



FINANCIAL AUTOMATION

COMMUNICATIONS

MEMBER FIRM INTERFACE

TCP/IP COMMUNICATION MANUAL

FINANCIAL AUTOMATION COMMUNICATIONS	Section: Page: 1 Release: 1.0 Date: Sep 25, 2003 Rev. Date: Aug 6, 2012
MEMBER FIRM INTERFACE	TCP/IP COMMUNICATIONS MANUAL

FINANCIAL AUTOMATION COMMUNICATIONS MEMBER FIRM INTERFACE	1
1 EXECUTIVE SUMMARY	2
2 INTRODUCTION	4
2.1 OVERVIEW.....	4
3 MFI CHARACTERISTICS AND CONNECTIVITY OPTIONS.....	5
3.1 STANDARDS CONFORMANCE.....	5
3.2 PROPRIETARY FEATURE DISCLAIMER.....	5
3.3 CONNECTIVITY METHODS	6
3.3.1 <i>Method 1 – Client WAN Extension</i>	6
3.3.2 <i>Method 2 – Client LAN Extension (Optical Ethernet)</i>	9
3.3.3 <i>Method 3 - Extranet</i>	11
4 BACKGROUND AND SERVICE BASICS	13
4.1 CONNECTIVITY MIGRATION FOR EXISTING CLIENTS	13
4.1.1 <i>Group Contact Information</i>	13
5 INTERNET PROTOCOL OVERVIEW	13
6 RELIABLE STREAM TRANSPORT SERVICE (TCP).....	15
6.1 RELIABILITY.....	16
7 TECHNICAL INTERFACE SPECIFICATIONS.....	18
7.1 NFX STANDARDS	19
7.2 CLIENT REQUIREMENTS FOR CONNECTIVITY.....	19
7.3 CABLING SPECIFICATIONS	21
8 IP REQUIREMENTS	23
8.1 NETWORK ADDRESS TRANSLATION OVERVIEW.....	24
8.2 NFX IP TRAFFIC AND ROUTING POLICY	24
9 NETWORK SUPPORT	25

FINANCIAL AUTOMATION COMMUNICATIONS	Section: Page: 2 Release: 1.0 Date: Sep 25, 2003 Rev. Date: Aug 6, 2012
MEMBER FIRM INTERFACE	TCP/IP COMMUNICATIONS MANUAL

1 Executive Summary

The NASDAQ OMX Futures Exchange, hereafter referred to as NFX, continues its long-standing commitment to quality by continually evaluating access methods, systems and application architectures that only offer the highest degree of resiliency and availability. The cornerstone of this universal belief is based on the adoption of technology that has been field proven to provide carrier grade reliability by providing 99.999% uptime. The Exchange's communications architecture that provides external connectivity for all market data recipients, participants and member firms hinges on industry "best practice" design methodology and product selection criteria. The NFX Member Firm Interface, hereafter referred to as the NFX MFI, is a redundant fault-tolerant hierarchical network architecture based on "best of breed" networking and network security platforms. Each networking device in the MFI architecture meets the NFX concept of "hardness". Hardened platforms sport load-balancing power supplies in an N+1 configuration, redundant load-sharing CPUs, routing engines and switching fabrics.

The Exchange has witnessed many trends in computing, communications and networking in its long history of serving the financial community. To keep pace with the paradigm shift of application deployments using the TCP/IP protocol suite and to leapfrog their emerging requirements, the NFX mandates all network connectivity to the MFI be based on the 4th version of the Internet Protocol, known as IPv4. In addition to migrating the Exchange's legacy protocols to IPv4, this requirement empowers the NFX to improve efficiencies by: 1) providing richer services and content, 2) simplifying overall connectivity requirements and 3) affording greater maintainability and operational ease.

The balance of this document is geared towards the familiarization of connection methods and requirements for establishing secure and reliable connectivity to the NFX MFI. This specification is intended for clients of, and contributors to the NFX MFI.

- Clients are defined as, but not limited to: member firms and market data recipients.
- Contributors are defined as, but not limited to: market data vendors, financial information suppliers, foreign and regional stock exchanges, ECNs, clearinghouses, providers of post-trade services, market data recipients, extranet providers, service bureaus, and contracted service providers.

Throughout the remainder of this document, the term **Client** will be used to refer to both Clients and Contributors.

Connectivity to the NFX MFI can be established through a variety of methods. The first model supports direct client connectivity to the Exchange. This is typically accomplished by clients extending their private Wide Area Network (WAN) to the Exchange or utilizing a Metropolitan (Metro) or Long Haul (LH) Ethernet based service provided by an Ethernet Service Provider. The second model supports client connectivity through a third party service provider. Examples of this include: managed service providers, extranet service providers and IP aggregators. Please consult with the NFX Trading and Market Operations at +1 215 496 5473 for a list of NFX's certified authorized service providers in this space.

Although numerous access methods are at the client's disposal, NFX requires a redundant pair of Ethernet handoffs be provided at separate secured MFI Demarcation (DMARC) points. At this

FINANCIAL AUTOMATION COMMUNICATIONS	Section: Page: 3 Release: 1.0 Date: Sep 25, 2003 Rev. Date: Aug 6, 2012
MEMBER FIRM INTERFACE	TCP/IP COMMUNICATIONS MANUAL

time, only copper Ethernet handoffs will be accepted. The NFX MFI strictly adheres to the IEEE 802.3 10BASE-T and 802.3u 100BASE-TX specifications.

The sole intention of this document is to provide the technical specifications for connectivity to the NFX MFI. It is not intended to fully document the services available, (beyond a listing of available application services). The Exchange reserves the right to change or modify this specification to improve service, reliability or support new or modified application requirements.

FINANCIAL AUTOMATION COMMUNICATIONS	Section: Page: 4 Release: 1.0 Date: Sep 25, 2003 Rev. Date: Aug 6, 2012
MEMBER FIRM INTERFACE	TCP/IP COMMUNICATIONS MANUAL

2 Introduction

2.1 OVERVIEW

The NFX MFI network architecture is based on a universal connectivity model that allows clients to access all front and back office NFX IP based applications through a single common redundant infrastructure. Each redundant connection to the MFI allows clients to subscribe to multiple services over a single common network interface. The Internet Protocol (IP) serves as the Network Layer vehicle that allows the upper layer Transport Control Protocol (TCP) to virtually multiplex multiple logical streams of traffic onto a single common data path. NFX has mandated that IP packets at the MFI edge are presented to the client in encapsulated Ethernet Version II Frames.

The MFI architecture allows clients to consolidate multiple dissimilar connections to the NFX over a single pair of redundant connections. Client connections to the MFI can gain access to any or all of the following services:

Equity Options	Equity	Business Systems
AUTOM – Order Interface Specialized Quote Feed (SQF) Specialized Order Feed (SOF) Risk Management System	PACE – Order Interface Specialized Quote Feed (SQF) Specialized Order Feed (SOF) Remote Competing Specialist Risk Management System	Mainframe File Transfers General File Transfers

FINANCIAL AUTOMATION COMMUNICATIONS	Section: Page: 5 Release: 1.0 Date: Sep 25, 2003 Rev. Date: Aug 6, 2012
MEMBER FIRM INTERFACE	TCP/IP COMMUNICATIONS MANUAL

3 MFI Characteristics and Connectivity Options

3.1 STANDARDS CONFORMANCE

The MFI provides a stable, secure, redundant high capacity infrastructure that serves as a single ingress/egress access point for all client based data communications. Like many institutions in the financial sector, NFX has adopted the IPv4 protocol suite to standardize all network layer conversations between host systems across divergent heterogeneous autonomous systems. MFI adherence to the Internet Engineering Task Force's (IETF) Internet Protocol Standards (STD) and Request for Comments (RFC) document series ensures the highest degree of mature stable interoperability between the Exchange's network edge and the client's Customer Premise Equipment (CPE). It is NFX's primary mission to provide its client base with a flexible architecture that provides the following benefits:

- Rapid service fulfillment of accessing all of NFX's application systems.
- "One Stop Shopping" for all financial applications over a common infrastructure.
- Adherence to industry recognized standards to ensure interoperability and scalability.
- Flexibility in connectivity methods that allow choice for client service provisioning to the Exchange.
- Simplification of entitlement to additional NFX services.

