Execution Server (Ember)

Event Processor / OMS / EMS

Execution Server (codename “Ember”)

USE CASES

- Platform for complex BUY-side trading algorithms (HFT)
- Platform for Execution Algorithms and Smart Order Routing
- SELL side Market Making
- Platform for Exchange Matching Engines
- FIX Gateway for market data/order entry flows

Functional Characteristics

- OMS / EMS
- Position Manager that supports custom projections (System-wide, Per-Account, Per-Trader, Per-Exchange, Per-Order Source, Per-Contract, Per-Currency, Per-Product Root). Projections can be composed (e.g. Per-Trader and Exchange)
- Risk Limits can be defined at any projection:
    - Order Price/Size Checks
    - Long/Short Worst-Case Position Size
    - Order Submit Rate
    - Order Reject Rate
    - Gross Trading Volume
    - Max number of open orders
    - Max Order Lifetime
    - Max Order ACK timeout
- APIs:
    - RPC API (TCP/UDP/IPC)
    - FIX 4.4 Protocol API
    - APIs to develop custom components:
        - Trading and Execution Algorithms
        - Risk Rules
        - Trading Connectors
        - Order Router
        - Message Transformers
- Normalized Data Model:
    - Market-By-Level or Market-By-Order price feed (Level 2 / Level 3)
    - Submit/Cancel/Replace Order Entry
        - Mass Cancel
        - Mass Status Request
        - Optional Cancel on Disconnect

Execution Server (codename “Ember”)

STREAM PROCESSING FRAMEWORK DESIGNED FOR TRADING DOMAIN

- Designed for high availability (Replicated State Machine)
- Write-ahead logging of trading messages
- Synchronous low-latency message replication
- New Market Feed and Trading Data models
- 100..1000x better performance than Deltix 4.3
- New concurrency model (wait & lock free, asynchronous)
- Current major version: 15+ production deployments
- Linux/Windows, Cloud

- Lock-free parallelism. In-memory state, replicated on cluster, no state sharing between threads
- CPU Affinity for critical threads
- Zero memory allocation in hot cycles (No Java GC pauses)
- Non-blocking IO
- Custom IEEE 754 compliant DECIMAL64 data type for prices and sizes (wider range than fixed point decimal, better precision than ‘double’).

Execution Server Building Blocks

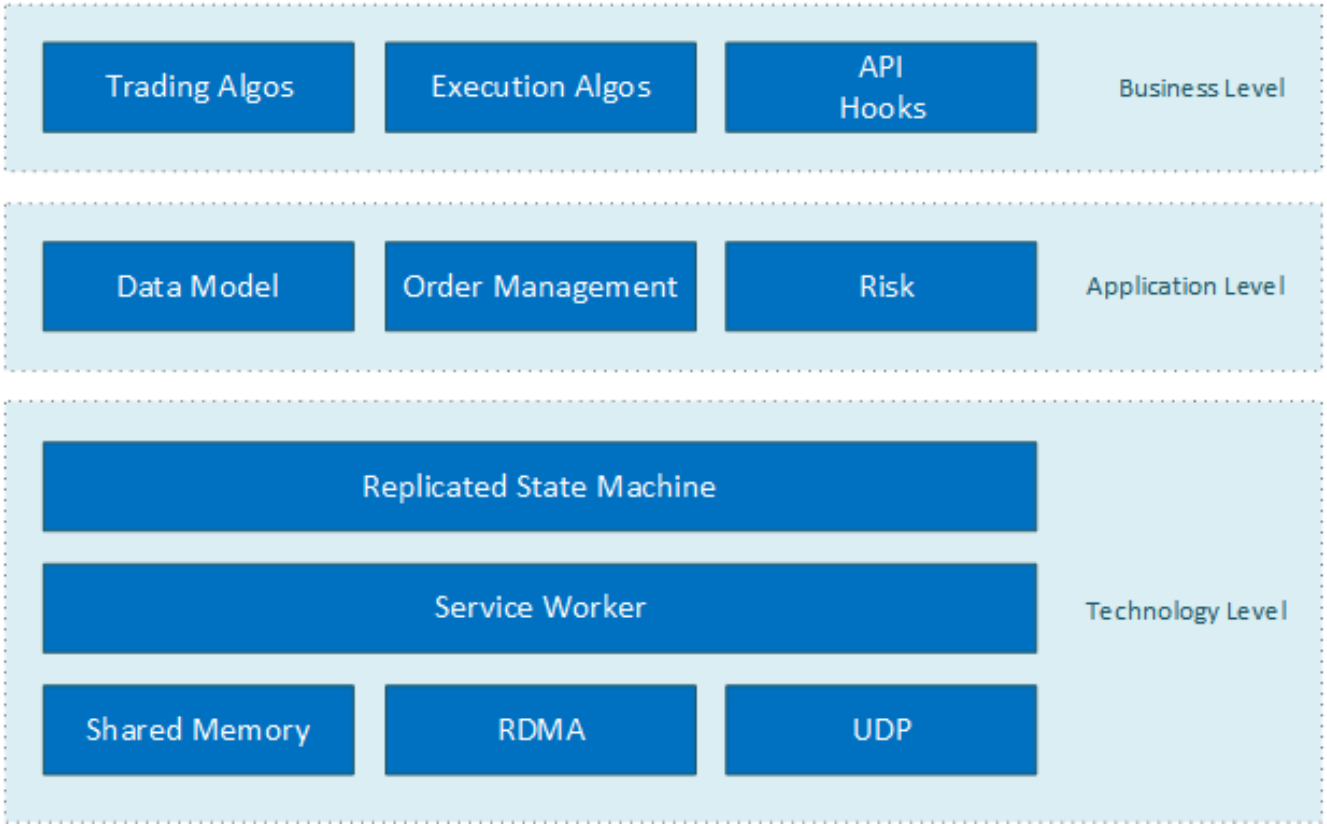

Execution Algorithms

- Built in Execution Algorithms:
  - TWAP
  - VWAP
  - PVOL
  - Iceberg
  - Smart Order Routing (SOR)
  - SOR with internal crossing
  - CLOB Matching Engine
- SDK to develop custom algorithms (Documentation, Samples, Test harness)
- Market Simulator
- Transaction Cost Analysis (TCA) Module

