One Platform, Many Workloads: Powering AI Applications with OceanBase on Kubernetes

Peng Wang

Global Technical Evangelist @OceanBase

2025/08/05

Contents

AI Workloads Are Changing the Game

IDC

The proportion of unstructured data globally reached 92.9% in 2023.

41% of executives see RAG architecture as essential, and 81% of IT leaders believe GenAI models using their own data offer a key competitive edge.

CLOUDERA

Enterprises are embracing agentic AI at scale — 66% are using enterprise AI infrastructure platforms, and 60% are embedding agent capabilities directly into their core applications.

- AI workloads (RAG, Q&A, semantic search, log analytics...) are becoming mainstream
- These workloads rely heavily on vector search, metadata, and hybrid queries
- Platform teams now need to support structured + unstructured data

The Stack Is Getting Complicated

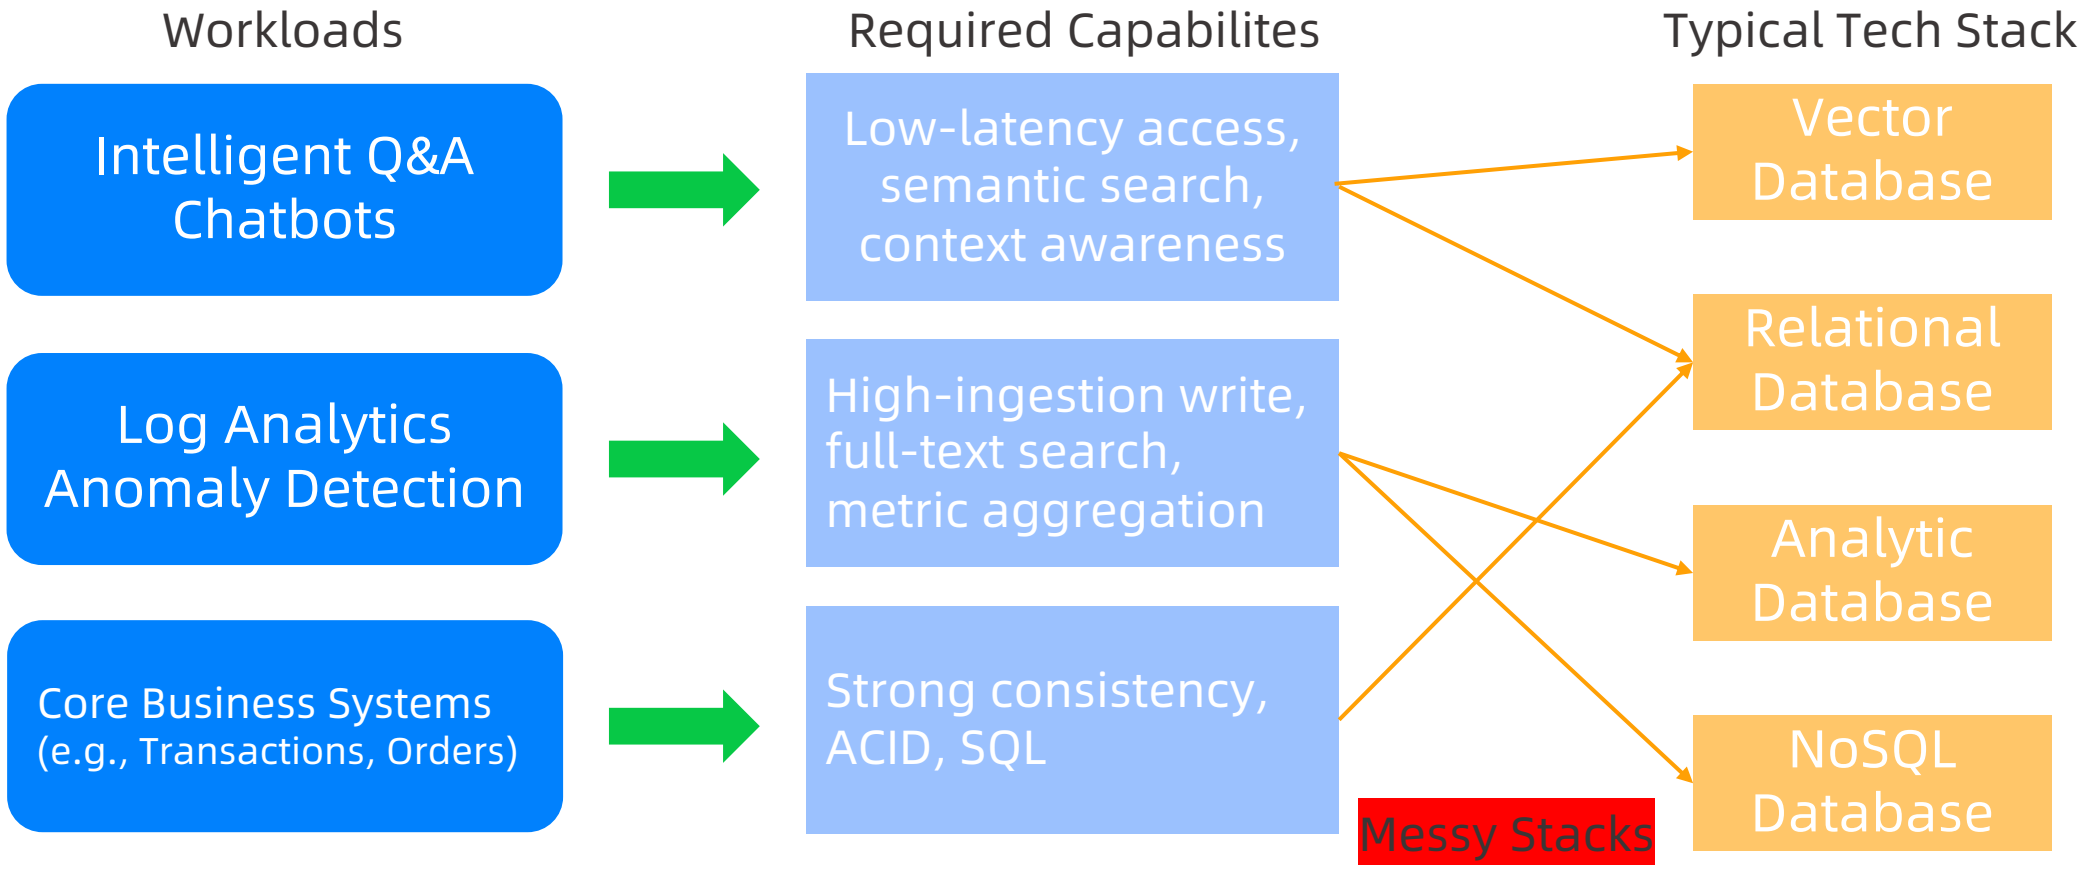

Supporting modern AI workloads requires stitching together fragmented infrastructure, which slows delivery and increases complexity.

