SAP S/4 HANA - Simple Logistics Training

By Praveen

www.praveensd.com

Introduction To HANA - Deep Dive

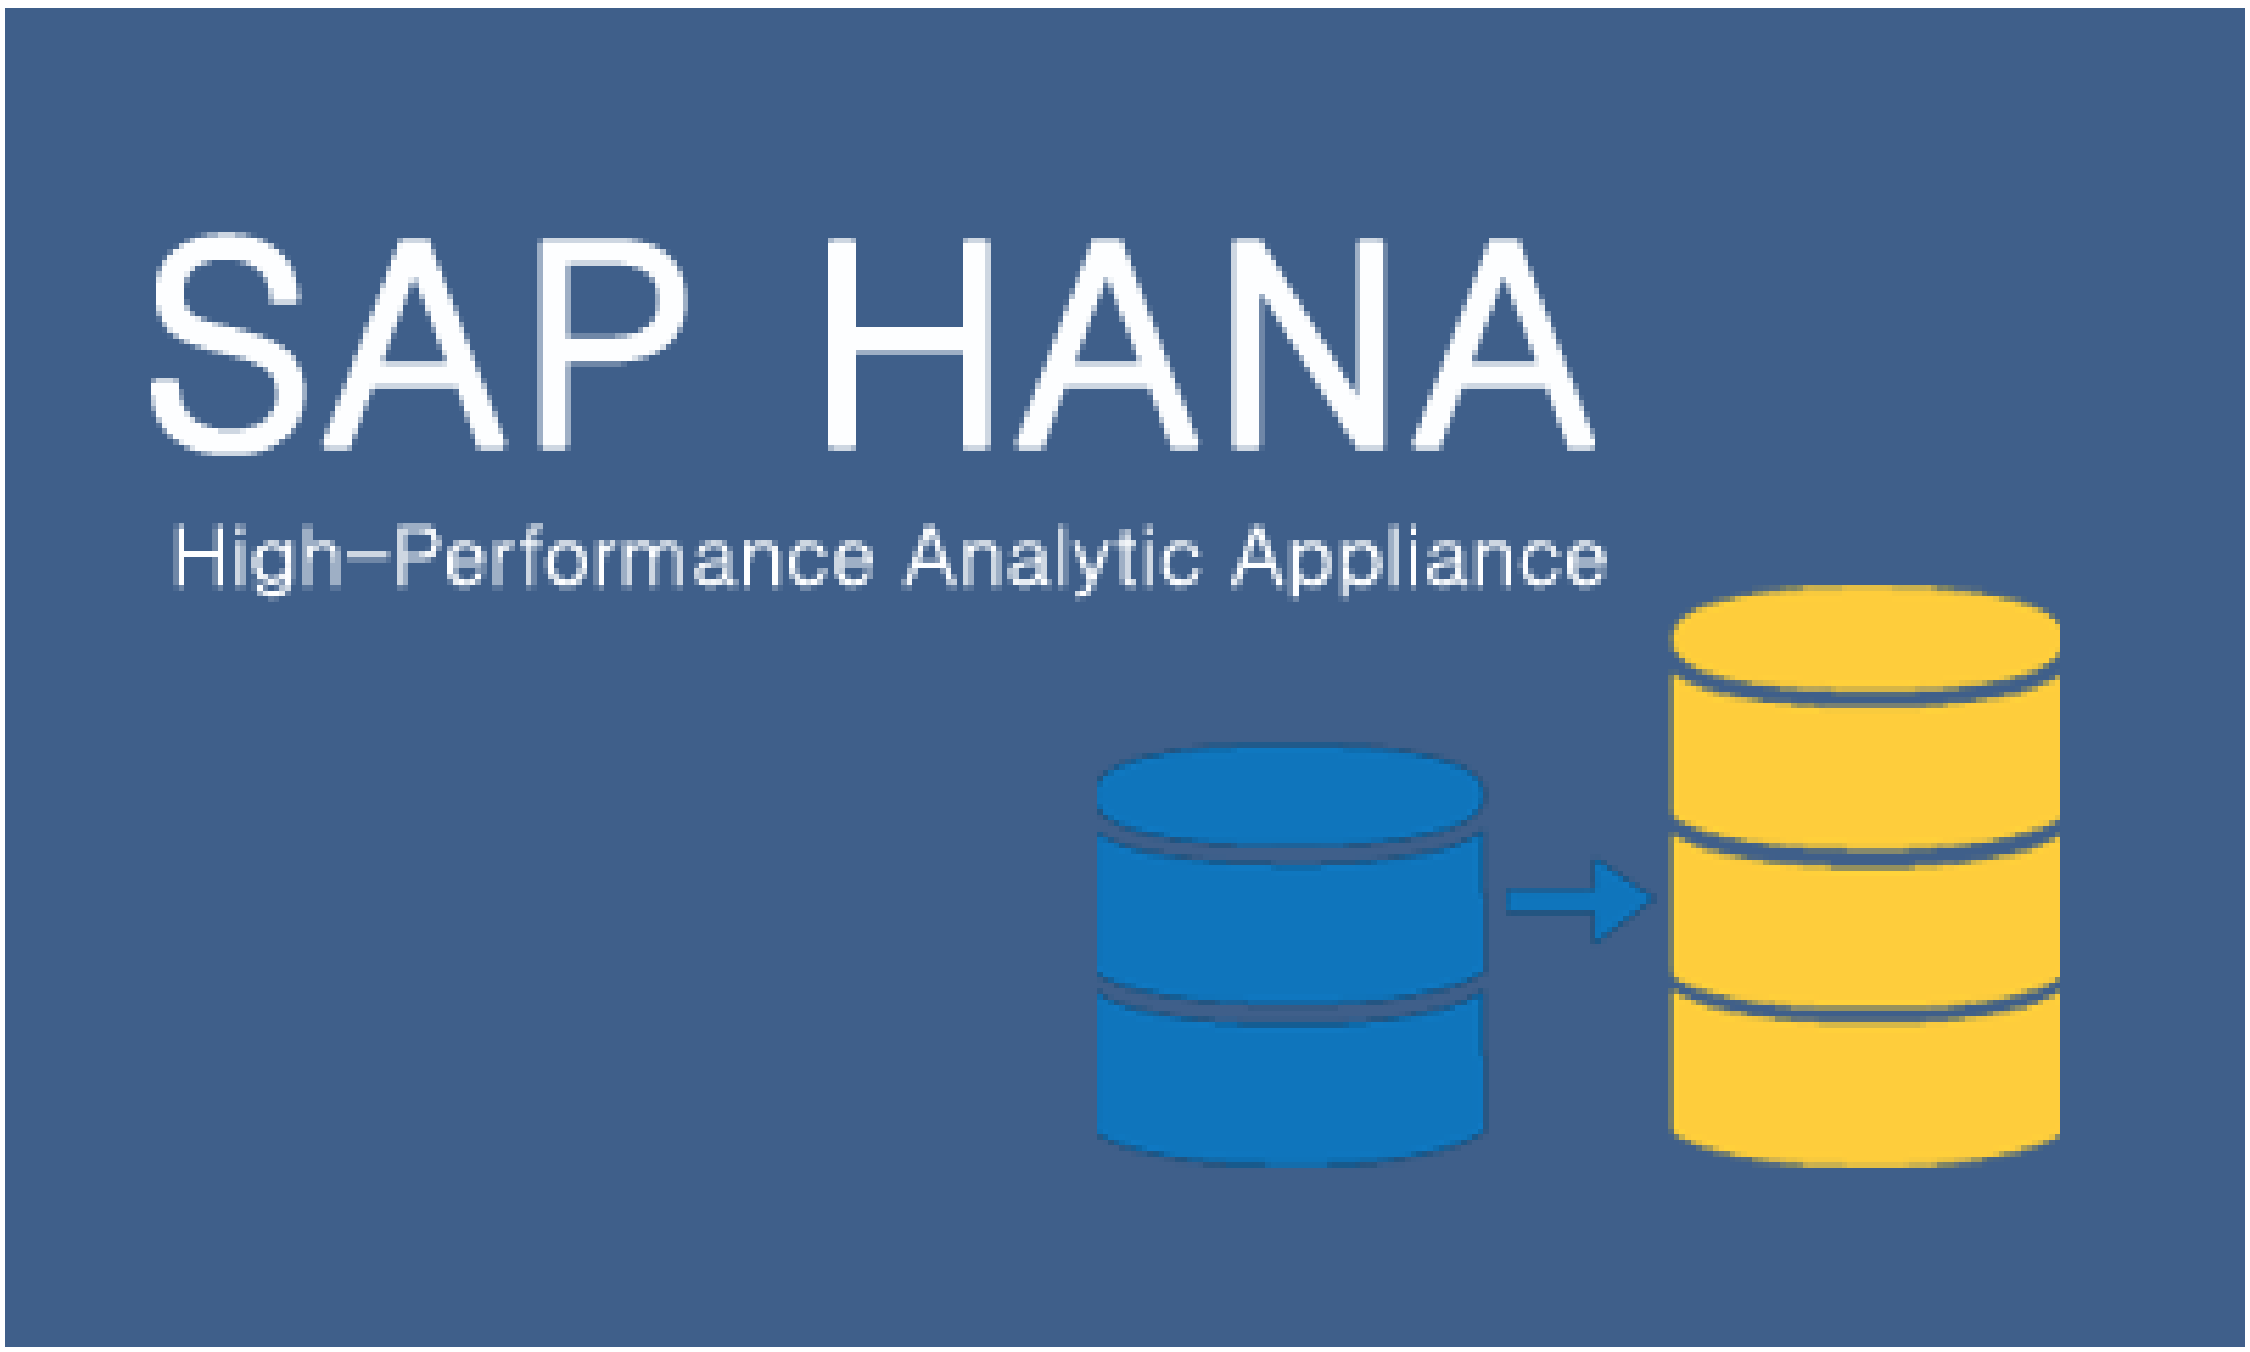

+91 9154555866 - praveen@praveensd.com - www.praveensd.com

Need of today data base features

- Real time data reporting for Analysis
- High speed processing
- Ability to process structed and unstructured data
- Ability to connect Bigdata
- Ability connect IOT and process real time data
- Ability to process Spacial data
- Ability to connect to different sources of data and report at same data with no data redundancy

History of SAP HANA

- SAP HANA is the synthesis of three separate products – TREX search engine, P*Time in-memory OLTP database, and MaxDB in-memory liveCache engine.

- In 1996, a student project at SAP, in collaboration with DFKI, began development of TREX (Text Retrieval and Information Extraction), a search engine.

- TREX became a standard component in SAP Net Weaver in 2000.

- In-memory attributes were added in 2002 and columnar data store was added in 2003, both as ways to enhance performance.

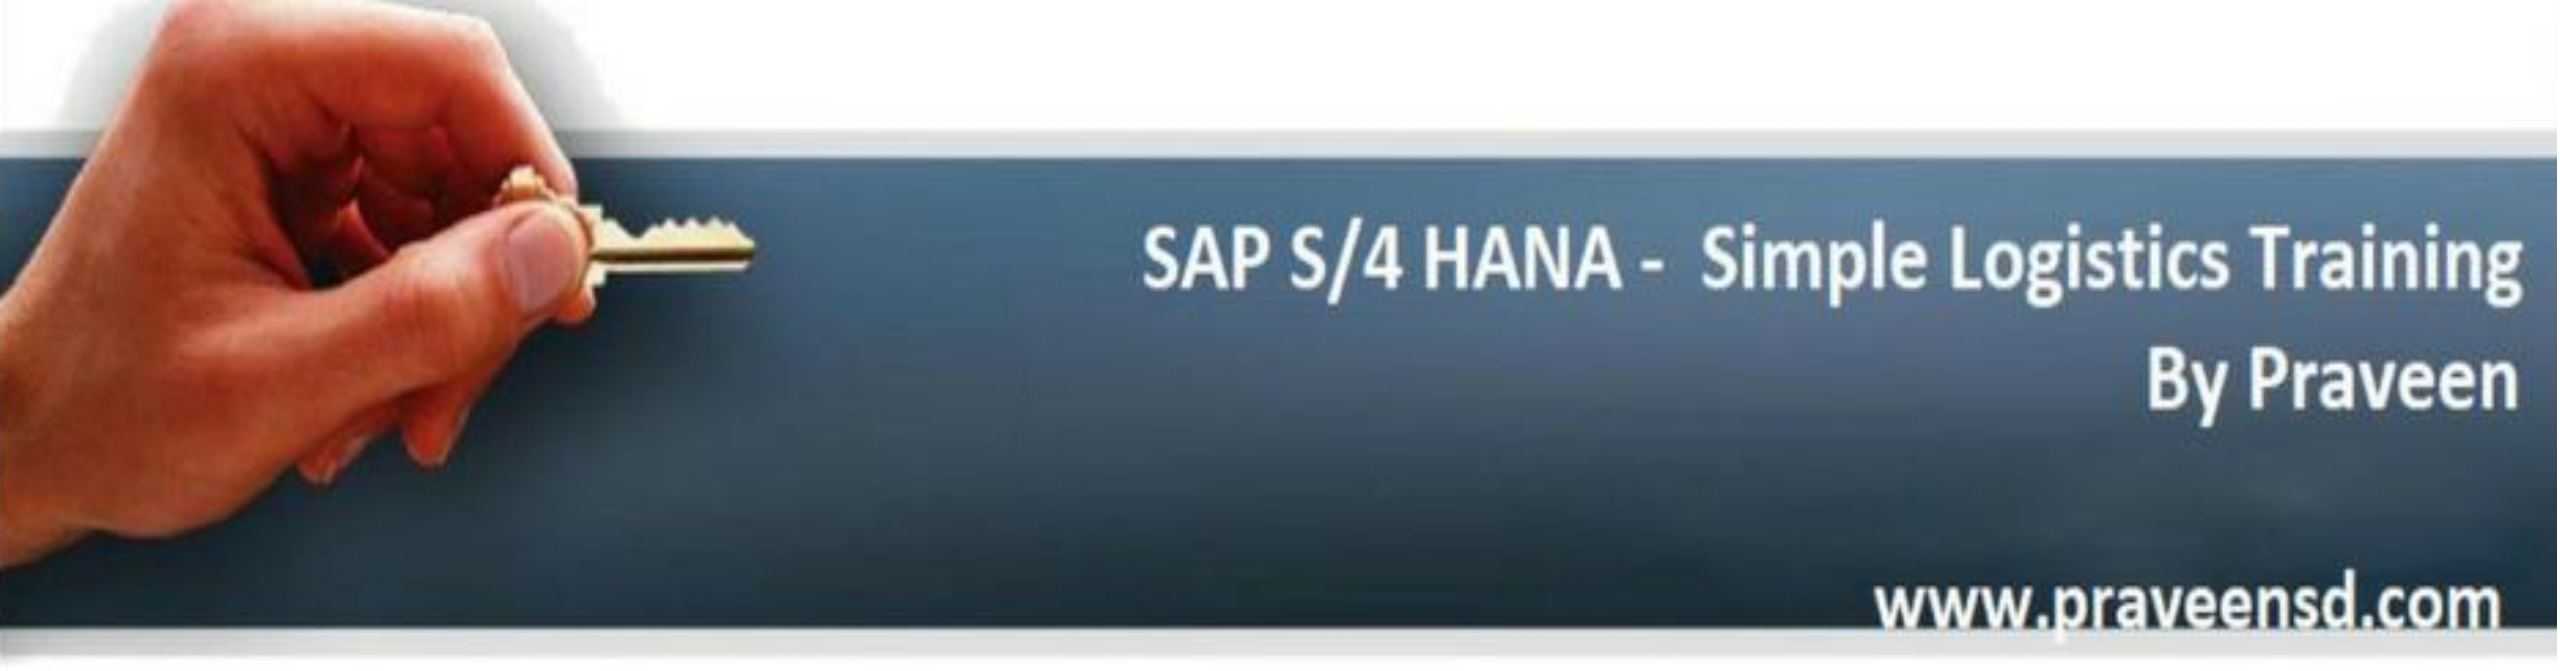

- In 2005 SAP acquired Menlo Park-based Transact In Memory, Inc. With the acquisition came P*Time, an in-memory, light-weight online transaction processing (OLTP) RDBMS technology with a row-bases data store.

- MaxDB (formerly SAP DB), a relational database, came from (first) Nixdorf, (second) Software AG (was named adabas D), and then (third) SAP. It was added to TREX and P*Time to provide persistence and more traditional database features, like backup.

- In 2008, teams working from Hasso Plattner Institute and Stanford University developed this “New Database” as it was called.

- First shipment was in November 2010, support for BW available in November 2011, support for ERP available in May 2013.

What is HANA - Is it a data base or ERP or Software ?

- HANA is a solution for in-memory computing, Acronym HANA means "**High Performance Analytic Appliance**".
- SAP HANA is a flexible data source-agnostic appliance that enables customers to analyze large volumes of data in real time.
- HANA DB takes advantage of the low cost of main memory (RAM), data processing abilities of multi-core processors, and the fast data access of solid-state drives relative to traditional hard drives to deliver better performance of analytical and transactional applications.
- It offers a multi-engine query processing environment which allows it to support both relational data (with both row- and column-oriented physical representations in a hybrid engine), as well as graph and next processing for semi and unstructured data management within the same system

Why to choose SAP HANA?

- SAP HANA is a next –generation in-memory platform. It accelerates analytics and application on a single and in-memory platform.
- Mentioned below are the few reasons why to choose SAP HANA – **Real Time** – SAP HANA provides Real – Time Data Provisioning and Real-time Reporting.
- **Speed -** SAP HANA provide high speeds processing on massive data due to In-Memory Technology.
- **Any Data/Source -** SAP HANA can access various data source including Structured and Un-Structured data from SAP or Non-SAP data source.
- **Cloud** – SAP HANA database and application can be deployed to the Cloud environment.
- **Simplicity** – SAP HANA reduce efforts behind ETL process, Data Aggregation, Indexing, and Mapping.
- **Cost** – SAP claims that SAP HANA software can reduce Total IT cost of a company.
- **Choice Option** – SAP HANA is supported by different hardware vendor and software provider, so based on the requirement, the user can choose the best option.

Advantages of SAP HANA

- By In-Memory Technology user can explore and analyse all transactional and analytic data in real time from virtually any data source.
- Data can be aggregated from many sources.
- Real-time replication services can be used to access and replicate data from SAP ERP.
- SQL and MDX interface from third party support.
- It provides information modelling and design environment.

Reasons To Choose SAP HANA

- Accelerates analytics – In future we can analyze huge Meter data fro smart meters and billing analysis
- Applications on a single, in-memory platform as well as combining databases, data processing, and application platform capabilities.
- Faster Business transactions
- Advanced analytics
- Social media – Text Analyse possible
- Mobile experience – SAP Fiori can be used and good number of APPS available

HANA Live

Specific Analytical UI
(HTML5-based, …)
Generic Analytical UI
(BOBJ, 3rd Party, …)
Web Resource
oData Service
XS Engine
Application Services
SAP HANA PLATFORM
SAP-delivered Query Views
Personalized Views
Personalized Views
Personalized Views
Customer Extensions
Reuse Views
Private Views
Virtual Data Model
HANA Live
Database Layer
Physical Tables

SAP S/4 HANA - Simple Logistics Training

By Praveen

www.praveensd.com

Spatial tables on HANA and its uses

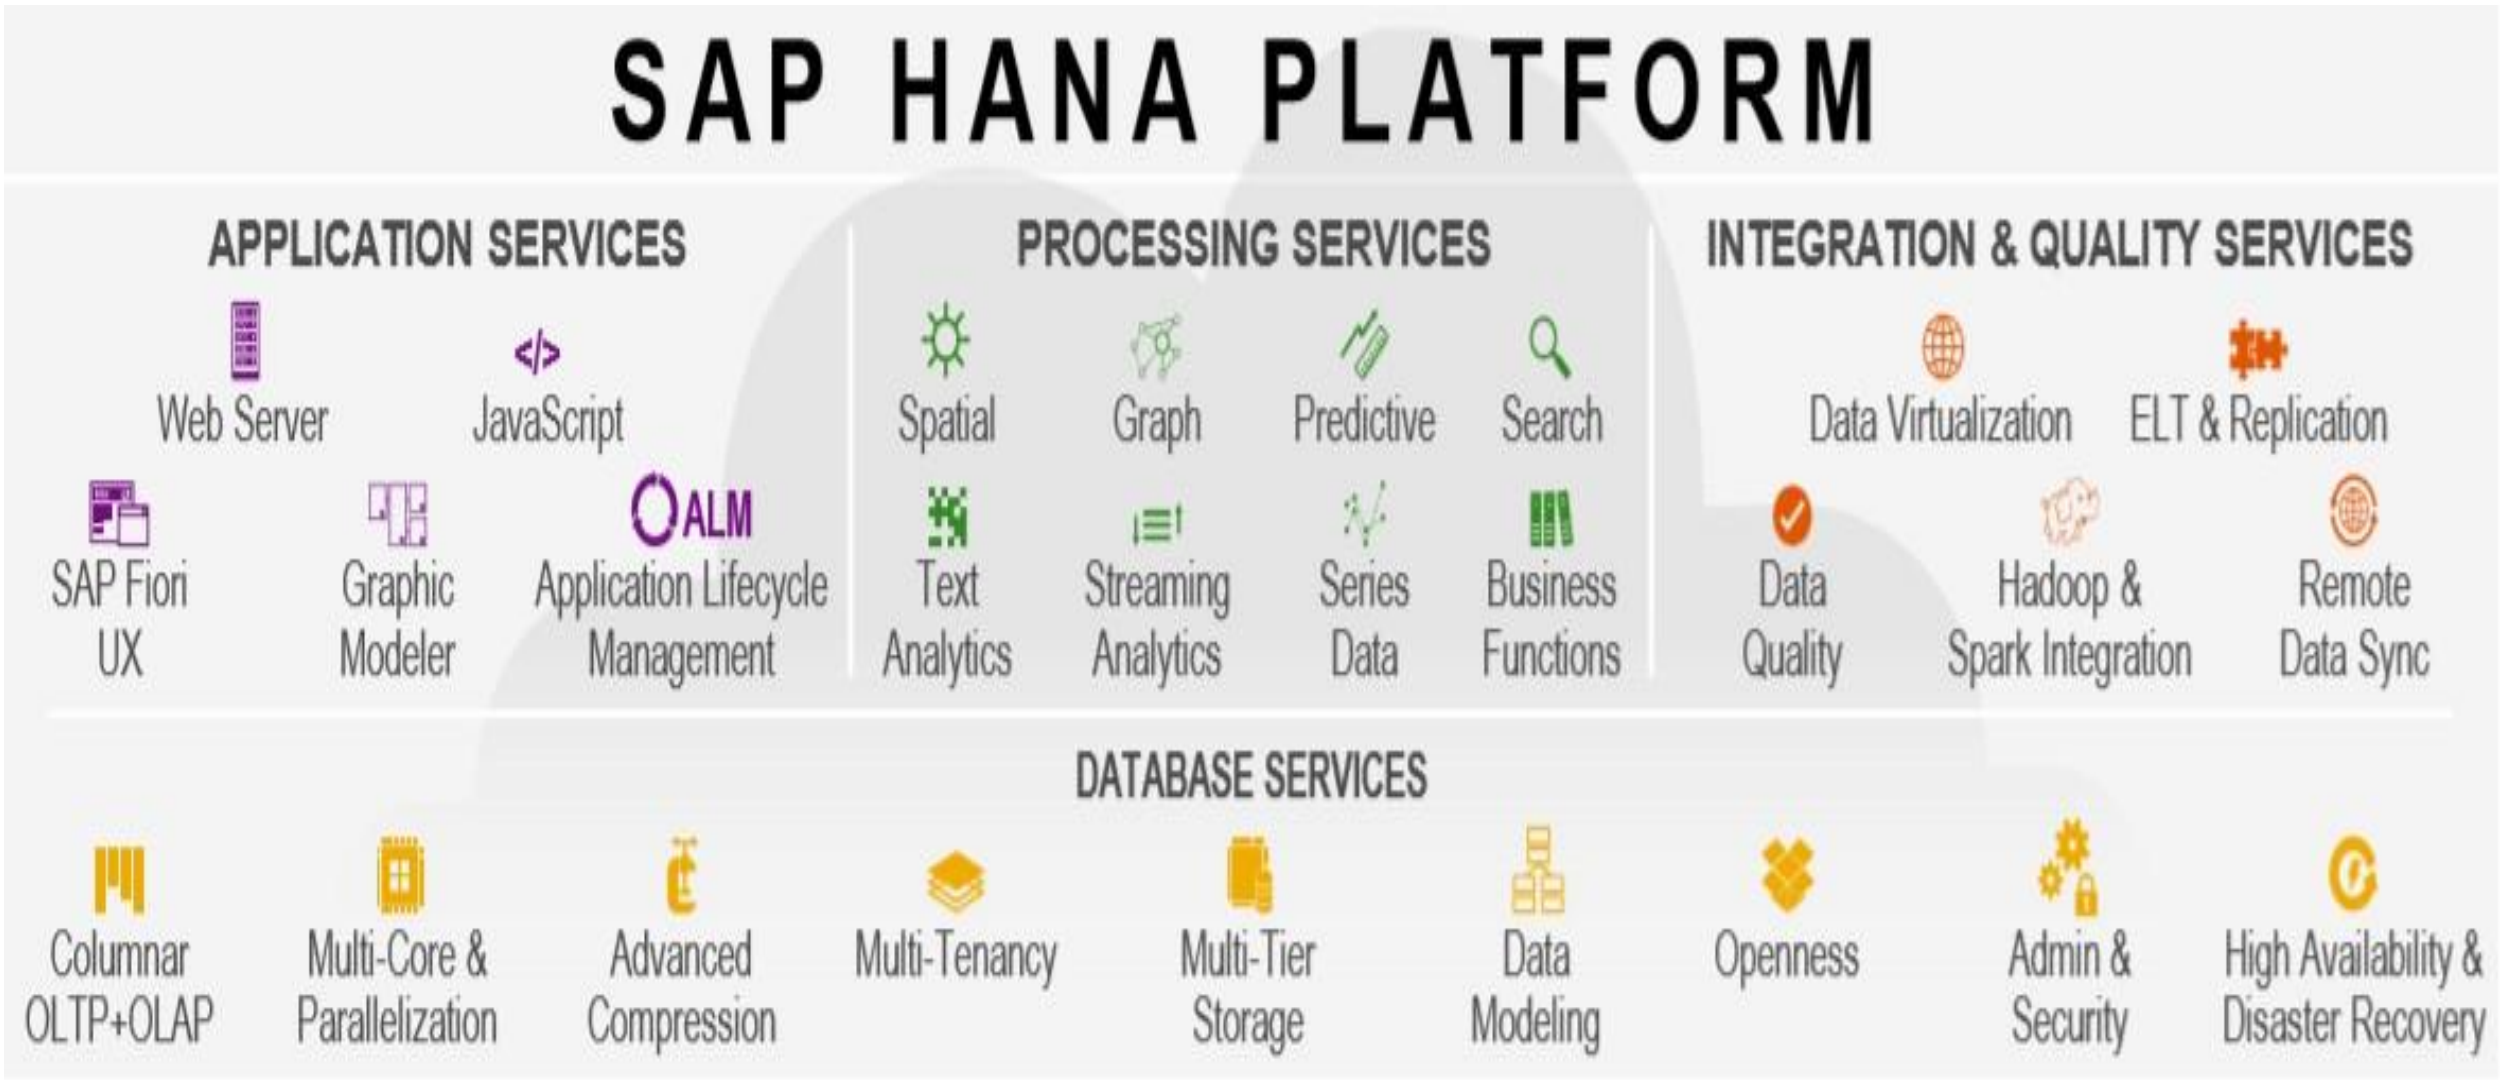

ONE Open Platform **OLTP + OLAP** **ONE Copy of the Data**

+91 9154555866 - praveen@praveensd.com - www.praveensd.com

SAP S/4 HANA - Simple Logistics Training
By Praveen

www.praveensd.com

SAP HANA Platform

Describe why SAP HANA powers SAP S/4HANA

SAP HANA PLATFORM

ON-PREMISE | CLOUD | HYBRID

Application Services

- Web Server
- JavaScript
- Fiori UX
- Graphic Modeler
- Application Lifecycle

Processing Services

- Spatial
- Graph
- Predictive
- Search
- Text Analytics
- Planning
- Data Enrichment
- Series Data
- Function Libraries

Integration Services

- Data Virtualization
- ELT & Replication
- Streaming (CEP)
- Hadoop Integration
- Remote Data Sync

Database Services

- Columnar OLTP+OLAP
- Multi-Core/ Parallelization
- Advanced Compression
- Multitenancy
- Multi-Tier Storage
- Data Modeling
- Open Standards
- High Availability/ Disaster Recovery

Hardware Architecture trends

•SAP HANA is the foundation for SAP S/4HANA and provides many of its critical services, so it is worth taking the time to learn a little about it.

