

Tech Overview

Q-Sensei

ADVANCING ANALYTICS

Ute Rother, CEO and Co-founder

September 2021

“Data is eating the World”

Challenge

Massive amounts of structured and unstructured data drive costs for data aggregation, cleansing, tagging, storage, and querying.



Modern Analytics for Everyone

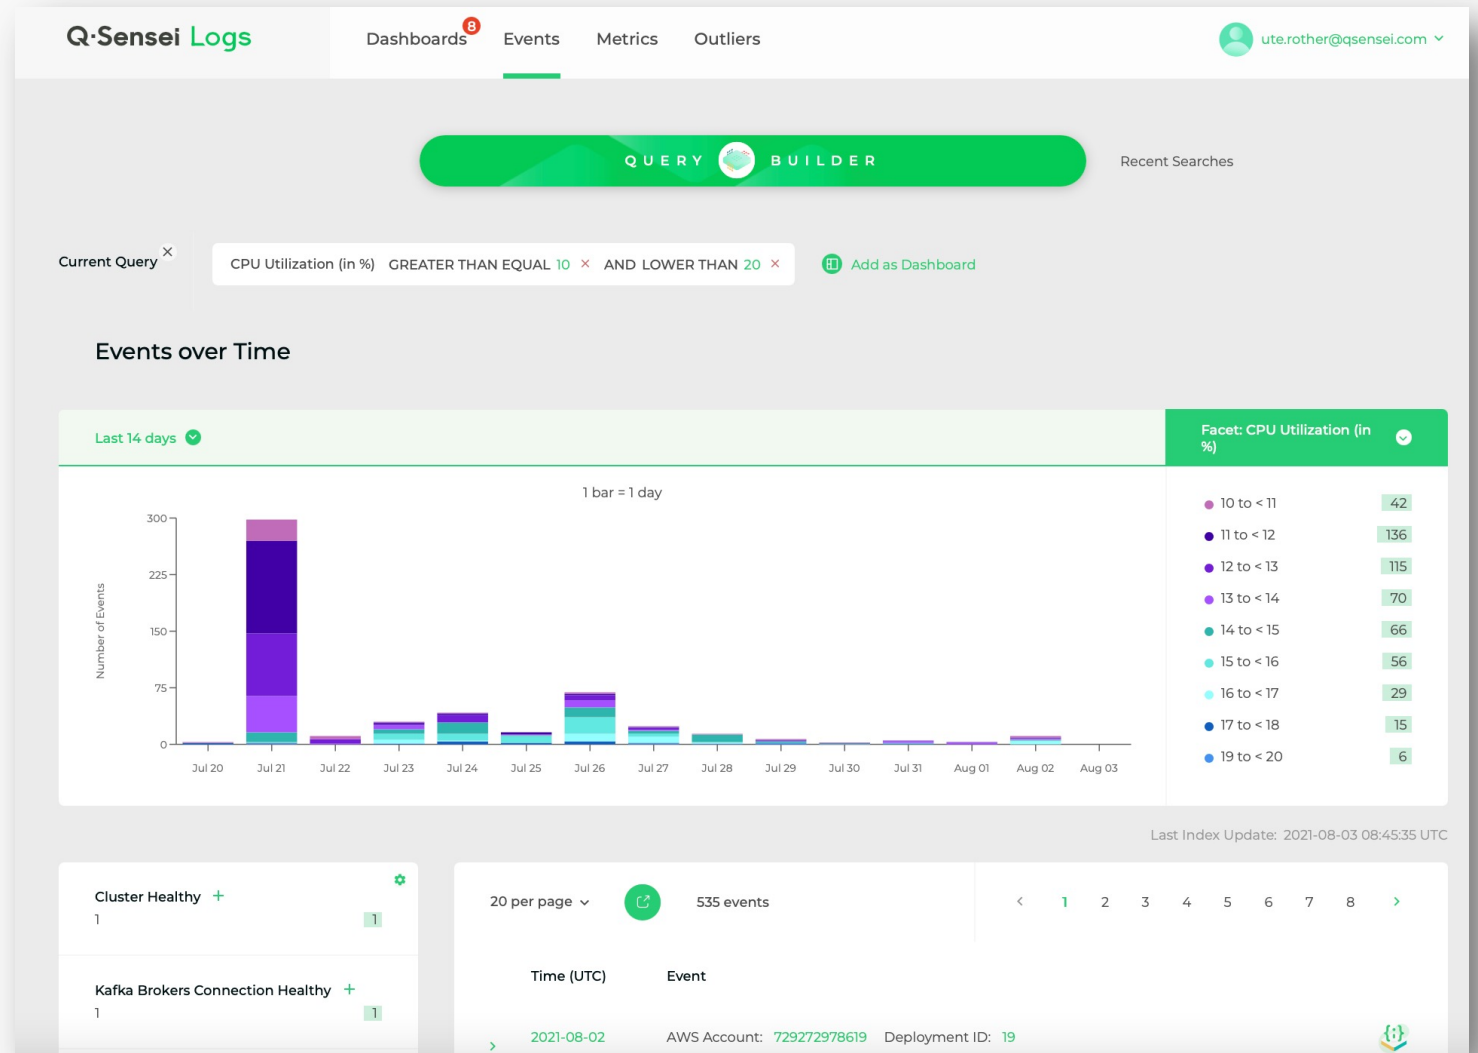
Q-Sensei Logs

Intuitive User Interface

For next-gen analytics featuring
Query Builder, Dashboards,
Advanced Log Aggregation and more

True Human-Data-Interaction

Accelerated processing of thousands of
indexes and billions of data points



Accelerating Analytics

Q-Sensei Fuse

Modular data platform

For data aggregation, cleansing,
and tagging

In-memory Index

For high-speed querying of data using