What Platform Engineers Want (But Rarely Get)

They want ...	But they get ...
One platform for all data workloads	Fragmented data systems (RDBMS + vector DB)
Strong consistency & high availability	Trade-offs between performance and consistency
Seamless integration with AI pipelines	Complex data movement between services
Elastic scaling on Kubernetes	Difficult scaling across stateful workloads
Unified access (structured + unstructured data)	Inconsistent query models & APIs
Multi-tenancy with resource isolation	Tenant interference and noisy neighbors

AI-Powered Knowledge Retrieval for Internal Teams

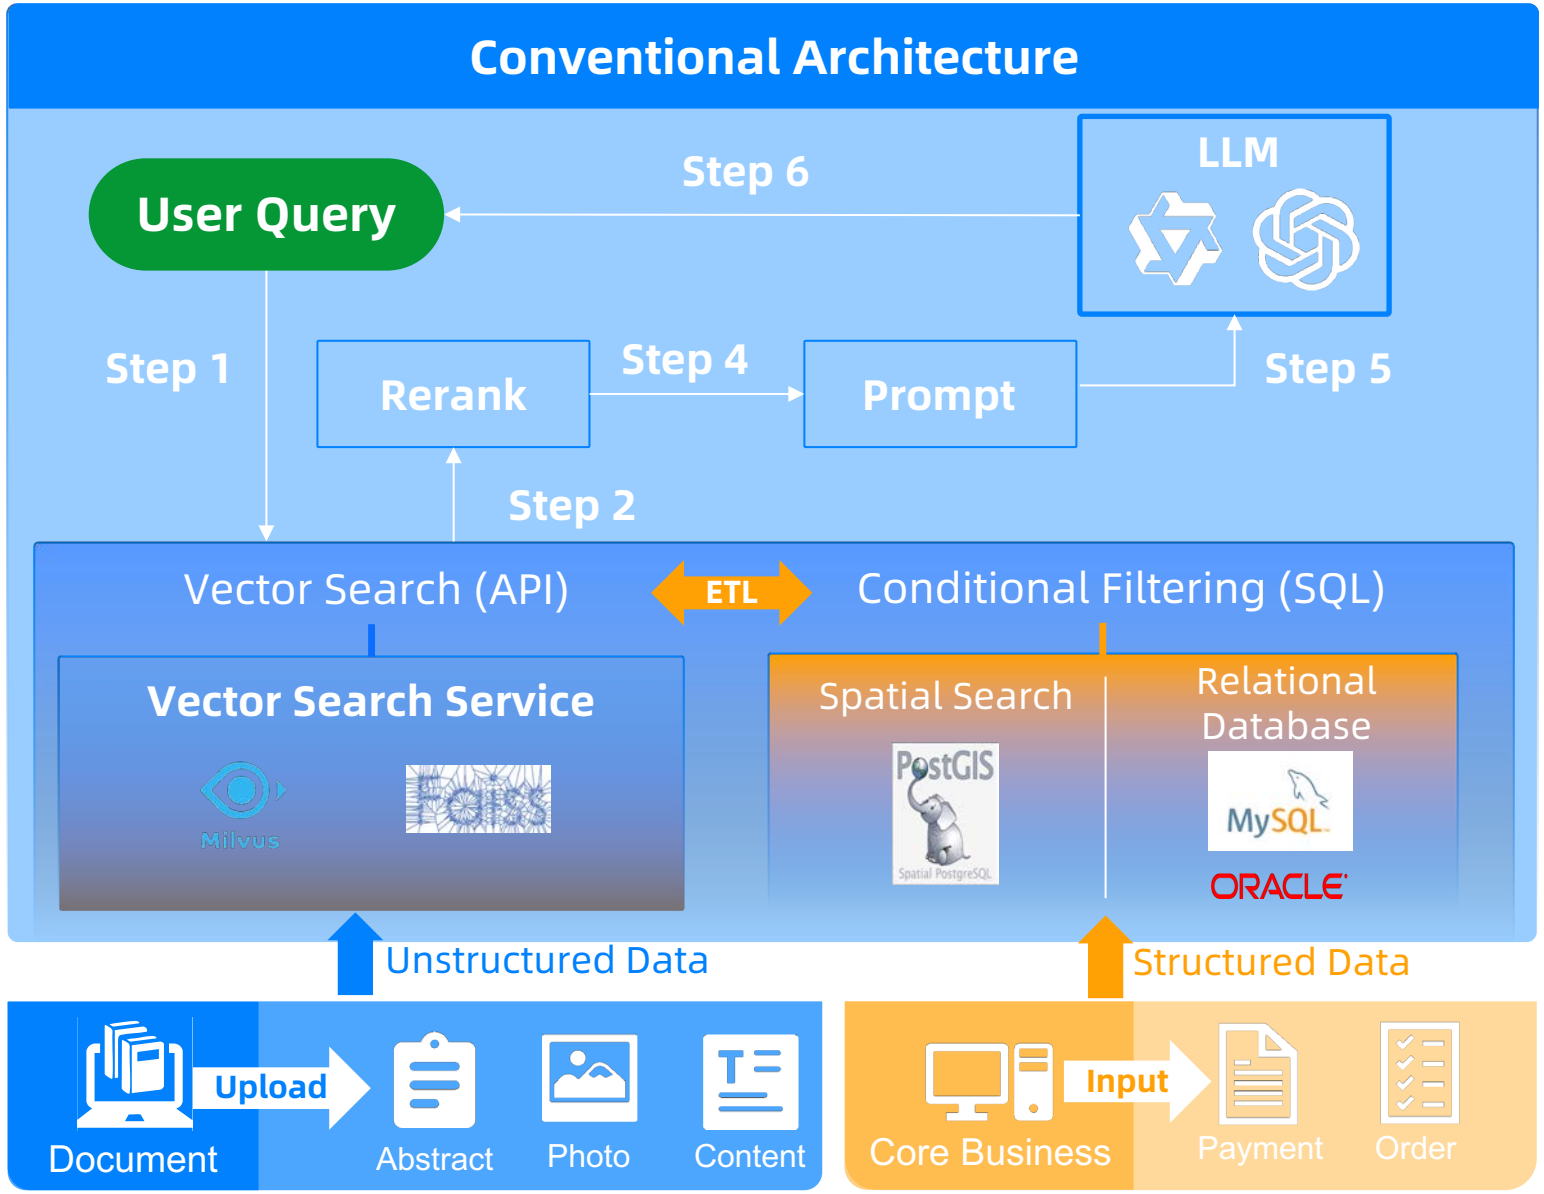

Pain Points

- Data scattered across multiple systems
- Hard to mix structured & unstructured data in one query
- High concurrency with low latency requirement

