Loopy∞

FRICTIONLESS DATABASE COST AND PERFORMANCE OPTIMIZATIONS

CONTACT: MARK.MOKRYN@GMAIL.COM

Founding team

MARK MOKRYN, CEO

30+ years of experience in product definition, strategy, and development

Key roles in both large corporates as well as strong startup experience

Founder of SANgate Systems (renamed Sepaton), acquired by HDS

Experienced in startup fundraising and M&As

DR. RAMI COHEN, CTO

25+ years experience in leading research, architecture activities, and R&D teams

Ph.D. in Computer Science from the Technion – Israel Institute of Technology

Numerous leading roles in startup companies, including VP of R&D in a cybersecurity firm
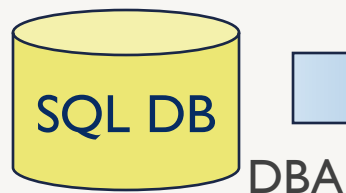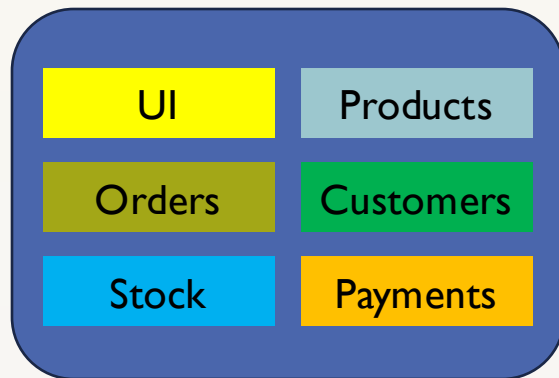
Data management crisis in modern apps

Cloud architecture drives up storage costs due to lack of centralized DB management

Loopy cuts database costs by over 50% by restoring control over data
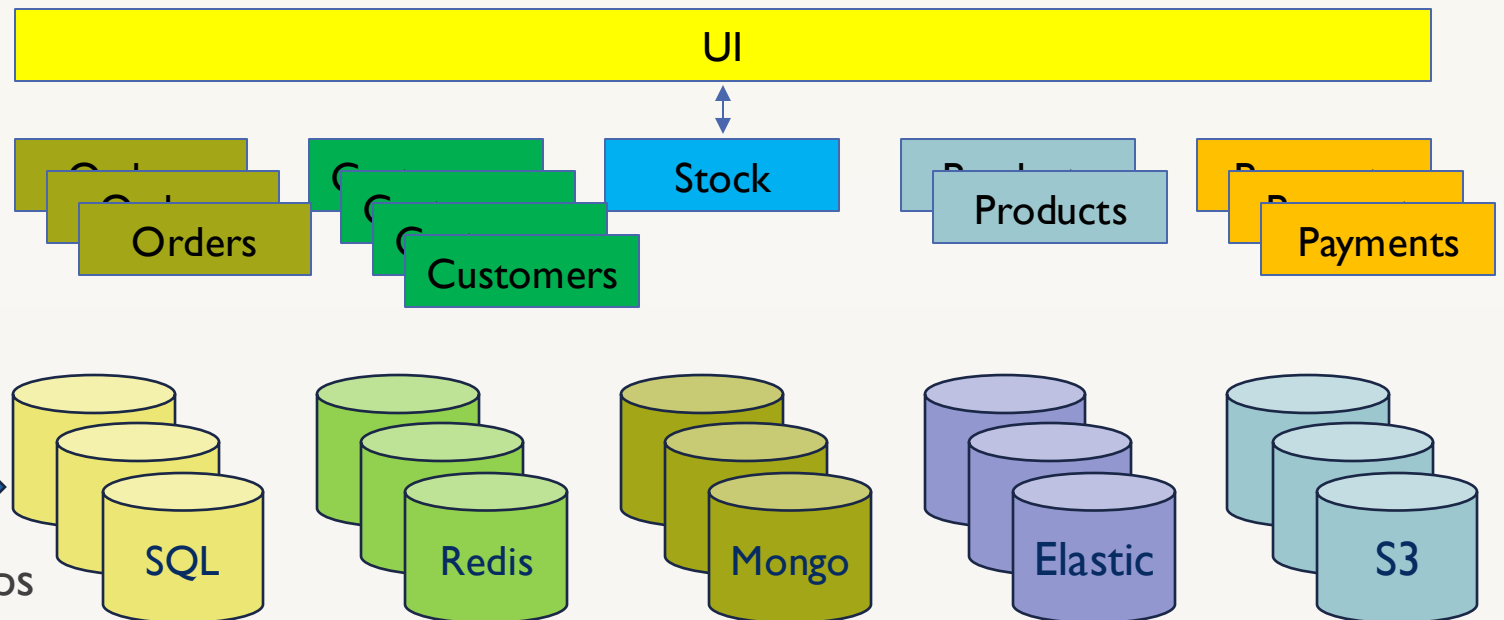
Scale-up app and storage