•S/4HANA is a business suite that has its own application server. The application server on which S/4HANA is based is SAP NetWeaver AS ABAP. This is the same application server as Business Suite, but is upgraded to suit S/4HANA.

•The application server sits on top of the database, in this case, and SAP HANA provides all the database services that S/4HANA requires.

•However, SAP HANA is far more than a database. It is an application and data management platform with a very large portfolio of capabilities that power the new applications that require real-time, instant responses on a variety of different data types.

•The figure illustrates the services provided by SAP HANA, and as you can see, there are quite a few

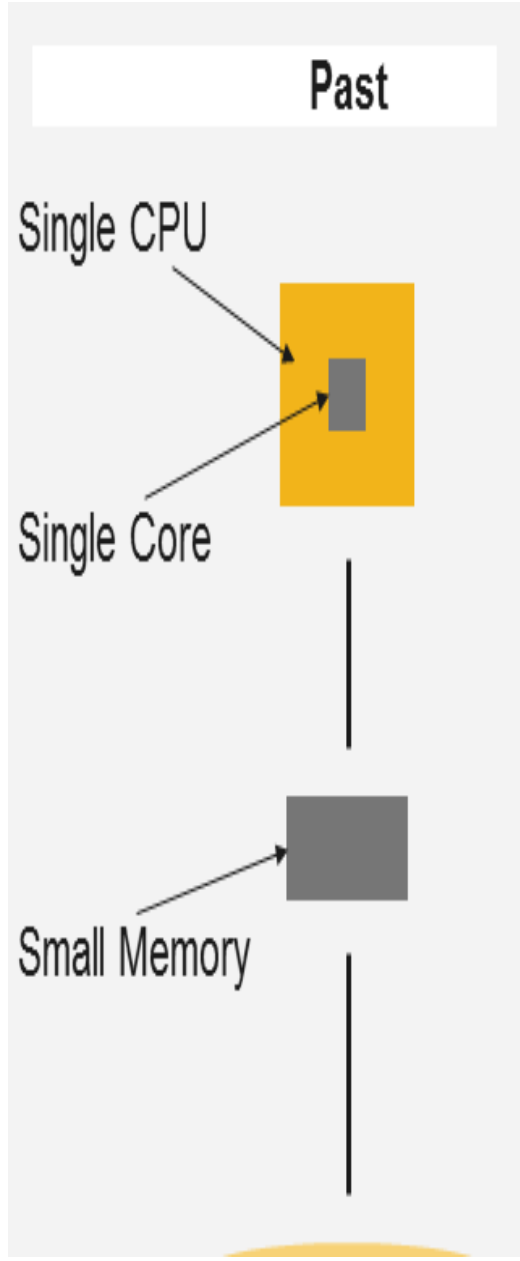

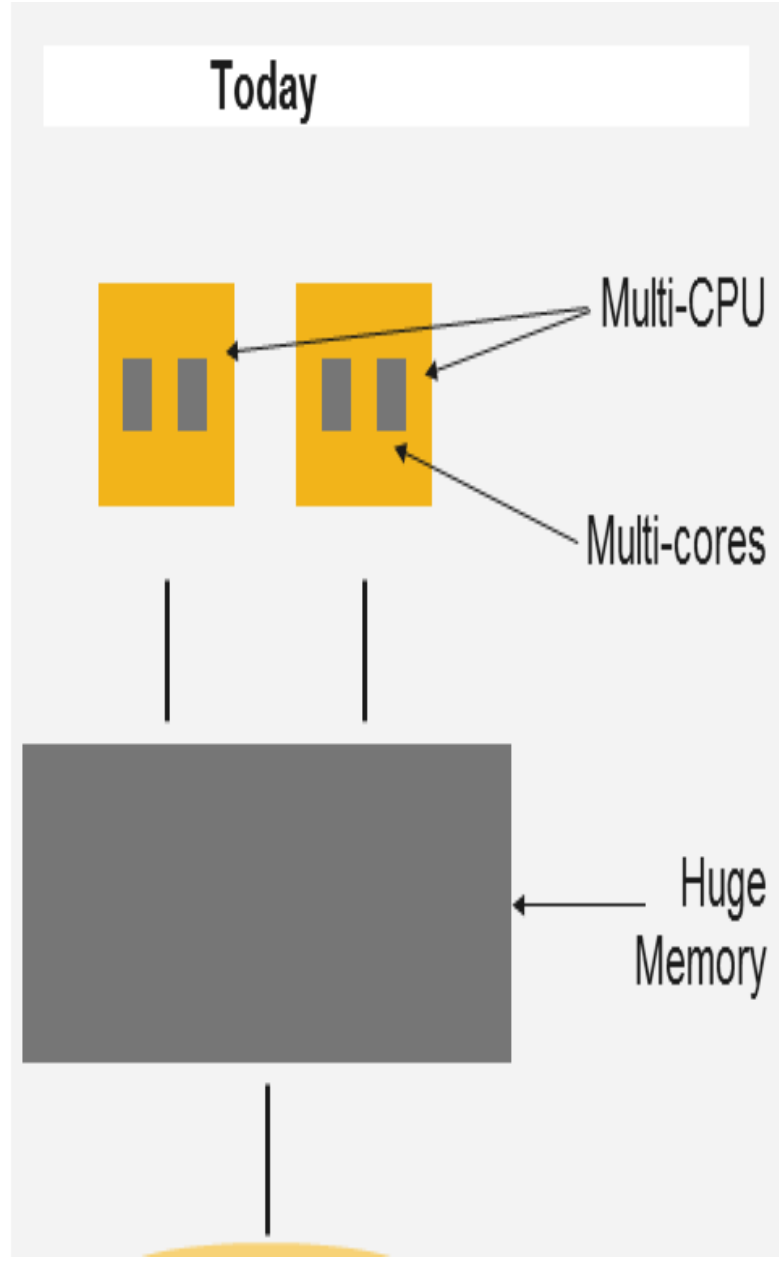

Recent trends in hardware evolution

•SAP HANA takes full advantage of the recent trends in hardware evolution.

•Historically, the high cost of memory meant that only small amounts were available. This caused a serious bottleneck in the flow of data from disk to CPU (see the figure), with the CPU waiting idle for data to arrive through the tiny gateway.

•Now with memory prices falling, we have access to huge amounts. SAP HANA runs on hardware with many terabytes of memory. In fact, with so much memory available, the entire database of even a large organization can be stored completely inside memory, so there is instant access to all data and wait times are eliminated. Memory is no longer the bottleneck it once was.

•In addition to huge memory, the processors continue to improve at a phenomenal rate. We have high-speed, multi-core processors that can take on complex tasks and process them in parallel. This means response times for even the most complex analytical tasks, such as predictive analysis, can be carried out in real time.

•SAP could have kept the same business applications produced 20 years ago, along with the traditional databases that supported them, and installed all of these on the new hardware. There would be some gains, but traditional databases and applications were designed around old, restricted hardware architecture. This means they would not be able to fully exploit the power of the new hardware.

•Put simply, the business software needed to catch up with advances in hardware technology, so a complete rewrite of the business suite was required.

Column Store And Row Store

Column Store and Row Store Tables

Row Store

Country	Person	Quantity
France	Ned	10
France	Tony	12
UK	Ned	8
Germany	Anita	4

Column Store

Country
France
France
UK
Germany

Person
Ned
Tony
Ned
Anita

Quantity
10
12
8
4

SAP HANA database Properties

•The SAP HANA database is fully in-memory, so it is very fast.

•The SAP HANA database is fully ACID-compliant. This means Atomicity, Consistency, Isolation, Durability. This is the mark of a database that is built to be 100% reliable for mission-critical applications, where fast, simultaneous read and write operations are applied to the same data sets. The ACID standard guarantees there will never be partially updated records. You can fully trust the data at all times.

•Most traditional business databases are row based. Some specialist analytical databases are column based. SAP HANA is built to support both types. Fast-moving transaction applications usually work better with row store tables, whereas analytical applications that perform a lot of aggregation work better with column store tables.

•Both storage types are needed in a system that handles both transactional and analytical applications in one platform, as is the case with SAP S/4HANA.

•Column store tables are incredibly efficient, especially for analytical applications where access to data sets is not predictable and we often do not know which columns are required. Column store tables work well with aggregation functions, such as sum, average, min, and max. Column store tables are automatically compressed, and can also be optionally partitioned. Column store tables are optimal for parallel processing. Why do we need row tables at all?

•The downside to column store is the cost of reconstructing complete records from the columns if all data is required by the application. This is the case when the application is transactional and all fields are needed for an update, insert, and delete. Additionally, for write-intensive applications, column store tables are not optimal compared to row store tables.

•Row storage is still needed to support fast-moving transaction processing when aggregation is not the main priority. Row store tables are not compressed and cannot be partitioned.

•SAP S/4HANA combines transactional and analytical applications, so it utilizes both column and row store tables.

Better Utilisation of Memory

Logical Table

Country	Person	Quantity
France	Ned	10
France	Tony	12
UK	Ned	8
Germany	Anita	4

Row Store

France
Ned
10
France
Tony
12
UK
Ned
8

Column Store

France
France
UK
Germany
Ned
Tony
Ned
Anita
10

Better Memory Utilization with ColumnStore

Better Memory Utilization with Column Store

UK	Ned	8
Germany	Anita	4

France
Ned
10
France
Tony
12
UK
Ned
8
Germany
Anita
4

France
France
UK
Germany
Ned
Tony
Ned
Anita
10
12
8
4

Classification of data and Memory Management

Disk is still required for logging and backup in case of power failure. However, the disk is also required to store data that has been displaced from memory.

Memory is now huge and relatively affordable. It is technically possible to store an entire enterprise database in memory, especially if you implement multi-terabyte memory. However, for most organizations, most of the data that they own is not frequently used, so they really do not need to implement such huge memory sizes. Only recent data is frequently used.

This may well be only 5-10% of the entire company's data, which is called hot data. The rest of the data, which makes up 90- 95%, is called warm data.

With SAP HANA, hot data is stored in memory, and warm data is stored on disk.

Whenever older data is needed by an application, it is loaded from disk to memory and the application reads the data from memory. This data may not be needed again for a long time, so it is displaced from memory at the moment when the memory is full and other, more recent data, replaces it. The older data then goes back to disk until it is needed again.

For row store tables, loading and displacement happens at the row level. This means all columns in the row, whether they are needed or not, are loaded to memory. For analytical applications that require only few columns, this is inefficient, as it involves moving all columns to memory, even those not used.

For column store tables, loading and displacement happens at the column and partition level. This means that only the required columns, and even better, only the required partitions in the columns, are loaded to memory. This is very efficient for analytical applications, which often only ask for small portions of data

Reducing the data Footprint

Logical Table

Order	Country	Product	Sales
456	France	corn	1000
457	Italy	wheat	900
458	Spain	rice	600
459	Italy	rice	800
460	Denmark	corn	500
461	Denmark	rice	600
462	Belgium	rice	600
463	Italy	rice	1100
…	…	…	…

Dictionary

1	Belgium
2	Denmark
3	France
4	Italy
5	Spain

5 entries →
3 bits to encode

Compressed column

1	3
2	4
3	5
4	4
5	2
6	2
7	1
8	4
…	…

Where was order 460?

Inverted index

1	7
2	5,6
3	1
4	2,4,8
5	3

Which orders in Italy?

Benefits of reduced data Footprint

•The data in the SAP HANA column store tables is automatically compressed in order to reduce the data footprint.

•The following are a number of benefits associated with a reduced data footprint:

•You can get more data into the CPU cache, and therefore reduce main memory access, in order to maintain high performance.

•You can fit entire enterprise databases into memory and avoid disk access.

•Operations such as backup and restore are speeded up as data sizes decrease.

•The amount by which data reduction can take place is driven by the shape of the business data. Compression is most impressive when there is a lot of data repetition in the tables. An example is a huge sales order table, in which the customer type is stored on each customer order, but there are only three customer types. The customer type is repeated many times across the table.

•Compression strips out the repetition and uses integers to represent the business values. Then it uses special dictionary tables to hold the distinct list of business values and the corresponding integers. This all happens in the background, and is not visible to the business user. It is also not something with which the developer needs to be concerned.

Parallel Processing

A key theme of SAP HANA is parallel processing. With the new hardware architecture, especially utilizing the new multi-core processors, you can ensure instant responses by spreading out the processing tasks across the cores.

SAP HANA automatically spreads the workload across all processors and ensures all parts of the hardware are contributing to the throughput.

SAP HANA is scalable, which means you can add more processors, as required, to increase the parallelization, and therefore the speed, of processing.

In addition, you can manually partition column tables to influence the parallelization based on common business values that are accessed frequently.

Parallel processing is a key enabler for real-time processing, on which many new S/4HANA applications are based.

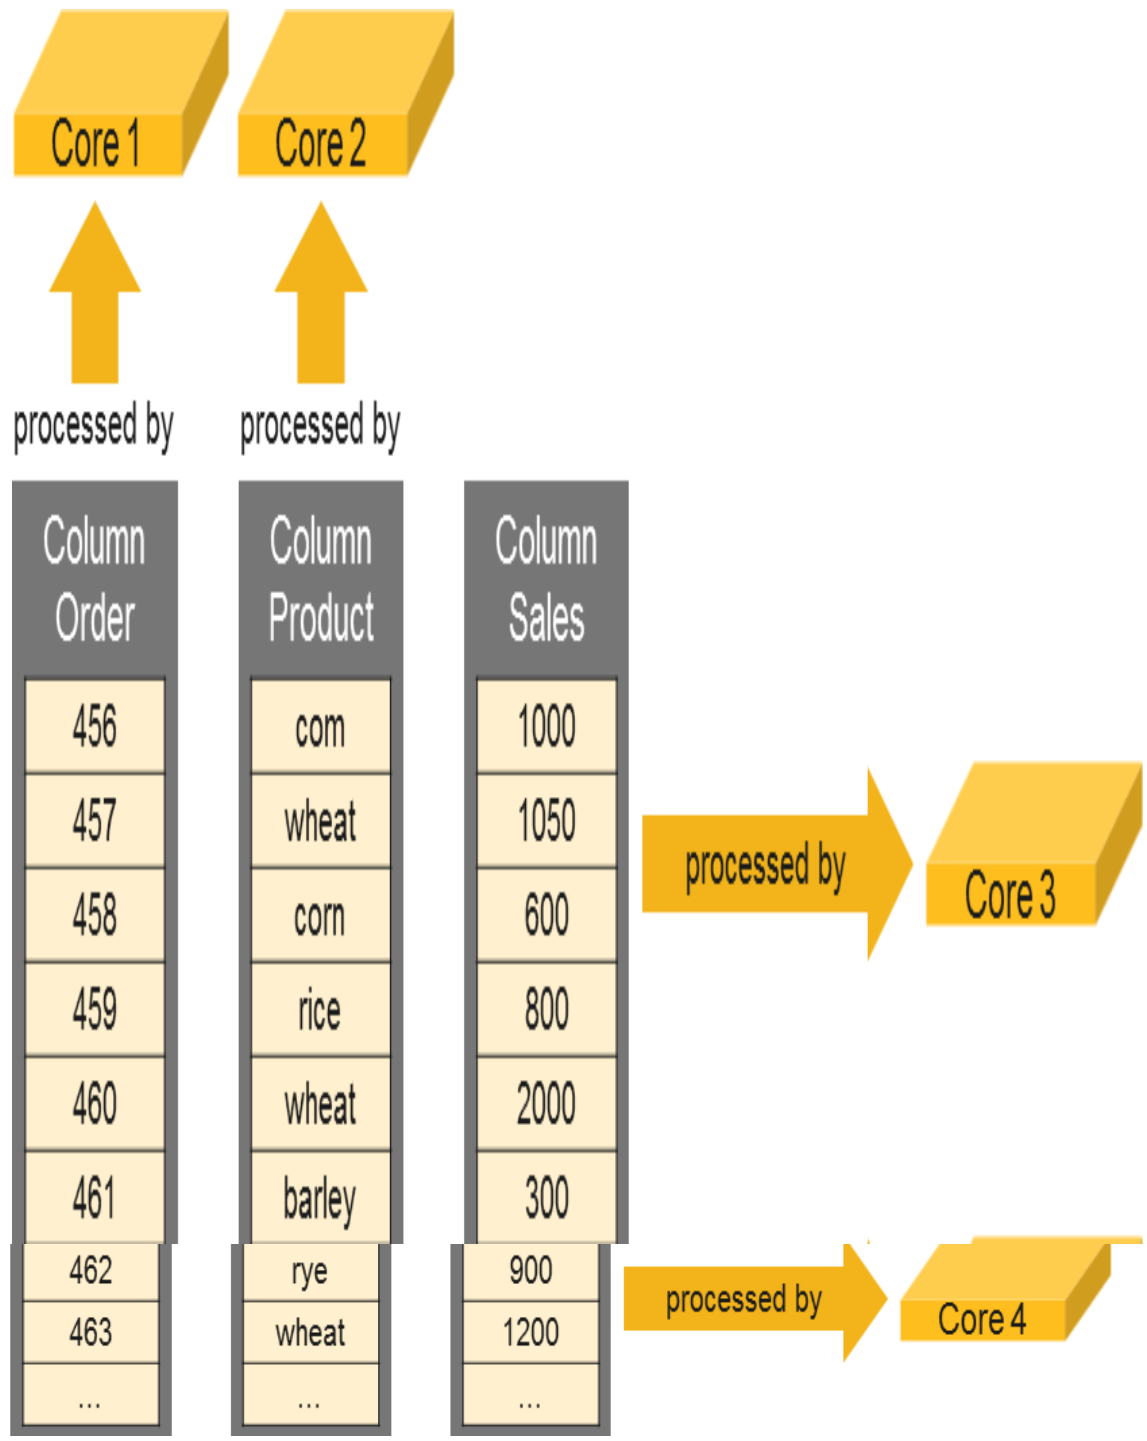

Push Down Processing to SAPHANA

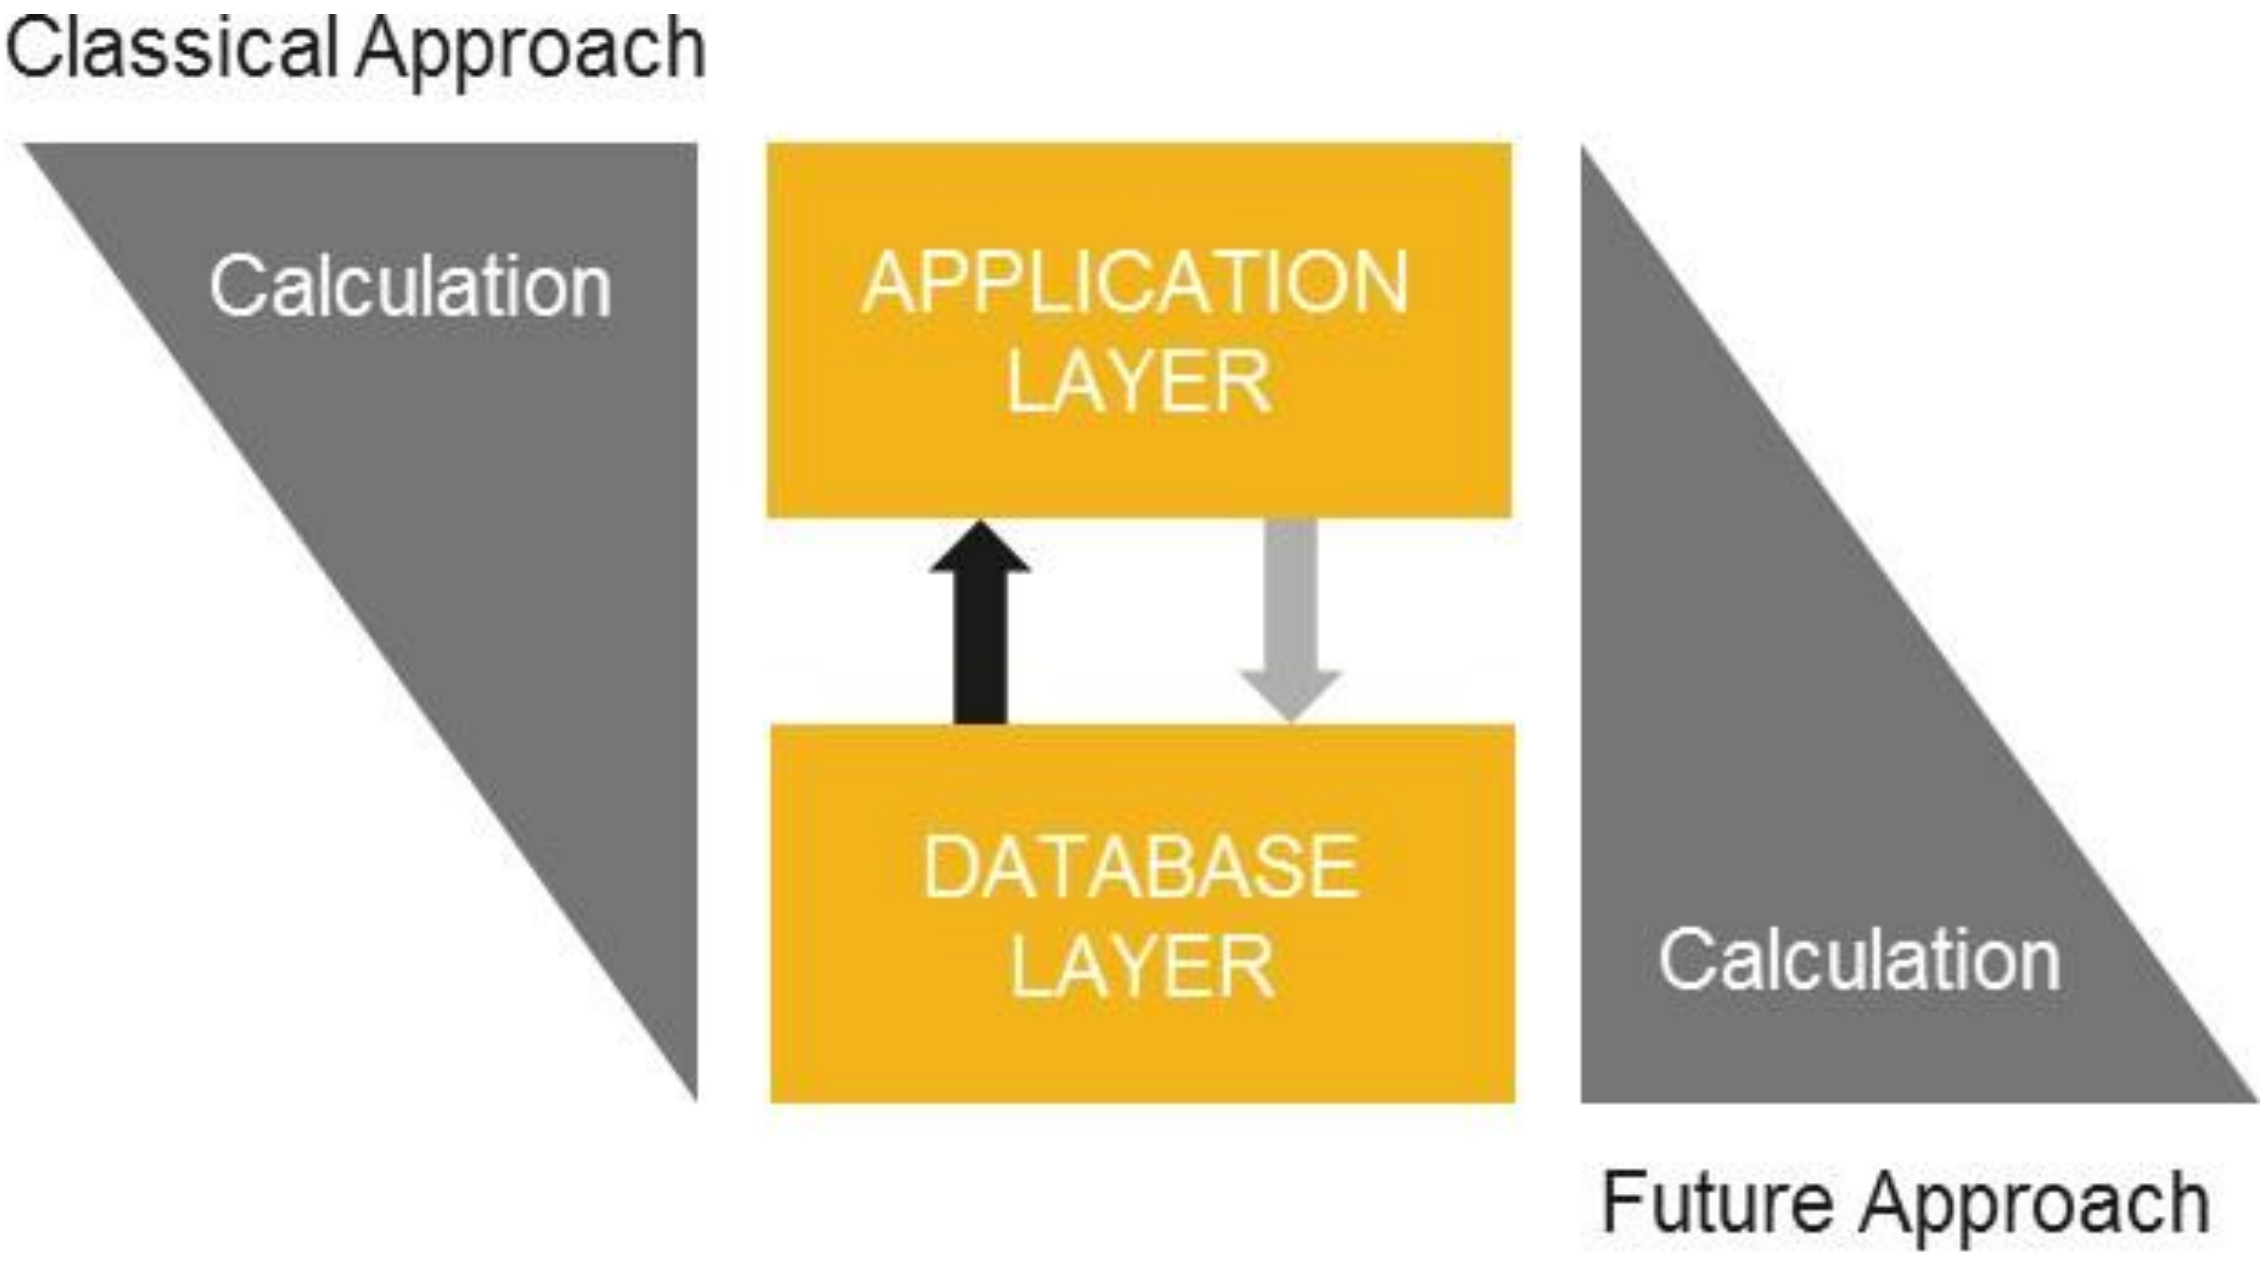

→ MOVE calculations to database

→ Only transfer RESULTS

Data Intensive Processing

•In the past, the key job of the database layer was to listen out for requests for data from the application server and then send that data to the application server for processing. Once the data had been processed, the results were sent back down to the database layer for storage.

•SAP HANA is capable of taking over many of the processing tasks from the application server. All data-related tasks, such as aggregation, filter, sort, calculate, and predict can be handled by SAP HANA.

•Now the application layer simply needs to tell SAP HANA what is to be done on the data, and SAP HANA processes the data and send only the results. This is done in memory, so speeds can be impressive. We call this code-to-data, as opposed to the traditional way, which was data-to-code.

•The application layer is still needed with SAP S/4HANA to handle the complex business logic that must be programmed in a business programming language. In the case of S/4HANA, this is ABAP. However, many simpler applications can be built directly on SAP HANA, with no need for an additional application server.

A key theme of SAP HANA is parallel processing. With the new hardware architecture, especially utilizing the new multi-core processors, you can ensure instant responses by spreading out the processing tasks across the cores.

SAP HANA automatically spreads the workload across all processors and ensures all parts of the hardware are contributing to the throughput.

SAP HANA is scalable, which means you can add more processors, as required, to increase the parallelization, and therefore the speed, of processing.

In addition, you can manually partition column tables to influence the parallelization based on common business values that are accessed frequently.

Parallel processing is a key enabler for real-time processing, on which many new S/4HANA applications are based.

