

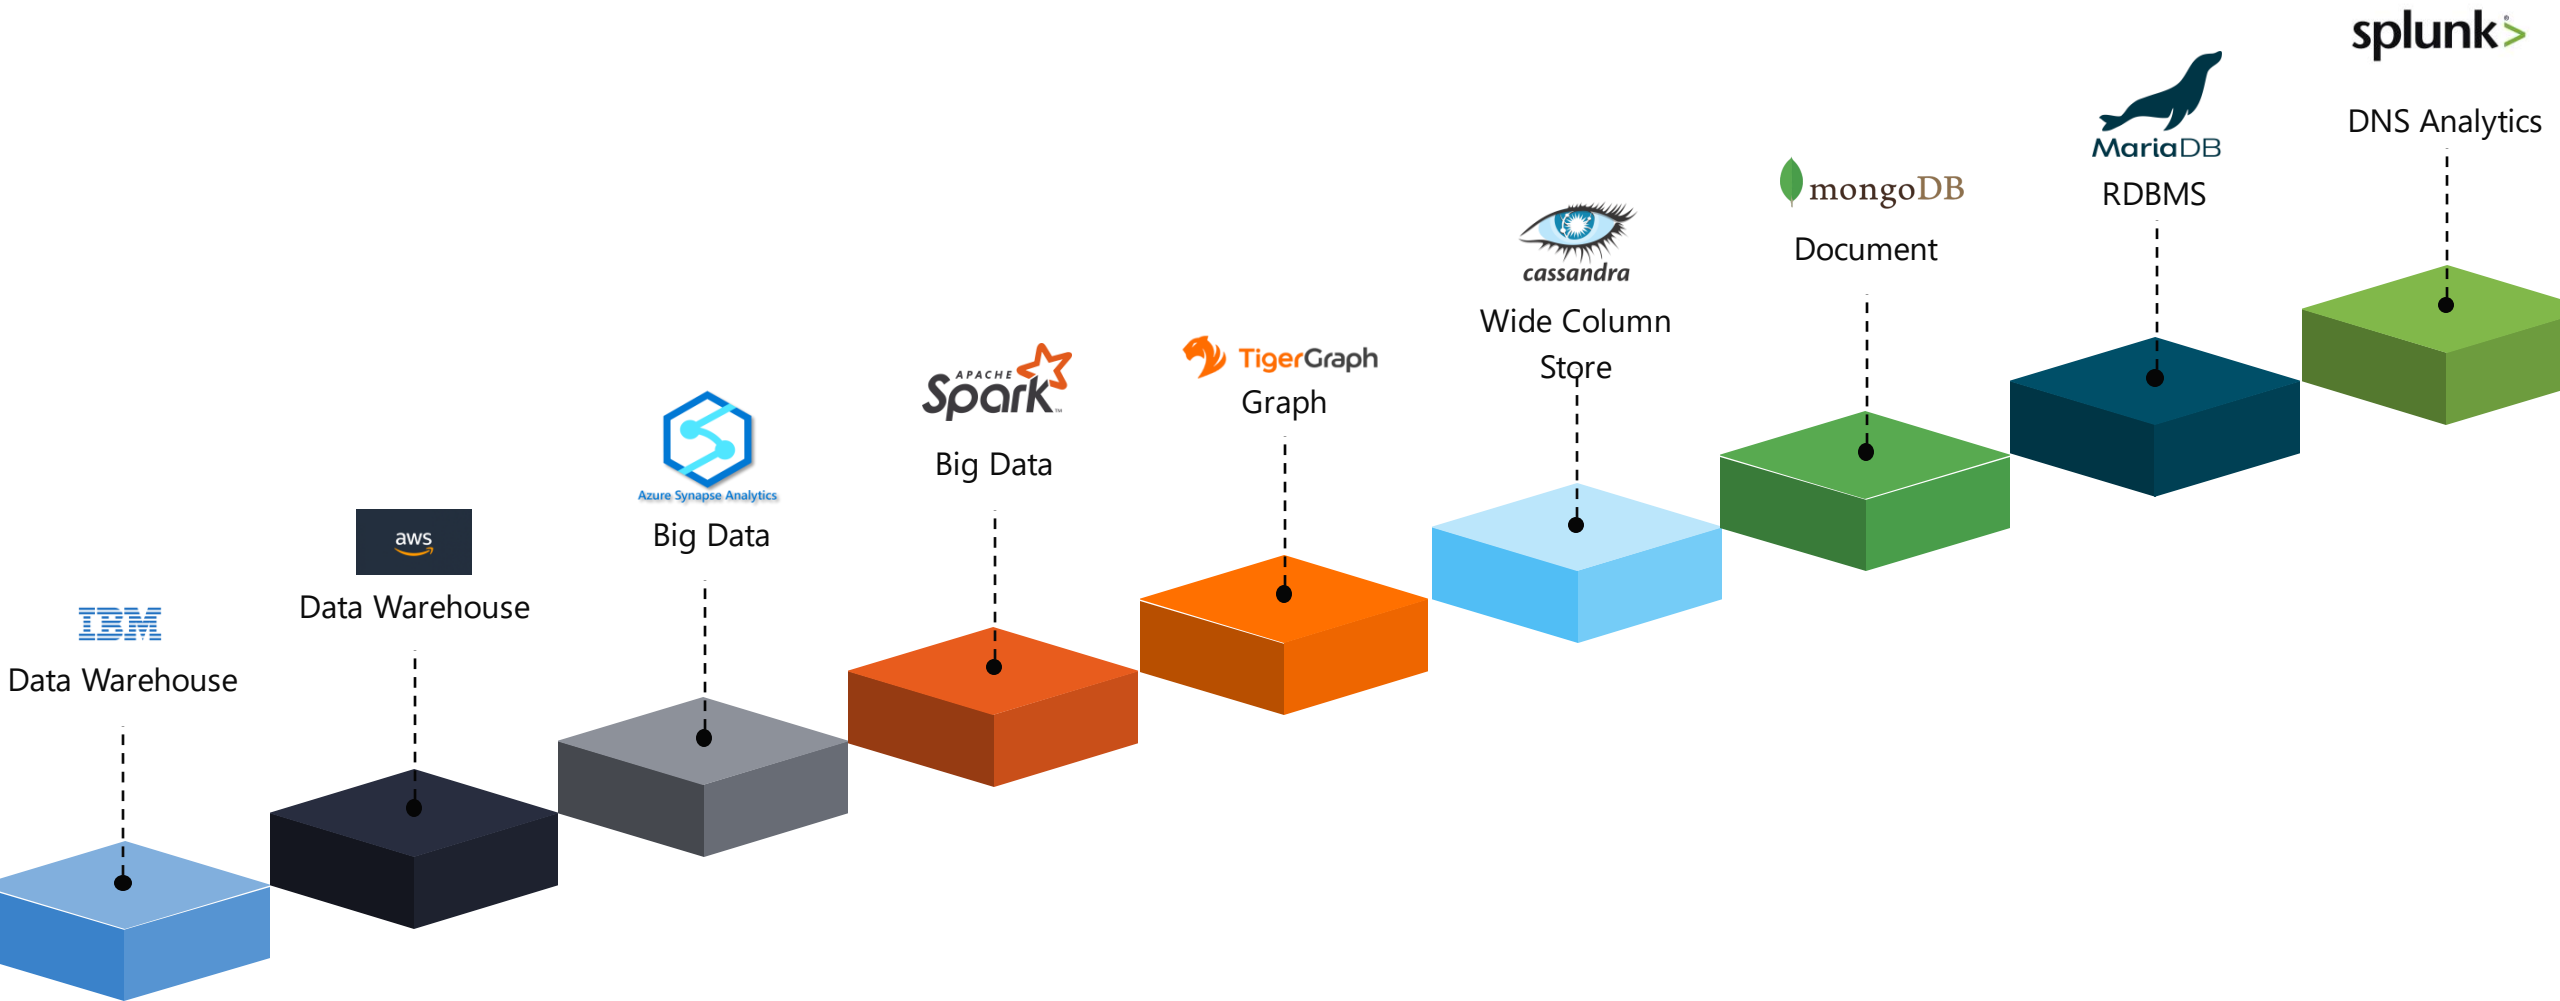


# Turbocharge Database Performance and Analytics with Xilinx

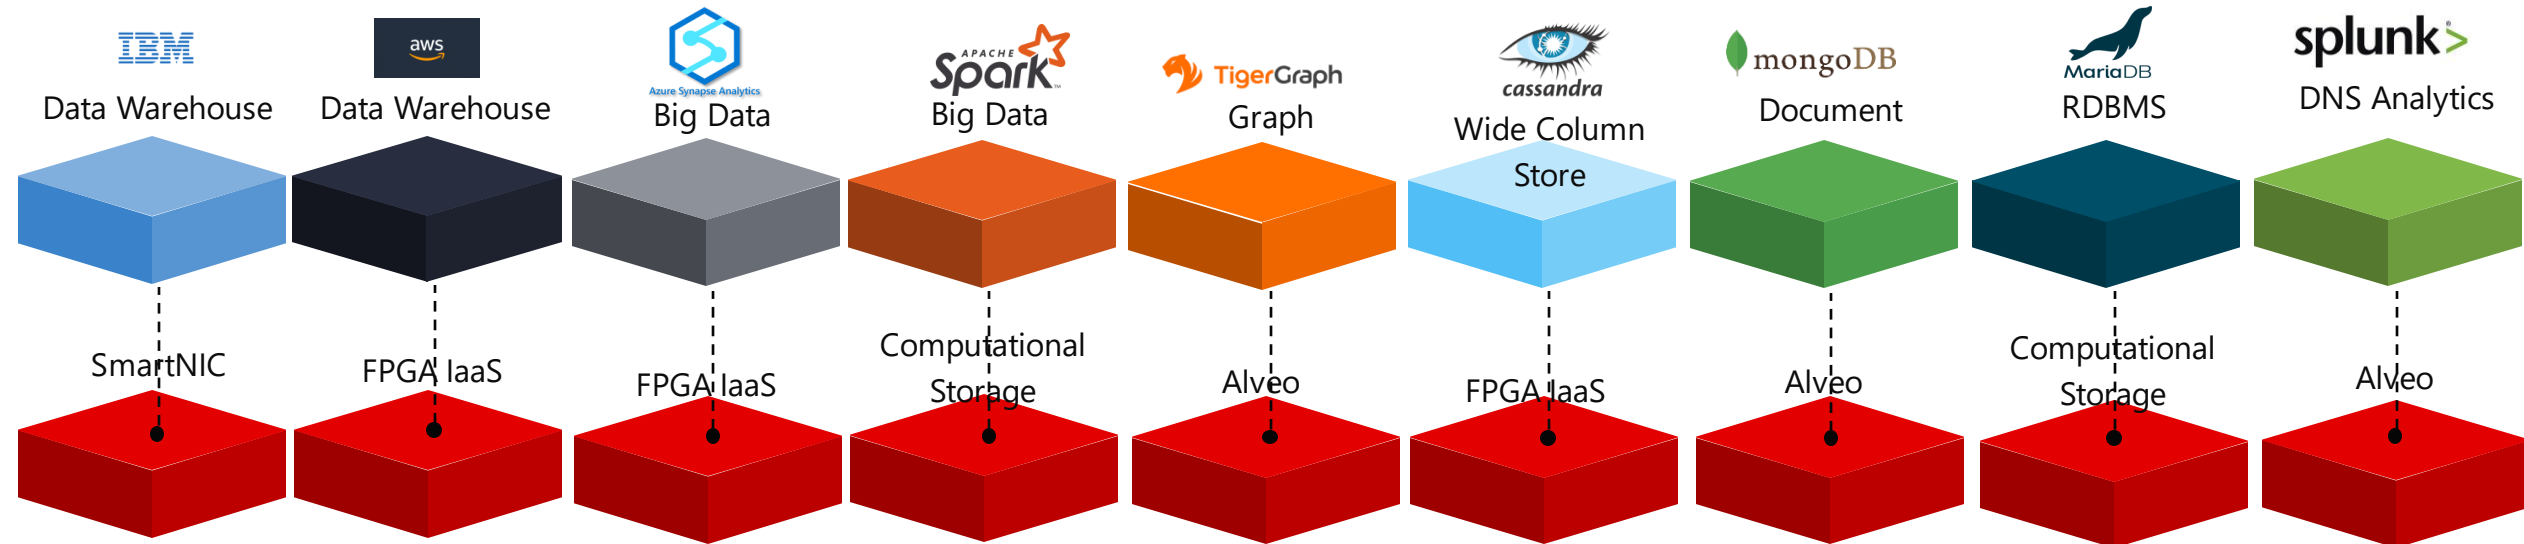
