



Luigi Troiano (M.Eng. 2000, Ph.D. 2004) is assistant professor and researcher at University of Sannio since 2006, where he is lecturing in Software Engineering and Intelligent Systems. He obtained Laurea (master degree) in IT Engineering at University of Naples “Federico II” in 2000, and Ph.D. in IT Engineering at University of Sannio in 2004. His research interests are mainly related to Computational Intelligence and Intelligent Systems, focusing on how to apply mathematical models and advanced algorithms to real-world problems. His competences refer to conception, data analysis, algorithm experimentation and validation, implementation of software systems and solutions, using the state of art of technologies. He earned an industrial experience working in Italy and abroad for multinational companies such as Pirelli Trelleborg (1999), Siemens ICN (2000-2001) and TotalFinaElf E&P Paris (2001-2002).

Maria Carmela Vitelli (Bc.Eng. 2010, M.Eng. 2012) is a PhD student at University of Sannio. She obtained Laurea (master degree) in IT Engineering at University of Sannio in Benevento in 2012. Her research is focused on Algorithms and Tools for Big Data Analysis with possible applications to fault detection, isolation and recovery (FDIR). Investigation is aimed at scaling machine learning techniques to the massive amount of data collected by telemetry sensors in order to identify patterns (generally unusual and sparse) which could provide early signals on something going to fail in the system. In particular, the approach she is investigating relies on Bayesian probabilistic reasoning and data mining techniques.



Enabling Big Data Architectures for the KAGRA Project

Luigi Troiano and Maria C. Vitelli
University of Sannio, Italy

- **Big Data**
- Hadoop & Map-Reduce
- HDFS vs. NFS
- Evolving towards a Big Data architecture
- FPGA co-processing
- Conclusions

- Very large, loosely structured data set that defies traditional storage
- Human and machine generated data
- Multiple sources
- Huge volumes of data that cannot be handled by traditional database or warehouse systems
- Mostly unstructured and grows at high velocity
- Big data doesn't always mean huge data, it means "*difficult*" data

The 4 Vs

Volume


Velocity

Variety

Veracity

- ◆ **Volume:** Data is too big to scale out
- ◆ **Velocity:** Decision window is small
- ◆ **Variety:** Multiple formats challenge integration
- ◆ **Veracity:** Same data, different interpretations

- Healthcare
- The public sector
- Retail
- Manufacturing
- Personal-location data
- Finance
...and science




Some examples:

- The 1000 Genomes Project is aimed to find most genetic variants that have frequencies of at least 1% in the populations studied. The genome of each human being is 100 GB long.
- Jack Gallant at UC Berkeley was able to recover what people were seeing by directly observing activity in their brains by using big data and statistical methods.
- The Large Hadron Collider (LHC) at CERN in Switzerland started to take data in 2009. The amount of data collected by CERN is about 25 PB a year.

- Built to run on a cluster of machines
- Scale horizontally
- Handle unstructured/semi-structured data
- Provide storage and computing

- Big Data
- **Hadoop & Map-Reduce**
- HDFS vs. NFS
- Evolving towards a Big Data architecture
- FPGA co-processing
- Conclusions



Apache Hadoop is an open source platform for data storage and processing that is

- ✓ Scalable
- ✓ Fault tolerant
- ✓ Distributed

Other features:

- Flexibility to store and mine any type of data
- Excels at processing complex data
- Scales economically



+



- Horizontal scalability
- Commodity hardware
- Fault tolerance
- Programming framework
- Organize multiple computers in a cluster in order to perform needed calculations
- Fault tolerance




- A programming model for data processing
- Consists of two phases: *Map* and *Reduce*
 - Take a large problem and divide it into sub-problems
 - Perform the same function on all sub-problems
 - Combine the output from all sub-problems
- Each phase has *key-value pairs* as input and output


- **Job**: unit of work that the client wants to be performed. It consists of the input data, the MapReduce program and configuration information
- Hadoop runs the job dividing it into **tasks**: map tasks and reduce tasks
- There are two types of nodes that control the job execution process: a **jobtracker** and a number of **tasktracker**
- Hadoop divides the input to a MapReduce job into fixed-size pieces called **splits**
- **Data locality optimization**: run the map task on a node where the input data resides in HDFS
- Map task write their output to the local disk, not to HDFS



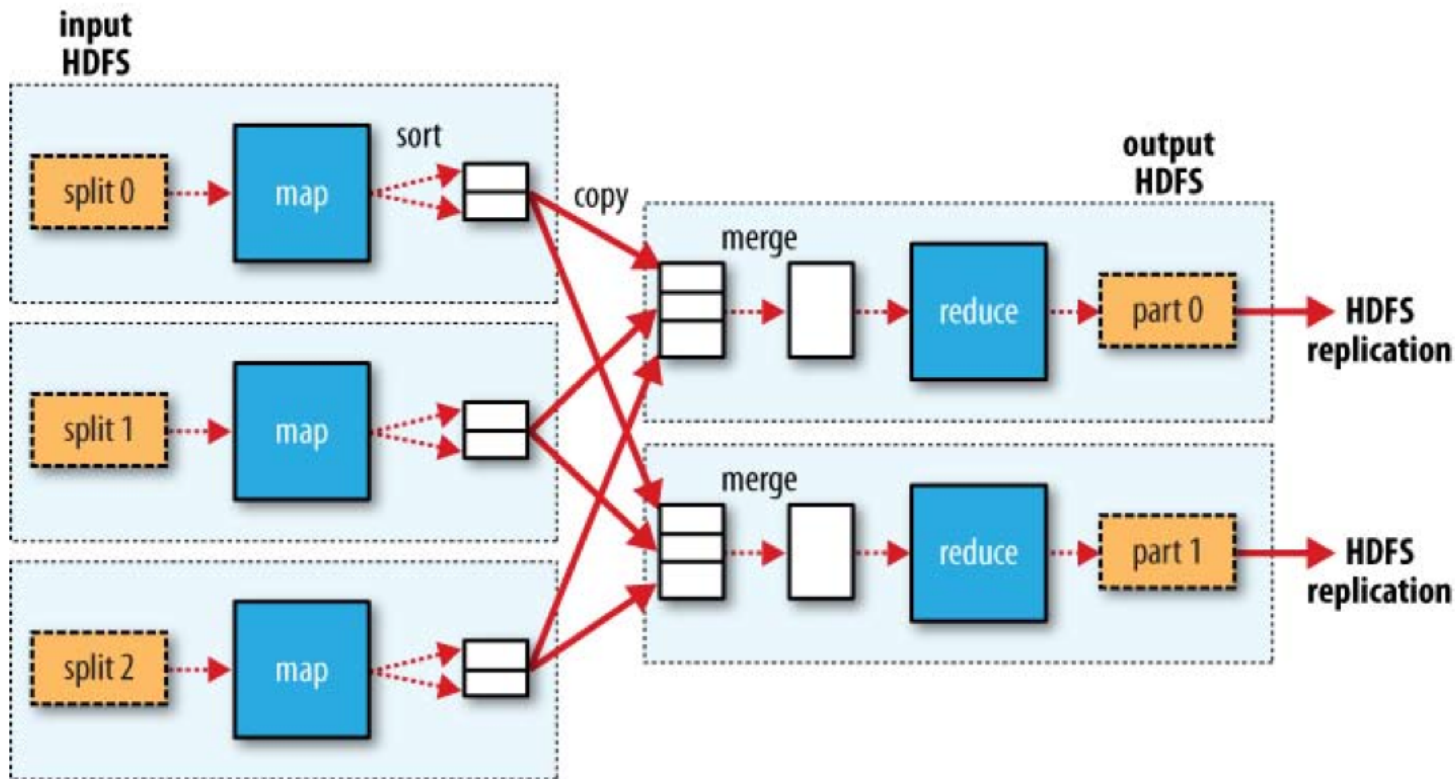
- JobTracker
- TaskTracker




- MapReduce data flow with a single reduce task



- MapReduce data flow with multiple reduce tasks
- Map tasks *partition* their output
- Data flow between map and reduce tasks is called “*the shuffle*”



- MapReduce data flow with no reduce tasks
- Appropriate when processing can be carried out entirely in parallel

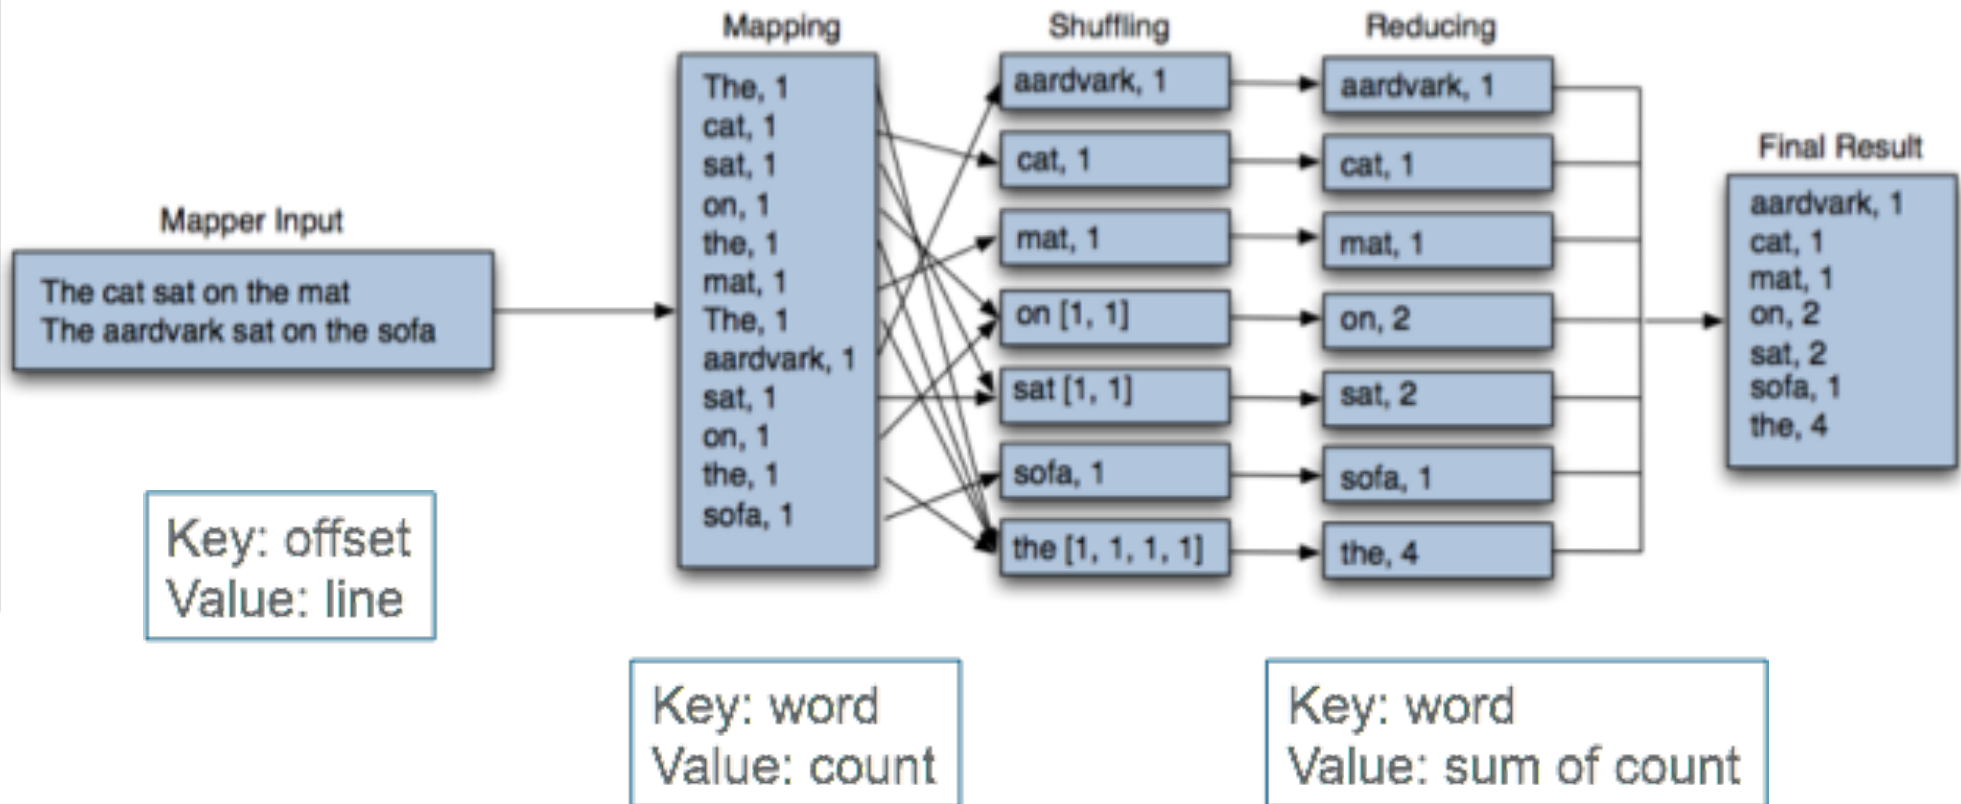


Count the number of occurrences of each word in a large amount of input data

```
map(String input_key, String input_value)
  foreach word w in input_value:
    emit(w, 1)
```

```
reduce(String output_key,
        Iterator<int> intermediate_vals)
  set count = 0
  foreach v in intermediate_vals:
    count += v
  emit(output_key, count)
```

The overall word count process

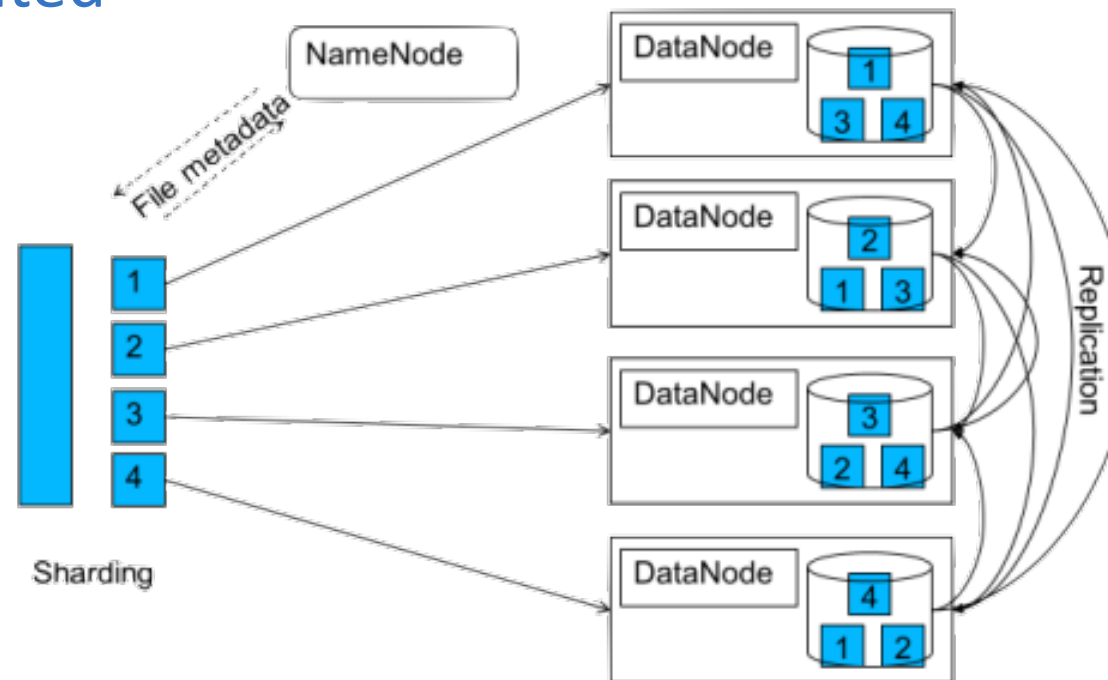


- Big Data
- Hadoop & Map-Reduce
- **HDFS vs. NFS**
- Evolving towards a Big Data architecture
- FPGA co-processing
- Conclusions

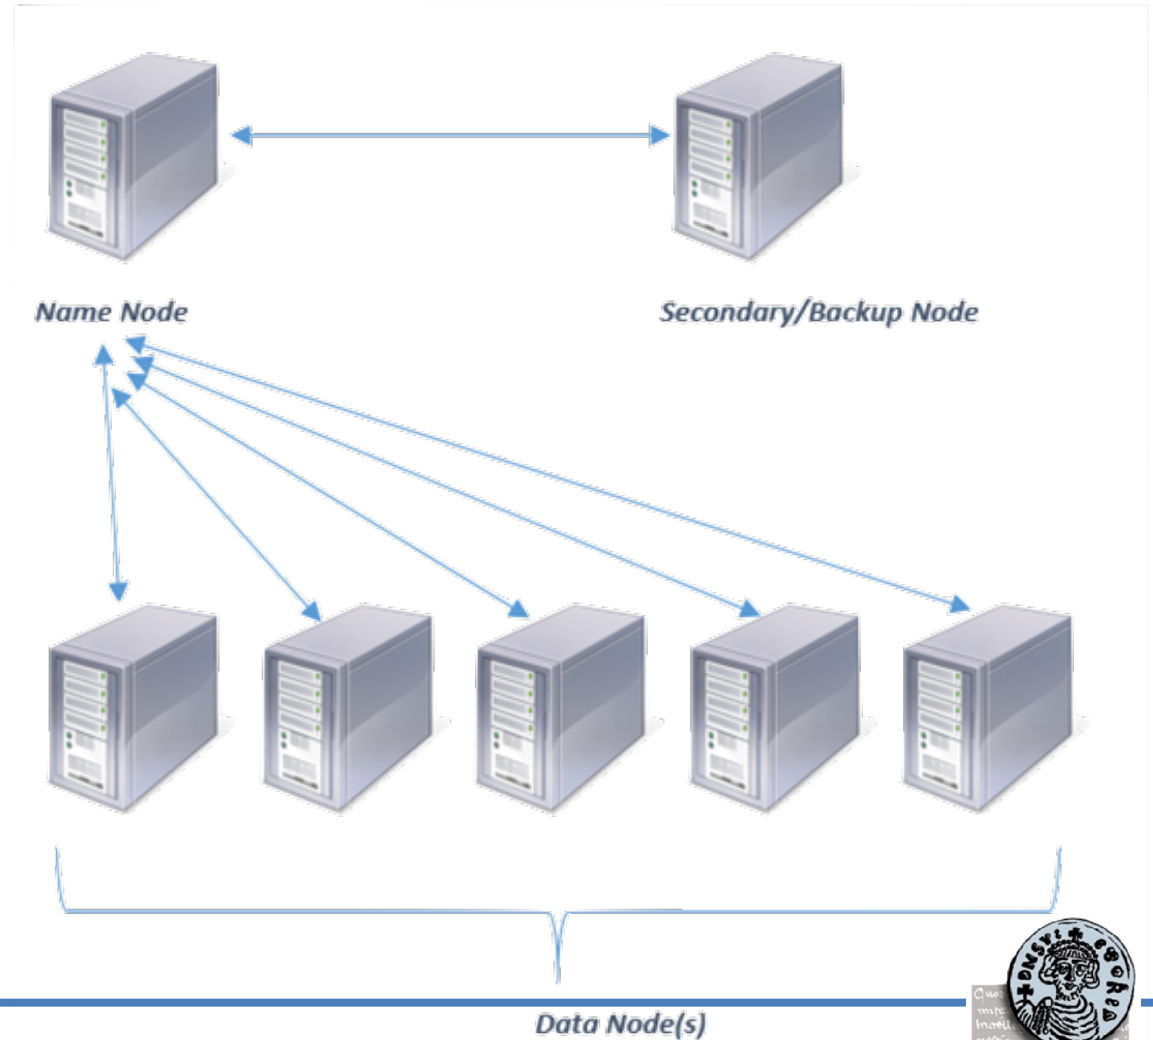


- Stores files by breaking it into smallest units called *Blocks*
- Default block size: 64MB
- Large block size to help in maintaining high throughput
- To ensure both reliability and availability each block is replicated across multiple machine on the cluster


- Master/Worker design
- HDFS is resilient (even in case of node failure)
- Data is replicated



- NameNode
- DataNode
- Secondary NameNode

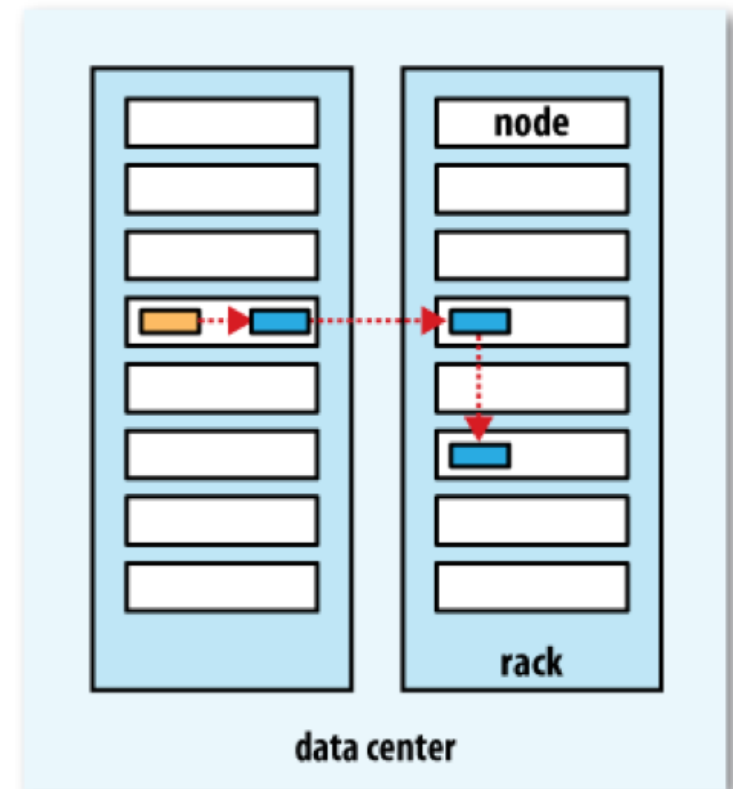


- Network is presented as a tree
- Distance between 2 nodes is the sum of their distances to their closest common ancestor



- Bandwidth available for each of the following scenarios becomes progressively less:
 - processes on the same node
 - different nodes on the same rack
 - nodes on different racks in the same data center

- Tradeoff between reliability and write/read bandwidth
- Default strategy is to place
 - the first replica on the same nodes as the client
 - the second replica on a different rack from the first, chosen at random
 - the third replica on the same rack as the second, but on a different node chosen at random
 - further replicas on random nodes on the cluster
- Block replication is across distinct datanodes



- “File System” or “Storage Layer” of Hadoop
- Designed for storing very large files (on a petabytes scale)
- Breaks incoming files into blocks and stores them redundantly across the cluster

Problem

Data is too big to store in one computer

Very high end machines are expensive

Commodity hw will fail

Hw failure may lead to data loss

Hadoop solution

Data is stored on multiple computers

Run on commodity hw

Sw is intelligent enough to deal with hw failure

Replicate (duplicate) data

- Design straightforward but also very constrained
- Provides remote access to a single logical volume stored on a single machine
- An NFS server makes a portion of its local system visible to external clients
- Important advantage: transparency
- As a distributed file system it is limited in its power: files in an NFS volume all reside on a single machine.
 - *No reliability*
 - *Possible server overload*
 - *Clients must copy data to their local machines before they can operate on it*

- Designed to store a very large amount of information
 - Spreading data across a large number of machines
 - Support much larger file sizes than NFS
- Store data reliably
 - Data should be available if individual machines in the cluster malfunction
- Provide fast, scalable access to information
 - Serve a larger number of clients by adding more machines to the cluster

HDFS


- **Distributed File System**, in which different machines are grouped for storing the data in a distributed manner
- The same data is stored in a distributed way on set of *commodity hardware*
- Allows *parallel processing* of data through MapReduce
- *Replicates the data* and thus provide *fault tolerance*

NFS

- **Network File System**, that provides a shared directory to a number of machines that can access the directory as good as a directory on local file system
- All data belonging to an entity (file or set of files) are stored on a single machine and thus require a *dedicated hardware*
- Since data is stored on a single location it has been *read sequentially*
- Since data in NFS is *stored on a single machine*, it is difficult to access/restore if the machine goes out of network

NFS	Hadoop
<i>If the volume is full, no more files can be stored</i>	<i>Information is stored across many servers and larger file sizes than NFS are supported</i>
<i>If a volume or a server fails, there is no built-in redundancy</i>	<i>Data are stored reliably</i>
Block sizes: 4-8KB	Block sizes: 64MB
<i>Most NFS administration is command line or included with overall system management tools</i>	<i>HDFS provides a web server to perform status monitoring and file browsing operations</i>

- Distributions provide easy to install mediums like RPMs
- Distros package multiple components that work well together
- Tested
- Support



