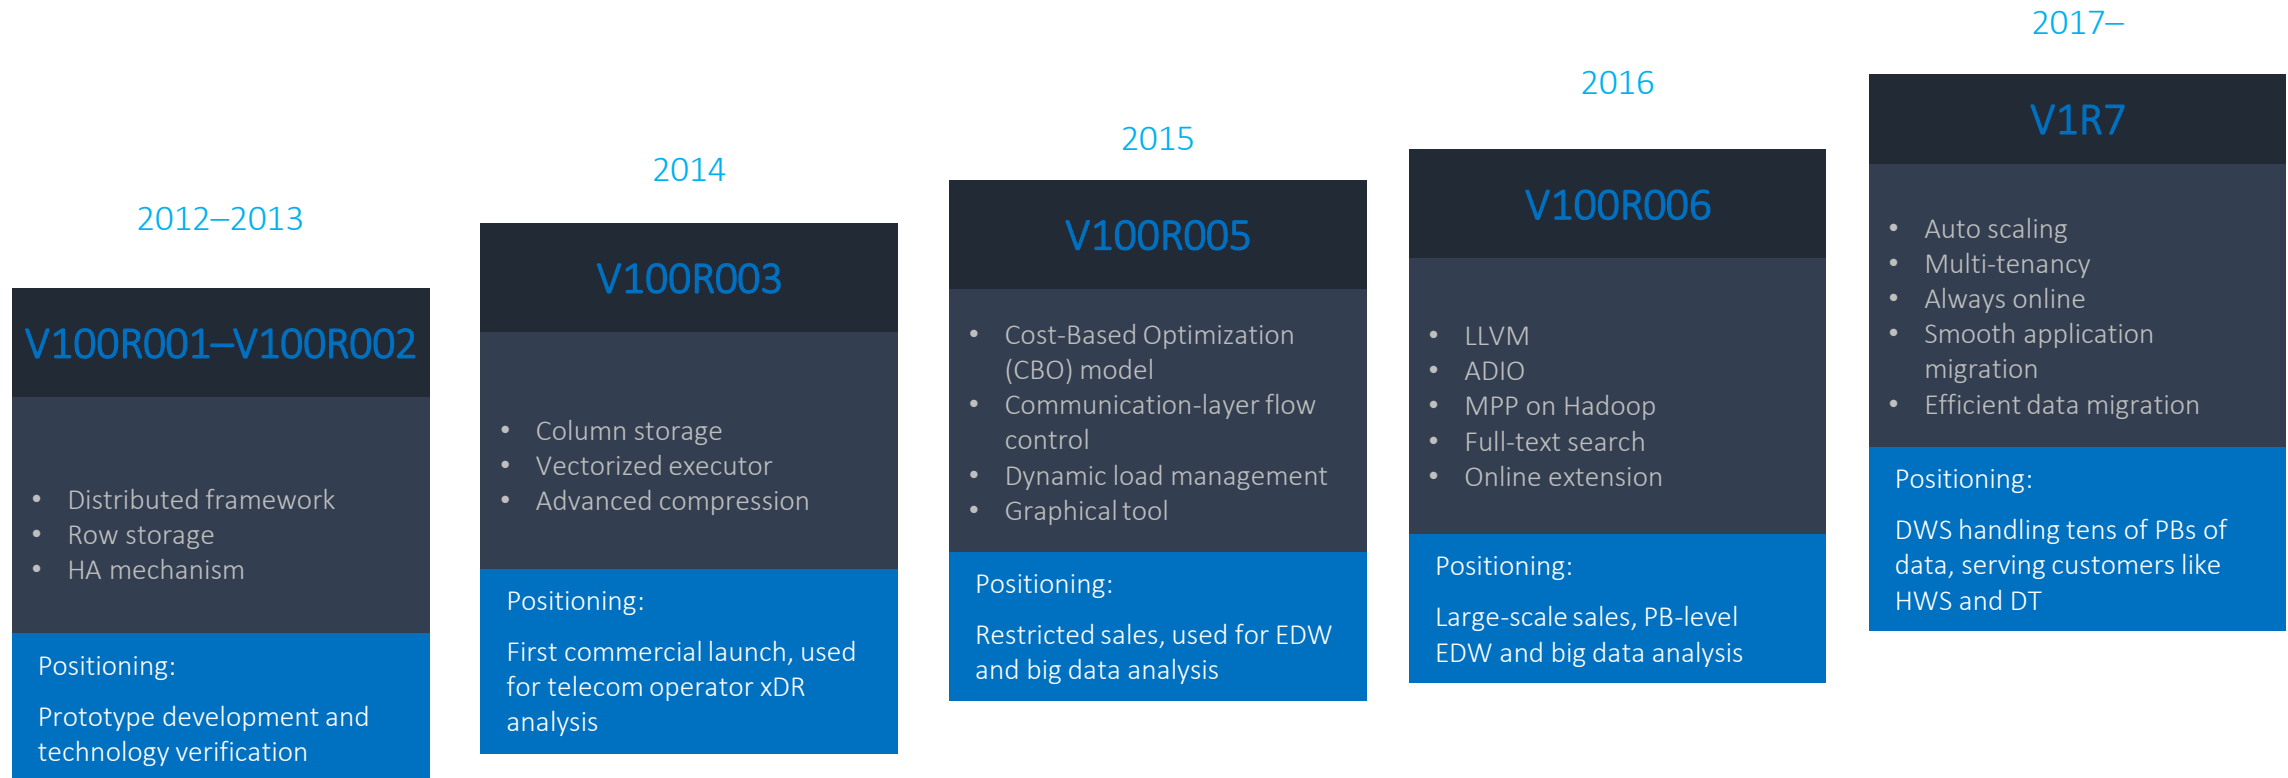


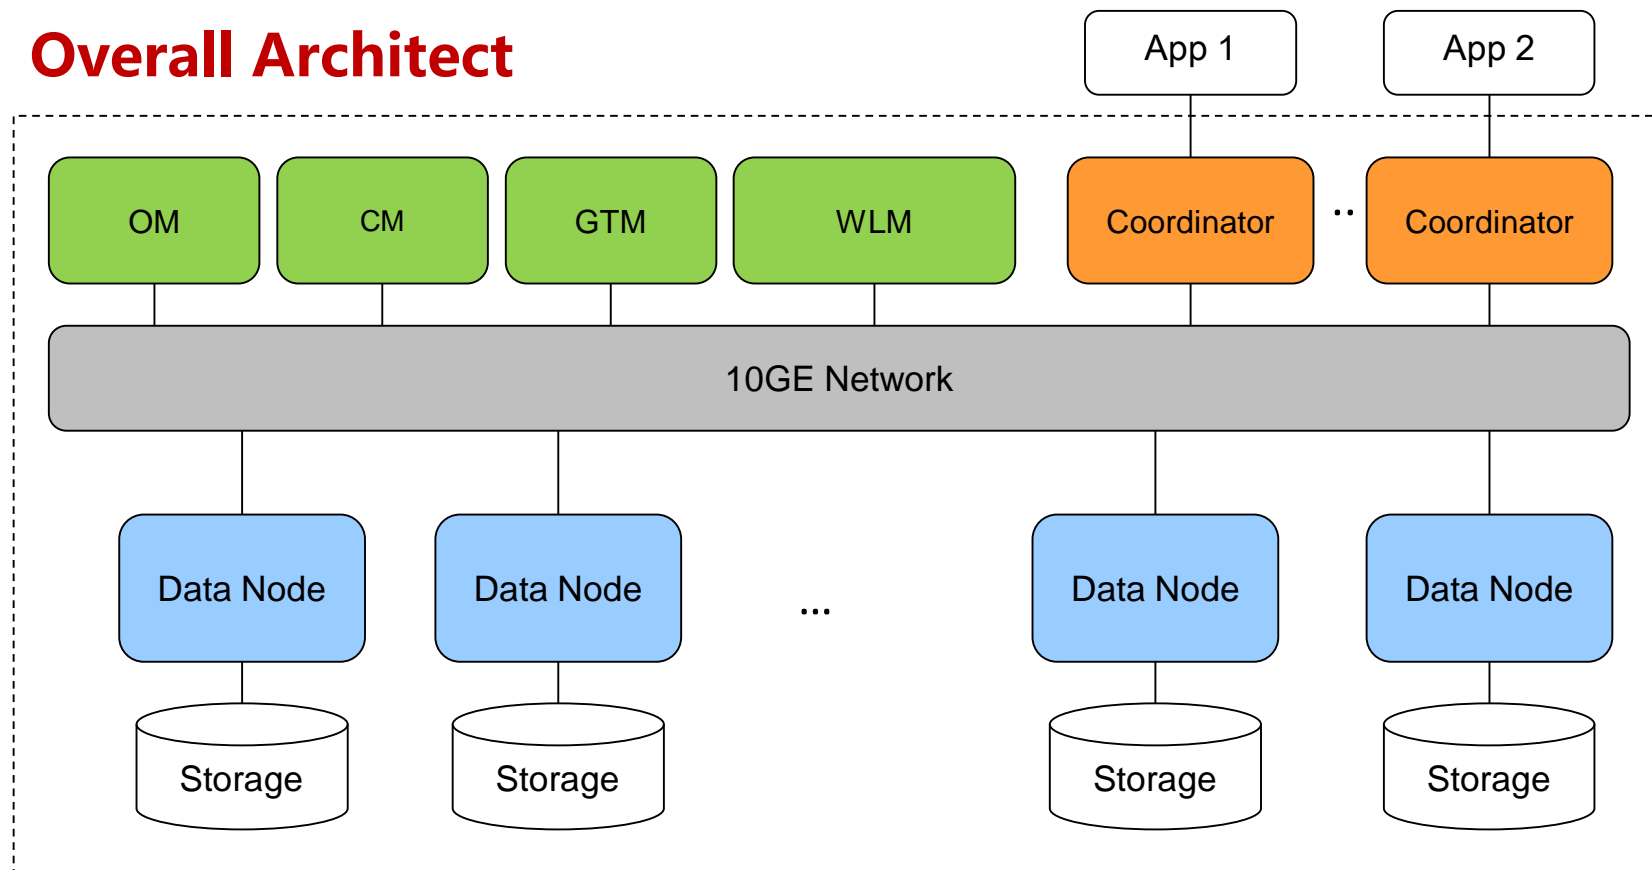
# Past and Future of Huawei DWS Kernel



Technology research project: continuous, pioneering technical exploration  
Machine-learning optimizer, Cloud-native architecture and real-time analysis