AI-Powered Knowledge Retrieval for Internal Teams

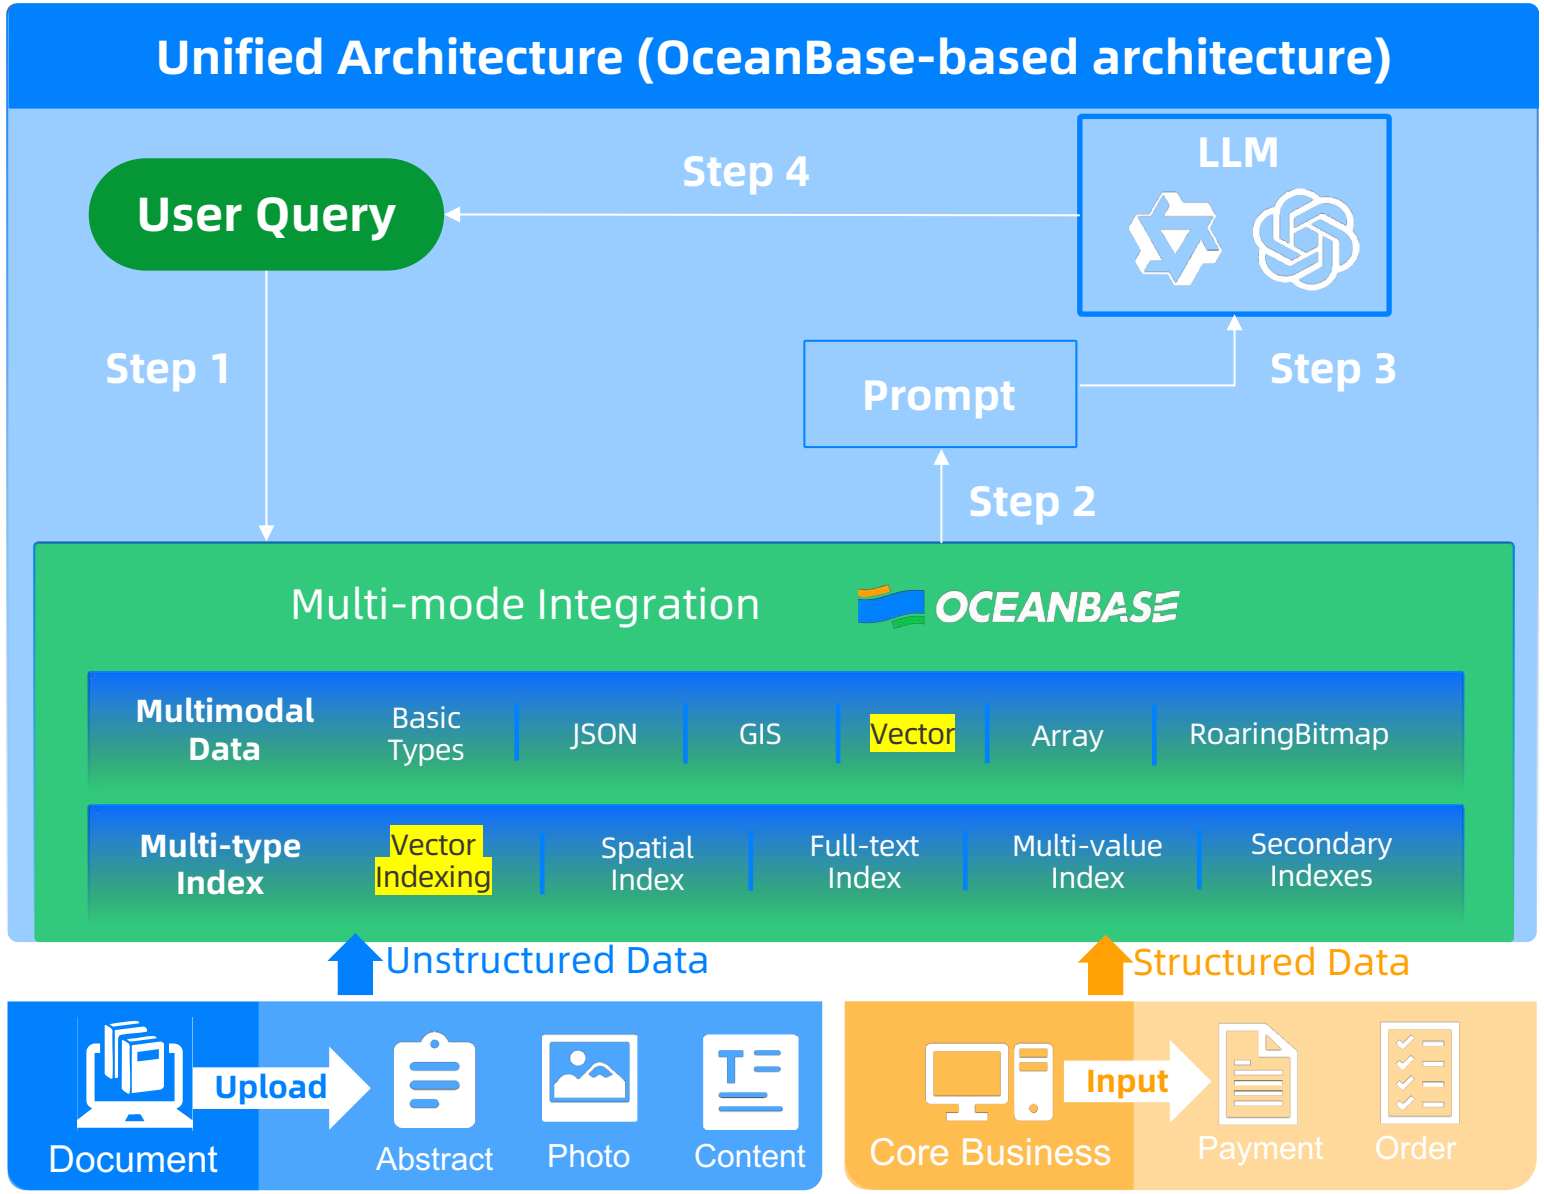

Why Engineers Finally Chose OceanBase

Platform Engineering Approach

- Build a unified data access layer (SQL + vector search)
- Elastic scaling on Kubernetes
- Minimize multi-system integration complexity

Outcome

- Query latency cut by over 70%
- Data freshness in minutes
- 40% less integration code to maintain

