


TAMING THE PETABYTE MONSTER

JASON HAMMONS, VP OF FIELD ENGINEERING




THE PETABYTE MONSTER IS GROWING



40 Years of Compounding Growth



New Approaches to Research




New Trading Instruments

40 YEARS OF COMPOUNDING GROWTH

- Quant Research Now 40 Years Old
- Datasets now growing 2x Y/Y
- VIX avg is roughly 2x since Covid

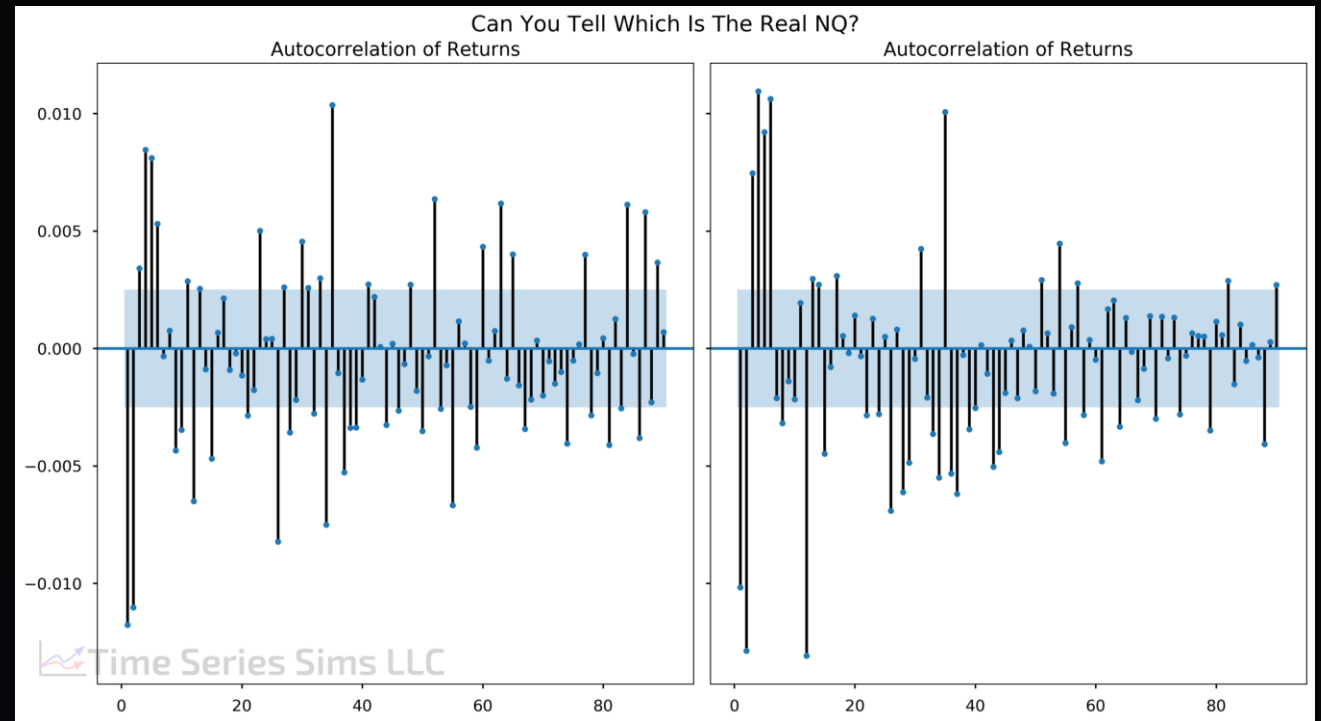


Scaling At One of VAST's Largest Quant Trading Customers



THE MACHINES ARE NOW CREATING DATA

- Monte Carlo is being challenged by lack of sophistication and auto-correlation
- More firms are turning to DL, but the “deep” in deep learning = lots of data
- Where to get the training data? **Make it!**



CRYPTO IS A DATA GROWTH CATALYST

- Full ownership lineage
- 500 exchanges trading coins
- More metadata to analyze



7x Larger Crypto
Trade Data
Payload

Options
Trade Data
Payload



**DATA IS GROWING AT
HISTORIC LEVELS &
MUST BE ACCESSED IN
REAL-TIME**


**BUSINESSES NEED TO
CONTROL THE COST
OF REAL-TIME
INFRASTRUCTURE**



ONE TIER TO RULE THEM ALL


NFS • SMB • S3 • GPUDirect Storage

1,000s of stateless containers
Zero east west traffic, linear scale



NVMe Fabric


Exabytes of low-cost flash
All-flash at archive price-points



DASE: the world's first partition-less scale out architecture

THE PATH TO FLASH AFFORDABILITY

VAST predictive placement saves 85%




Enable hyper-scale flash

Data protection saves up another ~30%



Locally-decodable erasure codes

Unprecedented storage efficiency (2:1 on compressed data)



Similarity-based data reduction



ALL-FLASH MARKET DATA ARCHIVES ARE WHAT WE DO

THANK YOU

www.vastdata.com