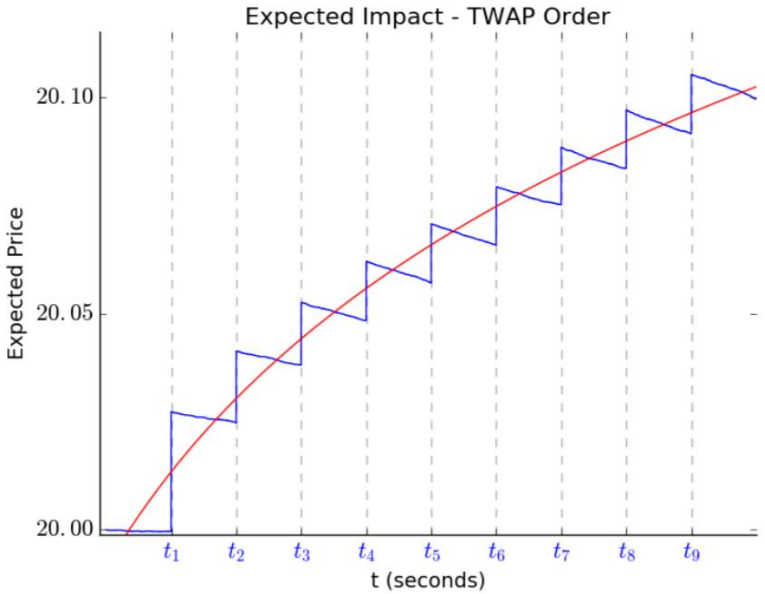

Ember – Standalone Mode

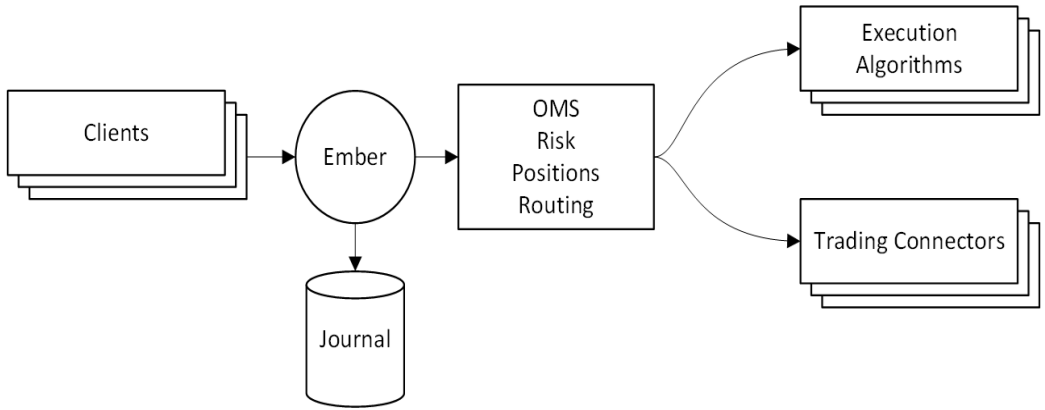

Ember – Active/Passive Mode

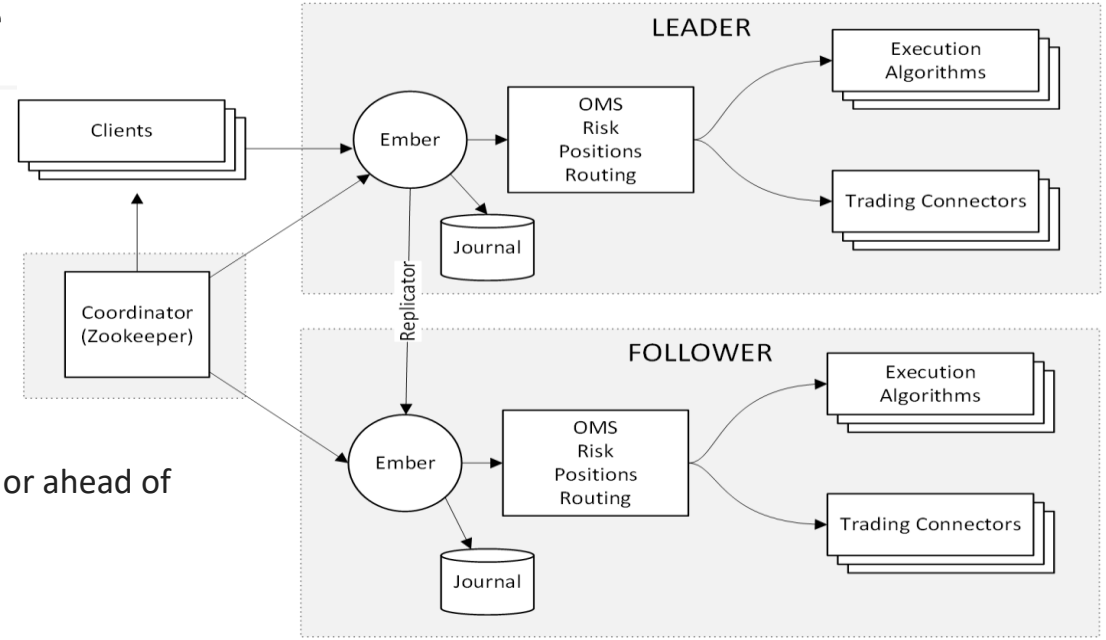

- Market data state – Follower May be behind or ahead of leader (delivered via TimeBase)
- Trade Orders state:
  - Includes all trading requests passed to execution venues
  - May lose some unacknowledged trading requests (lost in transmission during failover)
  - May not include last events from execution venues

