



FINANCIAL SERVICES CLOUD SYMPOSIUM | 2022

How JPMC modernized its hedging and risk management platform

Paul Bauerschmidt (he/him)

Managing Director
JPMorgan Chase

Agenda

- About JPMorgan Chase
- The history of Athena platform
- Cloud at JPMorgan Chase and why it is important
- Athena ML Hedging & Risk Management framework on AWS
- Challenges, accomplishments, and key takeaways

About JPMorgan Chase

Delivering at a global scale

\$33T

in assets under custody

450+

petabytes of data

250K

virtual workspaces

59M

active digital customers

50%

or 66 million US households have a relationship with Chase

>80%

of Fortune 500 companies do business with JPMC

>100

presence in markets globally

\$9T

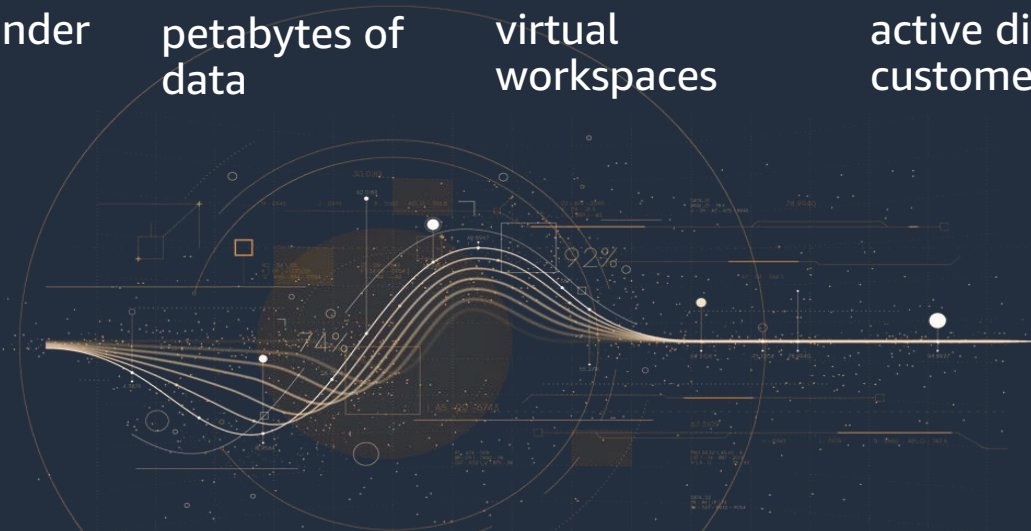
in payments processed daily

50K+

technologists

6.5K+

applications



What is JPMorgan's Athena?

Corporate & Investment Bank

#1 in markets revenue	100 markets globally
\$29.5B 2020 markets revenue	

Athena Ecosystem

4K+ Python developers	50M lines of code
20K production changes/week	One of the world's largest Python repositories

Athena Scale

5 Billion calculations per day	
10K batch jobs per day	50K Compute Grid

Athena Platform

Pricing

Trade Management

Risk

Data Science

Models

Frameworks

Data


Environment

Dev Tools



Athena hedging and risk management

Maturity Curve for Hedging Service



Key Performance Factors

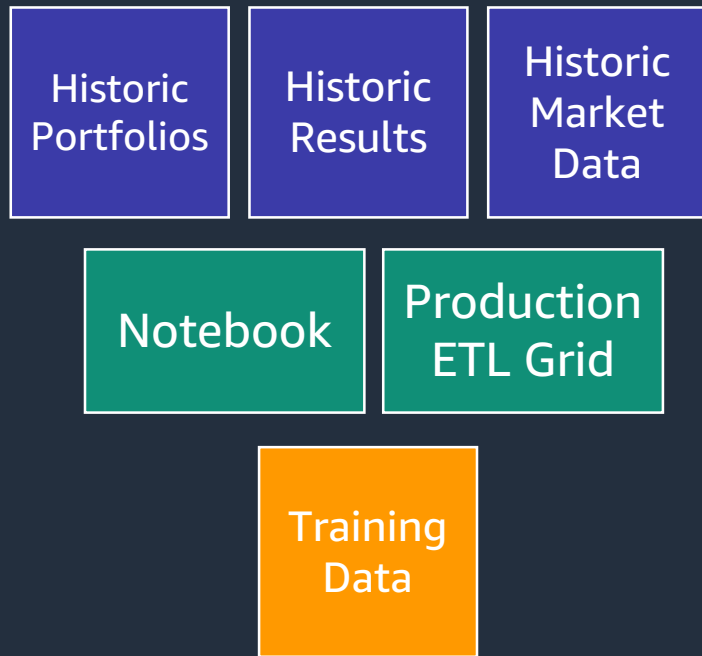
- Staff costs
- Time savings in market
- Operational risk and reliability
- Audit and traceability
- PnL!