Although NFX mandates all connectivity adhere to the RFC specifications provided in Appendix A, vendor specific (proprietary) features are not outside the scope of MFI support. Clients requiring support for vendor specific features should contact the NFX Network Engineering group at the following email address: NOCGROUP@nasdaqomx.com

3.2 PROPRIETARY FEATURE DISCLAIMER

The NASDAQ OMX Futures Exchange Financial Automation Communications Group reserves the right to test vendor specific proprietary features prior to certifying their approval for use in the production MFI network. Lead times for this process will vary based on the level of complexity and resource required for certification. All requests of this nature will be handled on a case-by-case basis. Acceptance testing of proprietary feature sets is not a guarantee of production network certification approval.

3.3 CONNECTIVITY METHODS

For NASDAQ OMX | NFX Use Only
Proprietary and Confidential – 2012

FINANCIAL AUTOMATION COMMUNICATIONS	Section: Page: 6 Release: 1.0 Date: Sep 25, 2003 Rev. Date: Aug 6, 2012
MEMBER FIRM INTERFACE	TCP/IP COMMUNICATIONS MANUAL

Connectivity to the NFX MFI can be established through a variety of methods. The first model supports direct client connectivity to the Exchange. This is typically accomplished by clients extending their private Wide Area Network (WAN) to the Exchange or utilizing a Metropolitan (Metro) or Long Haul (LH) Ethernet based service provided by an Ethernet Service Provider. The second model supports client connectivity through a third party service provider. Examples of this include: managed service providers, extranet service providers and IP aggregators.

Although numerous access methods are at the client's disposal, NFX requires a redundant pair of Ethernet handoffs be provided at separate secured MFI Demarcation (DMARC) points. At this time, only copper Ethernet handoffs will be accepted. The NFX MFI strictly adheres to the IEEE 802.3 10BASE-T and 802.3u 100BASE-TX specifications.

3.3.1 Method 1 – Client WAN Extension

Clients possessing their own private enterprise WAN may extend this service directly to the Exchange.

1. The client is fully responsible for converting the WAN based circuit into an Ethernet interface through an appropriate CPE device. The CPE device is typically a router. Clients may procure co-location space at the NFX and must coordinate this request through the [NFX Operations Desk](#) at +1 215 496 1471.
2. The client is fully responsible for managing the carrier or service provider for all aspects of the leased-line circuit installation, provisioning and commissioning cycles. NFX will act as the coordinator for change control windows to allow carrier and client personnel on-site for circuit and CPE installation and demarcation.
3. Please note, only full-duplex Ethernet configurations will be accepted.

A simplified single-threaded example of this configuration is presented in Figure 1.

Figure 1
Client Private WAN
Extension to NFX MFI

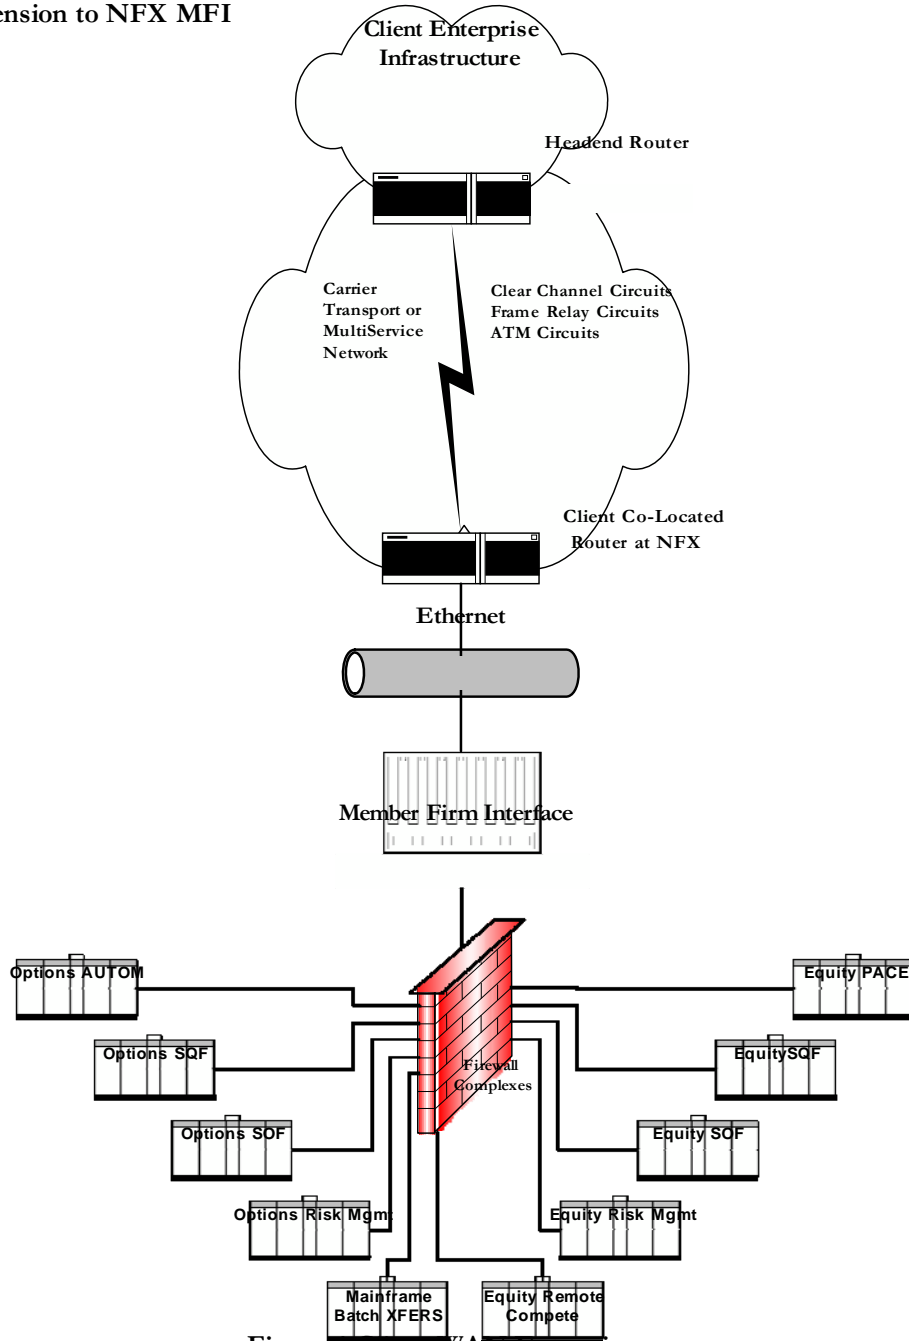


Figure 1 Client WAN Extension

FINANCIAL AUTOMATION COMMUNICATIONS	Section: Page: 8 Release: 1.0 Date: Sep 25, 2003 Rev. Date: Aug 6, 2012
MEMBER FIRM INTERFACE	TCP/IP COMMUNICATIONS MANUAL

Due to the mission critical nature of the traffic flowing between the NFX MFI and client infrastructure, NFX mandates that all client access be presented in a highly available redundant configuration; see Figure 2.