Ember Architecture

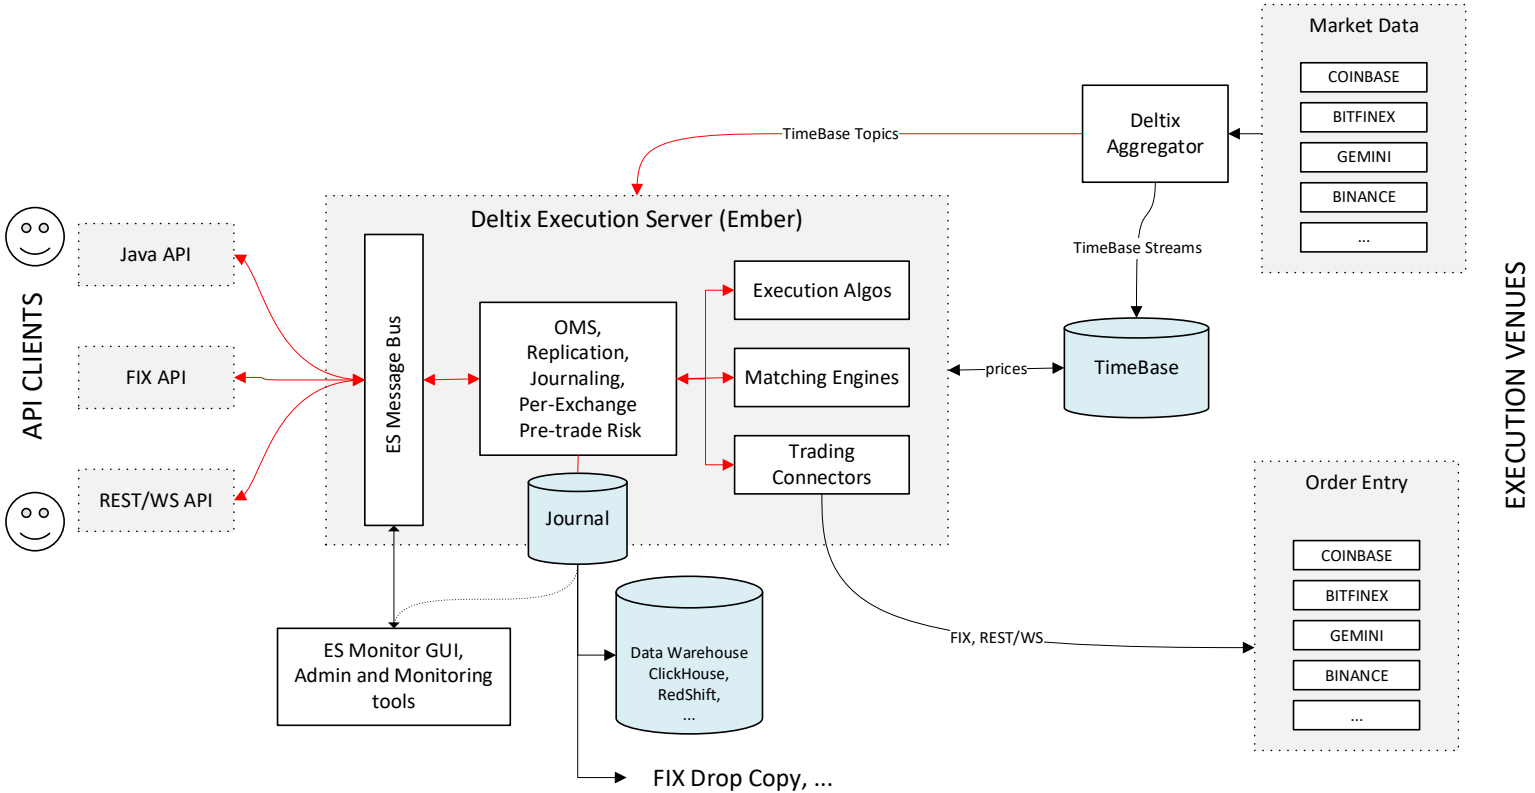

Ember Architecture

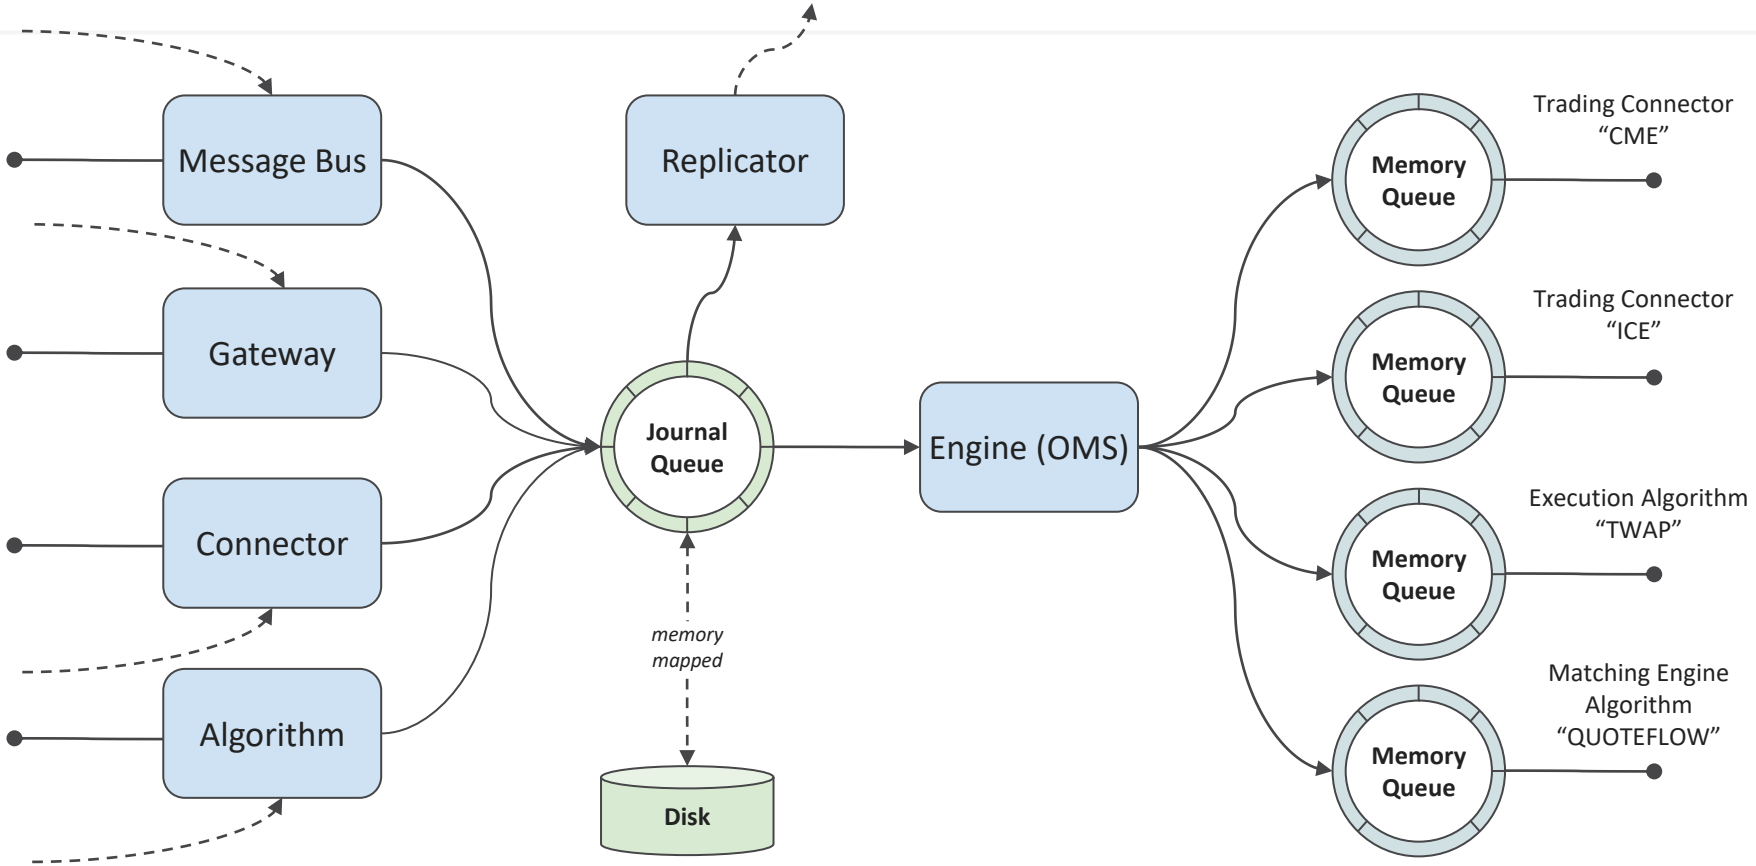

Journal Data Warehouse

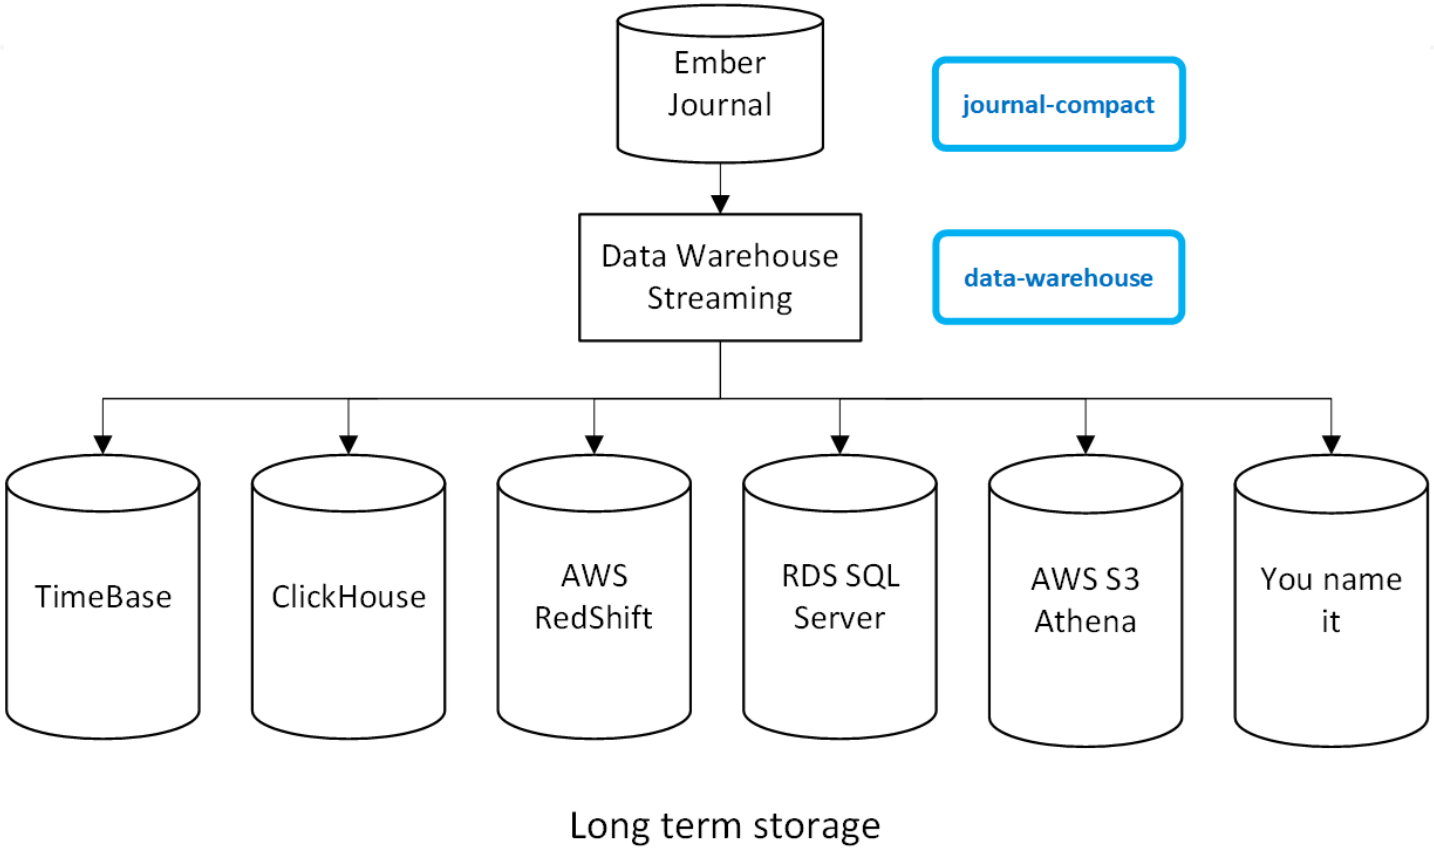

FIX Gateway to Matching Engine

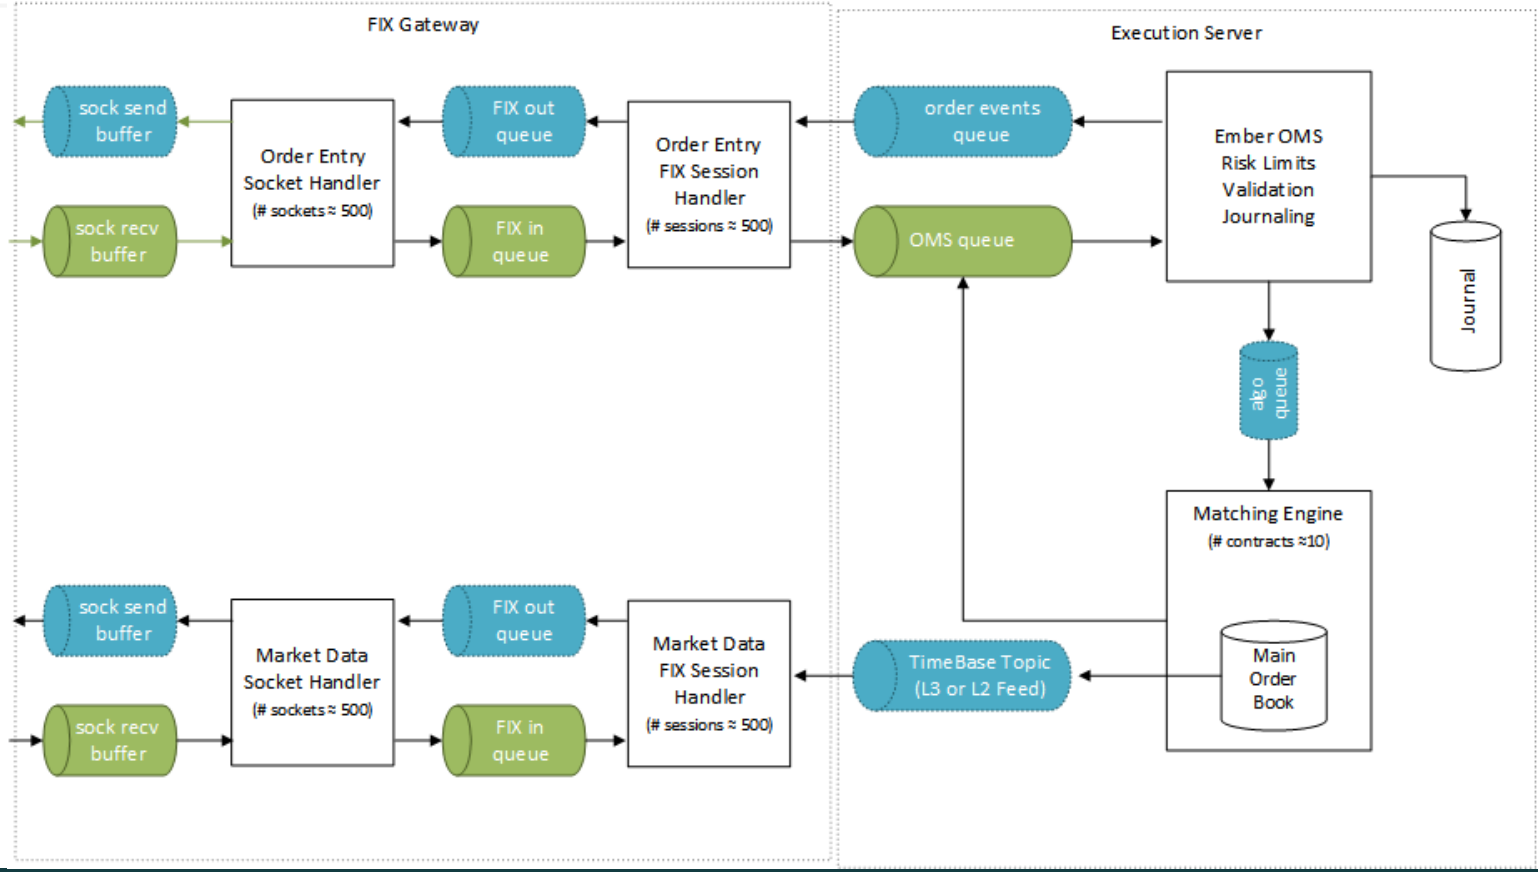

Security

- Development follows OWASP top 10 standard:
  - Spotbugs (with security plugin)
  - Dependency Check (scan of third-party dependencies against OWASP vulnerabilities database)
- SSL – data in transit encryption for FIX Gateway and Admin GUI
- HashiCorp Vault integration for secret management
- SSO (Ember Monitor GUI)
- UAC – user authentication and authorization (Ember Monitor GUI)
  - Users, Groups (LDAP/ActiveDirectory)
  - Permission Rules: { ALLOW/DENY, Principal, Operation, Resource }*