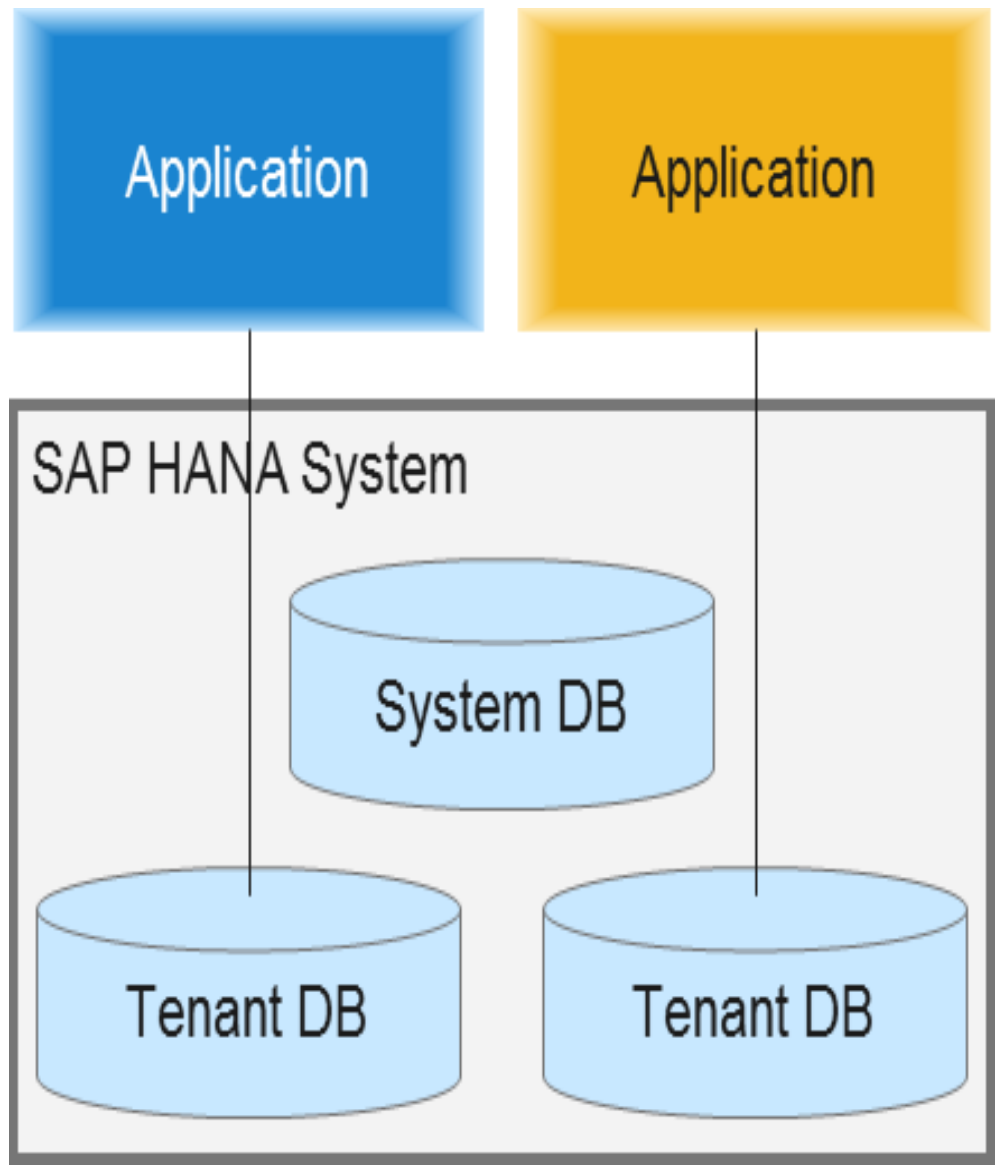

Auto recovery

•SAP HANA utilizes memory for storage and once the power is gone, we lose the data in memory.

•How does SAP HANA ensure we do not lose data when the power goes, and how does it get back up and running quickly? SAP HANA's solution for zero-downtime is based on a two-phase approach.

•Every few minutes, SAP HANA automatically takes a snapshot of the entire memory and stores this on a disk layer. This is called a savepoint.

•What happens if the power goes off between savepoints? Do we lose this data? We do not lose data because between savepoints, every committed transaction is also saved to a log area. This log area is often based on flash memory (SSD) to ensure lightning speed, so every update to the database is captured.

•When power is restored, SAP HANA automatically readies the last savepoint, and also re-applies the transactions from the log since that savepoint occurred, to ensure the system is exactly where it was when the power was lost.

•This all happens invisibly in the background.

Auto-recovery when power is interrupted

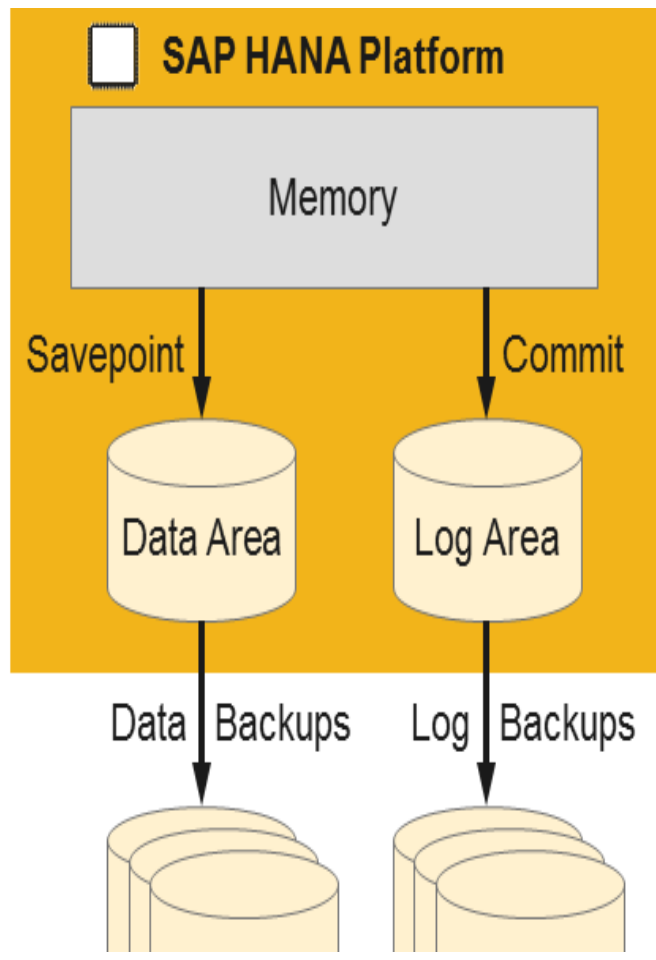

Auto Recovery

•If a server fails, SAP HANA can automatically swap it out to a standby server.

•Standby servers can be on warm standby, which means they are ready to go immediately and do not need to be started. The data is loaded to memory from a backup server. SAP HANA uses the savepoints and log, as mentioned previously, to bring the warm standby server up to date with the data.

•Standby servers can also be on hot standby, which means the standby server is always in sync with the live server, usually by continually replaying the database log. If it is necessary to swap over, this can happen with almost no interruption to processing.

•For mission-critical applications and where SLAs are implemented, you can ensure customers' systems are always running by implementing this approach. This auto-recovery approach is referred to as failover.

Auto-recovery when a server fails

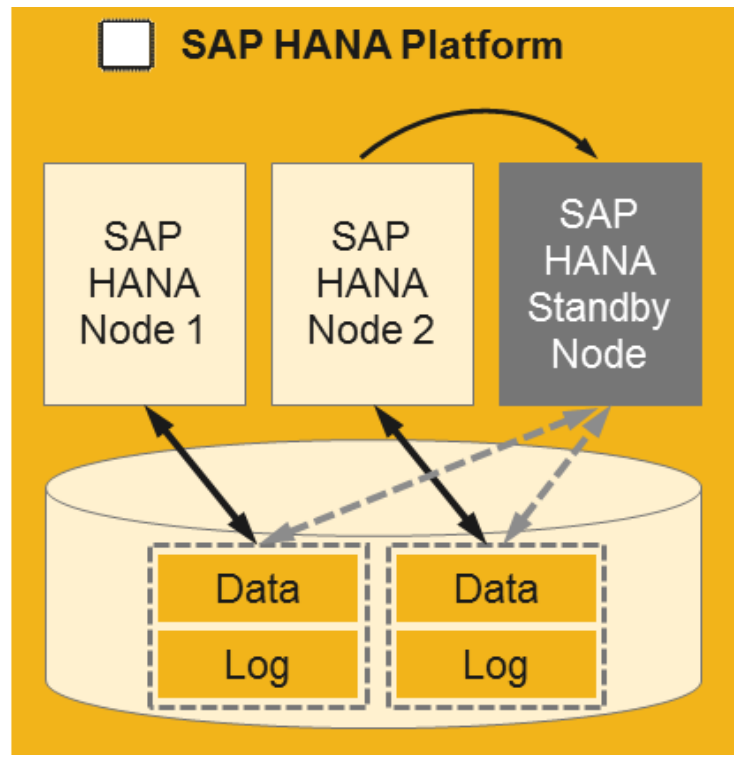

Bigdata

•We know the digital world is creating huge amounts of data. Do we just keep loading this data to SAP HANA?

•Technically, this is possible, but it would not be efficient. Most business applications refer to only a small subset of data, and this is usually the most recent data. You should not fill SAP HANA's in-memory database with data that is old and hardly used.

•SAP HANA allows us to classify our data as active and passive. We also use temperatures as a reference to how hot (useful) the data is. Active, or hot, data is data that is recent or perhaps the focus of a current analysis (even if it is old).

•Passive data is warm, or even cold,data that is older and less used

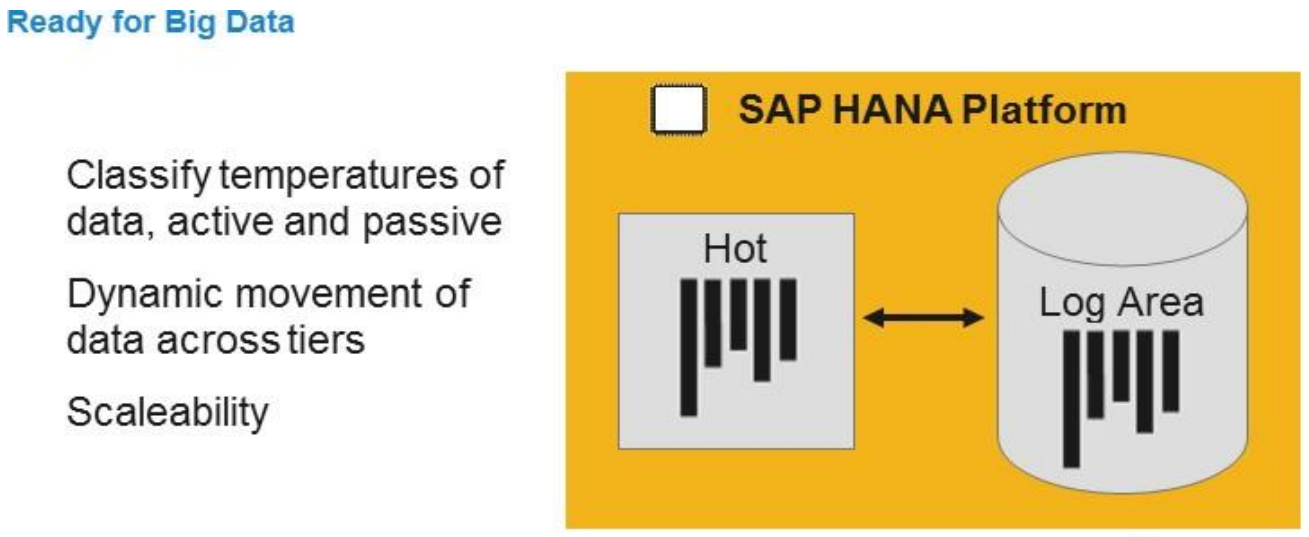

•storage options such as Big Data commodity server solutions (Hadoop) and data archive systems.

•A key point is that, regardless of where the data physically resides, all of it is still available seamlessly to SAP HANA applications. Application developers do not need to know where data is, as this is managed by HANA.

•SAP HANA moves data across the storage tiers automatically based on usage patterns and other programmable business conditions. This ensures a customer's SAP S/4HANA will always run optimally, with no older data clogging up the database.

Realtime Event Processing

•SAP HANA can consume data in many different ways.

•Real-time data can be consumed to power real-time S/4HANA applications. The Internet of Things (IoT) means we will connect large numbers of devices that transmit information continually.

•It is important to remember that real-time data streaming is not the same as real-time data loading. Often, once the data is consumed and processed, it is of no further interest and SAP HANA can ignore it.

Real-time event stream processing

1. **Capture** data arriving continuosly from devices and applications
2. **Act** on new information as soon as it arrives: alerts, notifications and immediate response to changing conditions
3. **Stream** information to live operational dashboards

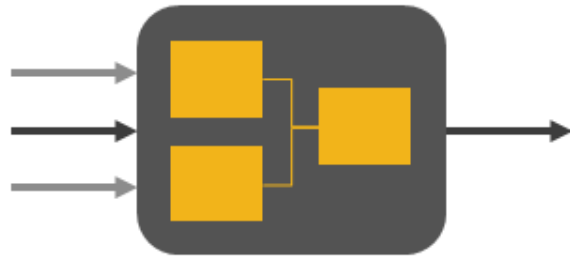

Highly scalable – process hundreds of thousands or even millions of events per second

•The following are examples of devices and activities that could stream real-time data to S/4HANA applications:
•Sensors – machines

•Clickstreams from Web activity

•Social media - respond to consumer sentiments (for example, Twitter)

•Market stock prices

•Energy consumption

•In-game sports analysis

IOT Data Synchronization

•SAP HANA can communicate with devices (IoT) using remote data sync. Often, such devices do not need to be continually online with SAP HANA. We call this "occasionally connected".

•Devices can collect data locally, with their built-in light databases, and SAP HANA can periodically collect this data. For example, every hour a vending machine passes its stock data to SAP HANA. When an item is running low, SAP HANA can pass back a message to the vending machine that a refill is on its way. Remote data sync is bidirectional.

•S/4HANA applications can communicate with IoT devices. There are many innovative enterprise applications that can benefit from communication with devices in the IoT.

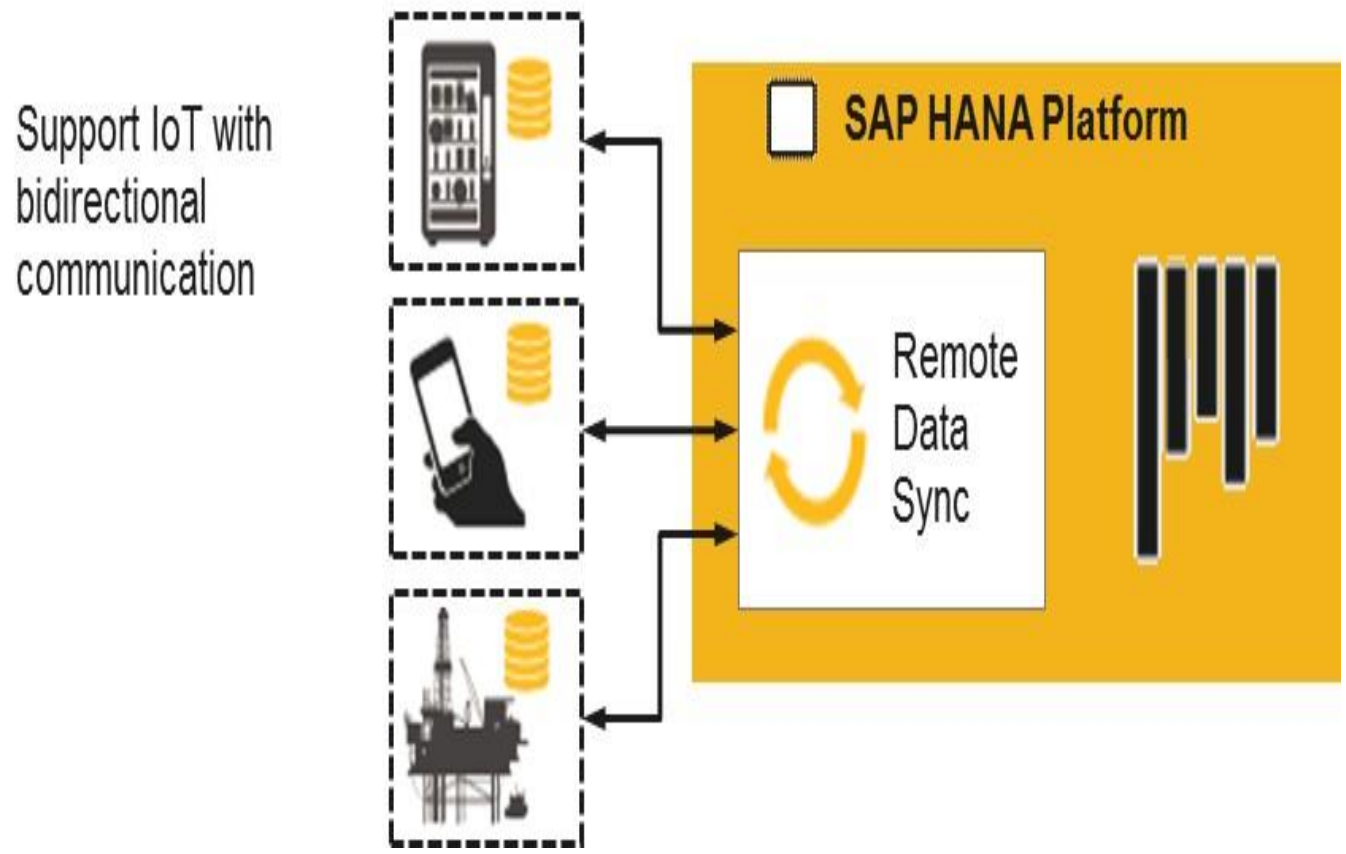

•The same technology is used to connect SAP HANA to remote environments that may operate in hostile conditions, or where the signal is not reliable, such as an engineer working in a lift shaft where the signal is poor, or an oil rig where a satellite passes only once per day to provide communication back to HQ

Data Access From Anywhere and Anytime

Real-time streaming and remote data sync, SAP HANA has many other options for data provisioning.

•Smart Data Access (SDA) allows SAP HANA to access remote database tables and files from any source, as if the data was loaded to SAP HANA. A great use case for this is the integration of Hadoop or data archives, where occasional access to data is required.

•Smart Data Integration (SDI) and Smart Data Quality (SDQ) provides real-time data replication from any source, with the option of enhancing the data quality during the loading process.

•SAP HANA is fully integrated with existing and well-known data-loading tools, such as SAP LT Replication Server and SAP Data Services for real-time and batch data loading.

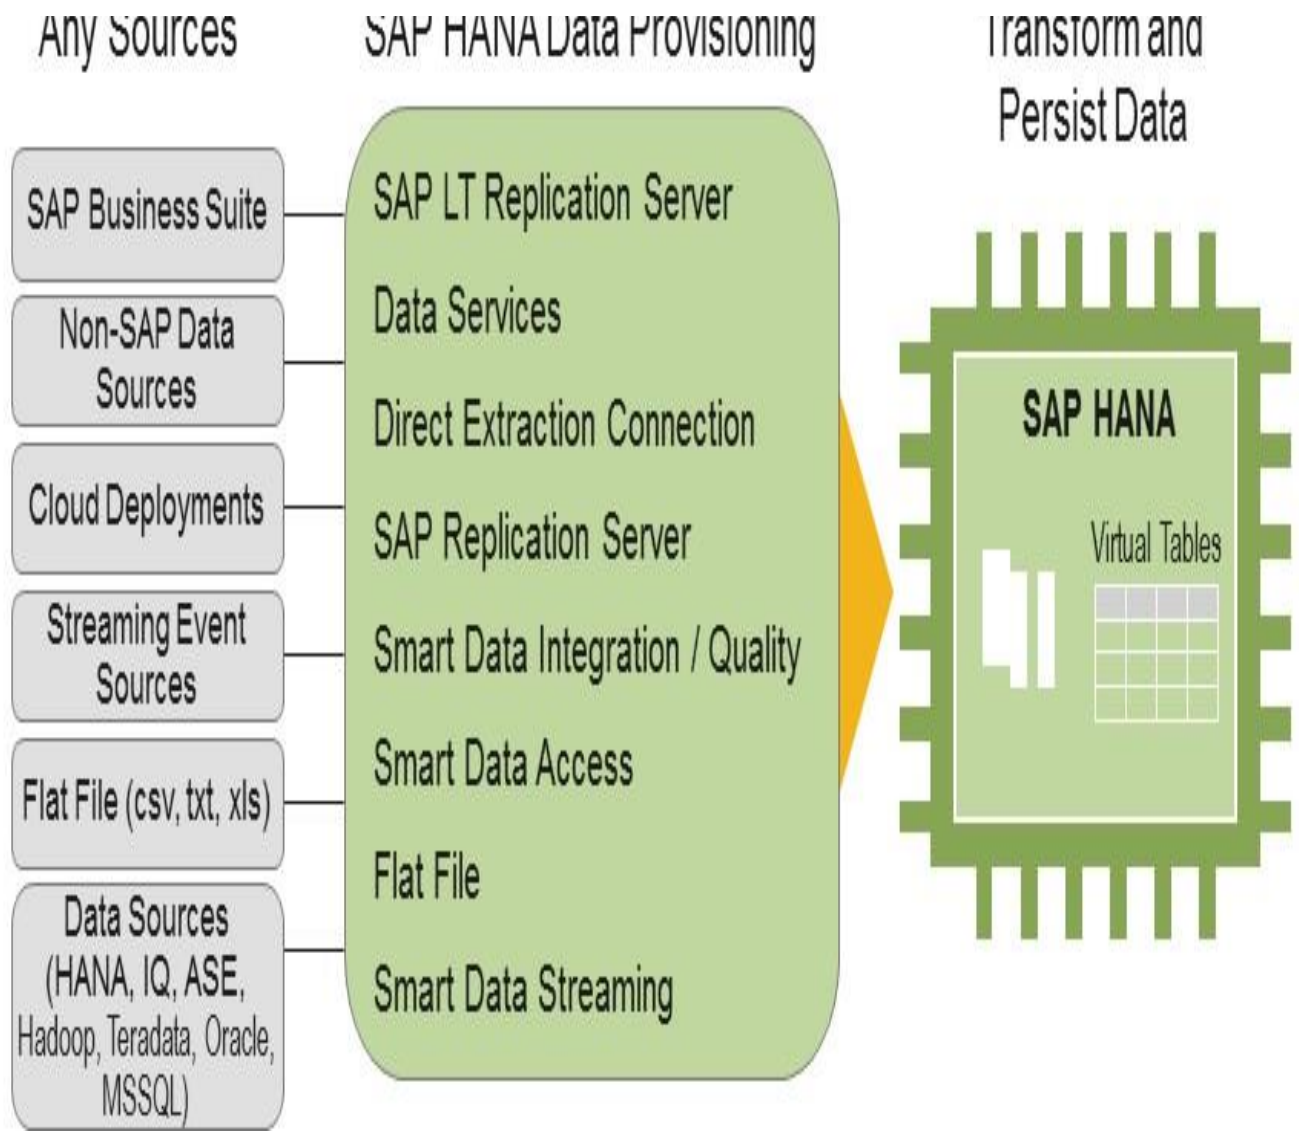

Data Access From Anywhere andAnytime

•If all we asked of SAP HANA was to support the database requests for S/4HANA, then we would be using only a fraction of SAP HANA's capabilities.

•SAP HANA is not just a database, it is also a powerful data processing engine with many built-in capabilities that can enable organizations to develop innovative applications integrating S/4HANA. We call this "extending the core", with the core as SAP S/4HANA.

•Companies implement a digital platform, such as SAP S/4HANA, not only to run their core business processes, but to take full advantage of the new digital world where innovative, disruptive applications can be game changing.

•This is why companies will look to exploit the full potential of the SAP HANA platform to move beyond the core

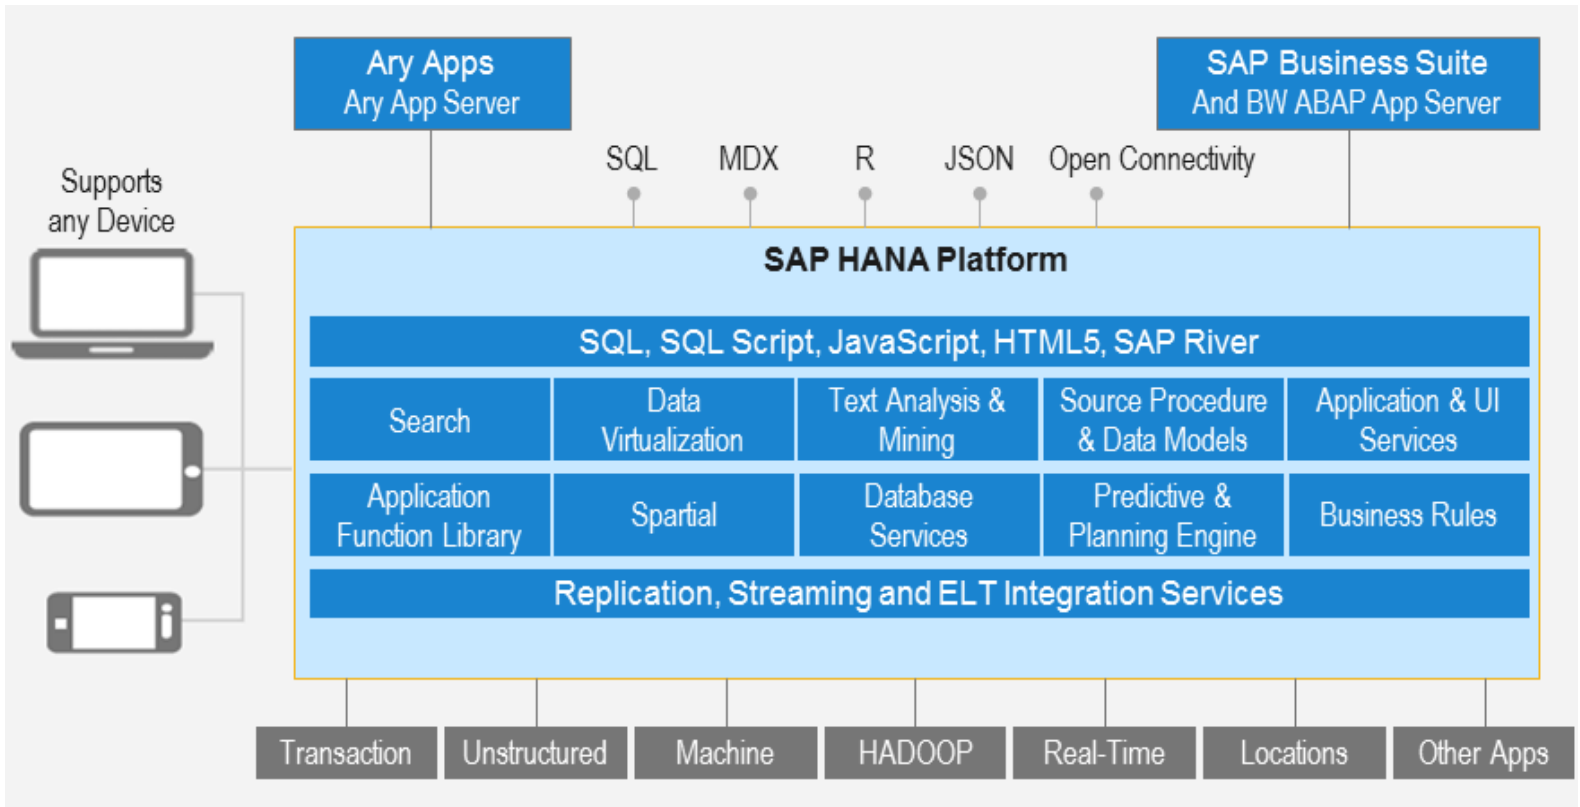

Text Processing With SAP HANA

•Between 70% and 80% of data in an organization is unstructured, and most of this unstructured data is text based?
•The majority of the most powerful and insightful business information is locked up in text. Unlocking it should be taken seriously. SAP HANA has native text-processing capabilities. These include the following:
•Text search: Fuzzy search (Google-like searching) helps users with fault-tolerant searches during data input. It helps to improve data quality by suggesting spellings and codes. It helps to avoid duplication by suggesting similar matches before a user creates another customer account.
•Text analysis: Identifies key entities in text. For example, how many times was company x mentioned this week in tweets that also mentioned words relating to acquisition? Aggregated sentiment analysis of a new product helps you to learn what consumers think, so you can react and make improvements.