Apache





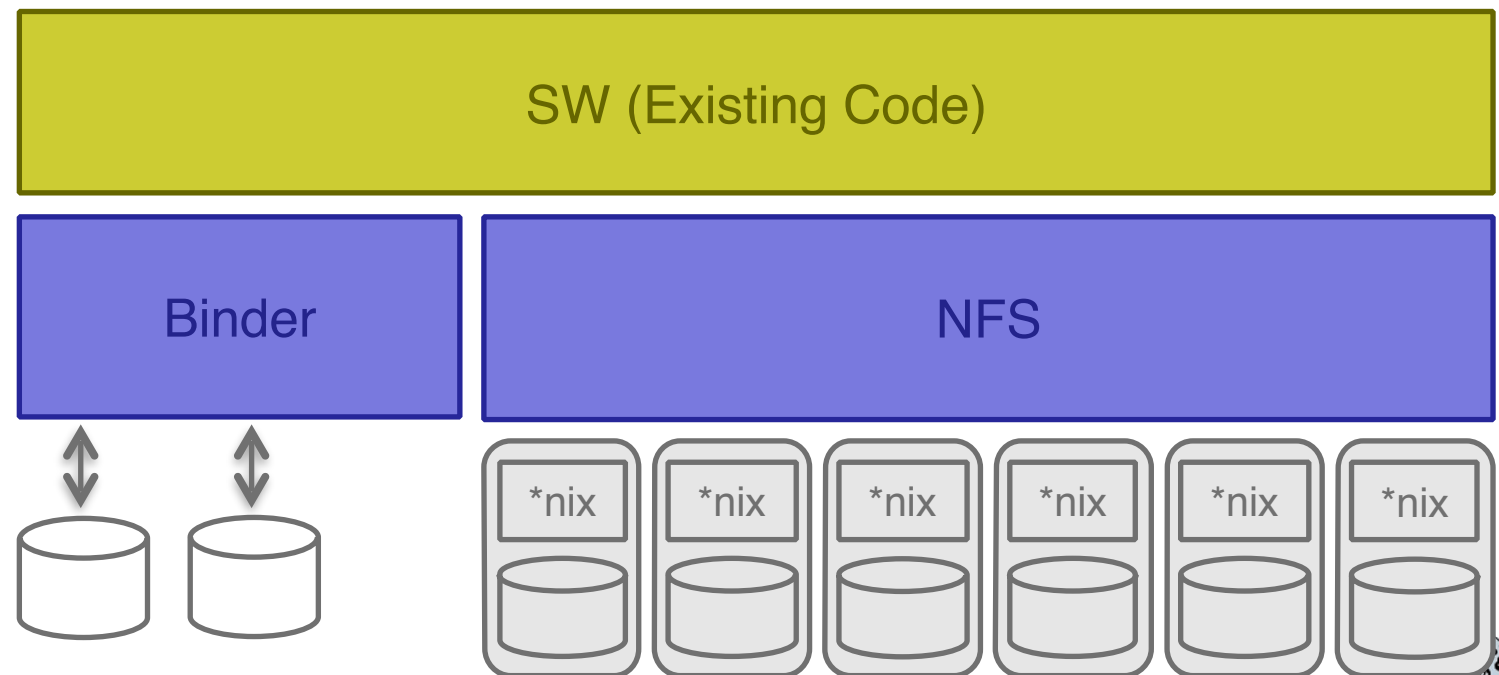
The Hadoop Ecosystem (partial)



- Big Data
- Hadoop & Map-Reduce
- HDFS vs. NFS
- **Evolving towards a Big Data architecture**
- FPGA co-processing
- Conclusions

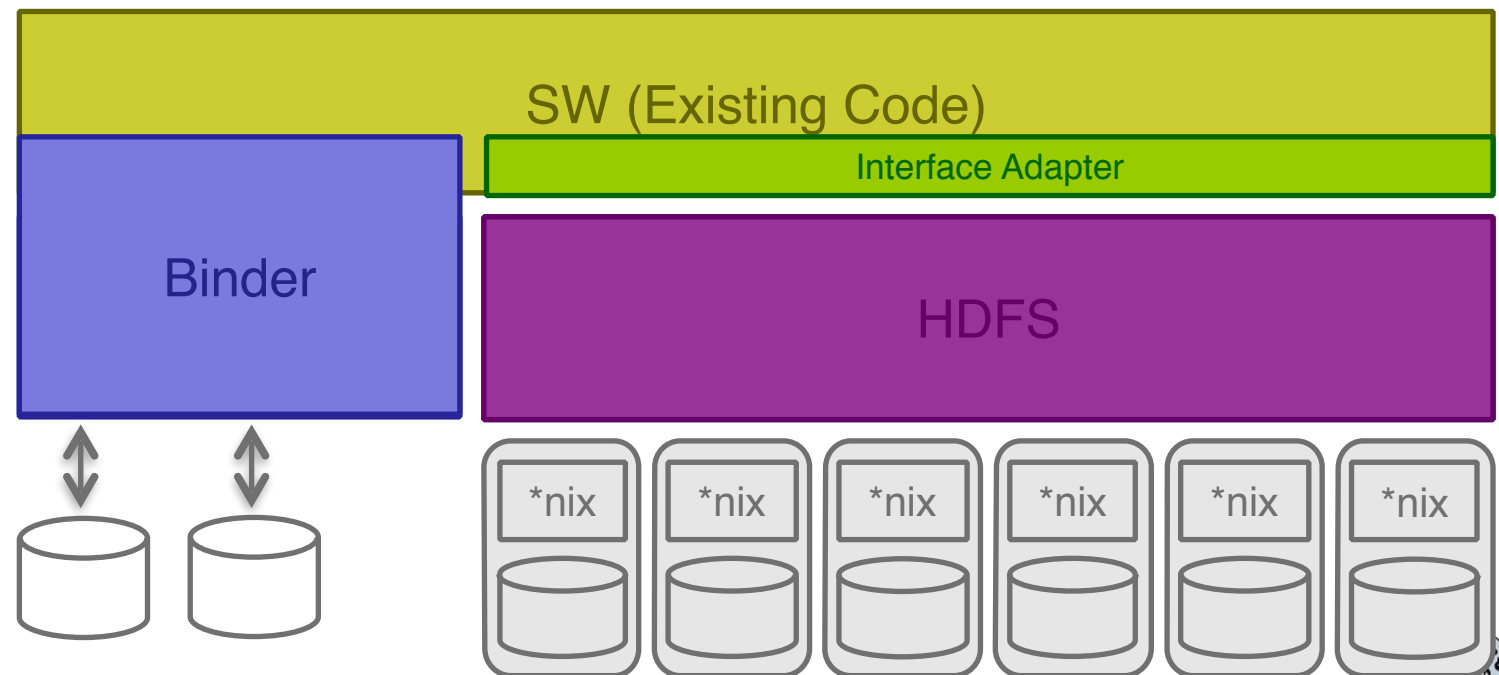
A software architecture, relying on:

- A set of software procedures
- A *nix cluster glued by NFS
- A binder to external sources

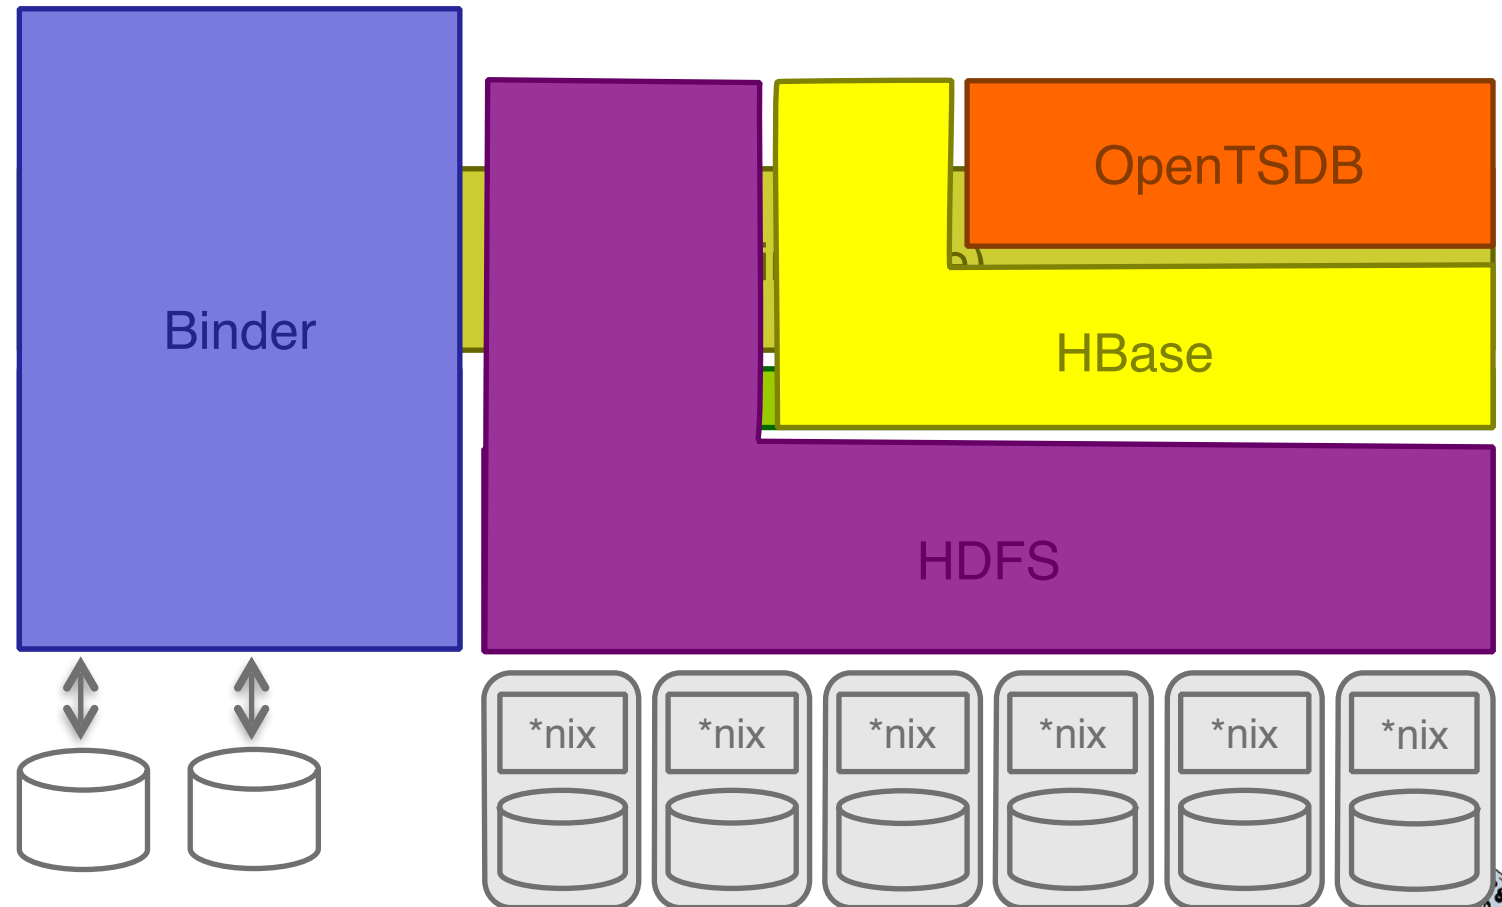


Move from NFS to HDFS:

- To achieve higher reliability
- To increase scalability



- Introduce HBASE and OpenTSDB



- Open source, distributed, versioned, column-oriented store
- Distributed Key/Value store
- Simple API (PUT, GET, DELETE, SCAN)
- Not a relational database
- Data changes through time
- Table rows are sorted by row key



Table

Row	Timestamp	Animal		Repair
		Type	Size	Cost
Enclosure1	12	Zebra	Medium	1000€
	11	Lion	Big	
Enclosure2	13	Monkey	Small	1500€

Region {

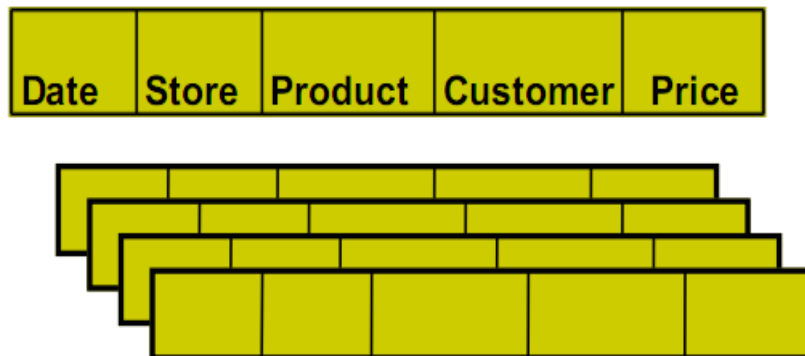
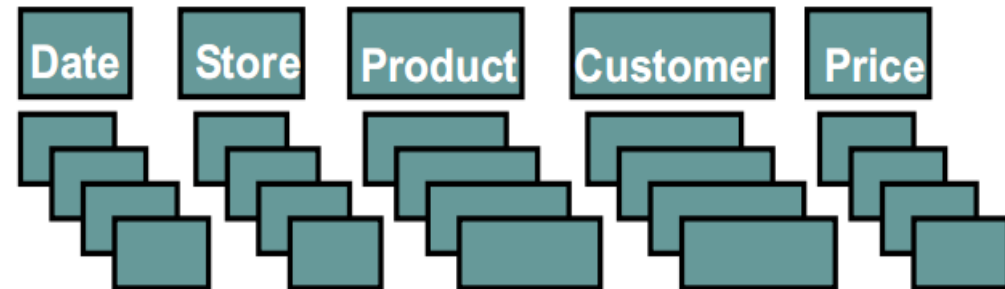
Key Column Family Cell

(Table, Row_Key, Family, Column, Timestamp) = Cell (Value)

- Row columns are grouped into *column families*
- All column families have a common prefix
- **Column-family-oriented store**: physically, all column family members are stored together on the filesystem (HFile)
- The basic unit of scalability and load balancing in HBase is called a *region*
- Regions are contiguous ranges of rows stored together
- Each region is served by one *region server*, and each of these region servers can serve many regions at any time

Column-oriented (a.k.a. vertical) databases store data with a focus on columns, instead of rows, allowing for huge data compression and very fast query times.

The downside to these databases is that they will generally only allow batch updates, having a much slower update time than traditional models.

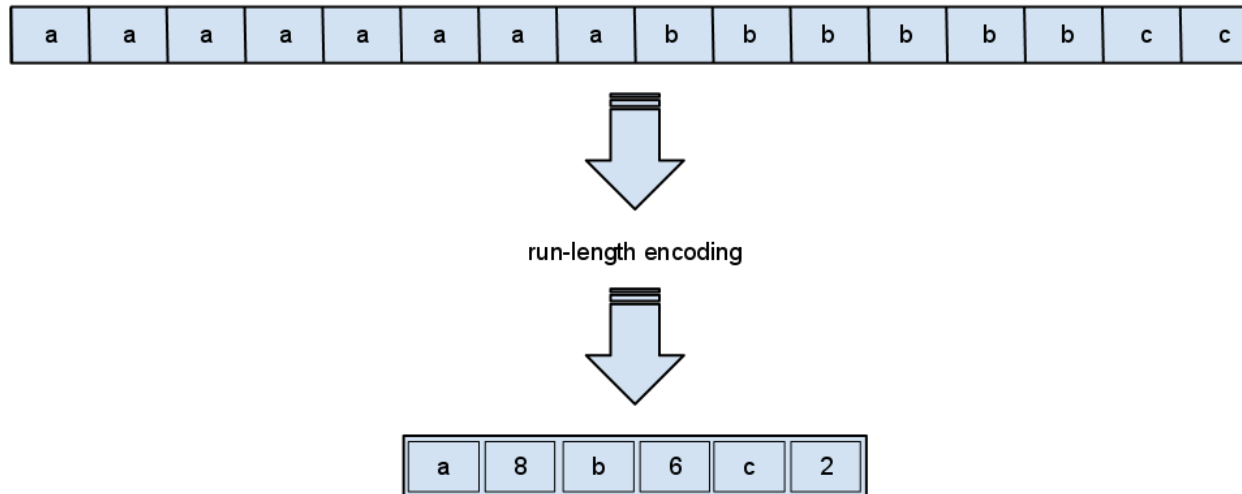
row-store**column-store**

- Column-oriented databases are suitable for read-mostly, read-intensive, large data repositories
 - OLAP, On-Line Analytical Processing
 - Big Data Analytics
- Row-oriented (conventional) databases are more suitable for accessing/update single transactions
 - OLTP, On-Line Transaction Processing
 - CRUD, Create/Read/Delete/Update activities

Row Store	Column Store
(+) Easy to add/modify a record	(+) Only need to read in relevant data
(-) Might read in unnecessary data	(-) Tuple writes require multiple accesses

- Column-oriented databases make large use of the following optimizations:
 - Compression
 - Late Materialization
 - Block Iteration
 - Invisible Join

- Low information entropy (high data value locality) leads to High compression ratio
- If data is sorted on one column that column will be super-compressible in row store
- eg. Run Length Encoding




- As result of queries we expect records
- So at some point of time multiple column must be combined
- One simple approach is to join the columns relevant for a particular query
- But further performance can be improve using late-materialization

- Delay Tuple Construction
- Might avoid constructing it altogether
- Intermediate position lists might need to be constructed
- Eg: `SELECT R.a FROM R WHERE R.c = 5 AND R.b = 10`
 - Output of each predicate is a bit string
 - Perform Bitwise AND
 - Use final position list to extract R.a

- Advantages

- Unnecessary construction of tuple is avoided
- Direct operation on compressed data
- Cache performance is improved

- Operators operate on blocks of tuples at once
- Iterate over blocks rather than tuples
- Like batch processing
- If column is fixed width, it can be operated as an array
- Minimizes per-tuple overhead
- Exploits potential for parallelism



- Invisible join is a late materialized join but minimize the values that need to be extracted out of order
- Invisible join
 - Rewrite joins into predicates on the foreign key columns in the fact table
 - These predicates evaluated either by hash-lookup
 - Or by between-predicate rewriting

```
SELECT c.nation, s.nation, d.year,  
       sum(lo.revenue) as revenue  
FROM customer AS c, lineorder AS lo,  
     supplier AS s, dwdate AS d  
WHERE lo.custkey = c.custkey  
      AND lo.suppkey = s.suppkey  
      AND lo.orderdate = d.datekey  
      AND c.region = ASIA  
      AND s.region = ASIA  
      AND d.year >= 1992 and d.year <= 1997  
GROUP BY c.nation, s.nation, d.year  
ORDER BY d.year asc, revenue desc;
```

Find Total revenue from Asian customers who purchase a product supplied by an Asian supplier between 1992 and 1997 grouped by nation of the customer, supplier and year of transaction

STEP 1

Apply region = 'Asia' on Customer table

custkey	region	nation	...
1	Asia	China	...
2	Europe	France	...
3	Asia	India	...

Hash table with keys 1 and 3

Apply region = 'Asia' on Supplier table

suppkey	region	nation	...
1	Asia	Russia	...
2	Europe	Spain	...

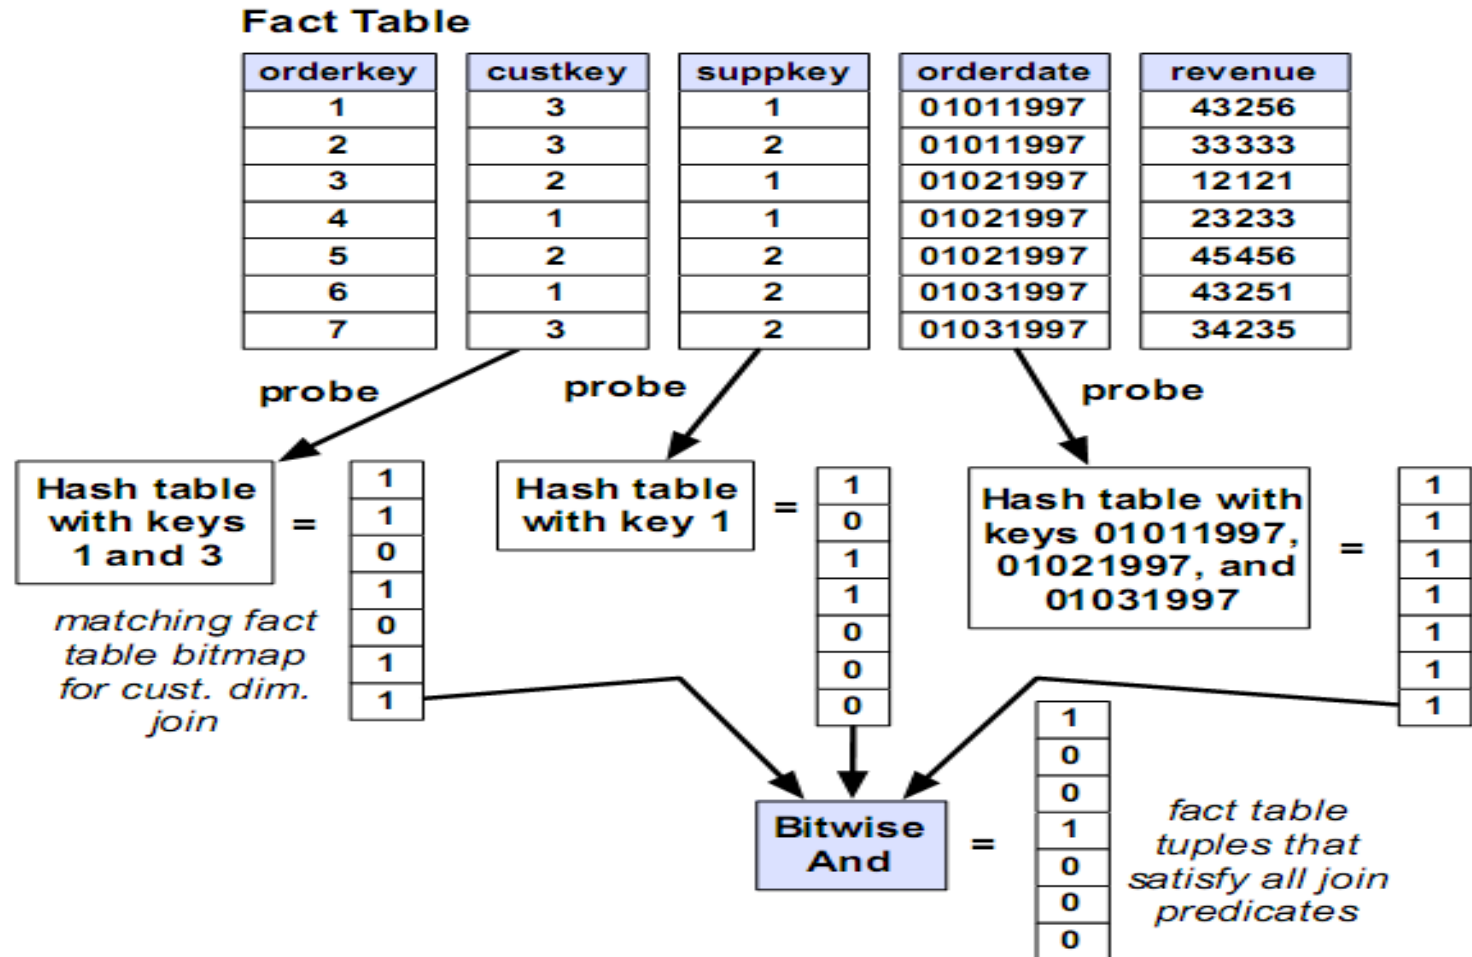
Hash table with key 1

Apply year in [1992,1997] on Date table

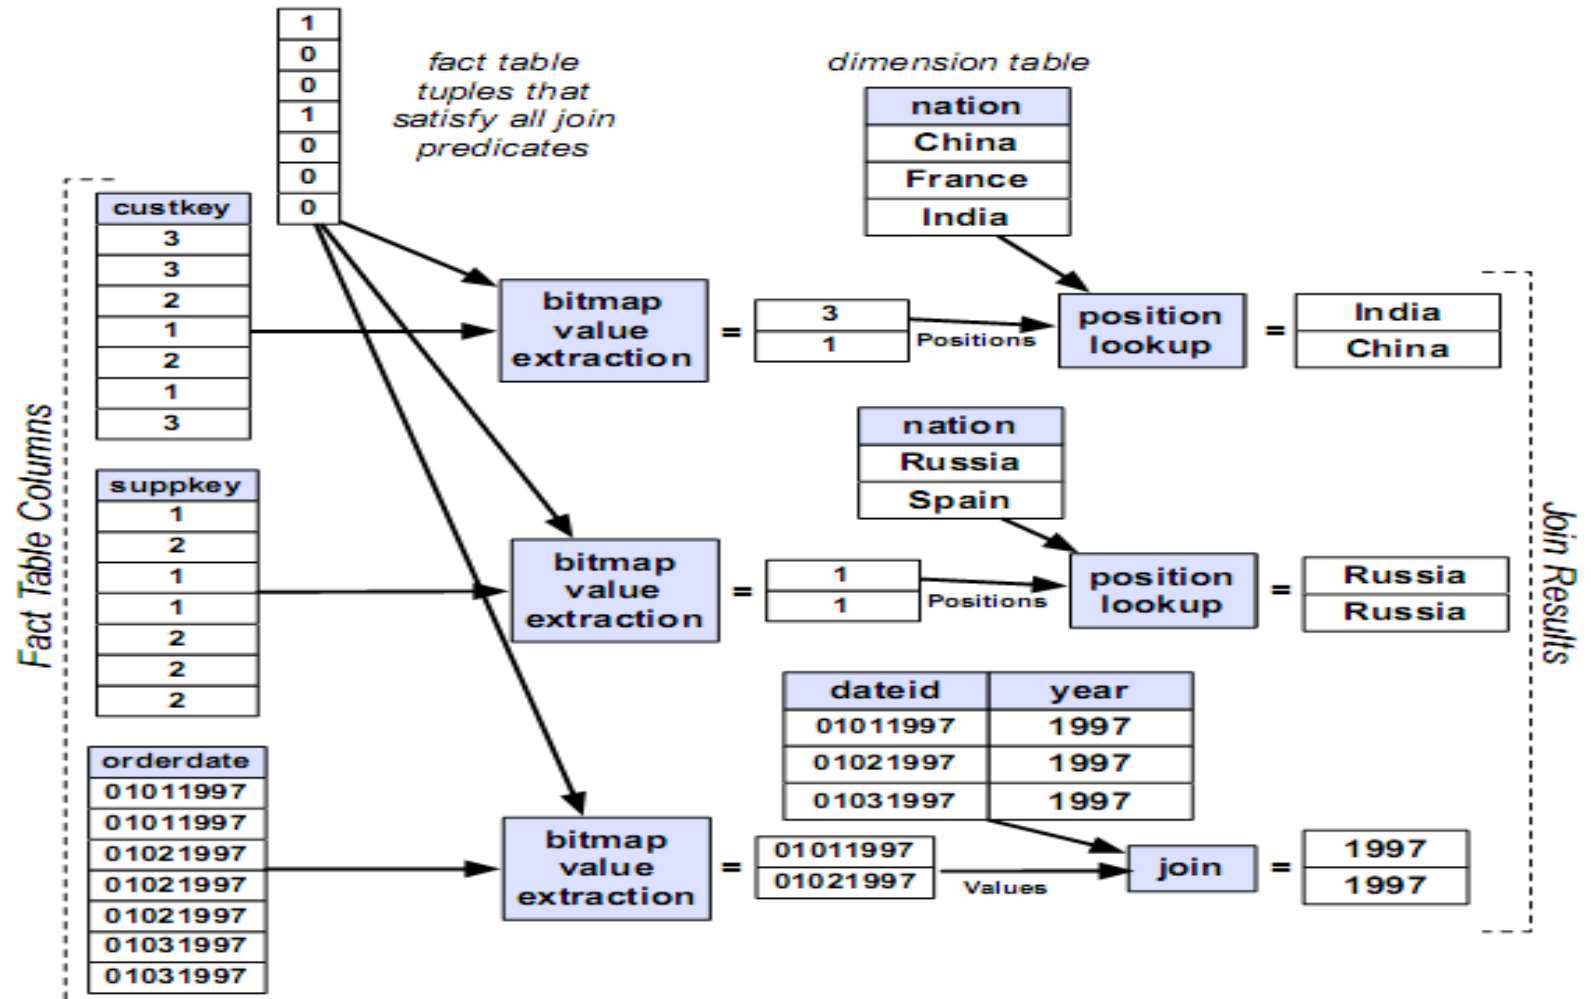
dateid	year	...
01011997	1997	...
01021997	1997	...
01031997	1997	...