Cloud design principles

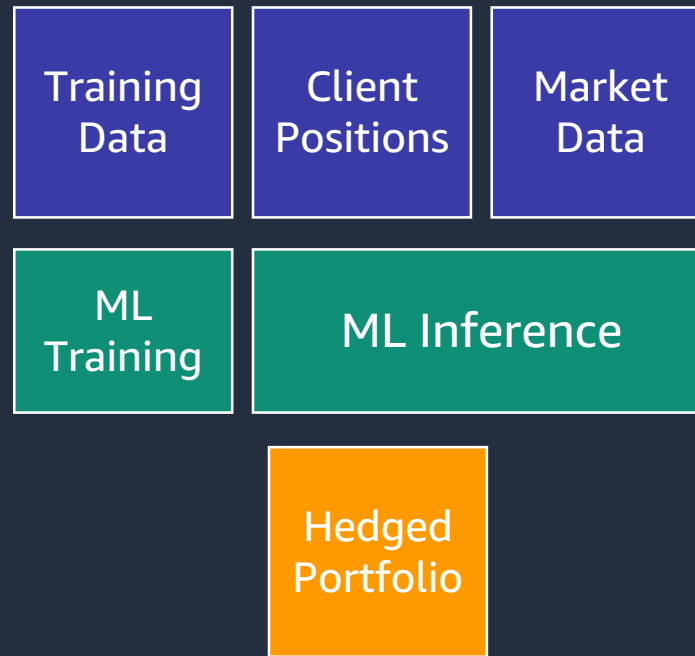


Athena ML hedging and risk management

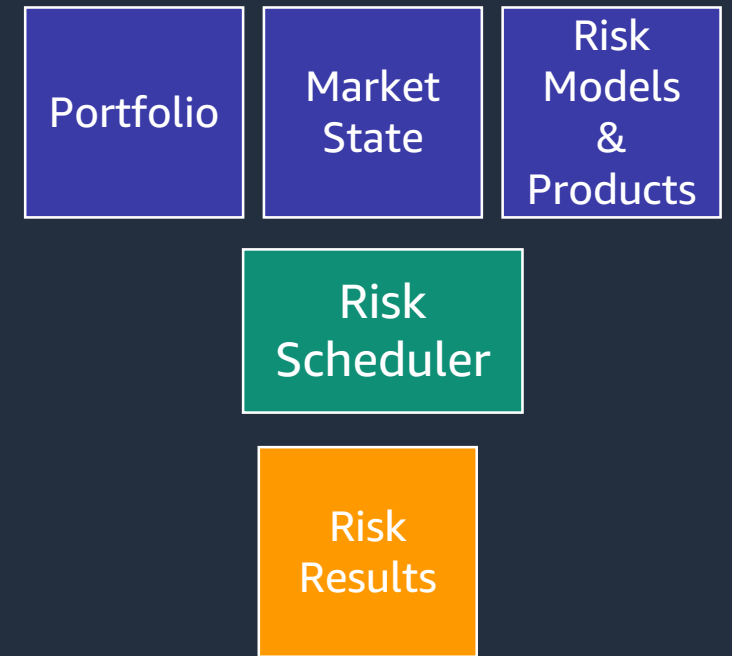
Data Pipeline



Hedging Engine



Risk Management




Elastic ML training and risk: requirements