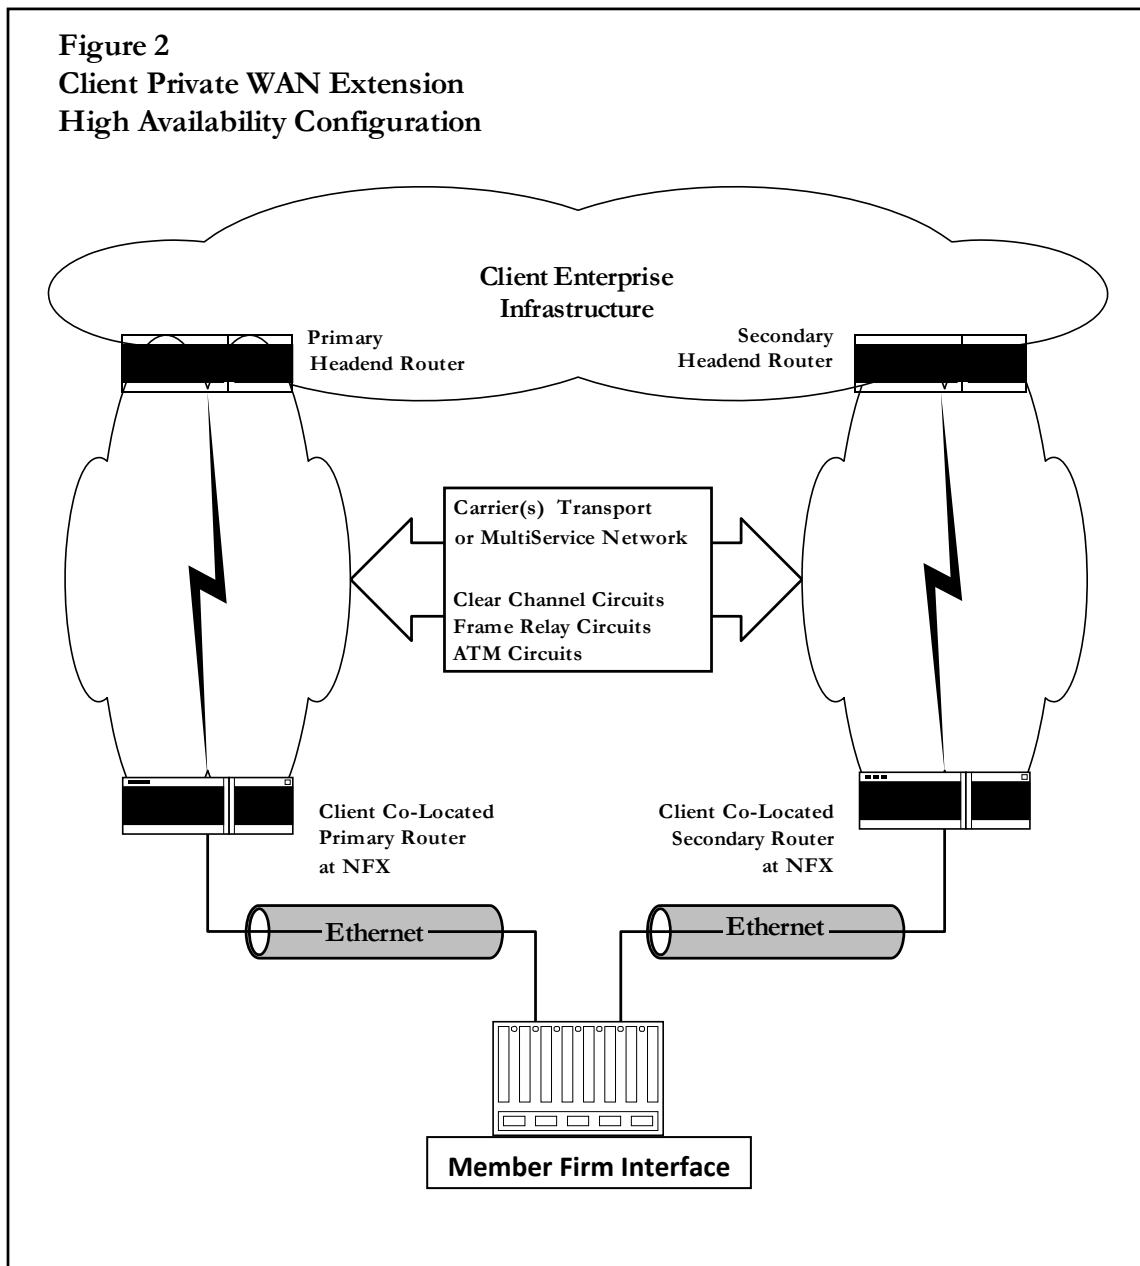


Figure 2 - Client WAN Extension High Availability Configuration

FINANCIAL AUTOMATION COMMUNICATIONS	Section: Page: 9 Release: 1.0 Date: Sep 25, 2003 Rev. Date: Aug 6, 2012
MEMBER FIRM INTERFACE	TCP/IP COMMUNICATIONS MANUAL

3.3.2 Method 2 – Client LAN Extension (Optical Ethernet)

Clients may establish connectivity to the NFX MFI with an Optical Ethernet Service provided by a metro or long-haul Ethernet Service Provider. The Optical Ethernet Service must provision an O-E (Optical to Electrical) conversion of the Optical Carrier at the NFX premise and provide NFX a copper based Ethernet interface handoff adhering to the 10BASE-T or 100BASE-TX standards.

1. The client is fully responsible for managing the carrier or service provider for all aspects of Optical Ethernet installation, provisioning and commissioning cycles. NFX will act as the coordinator for change control windows to allow carrier personnel on-site for Ethernet circuit installation and demarcation.
2. Please note, only full-duplex Ethernet configurations will be accepted.

An example of this configuration is presented in Figure 3.