Hash table with keys 01011997, 01021997, and 01031997

STEP 2



STEP 3





Baseline performance of C-Store “CS” and System X “RS”, compared with materialized view cases on the same systems.

- RS: Conventional Data Base System (Not Mentioned)
- CS: Base C-Store case
- RS (MV): System X with optimal collection of MVs
- CS (Row-MV): Column store constructed from RS(MV)

OpenTSDB

*“OpenTSDB is a distributed, scalable Time Series Database written on top of HBase. OpenTSDB was written to address a common need: **store, index, and serve metrics collected from computer systems at a large scale, and make this data easily accessible and graphable**”*

- Internal data architecture supporting very high-performance data recording

- ***Time series*** data is defined as a sequence of data points measured typically at successive times spaced at uniform time intervals
- With time series data, not only it is possible to determine the sequence in which some events happened, but it is also possible to correlate different types of events or conditions that co-occur

- Optimized for best performance for queries based on range of time
- NoSQL approaches
- Advantages in flexibility and performance
- There are different methods to store and access TS:
 - Flat files
 - RDBMS
 - NoSQL non-relational db

- **Parchet** is an effective and simple, modern format that can store the time and the number of optional values
- Problem: as the number of time series in a single file increases the fraction of usable data for any particular query decreases

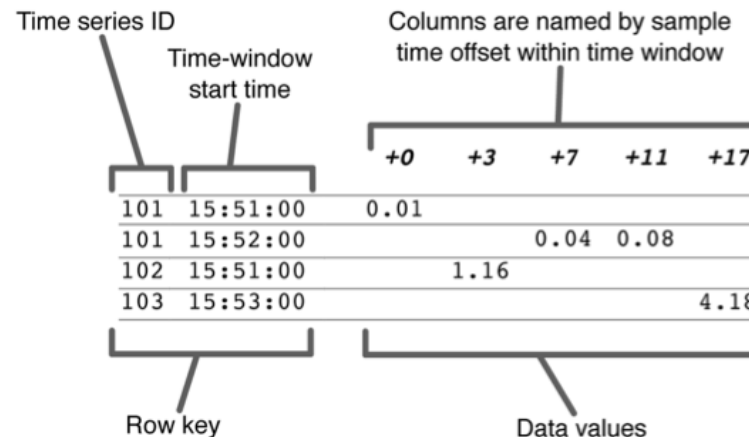
```
message simpleSeries {  
  repeated group sample {  
    required float t;  
    optional float tempIn;  
    optional float pressureIn;  
    optional float tempOut;  
    optional float pressureOut;  
  }  
}
```

```
message fancySeries {  
  repeated group block {  
    repeated group tags {  
      optional string name;  
      optional string value;  
    }  
    repeated float time;  
    repeated float value;  
  }  
}
```

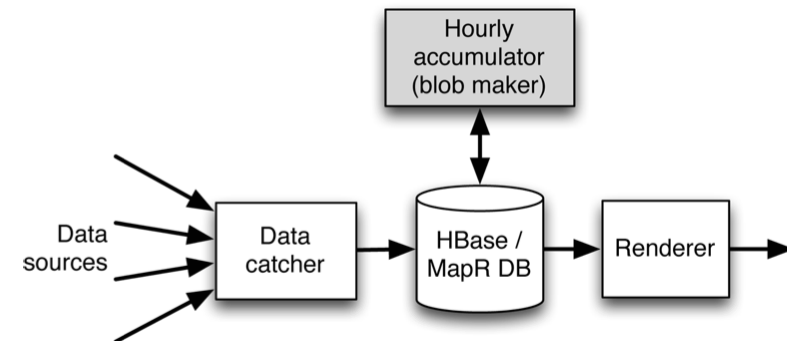
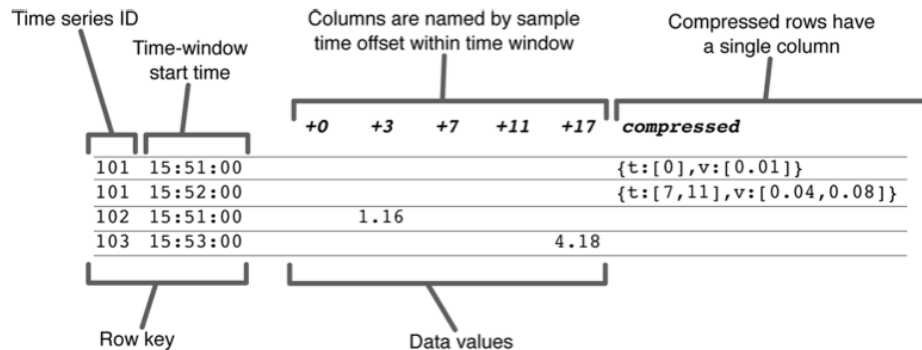
- Time, series ID and a value are stored
- Details of the series are stored in a dimension table
- Problem: use of one row per measurement

Time	Time series ID	Value
15:51:00	101	0.01
15:51:03	102	1.16
15:52:07	101	0.04
15:52:11	101	0.08
15:53:17	103	4.18

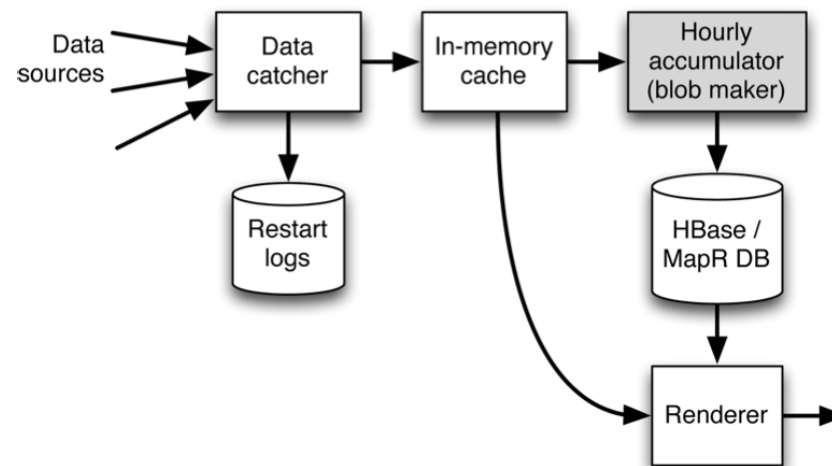
- Store many values in each row
- Allowing data point to be retrieved at a higher speed
- Rows containing data from a single time series to wind up near one another on disk



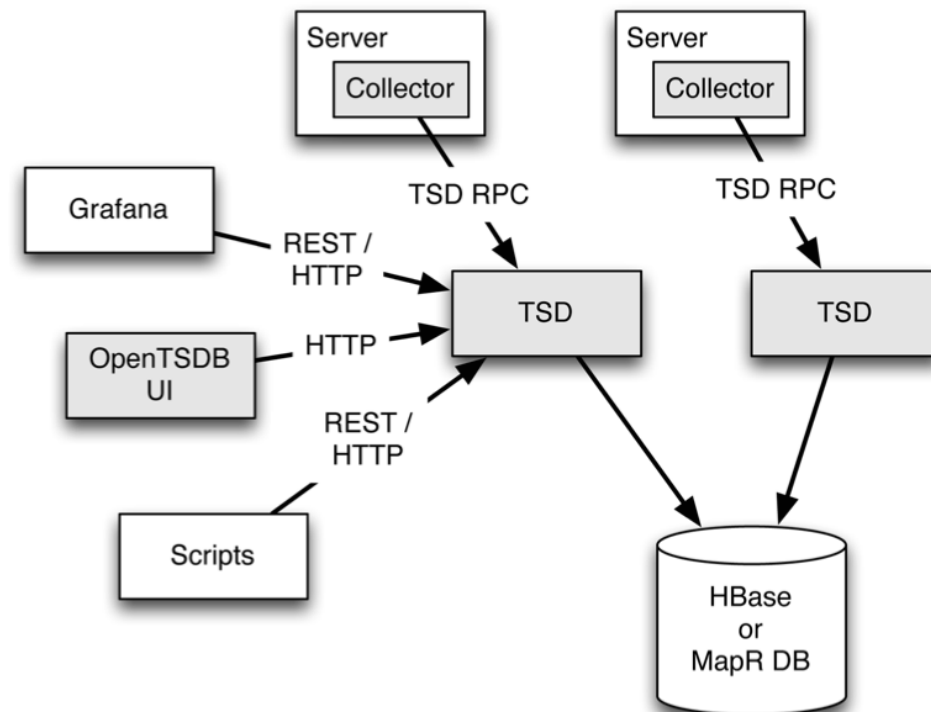
- Collapse all of data for a row into a single data structure known as blob
- Blob can be highly compressed so that less data needs to be read from disk



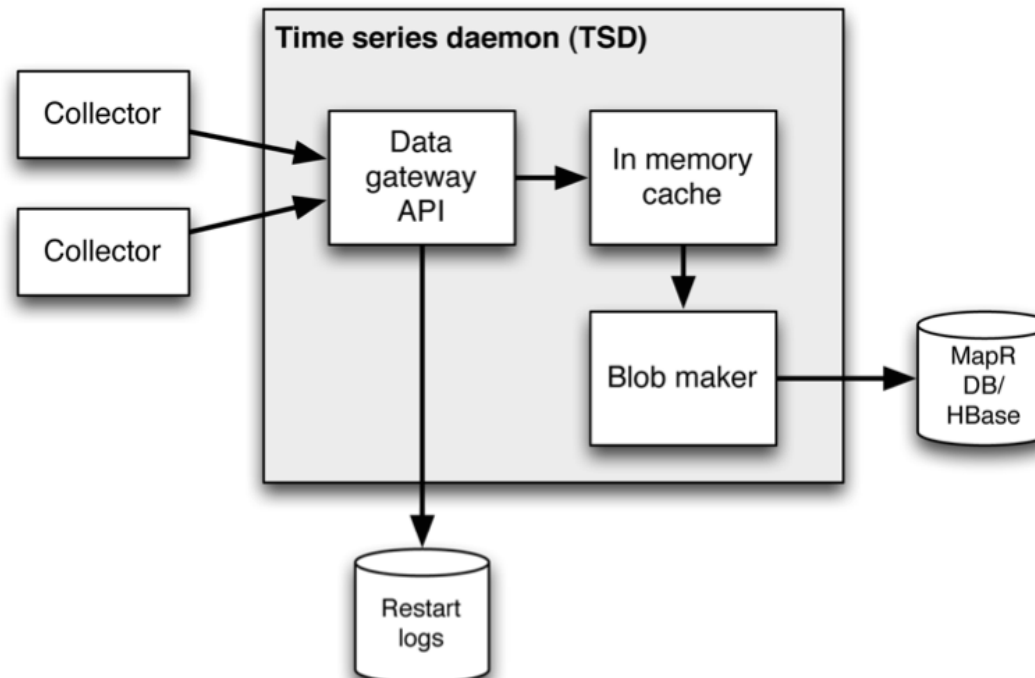
- The blob maker uses incoming data from a memory cache
- The full data stream is only written to the memory cache
- Data is not written to the storage tier until it is compressed into blobs



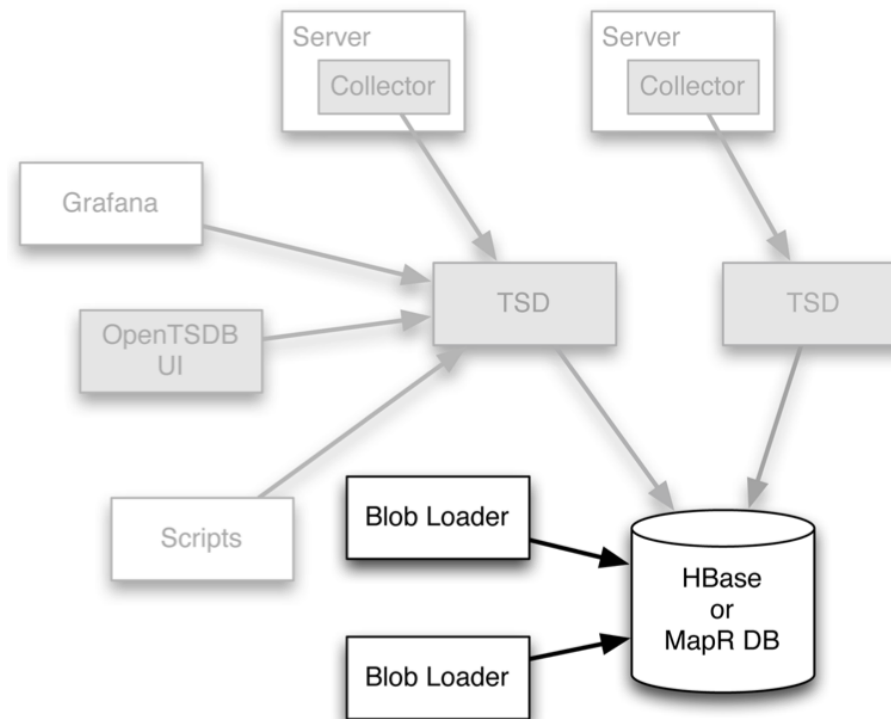
- In OpenTSDB, a time series data point consists of:
 - A metric name.
 - A UNIX timestamp (seconds or milliseconds since Epoch).
 - A value (64 bit integer or single-precision floating point value).
 - A set of tags (key-value pairs) that describe the time series the point belongs to.
- Data collectors
- Time-series daemons (TSD)
- User interface functions
- Scripts/Tools



- Data is ingested initially to the storage tier in the blob-oriented format that stores many data points per row

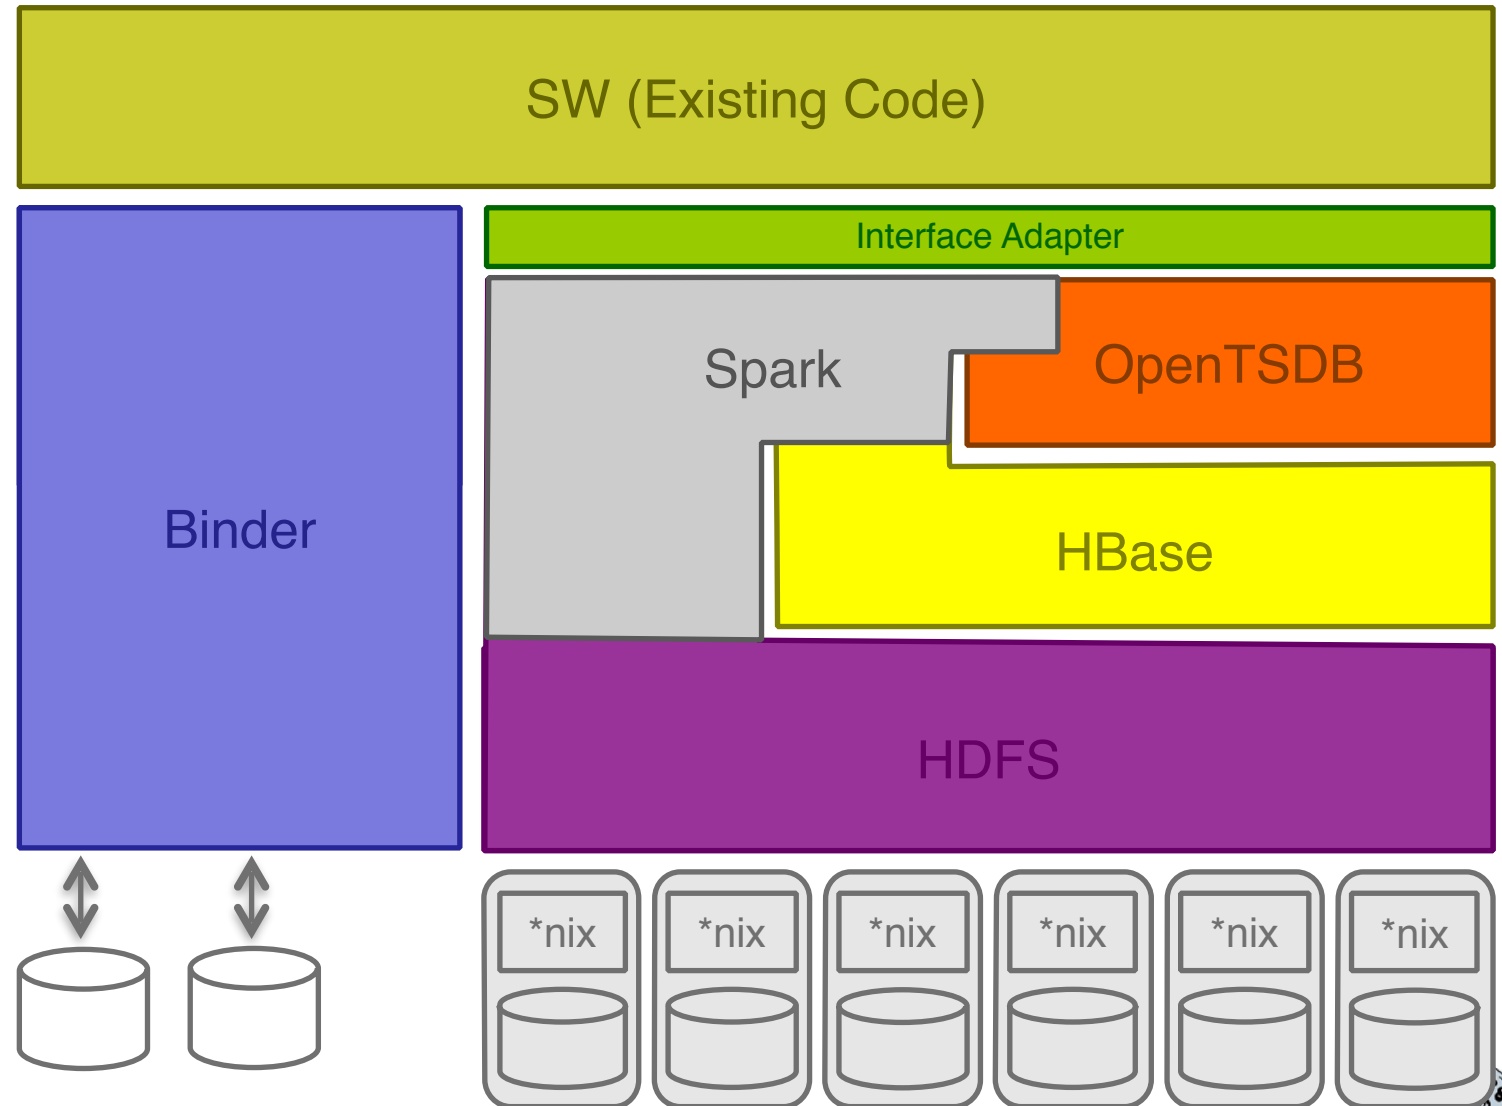


- Separate data flow that loads data in blob-style format directly to the storage tier independently of the TSD



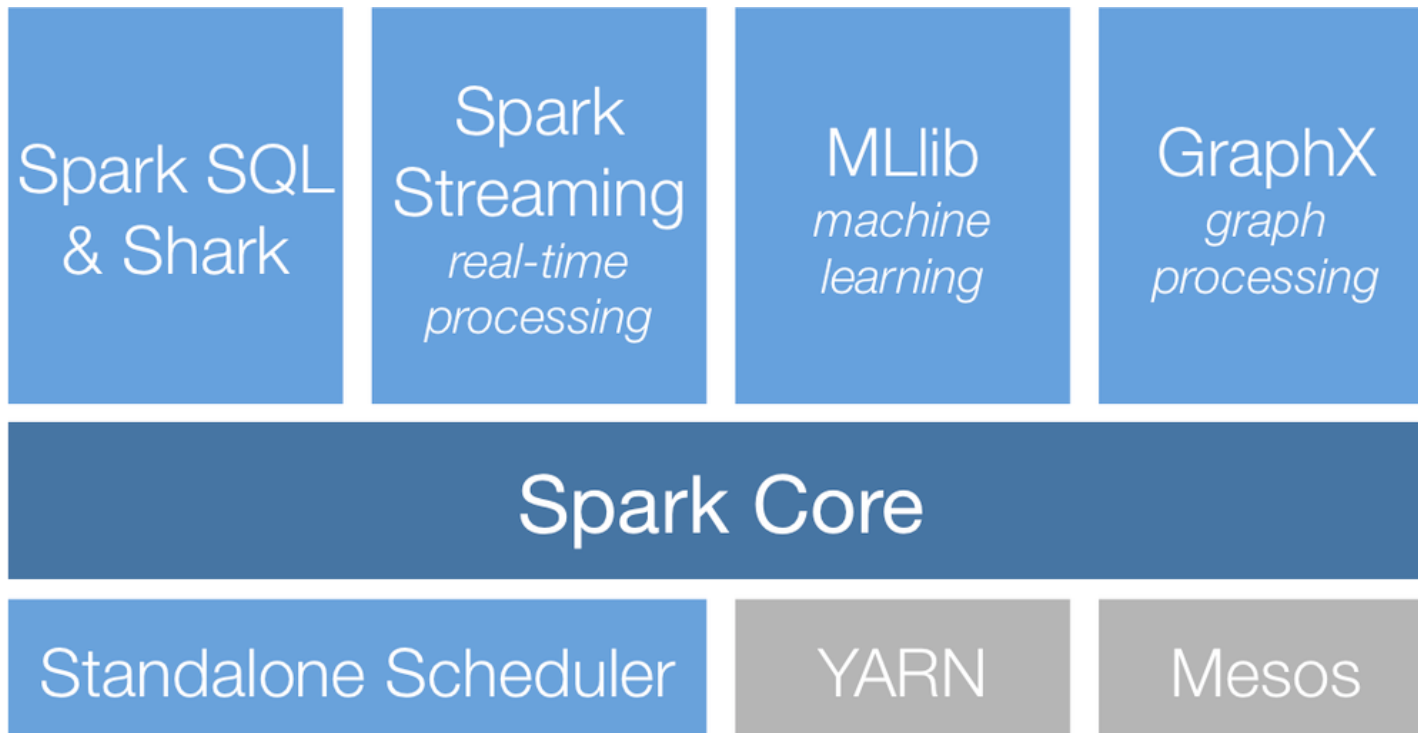
- Steps to transform raw data into processed data:
 - selection
 - grouping
 - down-sampling
 - aggregation
 - interpolation
 - rate conversion