- Text search
- Text analysis
- Text mining

•Text mining: Which documents cover similar topics? What is the key subject being discussed in a series of documents or emails?

•SAP HANA text processing handles multiple languages. It can identify the language automatically from the text and apply appropriate linguistic rules.

•SAP HANA text-processing capabilities are already very well exploited in S/4HANA applications, and customers can develop their own applications using the same

Special data Analysis with HANA

•SAP HANA can store and process spatial data. For example, we can identify the exact location of each customer and when the customer is browsing our online catalog we can suggest the nearest pickup location.

•SAP HANA is fully integrated with industry leading partners who specialise in spatial processing. These include Google, ESRI, Pitney Bowes and Tom.

•There are many use cases for spatial data. These include:Live Traffic information - communicate to emergency services driver

•Sport - In-game football analysis - add geo sensor to ball and players and track movements, distances, contacts etc.

•Energy companies - map their pipes, cables, identify closest engineer, or identify nearby assets that could also be cleaned / maintained to save on separate call out

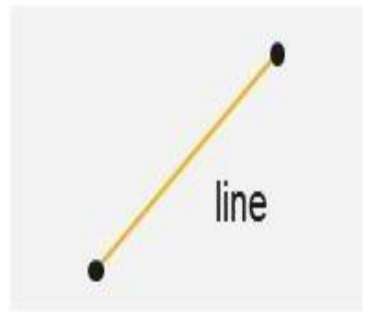

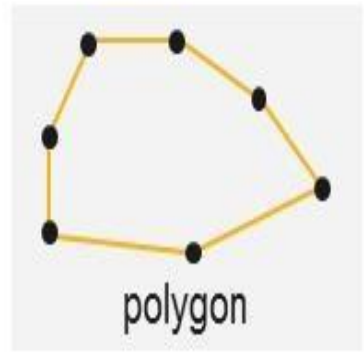

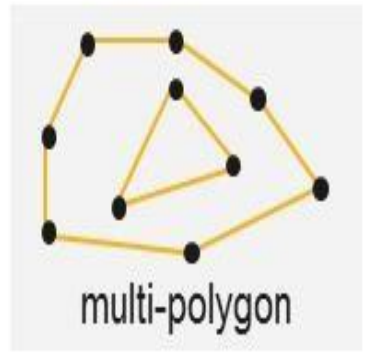

- Geo
- Text analysis
- Text mining

Predictive Analysis with SAP HANA

Predictive Analysis

•A key theme running through SAP S/4HANA is embedded analytics. In many cases, this means adding in predictive capabilities to a transaction flow. Customers can continue to build their own applications that embed predictive capabilities.

•For example, an administrator is providing security clearance to sensitive data for a new employee. However, during the clearance process, SAP HANA identifies and alerts the administrator to a suspicious pattern of system access by the employee that does not fit the profile of this type of worker.

•SAP HANA has an extensive built-in library of powerful predictive algorithms and business functions to suit different analysis scenarios, as shown in the figure.

•In addition to the built-in algorithms, SAP HANA is integrated with the 'R 'public libraries, where
thousands of additional algorithms can be found.

•With SAP HANA's ability to manage huge data volumes, and at speed, real-time predictive analysis is possible and can add huge value to business transactional processing to offer decision support in-line. You can find many examples of embedded predictive analysis in S/4HANA applications.

Reasons To Choose SAP HANA - IOT, Big data, Social media, Real time data

- Smart Meter Data Analytics Improves Grid Reliability
- A Single View of Assets to Optimize Grid Operations
- Predictive Equipment Maintenance to Prevent Blackouts

Reasons To Choose SAP HANA - IOT, Big Data, Social media, Real time data

All Major Industries now collect and process large volume of data as business models have changed .

Example : Utilities – electricity

Past : No account on How much bill and consumption of power from customers . Power was shortage . Now On Oil industry is there and on account go suppress power, alternative entry .

Implementation of smart meter , Smart Grid , Feeder monitoring system and SCDA for substation maintenance etc are generating huge data . This data is being analyzed on real time in dash boards.

To process this ,AP Transco already implementing SAP S/4 HANA and other discussions to follow.

To manage the data real time and huge volumes the solution is Big data on HANA.

SAP S/4 HANA - Simple Logistics Training

By Praveen

www.praveensd.com

What is OLAP & OLTP

SAP ERP / SAP CRM

SAP BDW

OLTP Systems

OLAP Systems

SAP HANA
(Single system for both OLTP and OLAP)

+91 9154555866 - praveen@praveensd.com - www.praveensd.com

How data Gets into OLTP

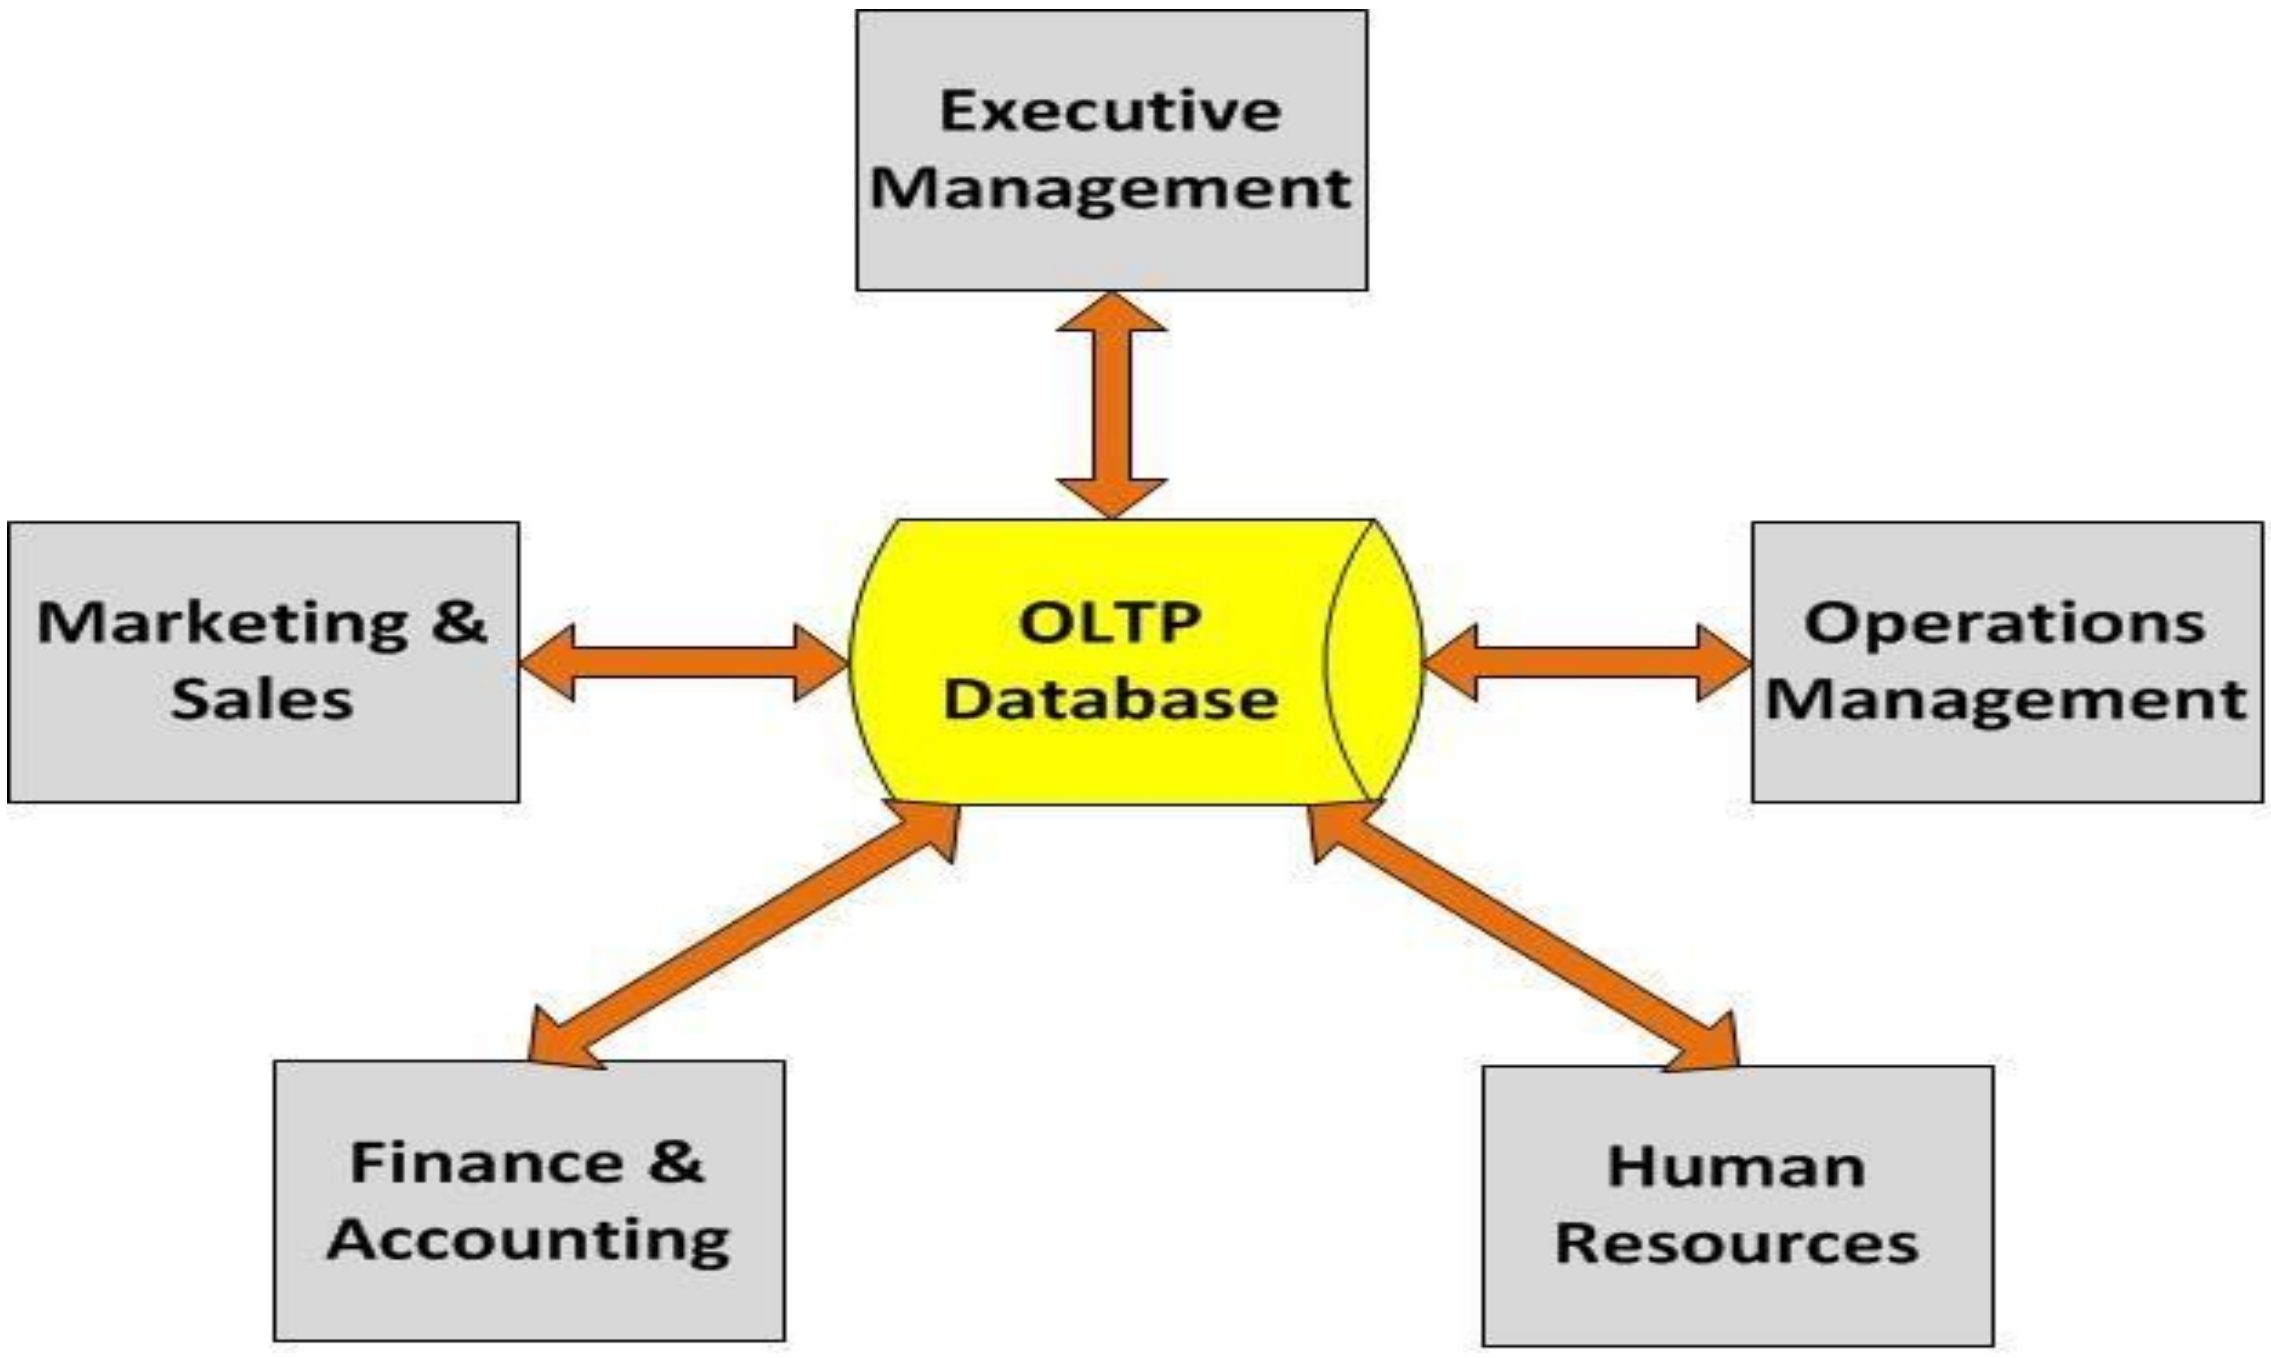

SAP S/4 HANA - Simple Logistics Training

By Praveen

www.praveensd.com

OLAP VS OLTP

Separated transactions + analysis
+ acceleration processes

- Three copies of data in different data models
- Inherent data latency
- Poor innovation leading to wastage

VS

One in-memory atomic copy of data for
transactions + analysis

- Elimination of unnecessary complexity and latency
- Less hardware to manage
- Acceleration through innovation and simplification

OLTP & OLAP optimization

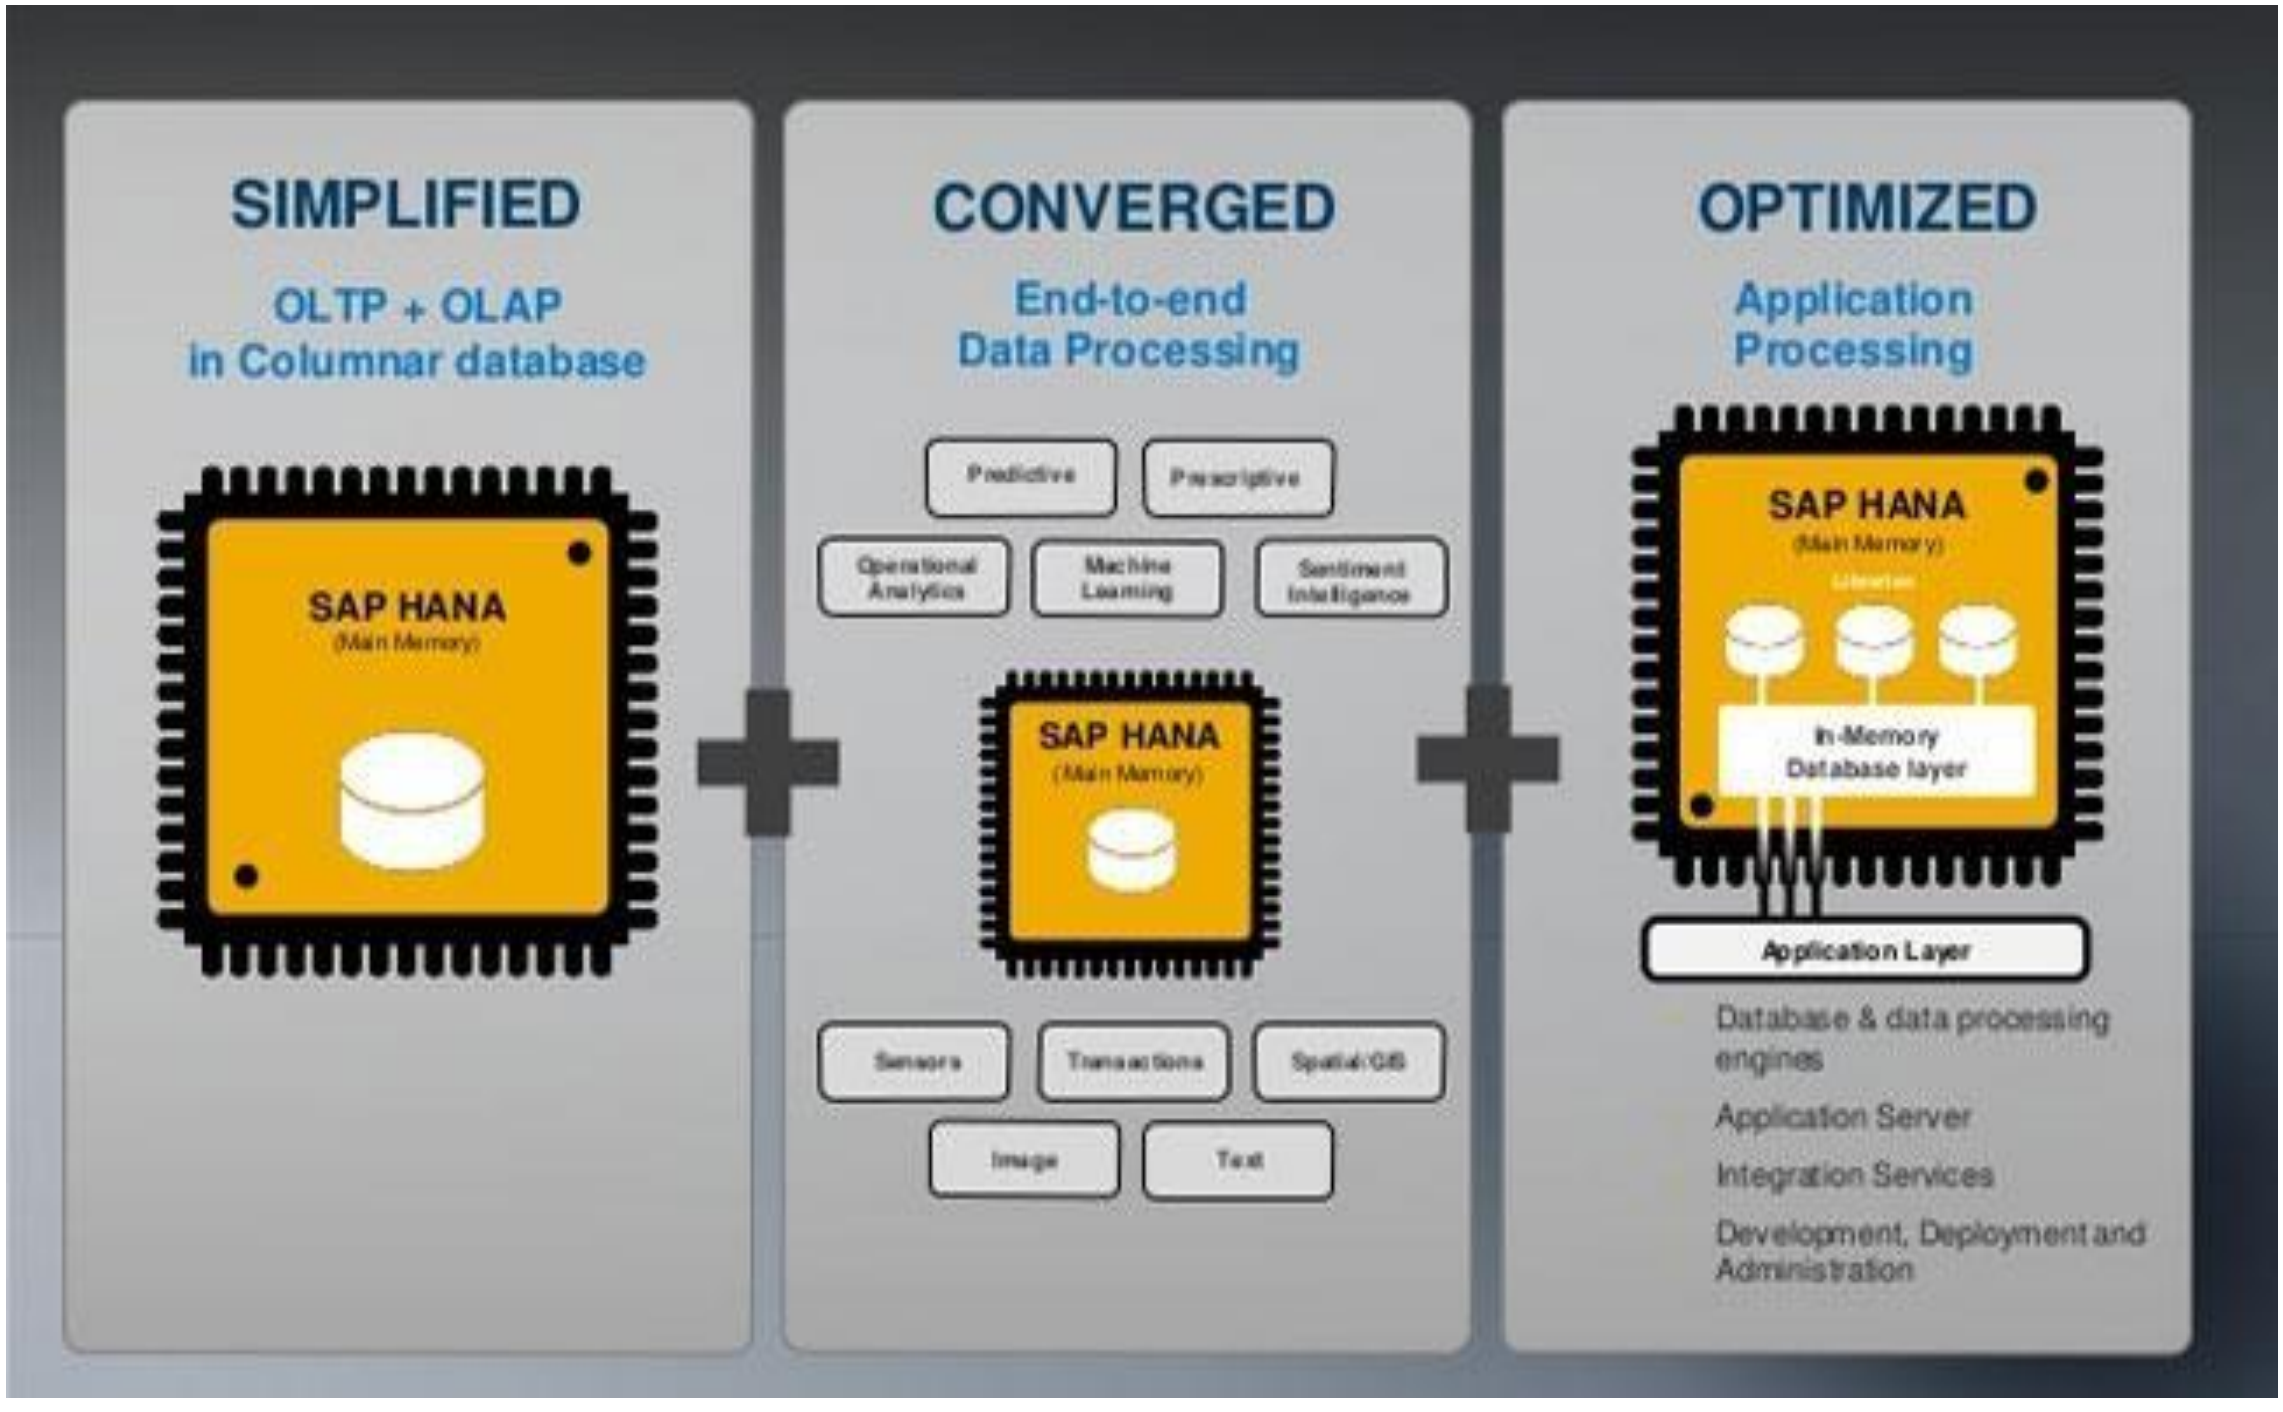

Column based Tables Vs Row based

- Row Storage - It stores table records in a sequence of rows.
- Column Storage - It stores table records in a sequence of columns i.e. the entries of a column is stored in contiguous memory locations.

  Traditional databases store data simply in rows. The HANA in-memory database stores data in both rows and columns. It is this combination of both storage approaches that produces the speed, flexibility and performance of the HANA database.
- Minimum Expected Compression of database is 5 time . Our current 1 TB OLTP database may become 200 GB

- Monthly data growth may be 20gb instead of 100 GB in the traditional database

Feature of In -memory database

- Mostly Colum based tables
- Should be able to support both row based and Colum based tables
- Low data footprint
- Realtime data reporting

What is In-memory data base?

- An in-memory database system is a database management system that stores data entirely in main memory. This contrasts to traditional (on-disk) database systems, which are designed for data storage on persistent media.
- Because working with data in memory is much faster than writing to and reading from a file system, IMDSs can perform applications' data management functions an order of magnitude faster.
- Because their design is typically simpler than that of on-disk databases, IMDSs can also impose significantly lower memory and CPU requirements.