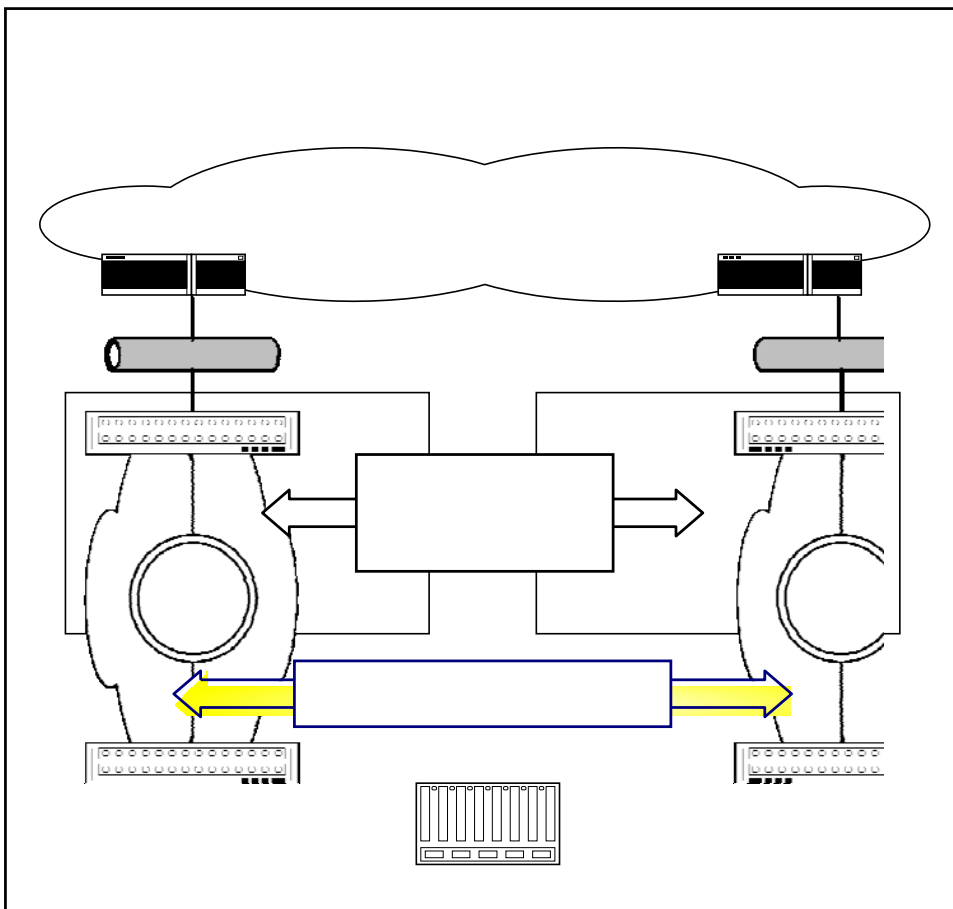
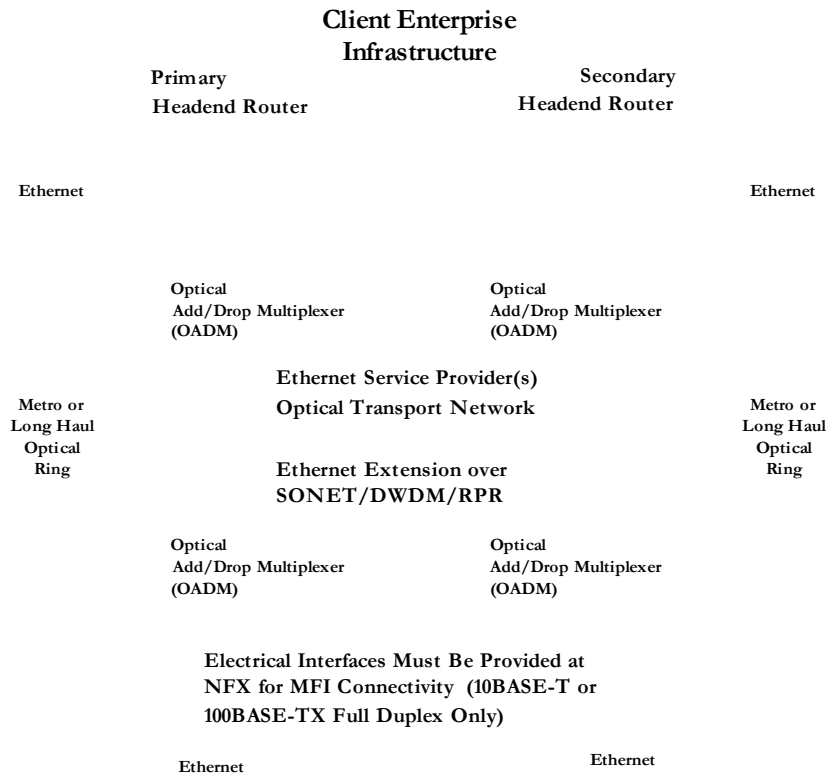


Figure 3

FINANCIAL AUTOMATION COMMUNICATIONS	Section: Page: 10 Release: 1.0 Date: Sep 25, 2003 Rev. Date: Aug 6, 2012
MEMBER FIRM INTERFACE	TCP/IP COMMUNICATIONS MANUAL

**Client LAN Extension
Over Optical Ethernet Service
High Availability Configuration**



Member Firm Interface

Figure 3 – Client LAN Extension

FINANCIAL AUTOMATION COMMUNICATIONS	Section: Page: 11 Release: 1.0 Date: Sep 25, 2003 Rev. Date: Aug 6, 2012
MEMBER FIRM INTERFACE	TCP/IP COMMUNICATIONS MANUAL

3.3.3 Method 3 - Extranet

Clients may establish connectivity to the MFI through an extranet service or IP VPN from a 3rd party service provider. These providers are compelled to adhere to the same requirements previously mentioned throughout this document and will meet the same high availability and redundancy specifications as a directly connected client.

1. It is the client's responsibility to deal directly with the service provider on interface specification requirements between the client's network edge and the service provider's DMARC.
2. All requirements for bandwidth capacity planning will be outlined in a separate subscription package based on the application service requested.

A sample of this configuration is depicted in Figure 4.

Figure 4
Extranet Service
High Availability Configuration

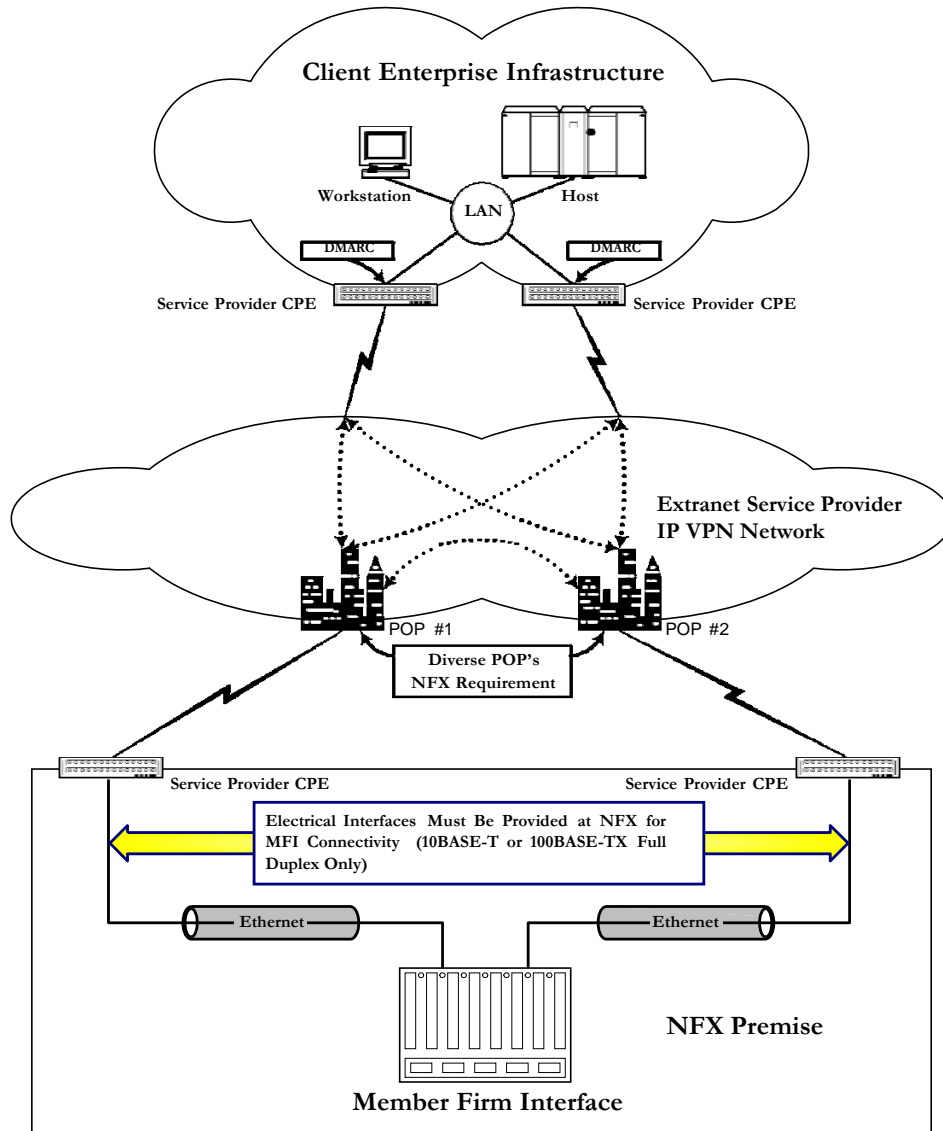


Figure 4 - Client Extranet

FINANCIAL AUTOMATION COMMUNICATIONS	Section: Page: 13 Release: 1.0 Date: Sep 25, 2003 Rev. Date: Aug 6, 2012
MEMBER FIRM INTERFACE	TCP/IP COMMUNICATIONS MANUAL

4 Background and Service Basics

4.1 CONNECTIVITY MIGRATION FOR EXISTING CLIENTS

Clients with existing connections that do not meet the specification outlined in the document should contact [NFX](#) to discuss the options available in migrating to a compliant method of connectivity. NFX will be able to assist you with planning services which may include network link consolidation, capacity planning, migration planning, coordination and execution.

