BROADPEAK PARTNERS, INC.

High Volume Scaling With K3

Client: Hedge Fund

Issue: Establishment of a High Volume Array on K3

Client Custom Code Required: None

Overview

A Hedge Fund client trades approximately 8000 trades per day across ICE and NYMEX. Due to the initiation of algorithmic trading and other initiatives it is possible that trade volume may increase three fold or more. BroadPeak was asked to deploy K3 to meet the following requirements:

1. Maintain high volume throughput in K3 from exchange fix message receipt through persistence and loading into downstream applications.
2. Ensure high volume throughput even though load time to downstream systems may be as slow as one trade leg per second.

K3 - Architecture Background

The fundamental architecture is designed to scale with large operations.

K3 is actually designed as **two independent server operations**: "K3 Server" and "K3 Ref Server."

The fundamental architecture is as follows:

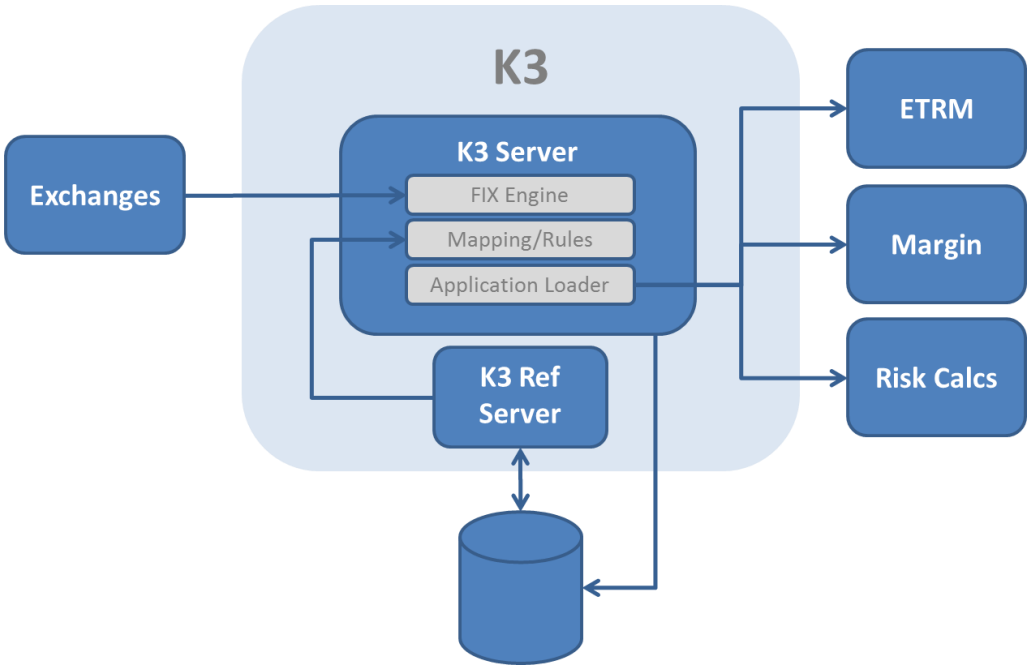

K3 Server is the server that is designated as the primary vehicle for receiving, onboarding and viewing messages. K3 Trade server is the portion of the application where the FIX engines (FIX Message Connections), Feed Manager (viewing trades) reside.

K3 Ref Server is the server that is designed to handle operations such as mapping, rules, and feeds to downstream systems.

In other words the operation of K3 as a whole is bifurcated into two servers: one to manage the inbound trade flow. A second server to handle Mapping, Rules and Loading to downstream trading systems. Thus in terms of a discussion of scaling there are two primary elements:

1. Management of Inbound Trade Flow
2. Management of K3 Throughput and Loading to downstream applications.

Management of Inbound Trade Flow

There are many challenges to opening the gates to receive messages of executed trades and these challenges vary by exchange.

For example, the ICE connection is a "push" connection, meaning trades are pushed from ICE to K3. This push is performed over a TCP connection with typical bandwidth capacity of 100MB/s. While utilization of the full 100MB/s connection would result in a total potential constant trade flow of approximately 24,000 trades per second, the connection requires encryption (HTTPs) which naturally degenerates performance of inbound trades. We estimate that the maximum possible inbound capacity is about 15,000 trades per second, assuming that there is no degeneration in bandwidth.

CME, on the other hand, has a different approach. CME CT API is a "Poll" model where trades are polled. CME does not allow market participants to poll more than once every three seconds. This feed is also encrypted using HTTPs. So, assuming the same 100MB/s bandwidth the maximum amount of trade legs received by CME is 5,000 trades per second.

The logical solution to this is to source higher and higher bandwidth to accommodate greater trade input. On this point K3 has throttle max capacity of 130,000 trades per second. We think that due to existing infrastructure at ICE and NYMEX it is unlikely that the exchange will be able to send more than this amount of messages of executed trades. (This is different than order execution which has a different structural makeup.)

Simple Scaling for High Performance | "Scaling Up"

Scaling Up refers to K3's architectural ability to natively scale upward and process more messages.

K3 as a whole is designed from the ground up with multi-threading. In other words, the more processing power the resident server has, the faster K3 will operate. For example, if K3 is operating on a 16 Core Server, all 16 cores are utilized to map, rule and load the messages into the target application.

