

GPU DATA WAREHOUSE FOR MASSIVE DATA ANALYTICS

Rafi Glickman, APAC Sales Director



CORPORATE PROFILE FOUNDED IN 2010

Sigma 75+
EmployeesImage: 10
Patents

Strategic Partnership with Alibaba Cloud

O HQ in 7 WTC New York | R&D in Tel Aviv



A STATE OF A STATE

Agenda

- Data is growing exponentially, but your database can't keep up
- Why Big Data Analytics?
- The Challenge-Extracting Value out of Big Data
- Why existing solutions are not good enough?
- SQream Value Proposition
- Case Studies





Customers are looking for more and NEW Insights



Why Big Data Analytics?

- Data is the new oil
- Not a nice to have a MUST have
- More data , more dimensions, time period more insights, more value!





What customers are telling us

- We want to become a data driven organization
- We have a lot of data, but analyze only a fraction!
- Our systems were not designed to handle the amount of data
- Time of data preparation and modeling is growing constantly
- Maintaining and upgrading our existing data warehouse is increasingly expensive



WHY is it so challenging?

Data Sources

Customers

Transactio

ns

Positions

Activity

???

ETL

Pipelines

MARKETING Customer 360^o and multi channel experience

<u>FRAUD</u> Enhancing the fraud detection algorithms

<u>COMPLIANCE</u> Producing capital adequacy and Risk

PAIN examples

PERFORMANCE: Slow / Stuck processes -> Long time to Insight -> Loosing opportunities

SIZE: Analyzing only part of the data -> less accurate insights -> Loosing money

COST: Complicated and costly systems -> No scalability / Missing Insights -> Penalties





Extracting Value out of Big Data is not that simple !

- Most databases can't keep up with the current amounts of data
- Cost
- Time to Value
- Knowledge
- Technical Resources
- Technical barriers



BUT YOUR DATABASE WASN'T BUILT TO HANDLE THIS LEVEL OF DATA



VALUABLE INSIGHTS GO UNDISCOVERED





<10%

Data Analyzed

Valuable Insights go undiscovered

	Potential impact:		Value captured %
Location-based data	 \$100 billion+ revenues for service providers Up to \$700 billion value to end users 		rs 50– 60
US retail ¹	 60%+ increase in net margin 0.5–1.0% annual productivity growth 		30- 40
Manufacturing ²	 Up to 50% lower product development cost Up to 25% lower operating cost Up to 30% gross margin increase 		it 20- 30
EU public sector ³	 ~€250 billion value per y ~0.5% annual productivi 	ear ty growth	10-20
US health care	 \$300 billion value per ye ~0.7% annual productivi 	ar ty growth	10-20
			SOREAN



The more data you have , the less % can be analyzed

Data Lake	Analytic Data	% of Data
100 TB	10 TB	10 %
500 TB	20 TB	4 %
1 PB	30 TB	3 %
10 PB	50 TB	0.5 %

- Do you believe history has value ?
- A full cycle can take ~ 7 years so to get the real analytics you need to analyze 14 years
- Any hiccup may take you to the wrong direction





The value in your Big Data

Potential economic impact of \$110 billion in retail banking industry

- Enable better cross-selling
- Personalized products
- Dynamic pricing
- Better risk assessment
- More effective marketing

10-20% of the potential value has been captured in the public sector and health care



Why existing solutions are not good enough

Hadoop environment

Appliances environment







Challenges within Hadoop environment

- A great file system (and not a Data Base), but not an analytic platform!
- Not designed to solve the Big Data analytics
- Complex environment
 - Variety of technologies
 - Army of programmers with larger range of skills
 - Maintenance of the large amount of code
 - Many nodes are hard to maintain and manage
- As a result it does not meet the Big Data analytics challenge !



Challenges with common appliances (Exadata, Teradata, etc.),

- Expensive to start with
- Very costly to scale
- Compute and storage are linked
- Inflexible





Help your customers achieve more, better, faster and previously unobtainable BUSINESS and OPERATIONAL INSIGHTS

Eliminating Big Data Processing & Analytics SIZE, SPEED & COST limits



SQream Value Proposition

- Get more business insights by analyzing more data!
- Focus on business insights and not infrastructure
- Cost effective
- Much faster!
- With minimal footprint and administration



INTRODUCING SQREAM DB GPU-ACCELERATED DATA WAREHOUSE





SQream complements MPP systems

- Bridge the gap by bringing massive data analysis capabilities to traditional MPP ecosystems.
- Designed for large-scale analytics
- Based on GPU technology when scaling with CPUs is not cost-effective
- Provides a fast, cost-effective and energy efficient environment
- Boosts analytics performance through **MPP** (massively parallel computing).



