in which the receiving system will send an IP acknowledgment packet back to the sending system upon reception of one or more packets. When a packet is received with errors, the acknowledgment packet can also request a retransmission from the source. Conversely, if the source system does not receive an acknowledgment for packets sent, it will assume they were never received by the destination system and will re-transmit them.

The TCP handshake mechanism also performs implicit data pacing between source and destination guaranteeing that the source system will not overload the input of the receiving system. This occurs because the acknowledgment messages have the ability to signal the source of a busy/clear condition, which effectively paces the packets across the network. The guaranteed delivery and pacing occur transparently to the applications running on both the sending and receiving systems.

3. Existing NMS Multicast Environment

Figure 1 (on next page) provides a high level view of the NMS IP environment. For NMS market data dissemination the NMS hosts (OPRA, CTS and CQS) output their messages encapsulated or wrapped in a series of IP/UDP packets. These packets are forwarded into the IP distribution network with an IP multicast destination address. The multicast address selected per packet is based on the unique data line and system that the message data supports.

For a recipient system to receive a particular data line, it would send a multicast subscription message containing the mapped multicast address to their recipient router, which is connected to SIAC’s distribution network. On reception of the subscription message, the network will replicate the requested IP multicast data stream and forward it on to the recipient’s network. The Figure also illustrates how the participating market centers can connect via TCP/IP to the NMS system.

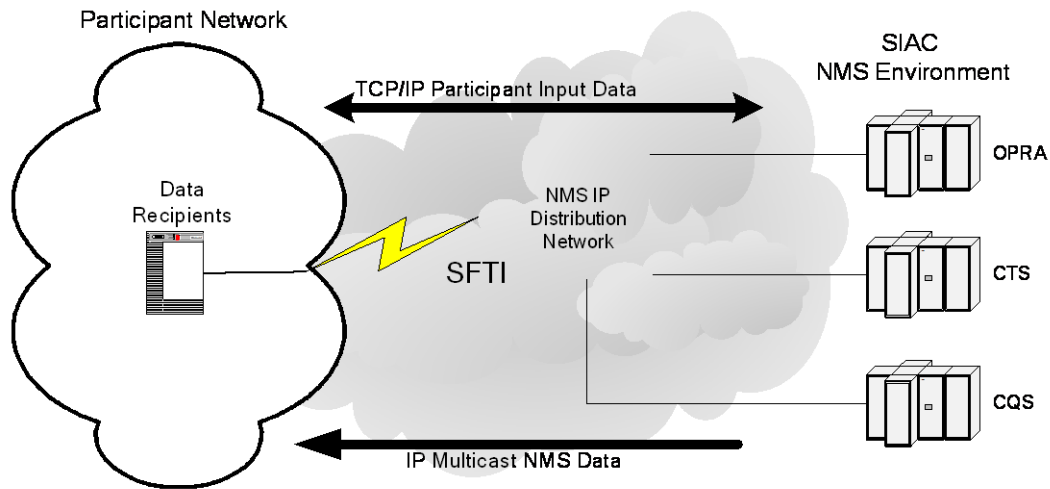


Figure 1

Overview of connecting to NMS hosts via IGN (formerly SFTI) and the NMS network

4. Application Overview

4.1 Architecture


For the applications to support the TCP/IP protocol the following must be considered:

- The message formats for CTS/CQS/OPRA can be found in the Participant Communications Interface specifications.
- The number of logical connections to each system by each Participant will be based on a Participant's projected maximum input rate (in messages/second) and maximum data rate that can be supported on a logical connection.
- The Participant will establish the TCP/IP connections. SIAC will always be the listener.
- A set of IP addresses will be assigned for each Participant's access to each system (CTS, CQS and OPRA), with one or more TCP/IP ports associated with each of these IP addresses.


The Participants' application systems will communicate with SIAC's NMS (CTS, CQS and OPRA) systems via TCP/IP socket connections. Sockets are established by the Participants' application system connecting to a pre-determined unique application port number in the NMS system. Socket addresses (IP address and port number) for all logical connections will be pre-assigned by SIAC for each Participant.

Each Participant is given access to multiple physical interfaces that are shared among multiple Participants. Participants are assigned a unique TCP port number to which only their applications will be permitted to connect. The TCP/IP socket connections are established by the Participants' application systems.

