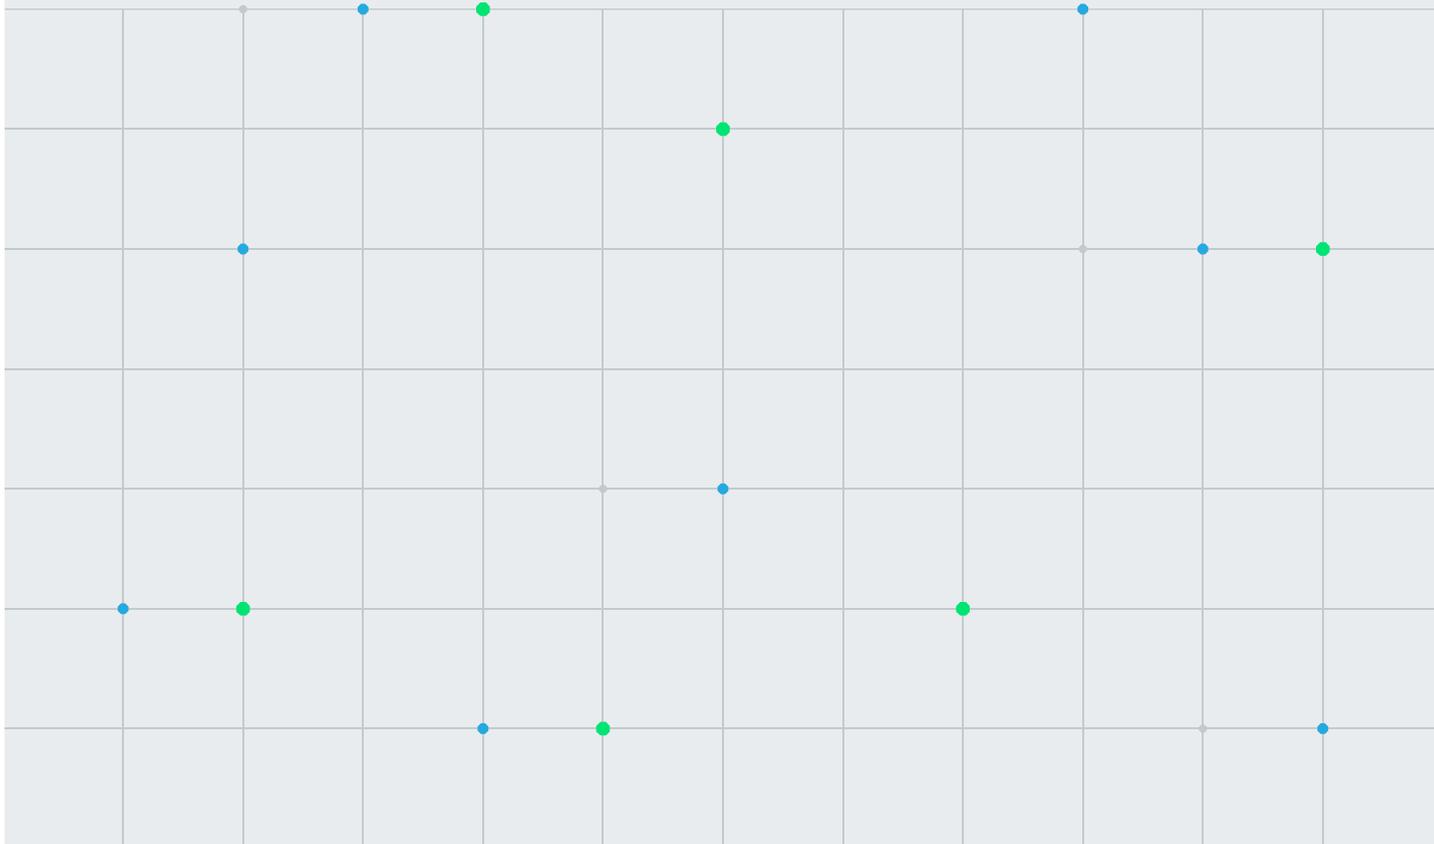


Supporting FinTech Startups via Cloud Distribution Solutions

How skew went from scratch to live trading application in a few weeks using CME Smart Stream on Google Cloud Platform solution.