SQream main advantages





Fast to get data in - Fast to get data out

- Use GPU for loading
- Over 5,000 GPU Compute core in every GPU
- 900 GB/s Memory Bandwidth
- Compress all the data by 1:5
- Collect metadata





Real Data Ingestion Examples



	rvc_raw
CSV Statistics	
Number of Files	913 files
Size of Single File	2 GB
Rows Per File	9,000,000
Single file Ingestion	5.5 sec
Total Ingestion Size	1860 GB
Ingestion Time (from External Table)	1664 sec
Table Record Count	8,825,843,568
Table Size on Disk	326 GB
Compression Ratio	1:5.7
Ingestion Rate	4 TB/hour





TPCH Ingestion Examples

Solara Vinana



Table Name	Row Count	Raw size	Time Elapsed(secon ds)	Effective Ingest Rate(GB/h, average)
PART	2,000,000,000	230 GB	324.00	2,456
SUPPLIER	100,000,000	14 GB	23.97	2,103
PARTSUPP	8,000,000,000	1.2 TB	1599.92	2,700
CUSTOMER	1,500,000,000	230 GB	360.00	2,300
NATION	25	2.2 KB	-	-
REGION	5	384 B	-	-
LINEITEM	59,999,994,267	7.2 TB	9922.60	2,612
ORDERS	15,000,000,000	1.7 TB	2343.75	2,611
Total	86.6 B	10.5 TB	4h3m	2.6 TB/h



Typical DBMS Loading



Data not ready for querying. Still need to create projections, materialized views, indexes, etc.



SQream DB Loading

SQream DB reads less data from disk, because of 1:5 compression



- SQream DB's GPU compression is automatic
- Less data read = less I/O = faster query
- Access raw data directly, without cubes, indexes







SQream minimizes time to analysis, eliminates complex data pipelines, and addresses the poor scaling of traditional solutions.







Saving hours on Queries and reports

Replacing legacy MPP with a faster, easier to use SQream GPU-powered analytics database



Cost Effectiveness

- We've decoupled storage and compute, with GPUs
- Just add more compute power (GPU) or storage
- 1U Server can have 20,000 cores!



COMPACT SIZE - HUGE POWER !

- Access with easy-to-use SQL
- SQream DB is MPP, but in a small package
- Decoupled compute and storage
- Small footprint!



■ 1U server with <a>DVIDIA GPUs can replace ~100TB of 42U legacy data warehouse racks



Minimal Scaling Impact!

- Cost
- On performance
- Hardware cost
- Complexity





Accelerating Hadoop



Hadoop - A GREAT IDEA!

- Designed to store all 'big data' on cheap, commodity hardware
- Designed to solve 'Google-scale' issues for search engines
- Deploy code and jobs, rather than moving data. Run code directly on data



... WITH SOME PAIN FOR ANALYTICS



SOREAN

• Joins very difficult

- Finnicky SQL interfaces
- Requires careful management of data

BARRIERS FOR ANALYTICS

Built for scaling, not for speed

New queries can take days to prepare



COMBINING THE BEST OF BOTH WORLDS



- Distributed, cost-effective storage
- Well established as a data lake
- Already deployed successfully
- Well-integrated



- Scalable in both compute & storage
- High-performance Ad-hoc queries
- Leverages existing SQL skills
- Easy to use



COMBINING THE BEST OF BOTH WORLDS



HIGH PERFORMANCE WITH EXCELLENT SCALABILITY

Delivering on the Hadoop promise of scaling, without the complexity





The value of sqream db



Ad-Hoc analytics at scale, enabling better business decisions, with fully-featured SQL



More and new revenue streams from the increased insights



Empowering data consumers, simplifying procedures, streamlining operations and increasing employee productivity



make the most of YOUR HADOOP INVESTMENT







Maximize Hadoop investment

Combining Hadoop with SQream DB maximizes economy of Hadoop with performance of SQream DB

Full SQL support

Fully-featured SQL support to empower data owners and business users

Fast, performant, scalable

Analyze more data for critical insights faster - with SQream DB



Why SQream?

With the power of the GPU, SQream can do in a single node what others will need a cluster for :

- Massive power in a single server (\rightarrow no need in data distribution)
- Very fast data ingestion rates with high ratio of Adaptive
 Compression while gathering metadata statistics (-> small footprint, no need in indexes/prebuild aggregation etc.)
- Very fast joins (→ full normalization, no need in pre modeling based on the final queries)





Customer Examples

• PubMatic : Offloading from Hadoop (via Spark) into CSV/Parquet and load into SQream.