High Latency in Conventional Architecture

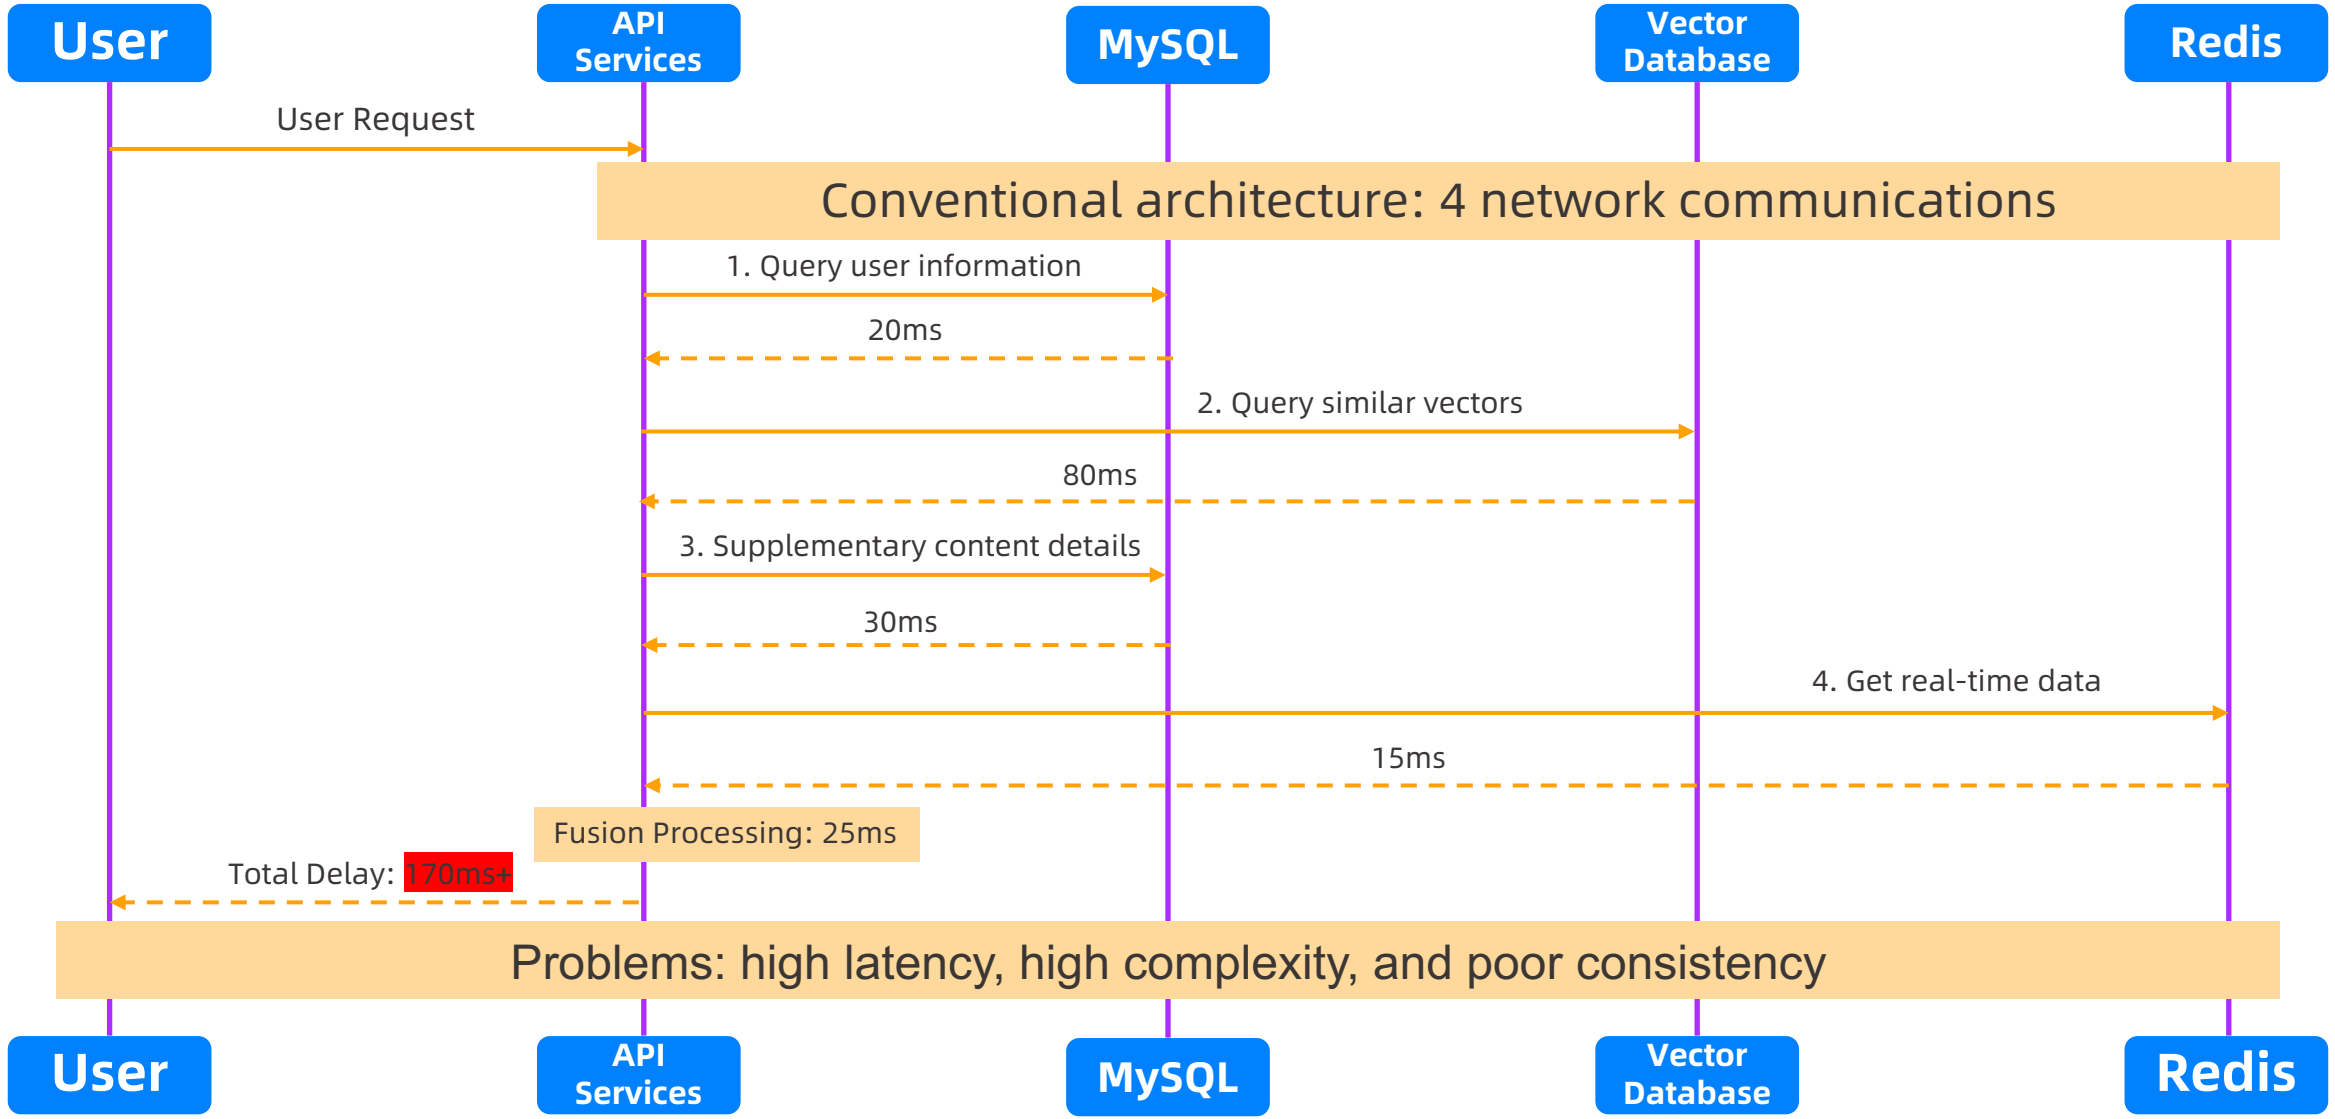

Low Latency with OceanBase Unified Architecture

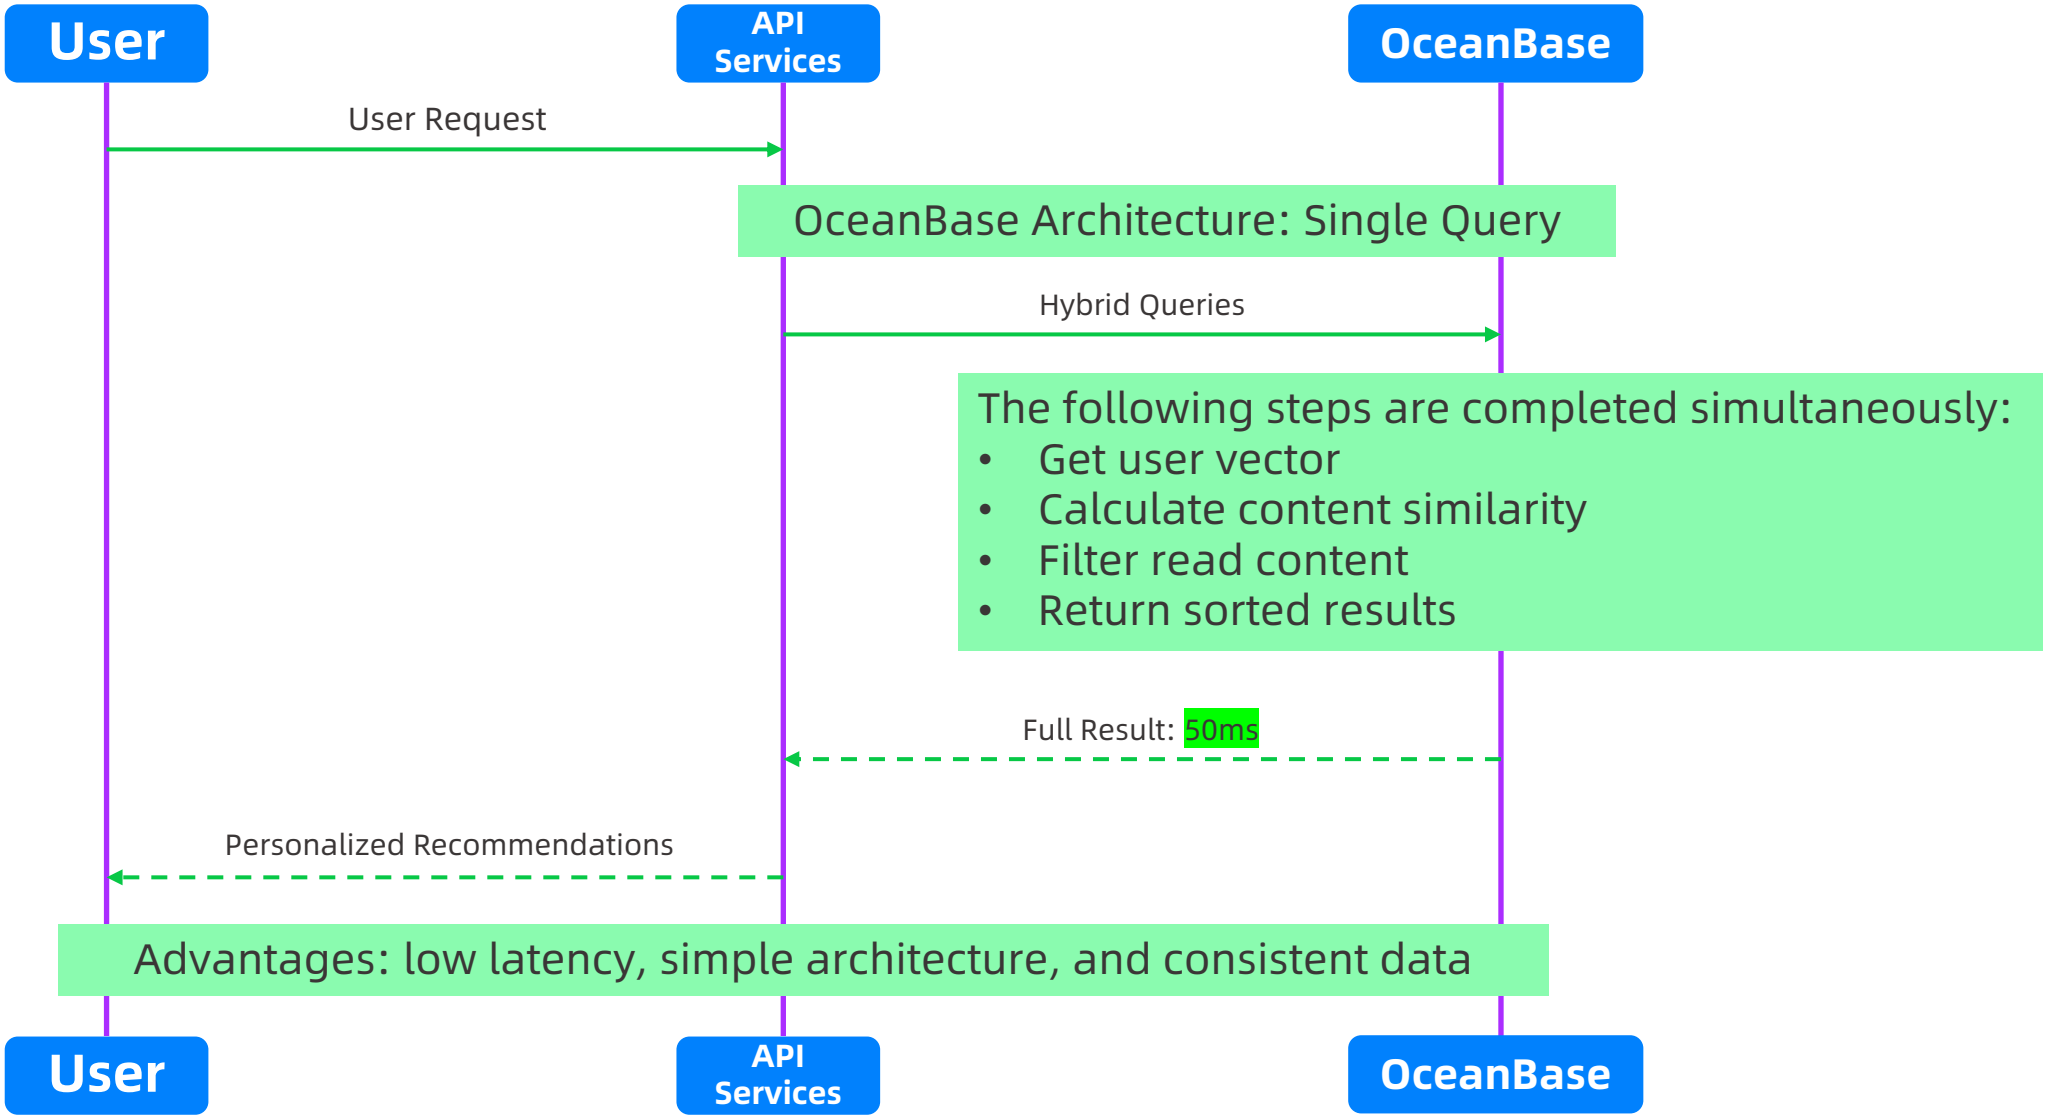

Executing Hybrid Queries in OceanBase