Performance - Summary

Throughput

- FIX API Gateway: 120 000 order request/sec
- RPC API Gateway: 210 000 order request/sec

Tick-to-Order

- FIX API: Network packet-to-packet latency:
  median: 9.5 microseconds, 99P: 25 microseconds
- RPC API: RTT latency (measured on local client)
  median: 5.1 us, 99P: 8.9 us

Execution Server latency (IPC clients)

RTT Latency (in nanoseconds) depending on Order Request rate (msgs/sec)

Percentile (%)	1K	10K	50K	100K	200K
0.000	4144	3836	3630	3616	3674
50.000	4783	4359	4191	4187	4211
90.000	7099	4679	4387	4395	4407
99.000	9295	7499	6151	5699	5795
99.900	12535	10463	10263	9735	11255
99.990	15191	13439	13479	16431	15607
99.999	26831	23199	31423	145023	32479
100.000	26831	45855	84735	194687	72511

Round Trip latency for co-located client (shared memory API).
AWS c5.9xlarge instances with CentOS 7.4
No CPU isolation

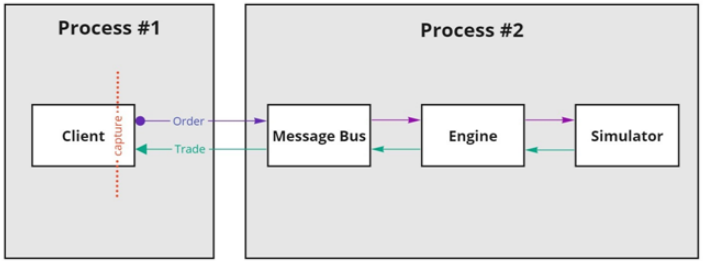

Execution Server latency (UDP clients)

RTT Latency (in nanoseconds) depending on Order Request rate

Percentile (%)	1K	10K	50K	100K	200K
0.000	47040	45632	46464	48896	53344
50.000	52127	50271	60223	76927	71231
90.000	56479	57951	66303	101503	90111
99.000	65343	63231	84031	127679	103743
99.900	95807	89023	100159	153727	134783
99.990	184703	182655	219135	249471	889343
99.999	421887	428543	517119	703999	5390335
100.000	421887	686591	847871	856063	5791743

Round Trip latency for closely-located client (UDP API).
AWS c5.9xlarge instances with CentOS 7.4
No CPU isolation

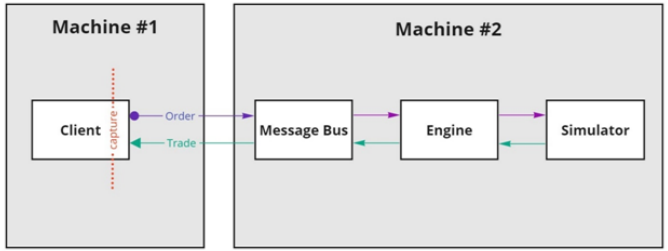

Tick-to-Order Latency

FIX Protocol:
Inbound market data message to outbound order

```
MIN        :         8 (microseconds)
50.0%      :        10
90.0%      :        12
99.0%      :        14
99.9%      :        16
99.99%     :        56
99.999%    :       108
99.9999%   :       277
99.99999%  :       433
99.999999%:       433
MAX        :       433
```

RPC IPC API: Inbound order to outbound order

```
0.000              3.168 (microseconds)
50.000             5.147
90.000             6.499
99.000             8.935
99.900            10.727
99.990            12.879
99.999            14.791
100.000           17.327
```

Network packet to packet latency (LIBPCAP)

Tick-to-Order RTT Latency

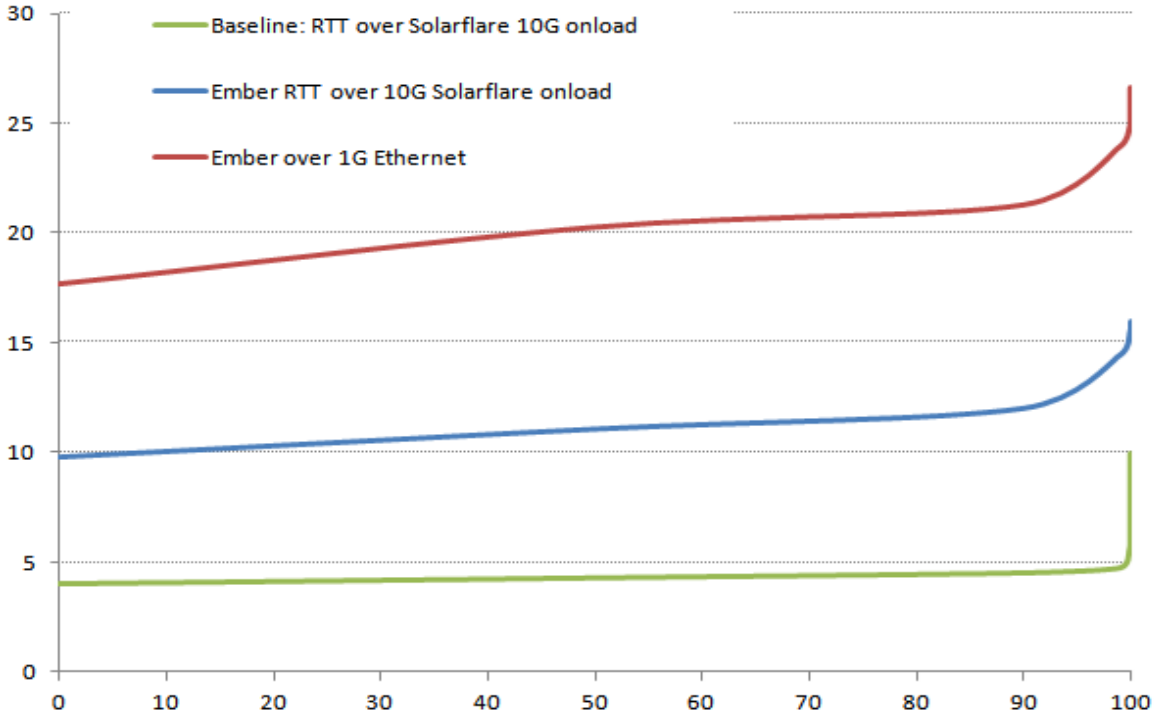

EPAM FIX Market Data Gateway – Performance Benchmarks

Number of FIX Gateways per server	Number of CPU cores utilized	Number of FIX Clients	Total throughput FIX messages/sec*	Order Book Depth	Network bandwidth used Gigabit/sec **
1	2	200	320 K	5	
2	4	200	520 K	5	
4	8	200	532 K	5	1.7
4	8	300	796 K	5	2.7
4	8	400	1064 K	5	3.7
4	8	300	1196 K	10	7.0

* No SSL encryption. FIX Gateway relies on third-party encryption solution. For example: STUNNEL or Load Balancer with SSL.

** Estimated network bandwidth, excluding TCP protocol overhead (headers, ACKs, etc).

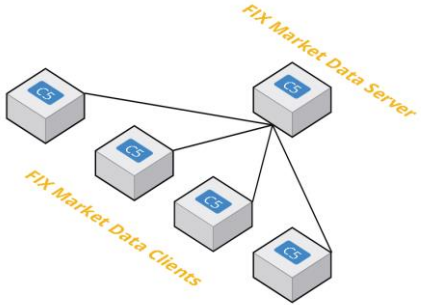

Server: 1 x AWS C5.9xlarge (18 cores, 26 threads) – some CPU cores were unused, this server spec guarantees 10G network bandwidth
Clients: 4 x AWS C5.4xlarge – some CPU cores were unused

Detailed AWS Instances Specs are here.
At its peak configuration was $4.25/hour

QUESTIONS?