Extensible to batch and real-time jobs



Resilient to component loss



Modern GPU compute resources



Minimize resource consumption



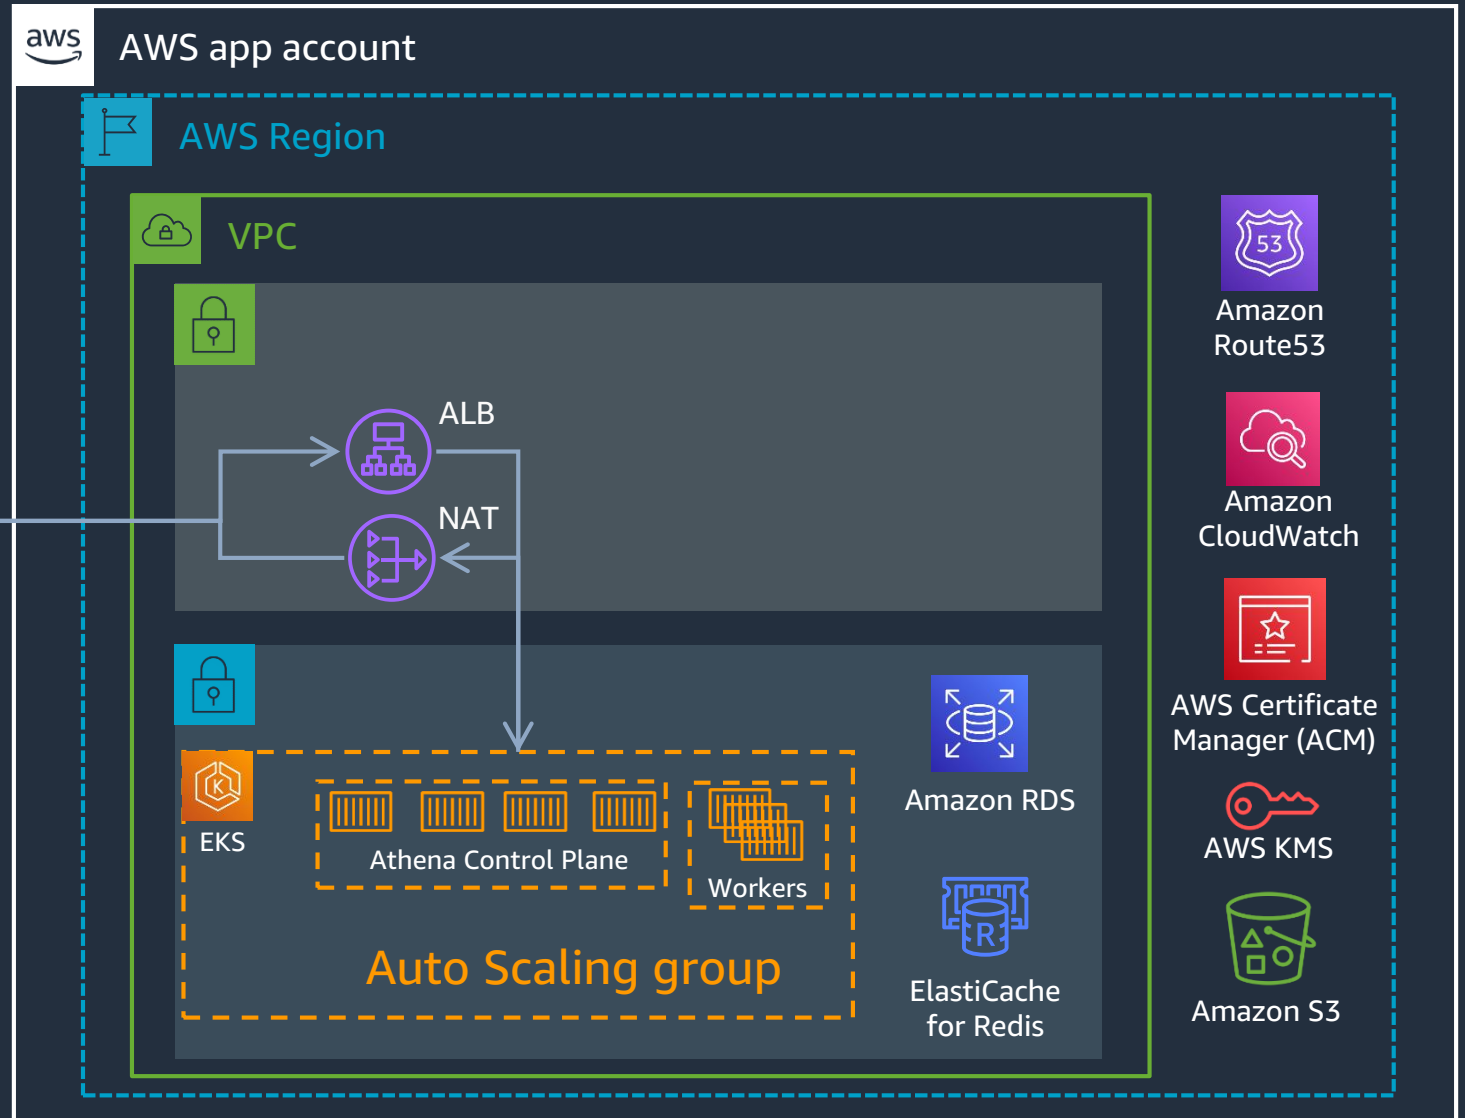
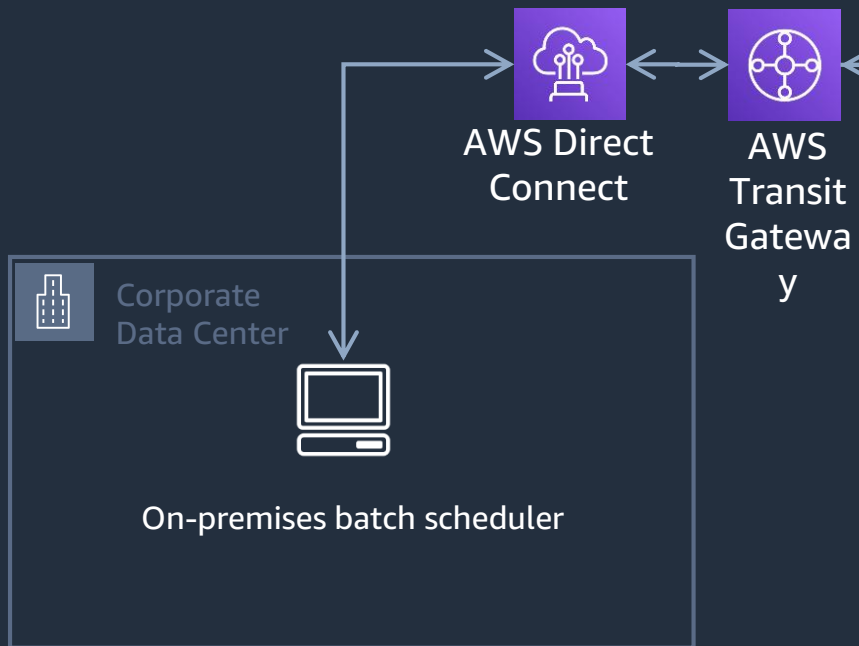
Minimize E2E job duration