Place a unified mediator to distribute the job across nodes and different layers

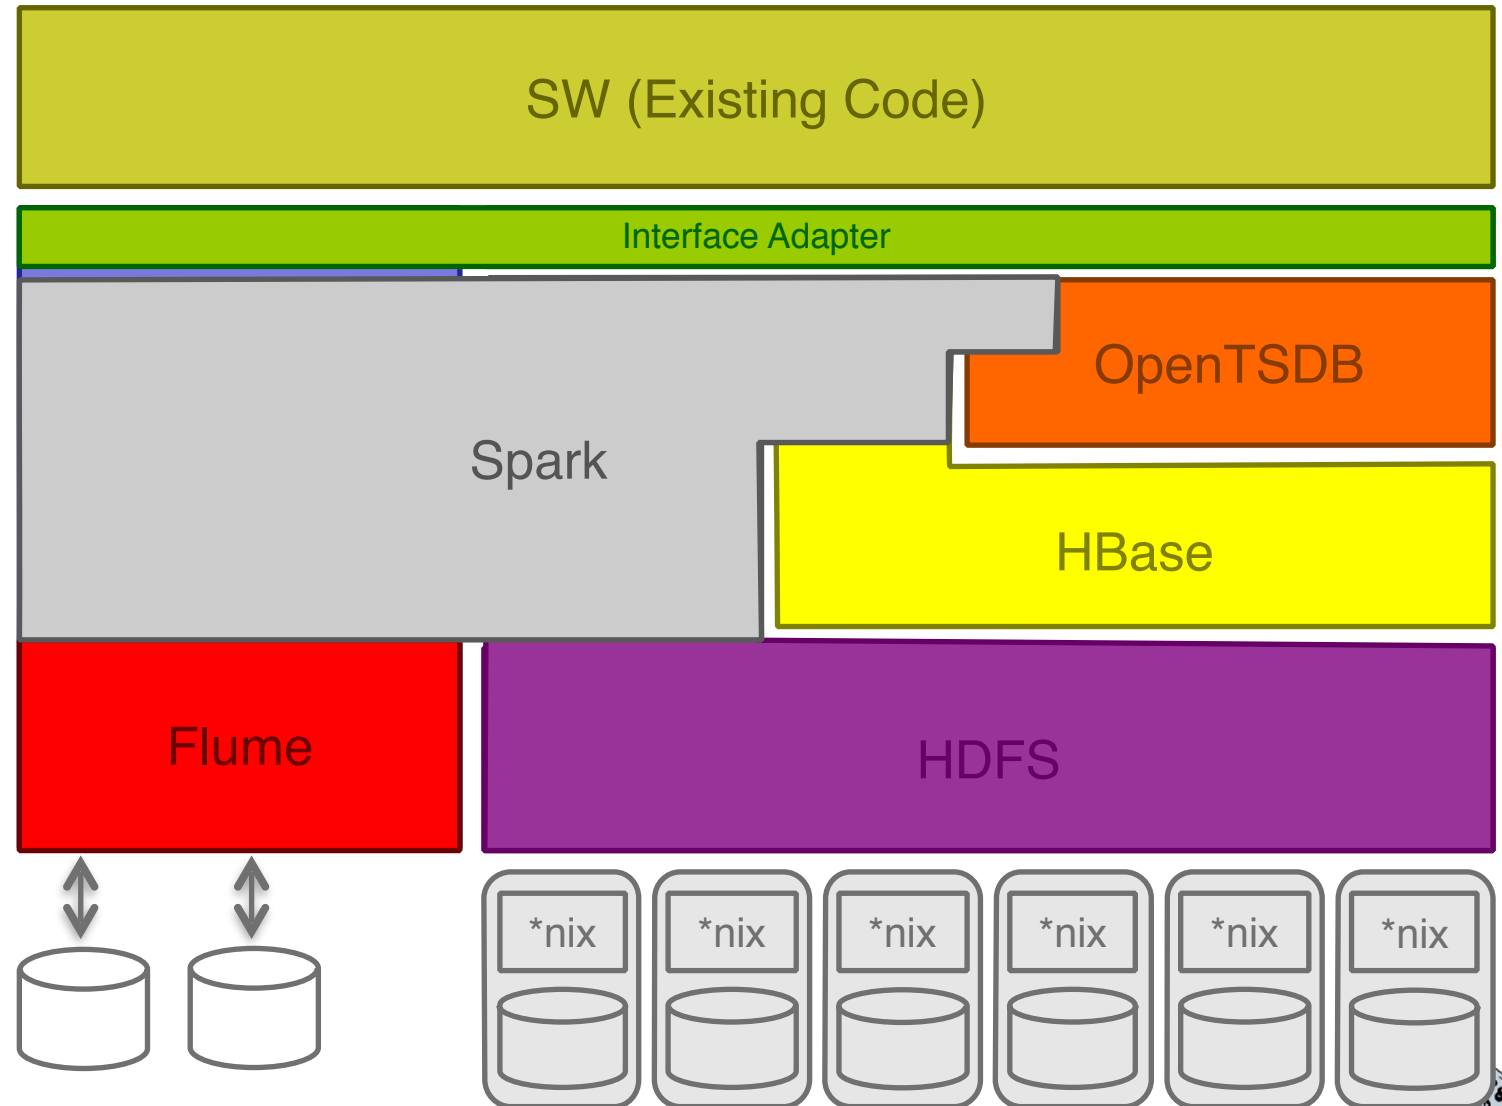




- Open source project at the U.C. Berkeley AMPLab
- Initially developed for two applications where keeping data in memory helps
 - iterative algorithms (common in machine learning)
 - interactive data mining
- Compared to MapReduce, Spark differs in two things
 - Spark holds intermediate results in memory (rather than writing them to disk)
 - Spark supports more than just map and reduce functions (greatly expanding the set of possible analyses that can be executed over HDFS data)





Make integration seamless




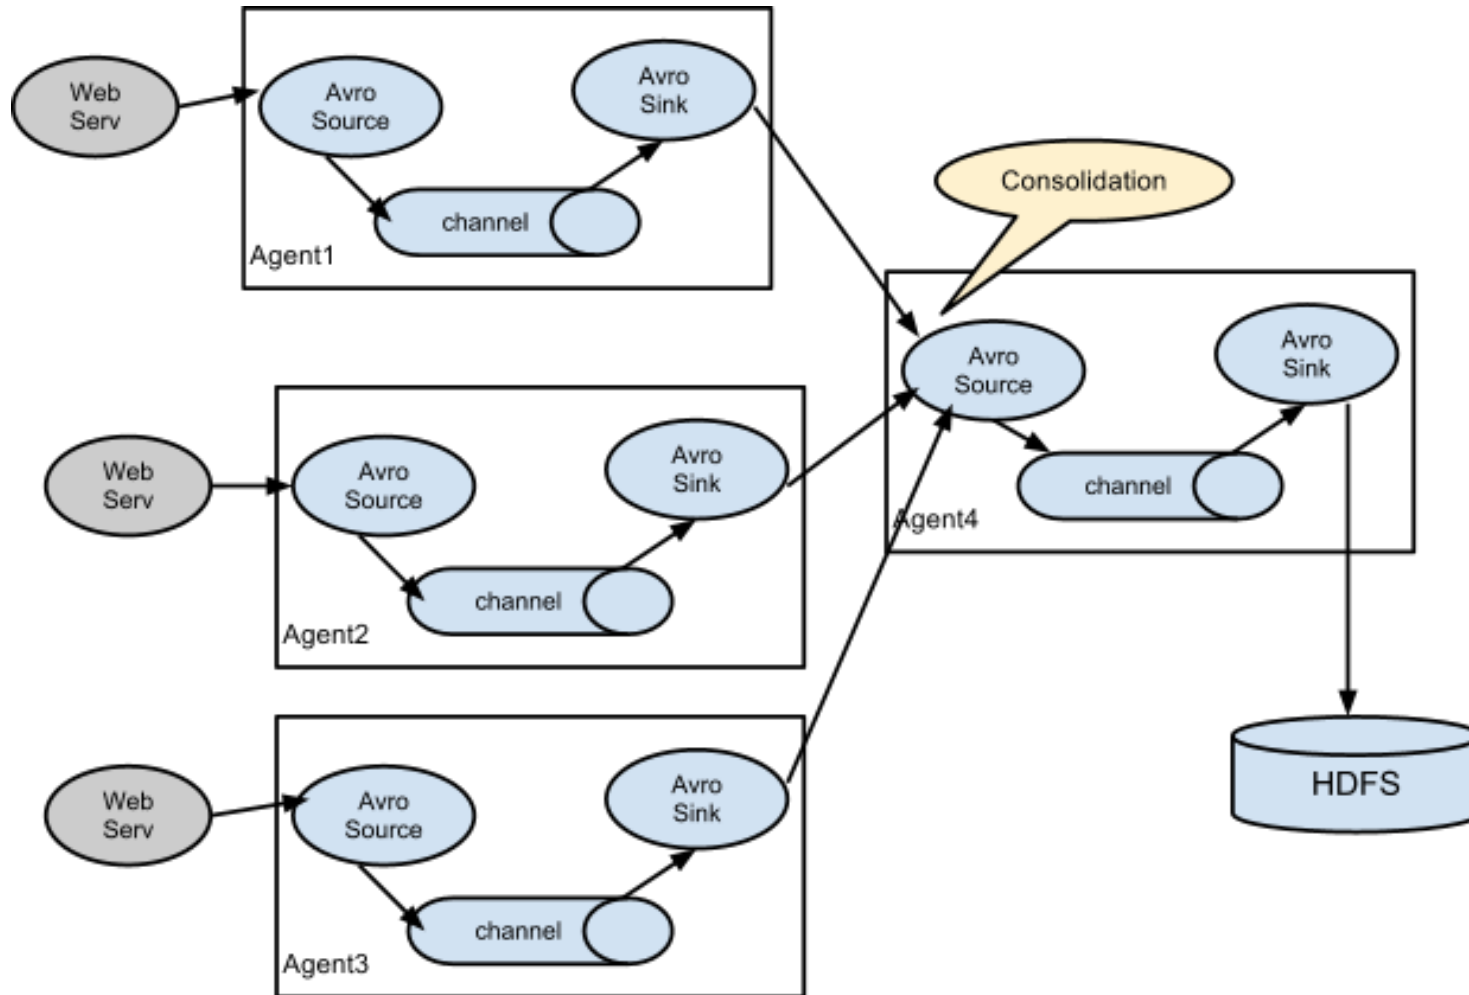
- Distributed, reliable, available service for efficiently collecting, aggregating and moving large amounts of data
- Simple and flexible architecture based on streaming data flows
- Components:
 - *Event* - data being collected
 - *Flume Agent* - source (where the data comes from), channel (repository for the data), sink (next destination for the data)

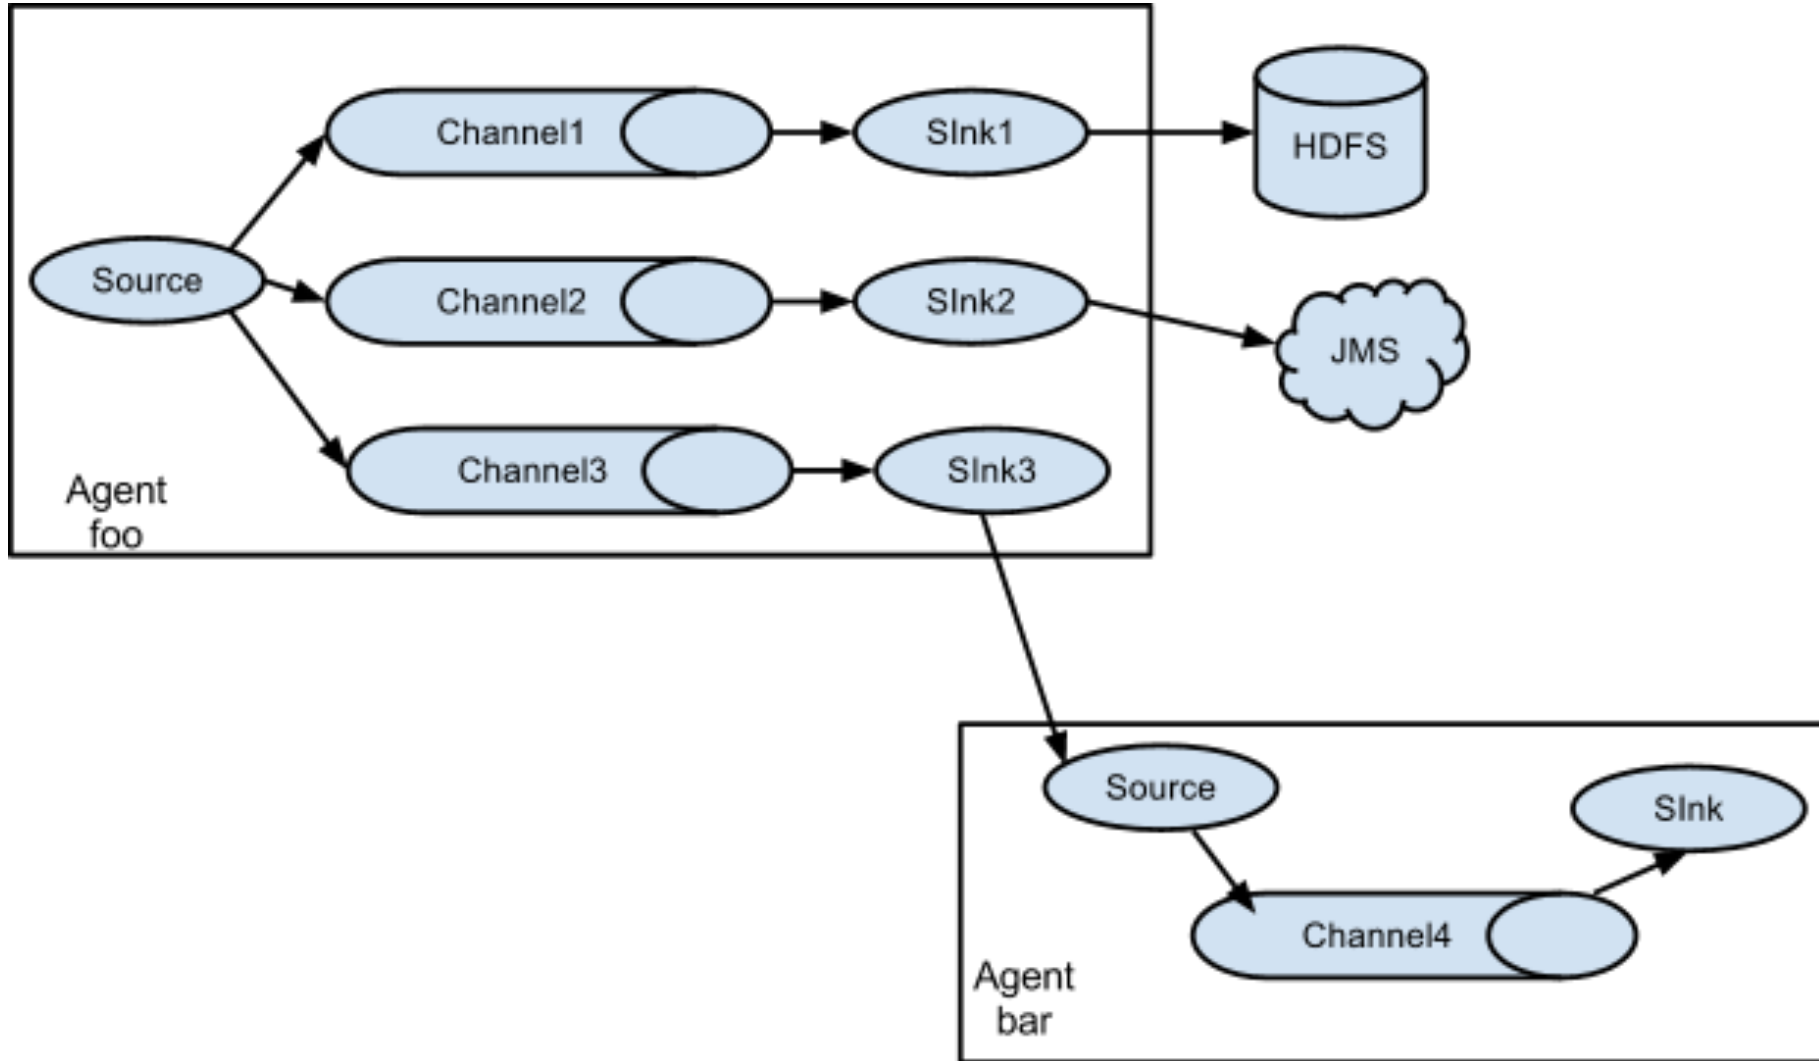






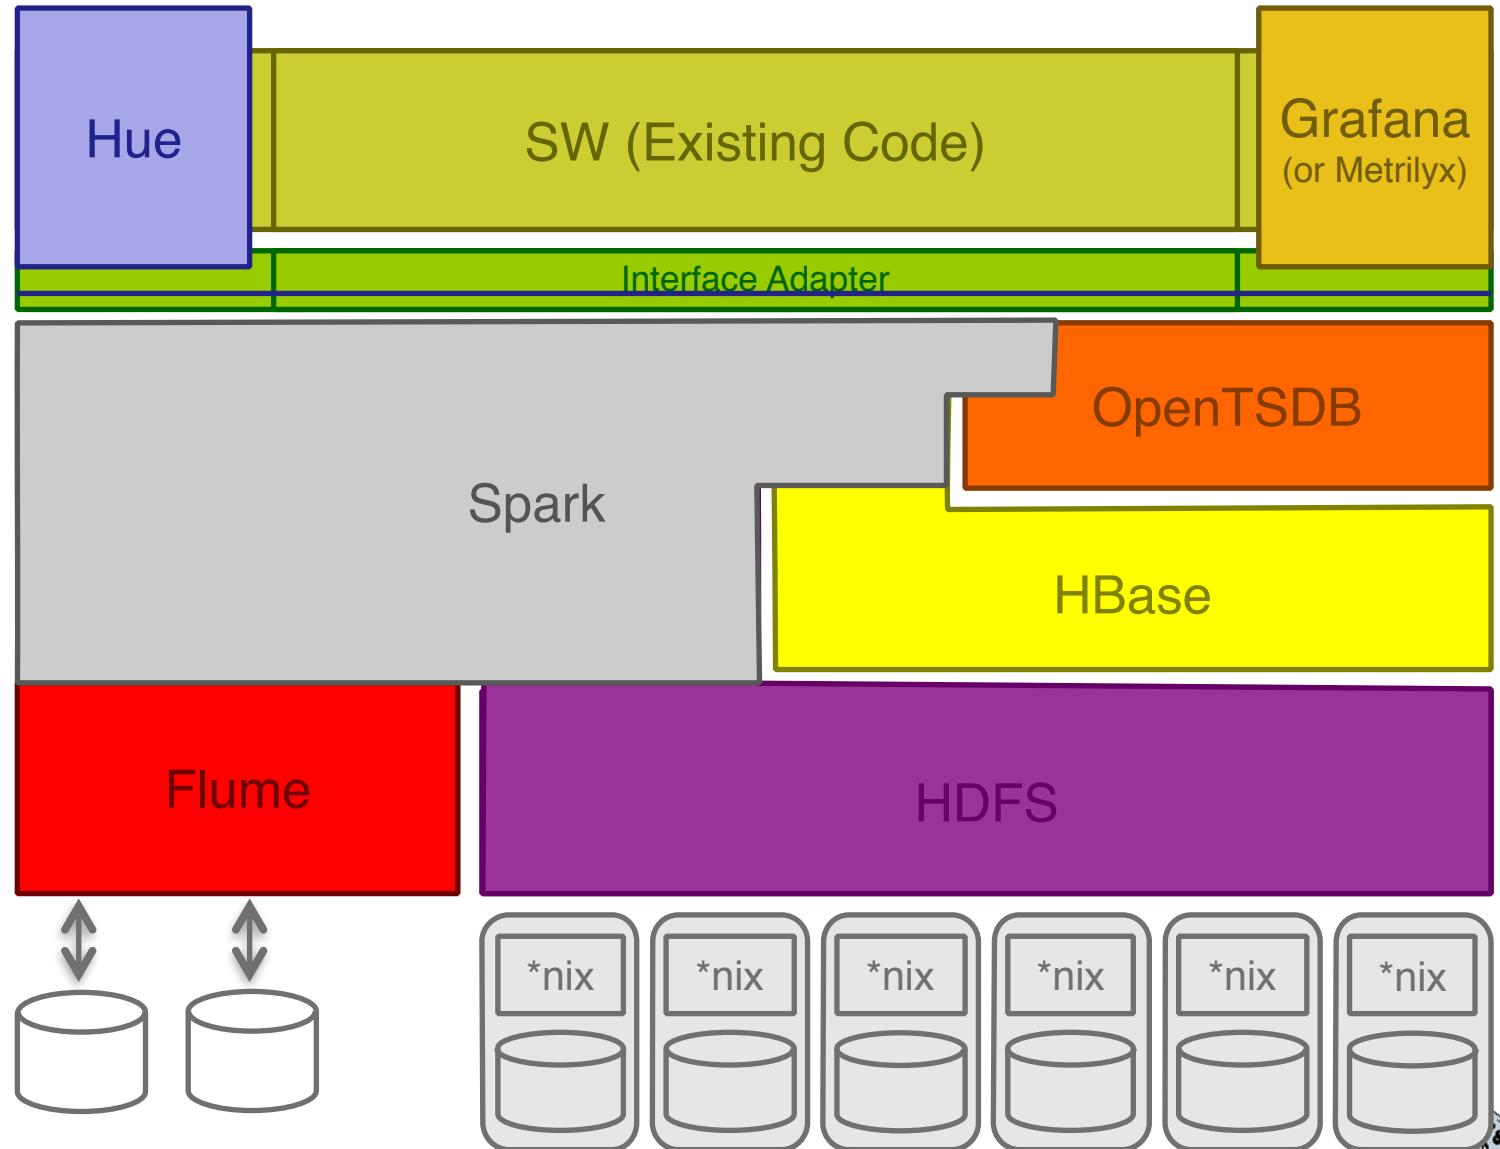






Adding top-level tools:

- Hue
- Grafana





- Hadoop User Experience
- Graphical front-end to developer and administrator functionality
- Developed by Cloudera and released as Open Source
- Extensible (publically-available API)
- Web based
- Integrated with Hadoop tools
 - i.e. HDFS file browser, HBase browser, ZooKeeper browser, Spark editor

HUE Home Query Editors Data Browsers Workflows Search Security File Browser Job Browser cloudera

File Browser

Search for file name Actions Move to trash Upload New

Files Zip/Tgz file Trash

Home / user / cloudera

Name	Size	User	Group	Permissions	Date
↑		hdfs	supergroup	drwxr-xr-x	January 25, 2015 01:09 AM
.		cloudera	cloudera	drwxr-xr-x	February 03, 2015 03:04 AM
.Trash		cloudera	cloudera	drwxr-xr-x	January 28, 2015 04:44 AM
.staging		cloudera	cloudera	drwx-----	February 04, 2015 02:31 AM

Show 45 of 2 items Page 1 of 1

HUE Query Editors Data Browsers Workflows Search File Browser Job Browser remain

Hive Editor Query Editor My Queries Saved Queries History

Navigator Settings

DATABASE default

Table name...

- page_view
- tweets
- business
 - city (string)
 - review_count (int)
 - name (string)
 - neighborhoods (string)
 - type (string)
 - business_id (string)
 - full_address (string)
 - state (string)
 - longitude (float)
 - stars (float)
 - latitude (float)
 - open (boolean)
 - categories (string)
- top_cool4_hbase
- top_reviews
- review
- top_cool
- top_cool_hbase
- timestamp_invalid_data
- test_partitions
- counties
- hankc

Sample: Salary growth Salary_growth (sorted) from 2007-08


```

1 SELECT s07.description, s07.salary, s08.salary,
2       s08.salary - s07.salary
3 FROM
4       sample_07 s07 JOIN sample_08 s08
5 ON ( s07.code = s08.code)
6 WHERE
7       s07.salary < s08.salary
8 ORDER BY s08.salary-s07.salary DESC
9 LIMIT 20
  
```

Execute Save Save as... Explain or create a New query

Recent queries Query Log Columns Results Chart

Chart type [Bar] [Line] [Map] X-Axis description Y-Axis salary



Profession	Salary (approx.)
Dentists, all other specialists	120,000
Surgeons	190,000
Oral and maxillofacial surgeons	180,000
Natural and managers	115,000
Physicians and surgeons	155,000
Orthodontists	185,000
Internists	165,000
Political scientists	90,000
Obstetricians and gynecologists	180,000
Chief executives	150,000
Rotary drill operators, oil and gas	45,000
Pediatricians	145,000
Sociologists	65,000
Family and general practitioners except epidemiologists	155,000
Medical scientists, and sports competitors	75,000
Athletes	70,000
Animal scientists	55,000
Dentists, general	145,000
Education administrators	85,000
Psychologists	80,000
all other postsecondary	85,000



HUE Query Editors | Data Browsers | Workflows | Search | File Browser | Job Browser | remain

Search: Apache Logs | impala | hive

Italiano (328) | Brazil (289) | Show more...