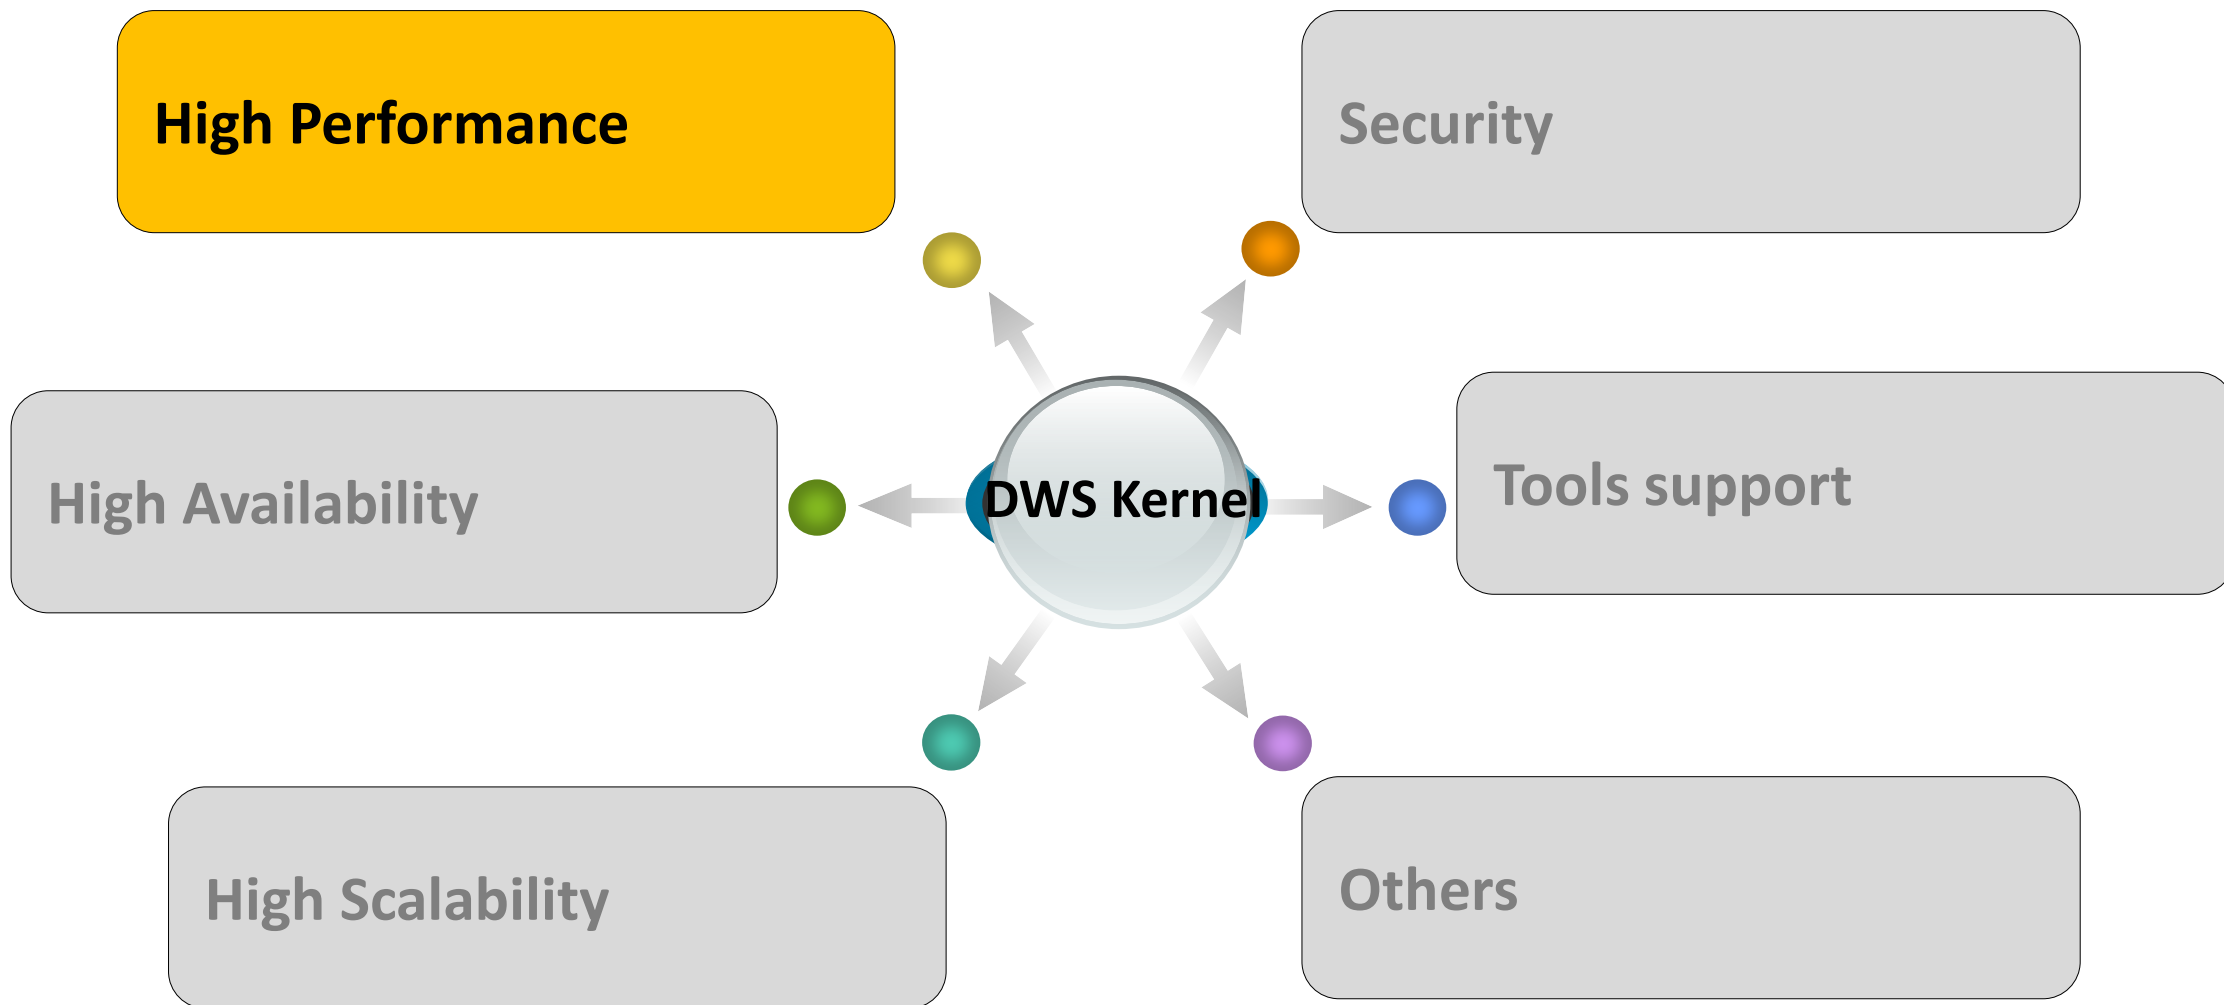
# DWS Kernel: A shared-nothing distributed data warehouse

## Overall Architect

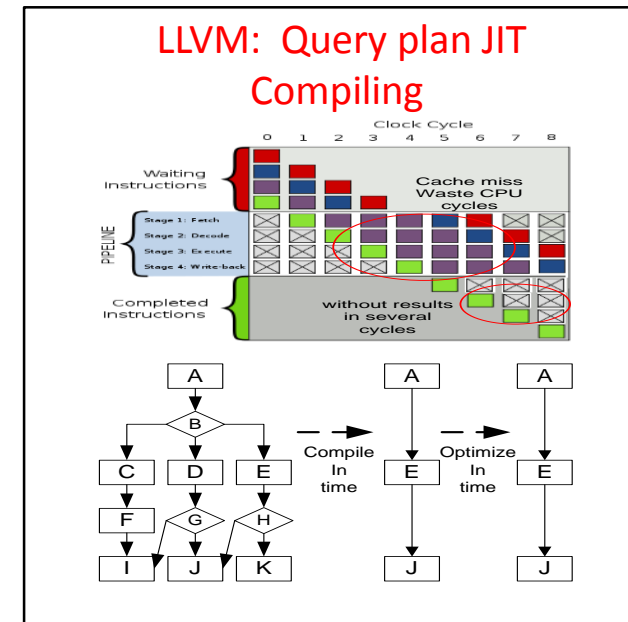
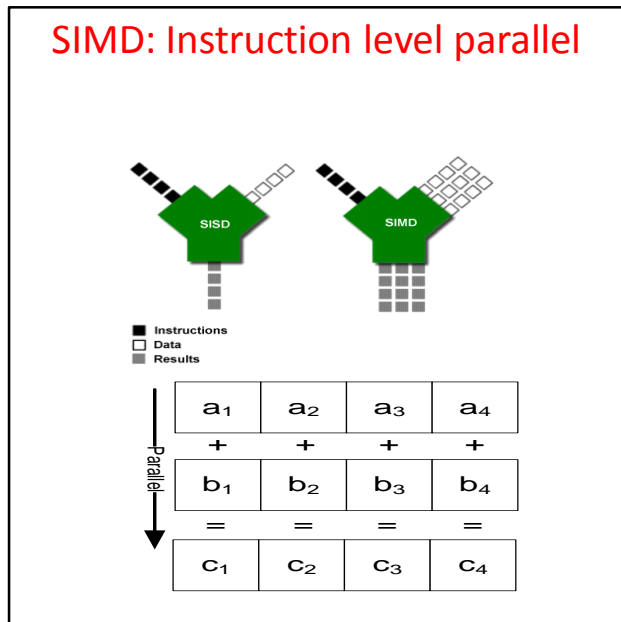
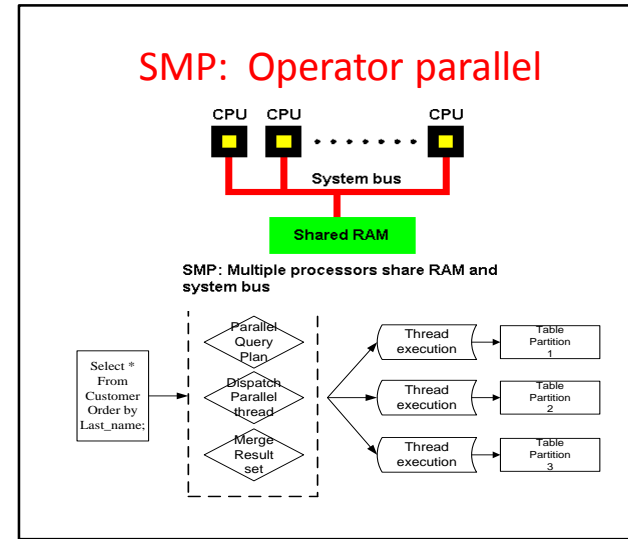
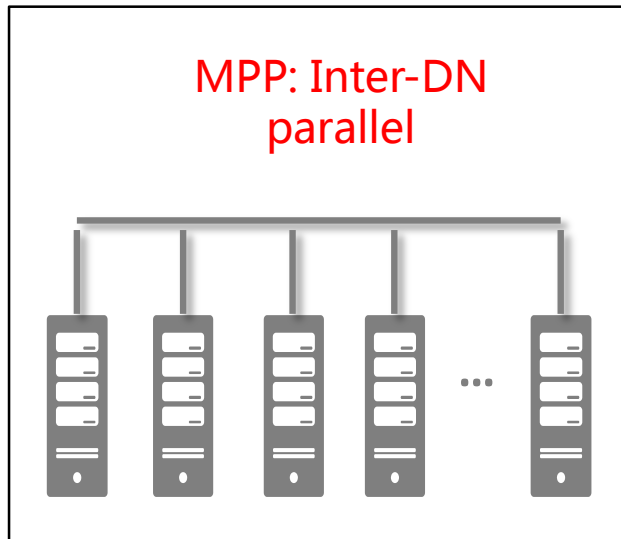