Dan Eaton & Vinay Singh



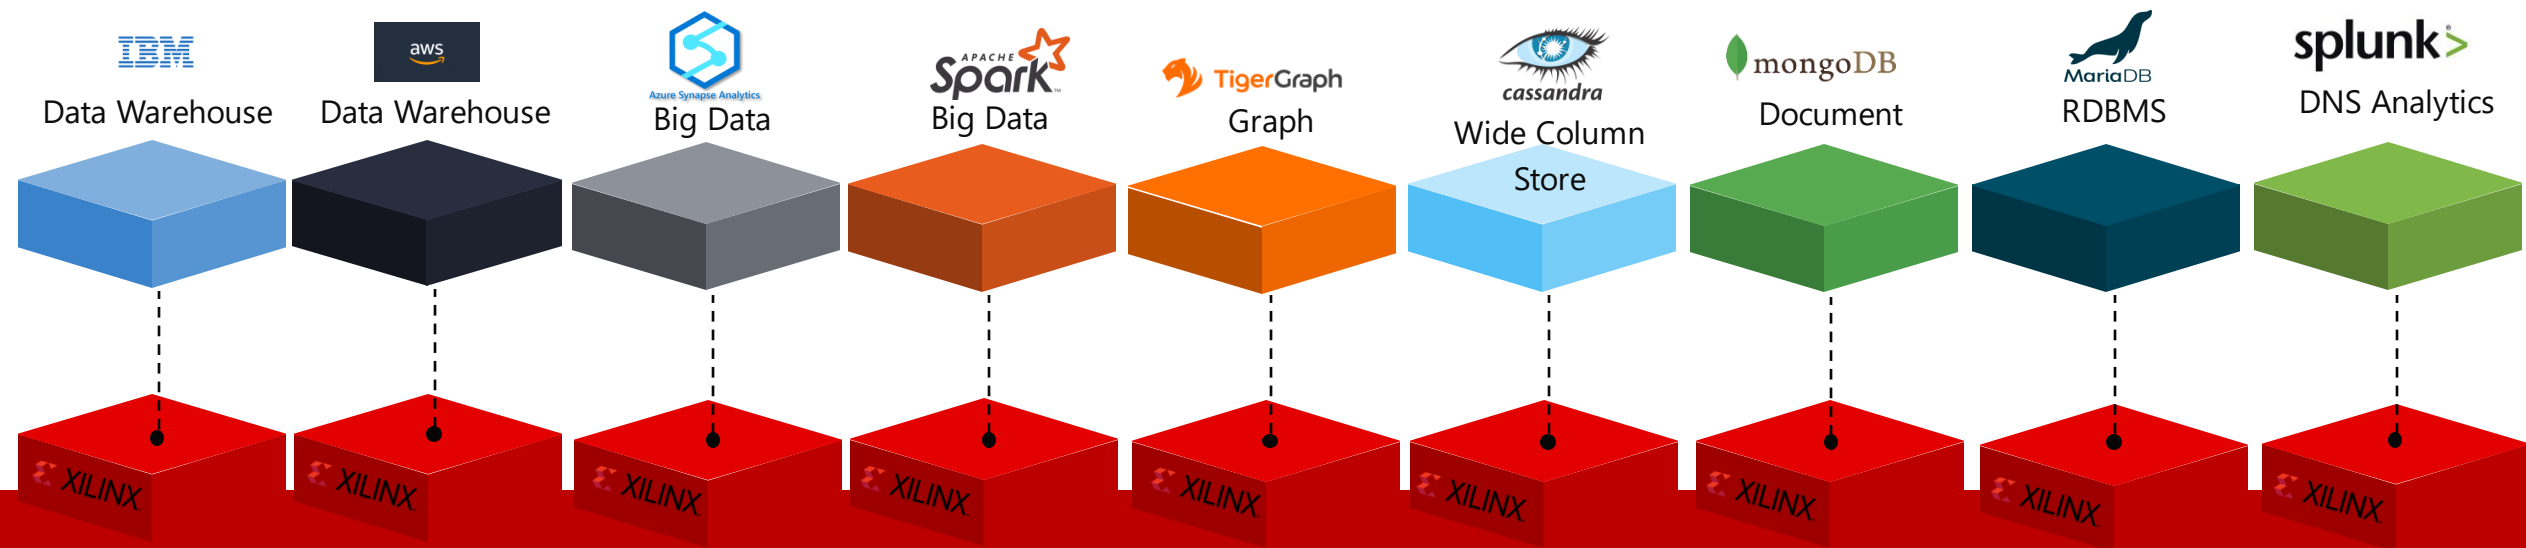
# Xilinx Accelerated Databases and Analytics