Execution Server Monitor: Telemetry screen

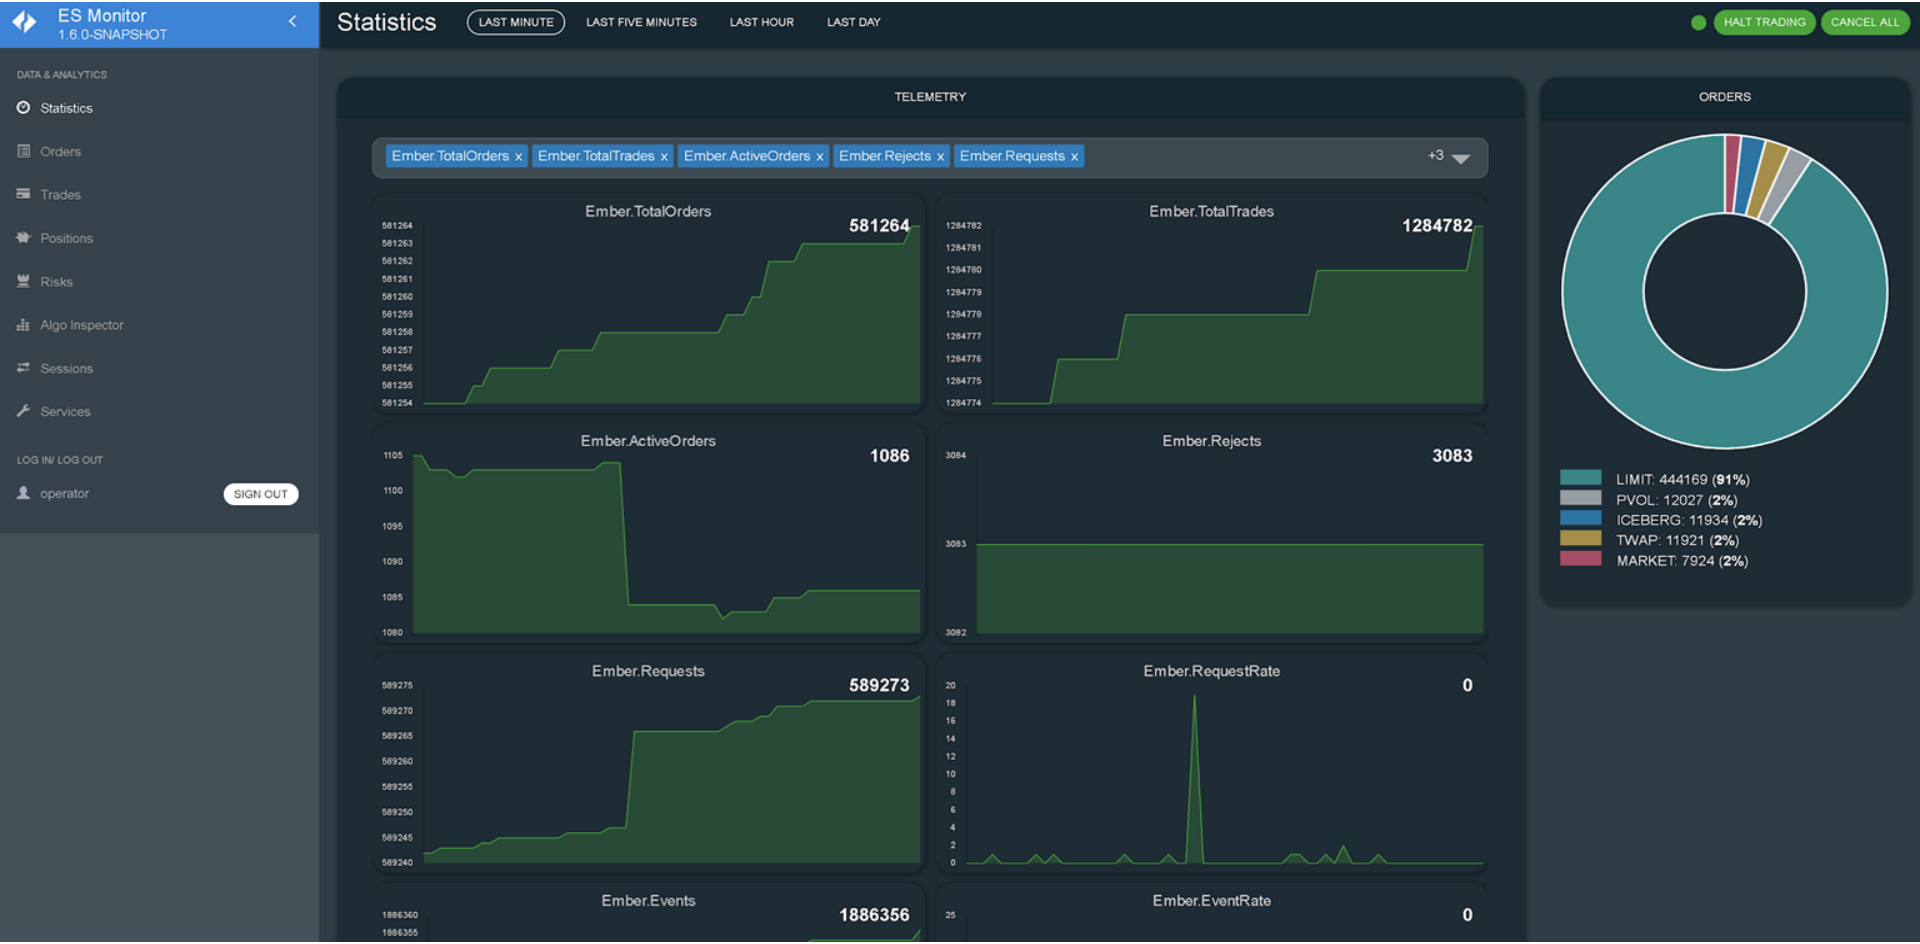

Execution Server Monitor: Orders blotter

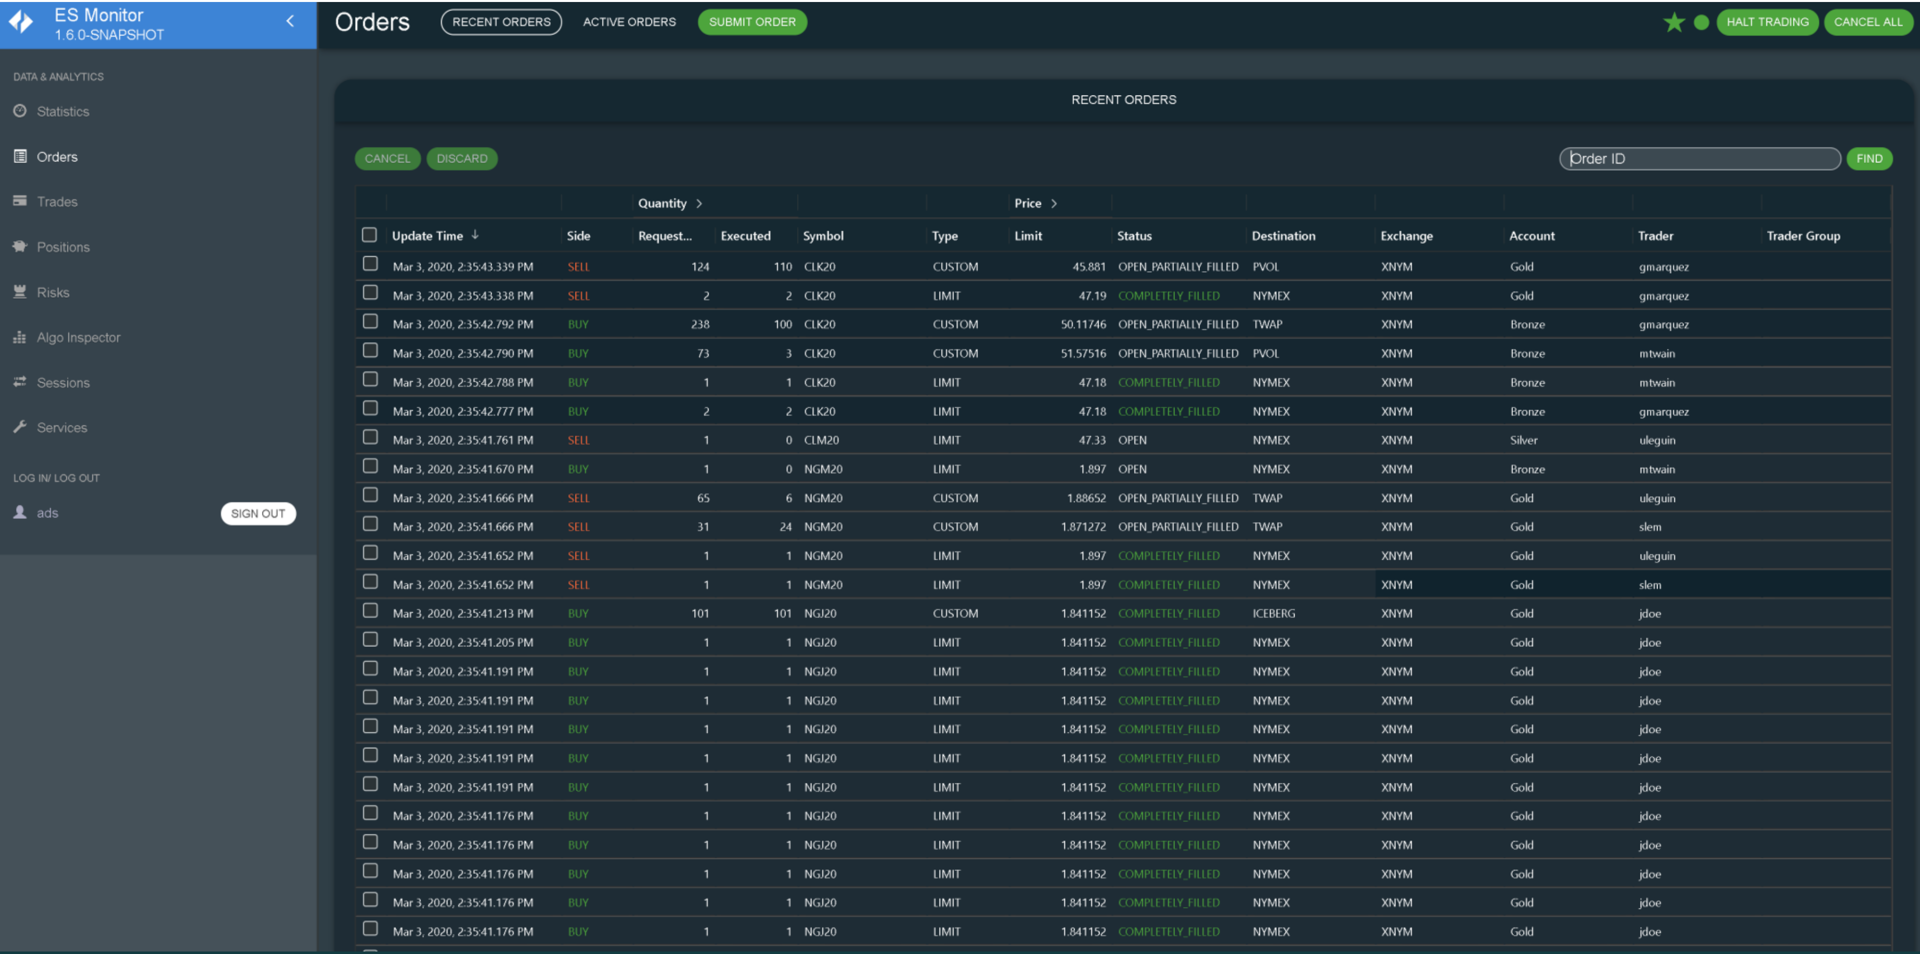

		Quantity				Price						
Update Time	Side	Request...	Executed	Symbol	Type	Limit	Status	Destination	Exchange	Account	Trader	Trader Group
Mar 3, 2020, 2:35:43.339 PM	SELL	124	110	CLK20	CUSTOM	45.881	OPEN_PARTIALLY_FILLED	PVOL	XNYM	Gold	gmarquez	
Mar 3, 2020, 2:35:43.338 PM	SELL	2	2	CLK20	LIMIT	47.19	COMPLETELY_FILLED	NYMEX	XNYM	Gold	gmarquez	
Mar 3, 2020, 2:35:42.792 PM	BUY	238	100	CLK20	CUSTOM	50.11746	OPEN_PARTIALLY_FILLED	TWAP	XNYM	Bronze	gmarquez	
Mar 3, 2020, 2:35:42.790 PM	BUY	73	3	CLK20	CUSTOM	51.57516	OPEN_PARTIALLY_FILLED	PVOL	XNYM	Bronze	mtwain	
Mar 3, 2020, 2:35:42.788 PM	BUY	1	1	CLK20	LIMIT	47.18	COMPLETELY_FILLED	NYMEX	XNYM	Bronze	mtwain	
Mar 3, 2020, 2:35:42.777 PM	BUY	2	2	CLK20	LIMIT	47.18	COMPLETELY_FILLED	NYMEX	XNYM	Bronze	gmarquez	
Mar 3, 2020, 2:35:41.761 PM	SELL	1	0	CLM20	LIMIT	47.33	OPEN	NYMEX	XNYM	Silver	uleguin	
Mar 3, 2020, 2:35:41.670 PM	BUY	1	0	NGM20	LIMIT	1.897	OPEN	NYMEX	XNYM	Bronze	mtwain	
Mar 3, 2020, 2:35:41.666 PM	SELL	65	6	NGM20	CUSTOM	1.88652	OPEN_PARTIALLY_FILLED	TWAP	XNYM	Gold	uleguin	
Mar 3, 2020, 2:35:41.666 PM	SELL	31	24	NGM20	CUSTOM	1.871272	OPEN_PARTIALLY_FILLED	TWAP	XNYM	Gold	slem	
Mar 3, 2020, 2:35:41.652 PM	SELL	1	1	NGM20	LIMIT	1.897	COMPLETELY_FILLED	NYMEX	XNYM	Gold	uleguin	
Mar 3, 2020, 2:35:41.652 PM	SELL	1	1	NGM20	LIMIT	1.897	COMPLETELY_FILLED	NYMEX	XNYM	Gold	slem	
Mar 3, 2020, 2:35:41.213 PM	BUY	101	101	NGJ20	CUSTOM	1.841152	COMPLETELY_FILLED	ICEBERG	XNYM	Gold	jdoe	
Mar 3, 2020, 2:35:41.205 PM	BUY	1	1	NGJ20	LIMIT	1.841152	COMPLETELY_FILLED	NYMEX	XNYM	Gold	jdoe	
Mar 3, 2020, 2:35:41.191 PM	BUY	1	1	NGJ20	LIMIT	1.841152	COMPLETELY_FILLED	NYMEX	XNYM	Gold	jdoe	
Mar 3, 2020, 2:35:41.191 PM	BUY	1	1	NGJ20	LIMIT	1.841152	COMPLETELY_FILLED	NYMEX	XNYM	Gold	jdoe	
Mar 3, 2020, 2:35:41.191 PM	BUY	1	1	NGJ20	LIMIT	1.841152	COMPLETELY_FILLED	NYMEX	XNYM	Gold	jdoe	
Mar 3, 2020, 2:35:41.191 PM	BUY	1	1	NGJ20	LIMIT	1.841152	COMPLETELY_FILLED	NYMEX	XNYM	Gold	jdoe	
Mar 3, 2020, 2:35:41.191 PM	BUY	1	1	NGJ20	LIMIT	1.841152	COMPLETELY_FILLED	NYMEX	XNYM	Gold	jdoe	
Mar 3, 2020, 2:35:41.176 PM	BUY	1	1	NGJ20	LIMIT	1.841152	COMPLETELY_FILLED	NYMEX	XNYM	Gold	jdoe	
Mar 3, 2020, 2:35:41.176 PM	BUY	1	1	NGJ20	LIMIT	1.841152	COMPLETELY_FILLED	NYMEX	XNYM	Gold	jdoe	
Mar 3, 2020, 2:35:41.176 PM	BUY	1	1	NGJ20	LIMIT	1.841152	COMPLETELY_FILLED	NYMEX	XNYM	Gold	jdoe	
Mar 3, 2020, 2:35:41.176 PM	BUY	1	1	NGJ20	LIMIT	1.841152	COMPLETELY_FILLED	NYMEX	XNYM	Gold	jdoe	
Mar 3, 2020, 2:35:41.176 PM	BUY	1	1	NGJ20	LIMIT	1.841152	COMPLETELY_FILLED	NYMEX	XNYM	Gold	jdoe	