low-cost commodity hardware

Comprehensive API Layer

For custom applications and AI/ML

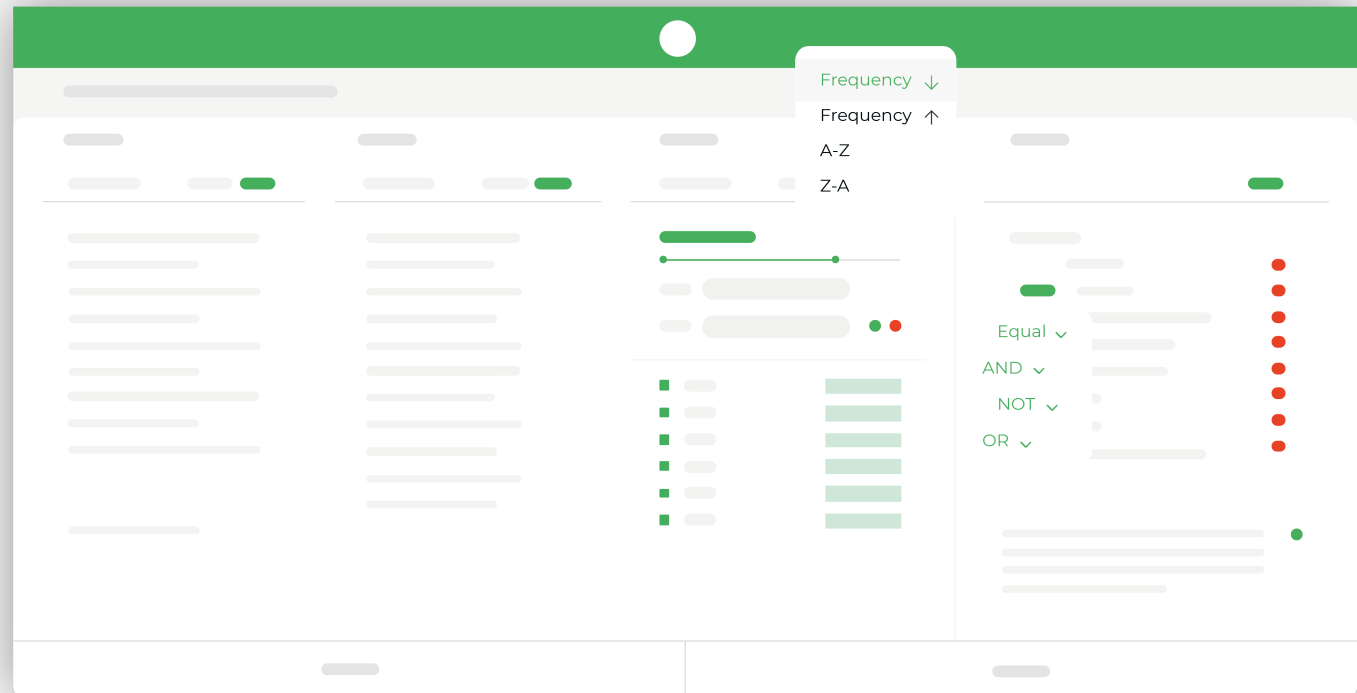
The screenshot shows the API Reference page for 'Combining Search Clauses'. The breadcrumb trail is 'API Reference > Queries > Combining Search Clauses'. The left sidebar contains a navigation menu with sections: FUSE GUIDE (Tutorials, Installing Fuse, Configuration, Schema, Adding Data, Querying Fuse), API REFERENCE (Queries, Facets, Contents, Search), and 'Combining Search Clauses' is highlighted under the 'Queries' section. The main content area shows a 'GET' method for the endpoint `/queries?q={string:q}&part={string:part1}&part={string:part2}`. Below this is a table of 'Query Parameters' and a 'Response JSON Object' table.