Column based Tables Vs Row based

Table

	Country	Product	Sales
Row 1	India	Chocolate	1000
Row 2	India	Ice-cream	2000
Row 3	Germany	Chocolate	4000
Row 4	US	Noodle	500

Row Store

Row 1	India
	Chocolate
	1000
Row 2	India
	Ice-cream
	2000
Row 3	Germany
	Chocolate
	4000
Row 4	US
	Noodle
	500

Column Store

Country	India
	India
	Germany
	US
Product	Chocolate
	Ice-cream
	Chocolate
	Noodle
Sales	1000
	2000
	4000
	500

SAP S/4 HANA - Simple Logistics Training

By Praveen

www.praveensd.com

Column based Tables Vs Row based

Table – SALES

	Date	Country	Product	Sales
Row 1	2013-01-01	India	Chocolate	1000
Row 2	2013-01-10	India	Ice-cream	2000
Row 3	2013-02-20	Germany	Chocolate	4000
Row 4	2013-03-01	US	Noodle	500

Column Operation: SELECT SUM(SALES) FROM SALES WHERE DATE > 2012-01-01

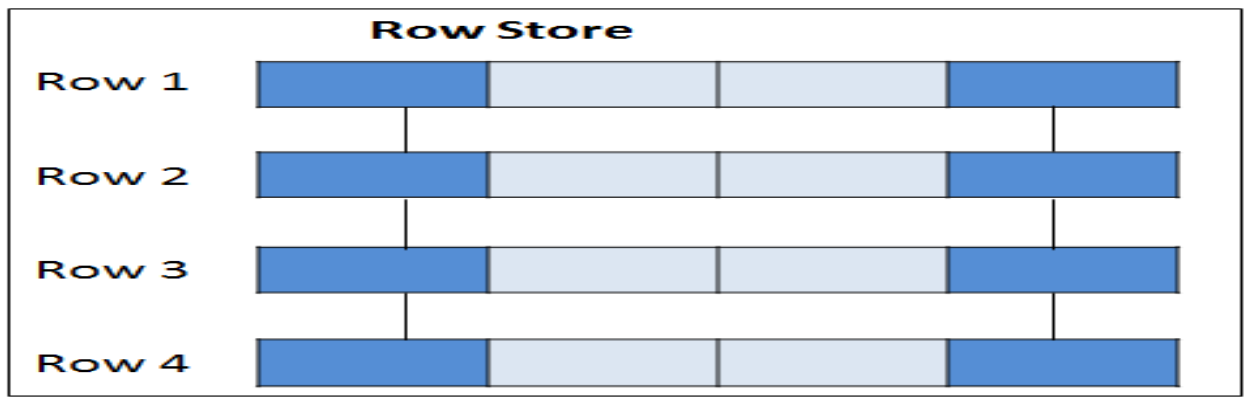

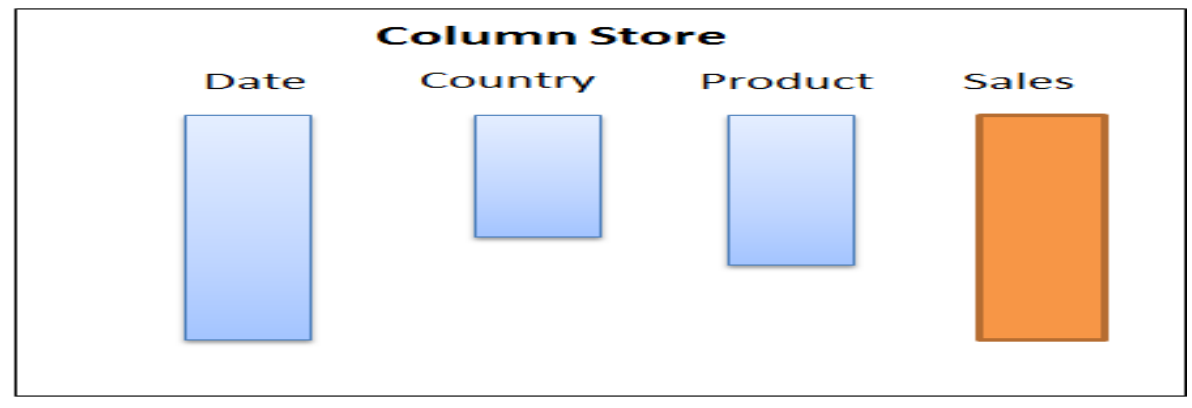

Row Operation: SELECT * FROM SALES WHERE COUNTRY = 'INDIA'

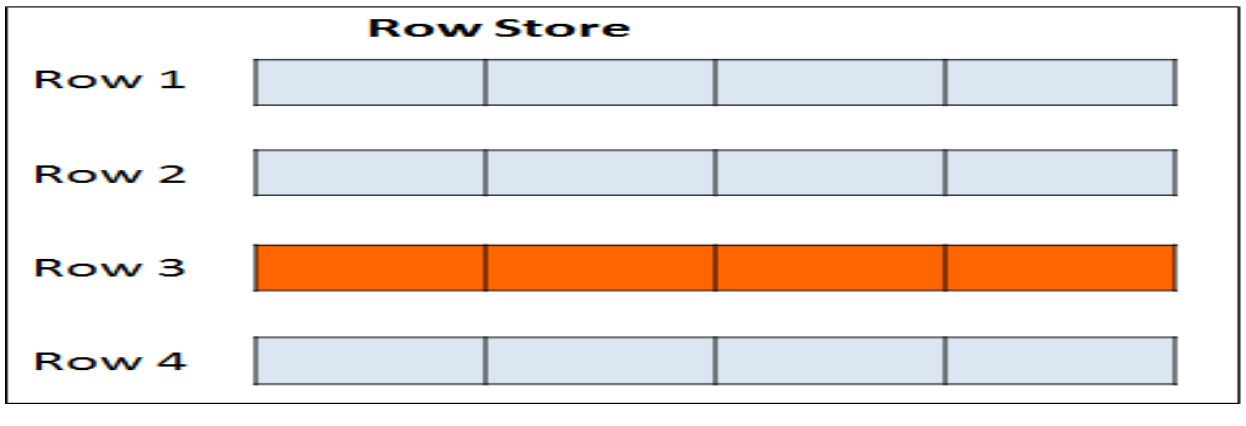

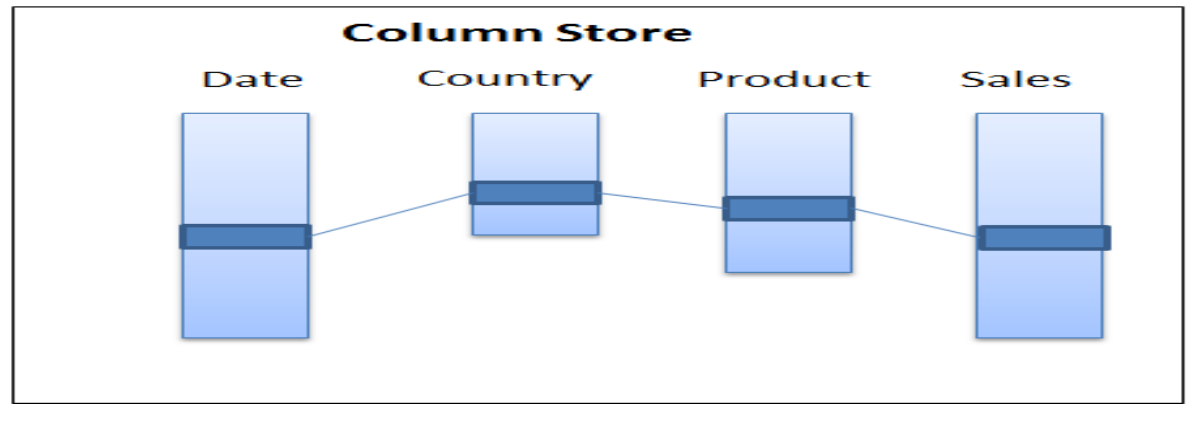

SAP S/4 HANA - Simple Logistics Training

By Praveen

www.praveensd.com

Which is suitable for OLTP and OLAP

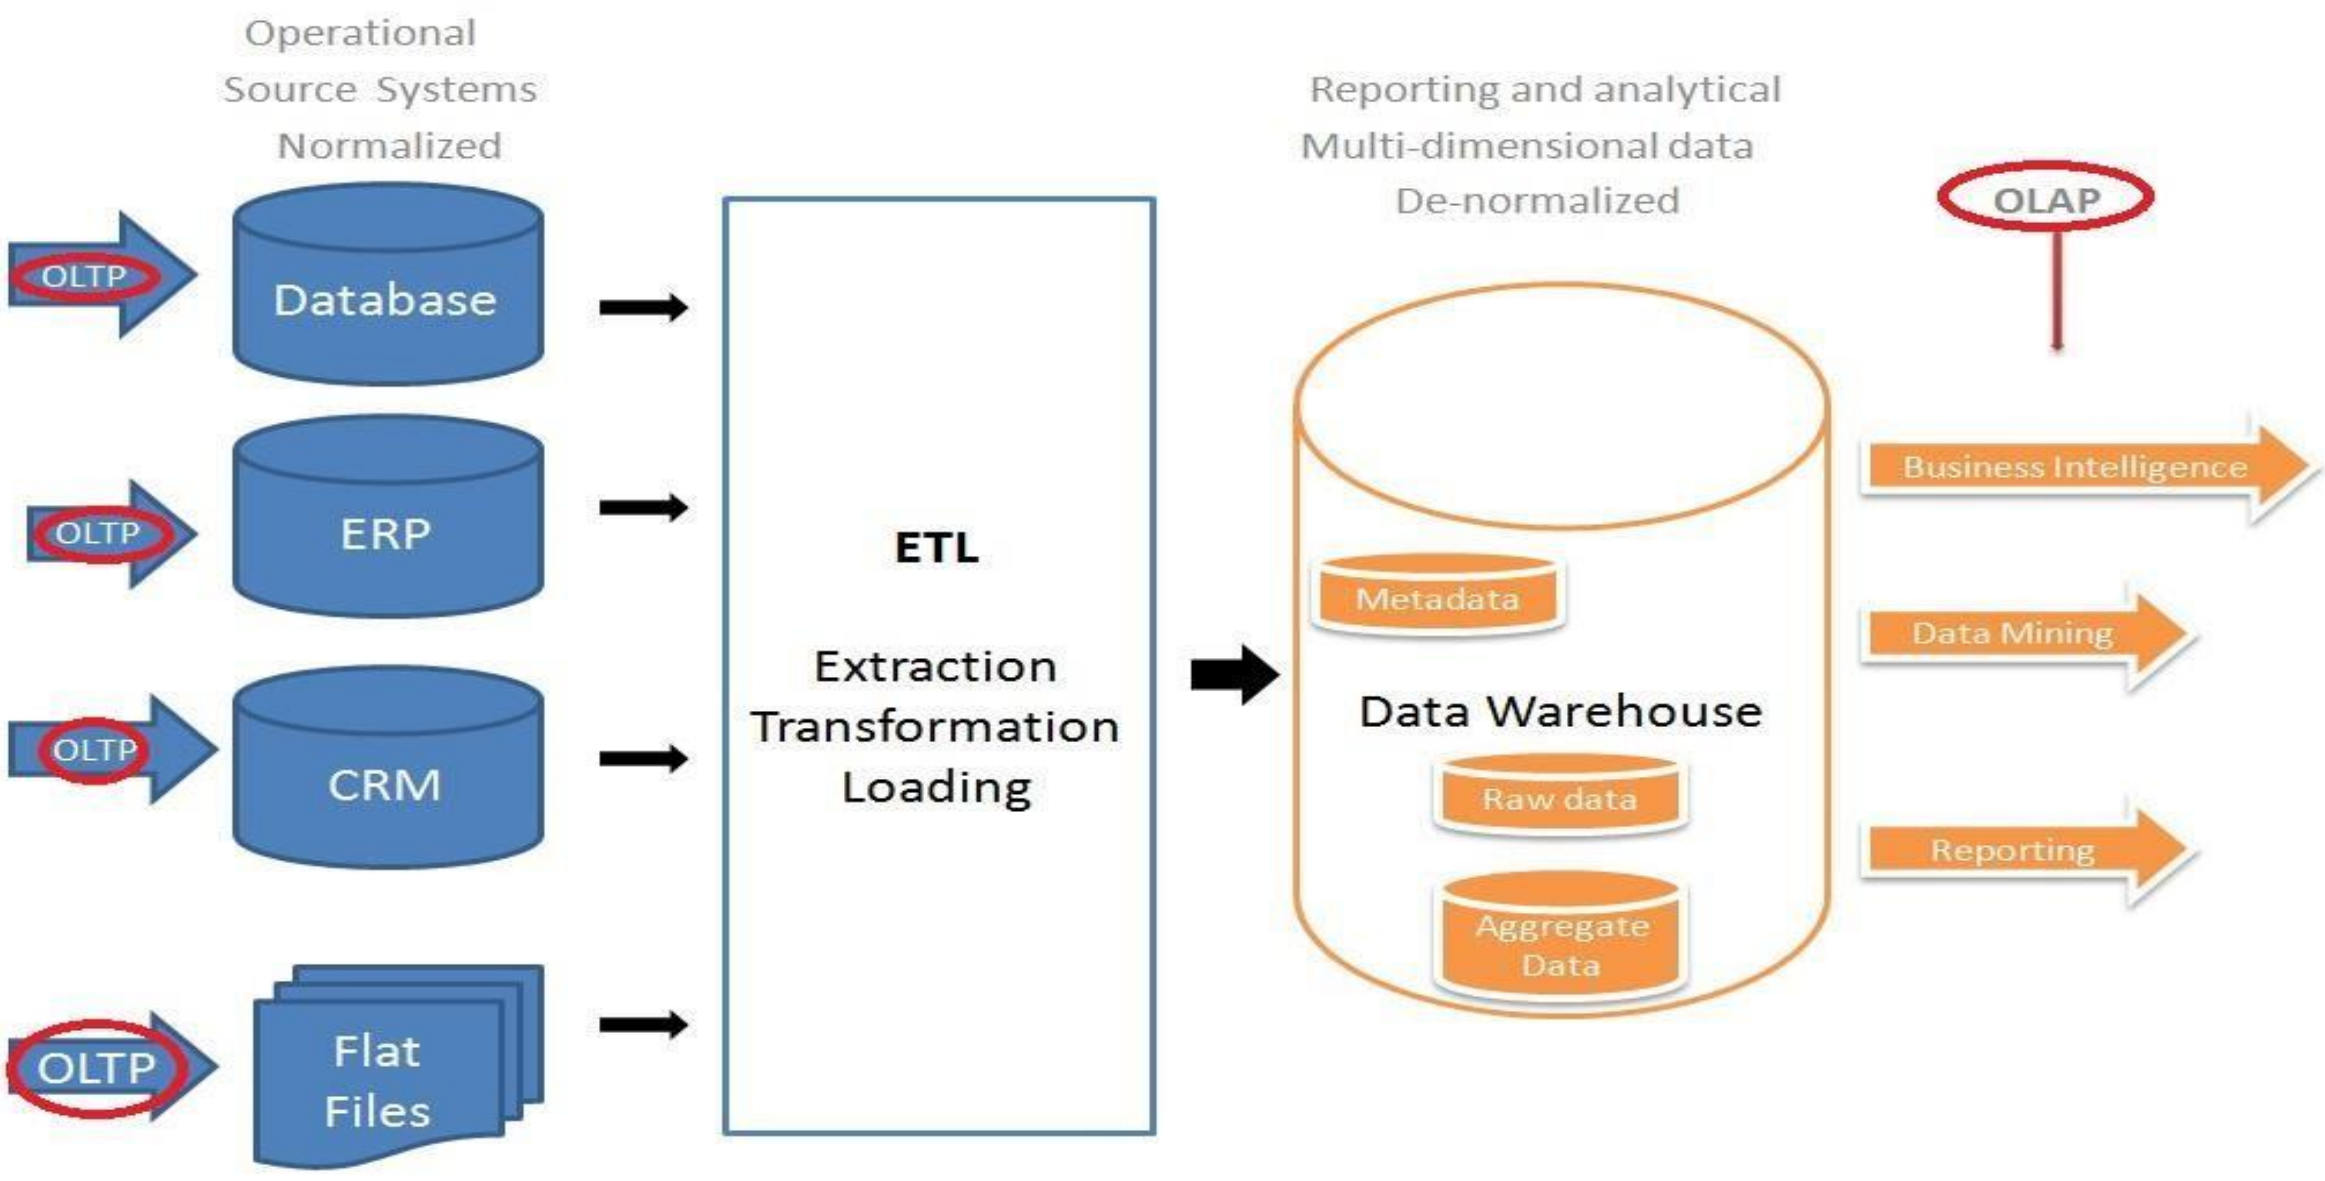

OLTP (Online Transaction Processing) versus
OLAP (Online Analytical Processing)

Column Vs Row

Columnar vs Row-based

Here is an example of translating a logical table schema. First the example table:

A	B	C
A1	B1	C1
A2	B2	C2
A3	B3	C3

In a row-based layout each row follows the next:

A1	B1	C1	A2	B2	C2	A3	B3	C3

While for a column-oriented layout it stores one column after the next:

A1	A2	A3	B1	B2	B3	C1	C2	C3

Row and Column with Delta Merge

SAP HANA

- Three-level column-oriented unified table structure

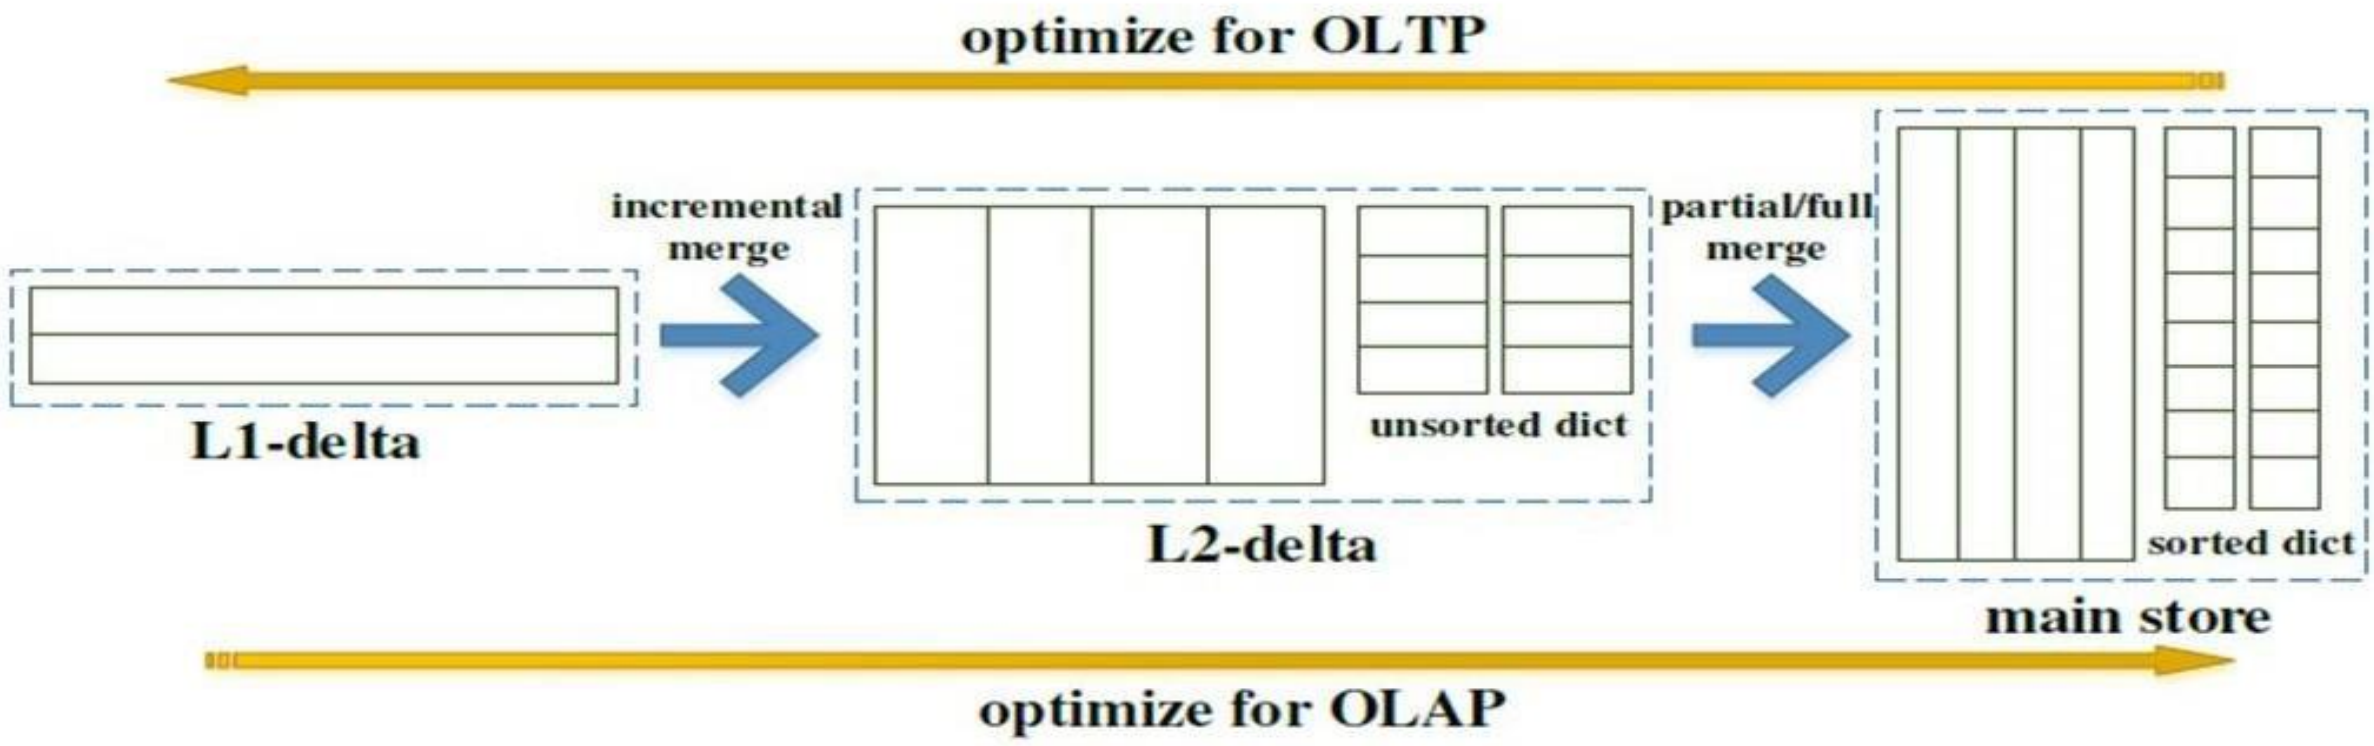

Row for OLTP and Colum for OLAP

Row/Column - OLTP/OLAP

- Row stores are good fit for OLTP
- Reading small portions of a table, but often many of the columns
- Frequent changes to data
- Small (<2TB) amount of data (typically working set must fit in ram)
- “Nested loops” joins are good fit for OLTP

Row for OLTP and Column for OLAP

Row/Column - OLTP/OLAP

- Column stores are good fit for OLAP
- Read large portions of a table in terms of rows, but often a small number of columns
- Batch loading / updates
- Big data (50TB-100TB per machine):
- Compression capabilities comes in handy
- Machine generated data is well suited

Column based Advantages

Column-based tables have advantages in the following circumstances:

1.Calculations are typically executed on single or a few columns only.

2.The table is searched based on values of a few columns.

3.The table has a large number of columns.

4.The table has a large number of rows and columnar operations are required (aggregate, scan etc.)

5.High compression rates can be achieved because the majority of the columns contain only few distinct values (compared to number of rows).

Row based advantages

Row based tables have advantages in the following circumstances:

1. The application needs to only process a single record at one time (many selects and/or updates of single records).

2.The application typically needs to access a complete record (or row).

3.The columns contain mainly distinct values so that the compression rate would be low.

4.Neither aggregations nor fast searching are required.

5.The table has a small number of rows (e. g. configuration tables).
To enable fast on-the-fly aggregations, ad-hoc reporting, and to benefit from compression mechanisms it is recommended that transaction data is stored in a column-based table.

The SAP HANA data-base allows joining row-based tables with column-based tables. However, it is more efficient to join tables that are located in the same row or column store. For example, master data that is frequently joined with transaction data should also be stored in column-based tables.

HANA Row Vs Column

SAP HANA Column Storage Vs. Row-Based Storage

Customer	Country	Product	Amount
100	DE	1	100
100	DE	1	110
200	US	1	120
300	US	2	130

ROW-BASED Storage

Tuple 1
Tuple 2
Tuple 3
Tuple 4

COLUMN-BASED Storage

Column 1 | Column 2 | Column 3 | Column 4

Row based storage in memory would look like this:
100-DE-1-100 100-DE-1-110 200-US-1-120 300-US-2-130

Column based storage in memory would look like this:
100-100-200-300 DE-DE-US-US 1-1-1-2 100-110-120-130

- Storing data in columns is not a new technology, but it has not been leveraged to its full potential, yet
- The columnar storage is read-optimized, that is, the read operations can be processed very fast. However, it's not write-optimized, as new insert might lead to moving a lot of data to create a place for new data
- HANA handles this well with delta merge. The columnar storage performs very well while reading and the write operations are taken care of by the In-Memory Computing Engine (IMCE) in some other ways

Limitation of standard data base and In-memory database

- Potential loss of data and limit on database size. When this is fine, you should certainly use RDM's in memory solution.

- three schemes for dealing with the durability issue. Each has advantages and challenges:

  1)Write data to disk as well as putting it in memory. The challenge here is whether you can write to the disk fast enough to keep up with the data that is being loaded into memory.

  2)Put the data in memory on two different machines. The risk here is if both machines go down, you lose the data.

  3)Use a combination of #1 and #2 above. Putting data in-memory on two machines provides some level of protection that allows time for a background process or asynchronous process to write data out to disk. In this case you need to understand what scheme a vendor is using and whether it meets your service level agreements.

- memory is much more expensive than spinning disk.

SSD data Store

- Like a **memory** stick, there are no moving parts to an **SSD**. Rather, information is stored in microchips.

Conversely, a hard disk drive uses a **mechanical** arm with a **read/write** head to move around and **read** information from the right location on a storage platter. This difference is what makes **SSD** so much faster.

Hard Disk or SSD or DRAM

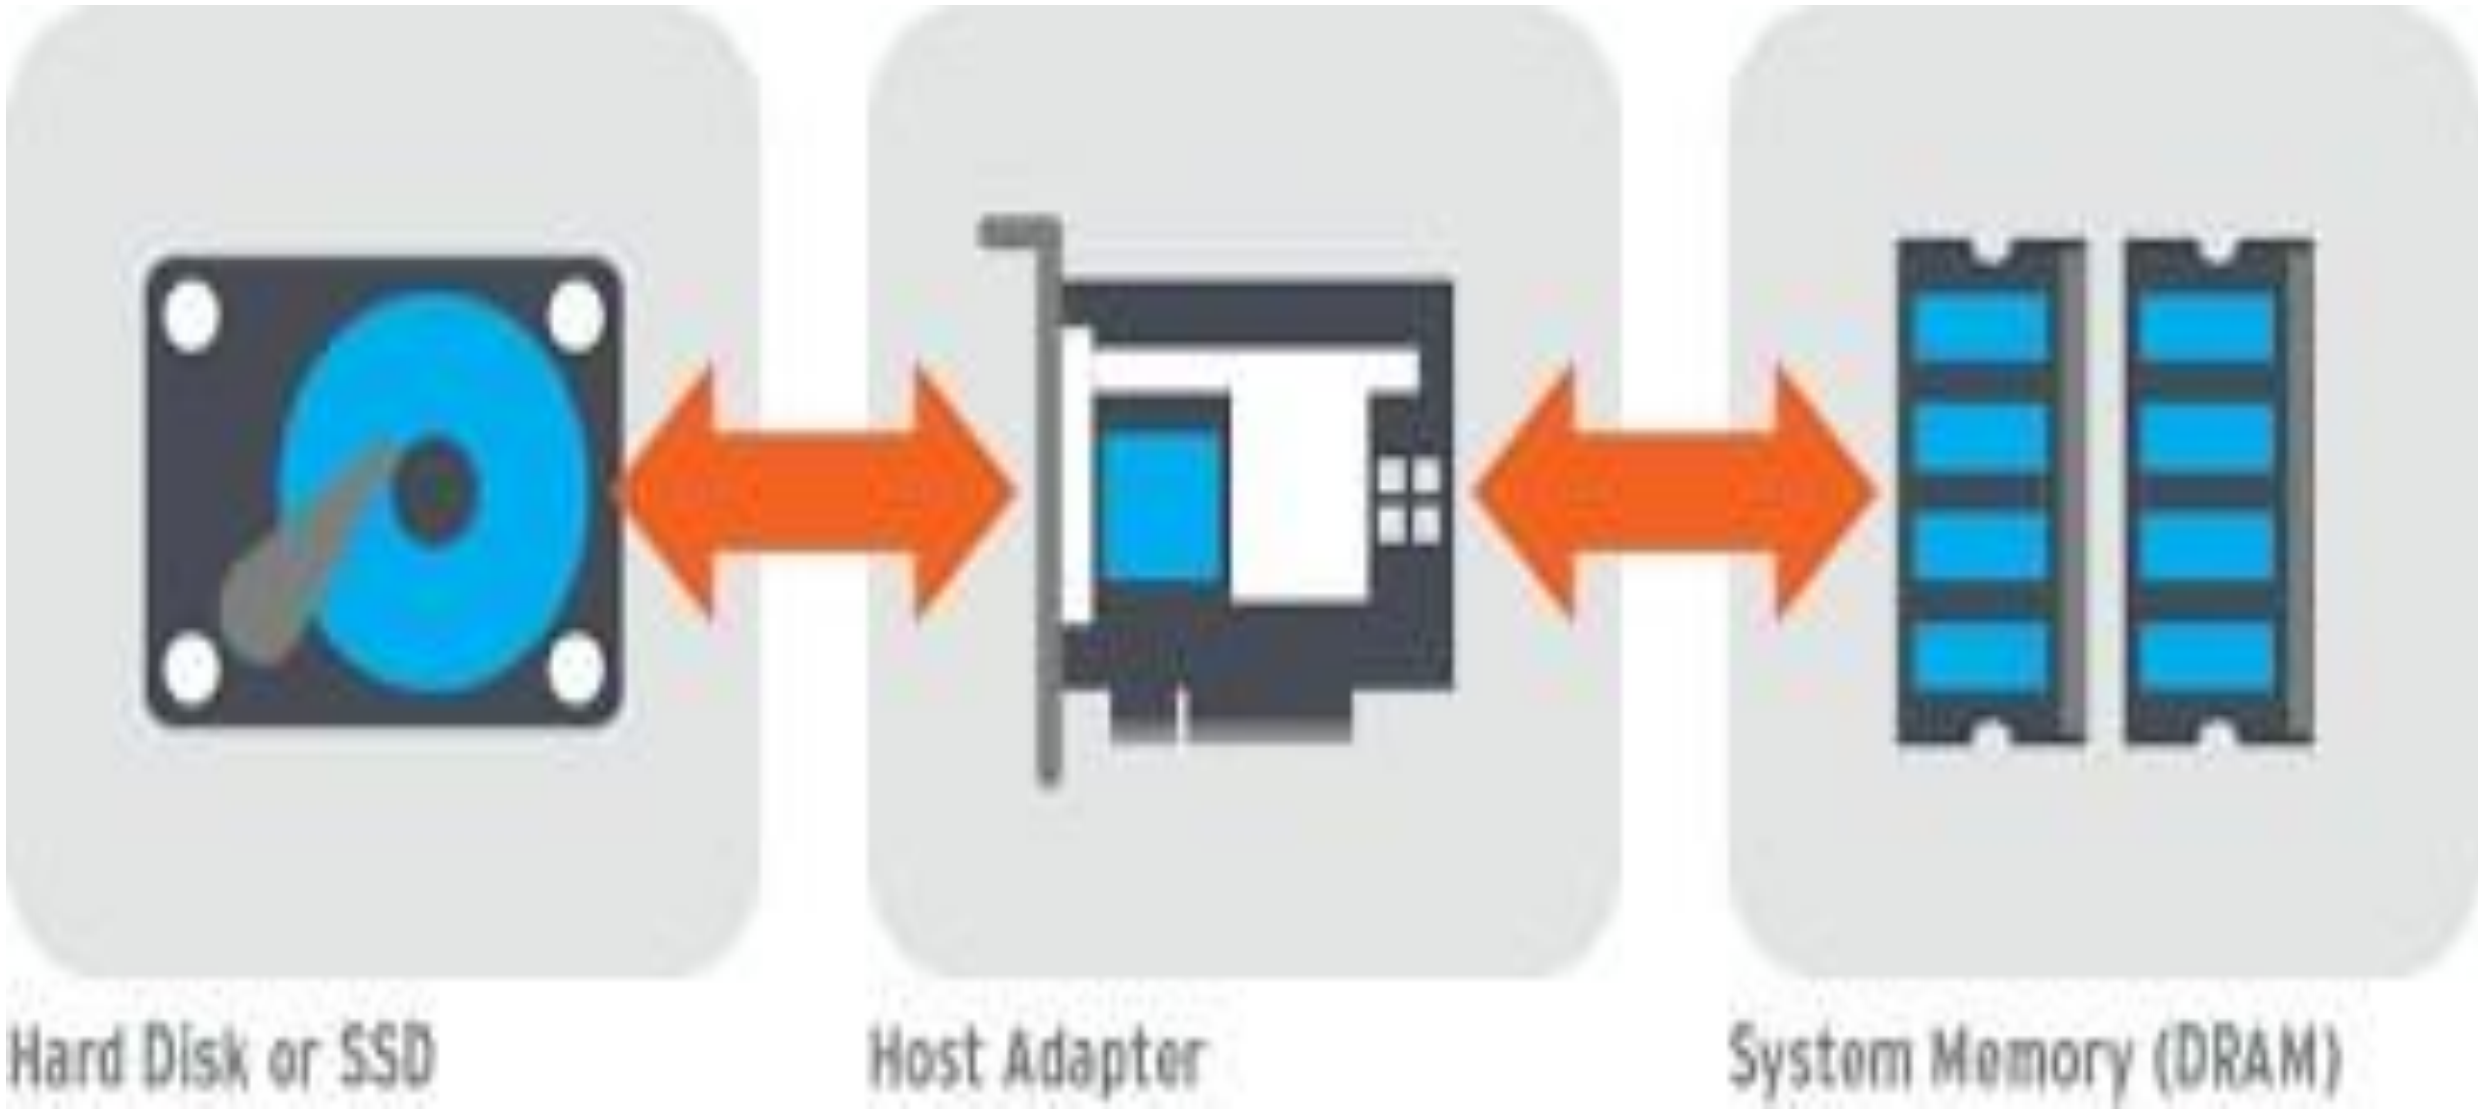

RAM Disk VsSSD

- RAM Disk vs SSD – Ten Times Faster Read and Write Speed via RAM Virtual Disk
- RAM Disk is a program that takes a portion of your system memory and uses it as a disk drive.
- This is same as in memory data base of Hana

In the age of SSD's why are in-memory databases like SAP HANA needed

- In-memory processing in SAP HANA is not only about storing data in RAM, but as well about cache-aware data structures and algorithms.