- 256 Physical Hosts/Cluster
- 5T/Host ~ 2PB/Cluster with data compression
- 1024 Hosts/Cluster under development

# Enterprise features






# Parallel Execution



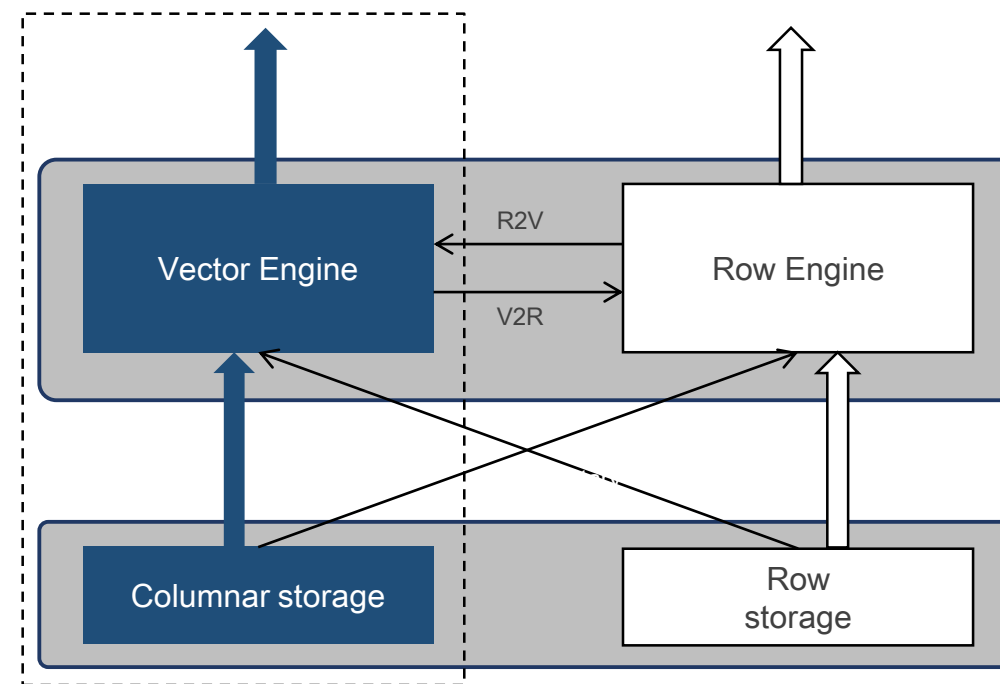
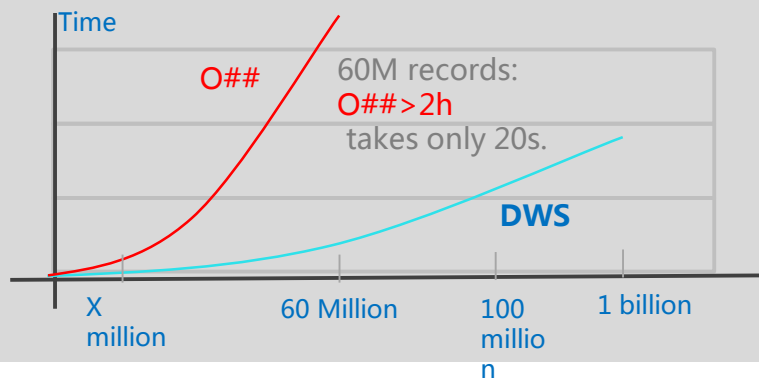
# Row and column hybrid storage

Analysis of tens of billions of data records within seconds.

For example, 6000w records and 1000 intelligent checkpoints in the city are used as an example. The query and analysis workload is heavy. 6000w x 6000w\*1000\*1000




When the data volume reaches 10 million, the query response time of the traditional database is extremely high.



# Global Statistics => Better CBO ==> 30% improve in performance

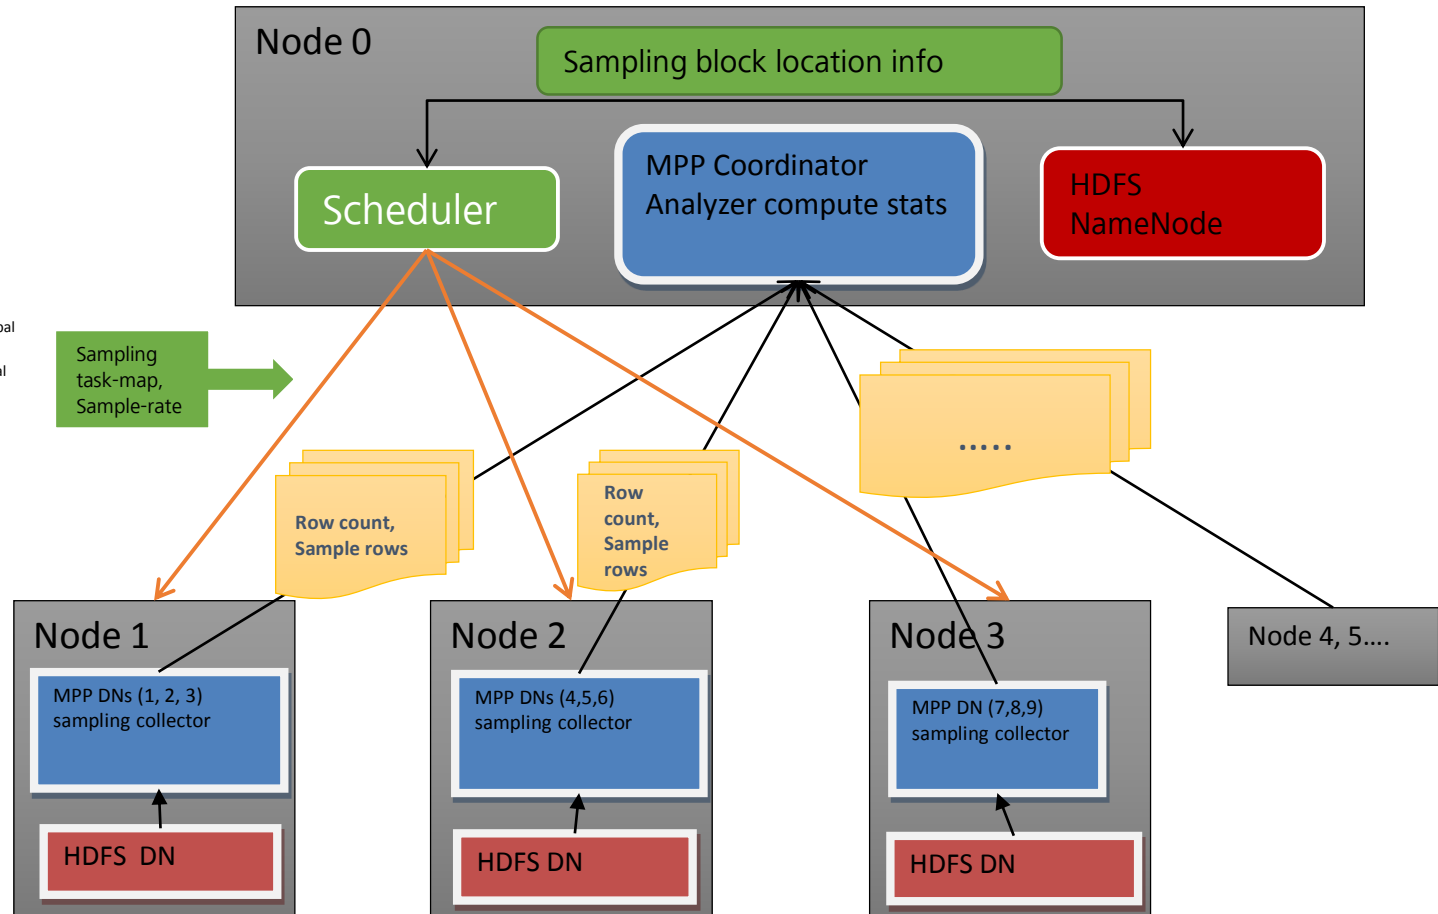

TPC-H 100X (3 node 18 DN cluster configuration)