PubMaticcreased Revenues by Improving Bids

PubMatic deployed SQream DB to unlock more insights from their Hadoop cluster.



- Increased bid accuracy results in increased earnings and yield
- 5 hours down to 5 minutes reduces time waiting for queries
- Leveraging existing SQL knowledge, no Scala or MapReduce needed

PubMatic creased Appetite for Data PubMatic expanded SQream DB for longer timeframes.



• The original system handled 85TB. The enhanced system handles 4.3x more data - 360TB

• 4x Dell PowerEdge servers and Qumulo filesystem-based deployment



Customer example 2

- Fraud detection use case: Data enrichments and new statistics to discover new possible deviations in purchases and possible frauds.
- Hadoop as the file system and Spark as the data processing unit.
- Basic everyday work is to add more statistics based on the existing data, in different distribution/point of view :
- → Statistics on all purchases based on time
- → Statistics on all purchases based product
- → Statistics on all purchases based on location
- → Statistics on all purchases based on payment details

For complex statistics, each will require different data modeling and very complex Spark operation.



Enhance Fraud Detection and Save Money

12TB Raw Data

	Queries in Spark	Queries in SQream
Cluster size	20 nodes Hadoop	1 node, V100 GPU
Generating statistics based on changed window functions Data period: small	34 minutes	31 minutes
Generating statistics based on changed window functions Data period: large	Never finished	65 minutes

- Better utilization of time and resources.
- Run the queries that are impossible on Spark in order to achieve unobtainable insights.



RETAIL

Monitor Competitors Customer Experience Operational Decisions



HEALTHCARE

Care Management IOT Devices Genomic Research

TELECOM

Customer 360 Competitive Analysis Network Optimization



FINANCE

Fraud analysis Risk consolidation Customized services



WHAT OUR CUSTOMERS SAY

"

"

SQream and Orange demonstrate 100x cost performance, removing limits of databases."

Pascal Déchamboux | Director of Software



AIS

SQream helps us keep pace with rapidly increasing data for real customer benefits."

Suppachai Panichayunon, Head Solution Architect



orange[™]

SQream is helping us to cut years of cancer research on large genomic datasets."

Prof. Gideon Rechavi, Head of Cancer Research



We saw a cost effective opportunity to obtain analytic capabilities we couldn't have before."

RF Group Leader

"



The leading mobile operator in Thailand 40 million subscribers



Main Challenges

- Couldn't get a holistic view of network subscribers with ad-hoc queries
- Dashboard took hours to show locations of key customers

Business Requirements

- Enable to easily analyze and troubleshoot the network quality in multiple locations in just a few clicks.
- Ability to drill down for much deeper data analysis
 - Speed Test Data in different locations to manage the quality of the network.
 - Decrease potential of competitor targeting



UNDERSTAND 40 MILLION CUSTOMERS

AIS

TELECOM





40 NODES 5 full racks 7600 CPU cores



HP DL380g9 with NVIDIA Tesla GPU 96 GB RAM + 6 TB storage





PROCESS TAKES A REALLY LONG TIME





Saving hours on Queries and reports

Replacing legacy MPP with a faster, easier to use SQream GPU-powered analytics database



46 seconds with SQream DB





Architecture before Sqream DB





Simplified with SQream DB







Orange France Telecom





orange[™]







The cost of performance

Benchmark performed by the operator

Analyze subscriber network usage - 4.3 billion call records, 30 million subscribers





PubMatic INCREASE REVENUES

AD-TECH

85 TB/day in ad impressions for constructing bidding histograms



PubMatic INCREASE REVENUES

AD-TECH

360 TB/day ingested to enhance bid histogram accuracy



Who is it for?

- BI Analysts
 - Need to analyze large amounts of data
 - Require fast response times to reports, dashboards
- Data Scientists
 - Easy connectivity with lots of different tools, including Python/Jupyter
 - Very straight-forward data exploration, with no need to pre-index the raw data
- Any company
 - With lots of data and difficult scalability of current solutions
 - Need to build a high performance data pipeline without spending millions of \$



SQream Integration with End to End Pipelines

- Connectivity ODBC, JDBC and specific connectors e.g. Spark, Tableau
- Can also use "External table" to query data from source



CONTACT

Rafi Glickman, APAC Sales Director