# Xilinx Accelerated Databases and Analytics



# Xilinx Accelerated Databases and Analytics



Public Cloud Partners


OEM Partners

Solution Integrators

ODM Partners



# Xilinx Accelerated Databases and Analytics



Public Cloud Partners

OEM Partners

Solution Integrators

ODM Partners



FPGAs

SmartNIC

Accelerated Computing

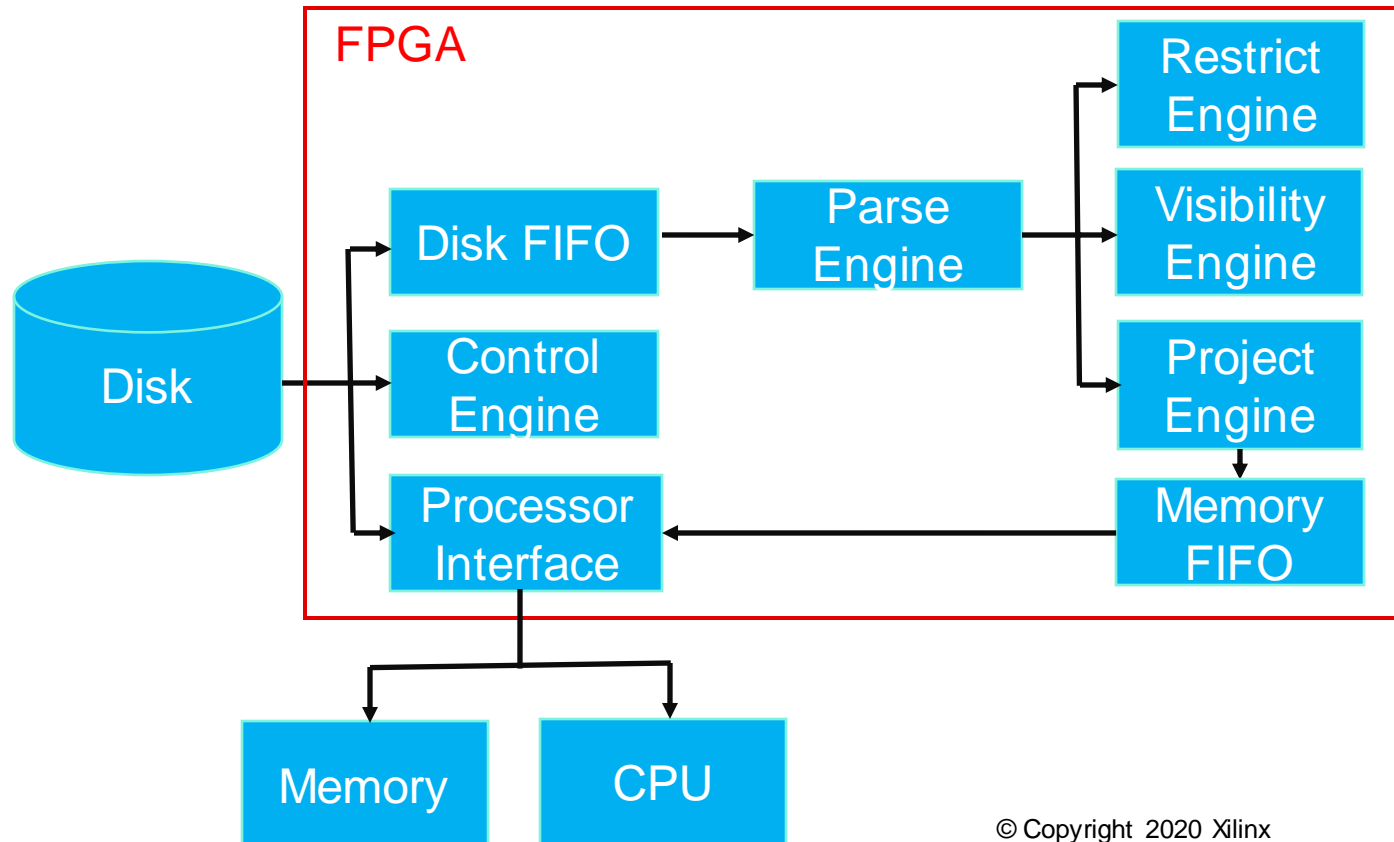

Computational Storage