Introduction

skew is a best-in-class provider of data analytics and trade execution services for cryptocurrency markets. These services include integration with CME Group services, such as CME Clearport and MDP (Market Data Platform) feeds.

skew was founded in 2018 and, like most fintech startups, builds and hosts its infrastructure in the cloud. Streaming market data from marketplaces, such as CME Group, would normally require implementing various low-level transport protocols and setting up expensive leased lines or co-location. This approach is not effective for cloud-native companies who are sensitive to costs, lead time and complexity before even being able to process market data messages.

CME Smart Stream on Google Cloud Platform (GCP) allows skew to easily stream native FIXBinary messages using standard cloud native APIs. This meant that skew could easily access Market Data from its cloud provider with very little additional setup and transport overheads. Using GCP, each FIXBinary message comes with a delivery guarantee meaning that complexity is further reduced instead of managing packet drops from standard UDP protocol.

The solution additionally provides great economical flexibility due to its on demand hourly usage-based fee model.

“Traditional” approach challenges

Traditional approaches to integrate with exchange market data feeds normally include managing data over two UDP streams and a TCP replay endpoint. Complex logic is required to switch between the two UDP feeds to ensure no data packets are dropped and any snapshot or message replay happens over the TCP transport (reference: <https://www.cmegroup.com/confluence/display/EPICSANDBOX/MDP+3.0++Dissemination>).

A custom implementation of such low-level protocols requires certification against the exchange to ensure the client software can correctly handle various scenarios. To reduce development time, pre-certified software can be purchased which manages all the low level transport complexities (see <https://www.onixs.biz/cme-mdp-3-0-market-data-handler.html>). In addition to the transport protocol complexities, access to these endpoints would normally include setting up cross-connection or even co-location in the exchange’s data center which incur a large lead time and cost.

Following the native connectivity approach provides the least latency as the client software interacts with the low-level transport but, as described above, comes with a connectivity and implementation cost. Small companies and startups, such as skew, may not have the technical resources to pursue this route. Additionally, many clients do not have single digit millisecond latency-sensitive requirements.

Cloud-based solution advantages and implementation

CME Smart Stream on Google Cloud provides a cloud native solution to access CME Group market data feeds without having to implement any low-level transport protocols or setting up dedicated connectivity. Using the standard GCP SDK, available in a variety of popular programming languages, software can be developed which can easily access the data from any client location over a single transport layer.

Once the initial integration with Google Cloud PubSub is complete, it becomes easy to subscribe to additional topics containing additional CME Group product data without the need to manage additional transport layer interactions. This puts skew in a position to quickly release new products in response to customer demand.

Large scale market data recovery is also simplified, since it is just another Cloud PubSub subscription to connect to, with the same delivery guarantees.