Participant A Systems

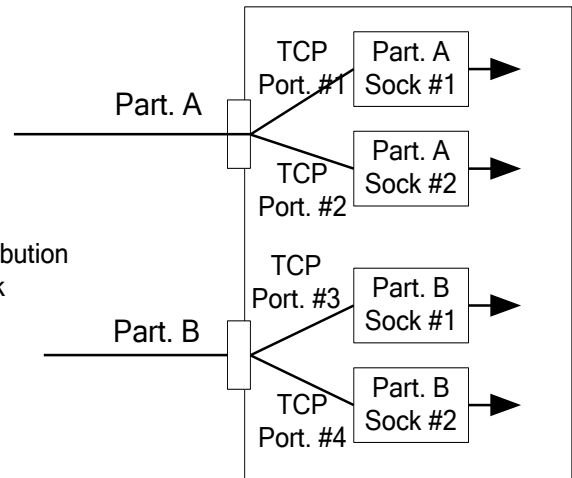


Participant B Systems



NMS IP Distribution Network

SIAC NMS Systems



CTS/CQS/OPRA

Figure 2
CTS/CQS/OPRA TCP/IP Architecture

4.2 Logical Connection Allocation and Fairness

For each host system, SIAC will be defining a Maximum Message Rate (MMR) allowed for a given logical connection (TCP/IP socket connection). This will be implemented in the form of an application throttle, which will keep the message rate at or below a preset rate. The MMR will be based on actual testing and known CPU capacity restrictions. It will also be system dependent (i.e., could be different for each NMS system). The application throttles will be set for each logical connection so that the total aggregate rate (across all connections) meets the Participants projected maximum message rate while not exceeding the systems capacity.

Fairness between Participants is defined as each Participant always being able to input their projected maximum message rate. The NMS systems are sized to support the projected message rate from all Participants concurrently.

5. NMS Input Network

5.1 TCP/IP Support

The NMS network is used to support TCP/IP input connectivity to the NMS hosts in addition to multicast distribution as shown in Figure 1. The network, regardless of whether the packets are IP multicast framed or TCP/IP framed, properly routes packets across the network. In the case of multicast packets, the packets are forwarded via a single path to all networks that have subscribed to the multicast group listed in the IP header of each multicast packet. TCP/IP packets follow standard unicast routing algorithms as dictated by the destination IP address of each TCP/IP packet.

When routing TCP/IP packets, the Participant's systems would encapsulate their application data within a TCP/IP packet and provide a destination IP address and destination TCP port number in order to direct the packet to a particular NMS host address and application port.

5.2 Redundancy

Though not shown in Figures 1 and 2, all Data Recipients and Participants will connect to the NMS network and its hosts via the ICE Global Network (IGN) (formerly SFTI). There is a separate network interface specification available that describes in more detail how to connect to the edge routers of IGN that are currently available for direct connectivity via several access centers in the United States. By connecting via the IGN Access Centers, customers are provided routed access into the NMS network's two data centers. The packets are routed based on destination address.

Participant application hosts will be responsible for initiating the connections to the NMS host systems.

With regard to the NMS hosts, multiple network-based connections on each system will be maintained to support redundant connectivity from the Participants in the event of a connection failure.

5.2.1 Redundancy Example

To illustrate the resiliency of the TCP/IP input environment, Figure 3 (on next page) shows a single Participant input system redundantly connected into the NMS networking environment. It can be assumed in this example that the customer has connectivity between its network and the IGN Access Centers. For more information, call US (770) 661-0010 option 1 | EU +44 203 808 6638 option 1, or visit the [Data Services Dashboard](#) portal.

5.3 Security

Network security within the NMS network is provided with the use of router filters and IP Tables installed in the NMS and IGN routers at appropriate points in the network path. These filters create an access policy of sorts between each Participant and their logical connections to the NMS hosts.

SIAC is protecting its network and hosts using packet level filtering on the IGN and NMS routers. Routing policies are implemented on SIAC managed routers to insure strict control over the IP routing table. These security measures maintain the integrity of network routing tables and also protect SIAC's network and hosts from intentional or accidental access by a Participant network. These measures are in no way intended to provide the same level of security to the Participants themselves. If a Participant believes that additional security is required to protect their network then it is left to them to implement it.