About 40% query plans improved. Geomean improved 30%

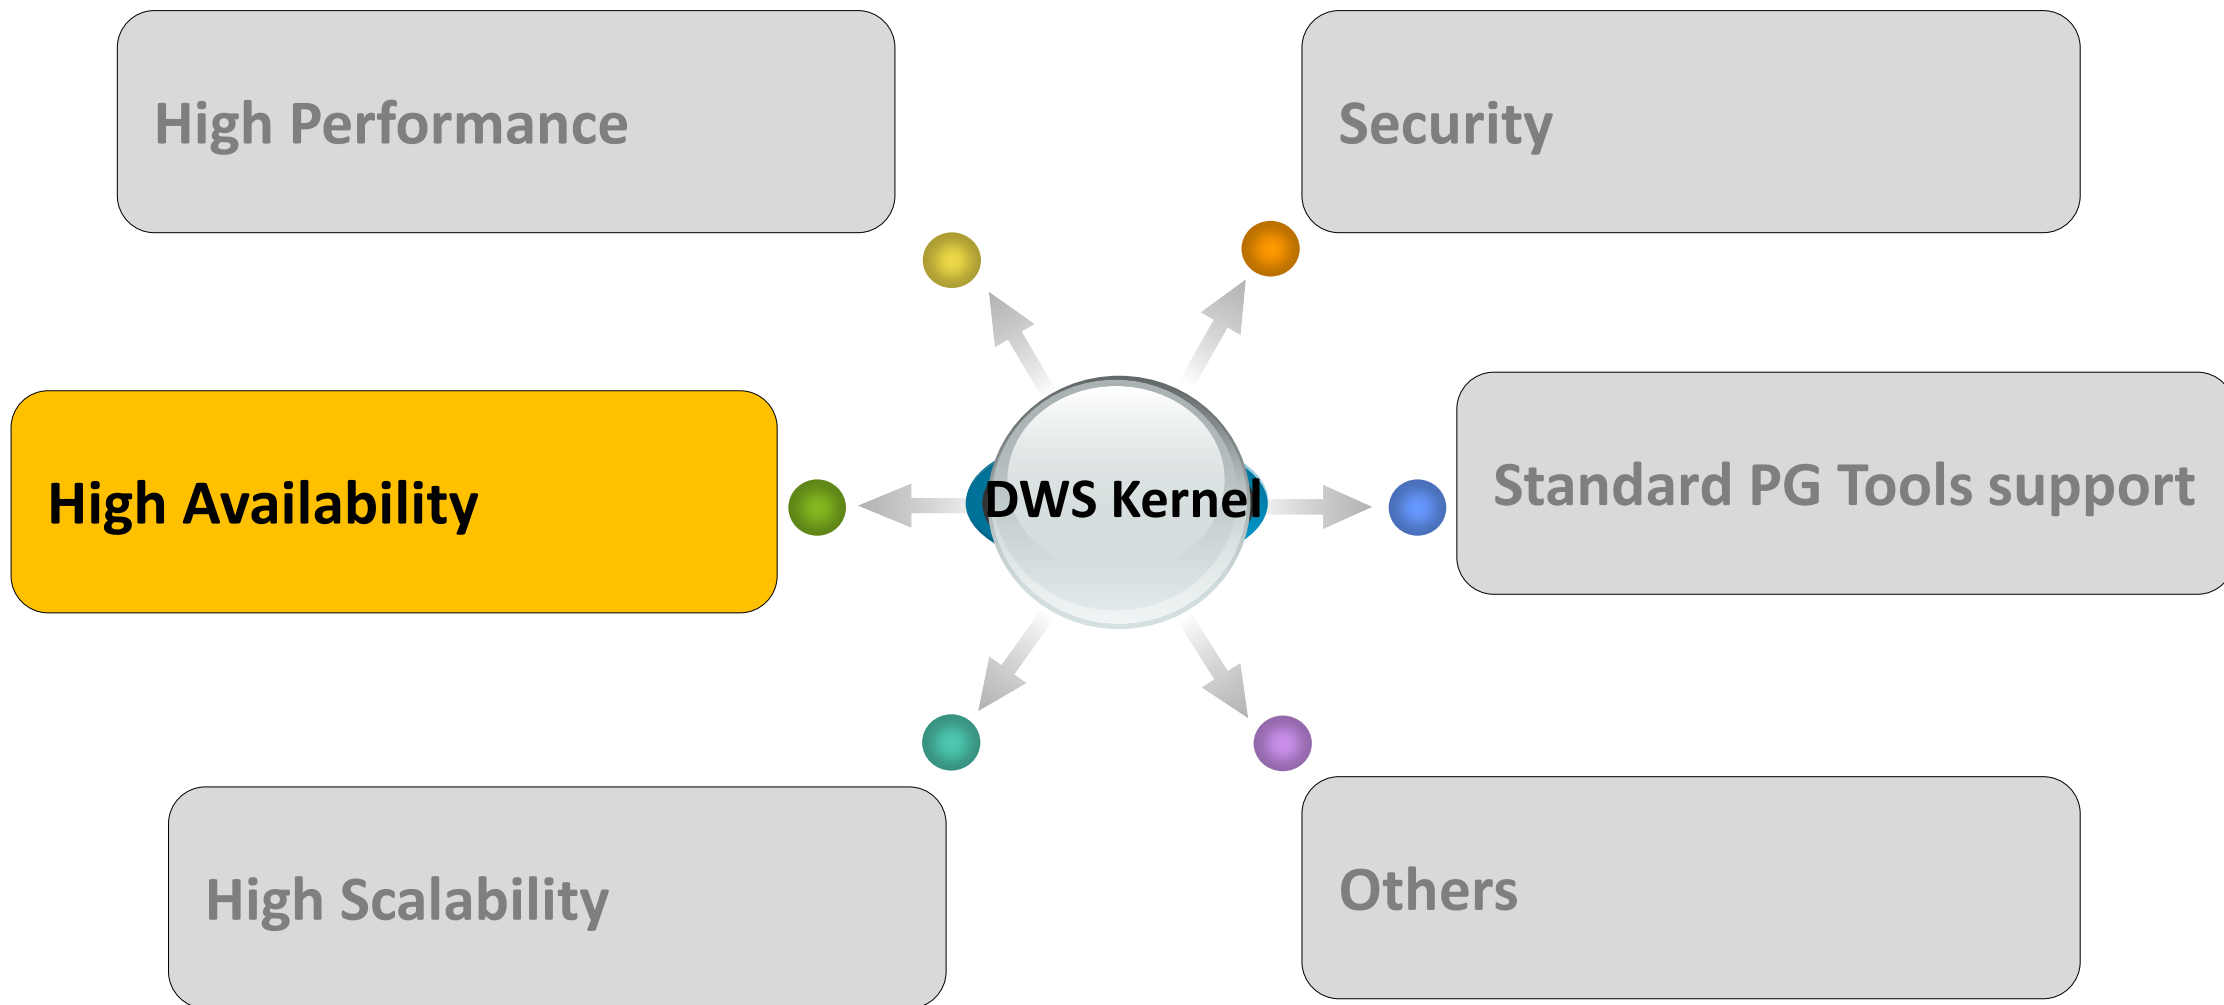


TPC-H 100X (3 node 18 DN cluster configuration)

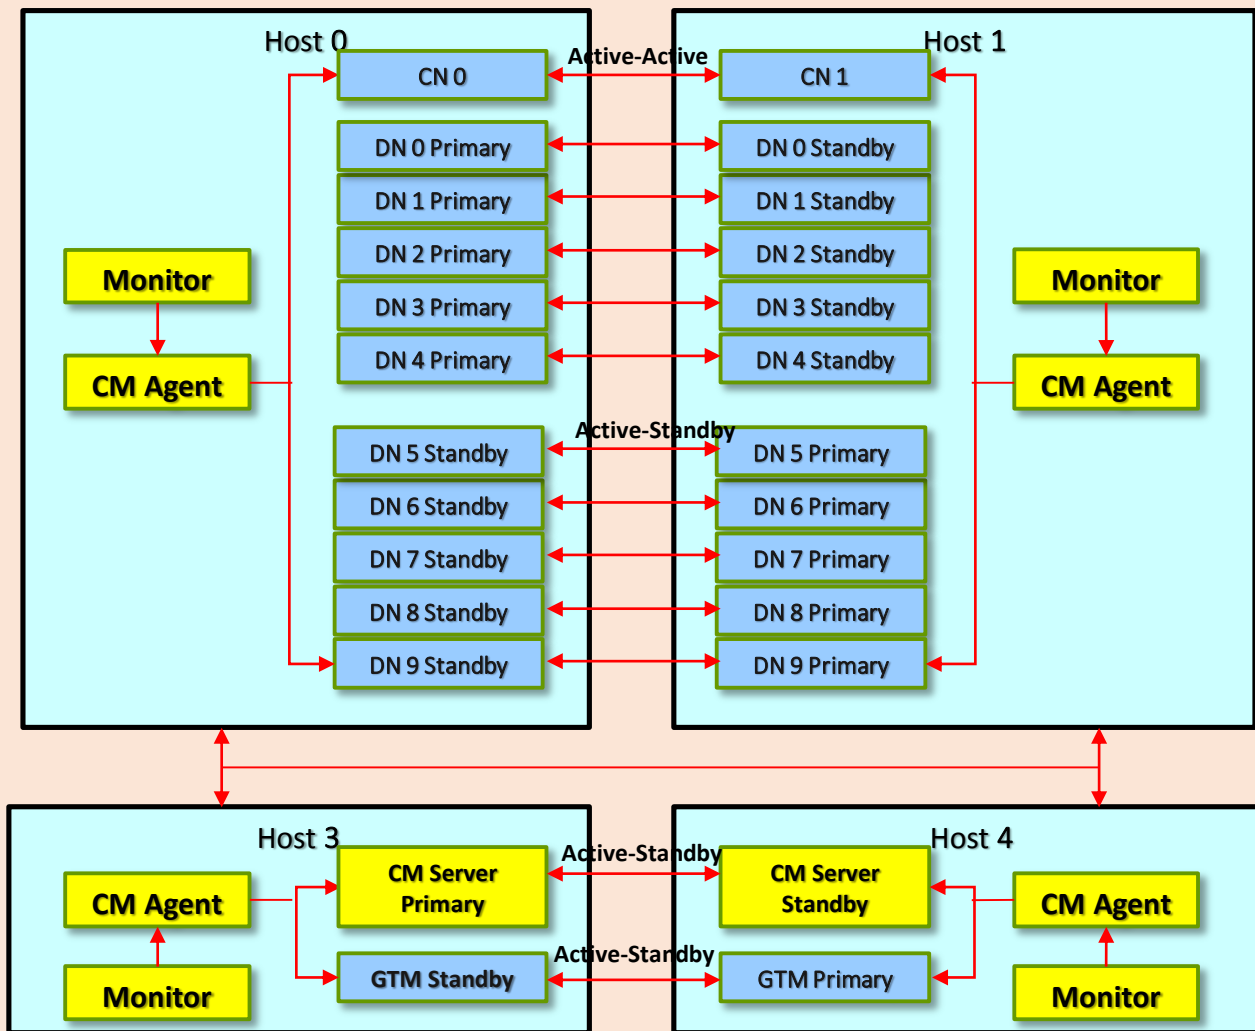
About 40% query plans improved. Geomean improved 30%