Certification

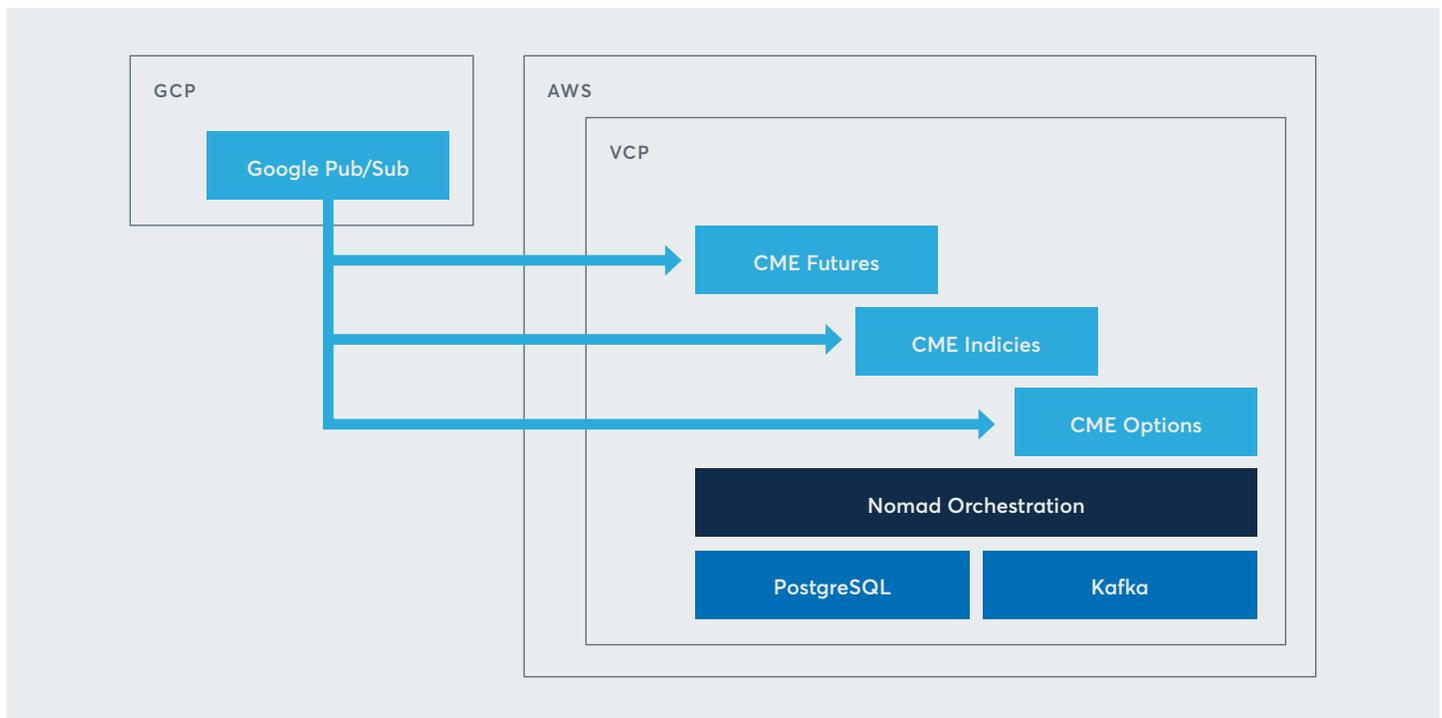
Before processing CME Group market data, client systems must certify with Autocert+. CME Group provides Cloud PubSub topics to accomplish this. Using these certification topics, and thanks to the integration of CME's Autocert+ with Smart Stream, skew was able to certify with CME Group in a very short period of time.

To bootstrap the development process, CME Group provides a sample project which consumes data from the topics and ensures in-order processing. The project can be found at <https://github.com/CMEGroup/CMESmartStreamonGCP-BinaryData>

Multi-cloud, networking and architecture

The service is also easily accessible cross-cloud. skew hosts AWS production and staging environments in Ohio with Pub/Sub latency of 13ms at the 95% level with a median of 2ms.

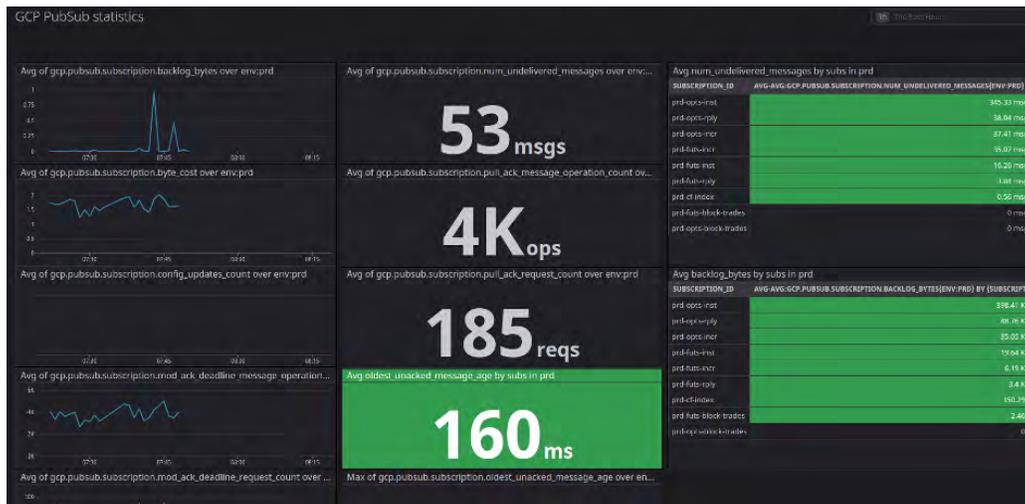
There is no requirement however to use AWS and, indeed, the skew stack has been run on GCP as a proof of concept. There are plans additionally to link resources in GCP and AWS using Consul Mesh Gateways.