Scale-out microservices and polyglot storage

UI | Products
Orders | Customers
Stock | Payments

UI

Orders
Orders
Customers
Customers
Stock
Products
Payments

SQL DB
DBA
15 years
Dev-ops
SQL
Redis
Mongo
Elastic
S3

The Loopy solution

Human or AI agent supervision

Loopy AI analytics engine

Loopy data path

kibana

Listen

Learn

Optimize

Transparently sample ingestions and queries

Out-of-band analytics engine for actionable info

Transformations on in-flight data

AEROSPIKE

amazon S3

elastic

mongoDB

redis

Go To Market

Target Markets 1

- No business-critical data → APM (logs, metrics) and RAG
- $10B+ TAM
- Popular DB with a vibrant easy-to-reach base: Elasticsearch

Beachhead 2

- Opensource for dev-ops and infrastructure managers
- Zero software integration
- Instant value: insights on ingested and queried data

Business 3

- Customer converts to paid version when she sees the potential $ saved
- Part of the DB budget
- Immediate insights into savings

Stickiness 4

- Removing Loopy removes the savings
- Custom API integrations when firmly established

Listen Learn Optimize
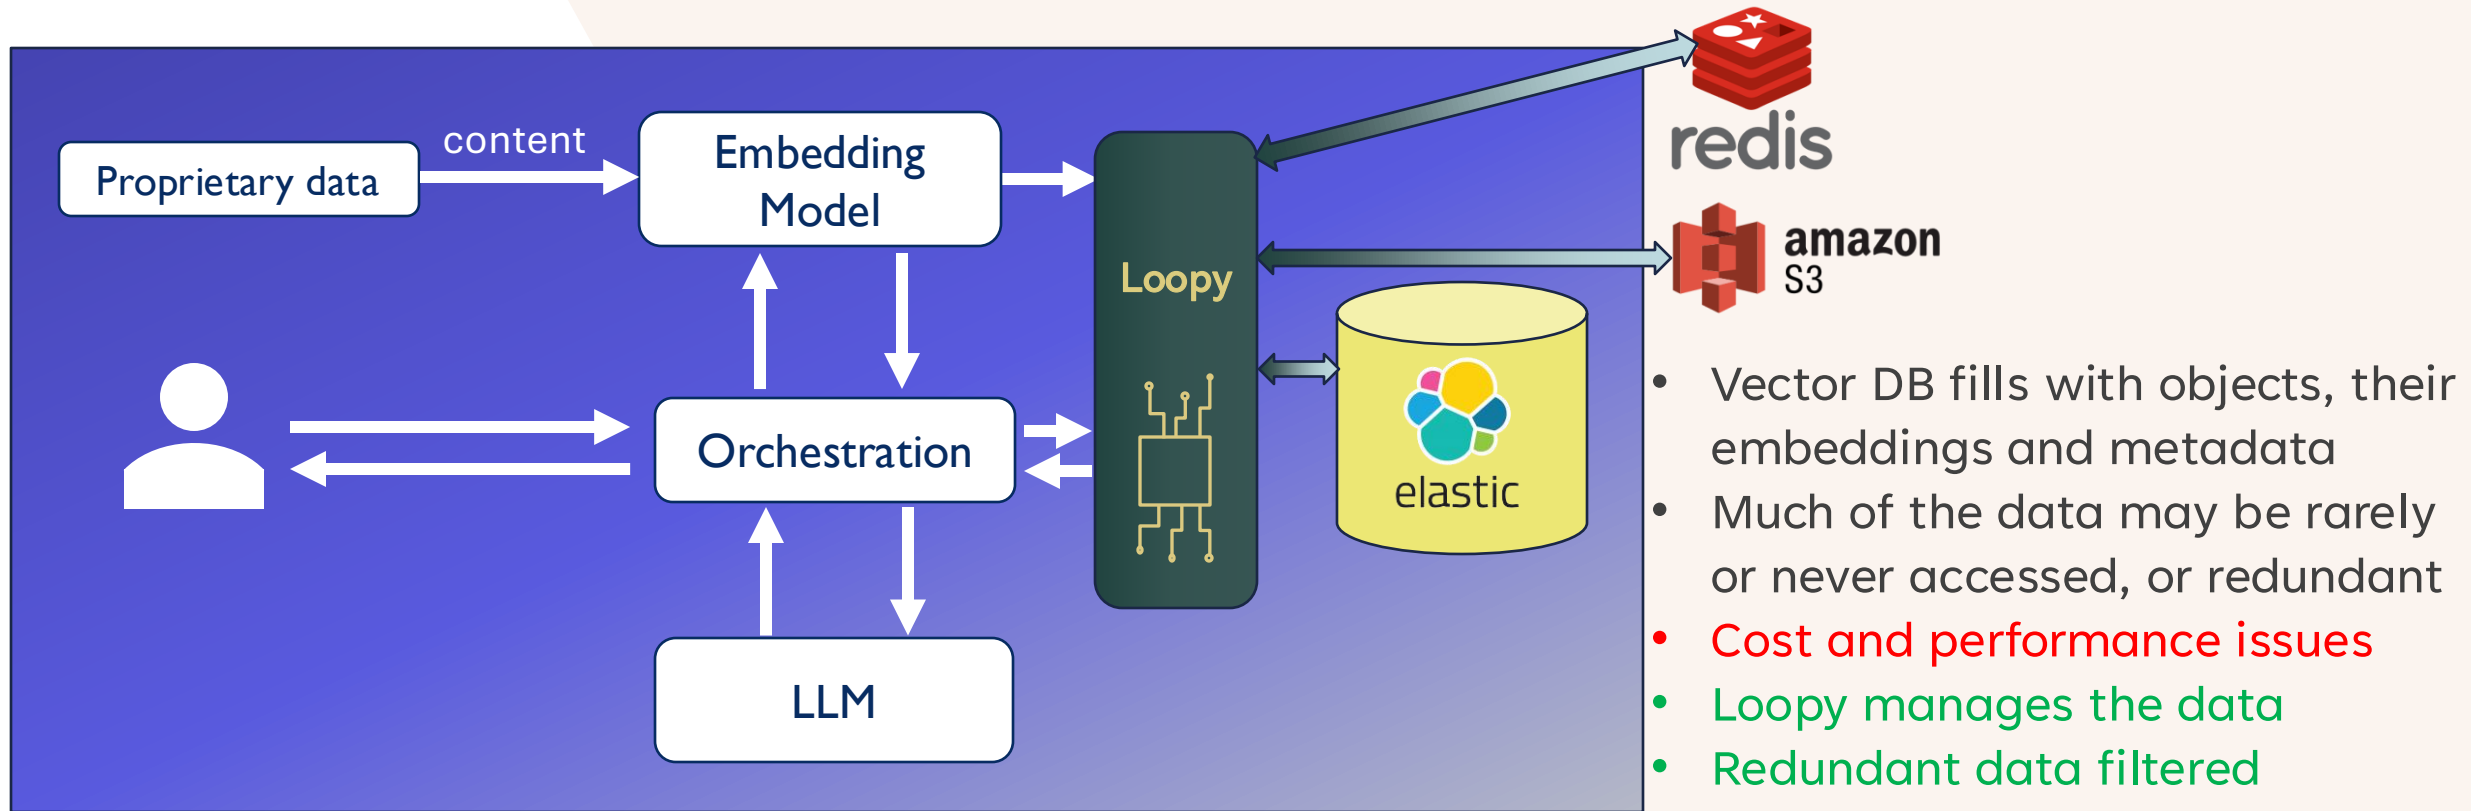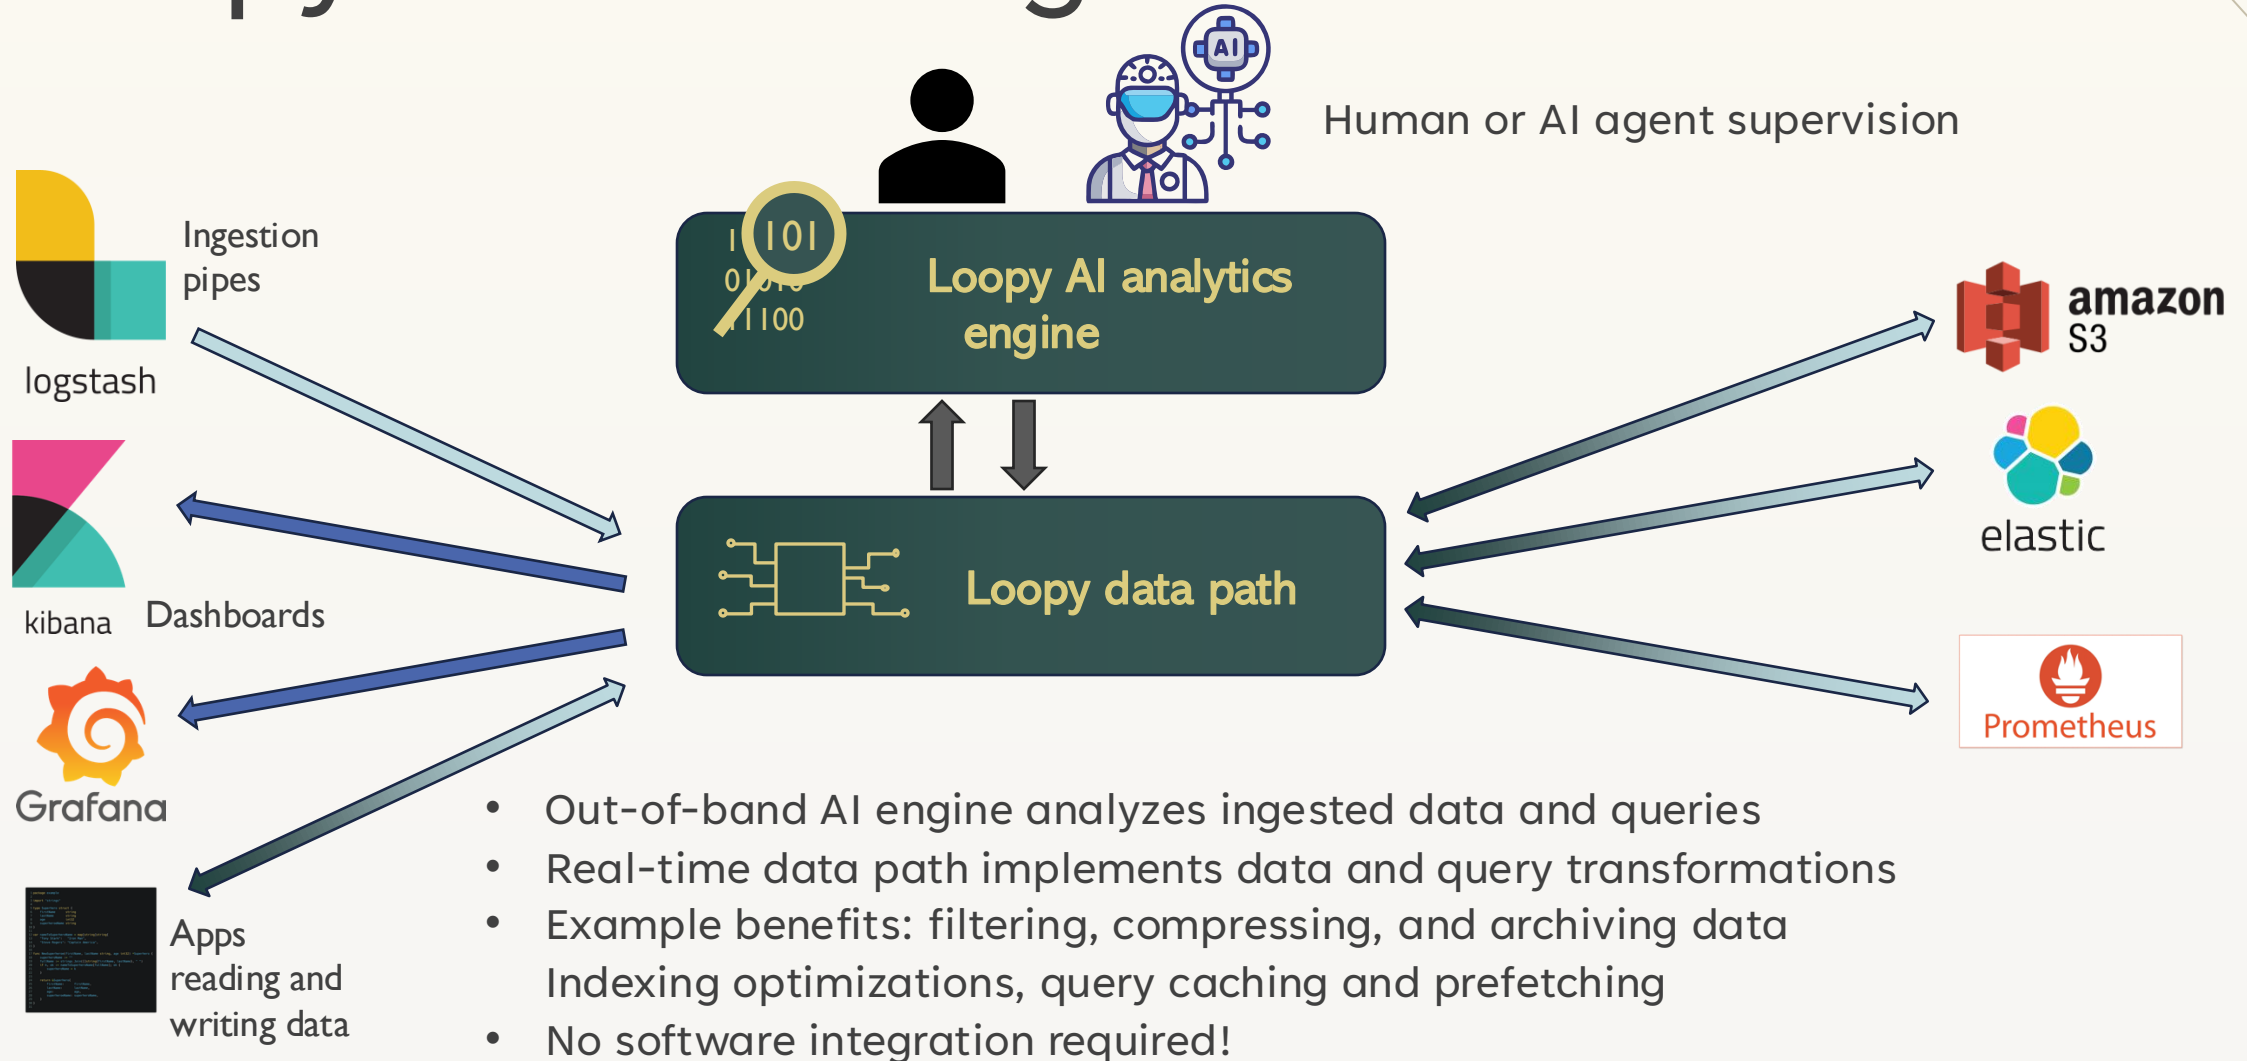
Optimizing data for vector DBs for RAG