4.1.1 Group Contact Information

For general questions regarding this specification and/or coordination of network and application service requests, you can contact the NFX via the following methods

E-Mail

NFX@nasdaqomx.com

Traditional Mail

NASDAQ OMX Futures Exchange
1900 Market Street
Philadelphia, PA 19103
Attention: Technical Account Management Group

Phone

+1 215 496 1571

For general questions regarding communication network connections and/or requirements, email NOCGroup@nasdaqomx.com.

5 Internet Protocol Overview

The NASDAQ OMX Futures Exchange has standardized on IPv4 as the network layer protocol upon which all application services will be delivered. This document is not intended as a full reference for IP and/or other protocols. MFI IP Interface technical requirements are located in Section X. The following information is being provided as a high level overview of fundamental Internet Protocol characteristics.

The vast end user community views an internal IP intranet or extranet extension as a single virtual infrastructure that interconnects all hosts and facilitates end-to-end communications between all end systems. The underlying architecture is often overlooked. In practicality, IP architectures that extend beyond corporate boundaries to make business-to-business arrangements possible are an abstraction of incongruent physical and logical network devices in mixed topologies that provide seamless transport for IP packet delivery. Higher levels of software in the Open Systems Interconnection (OSI) stack add most of the wealthy features that end users perceive.

For NASDAQ OMX | NFX Use Only
Proprietary and Confidential – 2012

FINANCIAL AUTOMATION COMMUNICATIONS	Section: Page: 14 Release: 1.0 Date: Sep 25, 2003 Rev. Date: Aug 6, 2012
MEMBER FIRM INTERFACE	TCP/IP COMMUNICATIONS MANUAL

Conceptually, TCP/IP architecture provides three sets of services:

1. Application Services
2. Reliable Transport Service
3. Connectionless "best effort" Packet Delivery Service

Although protocol software can be associated with each of these services, functional separation is achieved by the hierarchical encapsulation of each service in its entirety in the payload portion of the service directly beneath it. The concept of the hierarchy is depicted in Figure 5 and the stack encapsulation is depicted in Figure 6.

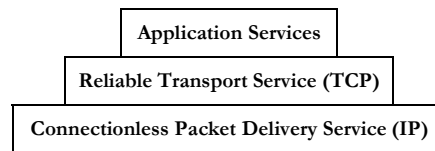


Figure 5 - TCP/IP Stack

Internet Protocol Packet Encapsulation **TCP Payload (Application Data)**

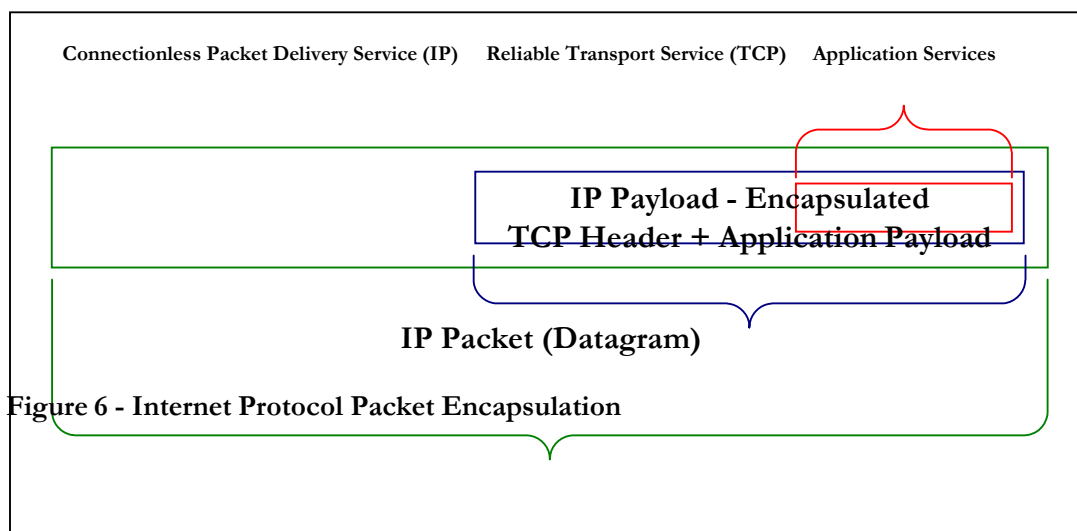


Figure 6 - Internet Protocol Packet Encapsulation

The robustness and adaptability of the stack hierarchy offers significant advantage since conceptual separation makes it possible to replace one service without disturbing others. The most rudimentary internet service consists of a packet delivery system. Technically, the service is defined as an unreliable, best-effort, connectionless packet delivery system. The service is

FINANCIAL AUTOMATION COMMUNICATIONS	Section: Page: 15 Release: 1.0 Date: Sep 25, 2003 Rev. Date: Aug 6, 2012
MEMBER FIRM INTERFACE	TCP/IP COMMUNICATIONS MANUAL

termed unreliable because packet delivery is not guaranteed. The packet may be lost, duplicated, experience delay (latency) or be delivered out of order. The IP protocol service will not detect such conditions, nor will it inform the sending or receiving host system. The service is characterized as connectionless since packets are independent of each other. A flow of packets sent from one host to another may transverse different network paths, experience some level of loss or degradation of service in that some packets are delivered, while others are lost. The service is coined “best-effort” because IP devices make every attempt to deliver packets, however unreliability occurs when system/device/network resources are exhausted or the underlying network is experiencing a failure.

IP performs three fundamental functions for enabling end-to-end data path communication. First, the IP protocol defines the basic unit of data transfer throughout an internet/intranet communications architecture. This data packet, also known as an IP datagram, specifies the exact message format of all data as it passes through the IP network. Second, IP performs a crucial routing function that directs traffic over a specific path for each path that exists between all devices in the overall infrastructure. Third, in addition to precise, formal functional specifications of data formats and routing protocols, IP includes a set of regulations that exemplify the notion of unreliable packet delivery. These regulations, known as Internet Standards (STDs) and Requests For Comments (RFCs) illustrate how host systems and routers should process and forward packets, how and when error messages should be generated by use of the Internet Message Control Protocol (ICMP) and the conditions under which packets can be discarded.

6 Reliable Stream Transport Service (TCP)

At the lower levels of the OSI stack, IP networks provide unreliable packet delivery. IP packets can be dropped or destroyed when transmission errors occur, hardware fails, or when traffic levels induce congestion and oversubscription of device resources. Networks that route packets dynamically can transmit them out-of-order, induce significant latency or duplicate data. Additionally, network topologies or devices may impose packet size restrictions or other traffic-based constraints in order to achieve efficient data transfer rates.

At the opposite end of the spectrum, applications often need to send large volumes of data from one host system to another. Using an unreliable delivery medium would require and place an inordinate burden on application developers to program reliability into each application developed for use with IP network architectures. The quintessential objective of the Transmission Control Protocol (TCP) is to create a general-purpose solution for providing reliable data transfer over an unreliable best effort delivery mechanism – IP. The TCP protocol resides at the Transport Layer of the OSI model and provides the crucial task of insulating the application from the details of IP networking and makes it possible to define a uniform interface for all applications to transfer data in a reliable, guaranteed-delivery stream transfer service. The TCP interface between application programs and reliable data delivery has five distinguishing characteristics:

1. Stream Orientation

When two applications (user processes) transfer data, the data is transferred in a stream of bits, divided into 8-bit octets also called bytes. The TCP stream delivery service on the destination host passes the bytes to the upper layer application in exactly the same sequence of bytes as which the sending host transmitted them across the underling infrastructure.