# IBM Netezza Performance Server


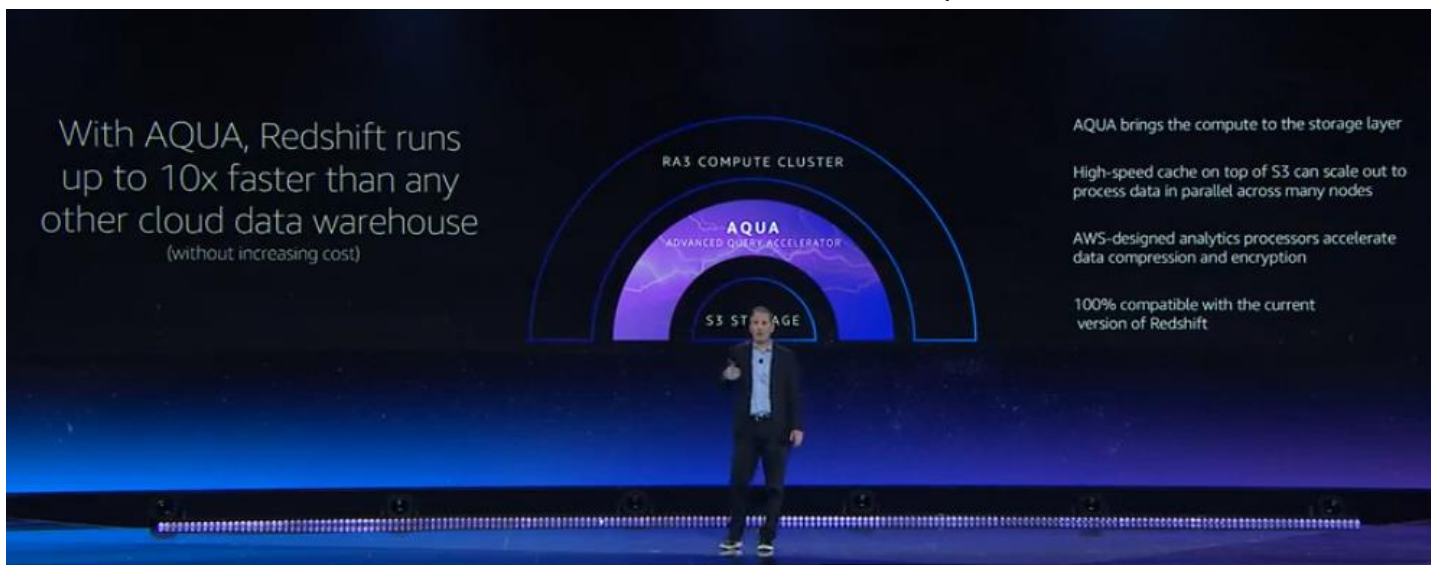


Data Warehouse

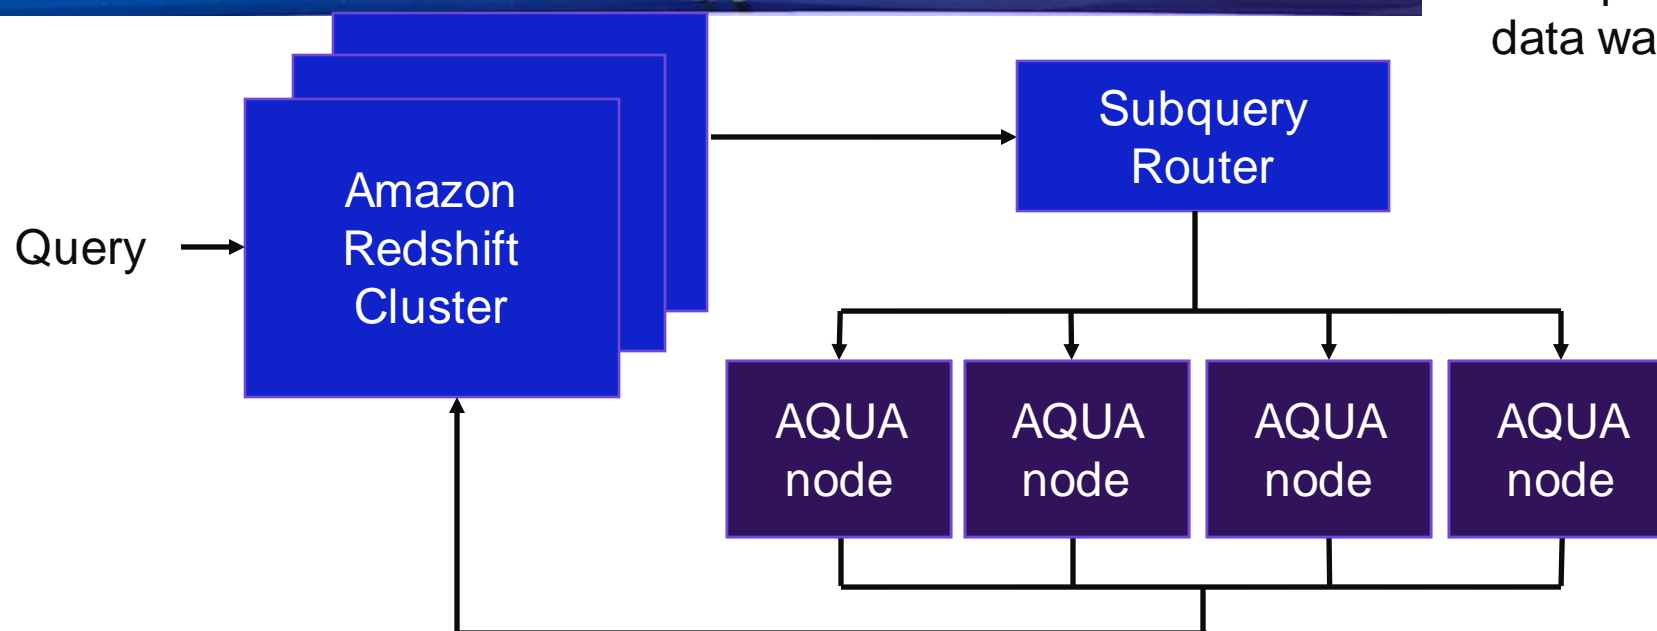


Netezza Performance Server leverages FPGAs to accelerated query processing and machine learning

# Amazon Redshift AQUA



AQUA is a new distributed FPGA-accelerated cache that enables Redshift to run up to 10x faster than any other cloud data warehouse.