- It is not only about how to retrieve data quickly from RAM, but as well how to retrieve and process data so that CPUs' caches are best utilized.

Advantages and Disadvantages of Standard Dbase and how did ERP survive till now using standard Database

- The Process used is Data redundancy
- Same data is kept in many places
- Aggregation tables are built
- Business warehouse is built for reporting
- Still the reports are old never real time
- Whereas in HANA the reports near real time

Single CPU Processing

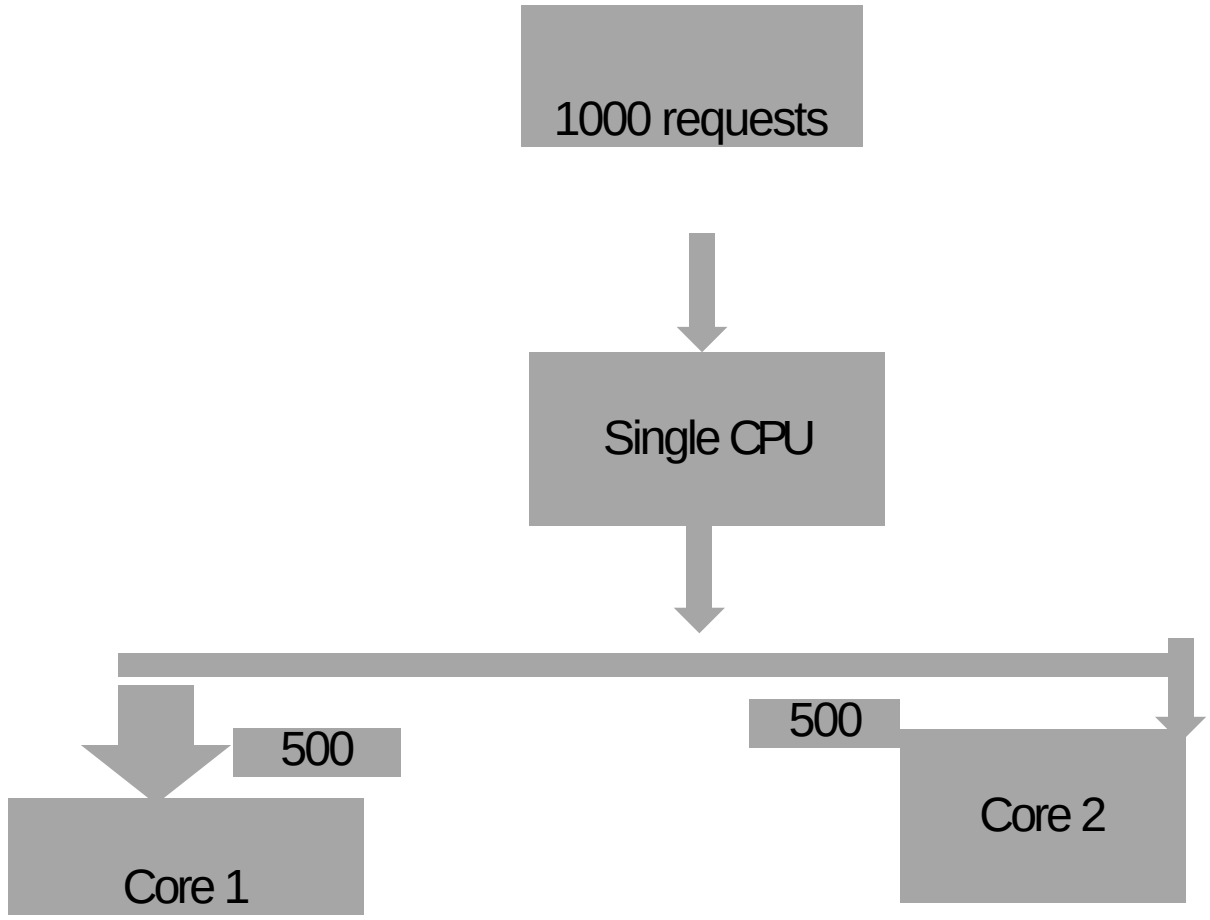

Multi CPU Processing Multi Core

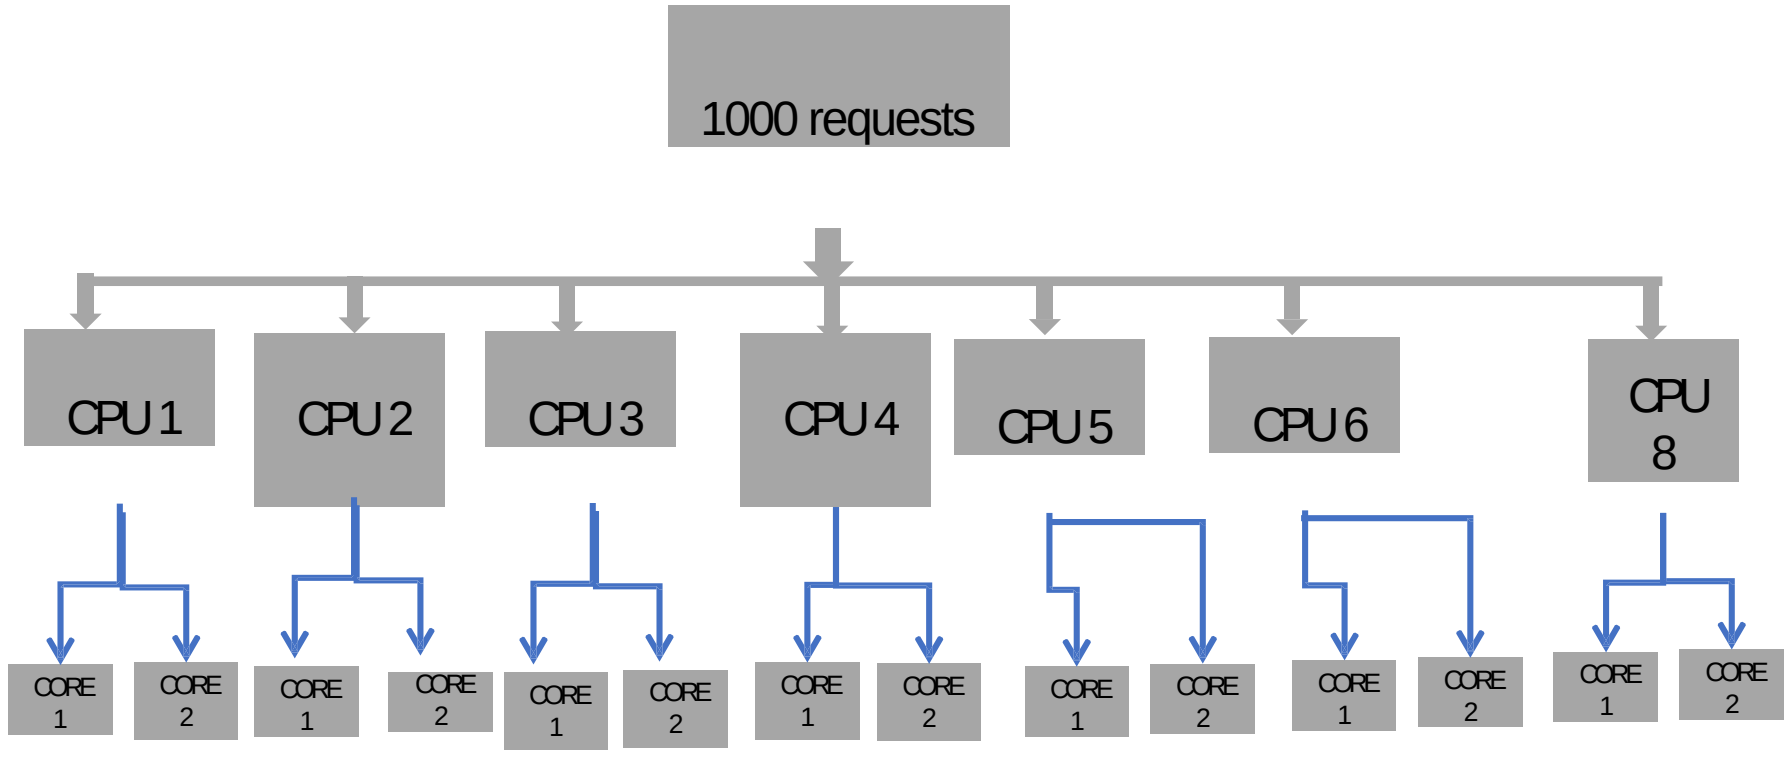

SAP S/4 HANA - Simple Logistics Training

By Praveen

www.praveensd.com

Hana Architecture

Architecture

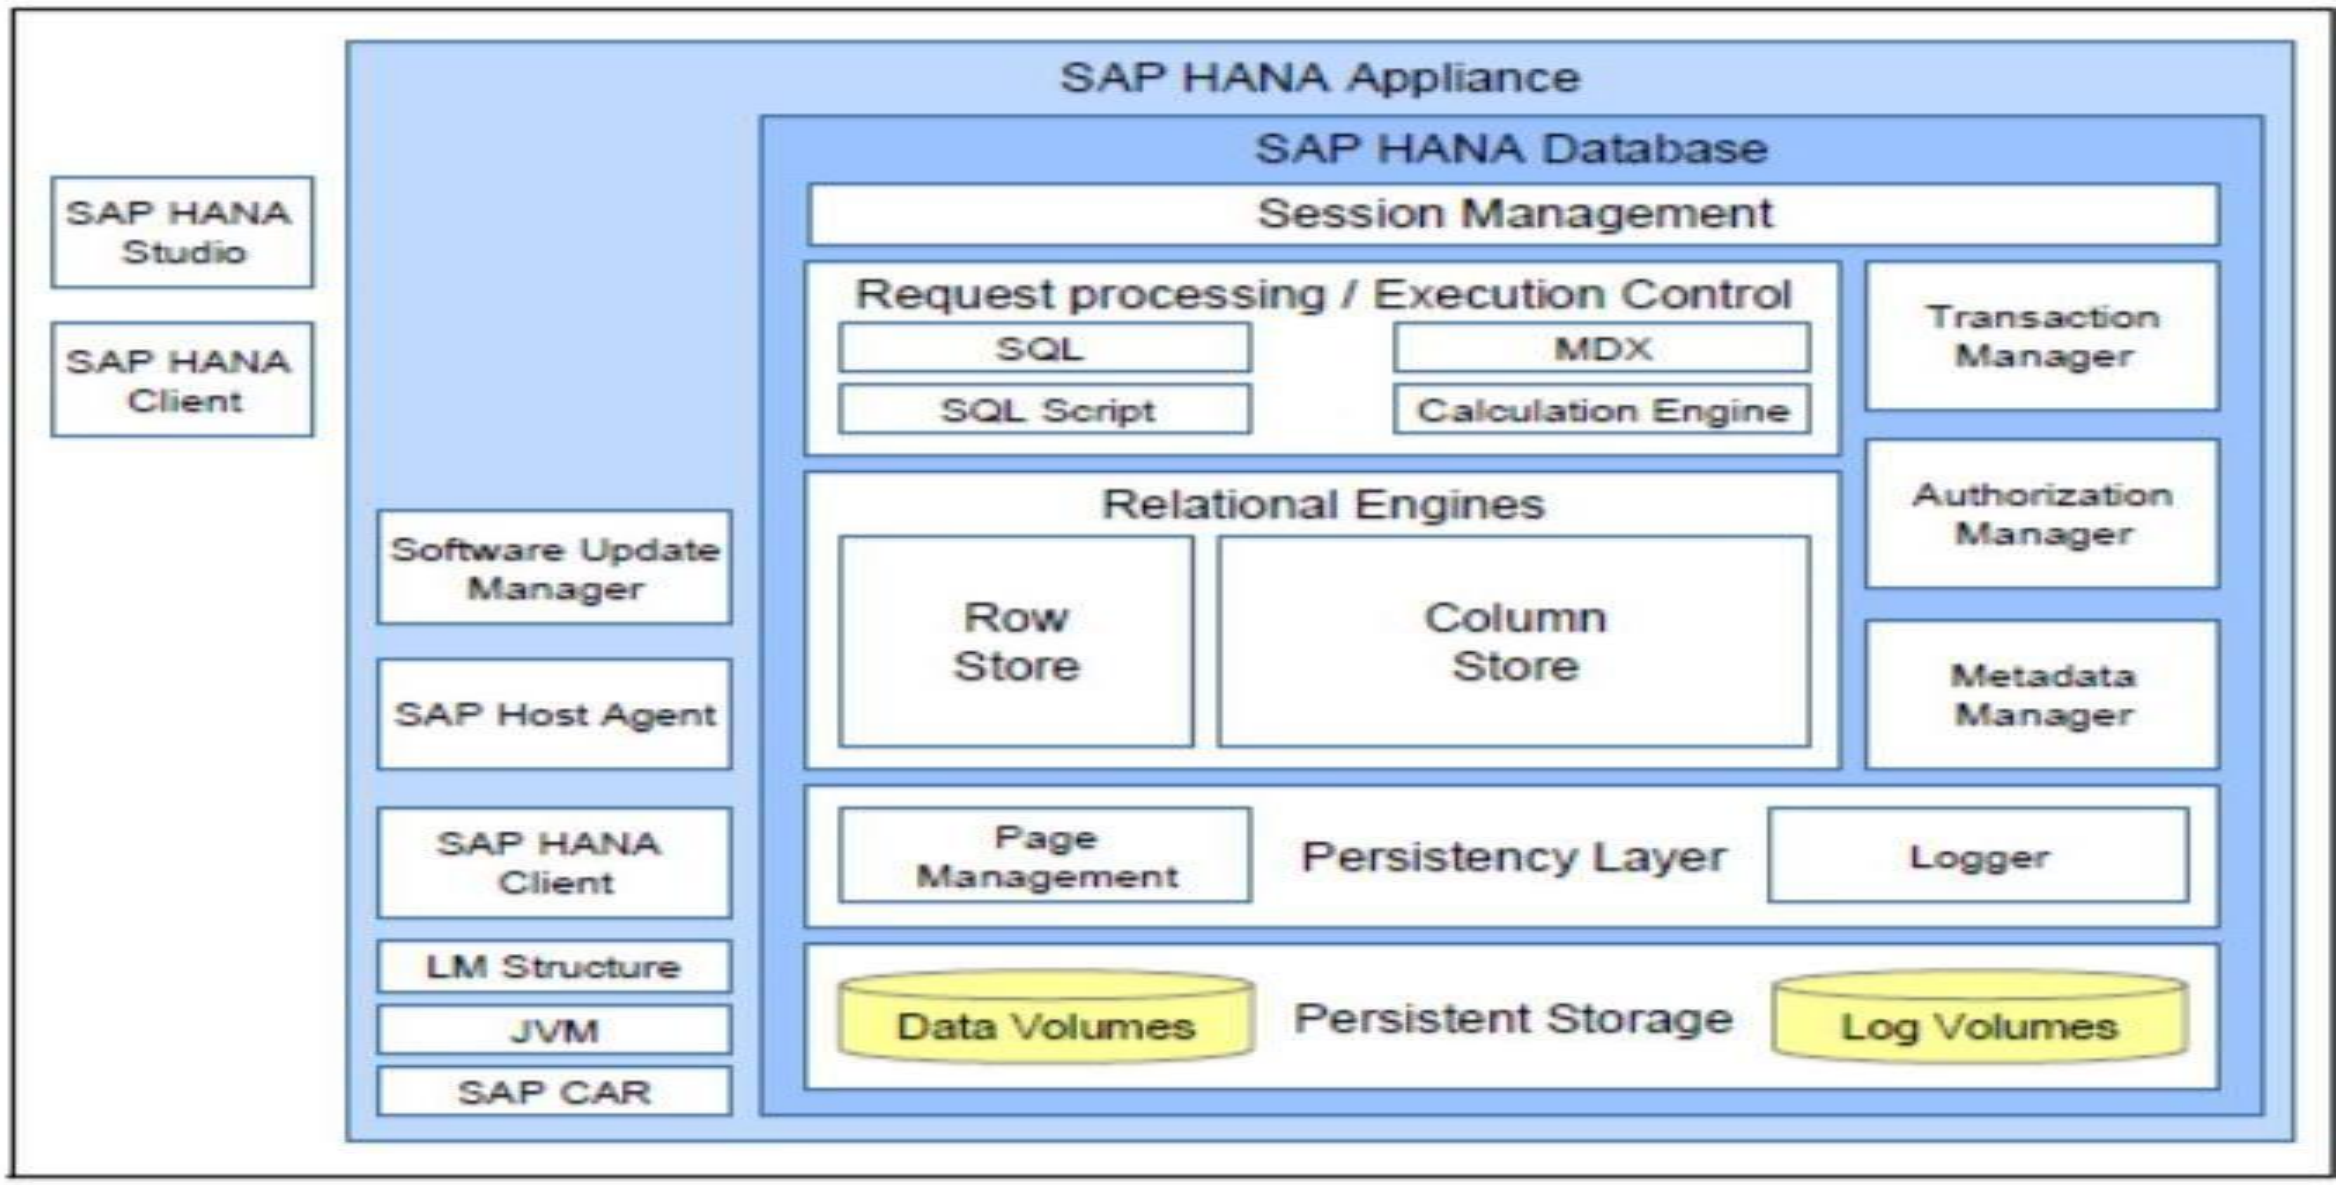

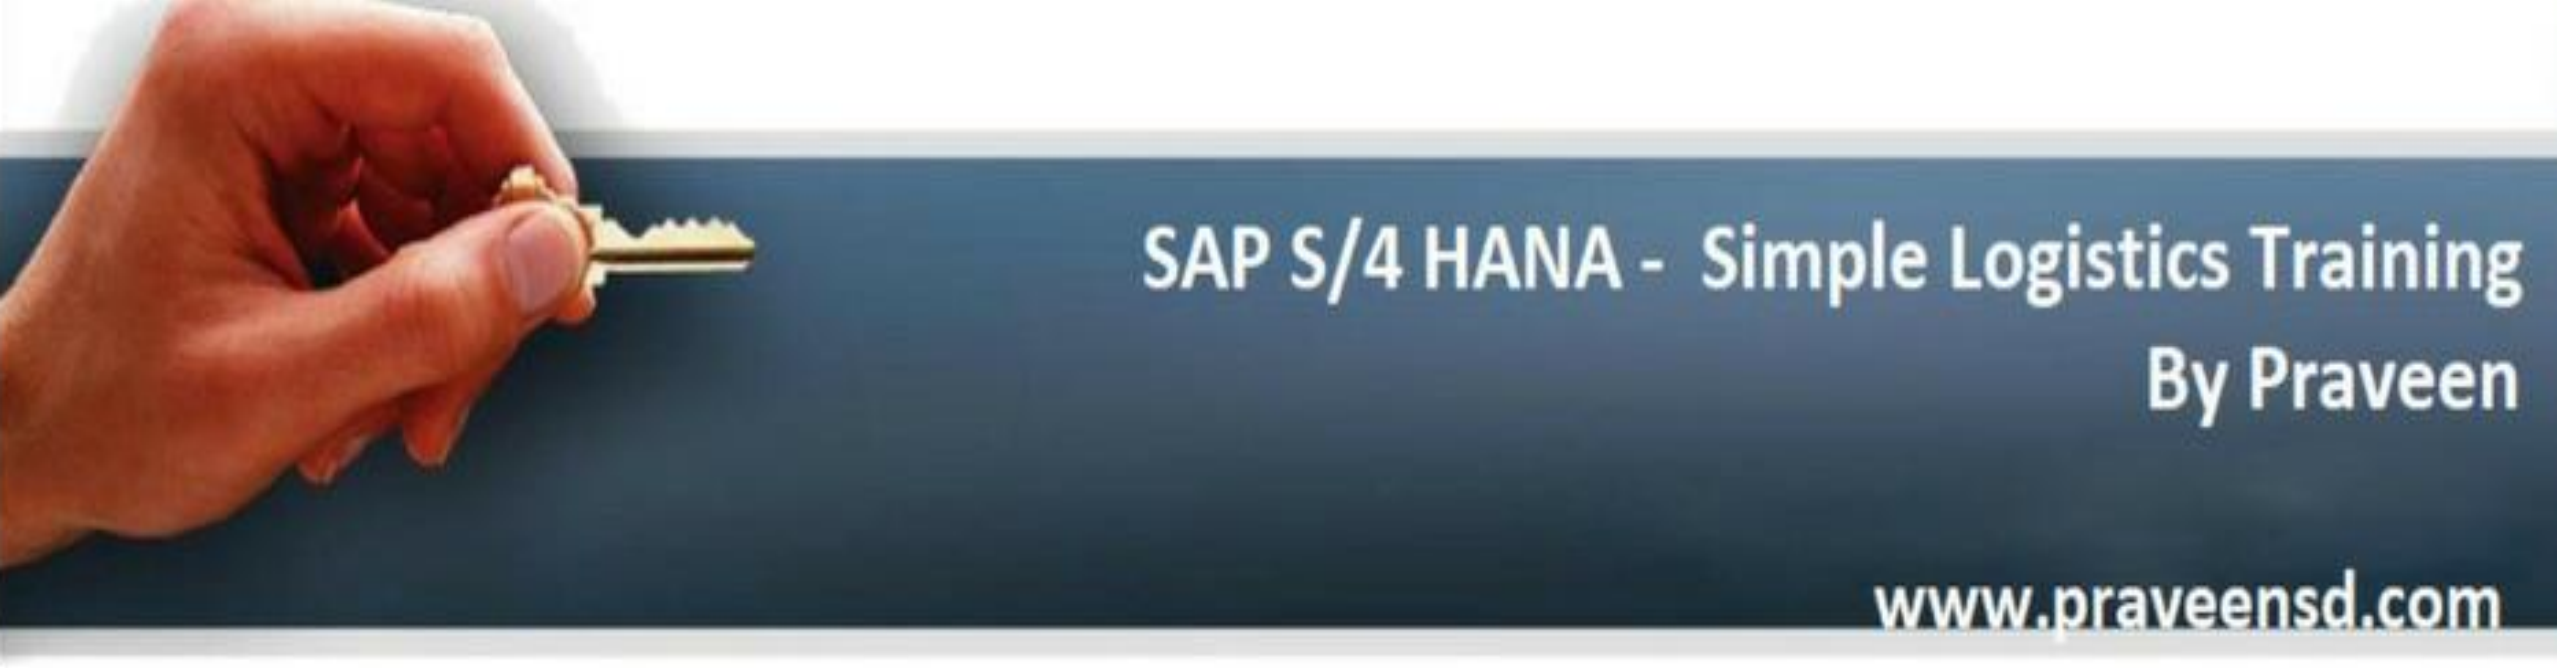

HANA Integration with IOT data or Big data in Utilities