Parameter	Type	Required	Description	More
q	string	required	First query.	
part	string	optional	Additional query to add.	>
op	string	optional	Query combining strategy.	>
sort	boolean	optional	Sort query parts.	>

Response Field	Type	Required	Description	More
expression	object	optional	Current query expression.	
negate	object	optional	GET /queries with the negative of	
parsed	object	optional	Parsed tree of the current query.	>

DEMO

Hyper-precise Querying



<https://www.qsensei.com/q-sensei-logs-demo>

Breakthrough in Interactive Analytics

10,000+

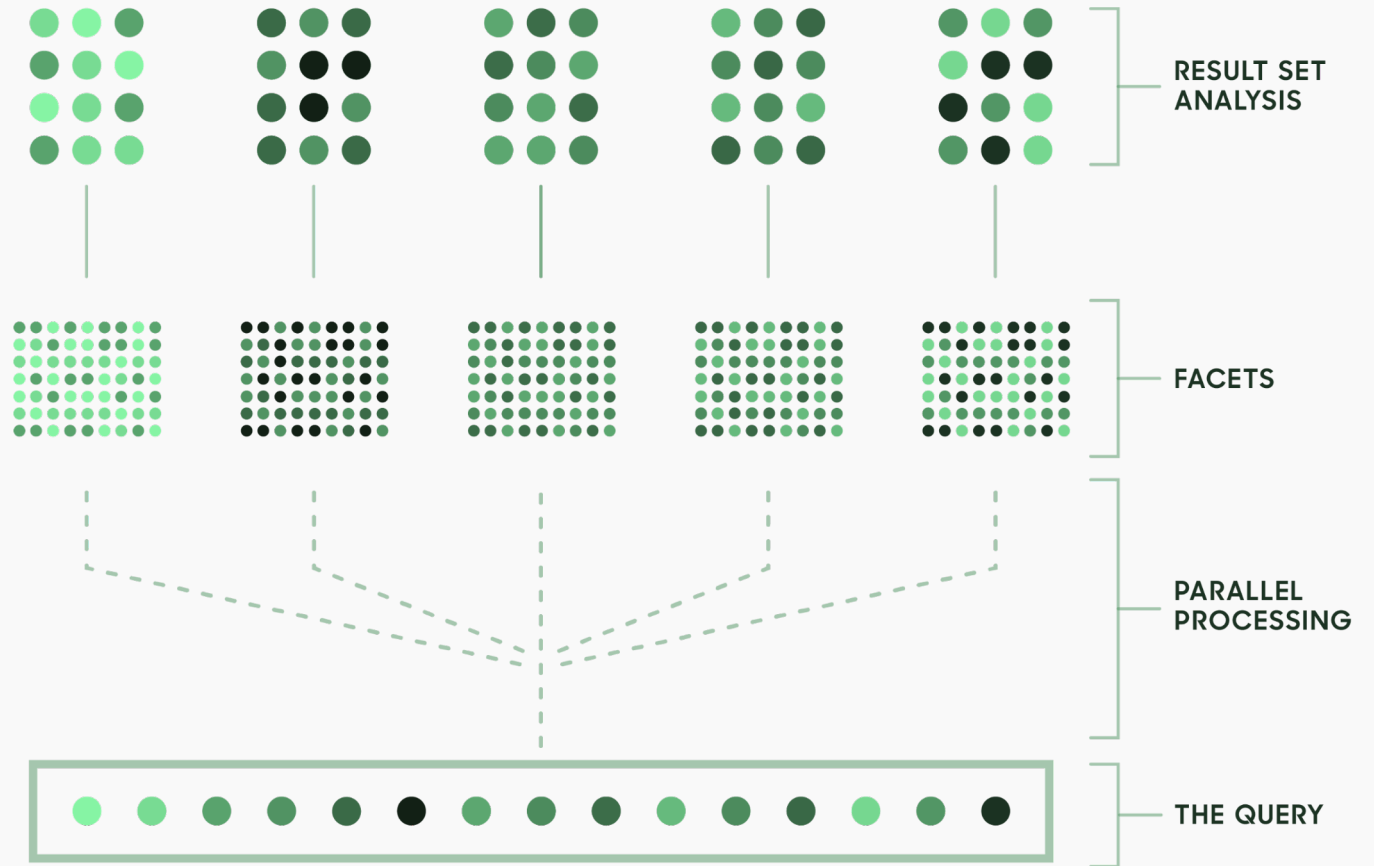
Indexes for
Human Data Interaction,
ML, and AI

