

Levyx: Trends in Low Latency Analytics and Computational Storage

October 2018
STAC Summit

Matt Meinel - SVP Sales and Solutions Architecture

For More Information:

Email: sales@levyx.com

Computational Storage in Oct 2018



Computational Storage SNIA®

The SNIA is in the process of forming a Computational Storage Technical Work Group (TWG) which is being created for the purpose of establishing architectures and software for disk and solid state drive based functionalities that allow them to be integrated with Computation in its many forms. The TWG will create software and standards that enable specific features for these drives that meet the requirements of customers with these computational needs.

Reducing Cost: Leverage 3D NAND for Risk Analytics

Cost Savings of 75+%

Backtesting App running Levyx/Spark on Cloud

Scale B: 1 million instrument simulations in 1 hour using a basket of the ideal size for the SUT	
Ideal basket size for this SUT (size that maximizes instrument simulations per second)	200 instruments
SWEEP speed on baskets of ideal size (from Table 3)	12.6 instrument simulations per second
Scale-up factor to achieve Scale B, assuming linear scaling	22.1 x
Price Performance for Scale B	215.37 USD per million instrument simulations

Backtesting App running with Spark Xenon on Vexata*

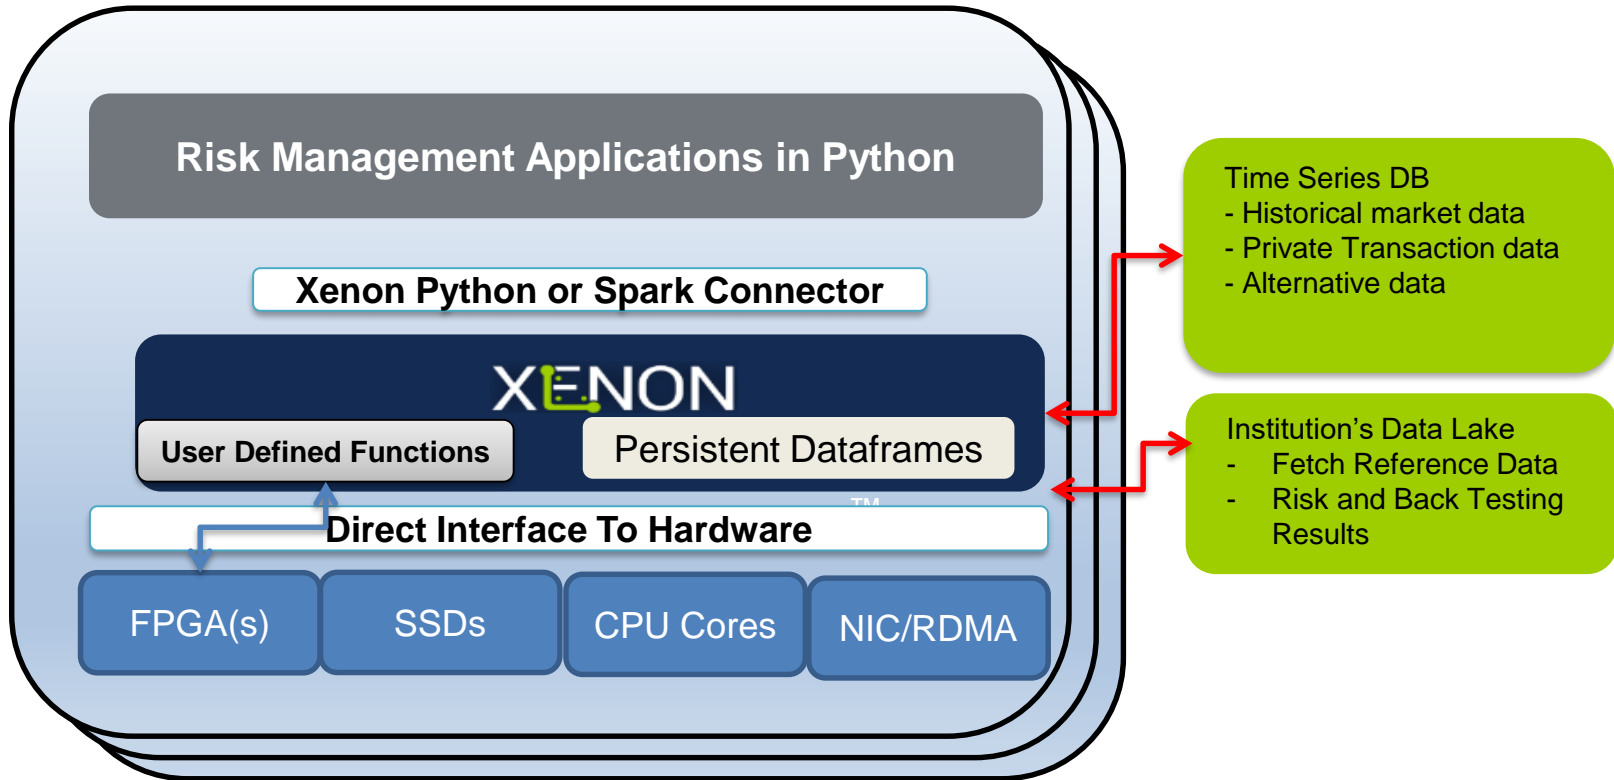
SCALE B CALCULATIONS		
Number of worker nodes in SUT		6.00
Number of worker nodes required for Scale B		24.00
Costs per hour per server + support		.062
	Total	\$14.88
Cost of Vexata Array + network		\$13.31
Enterprise Support		\$6.84
Solutions Architect Salary (.2 Headcount)		\$4.56
Master Server (one of the worker nodes)		\$0.00
	Total	\$24.71
Total cost of Scale B cluster over one hour		\$39.59