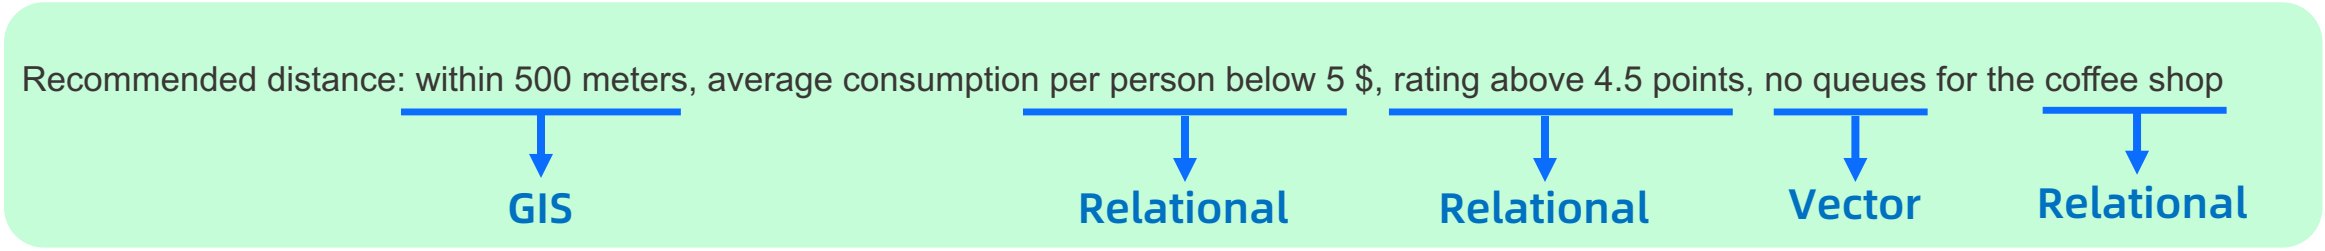

Perform mixed computing directly in SQL

```
SELECT *
FROM obAgent
WHERE st_distance(location, st_srid(point(@longitude, @latitude), 4326), 'metre') < @query_distance
      AND score > 4.5
      AND avgConsum < 5
      AND storeType = 'coffee shops'
ORDER BY  l2_distance(featureVec, @query_embedding) approximate limit 20;
```

OceanBase: One Unified Architecture. Multiple Capabilities (1)