extension

user_agent_family

bytes

- 0 - 900000 (8083)
- 900000 - 1800000 (0)
- 1800000 - 2700000 (0)
- 2700000 - 3600000 (0)
- 3600000 - 4500000 (0)
- 4500000 - 5400000 (0)

country_code3

region_code

Grid Results

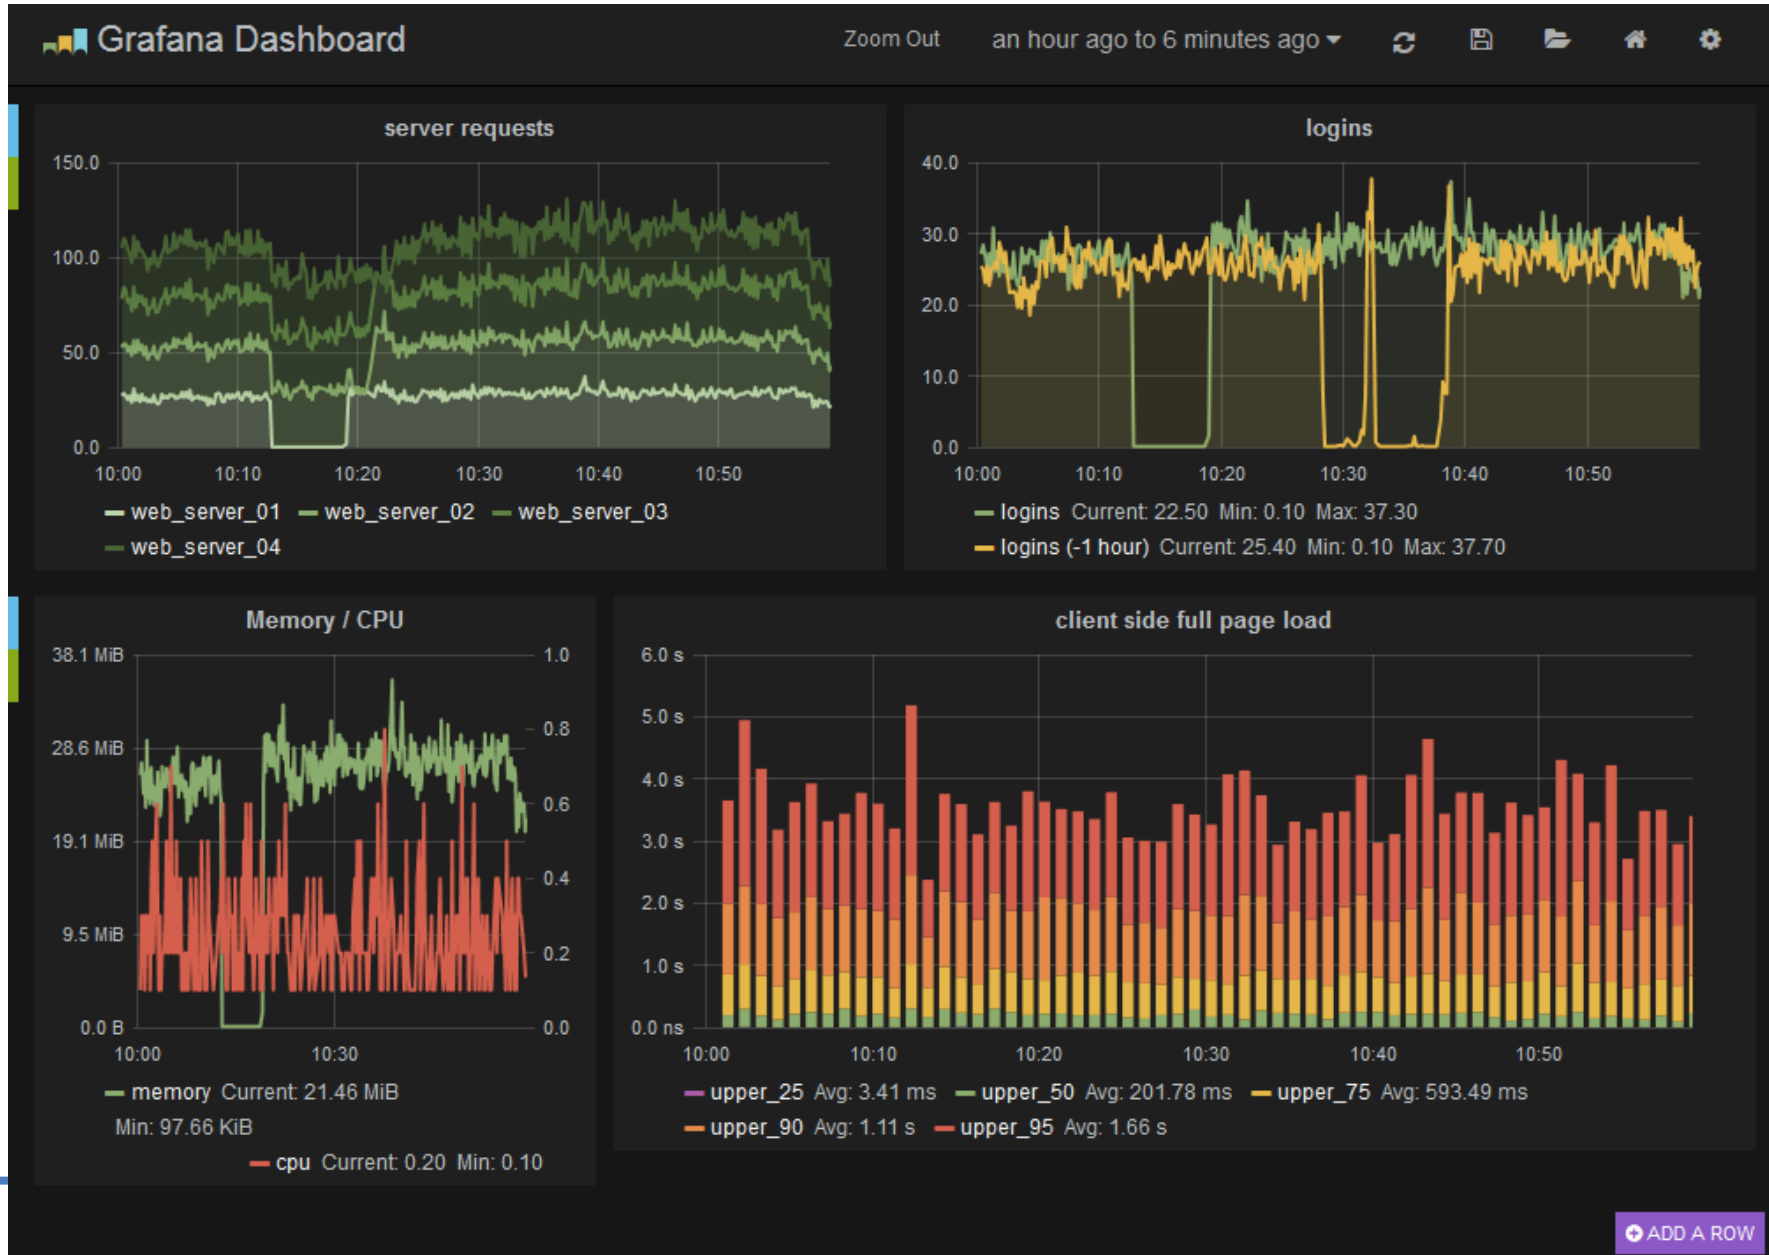
Showing 1 to 40 of 8083 results

code	app	request	bytes
▶ 302	impala	GET /impala/list_designs?q-page=1&q-type=impala HTTP/1.0	446
▶ 200	impala	GET /impala/list_designs?q-page=1&q-type=impala HTTP/1.0	10839
▶ 302	impala	GET /impala/query_history?q-type=impala&q-user=rw1hd7z&q-auto_query=off HTTP/1.1	549





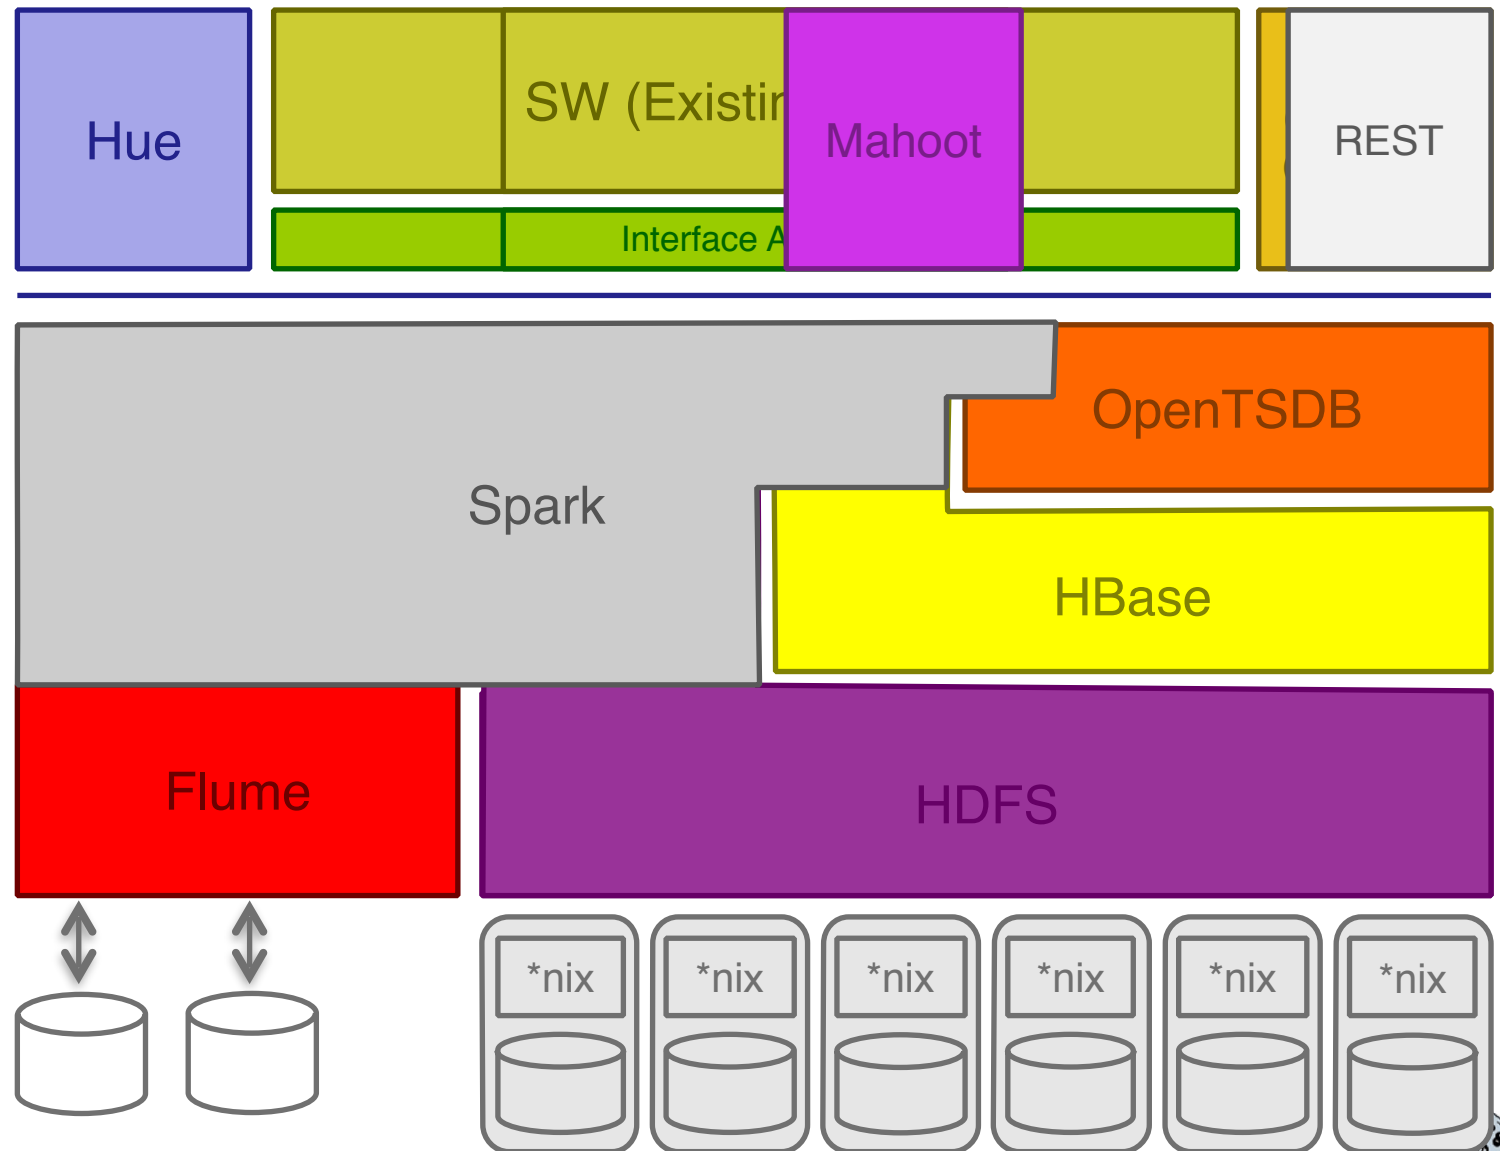
- Open source
- Feature rich metrics dashboards and graph editor for OpenTSDB
 - Rich graphing: fast and flexible client side graphs with a multitude of options
 - Mixed styling
 - Dashboards: drag and drop panels, change row and panel width easily
 - Annotations



Step 5

Empowering the architecture:

- Machine Learning
- Third-party Access
- ... and more



Distributed and scalable machine learning algorithms on the Hadoop platform:

- Collaborative filtering
- Clustering
- Classification
- Dimensionality Reduction



Use cases:

- Yahoo → spam detection
- Foursquare → recommendations
- Adobe → user targeting
- Amazon → personalization platform

Classification:

- Logistic Regression - trained via SGD
- Naive Bayes / Complementary Naive Bayes
- Random Forest
- Hidden Markov Models
- Multilayer Perceptron

Clustering:


- k-Means Clustering
- Fuzzy k-Means
- Streaming k-Means
- Spectral Clustering

Dimensionality Reduction:

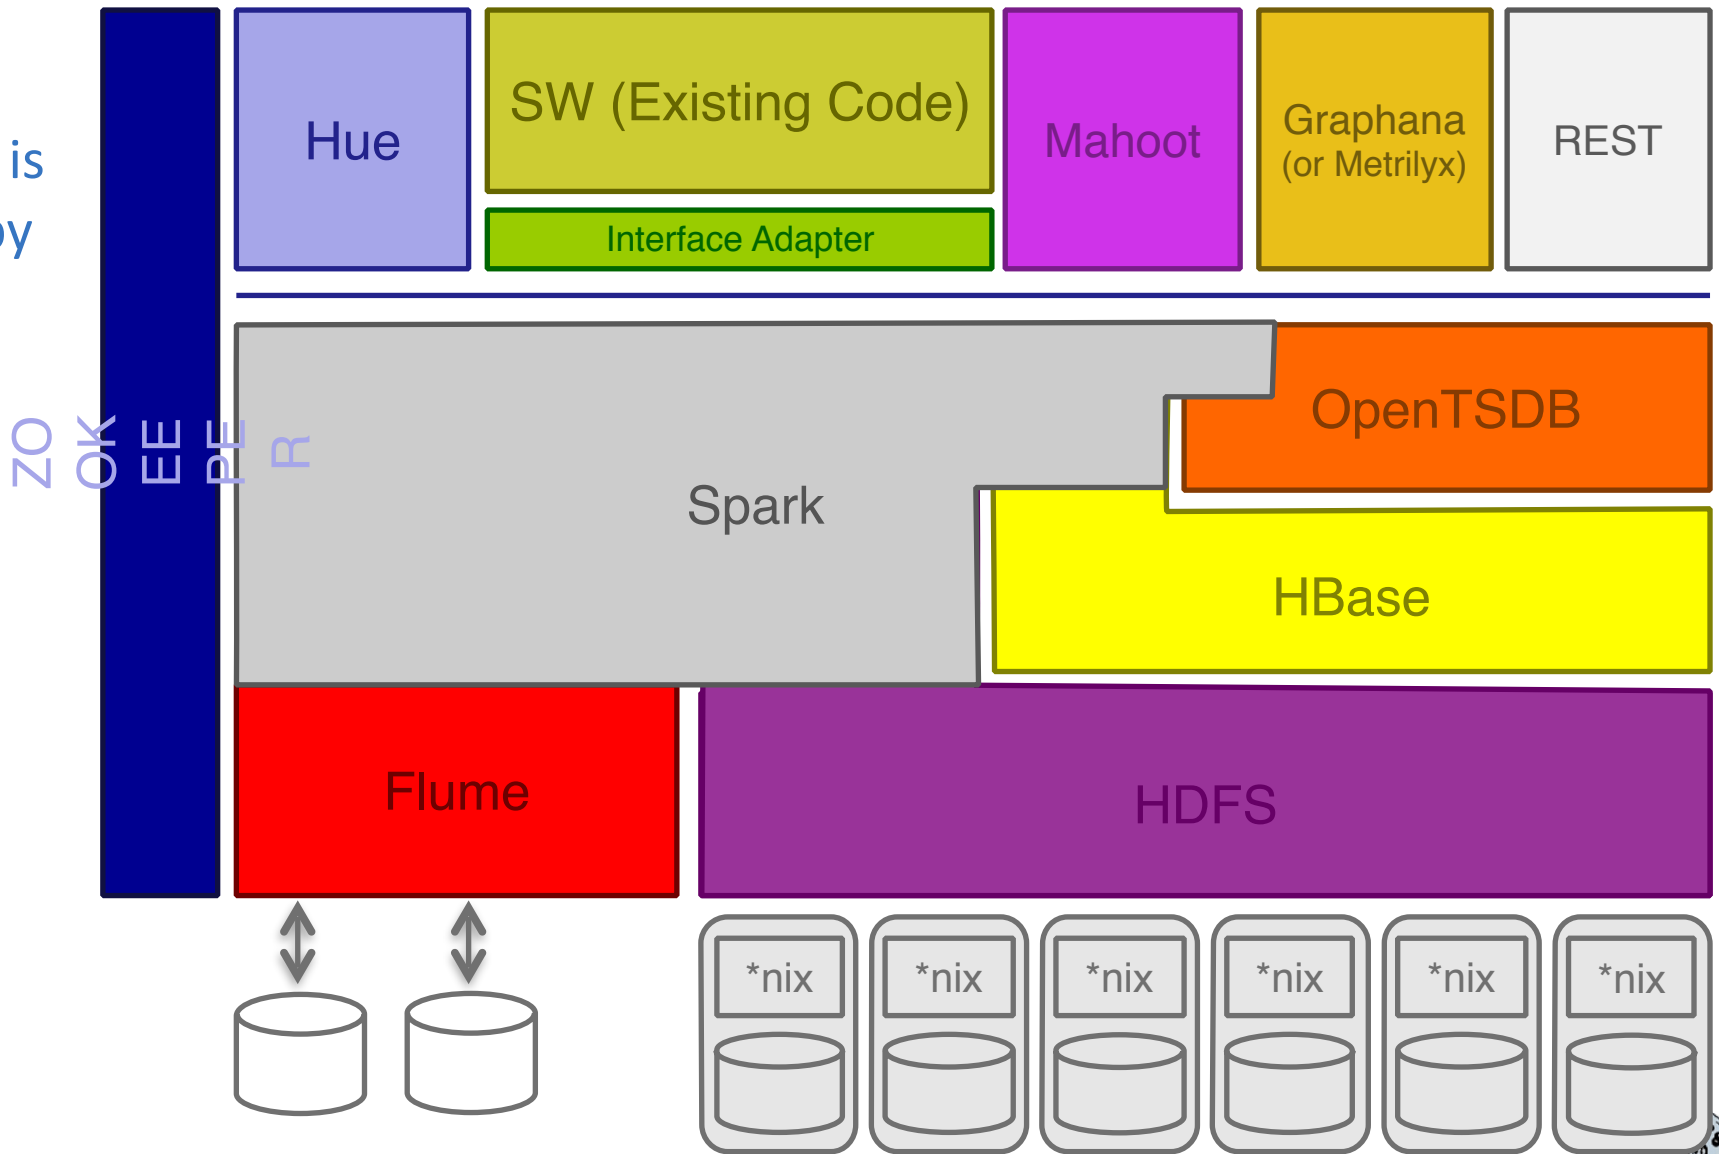
- Singular Value Decomposition
- Lanczos Algorithm
- Stochastic SVD
- PCA (via Stochastic SVD)
- QR Decomposition

Make easier the integration with third-parties, including:

- External software
- Scientific Partners
- Data Analysis Tools
- ...

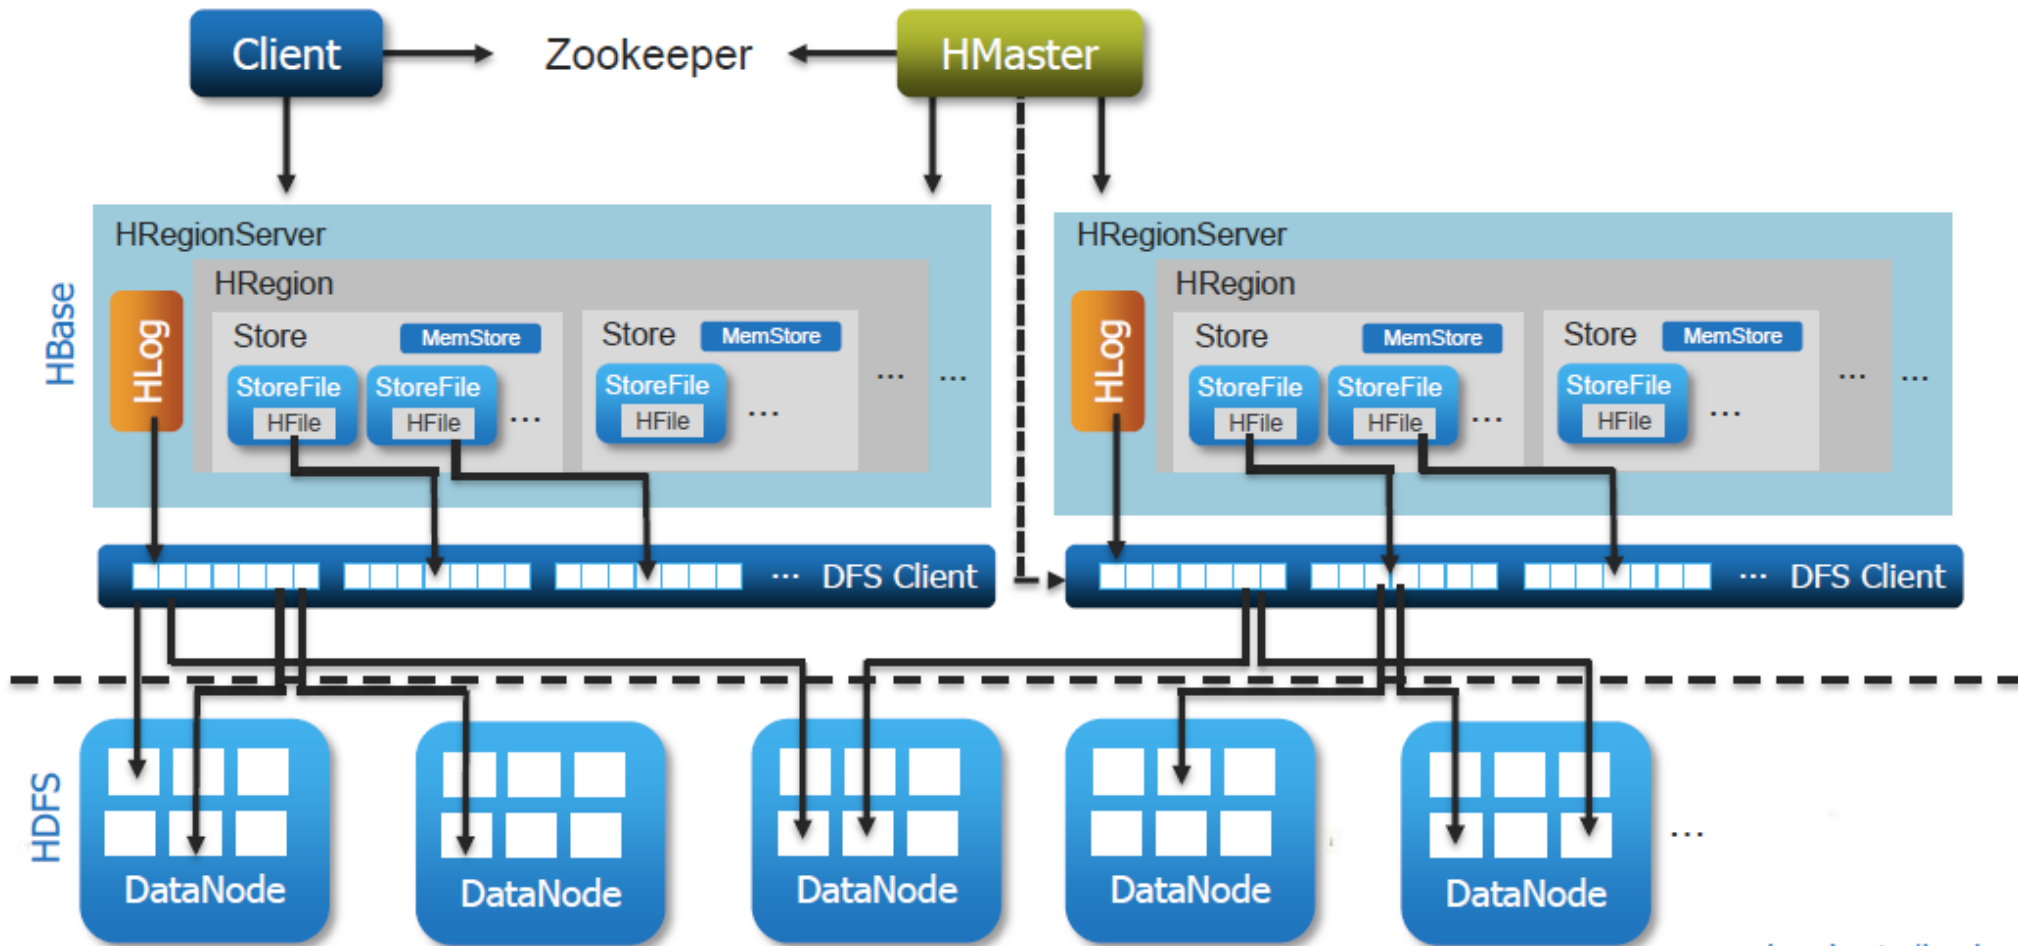


The overall architecture is maintained by ZOOKEEPER



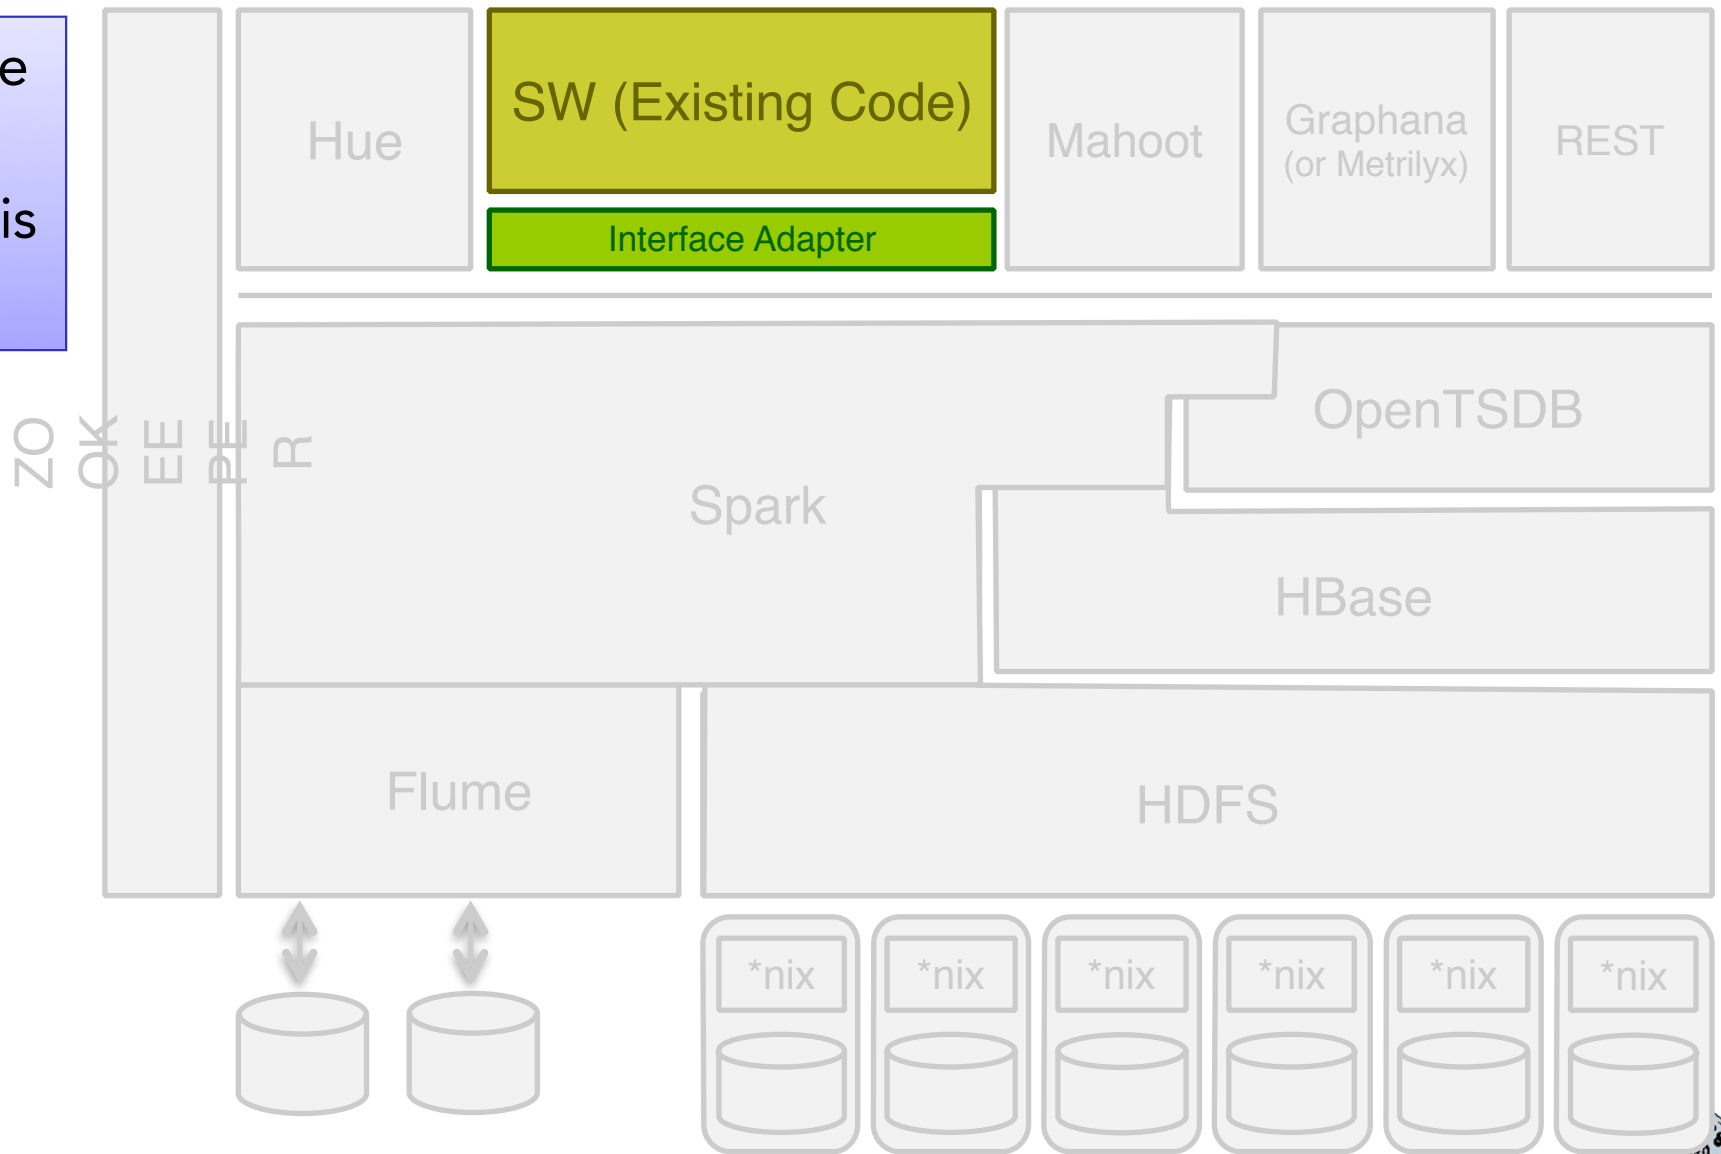
- *“ZooKeeper allows distributed processes to coordinate with each other through a shared hierarchical name space of data register” (ZooKeeper Wiki)*
- An open source, high-performance coordination service for distributed applications
- Service for maintaining configuration information, naming, providing distributed synchronization and providing group services