Integration with Deltix (EPAM) AXA – Transaction Cost Analysis dashboard

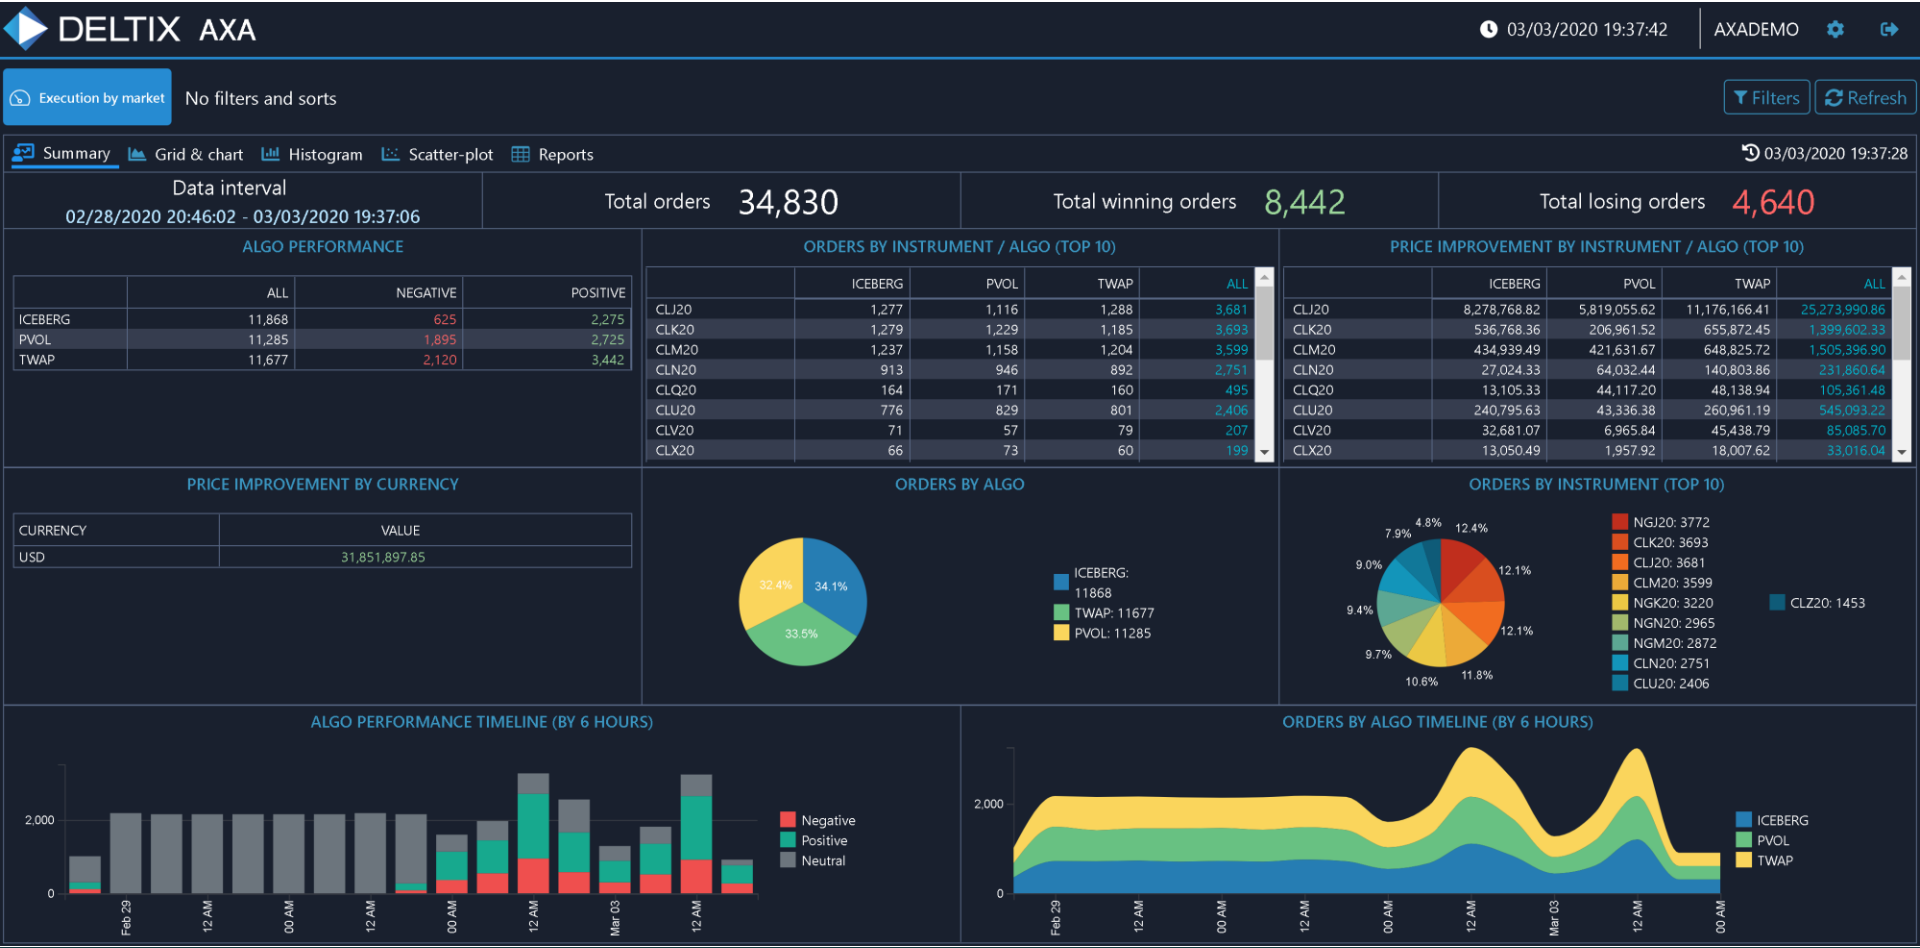

	ALL	NEGATIVE	POSITIVE
ICEBERG	11,868	625	2,275
PVOL	11,285	1,895	2,725
TWAP	11,677	2,120	3,442

	ICEBERG	PVOL	TWAP	ALL
CLJ20	1,277	1,116	1,288	3,681
CLK20	1,279	1,229	1,185	3,693
CLM20	1,237	1,158	1,204	3,599
CLN20	913	946	892	2,751
CLQ20	164	171	160	495
CLU20	776	829	801	2,406
CLV20	71	57	79	207
CLX20	66	73	60	199

	ICEBERG	PVOL	TWAP	ALL
CLJ20	8,278,768.82	5,819,055.62	11,176,166.41	25,273,990.86
CLK20	536,768.36	206,961.52	655,872.45	1,399,602.33
CLM20	434,939.49	421,631.67	648,825.72	1,505,396.90
CLN20	27,024.33	64,032.44	140,803.86	231,860.64
CLQ20	13,105.33	44,117.20	48,138.94	105,361.48
CLU20	240,795.63	43,336.38	260,961.19	545,093.22
CLV20	32,681.07	6,965.84	45,438.79	85,085.70
CLX20	13,050.49	1,957.92	18,007.62	33,016.04

CURRENCY	VALUE
USD	31,851,897.85

User Story –Market Maker in cryptocurrencies market

- [redacted] uses Market Maker solution built on top of Ember
- Ember OMS keeps ~ 500 active orders on several markets and processes million+ of trading orders per hour in 24/7 mode.

  FIX Logs alone accumulate 4+ GB per day.