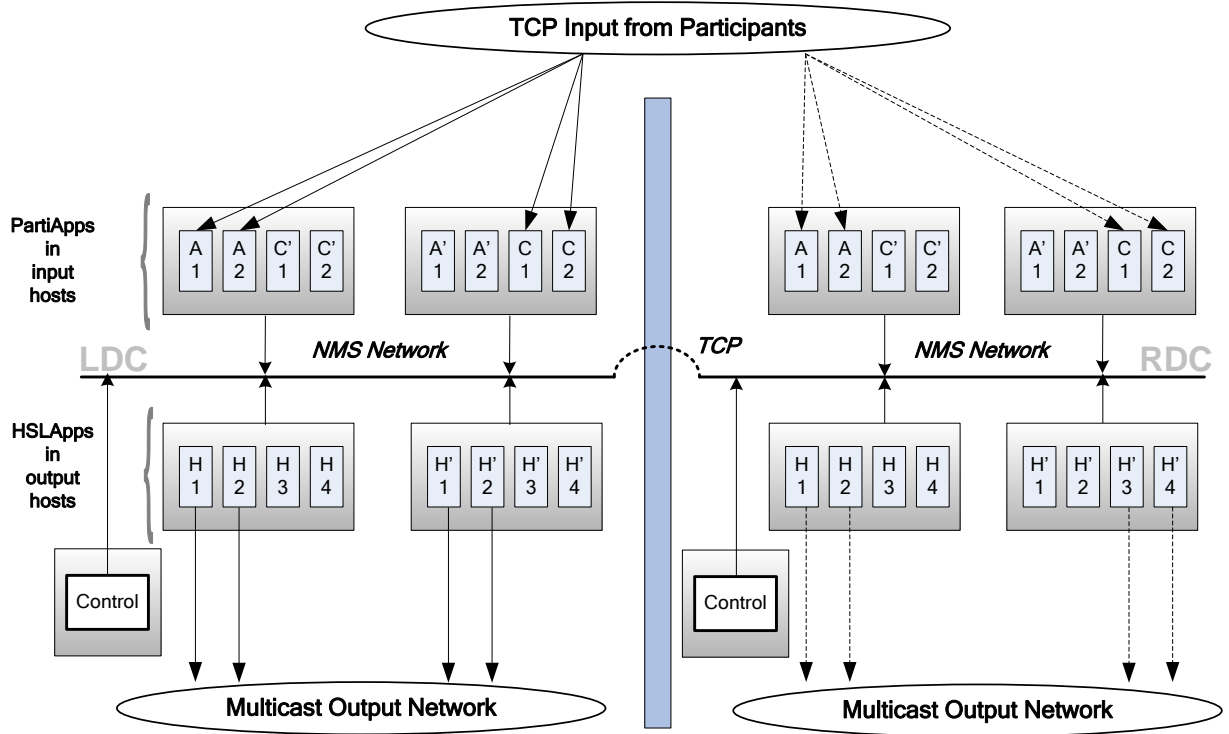


Figure 3

5.4 Data Pacing

The OPRA, CTS and CQS applications are set to read data received from Participants (via the TCP/IP stack) at a preset maximum rate to assure that each Participant is given appropriate and fair access to NMS system resources.

The TCP/IP protocol itself performs implicit data pacing between the Participants input system and the NMS hosts via the use of the TCP/IP window size advertisement and back-off timers. TCP/IP's window size pacing mechanism becomes active if a Participant is sending messages at a faster rate than the NMS application can process them and thus causing the NMS host's TCP/IP input buffer to fill up (the same is true in the other direction, i.e., NMS host to Participant host). The window size advertisement lets each end of the connection know how many bytes of data can be sent without losing data. Depending on a combination of factors including network WAN speeds, system capacities and network congestion, the TCP/IP window size adjustments can result in data pacing.

The protocol also uses packet acknowledgements for identifying when retransmissions are required. However if retransmissions occur frequently, the protocol's algorithms will automatically "back-off" from sending data at the existing rate, i.e., moving to a slower rate, until successful TCP packet acknowledgements resume and the back-off timers have readjusted to normal processing.

6. Participant Requirements

6.1 Applications

6.1.1 General Message Format

CTS/CQS/OPRA message formats are found in the respective CTS, CQS, and OPRA Participant Communications Specifications.

6.1.2 Safe-Store Exposure

When using TCP/IP, the application system has no control of the TCP/IP ACK messages. If the receiving system, experiences a problem processing input data, the messages that have been acknowledged by the TCP/IP stack but have not been passed to the application system would be lost and the Participant's system would not have this information.