AQUA accelerates encryption, compression, filtering, hashing and aggregation to maximize throughput



# Microsoft Azure Synapse



**Teams using FPGAs**  
with new language + toolchain

 <b>Azure accelerated networking</b> 400% better packet parsing throughput (vs. RTL baseline)	 <b>Azure storage protocol engine</b> 6% lower logic utilization 30% lower memory utilization 140% throughput (vs. RTL baseline)	 <b>Azure synapse</b> 4x+ throughput/\$ for data analytics tasks vs. CPUs alone	 <b>Azure quantum inspired algorithms</b> 100x-200x performance gain vs. 6-8 core CPU
--	---	--	--

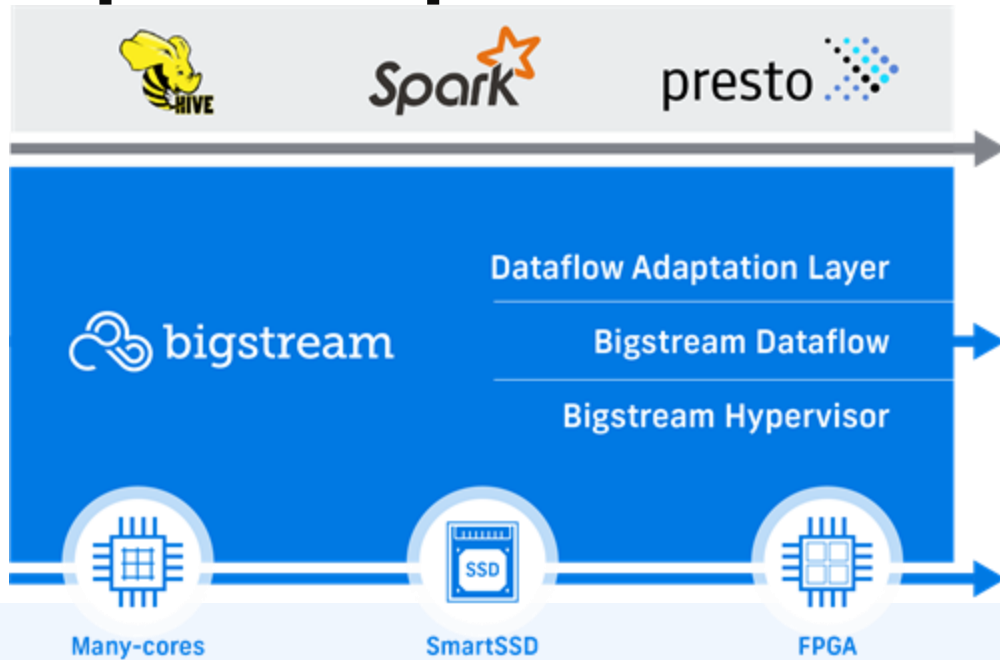
to improve the time to solution by orders of magnitude



The diagram shows a central blue cylinder labeled "SQL" and a blue square labeled "APACHE Spark" with an orange star. They are surrounded by several blue squares, each containing a small 3D cube made of colored blocks. A dashed line connects the "Big Data" text to a grey 3D cube on the right.



# Apache Spark



Zero code changes

Cross-platform versatility


Intelligent, automatic  
computation slicing

Cross-acceleration hardware

Up to 10x acceleration




Big Data

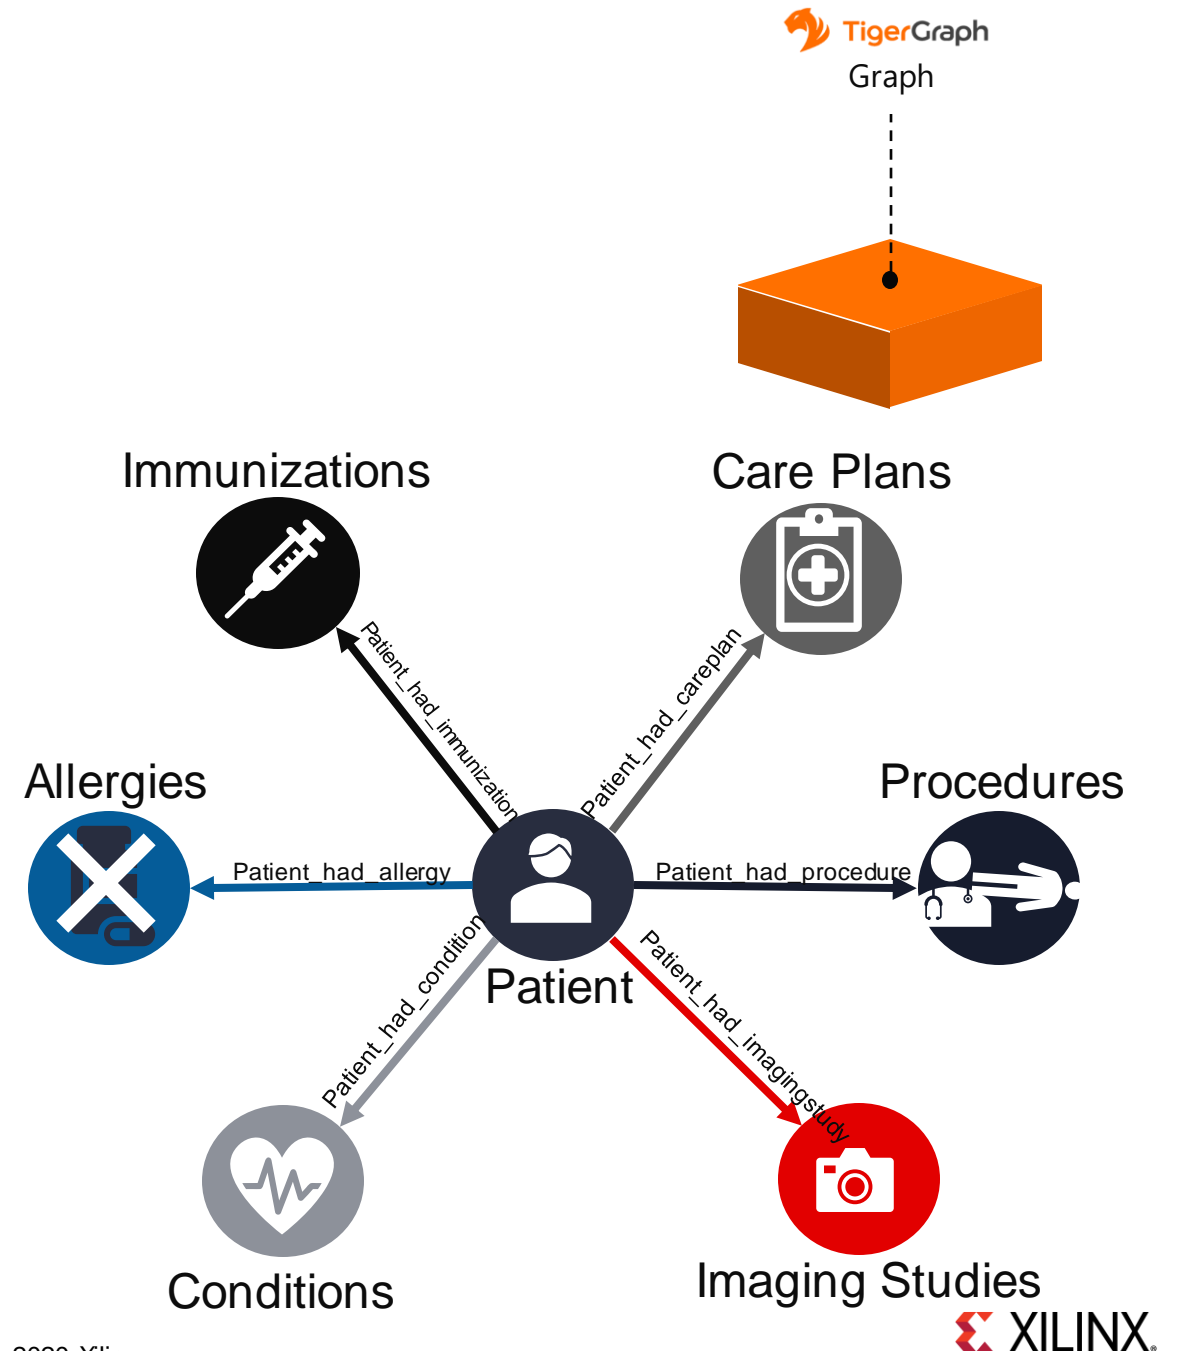
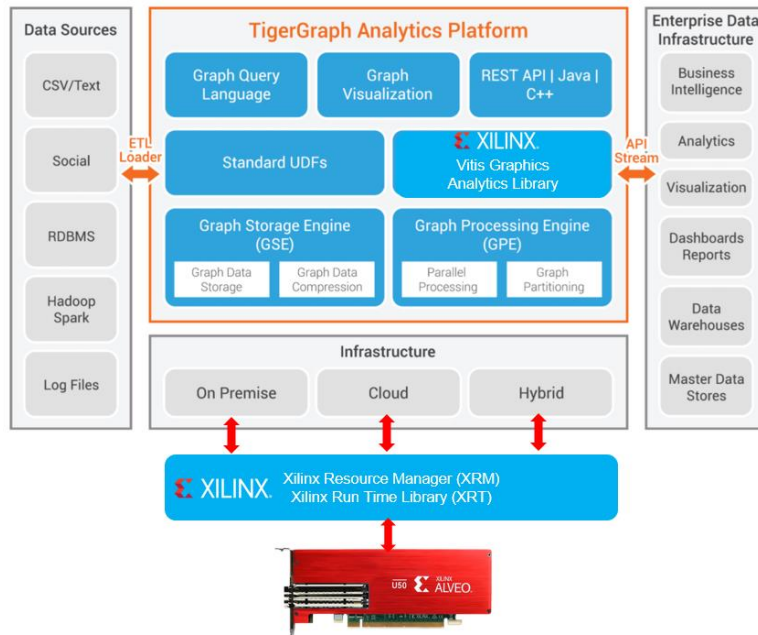