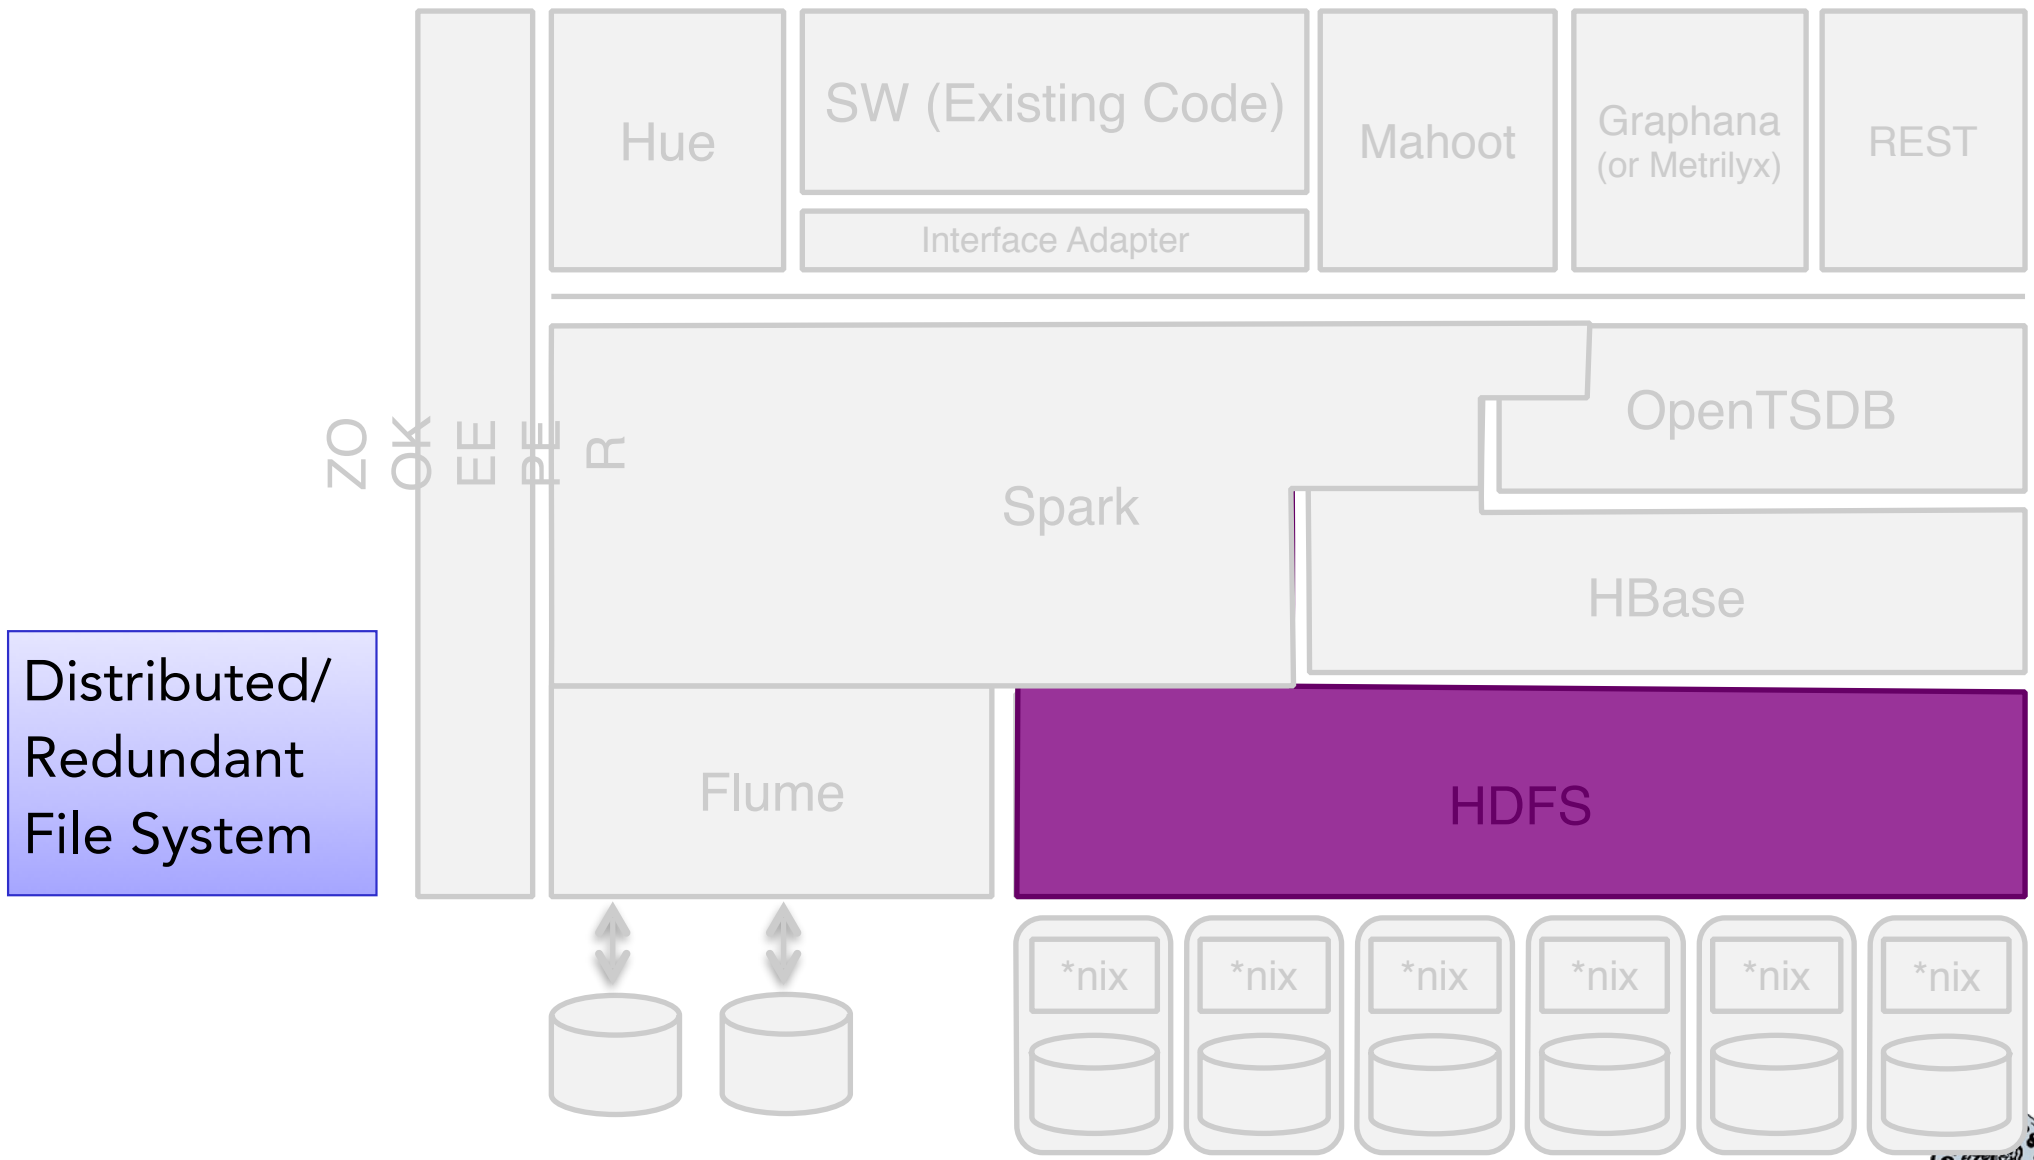


www.edureka.in/hadoop

This is where the domain knowledge is preserved.



