

Draft Kraken US Perps connectivity guide

Market maker onboarding | London and remote connectivity | US Perpetual Futures | 4 September 2026

Document set

This guide covers **how to reach us**: your connectivity options, the cross-connect process, and which service runs on which endpoint. Protocol detail lives in the companion references below. All are available from your account manager.

Document	Covers
This guide	Connectivity options, colocation and cross-connects, the LOA process, IP allow-listing, endpoints, onboarding sequence.
Kraken US Perps FIX protocol reference	Order entry and drop copy: session establishment, authentication, message-by-message field reference, enumerations.
Kraken public SBE protocol reference	Market data: session negotiation and authentication, subscription model, book recovery, message and enumeration reference.
Colocation High Level Network Design	Physical and logical design for direct cross-connects: topology options, addressing, BGP, VLAN layout. Issued under NDA.

The machine-readable artefacts — the FIX 4.4 XML data dictionary and the SBE schema XML — are released alongside the two protocol references.

Connectivity options

Kraken's core trading infrastructure is in **Equinix London**; the exact facility is disclosed under NDA. There are two ways in. You can sit inside the same Equinix campus and reach us over a **direct cross-connect**, or you can come in over **AWS London (eu-west-2)**. The hostname you connect to tells you which path you are on — see *Endpoints*.

Option	Latency to Kraken	What you need
Direct colocation cross-connect	< 1 ms	Signed NDA and an Equinix London presence. Diverse cross-connect pair provisioned against Kraken-issued LOAs.
Beeks Financial Cloud	< 1 ms	A Beeks VPS or bare-metal subscription. Served over the same colocation route, with no IP allow-listing required.
AWS London (eu-west-2)	~2 ms	An AWS presence in eu-west-2, reached by VPC peering, AWS PrivateLink, or IP allow-listed public ingress. CIDR and account details are exchanged with Kraken during onboarding.

Choosing between the low-latency options

Direct colocation and Beeks deliver the *same* latency — all cables terminate in the same meet-me room. Direct colocation buys you control over your own hardware, not additional speed. Beeks is the faster route to market and needs no IP allow-listing; direct colocation is the better long-term home if you want to own the stack. Beeks presence and Kraken's facility are not in the same IBX, so the inter-site cross-connect carries a small latency cost; a direct cross-connect from a Beeks-hosted server to Kraken is supported on approval.

Connecting from outside London

Kraken partners with **BSO** and **Avelacom** for inter-region connectivity. Indicative round trip: Europe to US approximately 60 ms, Europe to Tokyo approximately 140 ms. Note that a carrier path *within* London still traverses the provider core and typically adds 1–2 ms — for same-region traffic, a direct cross-connect or Beeks beats a carrier circuit.

Taking a cross-connect

If you are taking the direct colocation route. Read the *Colocation High Level Network Design* alongside these steps — it carries the addressing, BGP and VLAN detail.

1. **Choose a topology.** *Layer 3 routed / BGP* — dual cross-connects, four point-to-point networks, authenticated BGP sessions on each. **Kraken's preferred option.** Or *Layer 2 switched* — dual cross-connects with Prod and UAT VLANs trunked across both, static routing, Kraken hosting the gateway address in each VLAN.
2. **Send us your Equinix entity name, contact name and contact email.** This is what we need to issue the LOA.
3. **Receive two LOAs.** Kraken issues **two** Letters of Authorization so you can provision a **diverse pair**. Check with your carrier that both are ordered — a single cross-connect leaves you with no redundancy.
4. **Order and install.** Your carrier (Avelacom, BSO, McKay, Beeks or equivalent) orders against the LOAs and terminates in the meet-me room.
5. **Exchange BGP credentials.** For Layer 3 designs Kraken shares the BGP secret directly. **Confirm you have received it** before scheduling turn-up — this is a common source of delay.
6. **Test on UAT, then production.** The same physical pair carries both VLANs, so UAT testing runs over the production path.

Physical specification

Cross-connects	2, diverse
Speed	10 Gbps
Media	Single-mode fiber, terminated in the Meet Me Room
Encryption	MACSec optional
Segmentation	802.1Q VLAN tagging, separate VLANs for Prod and UAT
Client addressing	A /24 from RFC 6598 CGNAT space (100.64.0.0/10), split into two /25s for Prod and UAT
Path selection	Kraken does not enforce an optimal path; set BGP metrics on your side

IP allow-listing

Required only when you reach us over AWS London (eu-west-2). Traffic arriving over a direct cross-connect or from Beeks does not need allow-listing — the cross-connect itself is the trust boundary.

