



SMALL EXCHANGE

Infrastructure and Connectivity Guide

Version 1.6 | March 16, 2019

Table of Contents

[1. Introduction](#)

[1.1 SMFE Environments](#)

[1.2 SMFE Backbone - ASN 33016](#)

[2. Connecting to the SMFE Backbone](#)

[2.1 SMFE POP and Environment Locations](#)

[2.2 Colo Cross Connects](#)

[2.3 Connect via an extranet provider](#)

[2.4 Direct Peering Guidelines](#)

[3. Market Data and FIX Order Entry](#)

[3.1 Market Data - SMFE Market Data Feed](#)

[3.2 FIX Order Entry and Drop Copies](#)

[4. IP Addressing](#)

[4.1 Unicast IP Address Summary](#)

[4.2 Multicast IP Address Summary](#)

[4.3 Multicast Channels](#)

[4.3.1 Production](#)

[4.3.2 Disaster Recovery <Coming Soon>](#)

[4.3.3 Certification / Dev](#)

[4.4 SMFE Routing Policy Building Blocks](#)

[5. SFTP](#)

1. Introduction

1.1 SMFE Environments

The Exchange operates the following client facing environments:

- Production
- Disaster Recovery
- Certification / Dev

1.2 SMFE Backbone - ASN 33016

The SMFE Backbone is the logical network connection point between the Exchange and all external parties. All SMFE Environments are accessible via the SMFE Backbone. Regardless of where an environment exists physically, connectivity is provided by peering with the SMFE Backbone at any one of our POP (Point of Presence) locations.

The SMFE Backbone is comprised of highly reliable and geographically diverse network connections between our data centers. Parties are encouraged to connect to multiple POPs in order to ensure resilient access to the Exchange’s environments. Each POP supports redundant connectivity via diverse edge switch-routers.

Connected parties have the option to route traffic to the Exchange environments in a way that makes sense for their network/infrastructure and business needs. A BGP peering policy that provides preferred routing during normal operations and predictable routing in a failure scenario can be easily developed.

2. Connecting to the SMFE Backbone

2.1 SMFE POP and Environment Locations

The Exchange maintains Backbone POPs in the following locations:

SMFE POP Code	Location
AB2	CyrusOne Financial Campus, Building 2 - Aurora, IL
NY2	Equinix NY2 - 275 Hartz Way, Secaucus, NJ NY2/4/5 Campus
CH1	Equinix CH1 - 350 E. Cermak, Chicago, IL

The following table describes the physical location of the SMFE Environments and preferred POP that offers the lowest latency access to those environments.

Environment	Physical Location	Preferred POP
Production	AB2	AB2
Disaster Recovery	NY2	NY2
Certification and Conformance	NY2	NY2

2.2 Colo Cross Connects

In each location, 10G-LR optics over single-mode fiber cross connects is the standard way to connect to the Exchange. Please contact the Exchange to plan all physical and logical connectivity.

2.3 Connect via an extranet provider

The SMFE has partnered with the following extranet providers to offer access to the SMFE Backbone:

- ICE Global Network (SFTI)
- Options IT

2.4 Direct Peering Policy

The following high level peering policy applies for all direct peering relationships with the Exchange:

Peers must:

- Support BGP peering
- Support PIM Sparse-mode peering (if receiving market data)

Peers should:

- Peer with a public BGP ASN
- Advertise registered IP addresses to the Exchange

Peers may:

- Request a private ASN with which to peer with the Exchange
- Request RFC1918 allocation from the Exchange in lieu of using public IP addresses

3. Native Market Data and FIX Order Entry

3.1 Market Data - SMFE Market Data Feed

The Small Exchange proprietary Market Data Feed (MDF) provides a “Market by Order” view of market activity and is delivered via industry standard methods:

- PIM Sparse-Mode UDP Multicast
- Redundant and Diverse A/B channels
- SBE based proprietary protocol

The A and B multicast feeds are disseminated directly from the Matching Engines via physically diverse network paths maintained from the Matching Engine server (separate PCI cards) to “A” and “B” top of rack leaf switches, then through diverse spine layer switches, and then to the edge switches. The edge switches aggregate the A and B feeds, which offers flexibility in how the connected party chooses to receive the data. Since each edge switch is capable of delivering both the A and the B feed, the connected party has the choice as to how they wish to consume these feeds.