# Enterprise features



# High Availability



## HA

### Hardware HA:

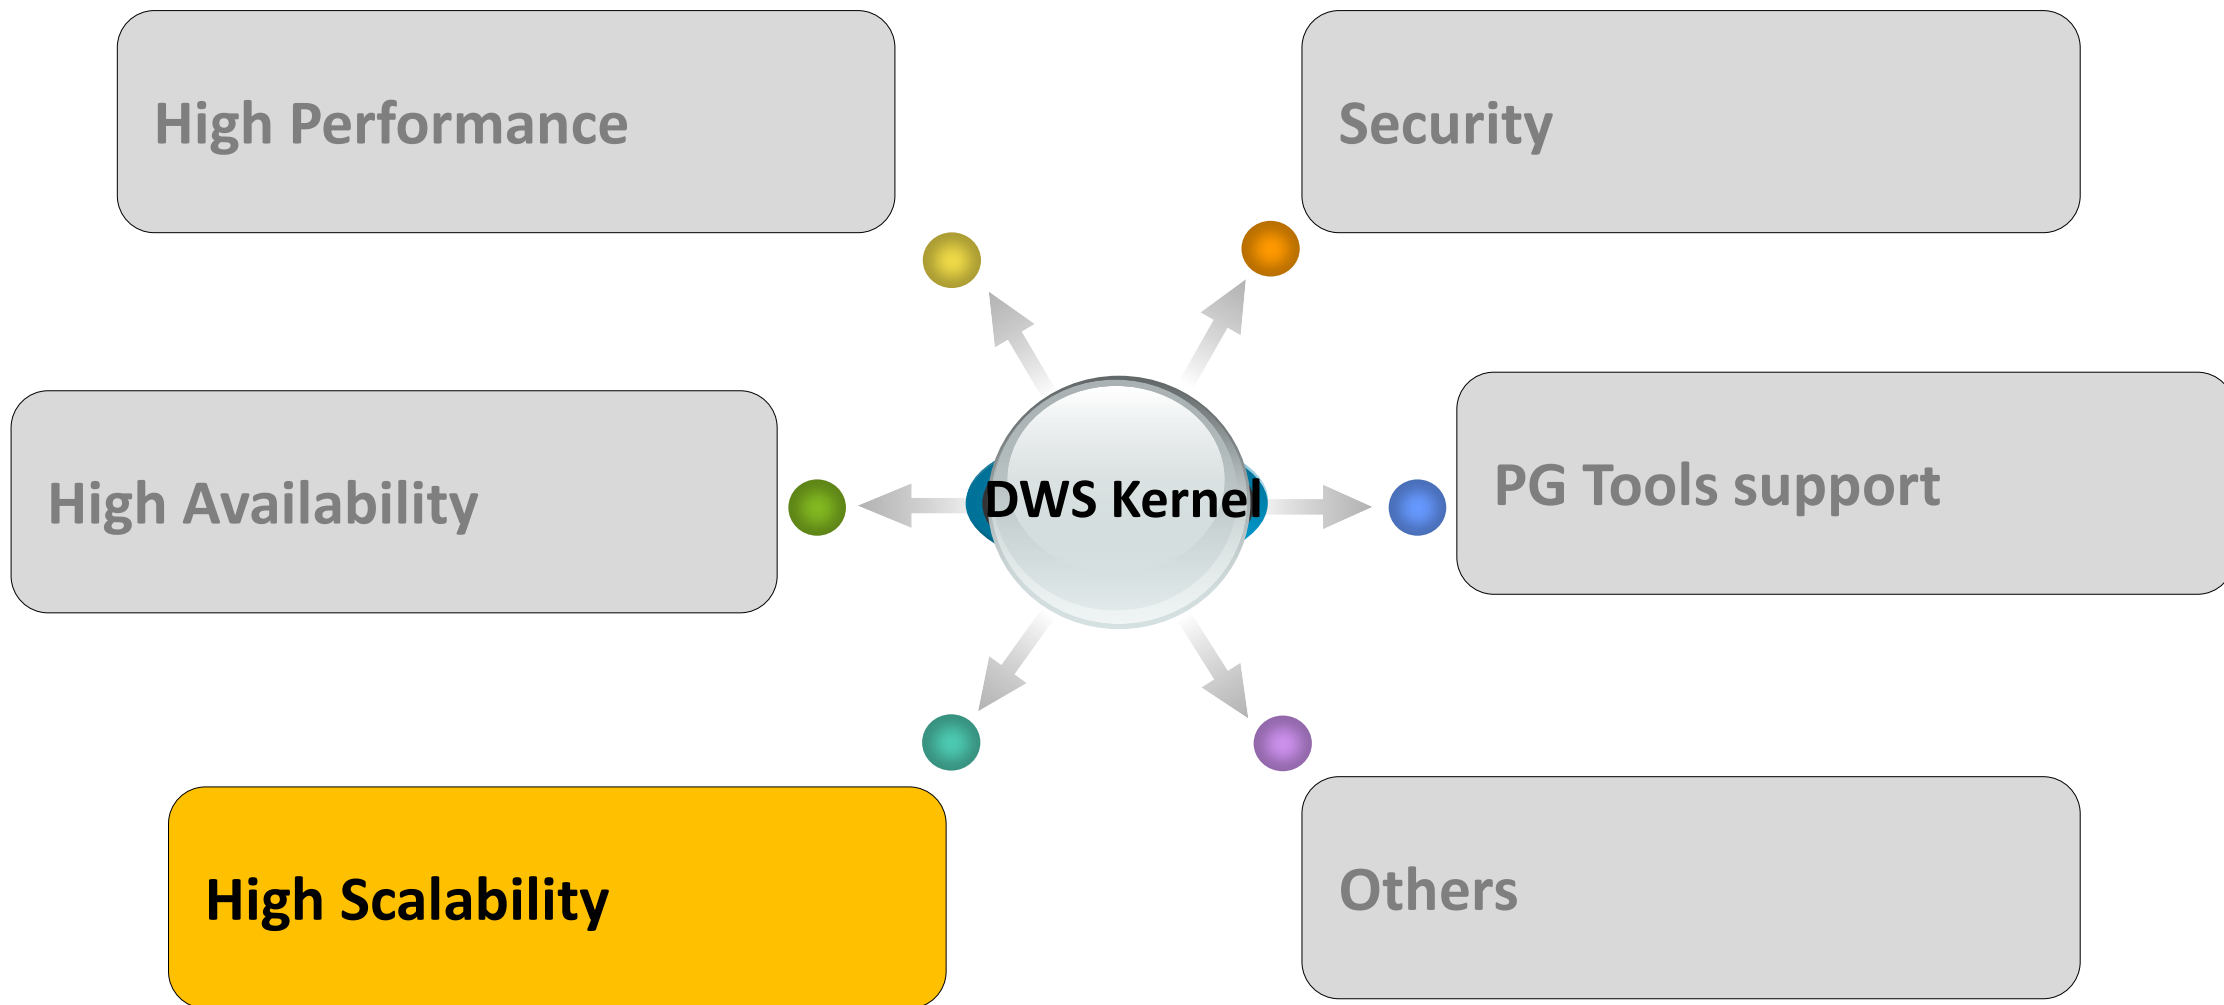
- ✓ Storage: RAID5 disk
- ✓ Network: Dual switches and Ethernet card
- ✓ Host: UPS Power

### Software HA:

- ✓ Coordinator node Active-Active
- ✓ Data node/GTM/ CM Active-Standby



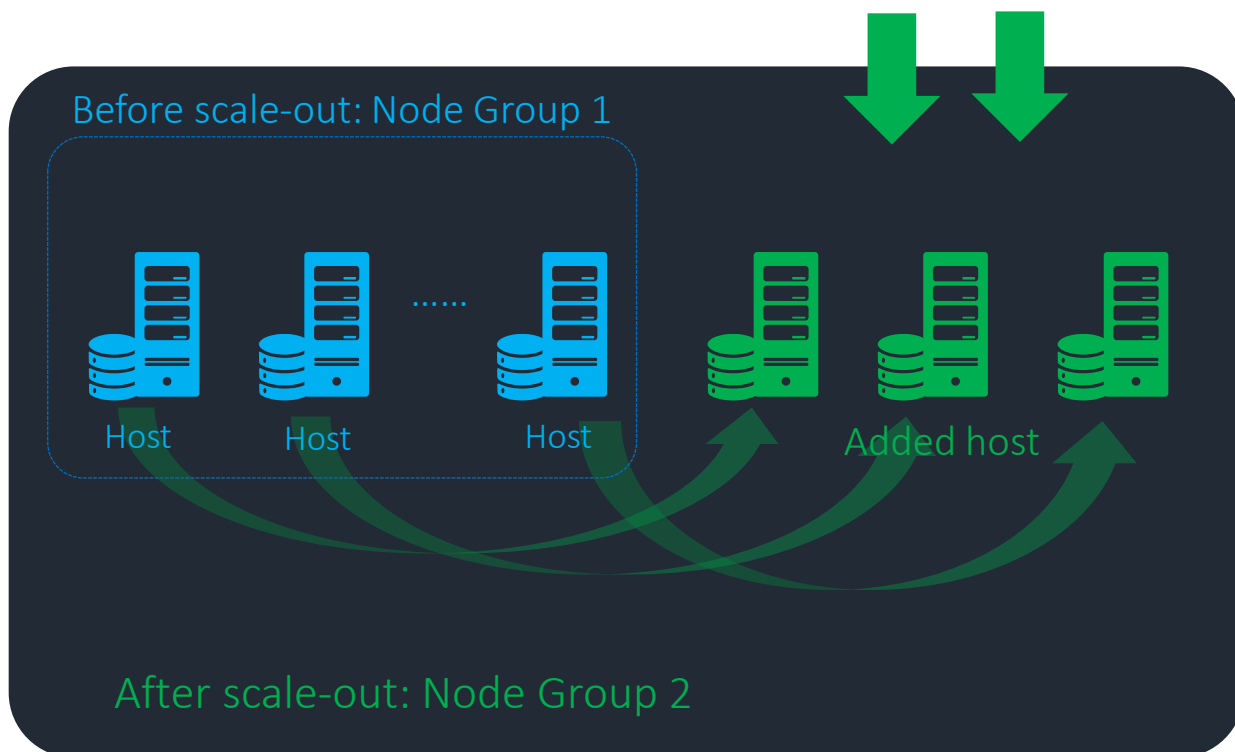
# Enterprise features



# Node Group Technology Avoids Service Interruption During Cluster Scale-Out

Online scale-out without service interruption:

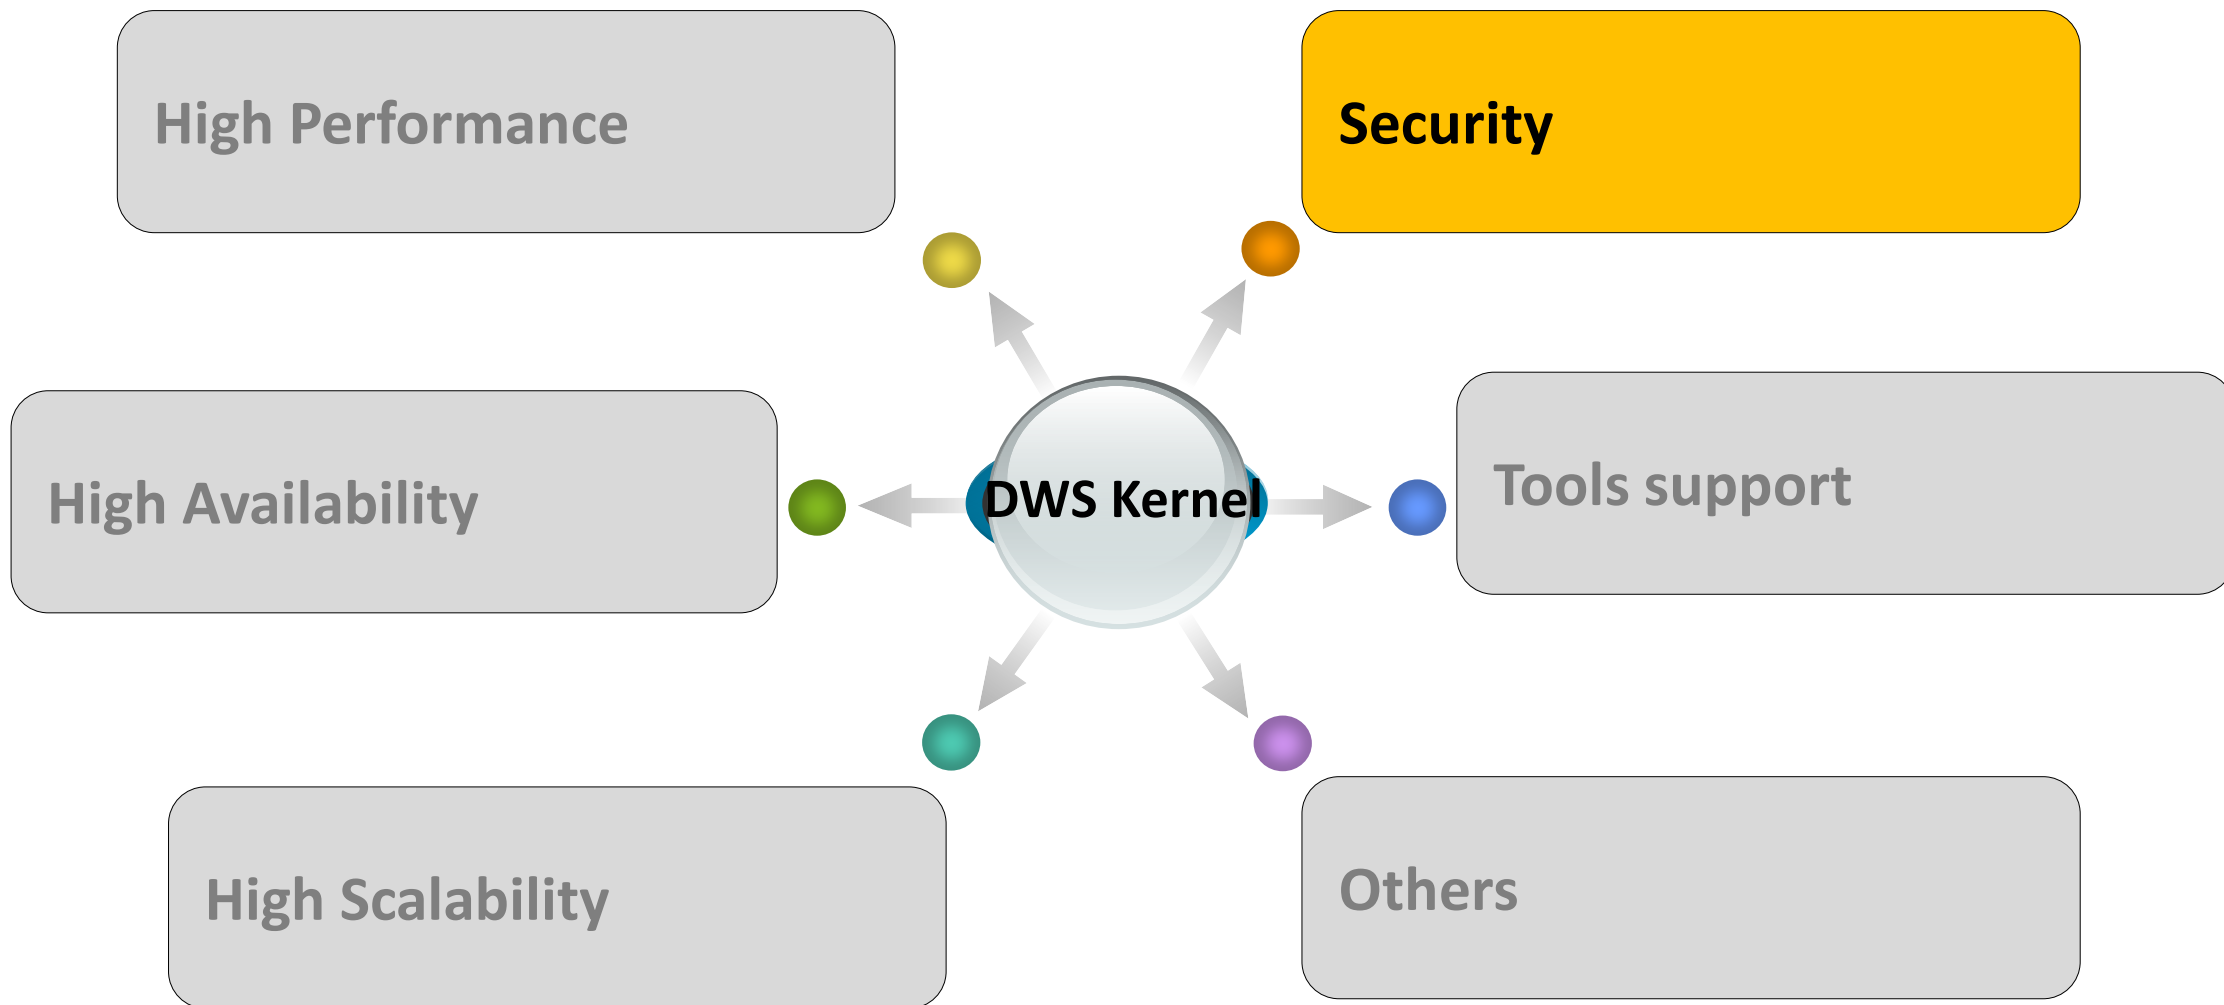
A cluster can be **smoothly scaled out to a maximum of 256 nodes**, improving storage and compute capabilities. Query and data import are not interrupted during the scale-out.



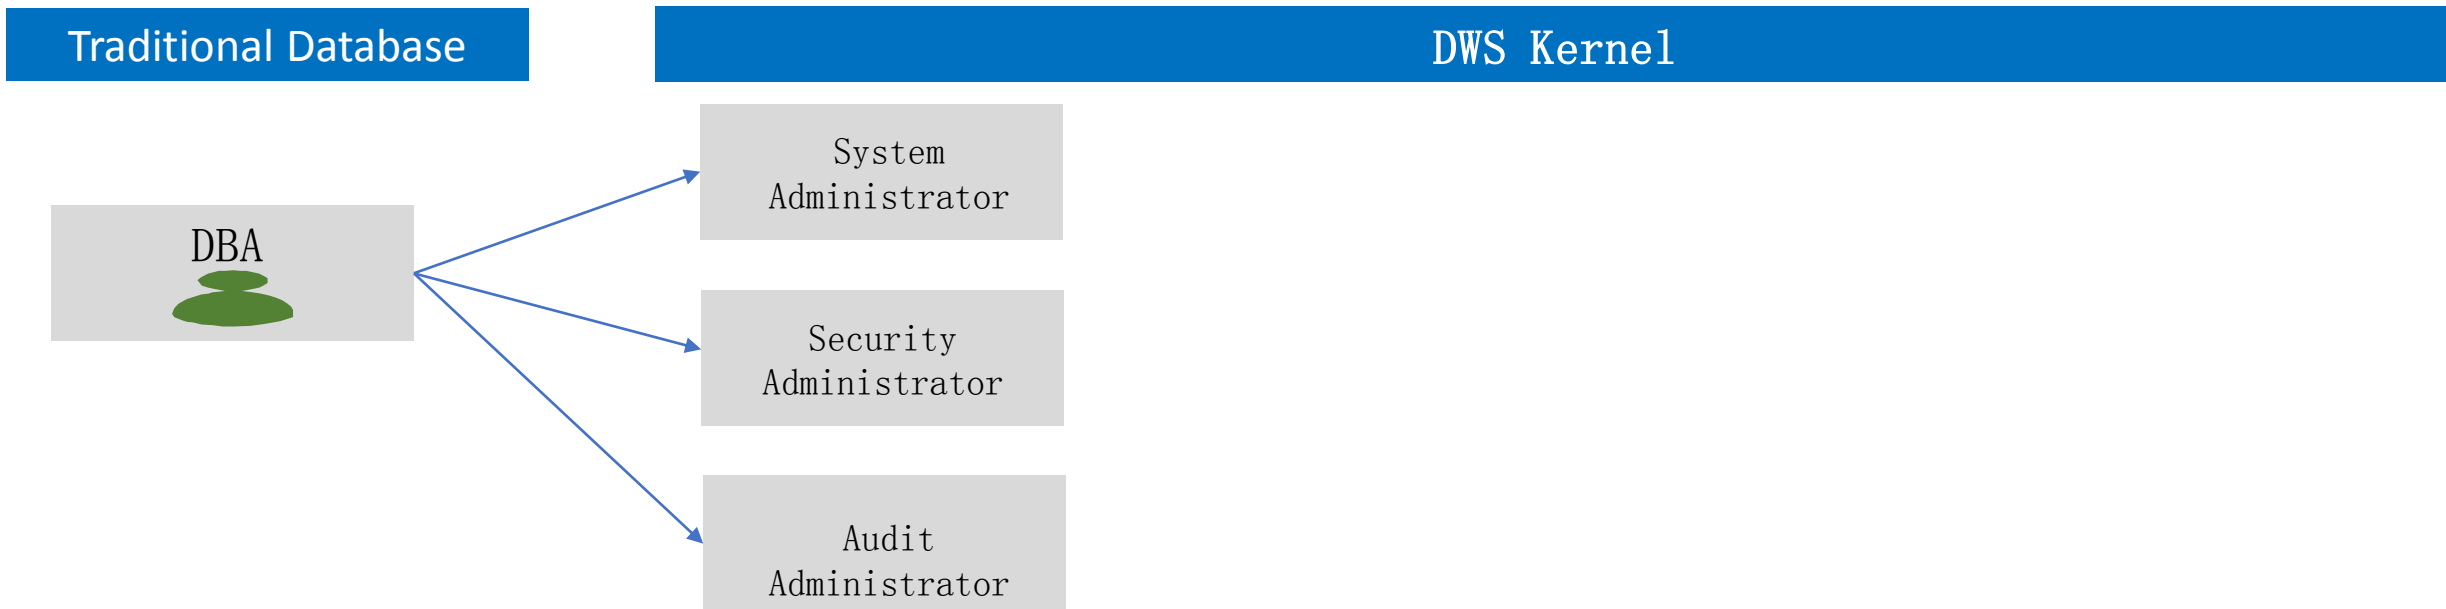
## Technical solution

- The **Node Group technology** ensures **queries in the service system are not interrupted** during the dynamic scale-out.
- **Consistent hashing** minimizes the volume of data to be **migrated** during redistribution.
- **Reentrant scale-out** uses transaction semantics. When a fault (such as a network or hardware fault, or manual cancelation of the scale-out) occurs during the scale-out, data consistency can be rolled back to the status before the scale-out using the transaction rollback mechanism in the database.