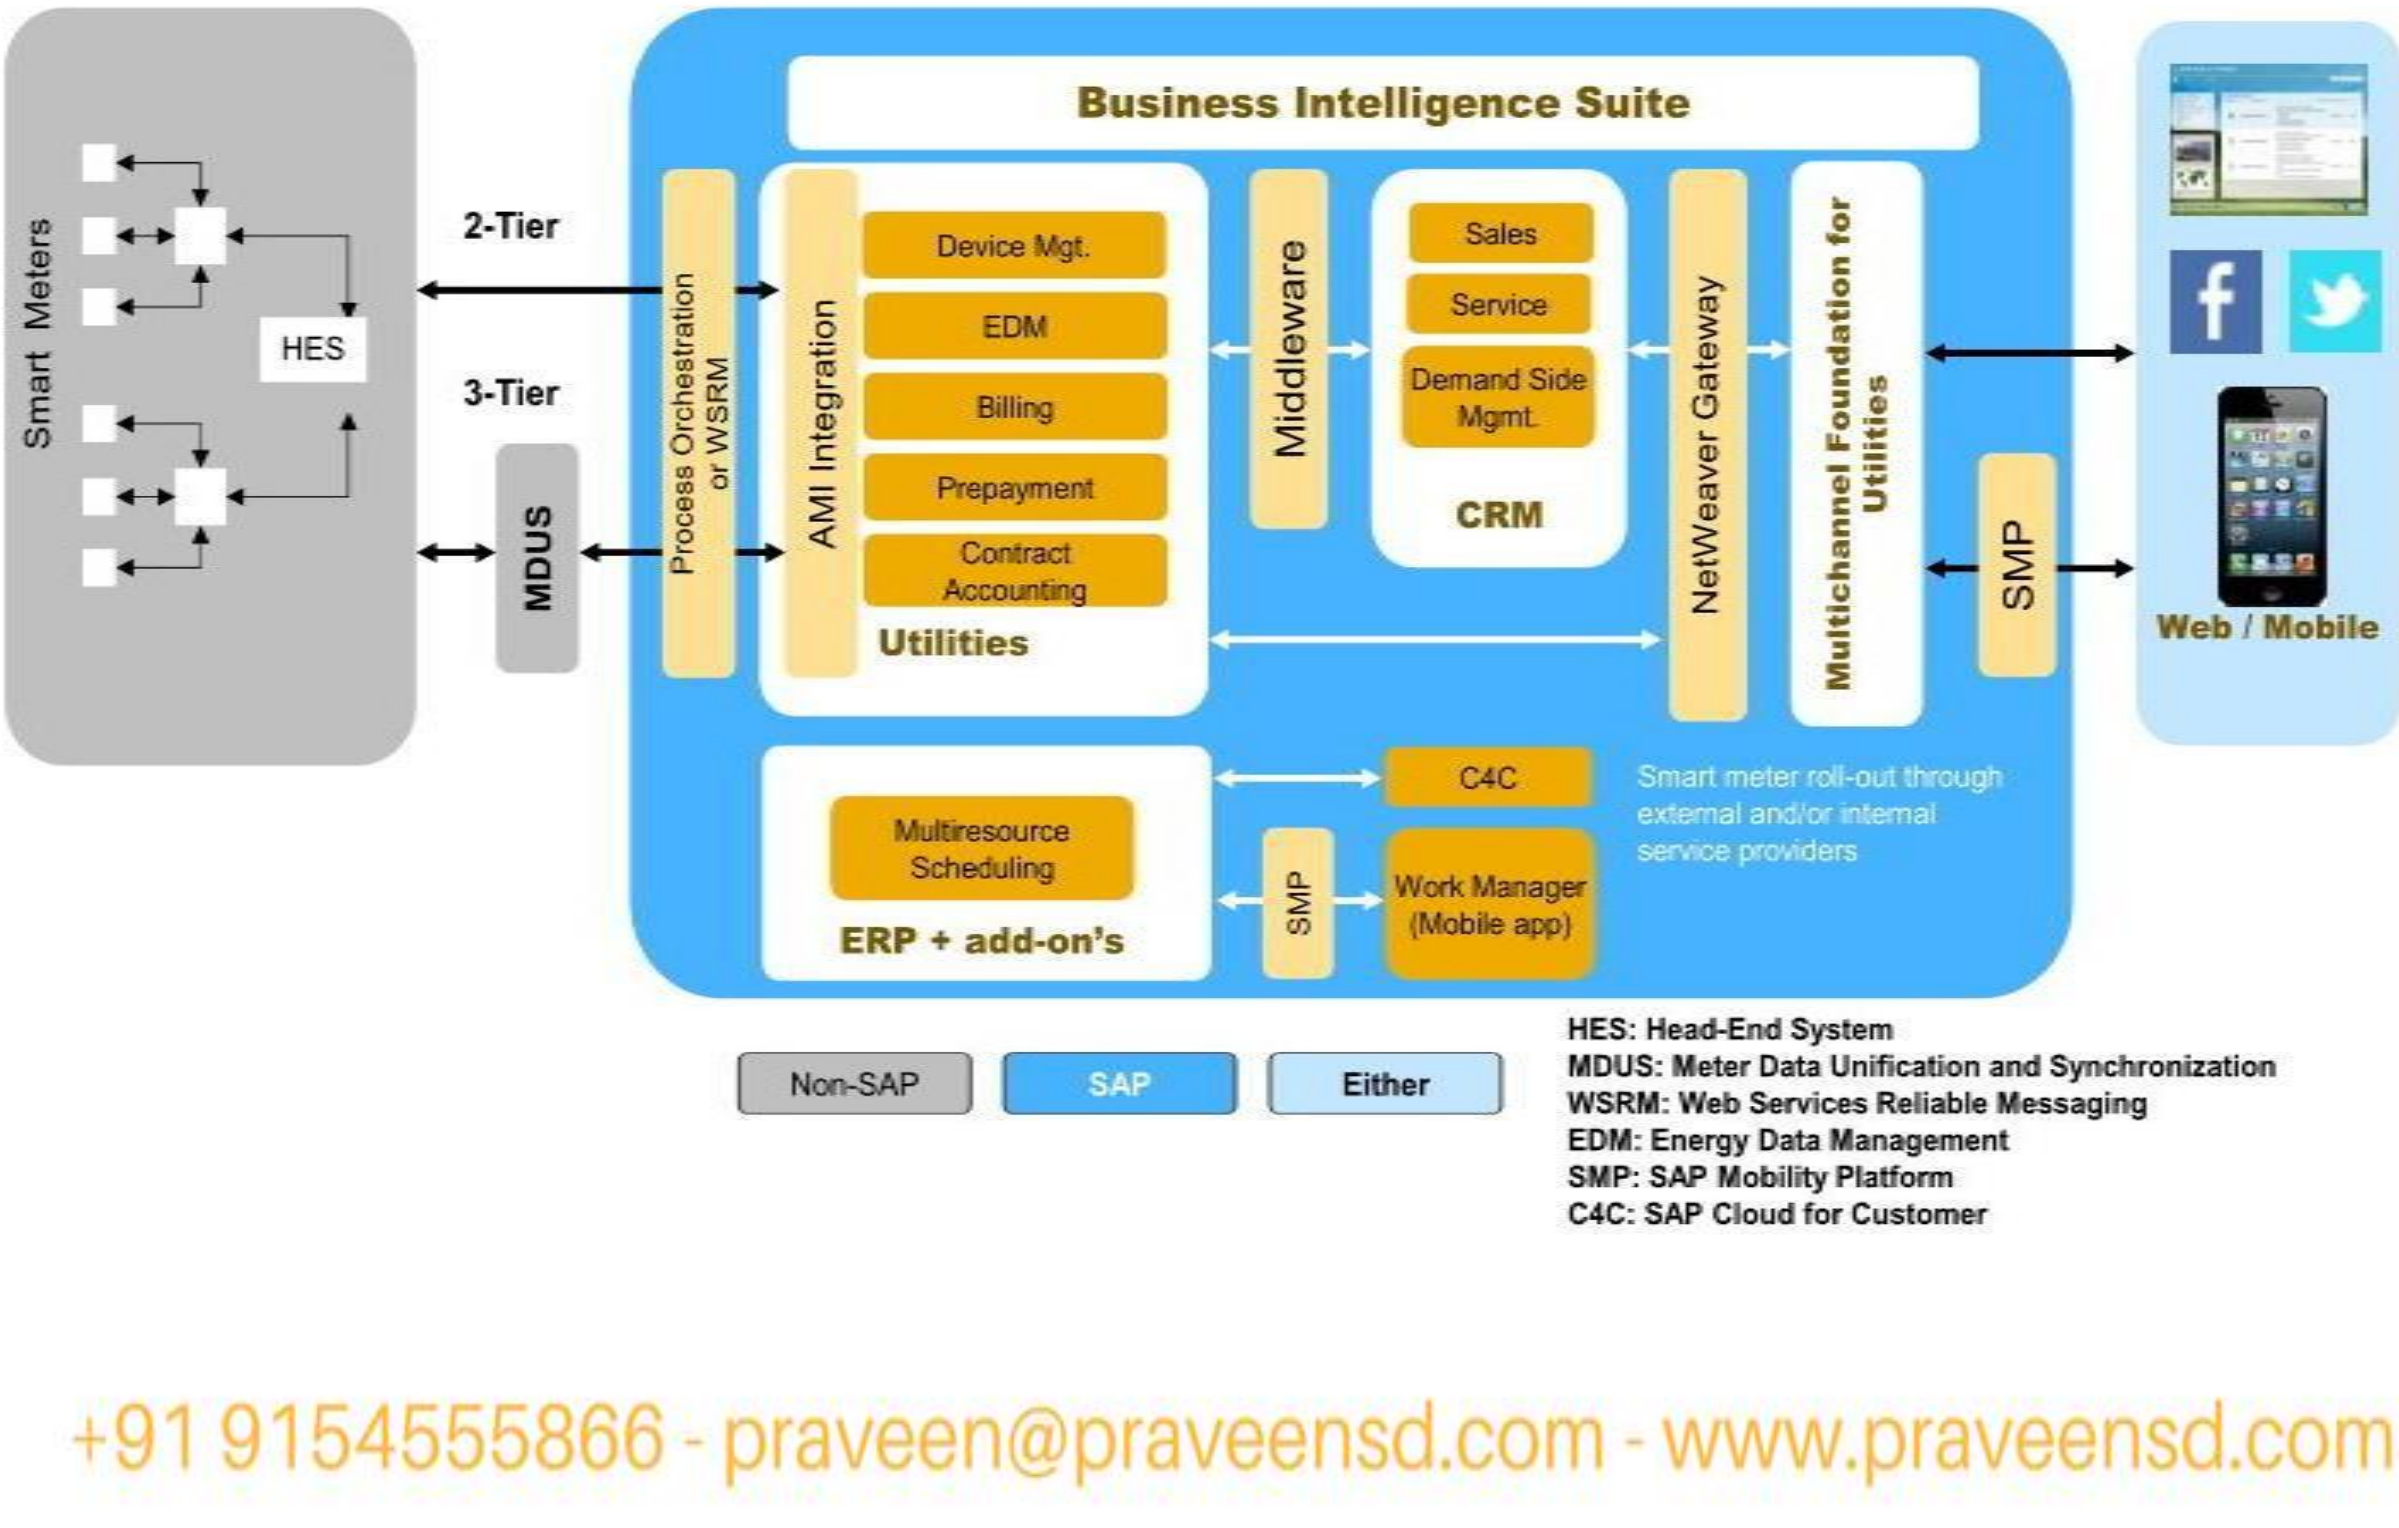

HES: Head-End System
MDUS: Meter Data Unification and Synchronization
WSRM: Web Services Reliable Messaging
EDM: Energy Data Management
SMP: SAP Mobility Platform
C4C: SAP Cloud for Customer

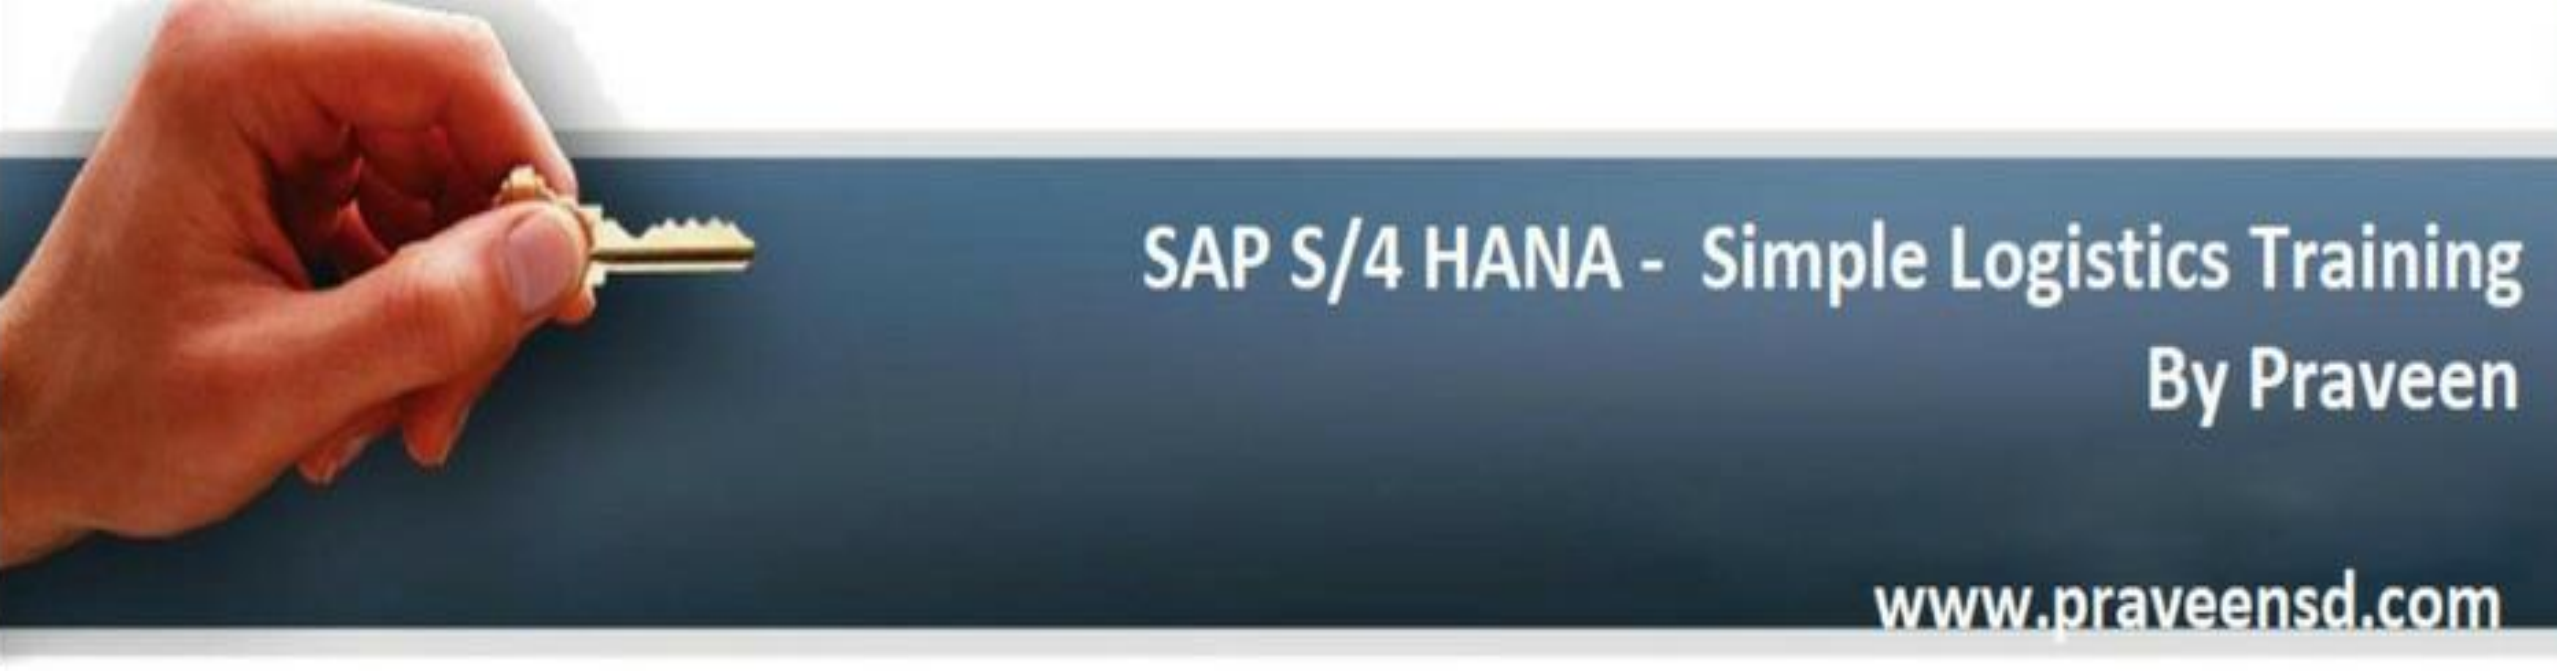
SAP S/4 HANA - Simple Logistics Training
By Praveen
www.praveensd.com

HANA IOT

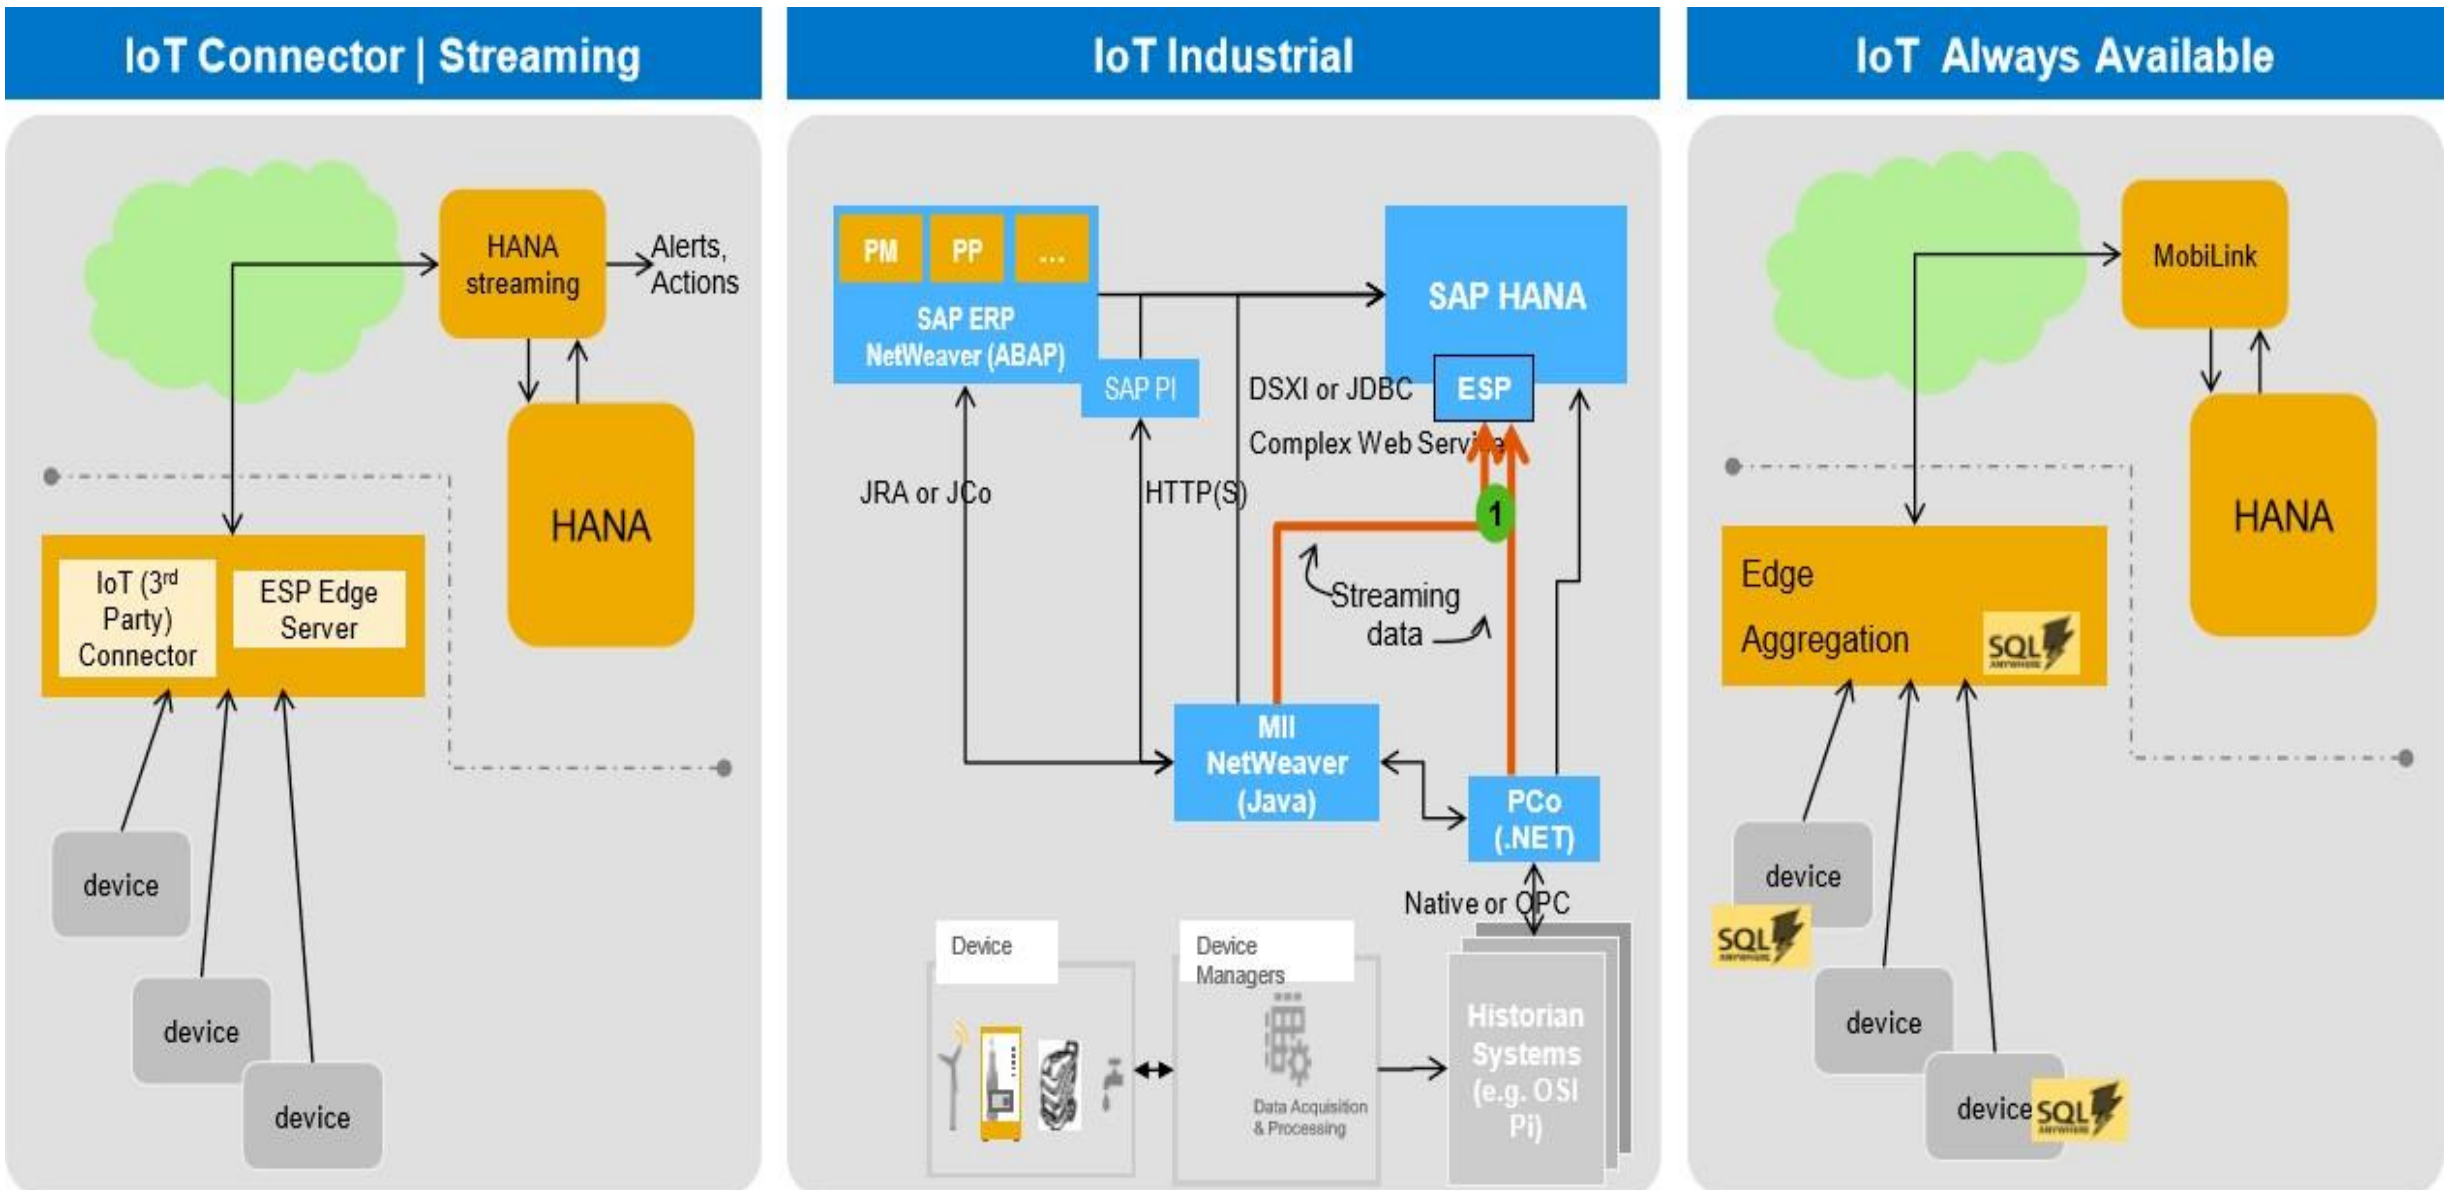
IoT Connector | Streaming
HANA streaming
Alerts, Actions
HANA
IoT (3rd Party) Connector
ESP Edge Server
device
device
device
IoT Industrial
PM
PP
...
SAP ERP NetWeaver (ABAP)
SAP PI
SAP HANA
DSXI or JDBC
ESP
Complex Web Service
JRA or JCo
HTTP(S)
1
Streaming data
MII NetWeaver (Java)
PCo (.NET)
Native or OPC
Device
Device Managers
Data Acquisition & Processing
Historian Systems (e.g. OSI Pi)
IoT Always Available
MobiLink
HANA
Edge Aggregation
SQL
device
device
device
SQL