FINANCIAL AUTOMATION COMMUNICATIONS	Section: Page: 16 Release: 1.0 Date: Sep 25, 2003 Rev. Date: Aug 6, 2012
MEMBER FIRM INTERFACE	TCP/IP COMMUNICATIONS MANUAL

2. Virtual Circuit Connection

Creating a stream transfer is comparable to making a phone call. Before the stream can start, both the source and destination host systems programs' must signal their intent of establishing a call with their own respective operating system. Theoretically, one application places a call, which must be accepted by the other. The call setup and establishment in TCP is referred to as a "TCP Three Way Handshake". Once the handshake has been deemed successful, the TCP protocol informs the application that a session between host systems has been established and that stream transfer can commence.

3. Buffered transfer

Programs transmit a data stream across the virtual circuit connection by constantly passing bytes of application data to the TCP protocol software for transmission across the IP network. When transferring data, each application is unique in that it sends the data in whatever size pieces it finds convenient to transfer, which can be as small as a single byte. To increase data transfer efficiency, TCP protocol software is free to divide or join stream data into packets independent of the pieces the application program transfers. For applications that require immediate data transfer, even though the stream has not filled the transmit buffer of a sending host, the TCP stream service provides a "push" mechanism that allows the application to force TCP to transmit the data. The TCP software on the receiving host system recognizes the urgency of the data since TCP set the Push bit to "on" on the outbound host. As a result, TCP immediately makes the data available to the application.

4. Unstructured Stream

The TCP/IP stream service does not honor structured data streams. For example, there is no way for an inventory application to have the stream service mark boundaries between shipping records, or to identify the contents of the stream as being inventory data. Application programs using the stream service must understand stream content and agree on stream format before they initiate a connection.

5. Full Duplex Connection

The TCP/IP stream service allows simultaneous bi-directional stream transfer for all TCP connections. The stream service allows an application process to terminate flow in one direction while data continues to flow in the other direction. This is referred to as half duplex mode operation. The advantage of full duplex operation is that an application host can push data to a receiving host that is only consuming, not pushing. The receiving host can still send control messages back in the other stream direction.

6.1 RELIABILITY

The TCP protocol provides reliable data transfer over an underlying unreliable packet delivery mechanism through a technique called positive acknowledgement with retransmission. This technique requires the receiving host to send an acknowledgement message to the source once it has received the TCP segment. The source host maintains a history of each TCP segment it

FINANCIAL AUTOMATION COMMUNICATIONS	Section: Page: 17 Release: 1.0 Date: Sep 25, 2003 Rev. Date: Aug 6, 2012
MEMBER FIRM INTERFACE	TCP/IP COMMUNICATIONS MANUAL

transmits and waits for an acknowledgement before sending the next packet. The source host also initiates a timer upon packet transmission and retransmits the packet if the timer expires before an acknowledgement is received. This simple form of reliability does not make for efficient use of network bandwidth since only one packet can be sent before an acknowledgement is received. To overcome this issue, TCP implements the concept of a sliding window. In a nutshell, this allows a host system to send multiple packets before it expects to receive the acknowledgement for the first packet transmitted. TCP constantly negotiates and shifts the window size of the number of packets that can be sent since the receiving host sends an acknowledgement with a window size advertisement that indicates the amount of memory available in its receive buffer. The concept of the simple acknowledgement technique is depicted in Figure 7 and the sliding window concept is depicted in Figure 8. Notice that a window size is equal to three since the source sends 3 packets before the first acknowledgement is received.

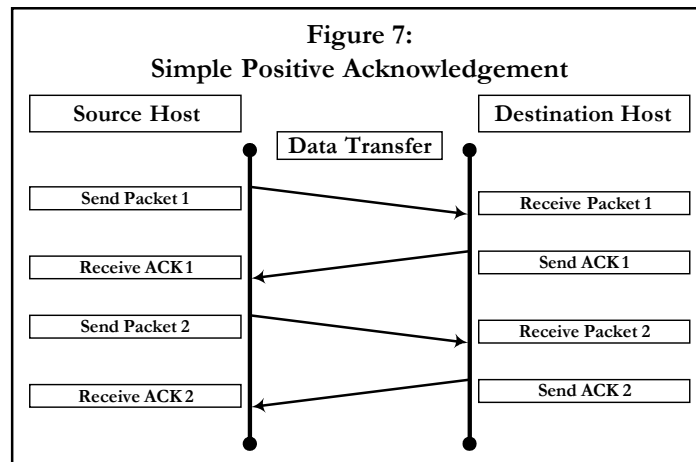


Figure 7 - Simple Positive Acknowledgement

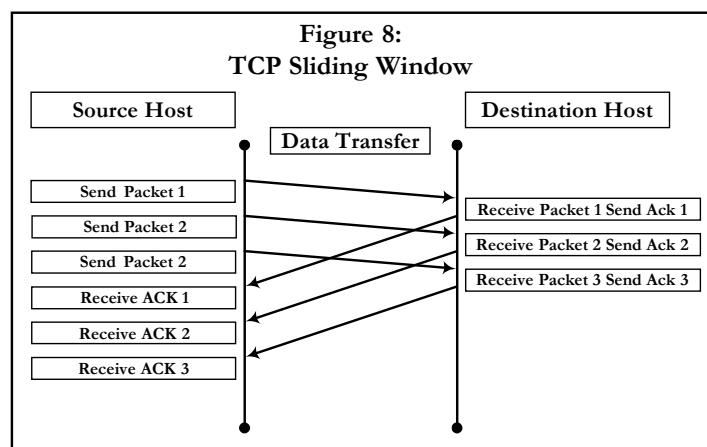


Figure 8 - TCP Sliding Window

FINANCIAL AUTOMATION COMMUNICATIONS	Section: Page: 18 Release: 1.0 Date: Sep 25, 2003 Rev. Date: Aug 6, 2012
MEMBER FIRM INTERFACE	TCP/IP COMMUNICATIONS MANUAL

7 Technical Interface Specifications

The physical interface to the NFX MFI is a 10/100 Megabit per second (Mbps) Category 5 Unshielded Twisted Pair (UTP) Ethernet Interface. The interface characteristics and supported industry standards are:

- o **Interface**
10BASE-T or 100BASE-TX
- o **Connector Type**
RJ-45 connector wired as MDI-X (Media Dependent Interface)
- o **Cable Specifications and Type**
Category 3, 4, or 5 UTP cable (10 Mb/s Operation)
Category 5 UTP cable (100 Mb/s operation)
- o **Maximum Distance**
328 ft. (100 meters)
- o **Standards Supported**
IEEE Std 802.3, 1998 Clause 4 Media Access Control CSMA/CD
IEEE Std 802.3, 1998 Clause 14 10 Base T New
IEEE Std 802.3, 1998 Clause 21 100Base T
ANSI/IEEE Std 802.1D, 1998 Media Access Control (MAC) Bridges
IEEE 802.1Q Virtual Bridged Local Area Networks
IEEE Std 802.3, 1998 Clause 28 (Auto-Negotiation)
IEEE Std 802.3, 1998 Clause 31 (MAC Control)
IETF Differentiated Services
- o **Data Rate and Encoding**
10 Mb/s Manchester encoding
100 Mb/s 4B/5B encoding

FINANCIAL AUTOMATION COMMUNICATIONS	Section: Page: 19 Release: 1.0 Date: Sep 25, 2003 Rev. Date: Aug 6, 2012
MEMBER FIRM INTERFACE	TCP/IP COMMUNICATIONS MANUAL

7.1 NFX STANDARDS

The NFX MFI Ethernet interface will be configured in the following manner:

1. IEEE Std 802.3, 1998 Clause 28 is administratively disabled. The NFX MFI will not participate whatsoever in any form of Auto-Negotiation for 10Base-T or 100Base-TX interfaces.
2. The NFX MFI will be administratively configured to accept Full Duplex Ethernet connections only. Half Duplex is not supported. As a by-product of this requirement, CSMA/CD is disabled.
3. All client-facing ports will be configured for type "access". The NFX MFI will not accept 802.1Q tagged traffic.
4. The NFX MFI will not honor any IETF Differentiated Services marked traffic, nor will it honor any 802.1p-prioritized traffic.
5. The NFX MFI will not participate, nor support the Spanning Tree Protocol.
6. NFX will only accept Category 5 UTP patch cables from client CPE

7.2 CLIENT REQUIREMENTS FOR CONNECTIVITY

It is the client responsibility to ensure the following connectivity requirements are met:

1. It is a requirement that from the client's CPE, two Category 5 UTP cables be provided to interface to the NFX demarcation point. These two connections must be equivalent at the Data Link layer (Layer 2) of the OSI model. The client is responsible for any Layer-2 or higher failover requirements on their equipment.
2. The Customer is responsible for all equipment, cabling, configuration and carrier services up to the NFX MFI interface.
3. The NFX MFI interface will be a female RJ-45 interface, wired to TIA-568 specifications. Two variations of TIA-568 exist, TIA-568-A and TIA-568-B. NFX recommends and utilizes the TIA-568-B specification for all internal wiring. See Figure 9.

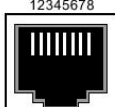
Connector	Pin number	Signal
 <p>12345678 8020EA</p>	1	Input receive data + (RX+)
	2	Input receive data - (RX-)
	3	Output transmit data + (TX+)
	6	Output transmit data - (TX-)
	4, 5, 7, 8	Not used

Figure 9 - RJ-45 MDI-X

- The MFI is configured as a MDIX (Medium Dependent Interface – Crossover) port, expecting to be connected directly to a MDI (Medium Dependent Interface) device via a straight-through cable. Check the CPE to determine which interface type the equipment specifies. See Figure 10.

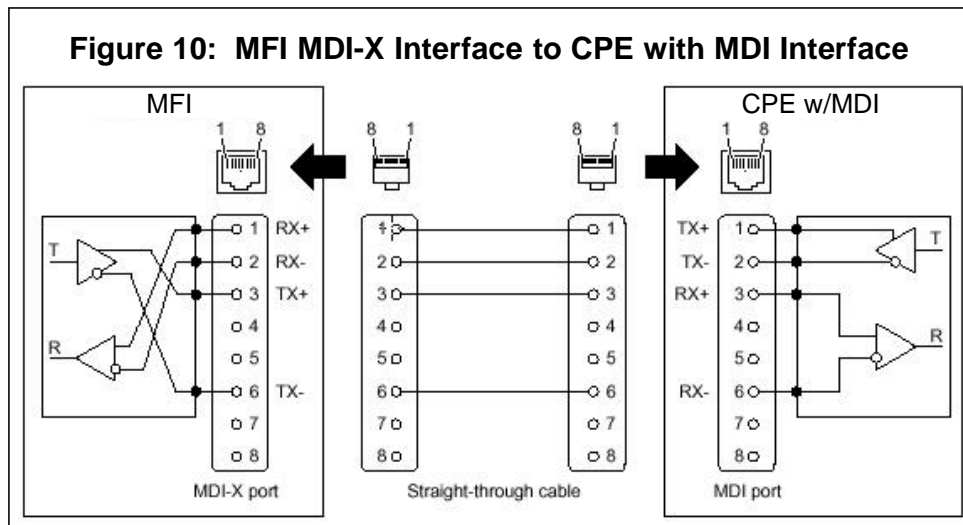


Figure 10 - MFI Straight-Through Configuration

- If connected to another MDIX device, (A Hub or Switch generally), the cabling must be crossed over to assure that transmit and receive pairs are swapped between the two like devices.

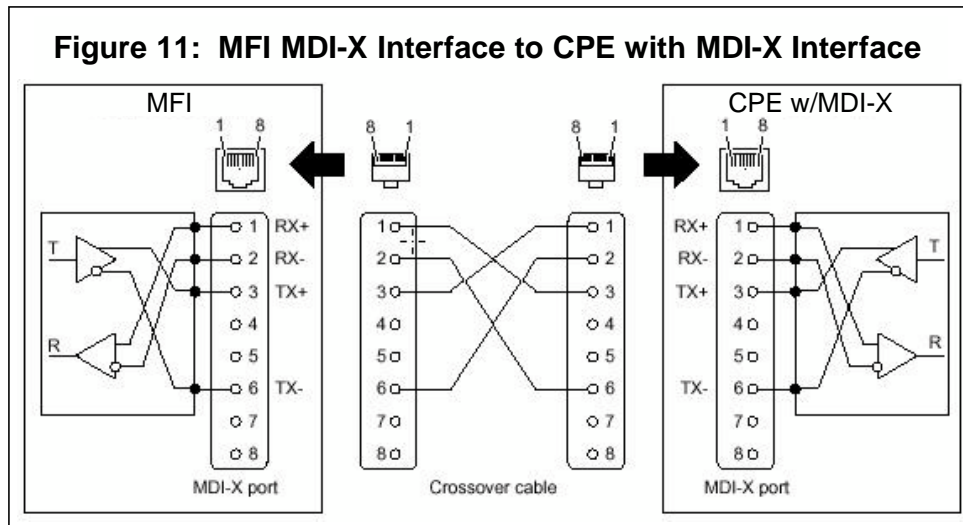
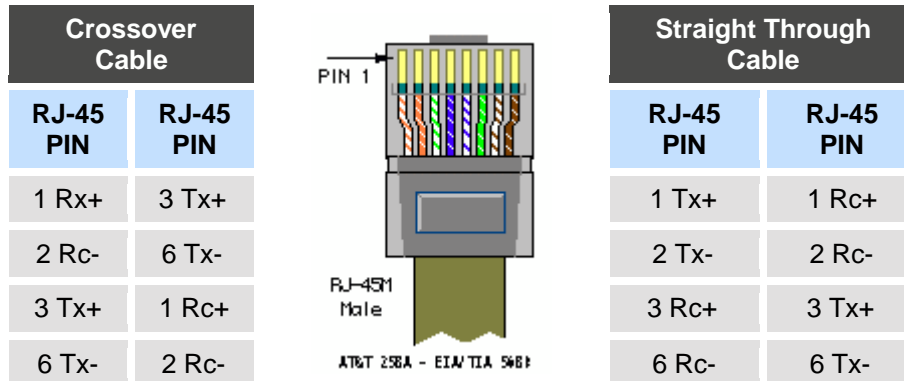


Figure 11 - MFI Crossover Configuration

FINANCIAL AUTOMATION COMMUNICATIONS	Section: Page: 21 Release: 1.0 Date: Sep 25, 2003 Rev. Date: Aug 6, 2012
MEMBER FIRM INTERFACE	TCP/IP COMMUNICATIONS MANUAL

7.3 CABLING SPECIFICATIONS



Note: The standard connector view shown is color-coded for a straight thru cable

Category 5 wiring standards:

EIA/TIA 568A/568B and AT&T 258A define the wiring standards and allow for two different wiring color codes.

Pin #	Signal	EIA/TIA 568A	AT&T 258A, or EIA/TIA 568B <i>NFX Recommended</i>	Ethernet 10BASE-T 100BASE-T
1	Transmit+	White/Green	White/Orange	X
2	Transmit-	Green/White or Green	Orange/White or Orange	X
3	Receive+	White/Orange	White/Green	X
4	N/A	Blue/White or Blue	Blue/White or Blue	Not used *
5	N/A	White/Blue	White/Blue	Not used *
6	Receive-	Orange/White or Orange	Green/White or Green	X
7	N/A	White/Brown	White/Brown	Not used *
8	N/A	Brown/White or Brown	Brown/White or Brown	Not used *

FINANCIAL AUTOMATION COMMUNICATIONS	Section: Page: 22 Release: 1.0 Date: Sep 25, 2003 Rev. Date: Aug 6, 2012
MEMBER FIRM INTERFACE	TCP/IP COMMUNICATIONS MANUAL

Only two pairs of wires in the eight-pin RJ-45 connector are used to carry Ethernet signals. Both **10BASE-T** and **100BASE-T** use the same pins, a **crossover cable** made for one will also work with the other.