1,000x

Boost in
Algorithmic Querying

450%

Higher
Relevancy of Results



Next-Gen Speed Using Low-Cost Commodity Hardware

Q-Sensei has the best price-performance ratio in the market:
 On a billion-row taxi data set, **Q-Sensei's PPR is up to 37,000 times better.**

Setup	Query 1	Query 2	Query 3	Query 4	Query Compute Costs
Q-Sensei Fuse 6.5.1 1 x c5.12xlarge	0.001	0.130	0.123	0.576	\$2.04
Q-Sensei Fuse 6.5.1 1 x c5.9xlarge	0.001	0.176	0.166	0.587	\$1.53
BrytlytDB 2.0 & 2-node p2.16xlarge cluster	0.009	0.027	0.287	0.428	\$28.80
OmniSci & 2-node p2.8xlarge cluster	0.034	0.061	0.178	0.498	\$14.40
ClickHouse, 3 x c5d.9xlarge cluster	0.241	0.826	1.209	1.781	\$5.18
Redshift, 6-node ds2.8xlarge cluster*	1.560	1.250	2.250	2.970	\$40.80
Spark 2.4 & 21 x m3.xlarge HDFS cluster*	2.362	3.559	4.019	20.412	\$5.59
Presto 0.214 & 21 x m3.xlarge HDFS cluster*	3.540	6.290	7.660	11.920	\$5.59
Elasticsearch (heavily tuned)*	8.100	18.180	N/A	N/A	\$0.50
Vertica, Intel Core i5 4670K	14.389	32.148	33.448	67.312	\$0.50
Elasticsearch (lightly tuned)*	34.480	63.300	N/A	N/A	\$0.50
PostgreSQL 9.5 & cstore_fdw*	152.000	175.000	235.000	368.000	\$0.50