What we need	The static source IP addresses your sessions originate from, separately for UAT and production
Where it applies	At network level and per API key . Both must be correct or you will not connect
Lead time	Changes are not instant. Submit them through your account manager ahead of when you need them
Capacity	Allow-list capacity is finite. Consolidate onto as few source addresses as your architecture allows

Endpoints

Four services, each on its own TCP endpoint. All are TLS-terminated. Entitlement is granted as a set — a firm enabled for US Perps receives access to all four.

The hostname prefix selects your access path. Use colo-london . if you reach us over a direct cross-connect or from Beeks. Use eu-west-2-1 . if you come in over AWS London. The ports are the same on both.

Direct colocation and Beeks

colo-london.vip-fix.uat.kraken.com:4010 FIX - order entry
colo-london.vip-fix.uat.kraken.com:4011 FIX - drop copy
colo-london.vip-sbe.uat.kraken.com:4500 SBE - market data, incremental feed
colo-london.vip-sbe.uat.kraken.com:4501 SBE - market data, snapshots

AWS London (eu-west-2)

eu-west-2-1.vip-fix.uat.kraken.com:4010 FIX - order entry
eu-west-2-1.vip-fix.uat.kraken.com:4011 FIX - drop copy
eu-west-2-1.vip-sbe.uat.kraken.com:4500 SBE - market data, incremental feed
eu-west-2-1.vip-sbe.uat.kraken.com:4501 SBE - market data, snapshots

Endpoint	What it does	Reference
FIX 4010 — order entry	Submit, amend, cancel and track orders. Receive execution reports. Instrument definitions and trading status come from SBE, not FIX.	FIX protocol reference
FIX 4011 — drop copy	Read-only execution report feed for reconciliation and compliance. 48 hours of recovery, at-least-once delivery, deduplicate on ExecID.	FIX protocol reference
SBE 4500 — incremental feed	Sequenced live market data: order book, trades, instrument definitions, security status, trading session status. Connect continuously.	SBE protocol reference
SBE 4501 — snapshots	Full order-book state, chunked. Connect at startup and on recovery only.	SBE protocol reference

Endpoint rules

- Connect to the hostname that matches your access path. Mixing them will not work — the two resolve into different ingress.
- All endpoints require **TLS**, with public DigiCert wildcard certificates.
- DNS carries a short TTL. Resolve at connection time and **never bind to a static IP address**.
- **Production endpoints have not yet been published.** They will follow the same naming pattern.
- Your account manager confirms which endpoints are enabled for your firm.

Why market data uses two endpoints

Market data is split across two connections with enforced roles. The **incremental feed** (4500) carries sequenced live updates and you hold it open. The **snapshot connection** (4501) returns full book state and you use it only at startup and on recovery. A snapshot request on the incremental connection is rejected, and so are subscribe and unsubscribe on the snapshot connection.

Two reasons for the split. First, a full-book snapshot is large and arrives in chunks; on a shared connection it would head-of-line block the incremental updates queued behind it. Second, UDP multicast — the eventual transport — uses separate incremental and snapshot channels, so building against two channels now makes that migration a transport change rather than a rewrite of your market data logic.

Startup and recovery

1. Subscribe on 4500 -> START BUFFERING incremental updates
2. Request a snapshot on 4501

3. Receive the snapshot. It arrives in chunks - wait for all of them. It carries LastRptSeqProcessed: the last incremental included.
4. DISCARD buffered updates whose RptSeq <= LastRptSeqProcessed
5. APPLY the remainder in RptSeq order, on top of the snapshot
6. Book is live. The snapshot connection is idle until needed again.

STEADY STATE -> 4500 only
GAP / STALE BOOK -> return to step 2

Buffer from 4500 before you request the snapshot. Requesting the snapshot first leaves an unrecoverable hole between the snapshot anchor and your first buffered update. Full detail, including the message and enumeration reference, is in the SBE protocol reference.

Onboarding

What we need from you

1. Account identifier and the first five characters of your FIX API public key, per session
2. Source IP addresses for UAT and production, if you are reaching us over AWS London
3. How many sessions and what for — order entry, drop copy, market data
4. Expected message rates and instrument counts, so we can size your configuration
5. Your chosen connectivity route, plus Equinix entity details if you are taking a cross-connect
6. Technical and operational contacts

Sequence

1. Eligibility and entitlement confirmed with your account manager and your FCM
2. Network route provisioned; IP addresses allow-listed if you are coming over AWS London
3. UAT credentials and CompIDs issued
4. Integration and certification against UAT
5. Production credentials issued, connectivity and access confirmed pre-go-live
6. A short post-connection, no-trading window in production before live trading opens

Market makers reach the venue through their FCM. Run the clearing arrangement in parallel with the technical build — it gates production access independently.