TPCDS 40KSF JSON (GZipped)

Acceleration of runtimes compared to spark

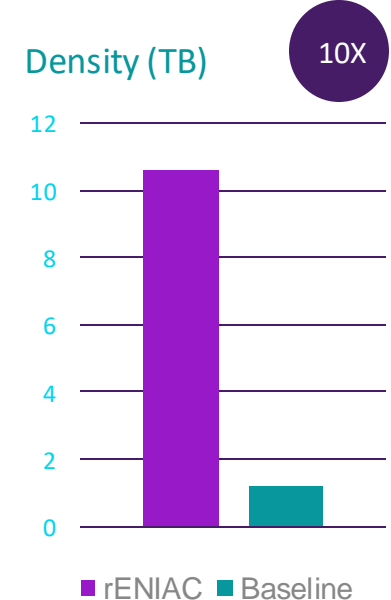
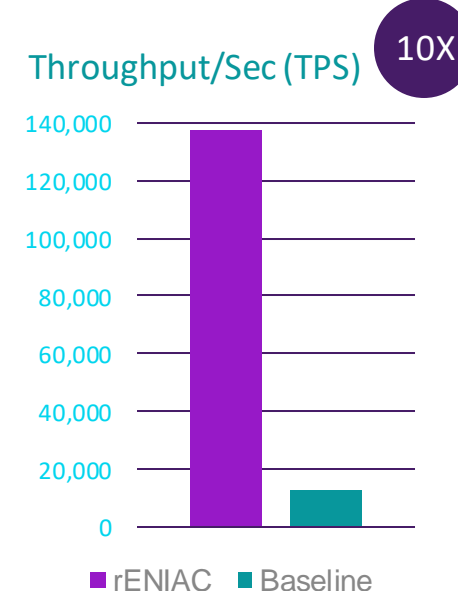

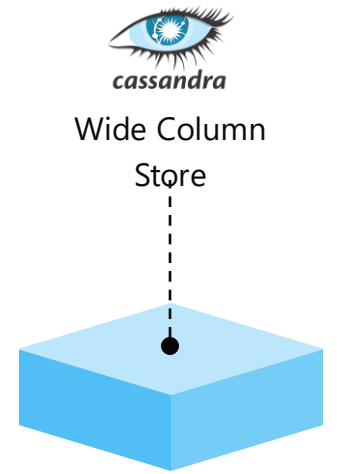


# TigerGraph



# Apache Cassandra

- ▶ Elastic performance to scale throughput and reduce latency on demand
- ▶ Use cases across retail, eCommerce, finance, telecommunications or PB scale real-time data processing
- ▶ Deploy in minutes with no application code changes required