Separate static “A” and “B” RPs and diverse A and B multicast source networks allow connected parties to easily receive, understand, and control the multicast traffic once it has entered their routing environment.

For the highest level of redundancy, it’s recommended to connect with two cross connects, taking A from one and B from the other. Details on the multicast groups, RPs, and source networks follow in section 4 “IP Addressing”.

3.2 FIX Order Entry and Drop Copies

Orders are entered using a typical FIX 4.4 based order management API. Order sessions can be requested via the Exchange’s external ticketing system.

FIX Drop Copy sessions are available on request. Drop Copy sessions can be requested via the Exchange’s external ticketing system.

Link to Small Exchange JIRA Ticketing portal:

<https://smallx.atlassian.net/servicedesk/customer/portal/5>

4. IP Addressing

4.1 Unicast IP Address Summary

The following are aggregate address blocks, per environment - these are not netblocks you should use directly in your routing policy. See below, in section 4.4, for a detailed breakdown of networks within these summaries - use that info to easily build your desired routing policy with the exchange.

Production	
204.137.28.0 /25	All unicast services (Order Entry, Drop copy, Admin and Risk Monitor, etc.)

Disaster Recovery	
204.63.168.0 /25	All unicast services (Order Entry, Drop copy, Admin and Risk Monitor, etc.)

Staging	
204.63.168.128 /25	All unicast services (Order Entry, Drop copy, Admin and Risk Monitor, etc.)

Certification / Conformance	
204.63.170.0 /27	All unicast services (Order Entry, Drop copy, Admin and Risk Monitor, etc.)

4.2 Multicast IP Address Summary

Use this section to create your static RP mappings. The multicast source IP blocks are summaries - these are not the actual prefixes that will be advertised via BGP. See section 4.4, below, for further breakdown of these summaries.

Production			
A Feeds - RP 204.137.30.253		B Feeds - RP 204.137.30.254	
A Multicast Groups	A Source Networks	B Multicast Groups	B Source Networks
233.247.221.0 /26	204.137.29.0 /26	233.247.221.64 /26	204.137.29.64 /26
Disaster Recovery			
A Feeds - RP 204.63.170.253		B Feeds - RP 204.63.170.254	
A Multicast Groups	A Source Networks	B Multicast Groups	B Source Networks
233.142.36.0 /26	204.63.169.0 /26	233.142.36.64 /26	204.63.169.64 /26
Certification and Dev			
A Feeds - RP 204.63.170.251		B Feeds - RP 204.63.170.252	
A Multicast Groups	A Source Networks	B Multicast Groups	B Source Networks
233.142.36.128 /28	204.63.170.32 /28	233.142.36.144 /28	204.63.170.48 /28

4.3 Multicast Channels

4.3.1 Production

Production Segment 1 SFX, SPRE, SMGO			Production Segment 2 SM75, S10Y		
A Feeds			A Feeds		
Incremental			Incremental		
233.247.221.1	21001	204.137.29.1 204.137.29.2 204.137.29.3	233.247.221.3	21001	204.137.29.4 204.137.29.5 204.137.29.6
Snapshot			Snapshot		
233.247.221.2	21002	204.137.29.1 204.137.29.2 204.137.29.3	233.247.221.4	21002	204.137.29.4 204.137.29.5 204.137.29.6
B Feeds			B Feeds		
Incremental			Incremental		
233.247.221.65	22001	204.137.29.65 204.137.29.66 204.137.29.67	233.247.221.67	22001	204.137.29.81 204.137.29.82 204.137.29.83
Snapshot			Snapshot		
233.247.221.66	22002	204.137.29.65 204.137.29.66 204.137.29.67	233.247.221.68	22002	204.137.29.81 204.137.29.82 204.137.29.83
Segment 1 Production Retransmission Servers			Segment 2 Production Retransmission Servers		
204.137.28.9	UDP 8916		204.137.28.10	UDP 8916	
204.137.28.69	UDP 8916		204.137.28.70	UDP 8916	