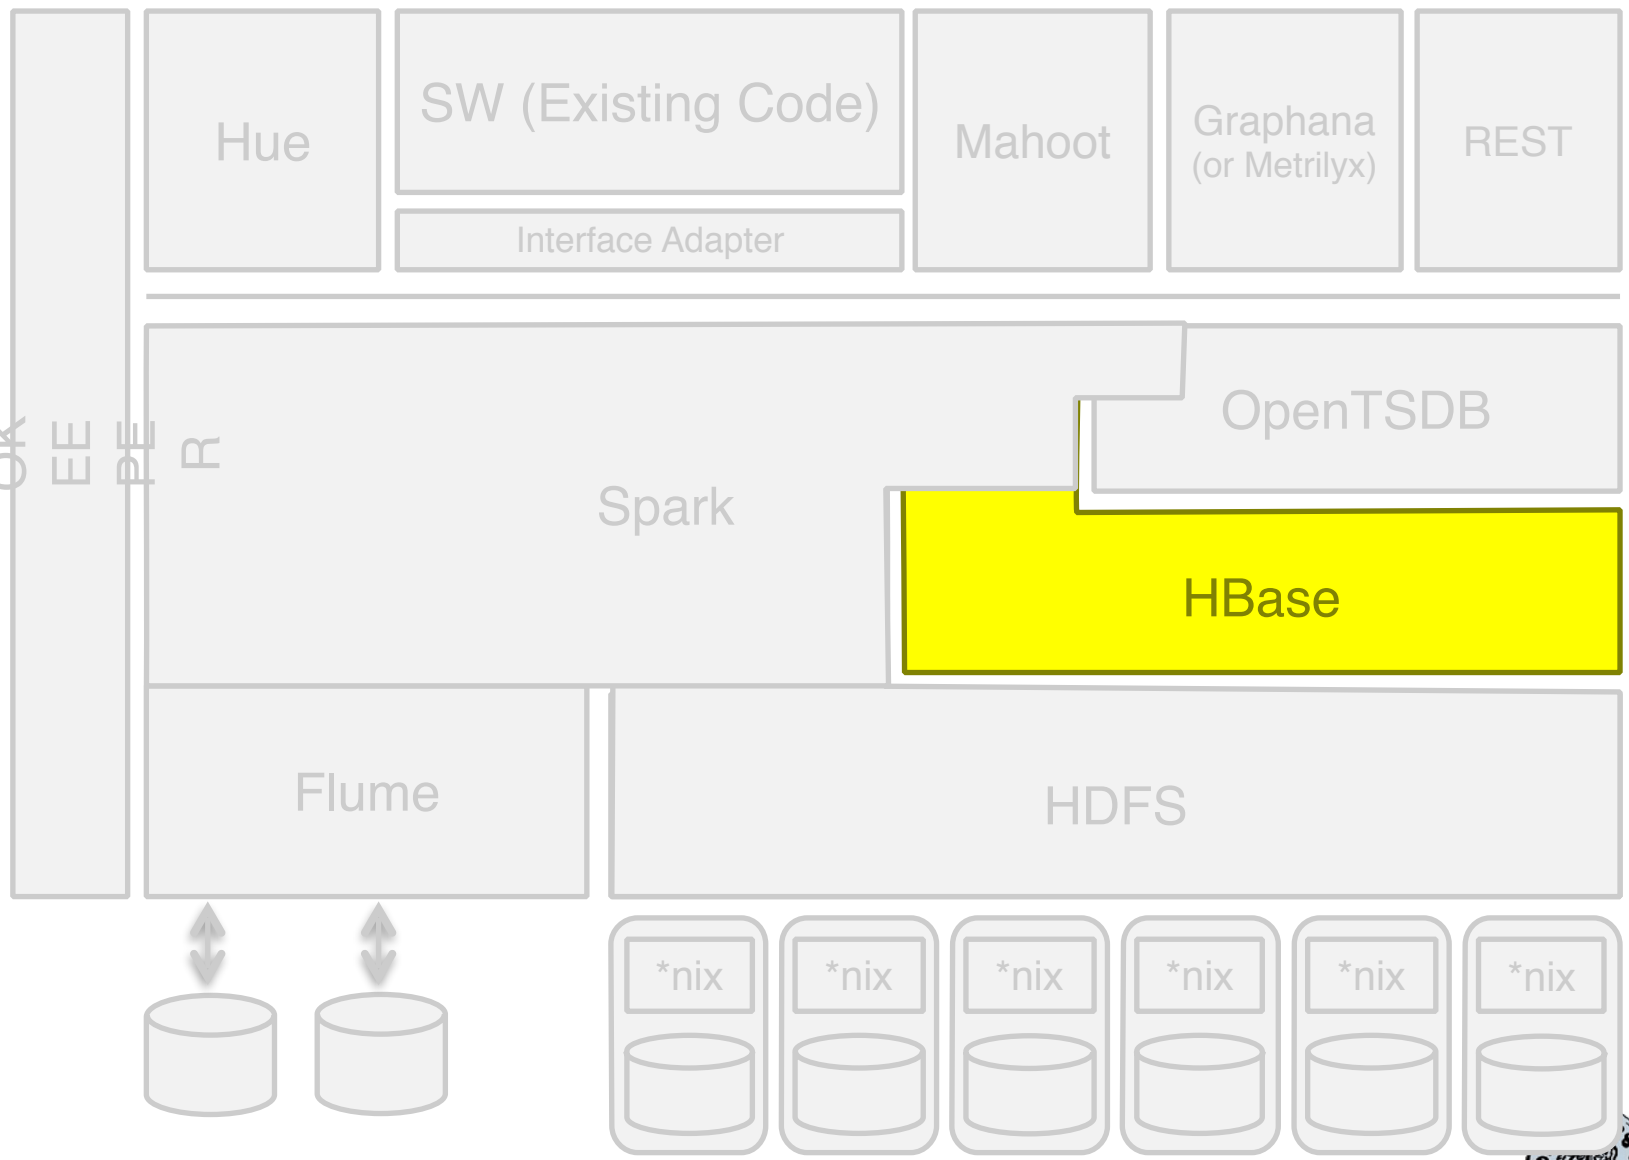
Overview: HDFS



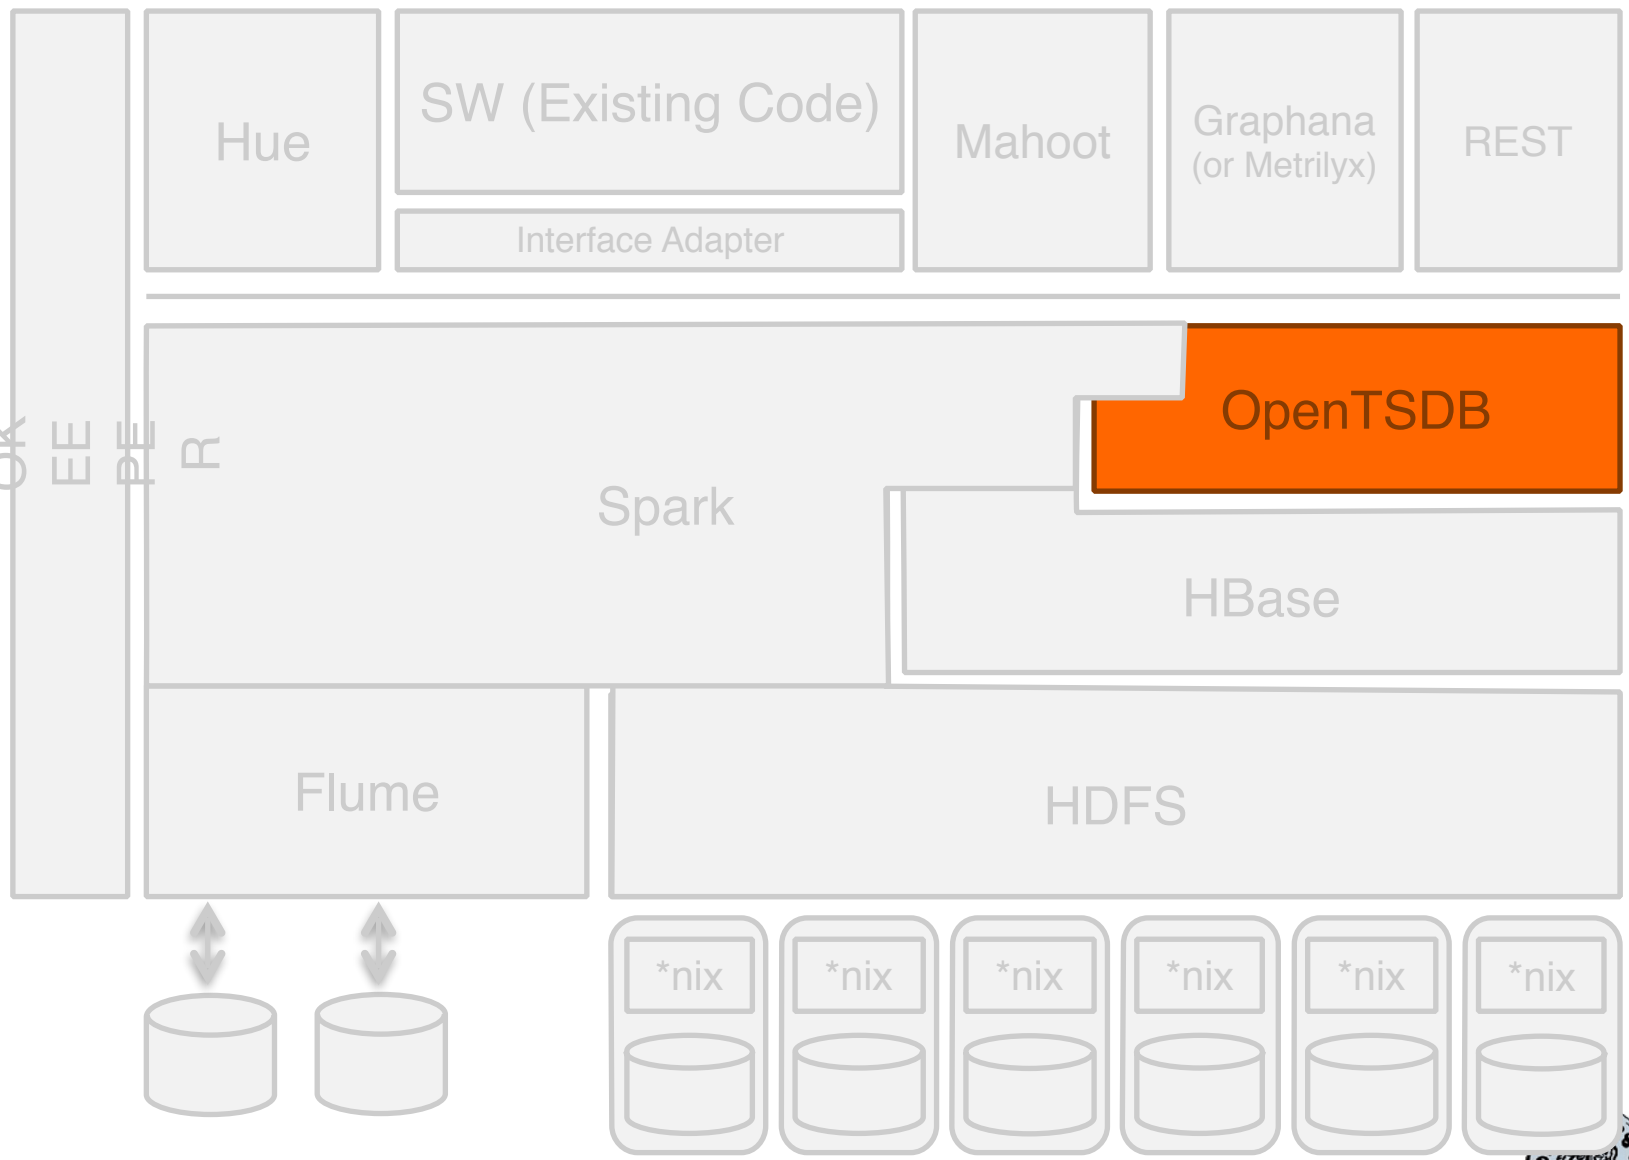
Overview: HBase

Vertical NoSQL Database

NO
OK
EE
DE
R

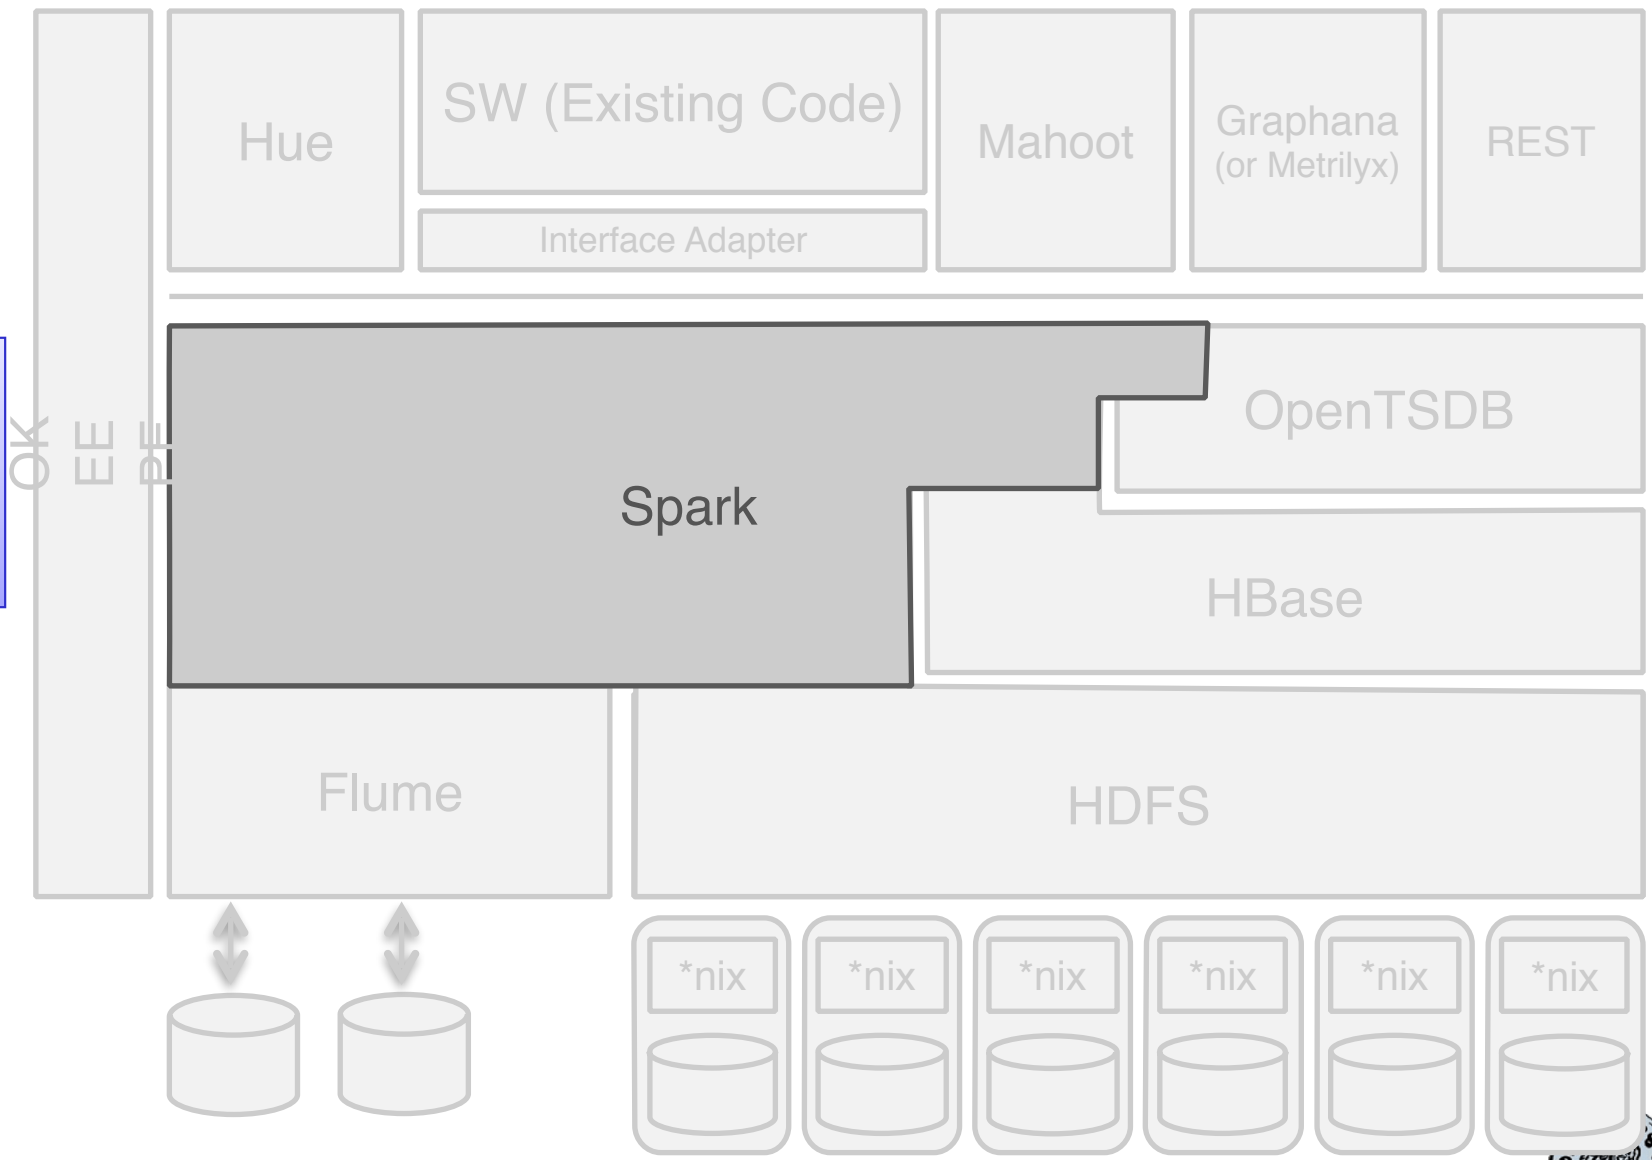


Time Series DBMS




Overview: Spark

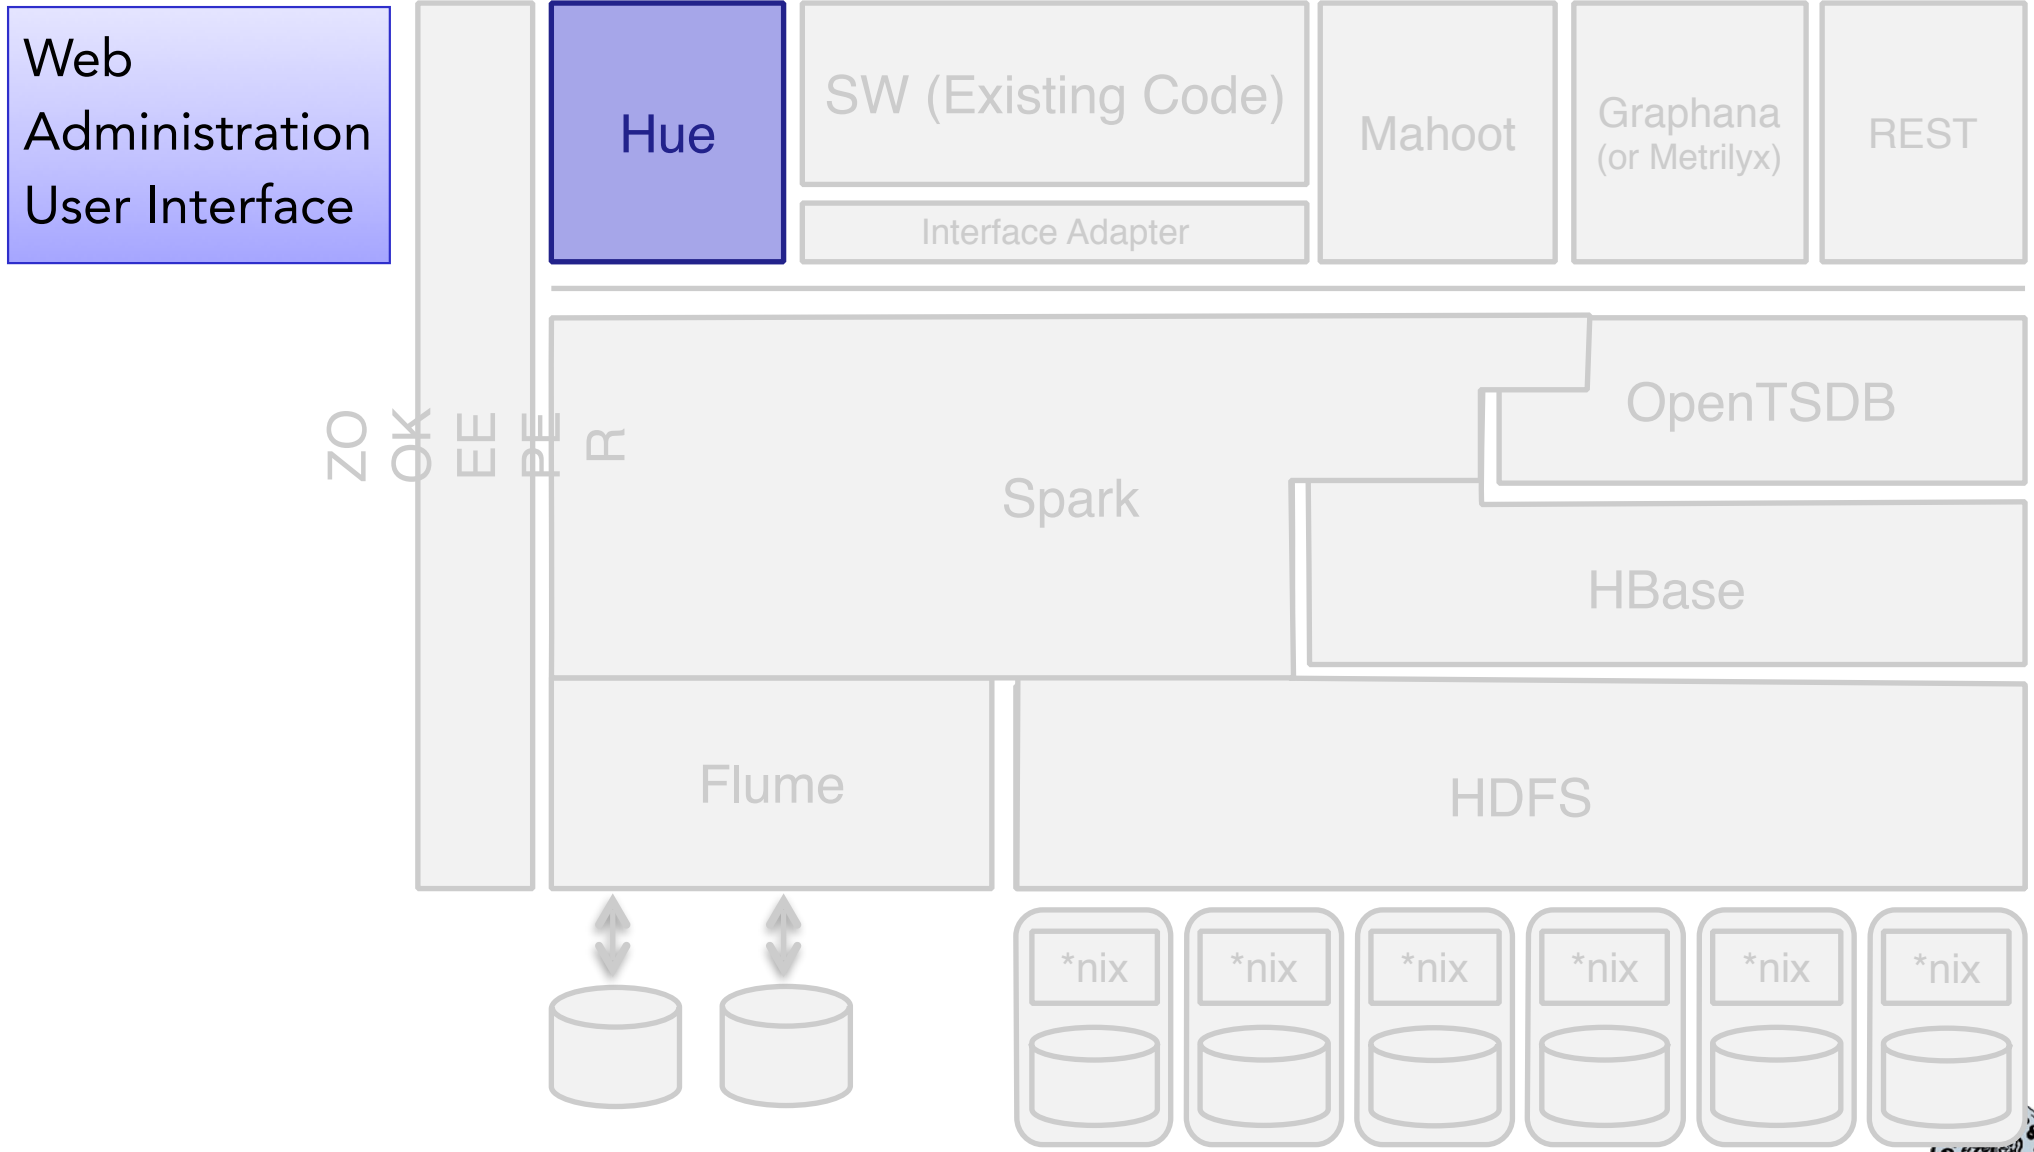
Enhanced Map-Reduce



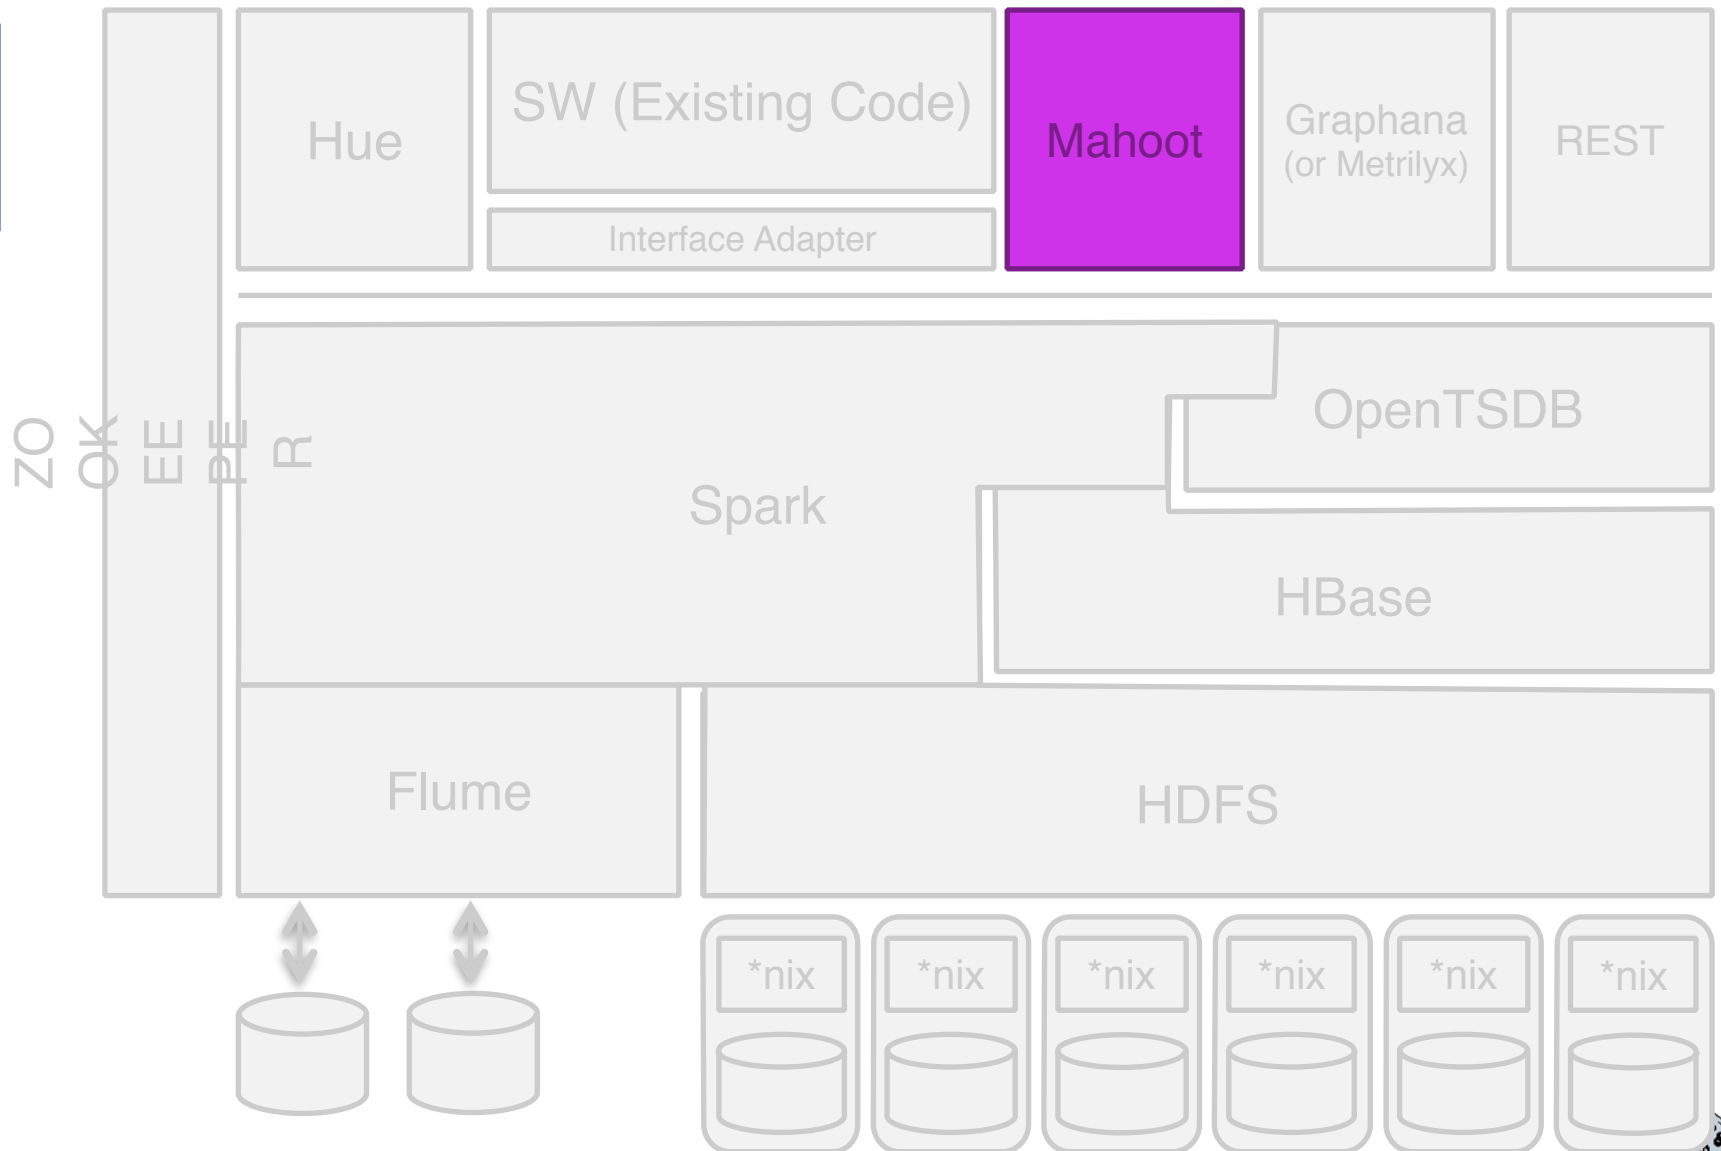
Overview: Flume



Overview: Hue

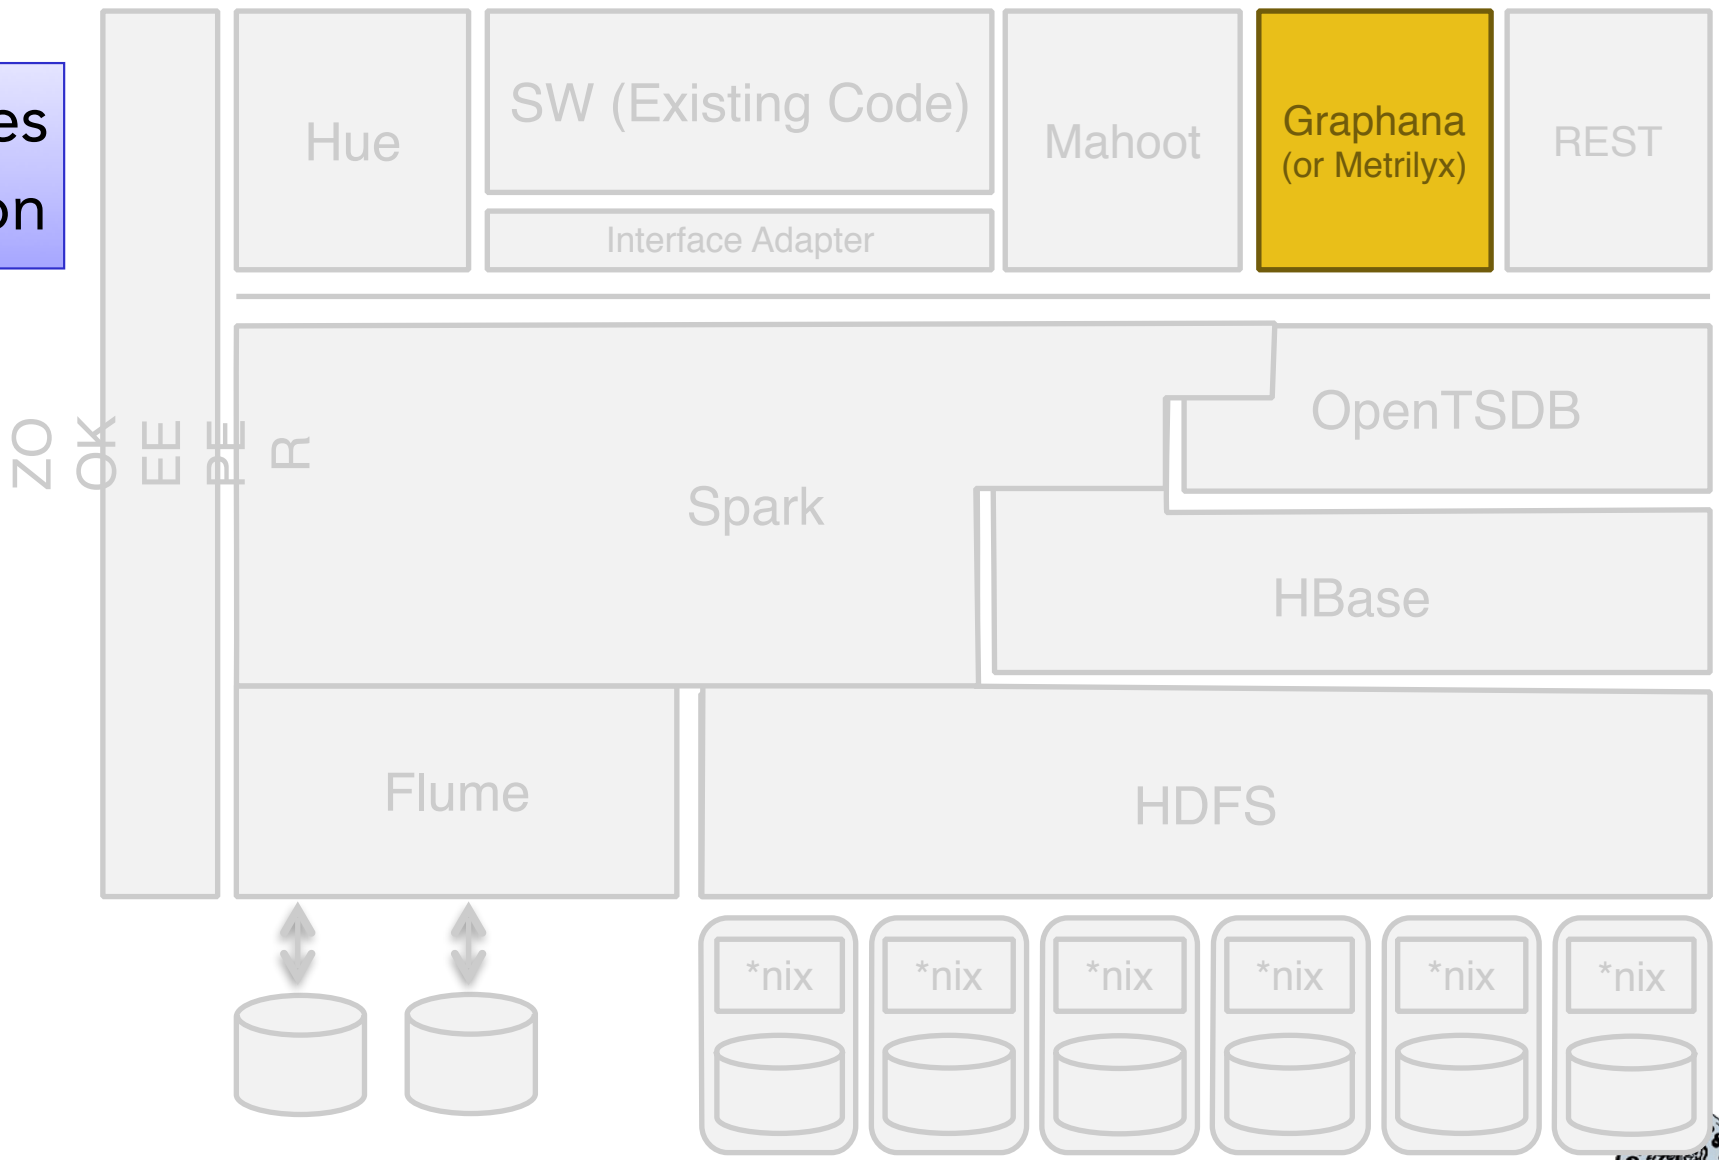


Machine Learning

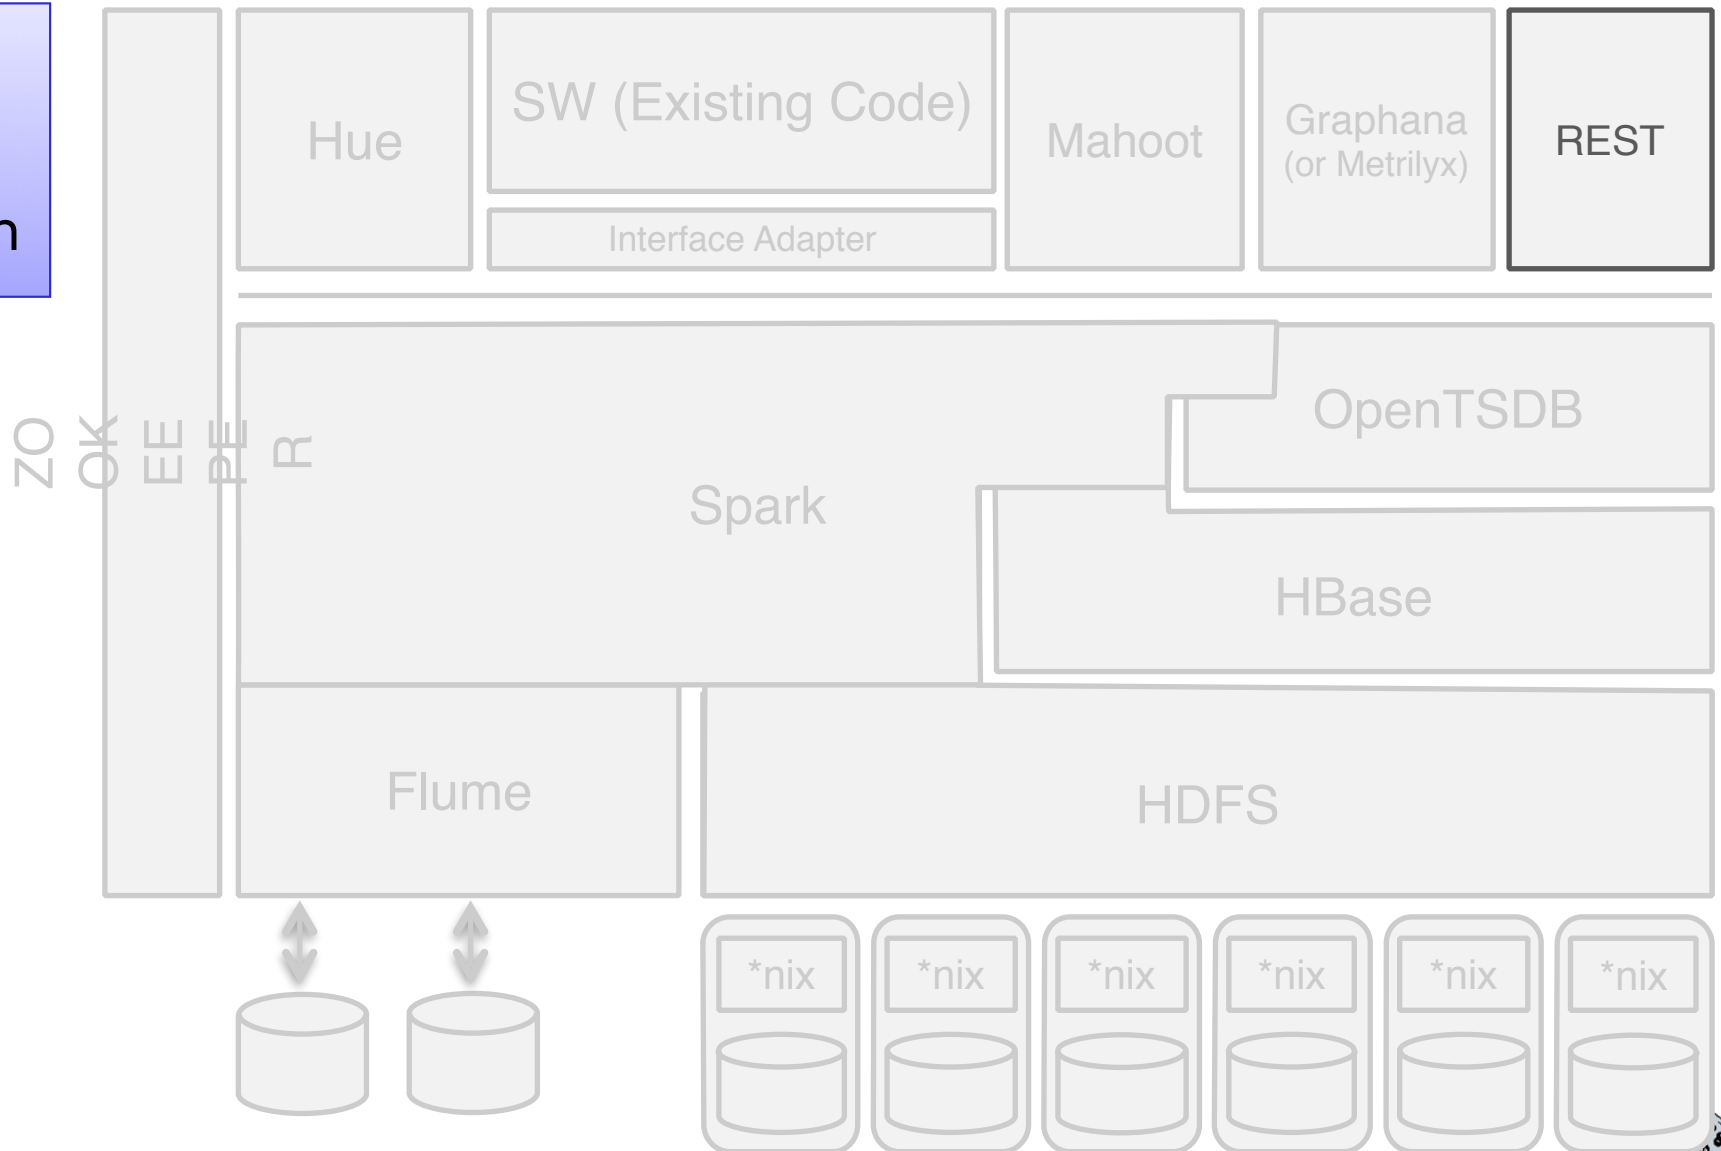


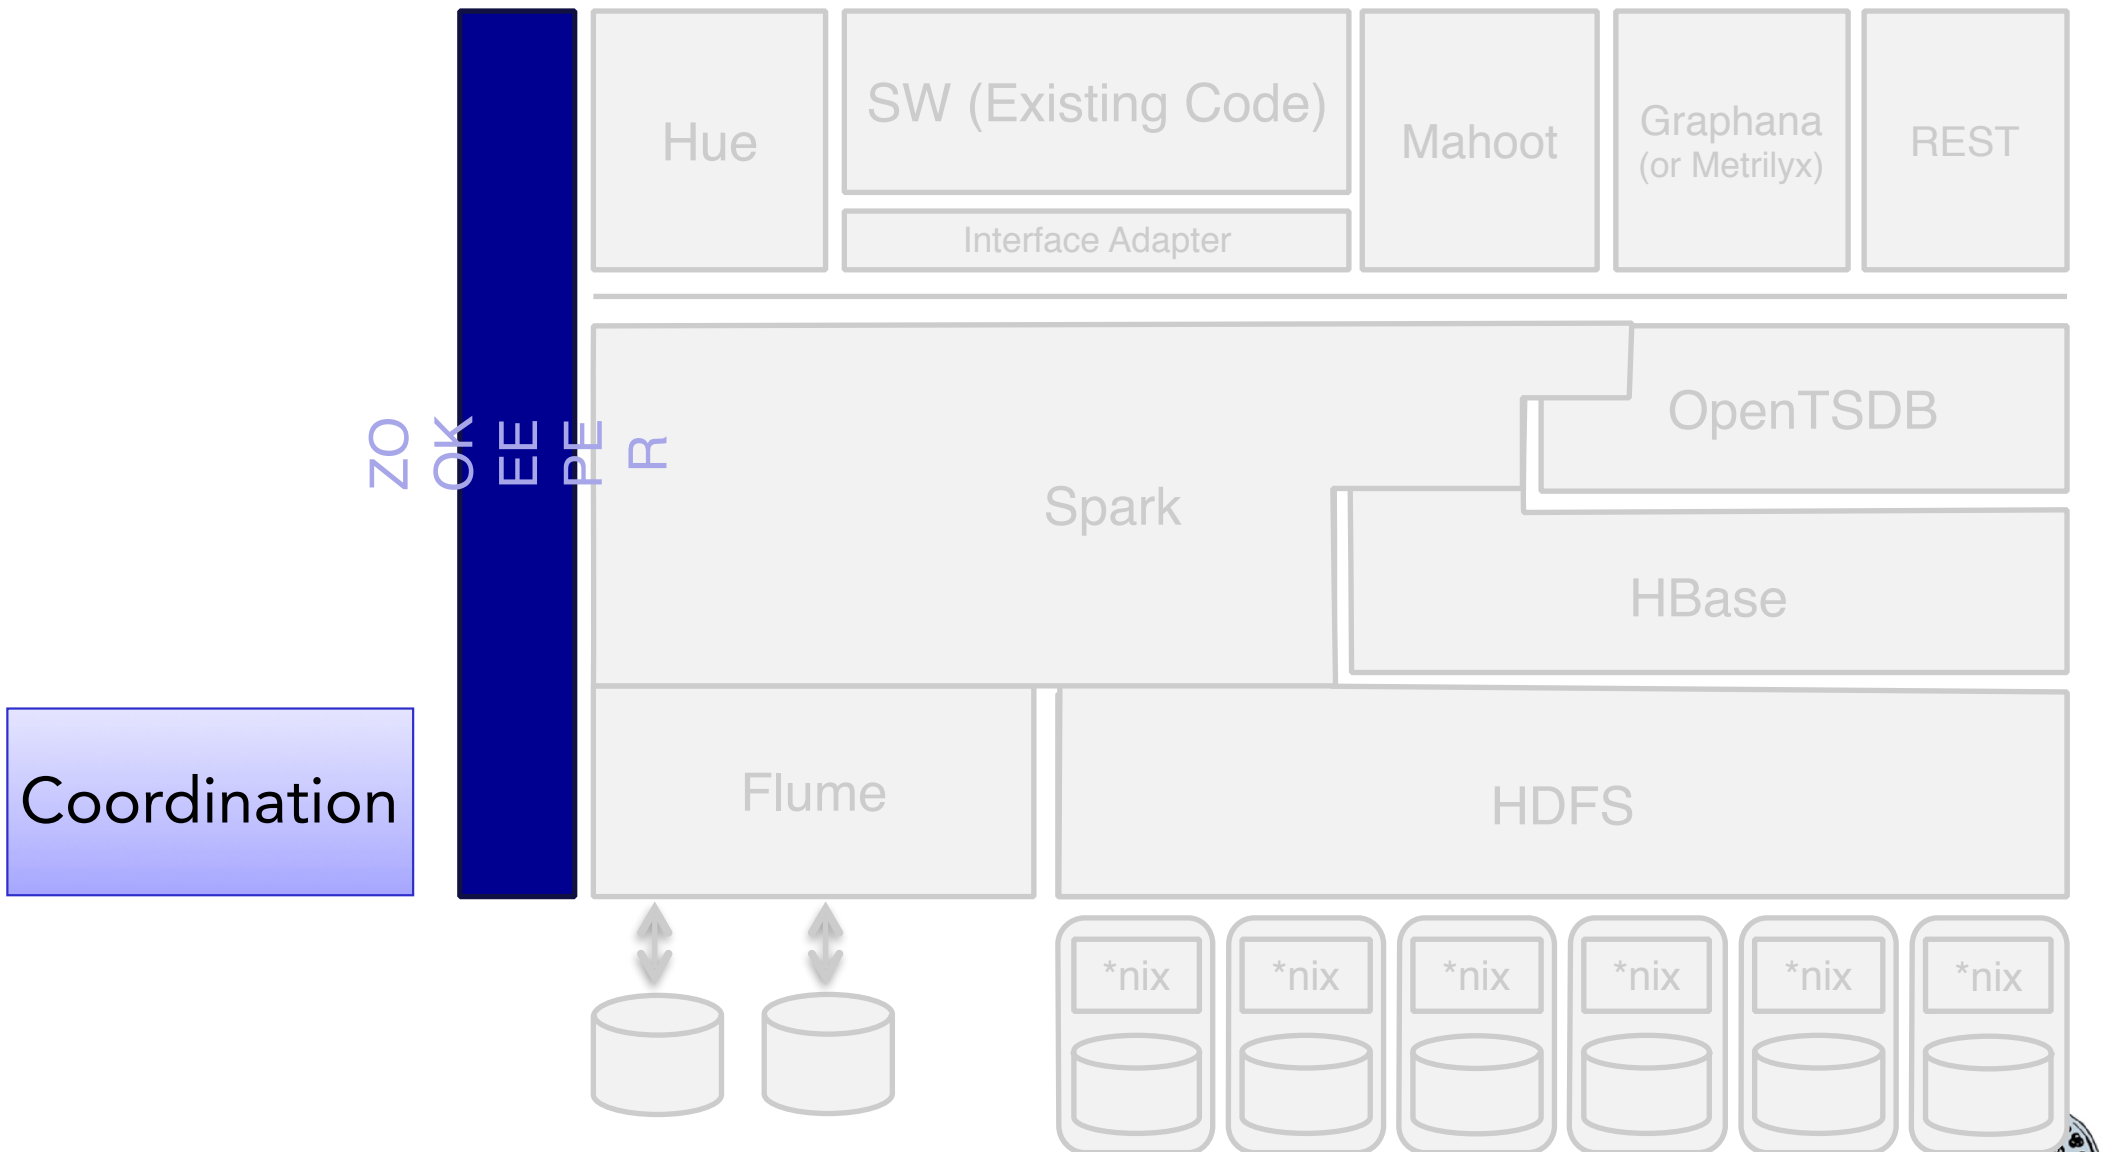
Overview: Graphana

Time Series Visualizaion



Third-Parties Integration







- Big Data
- Hadoop & Map-Reduce
- HDFS vs. NFS
- Evolving towards a Big Data architecture
- **FPGA co-processing**
- Conclusions

- We will need more and more computational power.


- More Moore:
 - + integration scale
 - + clock frequency



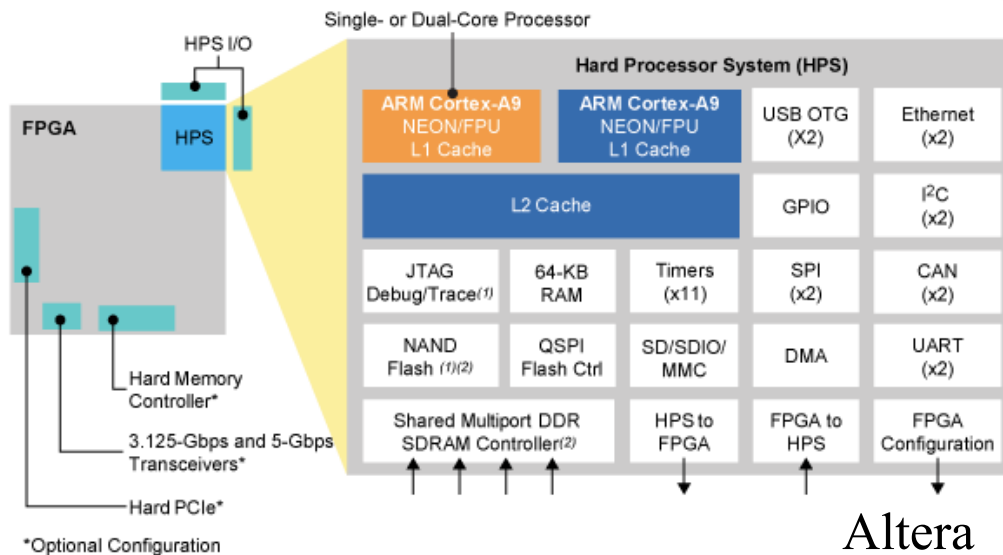
- More than Moore:
 - System on Chip (SoC)
 - FPGA




- Systems-on-Chips
 - Faster
 - Less power
- FPGA
 - Programmable Hardware



- Silicon Convergence
- Reconfigurable Computing
 - Make the device you need, when you need.



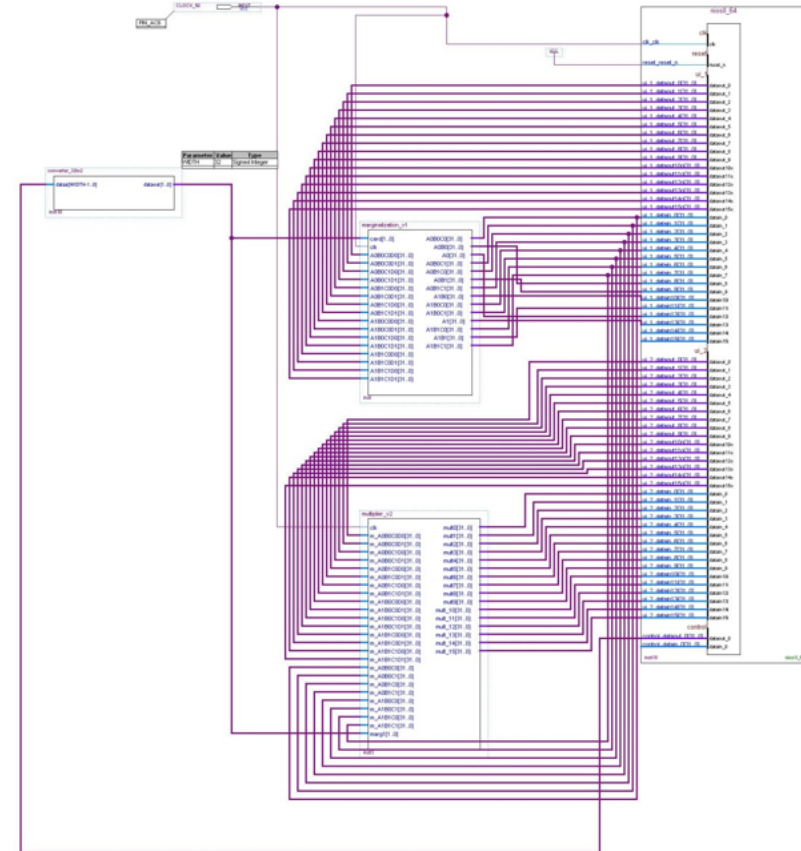


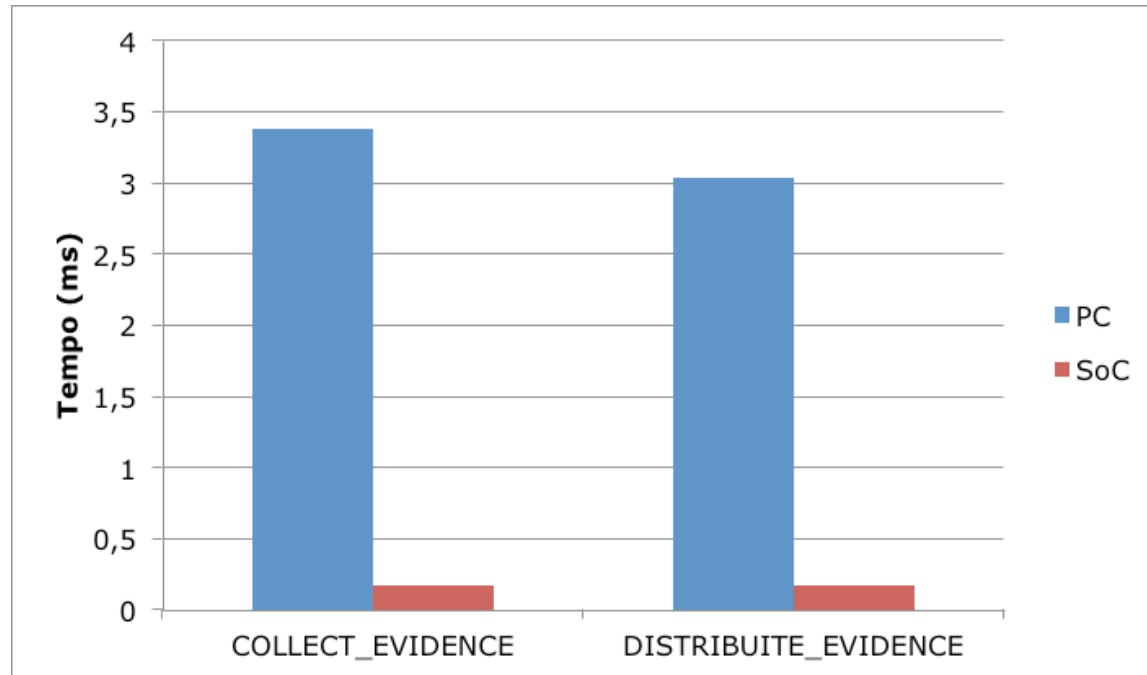
Programmable Logic Devices
Programmable "Logic"

All Programmable Devices
First Generation - 28nm

All Programmable Devices
Second Generation - 20nm

- A Bayesian coprocessor based on SoC in FPGA