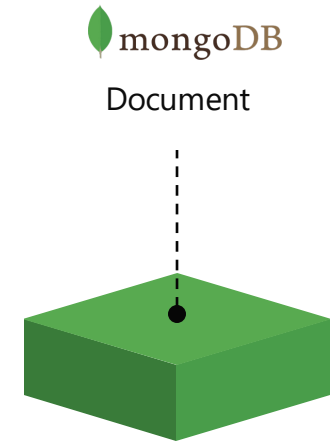
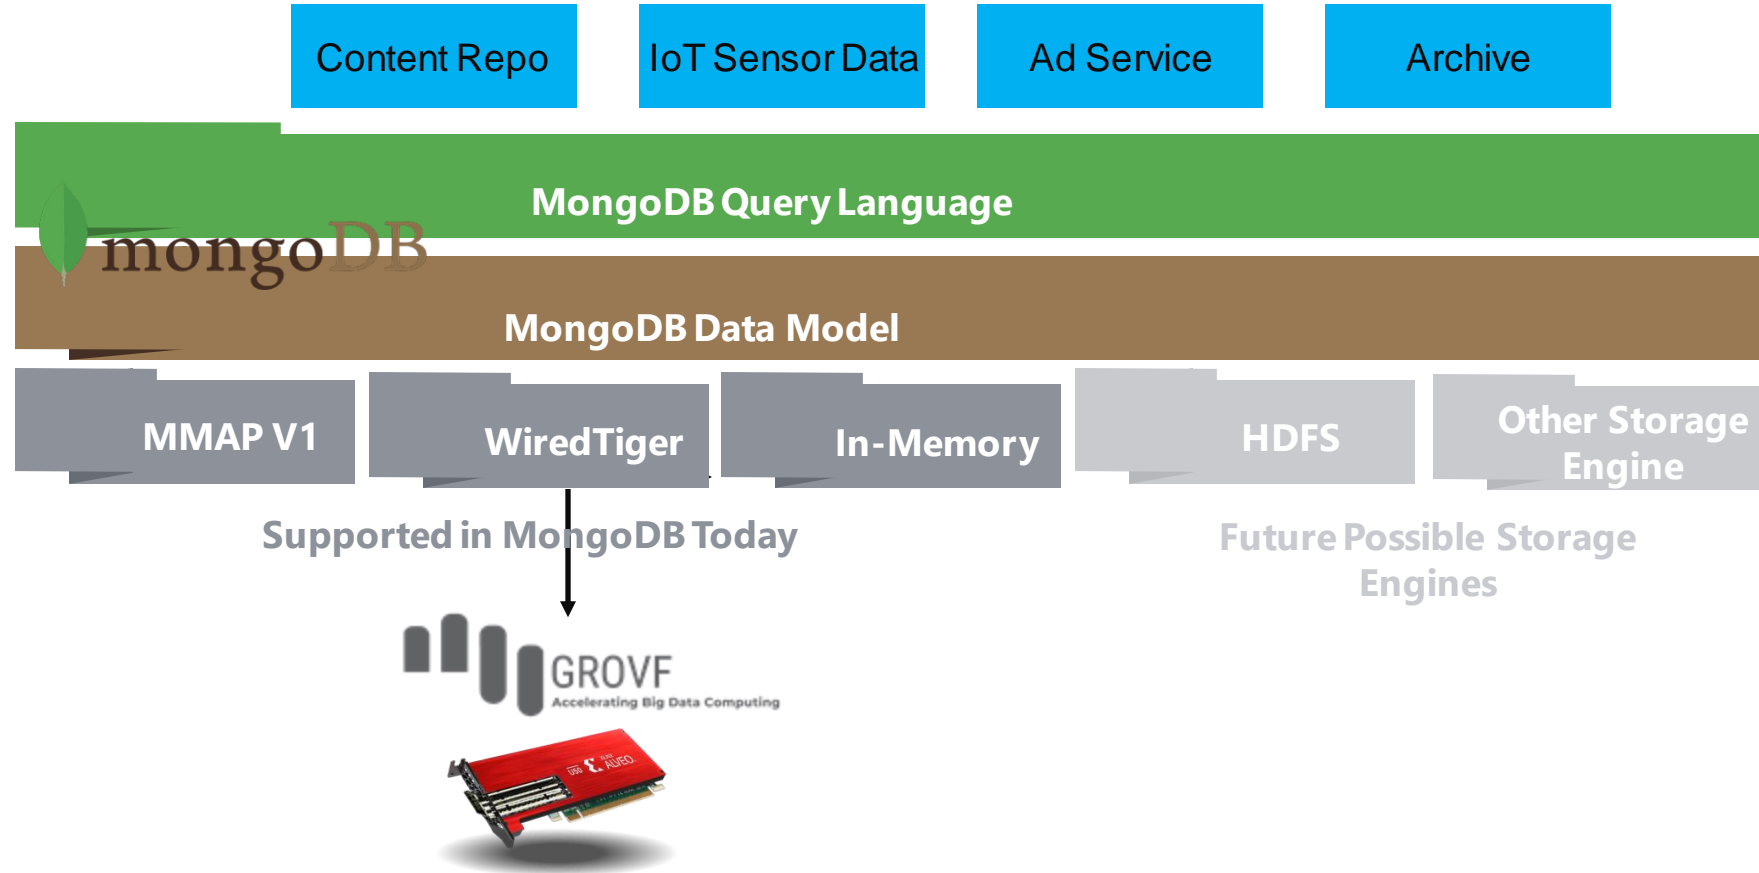
Measured for NOSQL DB: Cassandra

© Copyright 2020 Xilinx




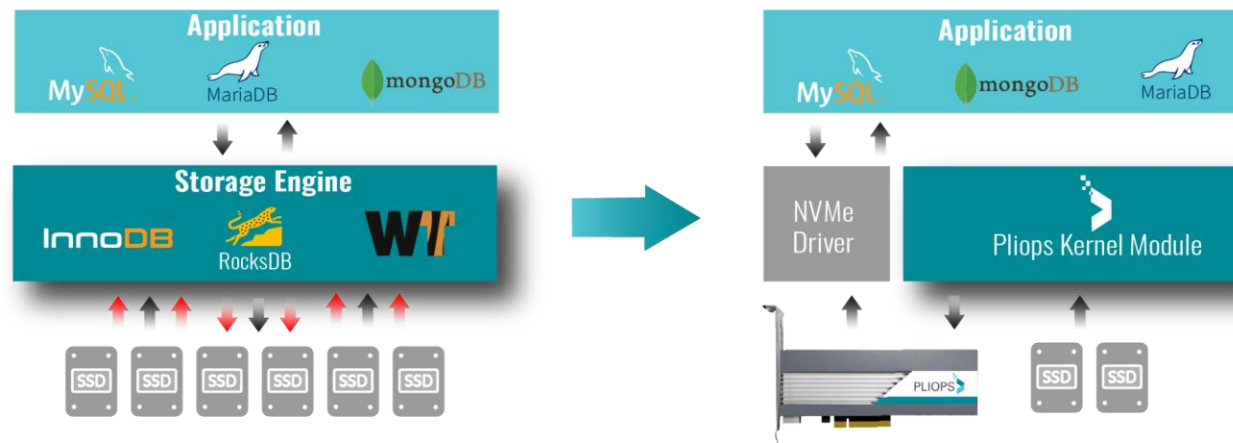
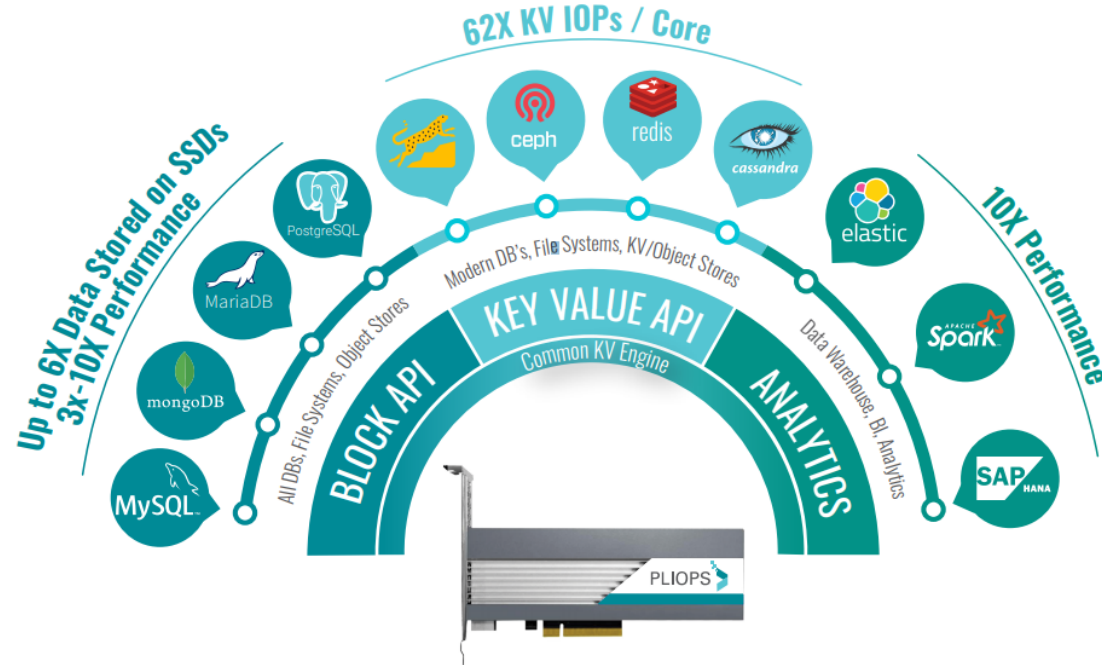
Wide Column  
Store

# MongoDB




MongoDB performance can be improved by 3.5x using MonetX and Alveo U50 as a memory extension technology.