Subscriptions, monitoring and costs

skew uses 41 subscriptions to nine unique topics across its development, staging and production environments. These can be configured via the GCP Console, the gcloud SDK or using Terraform. In July 2020, skew's subscriptions made 500 million requests at a cost of \$500.

Google Pub/Sub and, more generally GCP, can be monitored via a Datadog integration. skew uses Datadog to monitor its AWS and GCP infrastructure including its GCP subscriptions across different environments. In addition, monitoring alerts are triggered when configured thresholds are breached.



Lessons learned

Google PubSub provides guaranteed delivery but does not guarantee in-order delivery. skew requires in-order processing of CME Group market data in order to show clients correct and up-to-date pricing. The lack of in-order delivery from PubSub added some complexity to the client design, especially given the high volume of messages produced by CME Group. We would have faced a similar challenge using traditional solutions. Traditional UDP does not guarantee deliver or ordering.

This was slightly compounded by CME Group's certification environment not matching the throughput of its production environment. CME Group's certification PubSub topics have a throughput of a few hundred messages-per-second, whereas in production, throughput is closer to 15,000 messages-per-second for a single topic. skew's client systems had to be reworked to handle the increased throughput while also ensuring in-order processing of messages. Again, we would have experienced this with traditional solutions, but the cloud eliminated the potential of having to order and provision additional hardware once we optimized our code.

About skew

Founded in 2018 by two derivatives traders in London, skew provides leading data analytics and trade execution services for cryptocurrency markets.

skew's platform offers the most comprehensive real-time analytics for spot, futures and options on cryptocurrencies as well as access to block trade liquidity of CME-listed bitcoin derivatives contracts in regulatory compliant format.

For more information about skew please visit www.skew.com

Contacts

Daniel Cullender

Chief Technology Officer
daniel.cullender@skew.com

Tim Short

Head of Devops Engineering
tim.short@skew.com

James Asefa

Back-End Engineer
james.asefa@skew.com

About CME Group

For more information about CME Smart Stream on Google Cloud, please visit: cmegroup.com/cloud

As the world's leading and most diverse derivatives marketplace, CME Group (www.cmegroup.com) enables clients to trade futures, options, cash and OTC markets, optimize portfolios, and analyze data – empowering market participants worldwide to efficiently manage risk and capture opportunities. CME Group exchanges offer the widest range of global benchmark products across all major asset classes based on [interest rates](#), [equity indexes](#), [foreign exchange](#), [energy](#), [agricultural products](#) and [metals](#). The company offers futures and options on futures trading through the [CME Globex®](#) platform, fixed income trading via BrokerTec and foreign exchange trading on the EBS platform. In addition, it operates one of the world's leading central counterparty clearing providers, CME Clearing. With a range of pre- and post-trade products and services underpinning the entire lifecycle of a trade, CME Group also offers optimization and reconciliation services through TriOptima, and trade processing services through Traiana.

CME Group, the Globe logo, CME, Chicago Mercantile Exchange, Globex, and, E-mini are trademarks of Chicago Mercantile Exchange Inc. CBOT and Chicago Board of Trade are trademarks of Board of Trade of the City of Chicago, Inc. NYMEX, New York Mercantile Exchange and ClearPort are trademarks of New York Mercantile Exchange, Inc. COMEX is a trademark of Commodity Exchange, Inc. BrokerTec, EBS, TriOptima, and Traiana are trademarks of BrokerTec Europe LTD, EBS Group LTD, TriOptima AB, and Traiana, Inc., respectively. Dow Jones, Dow Jones Industrial Average, S&P 500 and S&P are service and/or trademarks of Dow Jones Trademark Holdings LLC, Standard & Poor's Financial Services LLC and S&P/Dow Jones Indices LLC, as the case may be, and have been licensed for use by Chicago Mercantile Exchange Inc. All other trademarks are the property of their respective owners.