Loopy applies optimizations on the data
Zero software integration by the customer

- Vector DB fills with objects, their embeddings and metadata
- Much of the data may be rarely or never accessed, or redundant
- Cost and performance issues
- Loopy manages the data
- Redundant data filtered
- Cold data moved to S3
- Hot data cached in Redis
- Transparent to the customer!

Loopy for APM logs

Human or AI agent supervision

Loopy AI analytics engine

Loopy data path

Ingestion pipes

logstash

Dashboards

kibana

Grafana

Apps reading and writing data

amazon S3

elastic

Prometheus

- Out-of-band AI engine analyzes ingested data and queries
- Real-time data path implements data and query transformations
- Example benefits: filtering, compressing, and archiving data Indexing optimizations, query caching and prefetching
- No software integration required!

Business Model:

OpenSearch pricing	$542/month	per node price
Customer Cost **without loopy**	**$54,200/month**	small 100-node deployment
Loopy cost	$5420/month	We charge 10% of the original cluster
Customer DB direct cost reduction*	$21680/month	Conservative assumption of 60% reduction**
Total customer cost **with Loopy**	**$27,100/month**	Removing Loopy removes the savings

* Plus indirect cost reductions such as DB management efforts

** Initial demos indicate that the reduction easily reach over 70%

Initial products address $10B+

Vector DB Market (RAG)
$1.5B 2023 → $4.3B 2028 (WSJ May 21 2024)
Elastic the current leader in this space
Pinecone, Milvus, Weaviate, Qdrant will be
supported based on customer priorities

SIEM logging space
Relevant once achieving customer trust
$3B for Elastic, OpenSearch, and infrastructure

**Logging and RAG DBs are the initial markets for first 3 years
The eventual target is operational savings and optimizations for
the overall NoSQL market which is well over $100B by 2030, and
far more when including infrastructure costs**