4.3.2 Disaster Recovery <Coming Soon>

4.3.3 Certification / Dev

EXISTING			<i>FUTURE - SM75 and S10Y will move here</i>		
Certification Segment 1 SM75, S10Y, SFX, SPRE, SMGO			Certification Segment 2 SM75, S10Y		
A Feeds			A Feeds		
Incremental			Incremental		
233.247.221.1	21001	204.137.29.1	233.247.221.3	21001	204.137.29.2
Snapshot			Snapshot		
233.247.221.2	21002	204.137.29.1	233.247.221.4	21002	204.137.29.2
B Feeds			B Feeds		
Incremental			Incremental		
233.247.221.65	22001	204.137.29.65	233.247.221.67	22001	204.137.29.66
Snapshot			Snapshot		
233.247.221.66	22002	204.137.29.65	233.247.221.68	22002	204.137.29.66
Segment 1 Certification Retransmission Server			Segment 2 Certification Retransmission Server		
204.63.170.12	UDP 8916			UDP 8916	

4.4 SMFE Routing Policy Building Blocks

The goal of this section is to illustrate the BGP policy building blocks you can use to identify and apply policy to the Small Exchange advertised prefixes. For example, in the first listing, "**AB2: Aurora: PROD**", we see the summary address 204.137.28.0/25, but the networks being advertised will all be /27's. So if you build your policy to accept all /27's contained within that /25, then you will always have all of the PROD unicast services networks (FIX services, market data retrans, etc...) Alternatively, you can use the BGP community attribute that such prefixes are tagged with. In this case, 63453:100.

AB2: Aurora: PROD

204.137.28.0/25 eq /27 = All PROD unicast networks (community = 63453:100)

204.137.29.0/26 eq /28 = All PROD A side multicast sources (community = 63453:200)

204.137.30.253/32 = PROD A side RP (community = 63453:200)

204.137.29.64/26 eq /28 = All PROD B side multicast sources (community = 63453:300)

204.137.30.254/32 = PROD B side RP (community = 63453:300)

NY2: Secaucus: DR

204.63.168.0/25 eq /27 = All DR unicast networks (community = 36388:100)

204.63.169.0/26 eq /28 = All DR A side multicast sources (community = 36388:200)

204.63.170.253/32 = DR A side RP (community = 36388:200)

204.63.169.64/26 eq /28 = All DR B side multicast sources (community = 36388:300)

204.63.170.254/32 = DR B side RP (community = 36388:300)

NY2: Secaucus: CERT/Dev

204.63.170.0/27 = CERT unicast network (community = 36388:160)

204.63.170.32/28 = A side CERT multicast sources (community = 36388:260)

204.63.170.251/32 = CERT A side RP (community = 36388:260)

204.63.170.48/28 = B side CERT multicast sources (community = 36388:360)

204.63.170.252/32 = CERT B side RP (community = 36388:360)

5. SFTP

The Exchange uses SFTP for general file passing between the Exchange and various participants. Request SFTP provisioning here:

<https://smallx.atlassian.net/servicedesk/customer/portal/5/group/18/create/100>

Here is basic information on how the service is provisioned.

1. Authentication
 - a. SSH public key authentication, RSA keys of at least 2048 bits, OpenSSH format
 - b. Use a separate key-pair per environment
 - c. Attach your public key files to the JIRA Service Desk request
 - d. Your username will be given to you
2. Source IP address(es)
 - a. Enter your source IP address(es) used to access the SFTP service in the JIRA Service Desk request.
3. The folder structure is like this:

```
/---|
  |
  |----public
  |
  |----FIRM_CODE
  |
  |-----in
  |-----out
```

Production/DR SFTP server (*Hostname and IPs do not change when moving to DR*)

Hostname: sftp.smallexchange.com

Port: 22

IP addresses: 3.20.74.92, 3.130.168.78, 3.12.1.167

Certification SFTP server

Hostname: cert-sftp.smallexchange.com

Port: 22

IP addresses: 3.12.100.237, 3.20.120.220, 3.12.106.248