Q-Sensei: Comparative Analysis					
Vendor	Setup	Price-Performance Ratio			
		Query 1	Query 2	Query 3	Query 4
Q-Sensei	Q-Sensei Fuse 6.5.1 1 x c5.12xlarge	1x	1x	1x	1x
Q-Sensei	Q-Sensei Fuse 6.5.1 1 x c5.9xlarge	0.75x	1.02x	1.01x	0.76x
Brytlyt	BrytlytDB 2.0 & 2-node p2.16xlarge cluster	127x	3x	33x	10x
OmniSci	OmniSci & 2-node p2.8xlarge cluster	240x	3x	10x	6x
Open Source	ClickHouse, 3 x c5d.9xlarge cluster	612x	16x	25x	8x
Elastic	Elasticsearch (heavily tuned)*	1,985x	34x	N/A	N/A
Vertica	Vertica, Intel Core i5 4670K	3,527x	61x	67x	29x
Apache	Spark 2.4 & 21 x m3.xlarge HDFS cluster*	6,472x	75x	90x	97x
Elastic	Elasticsearch (lightly tuned)*	8,451x	119x	N/A	N/A
Open Source	Presto 0.214 & 21 x m3.xlarge HDFS cluster*	9,700x	133x	171x	57x
Amazon	Redshift, 6-node ds2.8xlarge cluster*	31,200x	192x	366x	103x
Open Source	PostgreSQL 9.5 & cstore_fdw*	37,255x	330x	468x	157x

Sources: <https://tech.marksblogg.com/benchmarks.html>
 Q-Sensei performance test 08/2021
 Queries are shown in appendix
 Query speed is measured in seconds
 *Most often used

Queries Used in Benchmark Analysis

Query 1 `SELECT cab_type, count(*) FROM trips GROUP BY cab_type;`

Query 2 `SELECT passenger_count, avg(total_amount) FROM trips GROUP BY passenger_count;`

Query 3 `SELECT passenger_count, extract(year from pickup_datetime) AS pickup_year, count(*) FROM trips
GROUP BY passenger_count, pickup_year;`

Query 4 `SELECT passenger_count, extract(year from pickup_datetime) AS pickup_year, cast(trip_distance as int)
AS distance, count(*) AS the_count FROM trips GROUP BY passenger_count, pickup_year, distance
ORDER BY pickup_year, the_count desc;`