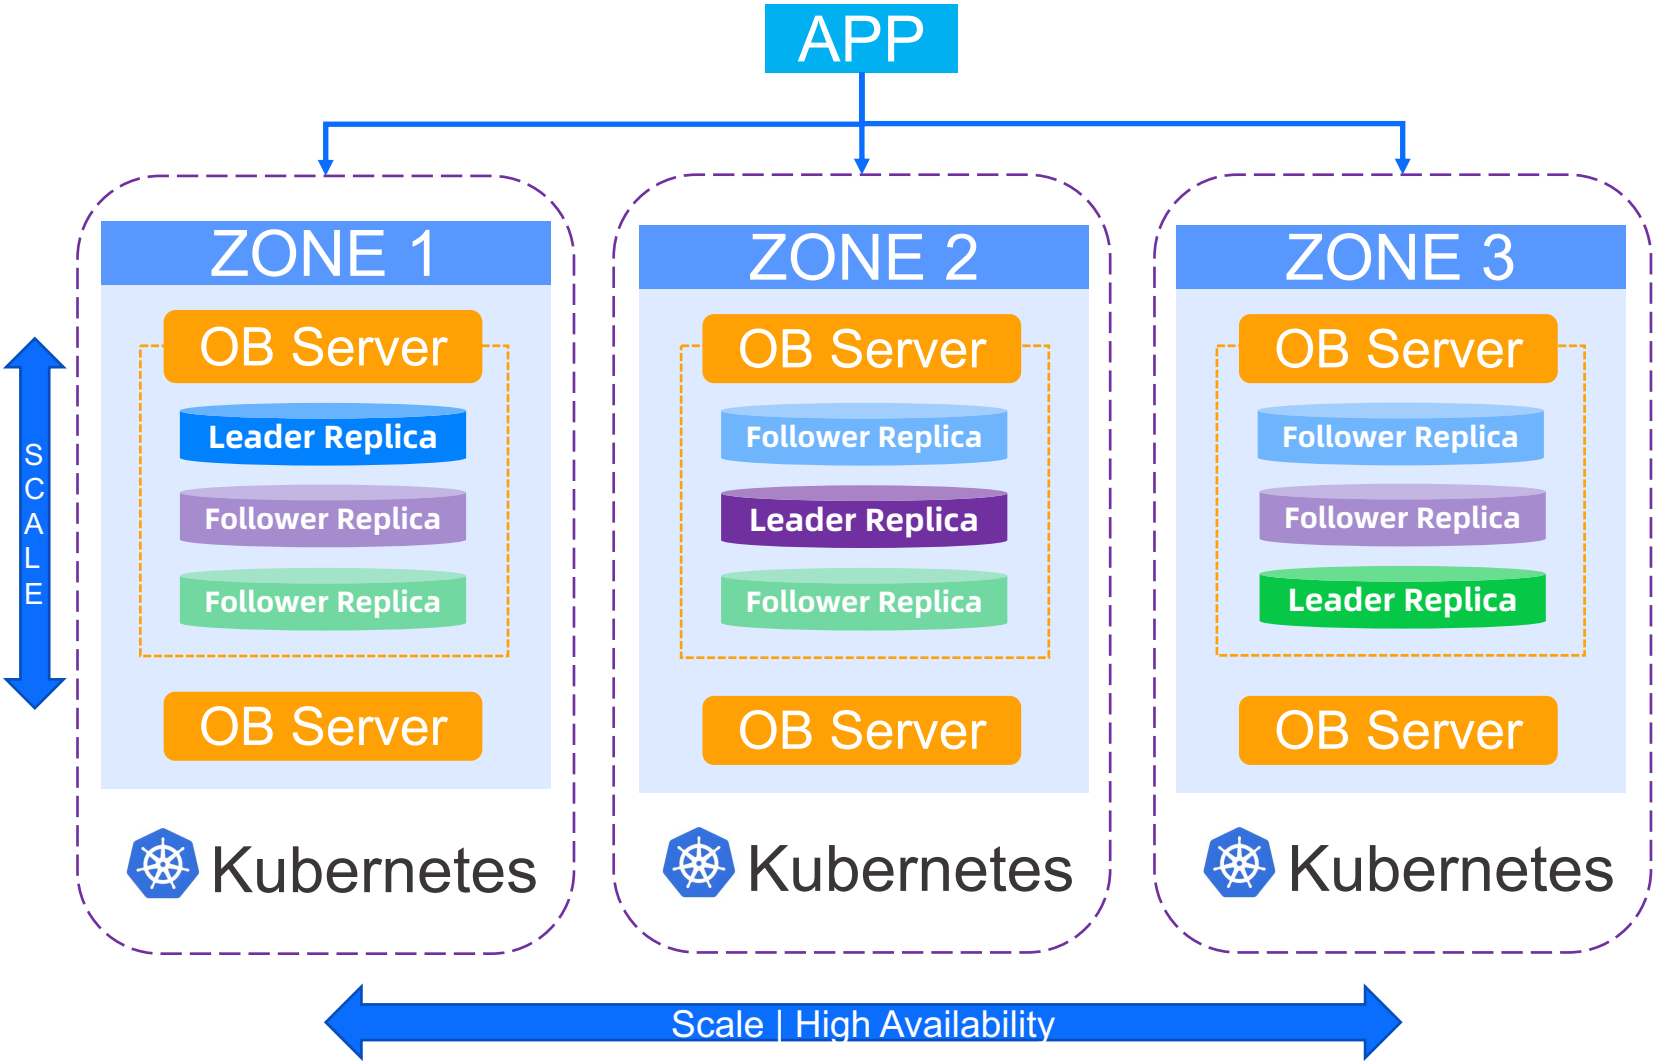

Elastic scaling on Kubernetes

- OceanBase is fully containerized, supports Operator management, and allows for elastic horizontal expansion of computing and storage.

Strong consistency & high availability

- Paxos consensus algorithm, triple replica mechanism, and automatic failover ensure high availability and strong consistency

OceanBase: One Unified Architecture. Multiple Capabilities (2)

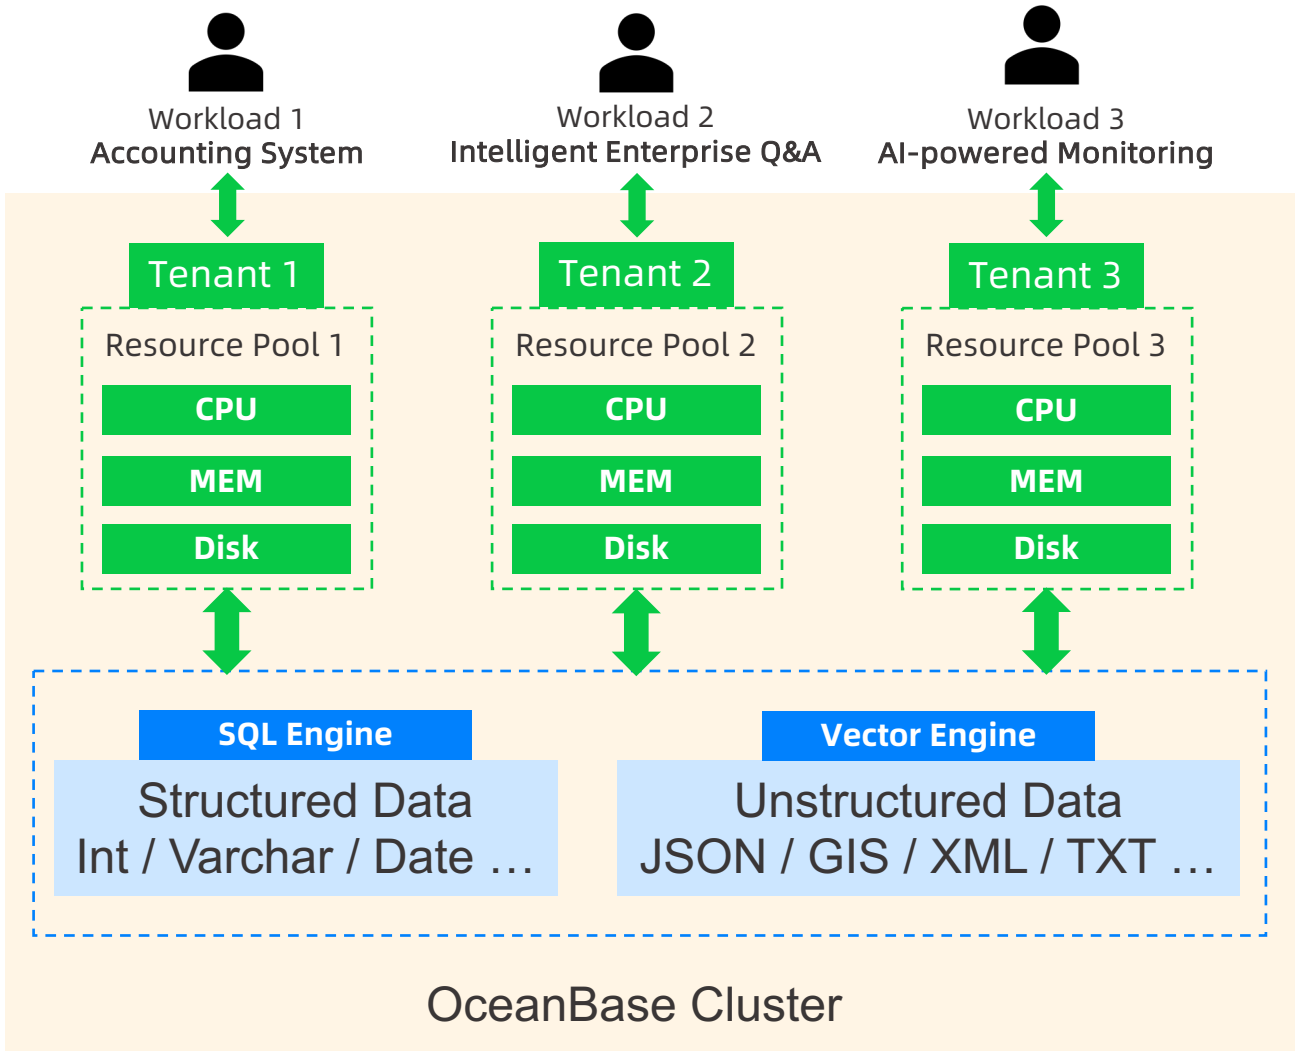

OceanBase provides native multi-tenant architecture, resource quotas, and data isolation between tenants

One platform for all data workloads

- One engine supports OLTP, log analysis, semantic retrieval, AI applications, etc.