# Enterprise features

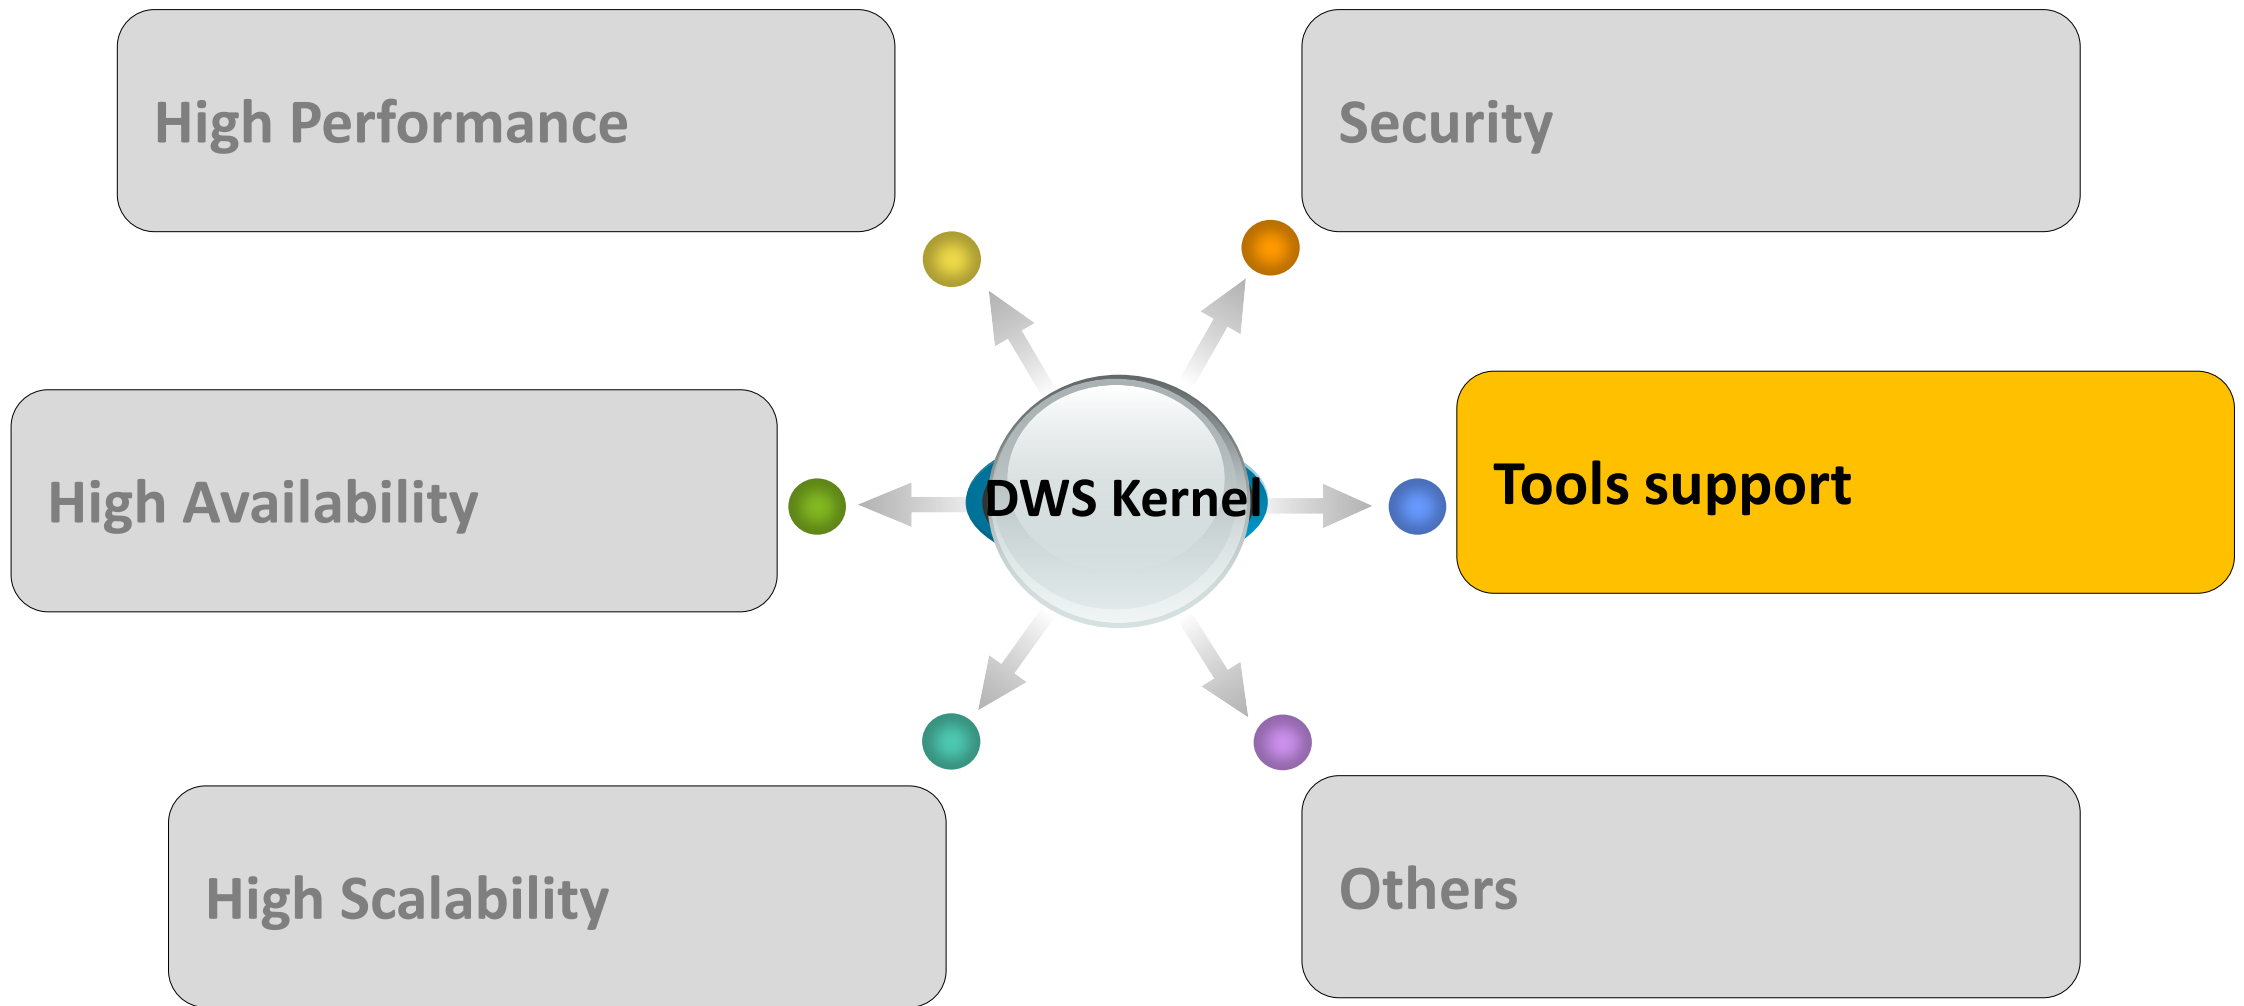


## Separation of Administration



Admin Role	Create User/Role	Audit	Admin(Create db , tablespace, etc)
System Admin	No	No	Yes
Security Admin	Yes	No	No
Audit Admin	No	Yes	No

# Enterprise features



# Ecosystem

- fully satisfy with the SQL 1992 standard and SQL 2003 standard.
- fully compatible with PostgreSQL protocol
- integrated with industry-leading tools for loading, transforming and visualizing data.