- o Pairs may be solid colors and not have the stripe.
- o Category 5 cables must use Category 5 rated connectors.
- o ****Note: Even though pins 4, 5, 7, and 8 are not used, it is mandatory that they be present in the cable.***

FINANCIAL AUTOMATION COMMUNICATIONS	Section: Page: 23 Release: 1.0 Date: Sep 25, 2003 Rev. Date: Aug 6, 2012
MEMBER FIRM INTERFACE	TCP/IP COMMUNICATIONS MANUAL

8 IP Requirements

Currently there are two types of Internet Protocol (IP) addresses in public and private networks: IP version 4 (IPv4) and IP version 6 (IPv6). IPv4 was publicly deployed on 1 January 1983 and is still the most commonly used version. IPv4 addresses are 32-bit numbers often expressed as 4 octets in "dotted decimal" notation (for example, 192.0.32.67). Public deployment of the IPv6 protocol began in 1999. IPv6 addresses are 128-bit numbers and are conventionally expressed using hexadecimal strings (for example, 1080:0:0:0:8:800:200C:417A).

NFX does not currently have a deployment schedule for IPv6, but efforts to research, test and support IPv6 are underway.

Both IPv4 and IPv6 addresses are assigned in a delegated manner. Users are assigned IP addresses by Internet service providers (ISPs). ISPs obtain allocations of IP addresses from a local Internet registry (LIR) or national Internet registry (NIR), or from their appropriate Regional Internet Registry (RIR):

[APNIC \(Asia Pacific Network Information Centre\)](#) - Asia/Pacific Region

[ARIN \(American Registry for Internet Numbers\)](#) - North America and Sub-Sahara Africa

[LACNIC \(Regional Latin-American and Caribbean IP Address Registry\)](#) – Latin America and some Caribbean Islands

[RIPE NCC \(Réseaux IP Européens\)](#) - Europe, the Middle East, Central Asia, and African countries located north of the equator

The IANA (Internet Assigned Numbers Authority) is the organization which allocates IP addresses from the pools of unallocated addresses to the RIRs according to their established needs. When an RIR requires more IP addresses for allocation or assignment within its region, the IANA makes an additional allocation to the RIR.

Due to the demands placed on the available pool of valid Registered IANA addresses, a pool of private addresses was set aside by the IETF (Internet Engineering Task Force – www.ietf.org) in

February of 1996 via RFC1918. (<ftp://ftp.rfc-editor.org/in-notes/rfc1918.txt>). This document set aside the following address space:

10.0.0.0	10.255.255.255	(10/8 prefix)
172.16.0.0	172.31.255.255	(172.16/12 prefix)
192.168.0.0	192.168.255.255	(192.168/16 prefix)

This set aside address space is utilized by many institutions internally, and the potential for conflict in connecting two or more institutions utilizing RFC1918 address space is significant. For this reason, NFX requires that all IP connectivity to NFX infrastructure be sourced by Registered Addresses. (This does not preclude the Customer's internal use of RFC1918 addresses, but the Customer is responsible for translating the sources to Registered Addresses, prior to handoff to the NFX demarcation point.)

For NASDAQ OMX | NFX Use Only
Proprietary and Confidential – 2012

FINANCIAL AUTOMATION COMMUNICATIONS	Section: Page: 24 Release: 1.0 Date: Sep 25, 2003 Rev. Date: Aug 6, 2012
MEMBER FIRM INTERFACE	TCP/IP COMMUNICATIONS MANUAL

8.1 NETWORK ADDRESS TRANSLATION OVERVIEW

In its simplest configuration, the Network Address Translator (NAT) as defined in RFC 1631 operates on a routing device connecting two networks together; one of these networks (designated as inside) is addressed with either private or obsolete addresses that need to be converted into legal addresses before packets are forwarded onto the other network (designated as outside). The translation operates in conjunction with routing, so that NAT can simply be enabled on a client-side CPE router when translation is desired. The goal of NAT is to provide functionality as if the private network (client infrastructure) had globally unique addresses and the NAT device was not present. Network address translation can be performed via two methods, one-to-one and many-to-one.

In a one-to-one translation, the private addresses are individually mapped to a registered address. The device performing the translation modifies the source or destination address as necessary, without modifying other frame data.

In a many-to-one translation, multiple private addresses are mapped to a single registered address (this is sometimes referred to a Port Address translation or PAT). To accomplish mapping the numerous connections from the various private addresses the device performing the translation not only modifies the source and destination address as necessary, but modifies the source or destination Transport layer (TCP or UDP) port number. Many applications will not function via a many-to-one or PAT translation configuration.

Customers will be required to identify the address space and mask for all sources which will access the NFX infrastructure.

8.2 NFX IP TRAFFIC AND ROUTING POLICY

The following is the NFX routing policy for MFI connectivity:

1. Routes Announced by NFX:

NFX does not administer a dynamic IP routing protocol on the MFI edge connection. All routing information exchange between the client and NFX is performed through static routing.

2. Accepting Routes from Clients

NFX will accept host based routes configured with a 32 bit binary subnet mask. At NFX's discretion, network based routes may be accepted, but must adhere to the summarization principles of Classless Inter-Domain Routing (CIDR), RFCs 1517, 1518, 1519 and 1520.

3. Multi-homing to the MFI

Although redundant client connections are required at the MFI interface, NFX will treat this network configuration in a primary/backup relationship. NFX will not perform load sharing across the redundant connections.

FINANCIAL AUTOMATION COMMUNICATIONS	Section: Page: 25 Release: 1.0 Date: Sep 25, 2003 Rev. Date: Aug 6, 2012
MEMBER FIRM INTERFACE	TCP/IP COMMUNICATIONS MANUAL

4. Sharing of 3rd Party MFI access

NFX reserves the right to permit or deny additional client access through a 3rd party provider that is acting as a primary agent on the behalf of an existing NFX client that has been contracted by that client for private network extension in which NFX has not authorized the 3rd party service provider as an IP aggregator for servicing multiple NFX clients.

5. IP Traffic Types

The current NFX IP connectivity specification does not permit any Tunneled or IP-encapsulated protocols, only native IPv4 frames.

6. IP Unicast Traffic

IP Unicast traffic may be filtered at the NFX ingress point by source address, protocol ID, source and/or destination port, or other criteria. In addition, NFX may filter outbound traffic to client connections.

7. IP Multicast Traffic

There are currently no NFX-sourced Multicast applications that are available for client subscription.

8. Network Security

The stability and security of NFX's network infrastructure and host systems are of the highest priority. Client's connecting to the NFX MFI are strongly encouraged to assess and adjust their security requirements on a pro-active basis. It is a "best current practice" that subscribers of external network services to their Autonomous System (AS) protect the ingress point of external traffic through a secured DMZ. It is highly recommended that the NFX client implement a firewall solution at the ingress of their network.

NFX reserves the right to restrict bandwidth, block protocols or ports, filter, capture and analyze data from or to any connected client connection. If necessary, NFX reserves the right to disable or disconnect any link which is believed to be compromising security in order to protect NFX resources and/or other connected clients.

9 Network Support

The NFX Operation Center (NOC) is available for telephone support between the hours of 8:00 am and 5:00 pm ET.

NFX NOC Telephone: +1 212 231 5049