* Based on inputting STAC-A3 performance results into Levyx TCO on-prem calculator

Sources: www.stacresearch.com/Levyx ; <https://stacresearch.com/LEVX180608>; <https://stacresearch.com/LEVX170603>



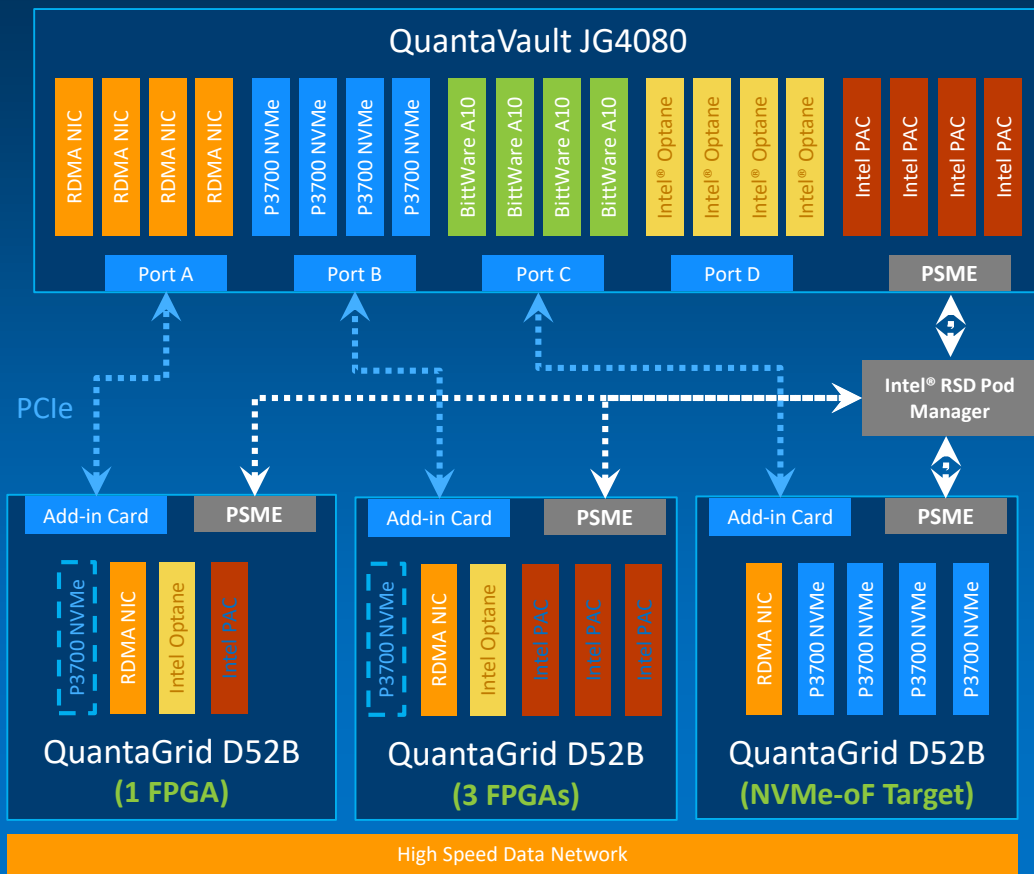
Using Levyx Xenon™ Persistent DataFrames in Risk Acceleration Framework