Unified access (structured + unstructured)

- SQL engine supports structured query, full text search and vector search for unstructured

Closing Thoughts & Takeaways

1 Simplify

- Reduce the number of systems and integration complexity

2 Unify

- The same platform supports both structured and unstructured data

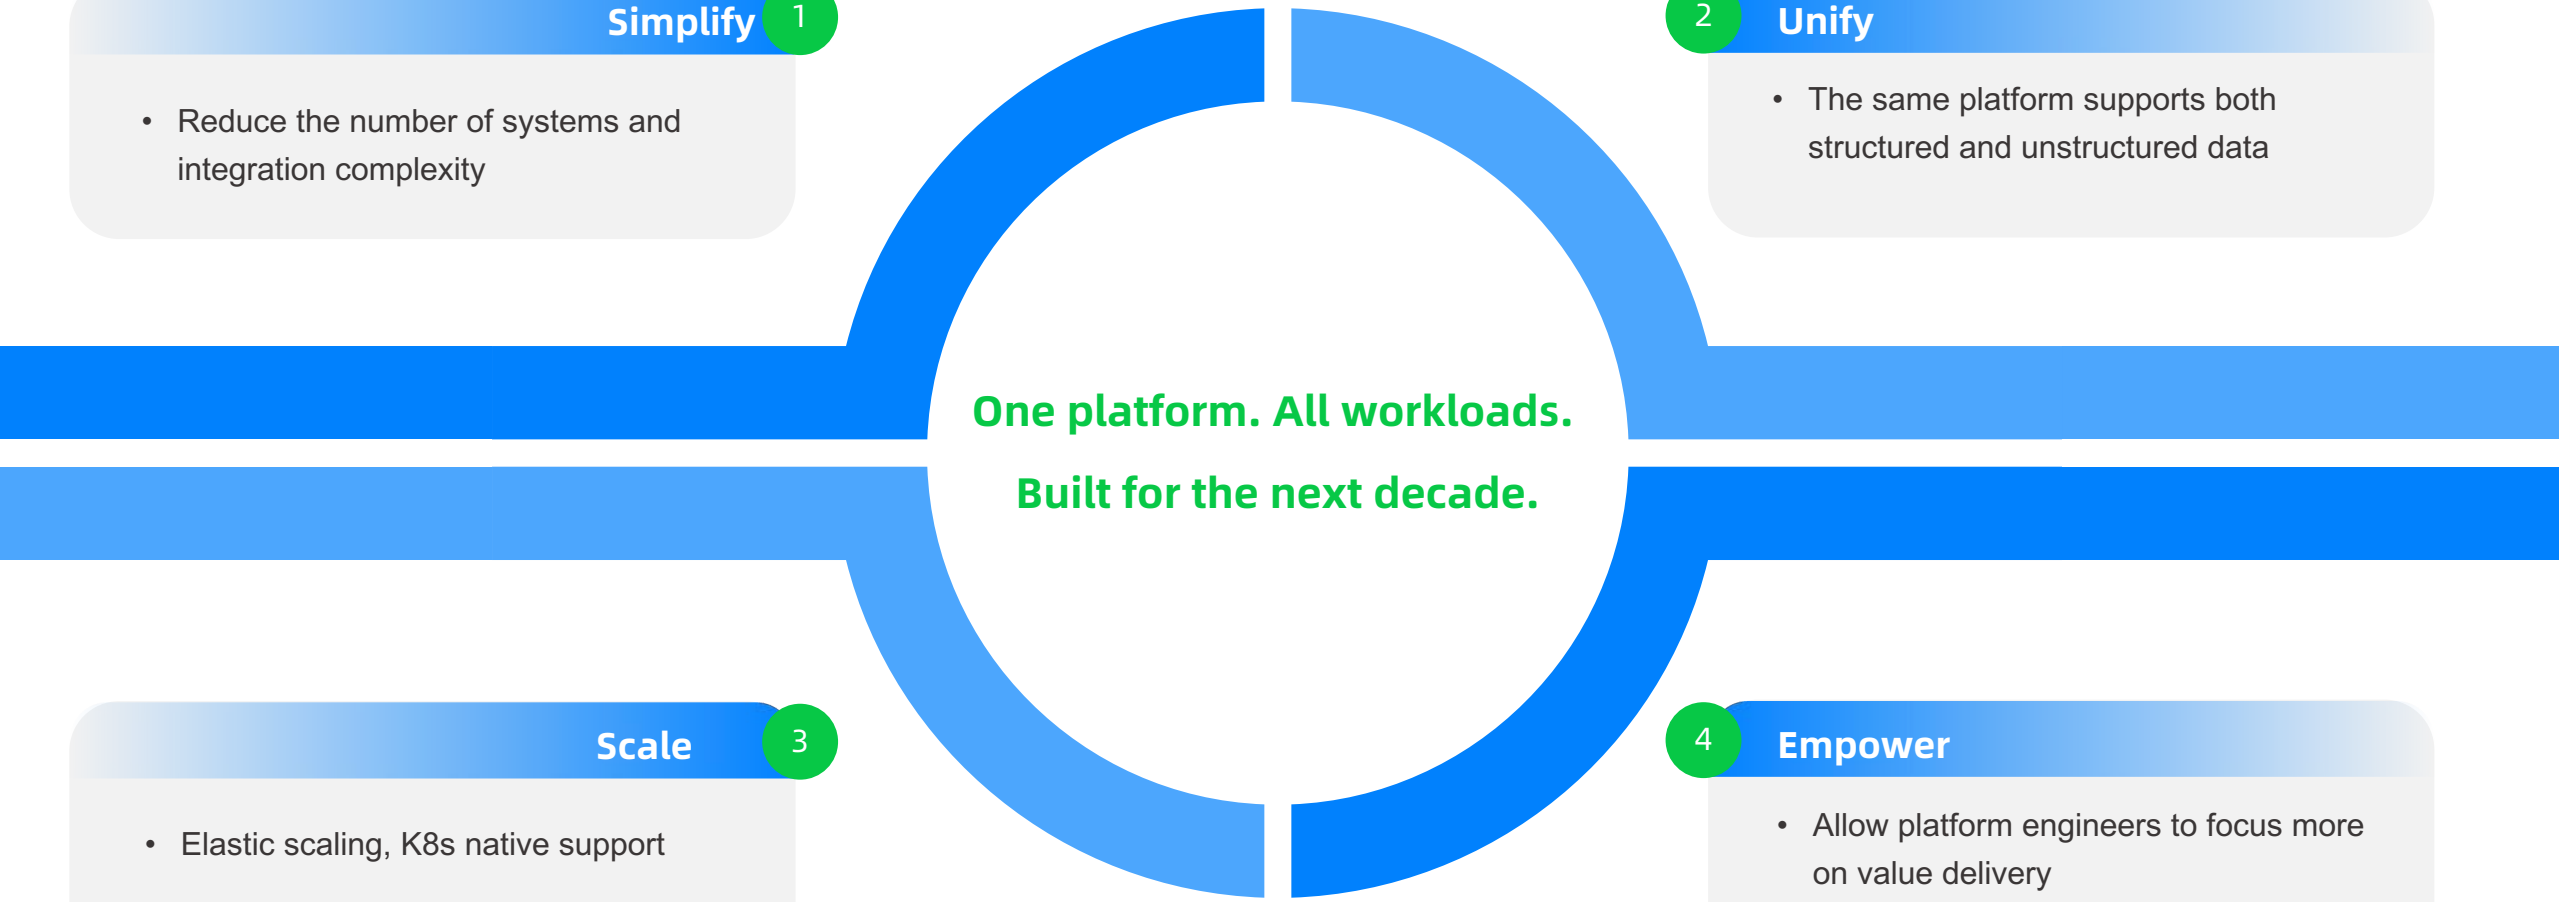

3 Scale

- Elastic scaling, K8s native support

4 Empower

- Allow platform engineers to focus more on value delivery

Thank you!