## Data Integration



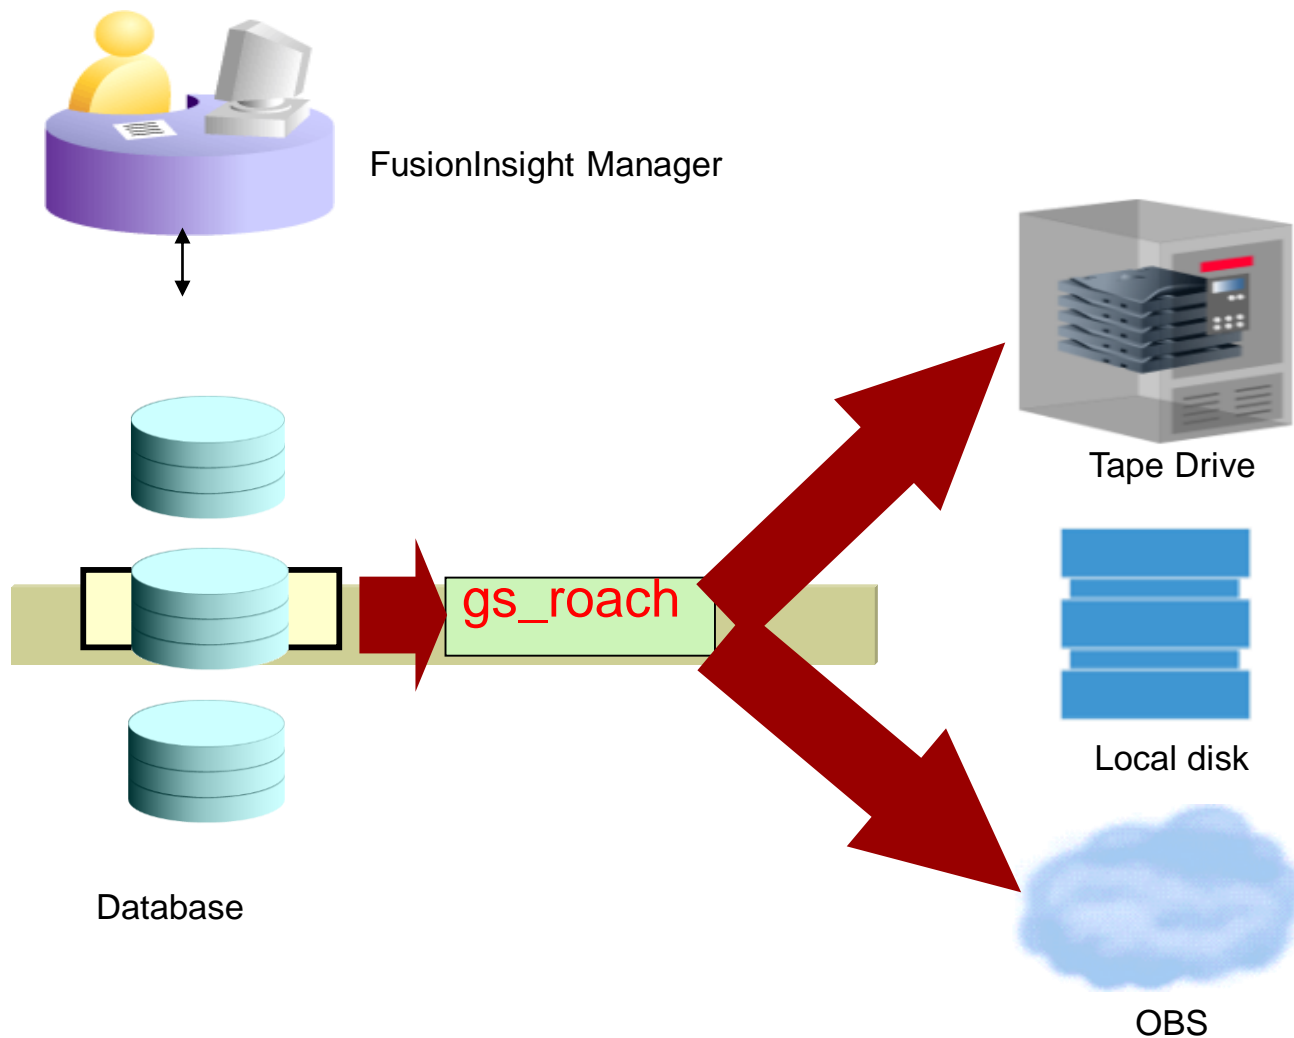
## Business Intelligence



## IDE and Others

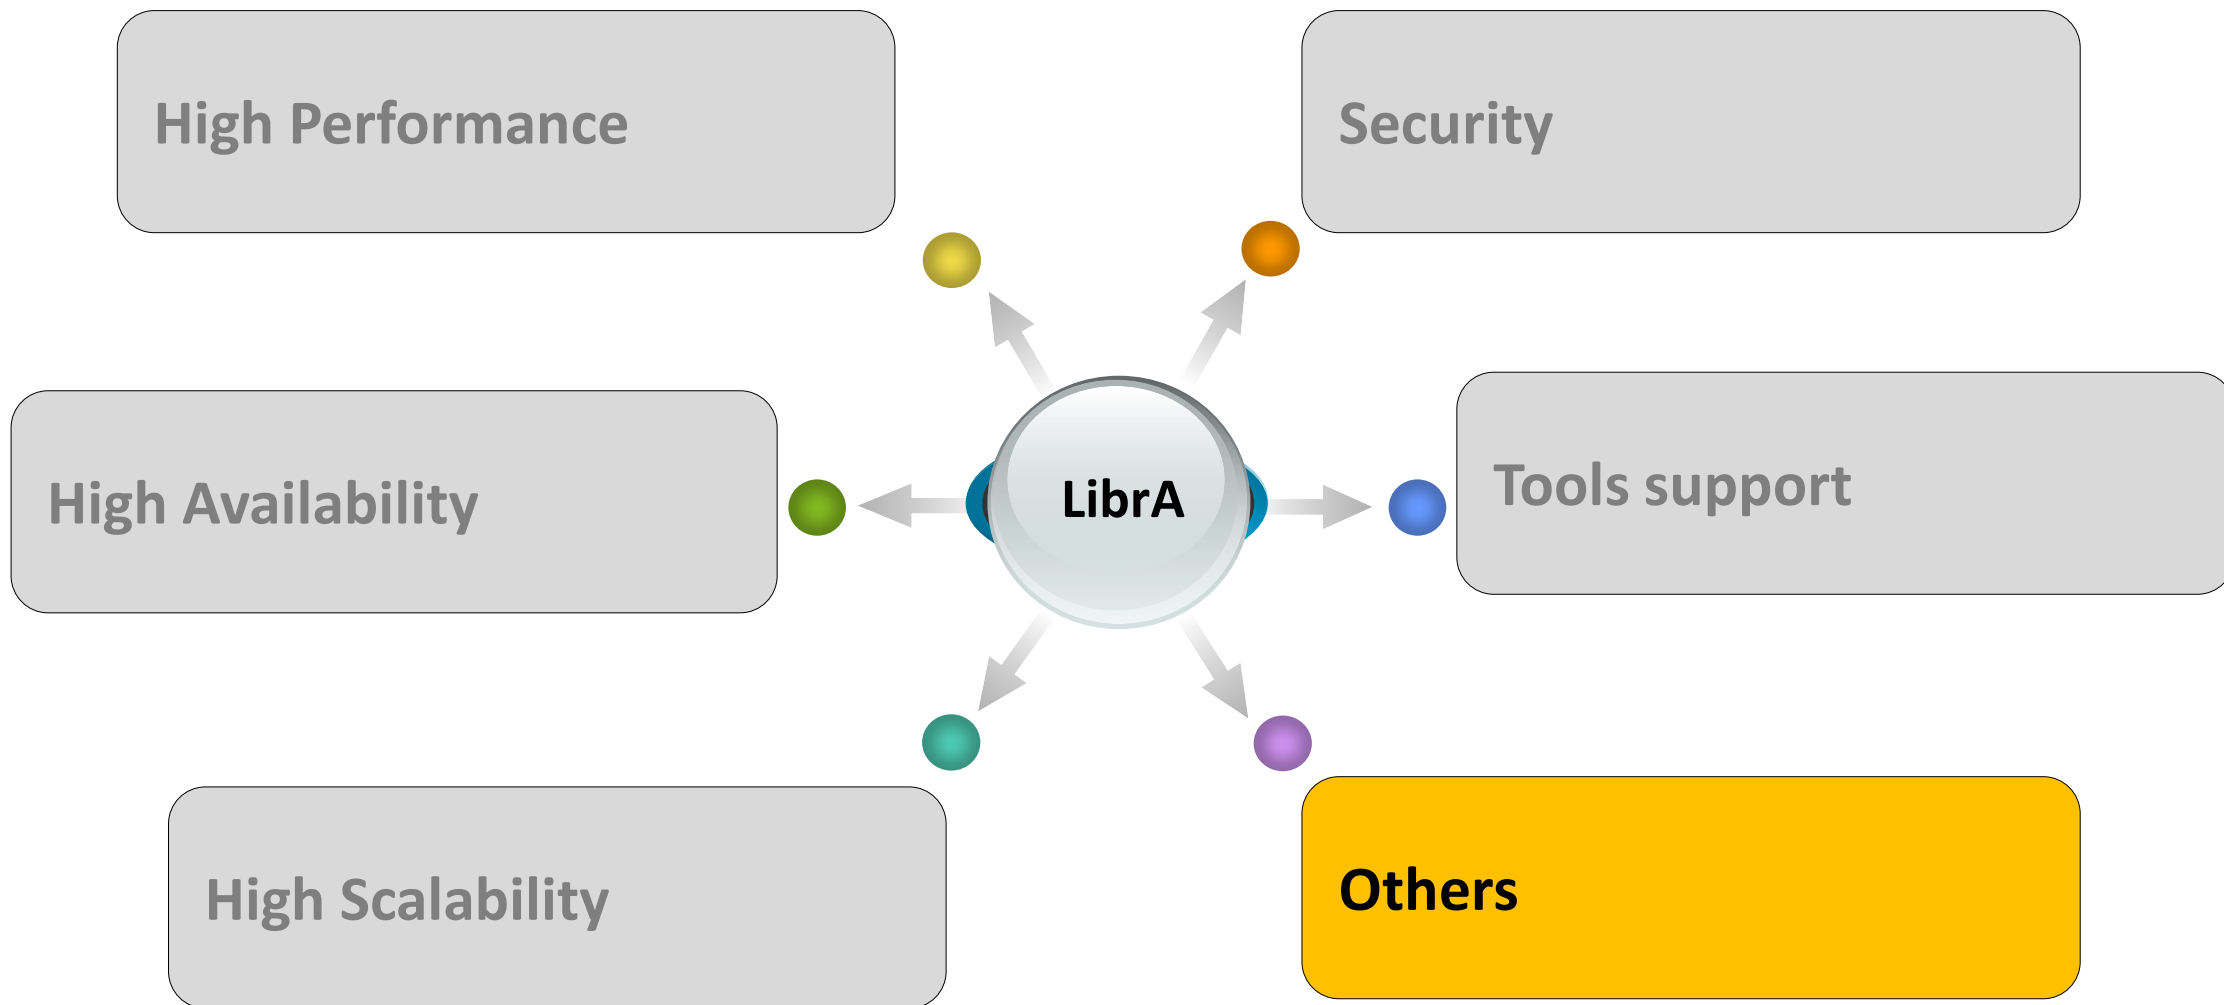


# Backup Tool (Roach)



- Back up data to tape, disk or OBS (online)
- Incremental backup/restore support

# Enterprise features





# Full-text Search

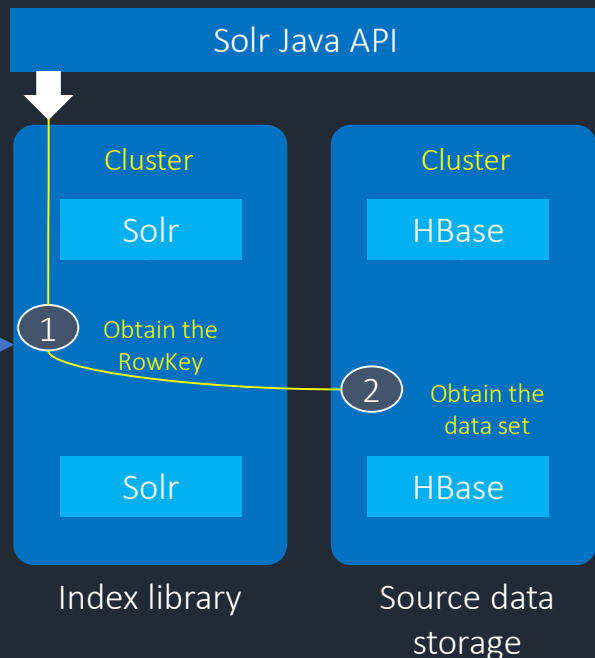
The Safe City big data project involves many full-text search scenarios, such as the police information search. Users need to search data on a big data platform like they do using Google.

**Performance:** Solr is slow in creating indexes. Chinese index occupies much more spaces. Certain scenarios cannot be covered, such as querying vehicle license plate numbers and customized ranking of search results.

Police information search

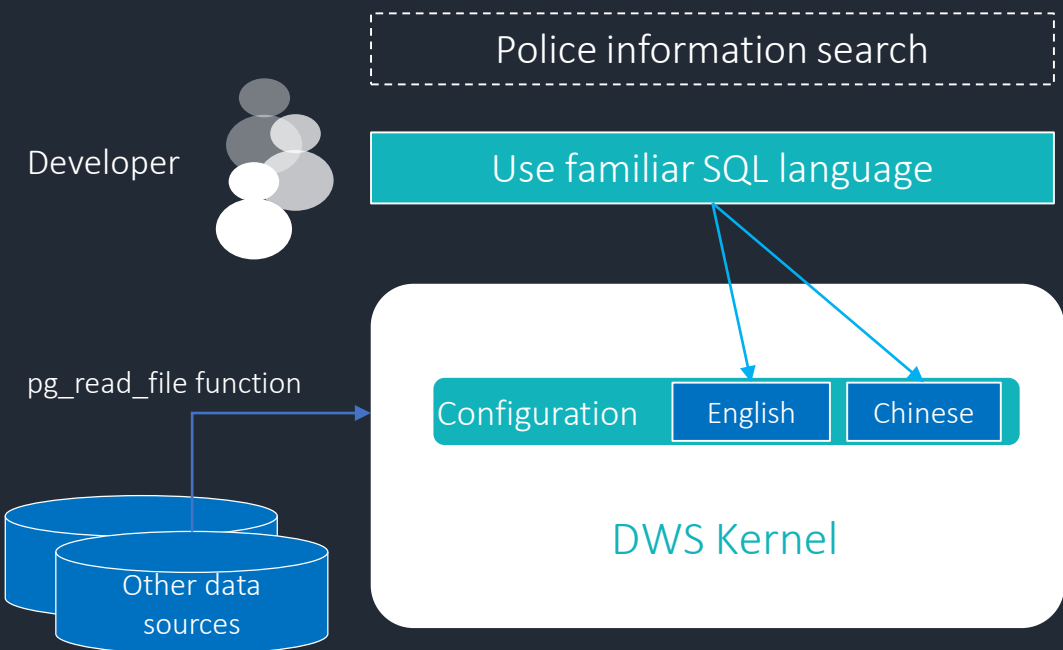
Two clusters incurring additional maintenance cost.

**Developer:** performs two queries undergoing a complex process.



**Low timeliness:** Maintaining Solr's full-text indexes is difficult when users add incremental data or delete data. Two clusters cannot synchronize data and indexes in a timely manner.

## DWS Kernel full-text search



### Customer benefit:

- Simple development: Customers do not need to introduce additional components. They can perform full-text searches on the big data platform simply by using SQL languages that the developers are familiar with.
- Low cost: Customers do not need to store additional data to Hbase.
- High timeliness: Customers can perform queries in real time after data is updated on the big data platform.

Thank you