QCT Rackgo-R based on Intel® RSD



FPGA
and
NVMe
pooling




Four hosts
and up to
twenty devices


High speed PCIe
end-to-end

Minimize TCO with Intel® RSD

A traditional data center might
deploy 3 single-FPGA servers

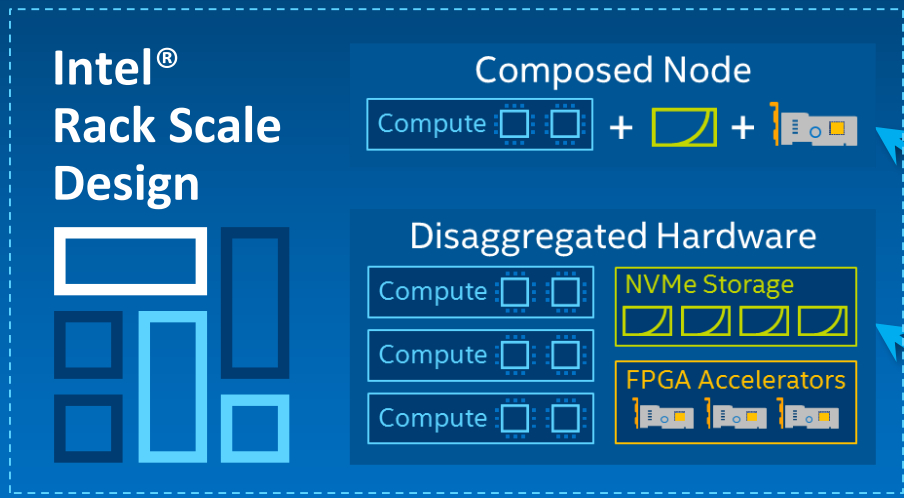


With Intel® RSD you can compose 1
server with 3 FPGAs



Workload is FPGA accelerated SLA
requires three FPGAs

FPGA & NVMe Pooling with Intel® RSD



What you will see...

Intel FPGA-accelerated Workloads



Resource Discovery, Composition, Management



Hardware based on Intel RSD



End-to-End Solution with Pooled Resources!

Risk Analytics Acceleration Framework – 2

Applications

- Accelerating financial options trading and computations:

Dynamically scale FPGAs on CPU

→ Performance benefits and TCO savings

- Accelerating financial stock backtesting operations:

Intel Optane™ versus Flash SSDs

→ Better performance and cost savings

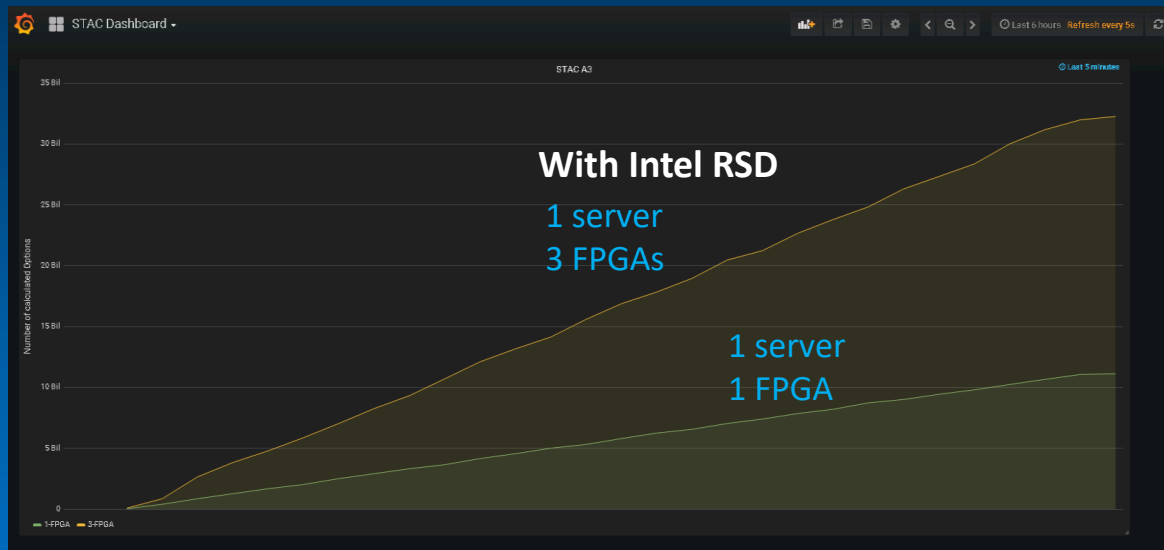
Powered by:



Levyx Financial Workload with Intel® RSD, Optane and PACs



3x FPGAs =
3x performance
but only
consumes
1 server with
Intel RSD



TCO Savings and Added Flexibility
Met SLA with less hardware!

* Not a STAC Benchmark

THANK YOU!

Please check the box for Levyx
and
Check out our Demo at the Intel Booth!

For more info: sales@levyx.com