- Based on Altera Stratix IV chipset and Nios Architecture
 - Memory on chip
 - Bayesian device (Memory Mapped)
- Two levels
 - Evidence propagation is controlled via software by Nios processor
 - Clique computation is perfumed by hardware

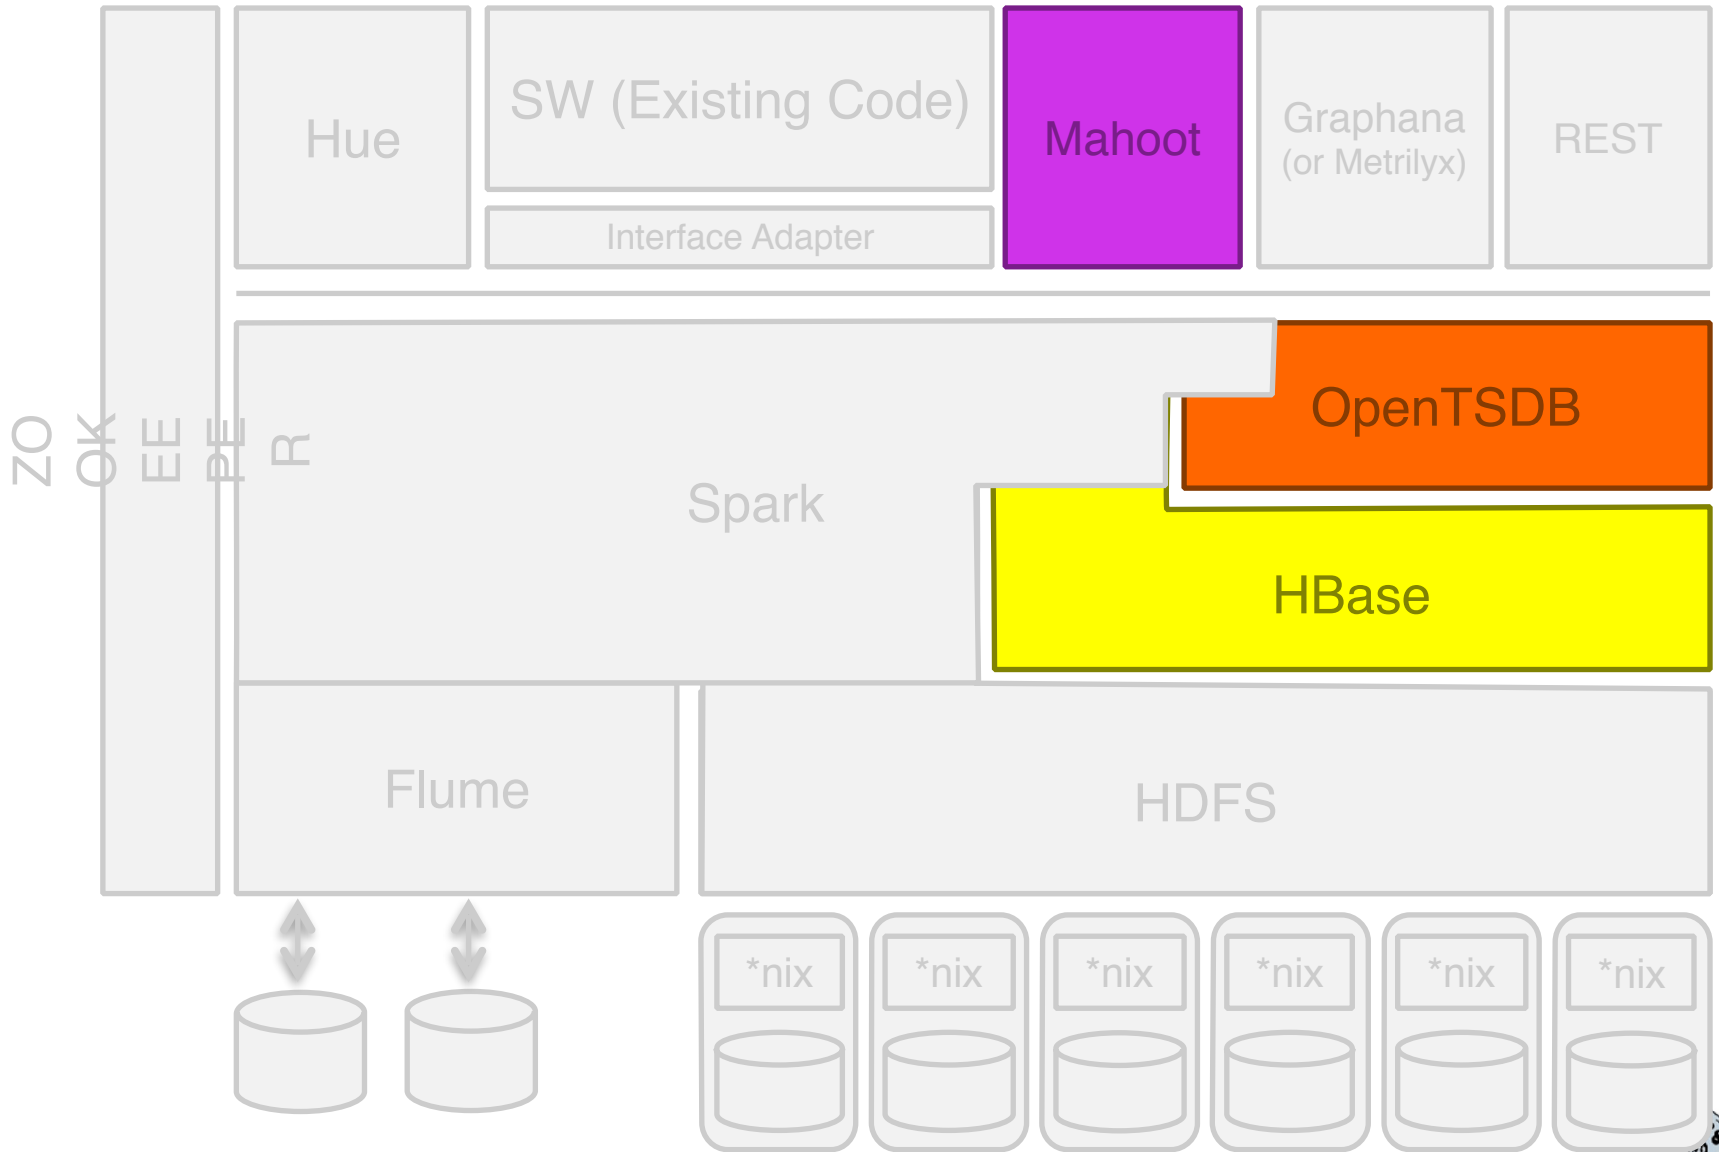




	PC	SoC
COLLECT_EVIDENCE	3.3765 ms	0.1731 ms
DISTRIBUITE_EVIDENCE	3.0313 ms	0.1742 ms
FULL STEP	6.4078 ms	0.3473 ms

HW Native Support for:

- Analysis
- Time Series Retrieval
- Data Base Queries



- We eager to establish a collaboration on this topic
- Expected Benefits:
 - To build a cutting edge solution
 - To train KAGRA people on Big Data IT skills
 - To face a challenging project in Big Data storing & retrieval

- Big Data
- Hadoop & Map-Reduce
- HDFS vs. NFS
- Evolving towards a Big Data architecture
- FPGA co-processing
- **Conclusions**

- Modern technologies for Big Data handling might empower iKAGRA platform
- Evolution towards an advanced architecture can be seamless, in order to:
 - Make the most value of current software
 - Gain experience and training
 - Tailor the solution to emerging needs
 - Mitigate the risk of failure
- Research ideas: can we boost the architecture by FPGA co-processing

Enabling Big Data Architectures for the KAGRA Project

Luigi Troiano and Maria C. Vitelli

University of Sannio

{troiano,vitelli}@unisannio.it



ありがとう
ございます。