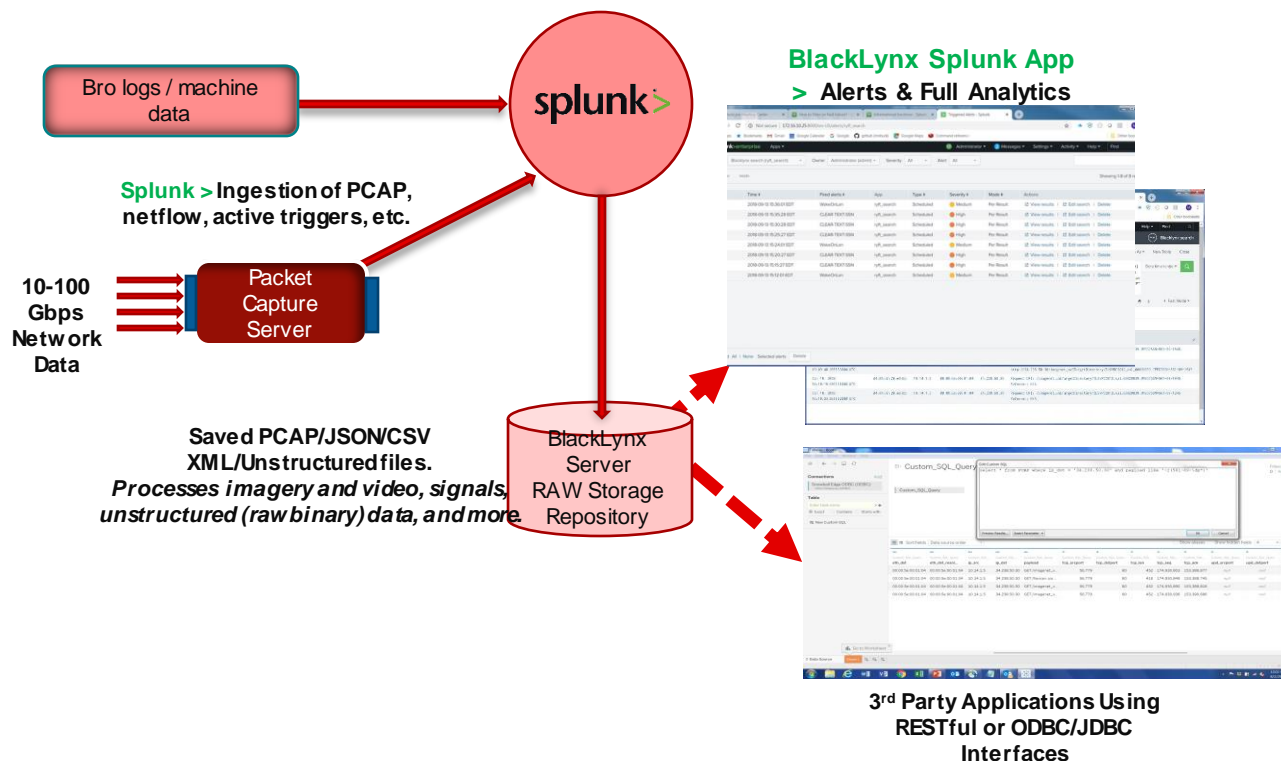
# MariaDB



MariaDB, PostgreSQL, MySQL can be accelerated by Pliops Data Processor



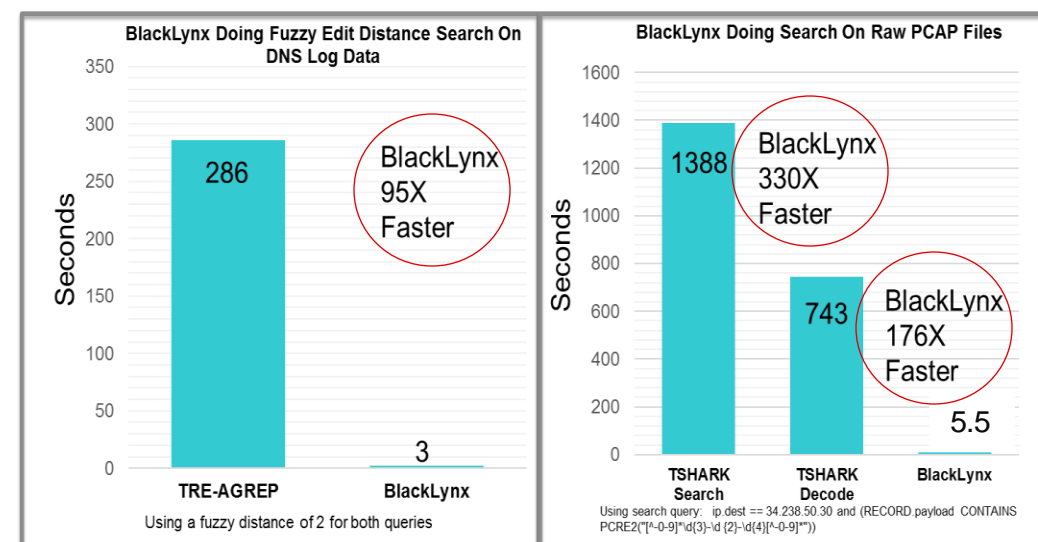

# Splunk



Leverage adding high end search capabilities (fuzzy search, regular expressions, raw PCAP) to handle the explosive growth of machine data in Splunk.

splunk>

DNS Analytics



Benchmark comparison for Fuzzy Edit Distance and PCAP primitives



# Xilinx Vitis Libraries





# Xilinx Accelerator Program

## Accelerator Program

Be part of the dynamic and growing ecosystem of Xilinx ISVs!

[Apply Now](#)



Support



Alveo Loaner Cards



Co-Marketing



Joint Sales Engagements

# Related Adapt Sessions



▶ **AWS and Xilinx: Unleash Performance and Unlimited Scale with AWS F1**













▶ **Azure and Alveo: Cloud Migration Made Easy**



▶ **Samsung and Xilinx: Delivering Composable Computational Storage**

▶ **Database and Analytics Ecosystem Presentations:**

 Key Value Store	 Key Value Store	 Apache Spark	 DNS Analytics	 Machine Learning
 Search Appliance	 Visual Search	 Packet Classification	 Apache Cassandra	 Machine Learning



---

Interested in learning more? [singhj@xilinx.com](mailto:singhj@xilinx.com)

# Thank You