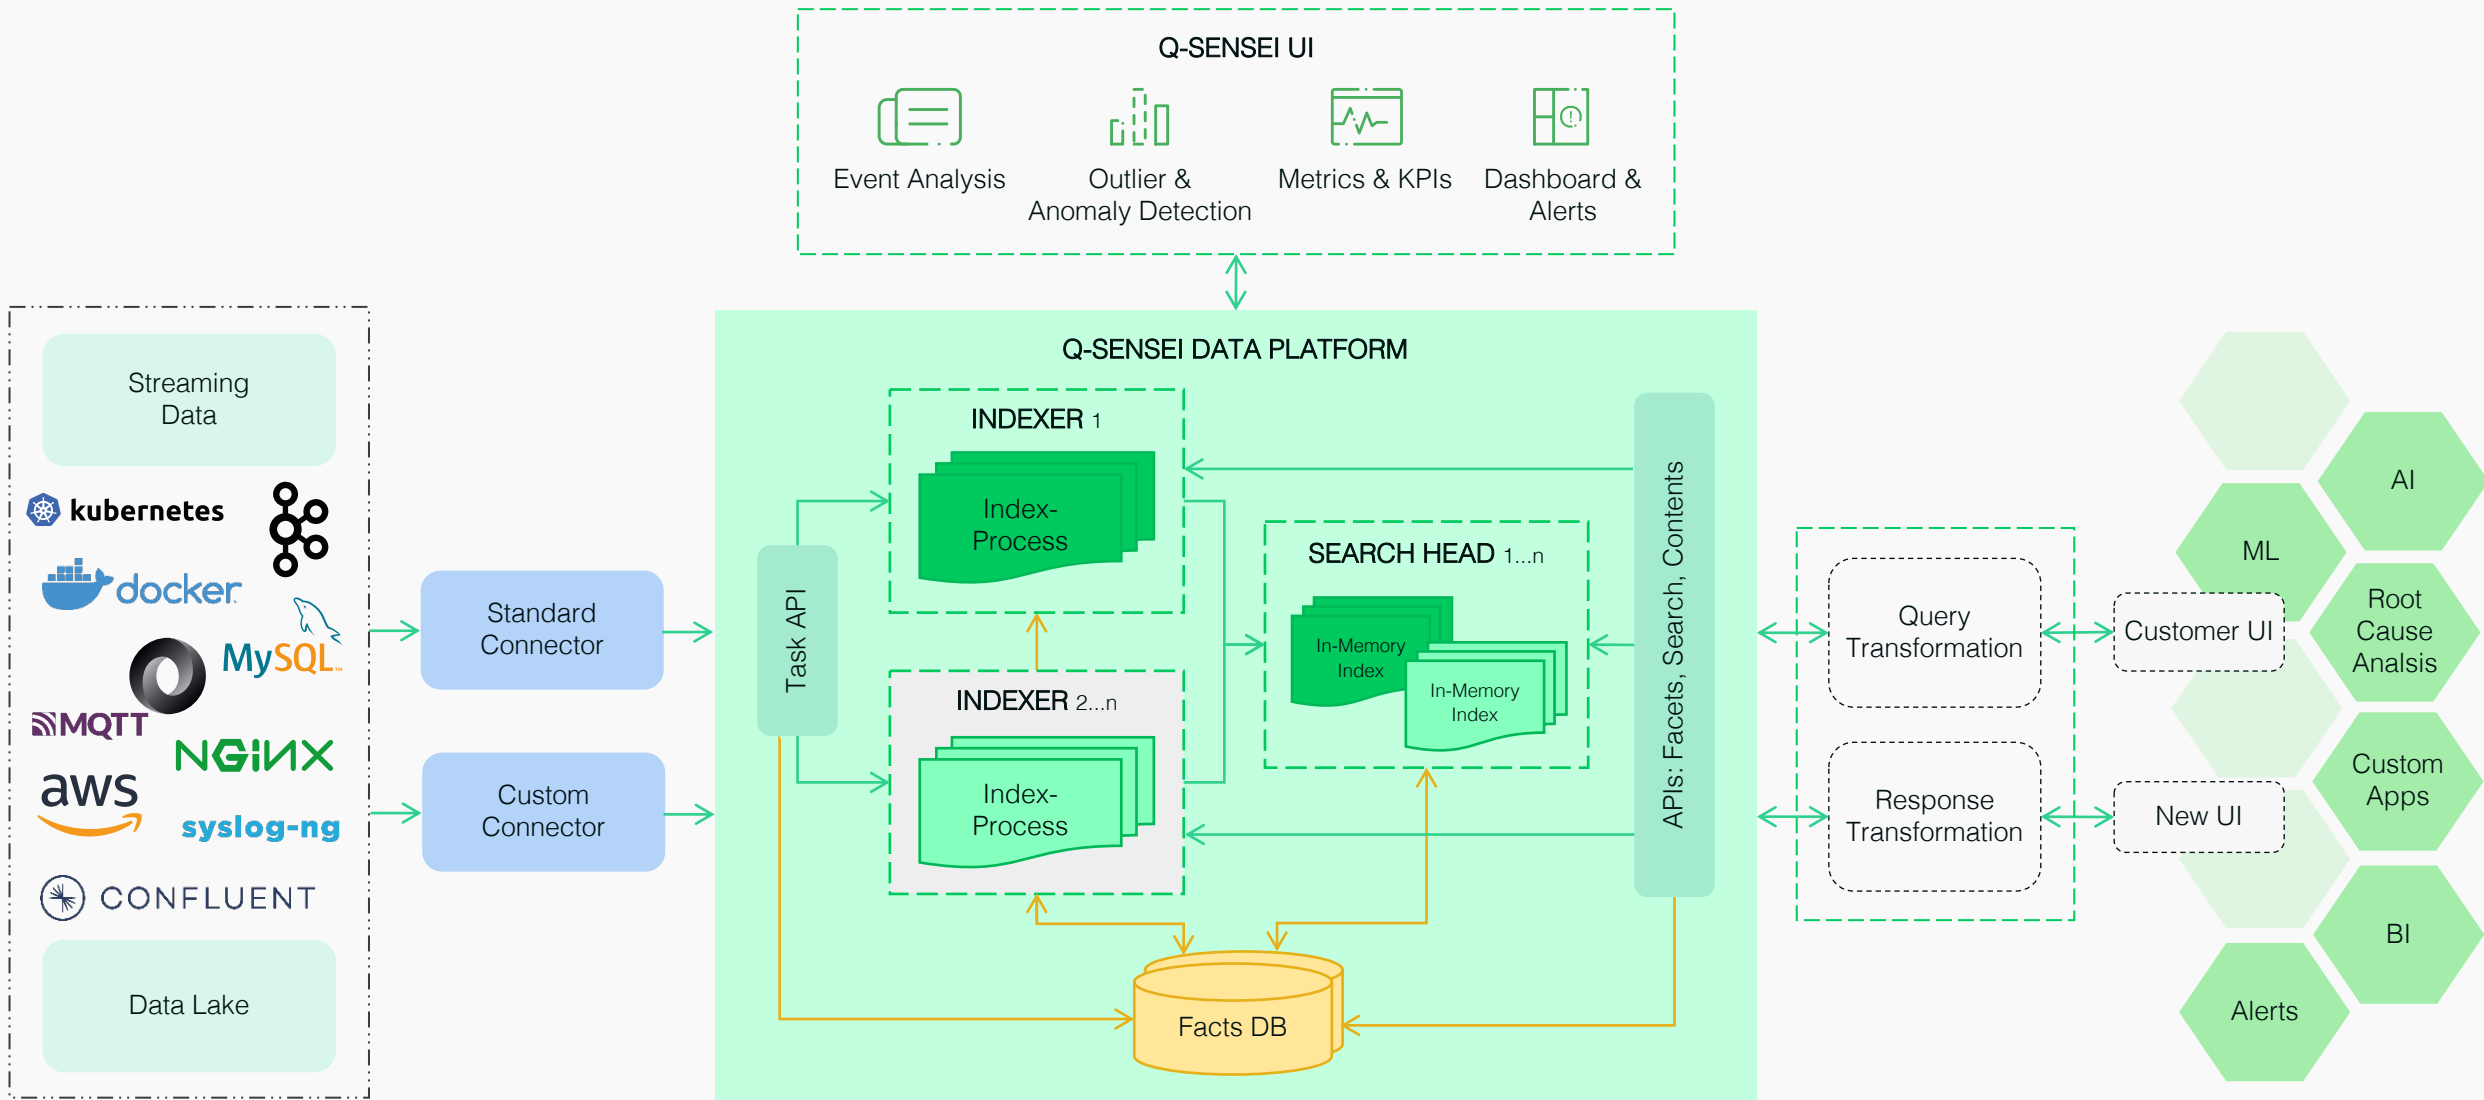
Comprehensive Set of Analytical Features

Q-Sensei Index



Protected by US Patents 7,080,059 – 7,680,777 – 9,690,824

Modular Platform



Use Cases




Product Development

Test Data Monitoring
KPI Comparison across Releases
Product Validation



Healthcare

Clinical Trial Analysis
ICU Monitoring
Patient Data Analysis



Telecommunications

Network Monitoring
Threat Analysis and Prevention
Bandwidth Optimization



Media

Audience Measurement
Ad Optimization
Predictive Programming



Financial Services

Real-time Cost Monitoring
Sales Analysis
Controlling and Auditing



Transportation

Real-time fleet monitoring
Optimized insurance packages
Smart cities

Company Snapshot

Management



Ute Rother
CEO, Founder



Wolfram Kerber
CTO, Founder



Richard Nottenburg
Chairman

Products

Q-Sensei Finance

Available in AWS Marketplace

Q-Sensei Logs

Available in AWS Marketplace

Q-Sensei Enterprise

Available in AWS Marketplace
and on-prem

Business

Founded 2007

US Patents

Search and Presentation Engine
7,080,059 – 7,680,777 – 9,690,824

Locations

USA, Germany

Industry Awards

FROST & SULLIVAN