APM Elastic/OpenSearch market: $7B+
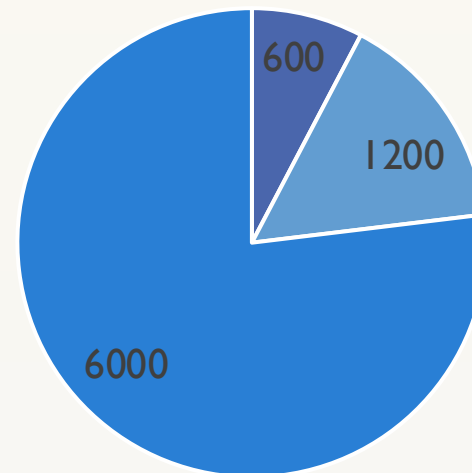
- Elastic observability revenue
- AWS OpenSearch obs. Rev.
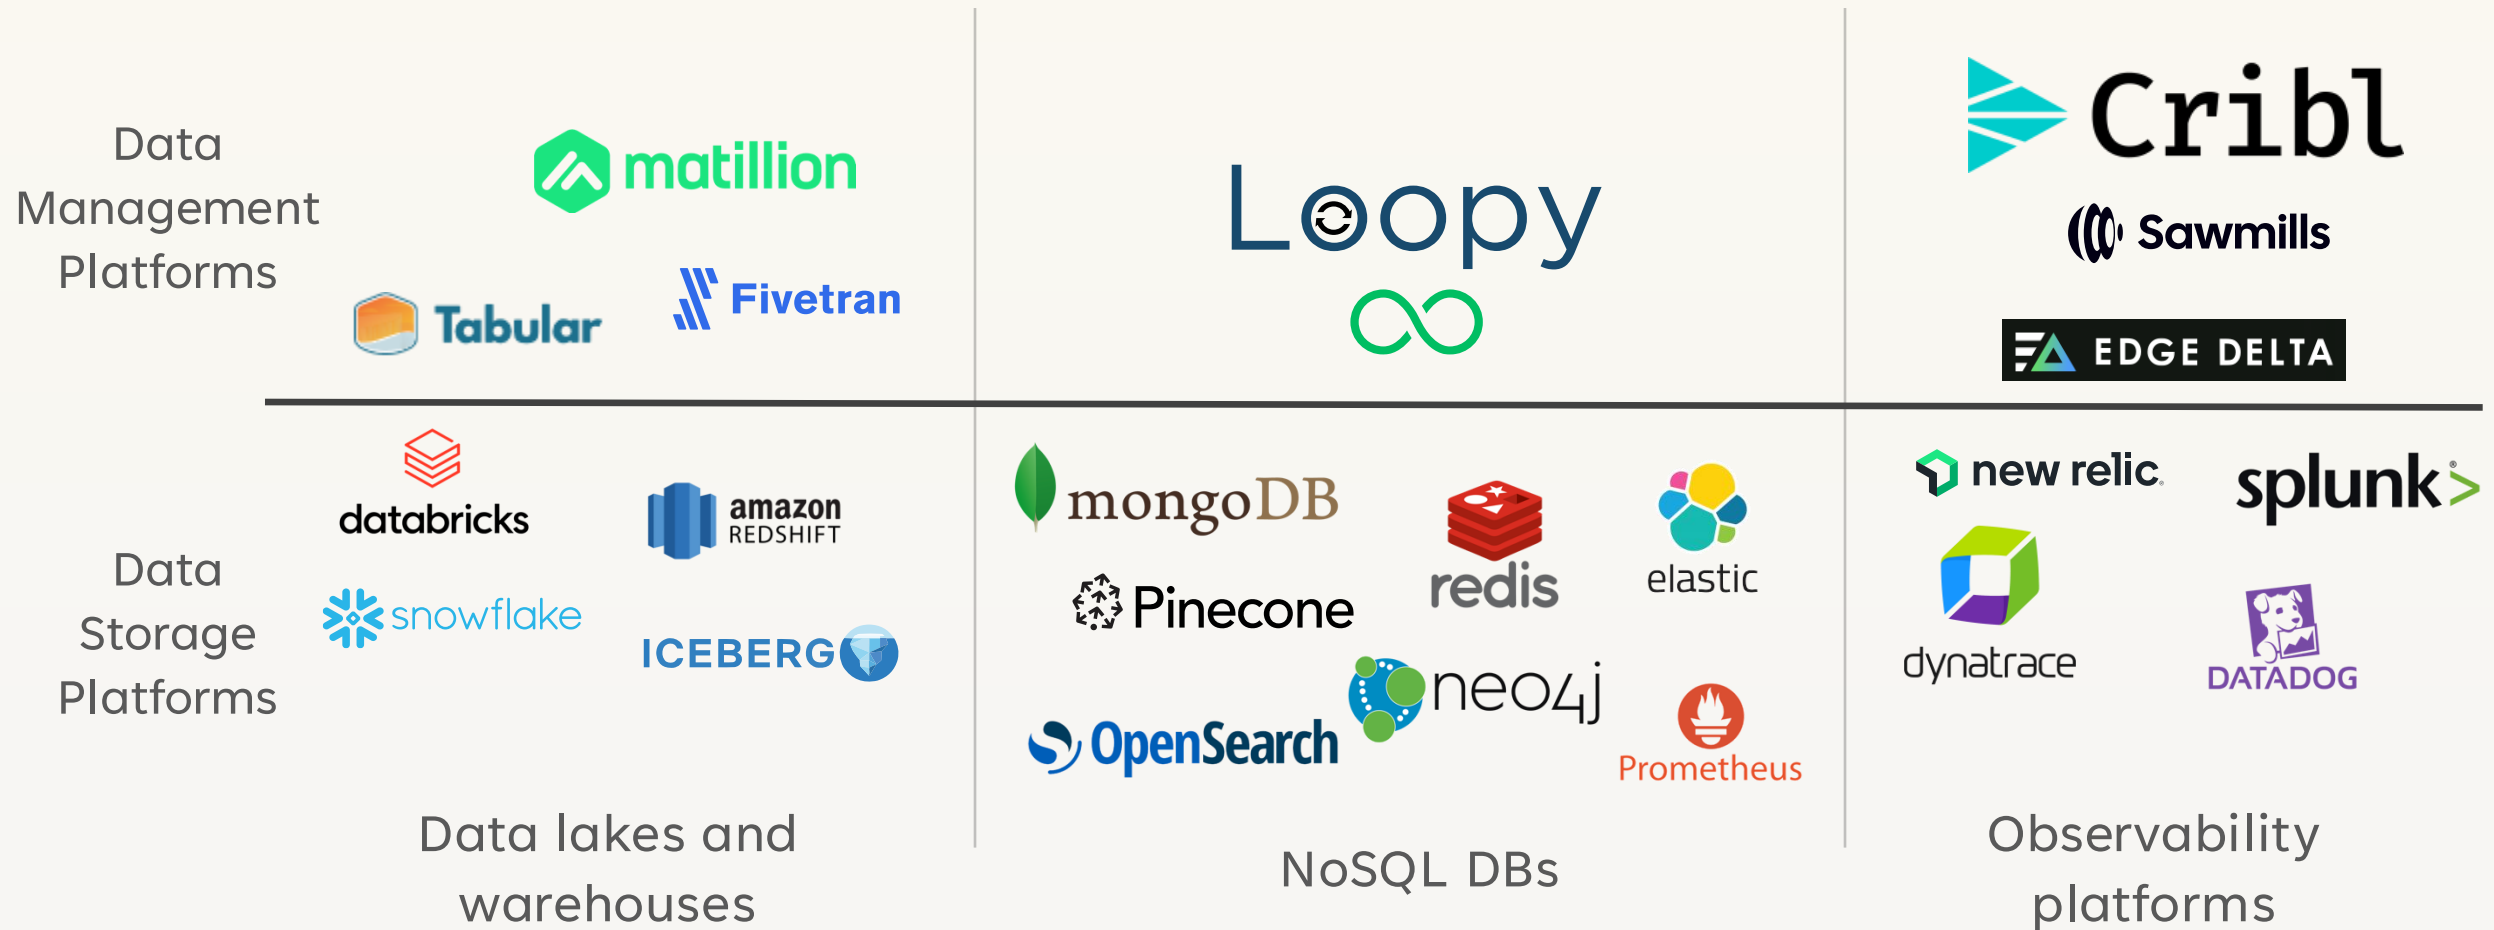- Infrastructure for licensed and OSS