The greatest challenge to throughput is load time to the downstream application. Load Time is entirely beyond the reach of K3 architectural design, and is entirely dependent on how well the target application can accommodate incoming messages/trades. In many cases this can be as slow as one trade leg per second. K3 approaches this problem from two fronts:

1. Additional RAM memory on the server will allow longer queues of data waiting to load into the downstream system. As a rule of thumb, each 5GB of data accommodates approximately 1 Million Trades (thus 24 GB of available memory accommodates roughly 5 Million trade legs.
2. Each Application loader utilizes our Failover component such that even if the connection is disrupted it will start up again where it left off.

Input to K3 from the exchange generally throttles (max capacity) at approximately 130,000 trade legs per second. Achieving higher volumes is entirely bandwidth and exchange proximity related. The higher levels of bandwidth and closer to the exchange the FIX engine is situated the higher the inbound capacity.

In short Simple Scaling Up only requires use of a larger server located proximate to the exchange. For example a 16 Core 24 GB server will be able to handle very high volumes.

Very High Throughput Scaling | "Scaling Out"

Very High Throughput Scaling is achieved by running "K3 Server" in Master and Slave mode; scalability is achieved by adding more slaves. This is a different configuration than described above, but requires no software modifications. It also incorporates the use of the client's choice of messaging technology.

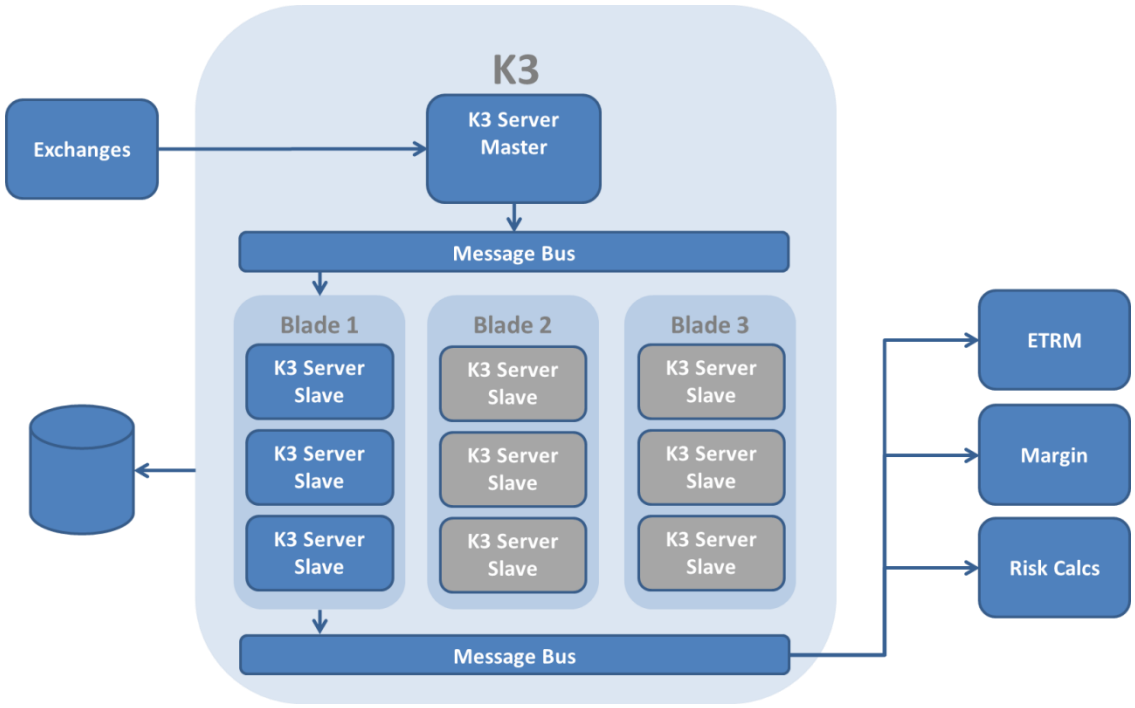

Here, the K3 Server "Master" receives all inbound trades. As with simple scaling the amount of available bandwidth and proximity to the Exchange are key to increasing the speed with which inbound trades are captured.

K3 Server uses its internal routing to write the FIX Messages to a standard queue. K3 is agnostic toward queues and any common queue is sufficient. These include, but are not limited to, TIBCO, Solace, Active MQ, and WebSphere. Writing to a queue here allows the K3 Server Master to easily feed multiple Slaves. Here the Blue K3 Server "slaves" are those that are live and running on a single Blade Server. The Gray Slaves are potential expansion slots. In other words, K3 can be run on as many servers as necessary to accommodate trade throughput.

Finally, the results from each of the slave K3 Servers are written to a second queue where they can live indefinitely while waiting to be loaded into the downstream applications.

This structure ensures that every message is processed in an expedited manner regardless of how trades are sourced from the exchange.

BROADPEAK PARTNERS, INC.

Conclusion

Either approach is entirely practical for K3. We feel that there is not a limitation from our end in terms of the throughput capacity of K3. Architecturally, we are able to receive trades and process them through feed manager, mapping, rules and other functions consistent with the volume of trades executed by the firm.

Need to trade high volumes with efficient infrastructure? Give us a call.

Phone: +1.646.370.6804

Email: info@broadpeakpartners.com