The number of messages that could be potentially lost is determined by the buffer space allocated for the specific TCP/IP socket on the SIAC system. To deal with this unlikely event, the NMS input applications already perform sequence number checking which will notify the Participant when a sequence gap is detected. This is an automated process.

7. Network Layer Connectivity

7.1 IP

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. Addressing between end stations is based on a source and destination IP address. Routing packets (data frames) between networks is a point-to-point operation; involving one source, one destination, and one or more intermediate routers.

7.2 IP Addressing

End stations participating in the NMS input network must use unicast IP addresses which include 1.0.0.0 through 223.255.255.255 (using standard IP address notation). When an IP router encounters a packet with a destination unicast address, the packet is routed using the standard IP routing table when determining the next hop interface or destination network to send the packet.

7.3 TCP/IP Segment

The application data is encapsulated in a TCP/IP frame as shown in Figure 7.3-1. The IP datagram includes the IP and TCP headers plus the application data block as described in section 6 above. 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 TCP 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 7.3-1 TCP/IP Datagram Format

7.3.1 IP Header Field

- **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 TCP header plus the application data (TCP data).
- **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.
- **Flags and Fragment Offset** - The combined 16 bit field is only used when an IP datagram is fragmented.
- **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.
- **Protocol** - This 8 bit field contains a value representing the next level encapsulated protocol. In this case it is TCP, which has a value of 6.
- **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 source system.
- **IP Destination Address** - This 32 bit field contains the IP address of the destination system.

7.3.2 TCP Header Field Description

SIAC will be expecting Participant applications to connect to pre-defined IP destination addresses and IP destination port numbers. The destination IP addresses will correspond to physical NMS system ports. The destination port numbers will be assigned to each Participant by SIAC. For detail regarding the TCP protocol and header information, the reader should refer to RFC 793.

- **TCP Source Port Number** - This 16 bit field identifies the sending process within the Participant source system. It is set by the source system.
- **TCP Destination Port Number** - This 16 bit field identifies the TCP process waiting for a connection request on the SIAC NMS system. Each Participant will be given a pre-defined set of TCP destination port numbers specific to that Participant. Participant systems must use these TCP destination port numbers when opening connections to the NMS system. Connections not matching the pre-defined values will be blocked by SIAC.

7.4 Dual Access Network

In order to provide a resilient/redundant networking environment the Participant is provided with the ability to connect to SIAC's data centers via multiple IGN Access Centers.

7.5 IP Network Considerations

The IP networks making up the connection between the Participant's location and the IGN edge router ports located at IGN Access Centers are the responsibility of the Participant and/or their third party access providers. With respect to the connection to the IGN edge port, the Participant is required to provide a publicly registered IP address for both ends of this connection, which includes the router port on the IGN edge router. If the Participant doesn't have a publicly registered address to use, SIAC will provide a range of private addresses to use. This and related IP address requirements are explained in more detail in the IGN Interface Specification for Directly Connected Customers.

8. Physical, Media Layer, and Network Connectivity via IGN

Participants must connect to the ICE Global Network (IGN) via one of the IGN Access Centers in order to access NMS systems. At each of the Access Centers will be IGN edge routers that provide customers with an Ethernet port to which they can connect.

The device connected to the IGN edge router must be a router or Layer 2 switch capable of tagging packets with VLAN membership tags as per the 802.1Q standard. Participants must configure their switches to perform VLAN tagging such that the IGN edge routers can recognize the tag and forward the incoming packets appropriately. *Reference Section 1.*

For additional information related to connectivity directly to IGN, customers are urged to review the IGN Interface Specification for Directly Connected Customers for more information, call US (770) 661-0010 option 1 | EU +44 203 808 6638 option 1, or visit the [Data Services Dashboard](#) portal.

9. Appendix A: References

IP related standards documentation:

RFC 768	User Datagram Protocol (UDP)
RFC 791	Internet Protocol (IP)
RFC 792	Internet Control Message Protocol (ICMP)
RFC 793	Transmission Control Protocol (TCP)
RFC 826	Address Resolution Protocol (ARP)
RFC 903	Reverse Address Resolution Protocol (RARP)
RFC 950	Internet Subnetting Procedures
RFC 1042	IP over IEEE 802 Networks
RFC 1075	Distance Vector Multicast Routing Protocol (DVMRP)
RFC 1112	Host Extensions for IP Multicasting (includes the Internet Group Management Protocol (IGMP))
RFC 1122	Requirements for Internet hosts - communication layers
RFC 1123	Requirements for Internet hosts - application and support