ICP: Elastic's 1400+ customers >$100K
Elastic: $1.5B run rate and 40% to APM
AWS OpenSearch estimated 2x of Elastic

Main customer and partner interactions

Customer	Main feedback
Redis	Spoke with VP of AI products: "When can I speak to customers about Loopy?" Using Loopy to integrate Redis vector cache into customer deployments with zero customer integration required
Spot	Discussions with their GM. They moved from self managed Elastic to Logz.io due to incapability to manage the logs. The are paying $800k a year to Logz.io. If using Loopy this cost can be easily cut to $150k
Bank Leumi	Spoke with the team managing Elasticsearch for bank's cybersecurity Confirmed the need and lack of available solutions Willing to be a design partner and even contribure to open source product
Verizon	Discussions with network and infrastructure architects, managing PB-scale logging systems Corporate mandate to eliminate or reduce Splunk usage due to cost, and move to Elastic "What you propose is not available in the ecosystem and is definitely of value to us" "No one enables analytics on our logging patterns to help us save cost" Very interested in working with us moving forward
Mobileye	Discussions with cloud infrastructure team, managing 100s TB/day Also moving from Splunk to Elastic "We need a solution that spans cloud and on-prem deployments" Very interested in use Loopy once a it will be ready.
Proofpoint	Discussions with Group VP managing 100 developers and $1M/mo AWS bill, 10% on Cloudwatch Logs Would switch to Elastic or OpenSearch if we could cut 50%+ on database costs
Cybereason	Discussions with dev ops lead Interested in managing logs in Elastic BUT more interest in managing 2,200 Elastic servers for production data! Upcoming discussions with architecture team

Competitive landscape

Data Management Platforms

matillion

Tabular

Fivetran

Loopy

Cribl

Sawmills

EDGE DELTA

Data Storage Platforms

databricks

amazon REDSHIFT

snowflake

ICEBERG

mongoDB

Pinecone

redis

elastic

OpenSearch

neo4j

Prometheus

new relic

splunk

dynatrace

DATADOG

Data lakes and warehouses

NoSQL DBs

Observability platforms

Product plan

Product	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8
Analytics for Elastic logging (# ppl)	5	5	5					
Active cost reduction for Elastic (#ppl)			1	3	3	4	4	4
RAG cost reduction for Elastic vector DB (#ppl)				1	5	5	5	5
SaaS (#ppl)				2	2	2	2	3
Total headcount	5	5	6	6	10	11	11	12
Pre-seed phase ($1.5M)								
Seed phase								

Product color denotes version availability