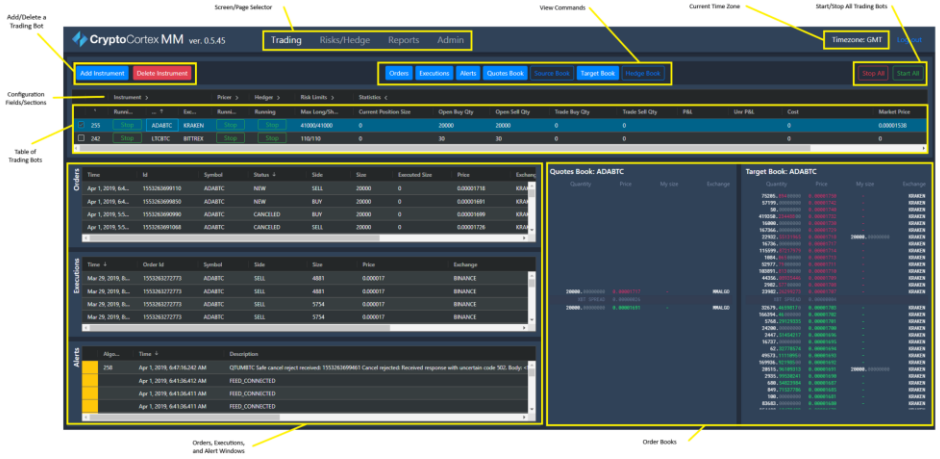

Ember inside CryptoCortex

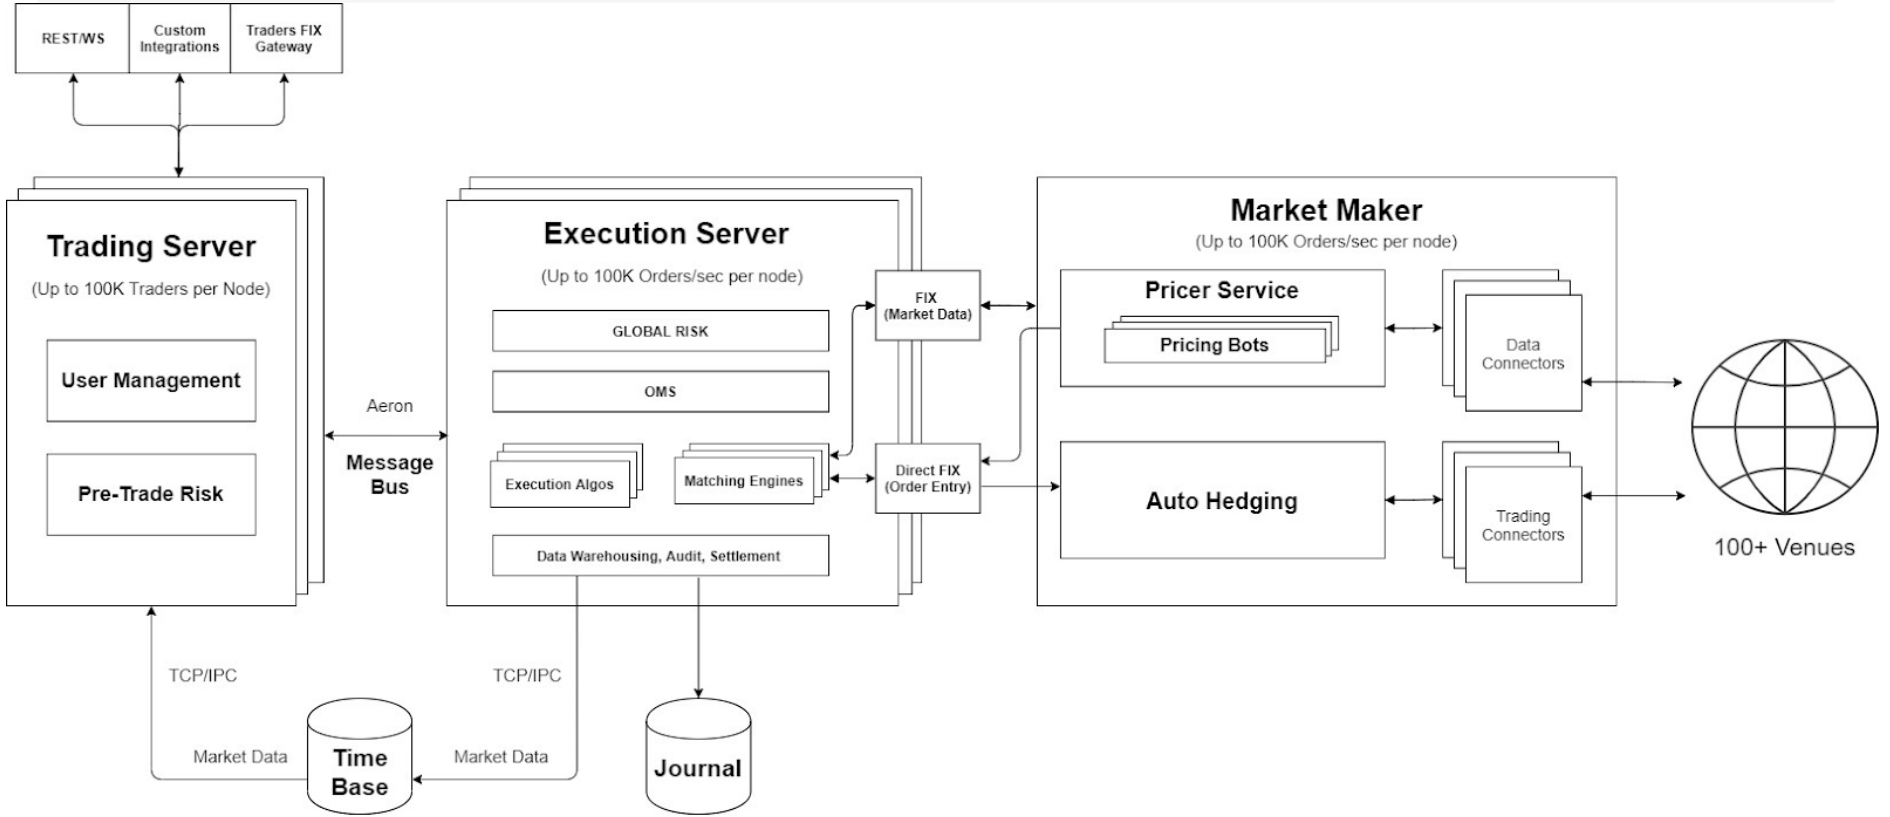

Stream Processing Latency

Trading Domain: Stream Properties and trade-offs

Timebase

Volume: Very high (2M+ msgs/sec per node)
Delivery: At most once, delayed data = junk, warm up mode
State: discardable
Source can replay events: Yes, but this is useless for RT applications

Ember

Volume: Moderate (up to 250K msgs/sec)
Delivery: Exactly once, Asynchronous
Guaranteed for ACKed orders
State: must be recoverable
Source can replay events: No! (downstream systems are not idempotent)

Ember Core – Spring 2017

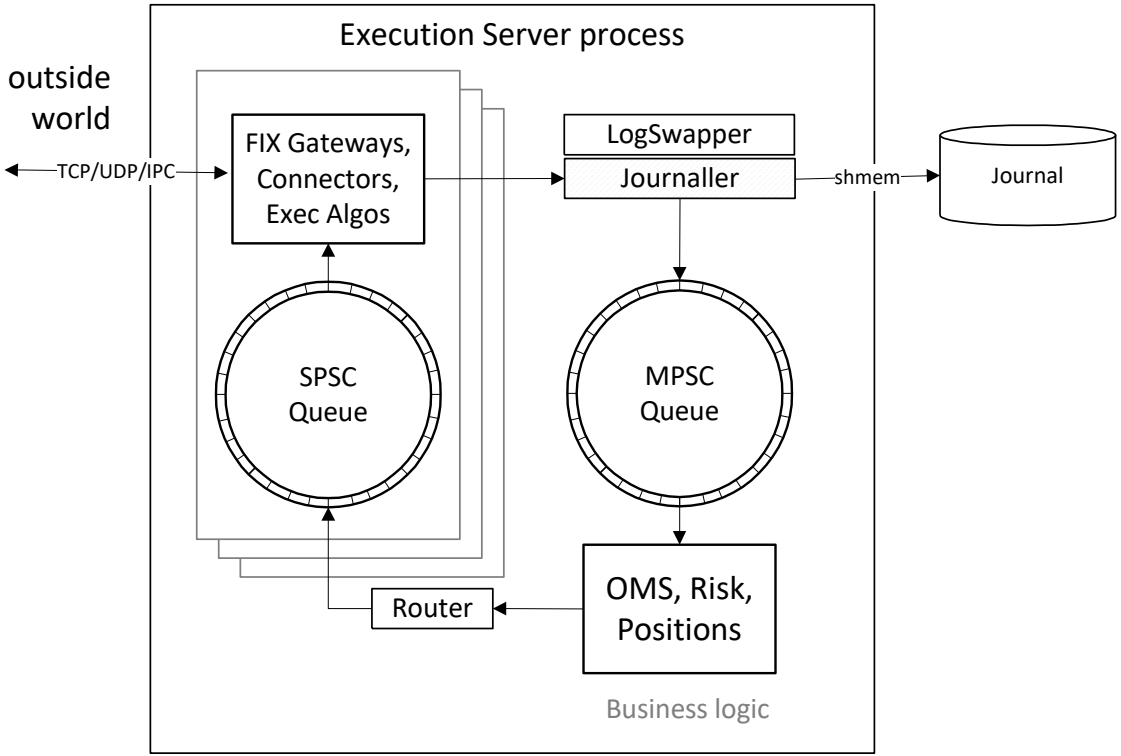

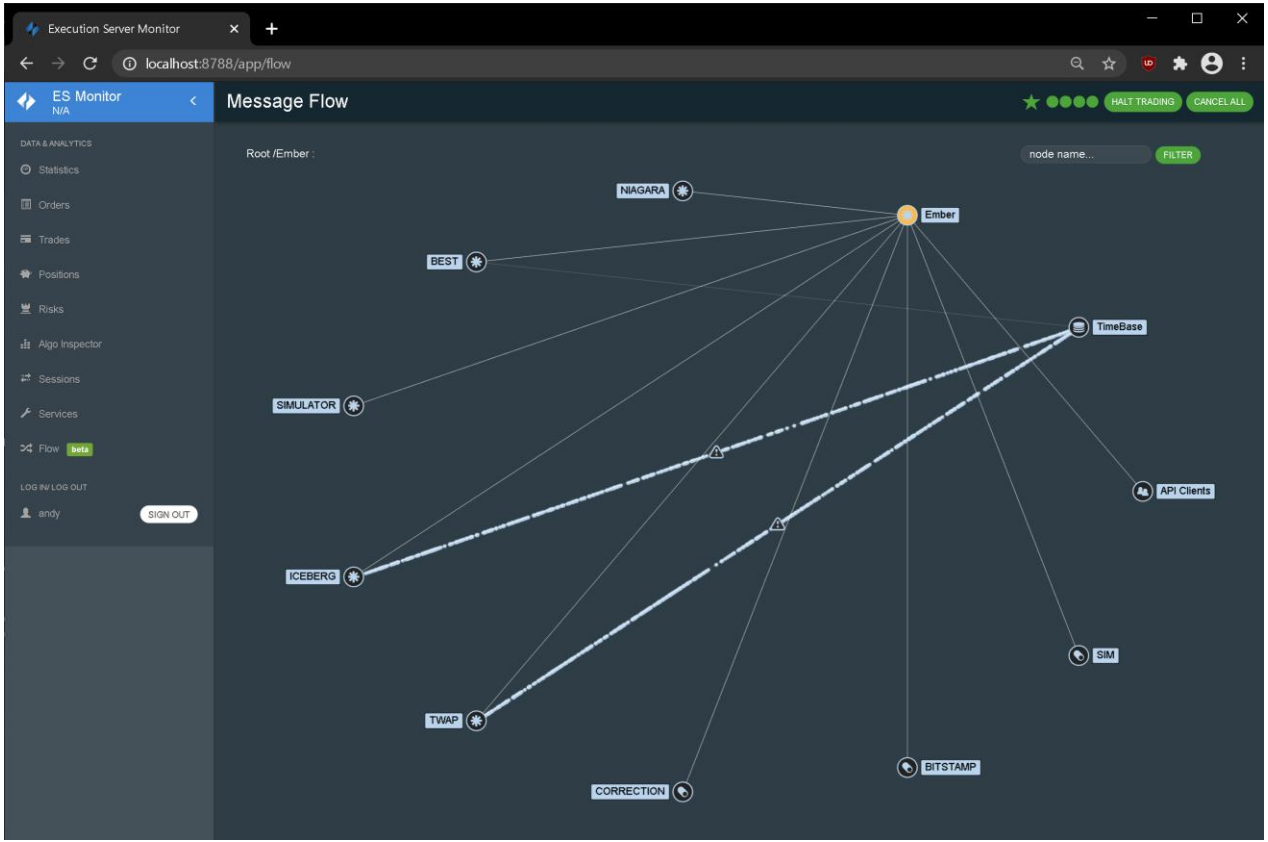
Execution Server Monitor
localhost:8788/app/flow
ES Monitor
N/A
DATA & ANALYTICS
Statistics
Orders
Trades
Positions
Risks
Algo Inspector
Sessions
Services
Flow
beta
LOG IN/ LOG OUT
andy
SIGN OUT
Message Flow
HALT TRADING
CANCEL ALL
Root /Ember :
node name...
FILTER
NIAGARA
Ember
BEST
TimeBase
SIMULATOR
API Clients
ICEBERG
SIM
TWAP
BITSTAMP
CORRECTION