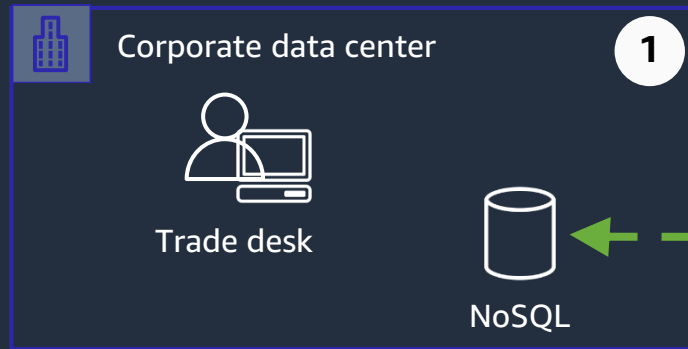
SLA management of jobs via dynamic priorities

Athena AWS architecture

- Subnet with ALB, NAT connected to premises
- Subnet with Athena control plane, Amazon EKS with Auto Scaling, Amazon RDS, and Amazon ElastiCache
- On-premises scheduler



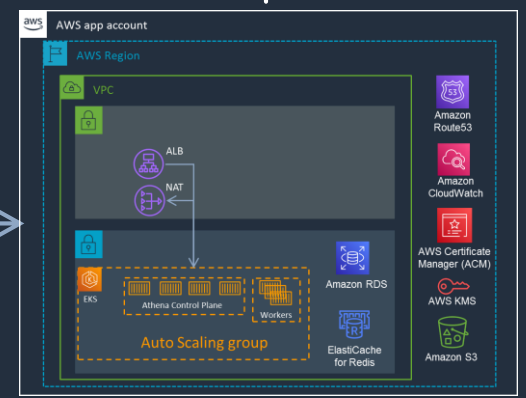
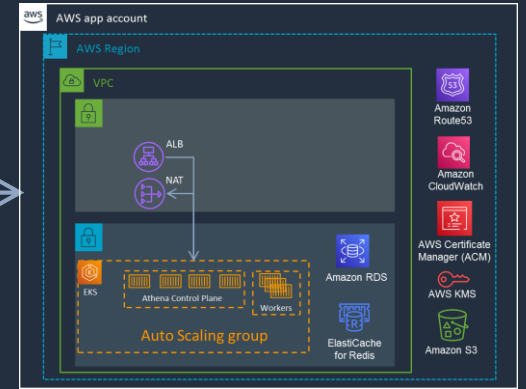
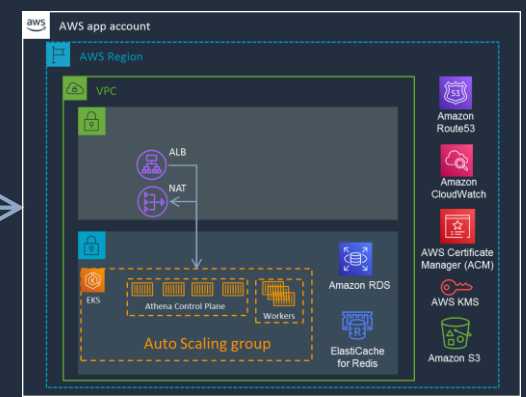
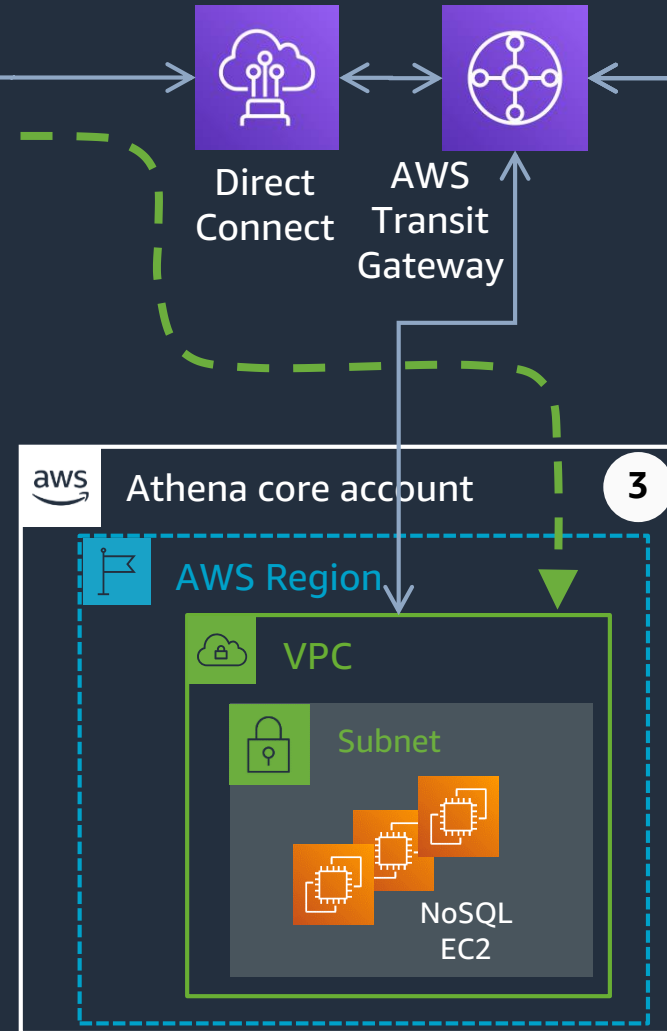
Athena cloud multi-account model



1. On-premises NoSQL database and trader desk hedging portfolio & consuming risk

2. Individual AWS accounts for each line of business

3. NoSQL read-only database running on Amazon EC2



Platform design goals for the ML use-case

1. Establishing an elastic Quant development environment with access to primary data
2. ML Pipeline which allows for full traceability and version management
3. Leveraging HPC with our existing quant libraries
4. Cost management
5. Operational management

Athena ML Pipeline



Corporate data center



Quant Research

Develop Data Pipeline and Model Architecture

Data Sources



Model Dev



Hedging Service

Production

Data Sources



Prod Servers



AWS Development Account



Training & Testing Data



Results



Jupyter Notebook Kernel Environment: Development



AWS Production Account



Prod Training & Portfolio Data



Hedge Proposals



Elastic Training & Inference Environment: Production



Implementation challenges and outcomes



Compute

- GPU reservations using On Demand Capacity Reservation
- GPU efficiency using Multi-Instance GPU (MIG) virtualization
- HW-aware kubernetes control plane



Data

- Very large data sets required for training, frequently updated



AWS Collaboration

- Hedging service in production
- Changes our business
 - Portfolios have increased hedge efficiency
 - Calculation time far quicker
- Using existing and new software
- Elastic and modern hardware, adaptable to new capabilities

Next steps

- Continued certification of new GPU categories
- Tighter auto-scaling for training, testing, and model inference
- Efficient use of reserved instances and support for multiple regions
- Cost reporting and allocations, setting limits in certain accounts



Thank you!

Paul Bauerschmidt