HANA Live

SAP HANA Live architecture
Foundation for new class of applications

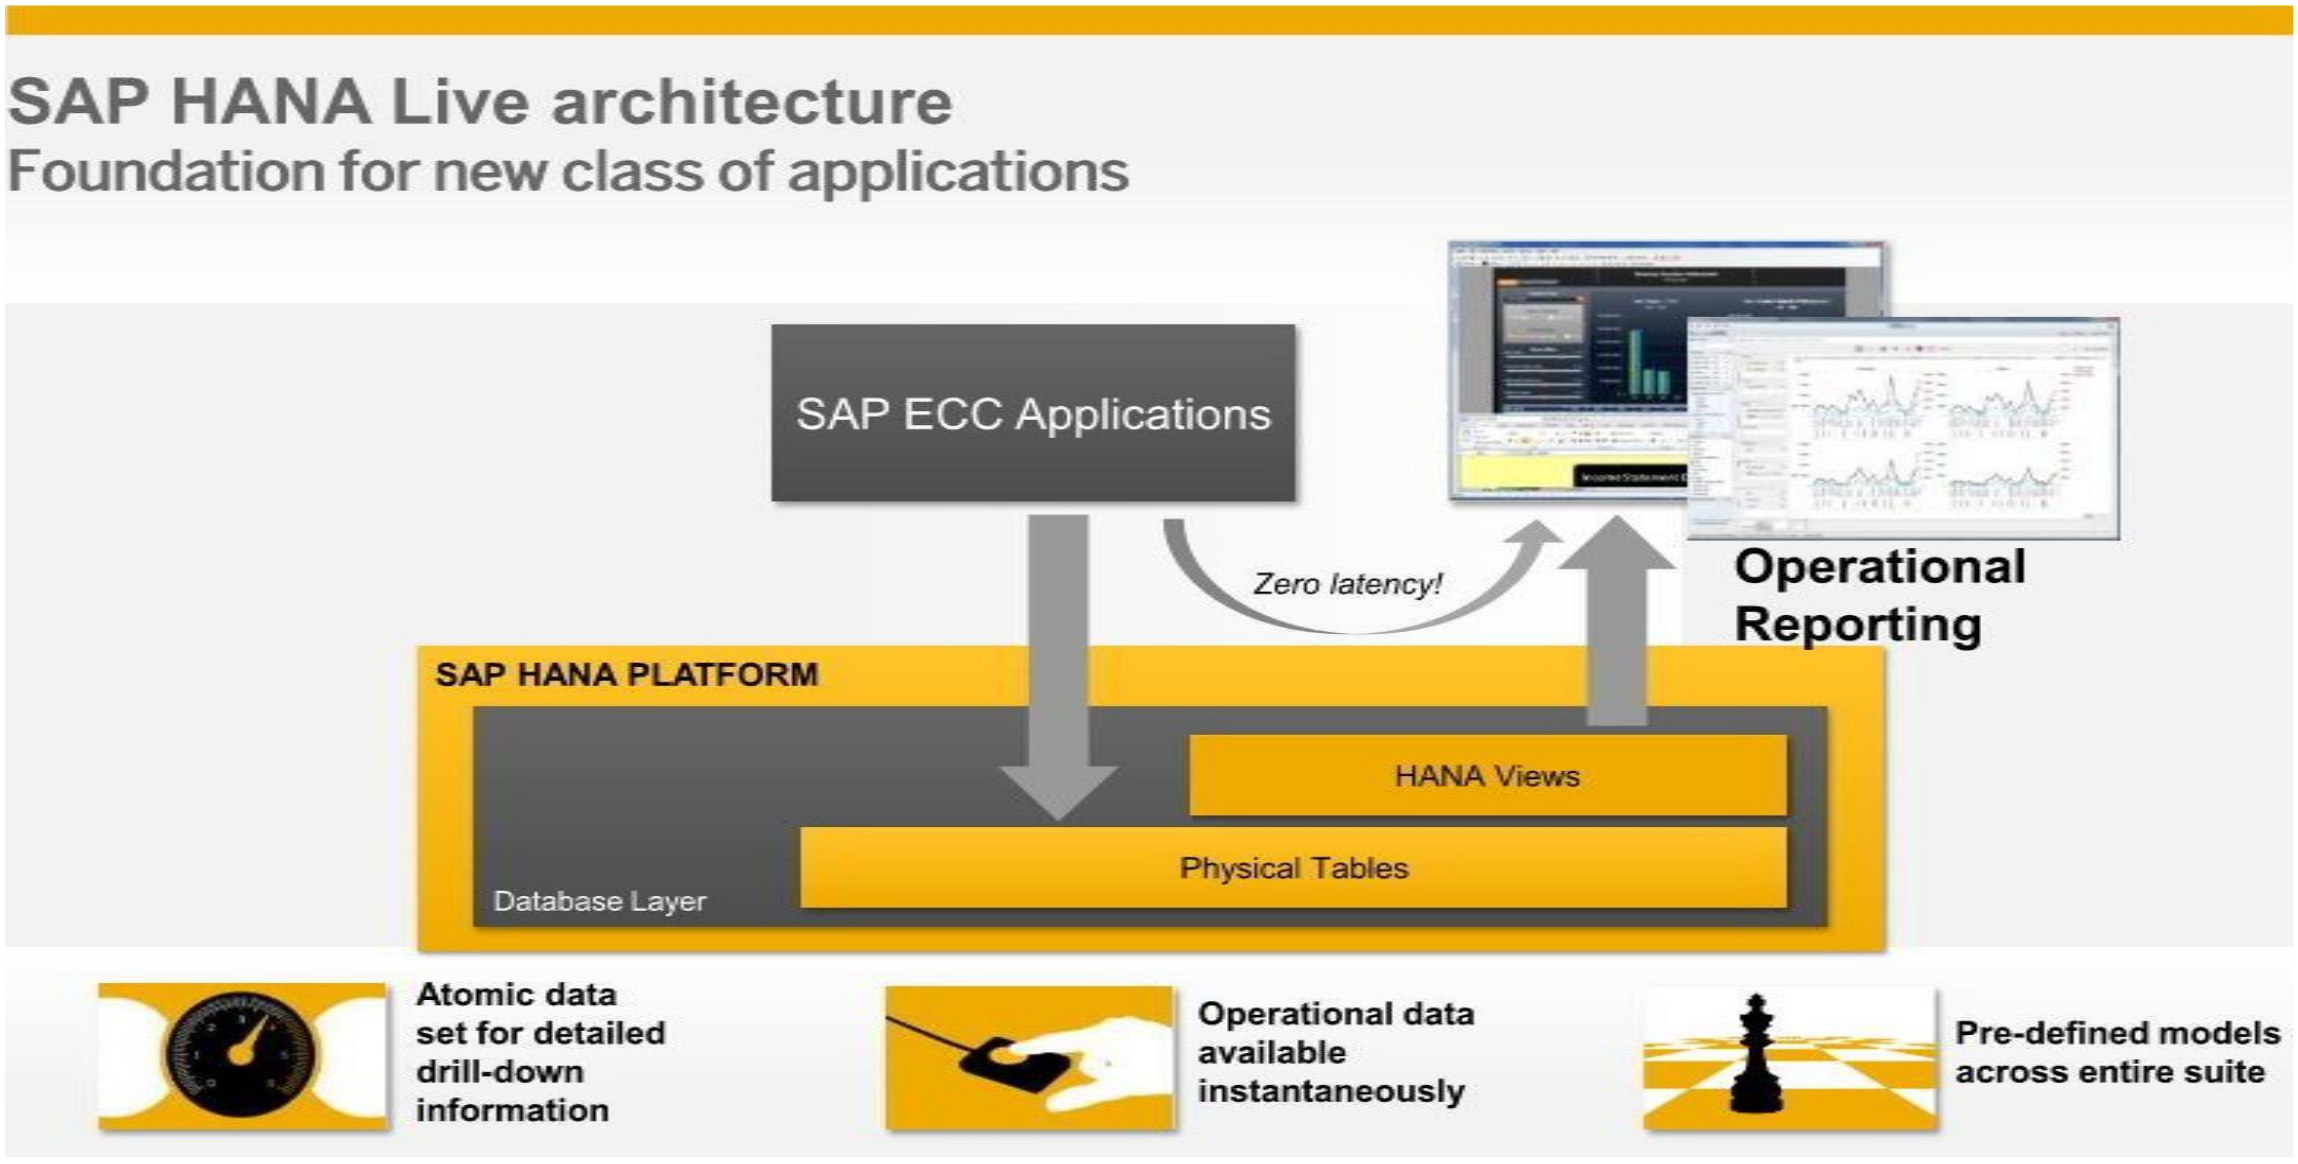

Hana Big data with Vora

- Big Data' as name implies is big volume of data that inundates a business every minute of every day. It could be in both structured and unstructured format. Big data can be analysed for insights that lead to better decisions and strategic business moves, and will fundamentally change the way businesses compete and operate.

- A simple example scenario is, while processing a sales order, the business want to give special discounts to customers based on their transactions history (say for last 30 years). The recent transaction data are available in Hana and the very old data are moved to Hadoop

- In fact the analysis has to be done over data made by correlating the recent data residing in Hana and old data residing in Hadoop and it has to be processed on-the fly in faster fashion. Being an in-memory database Hana runs very faster. SAP Vora is also an in- memory query engine and it can access data from Hadoop and process some of them in a faster way.

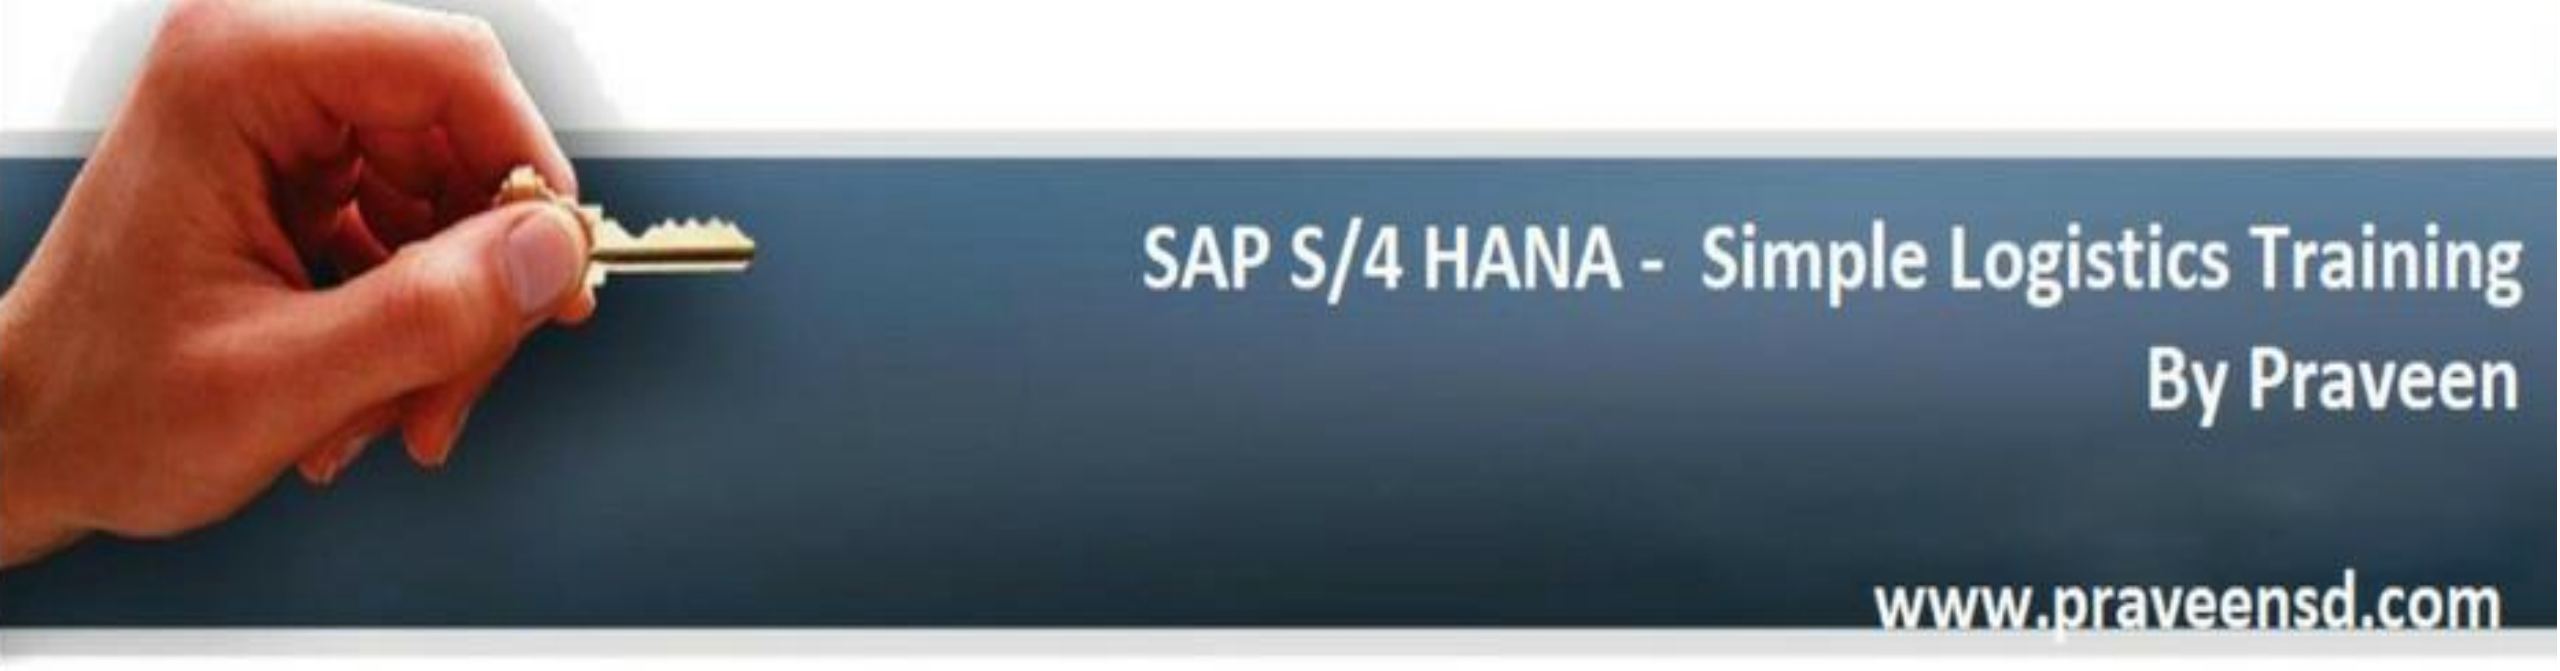

- SAP Vora builds structured data hierarchies for the unstructured data in Hadoop and integrates it with data from HANA, and then through the Apache Spark SQL interface it enables OLAP-style in-memory analysis on the combined data.

- Vora serves the purpose as a mediator when the compatibility question rises in between SAP HANA and Hadoop. Spark is not well compatible with HANA Systems and HANA Clouds, so SAP built something which follows the Spark framework and also have HANA Adapters for data connectivity.

- Typically, all this data in Hadoop is unstructured and SQL cannot be run immediately on top of that. And that's where Vora adds value and also could be a bridge between HANA & Hadoop

How In-memory data base will help

- Faster Processing
- Ability to give real time analysis
- Live dash boards
- Lesser data foot print on account no data redundancy
- Single source of truth

Compression

- Data is compressed by different compression techniques (e.g. dictionary encoding, run length encoding, sparse encoding, cluster encoding, indirect encoding) in SAP HANA Column store.

- When main memory limit is reached in SAP HANA, the whole database objects (table, view,etc.) that are not used will be unloaded from the main memory and saved into the disk.

- These objects names are defined by application semantic and reloaded into main memory from the disk when required again. Under normal circumstances SAP HANA database manages unloading and loading of data automatically.

- However, the user can load and unload data from individual table manually by selecting a table in SAP HANA studio in respective Schema- by right- clicking and selecting the option "Unload/Load".

What is compression ratio

SAP HANA: Dictionary Compression

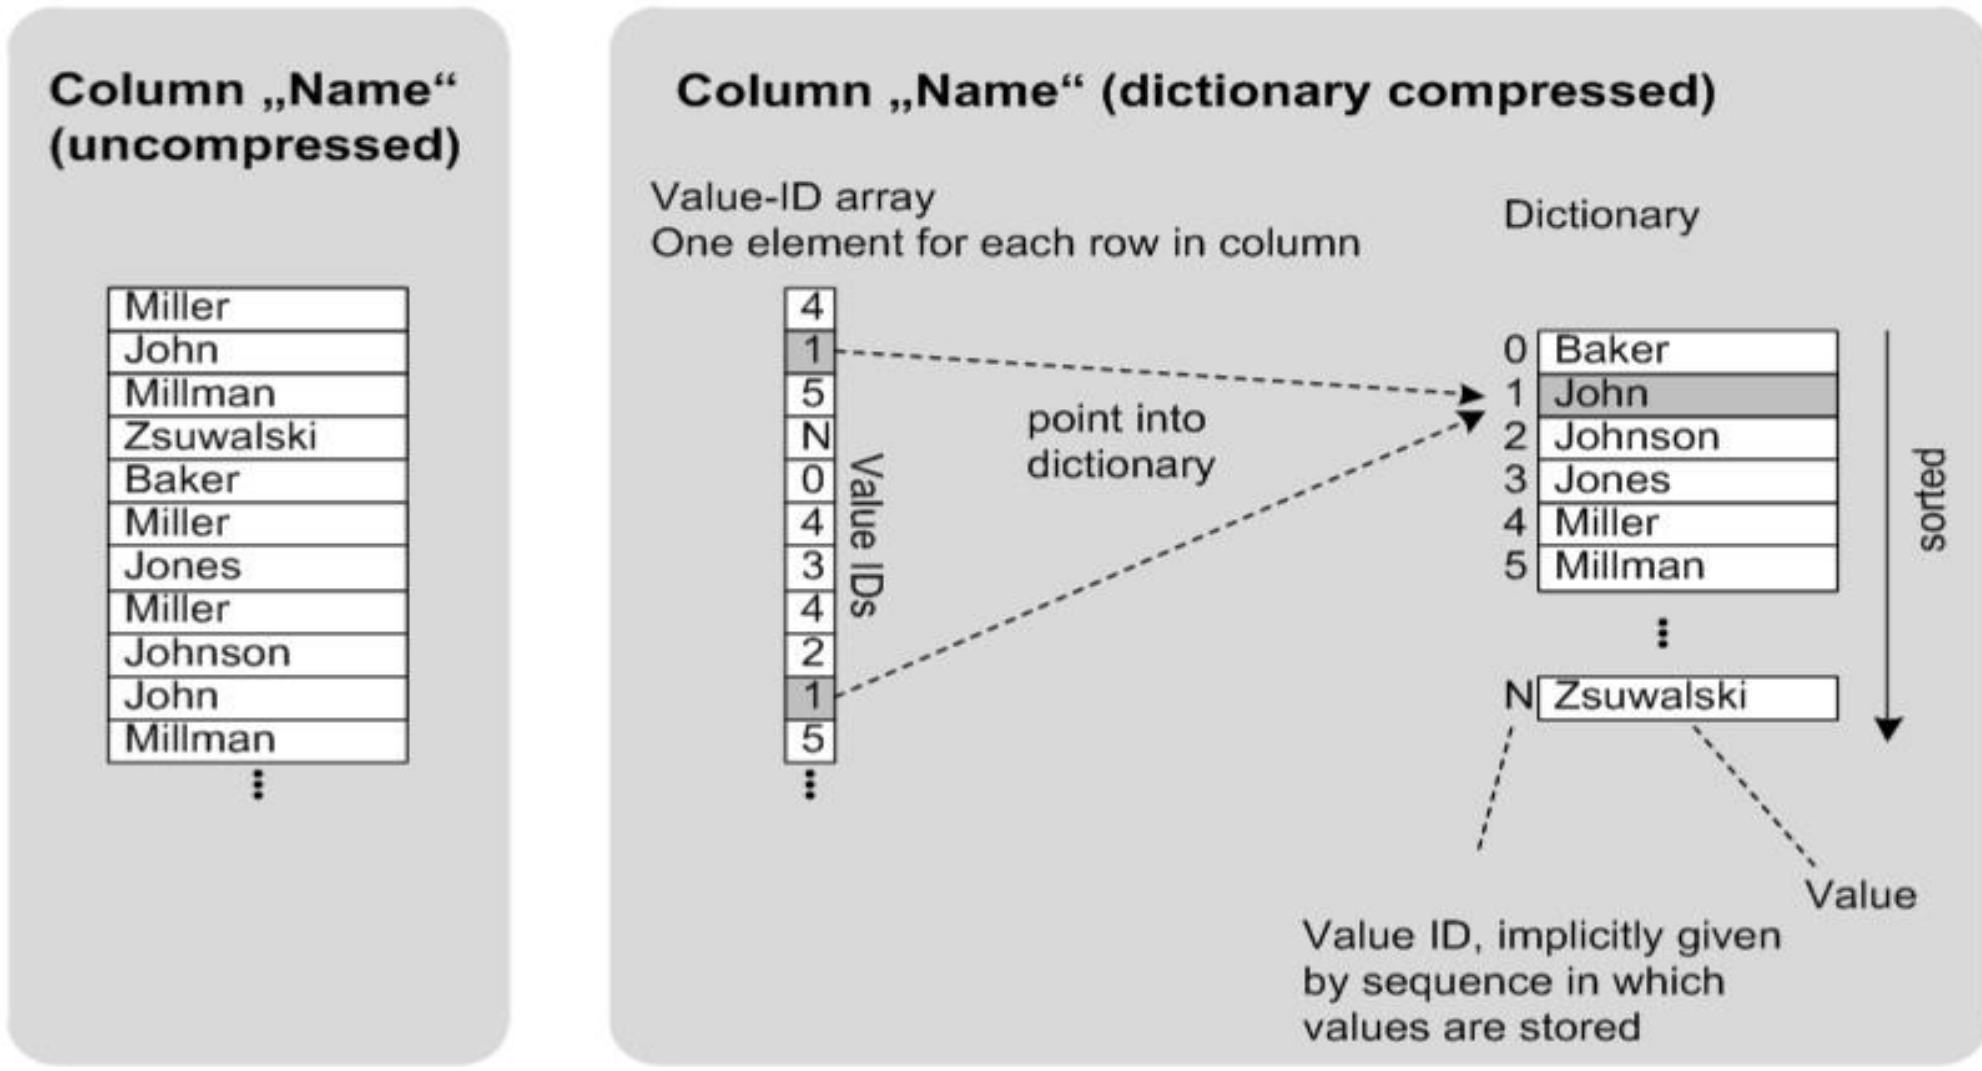

Compression Technologies

Additional Compression Technologies

1) Prefix Coding

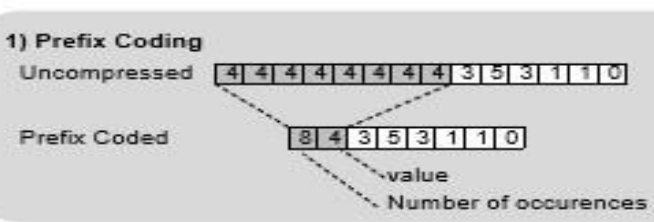

2) Run Length Coding

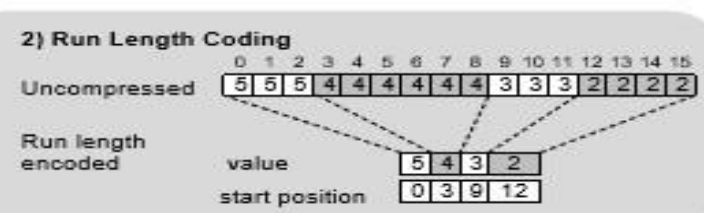

3) Cluster Coding

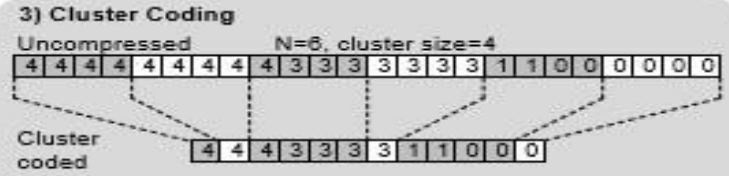

4) Sparse Coding

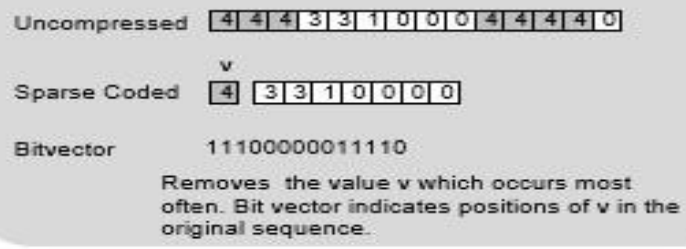

5) Indirect coding

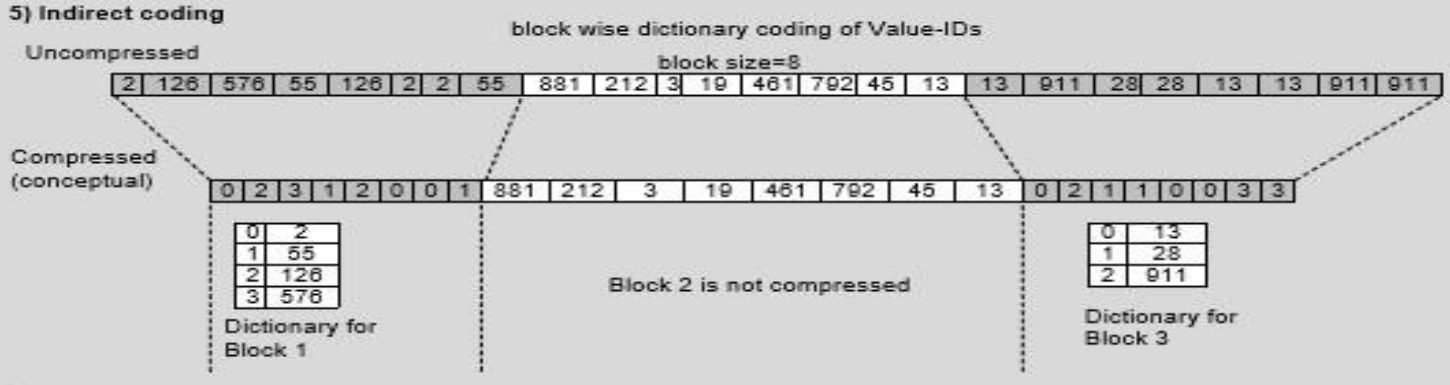

Multi -Core Parallel

SAP HANA: Multi-Core Parallelization

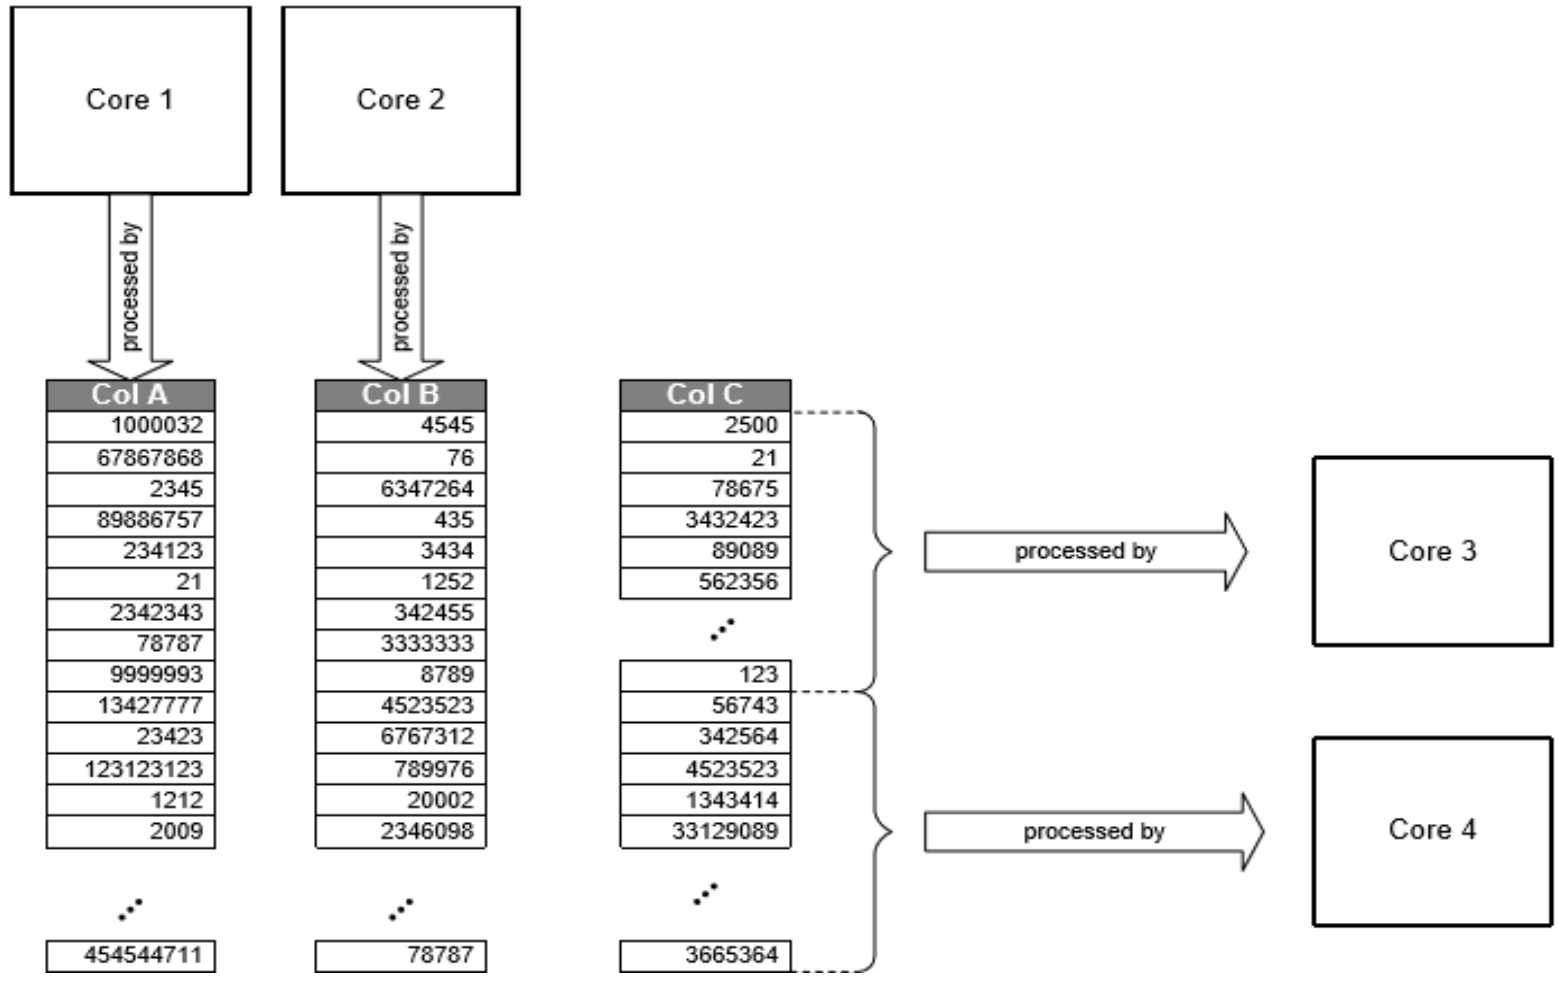

Multi User parallel

Parallelization at All Levels

- Multiple user sessions
- Concurrent operations within a query (… T1.A … T2.B…)
- Data partitioning on one or more hosts
- Horizontal segmentation, concurrent aggregation
- Multi-threading at Intel processor core level
- Vector processing

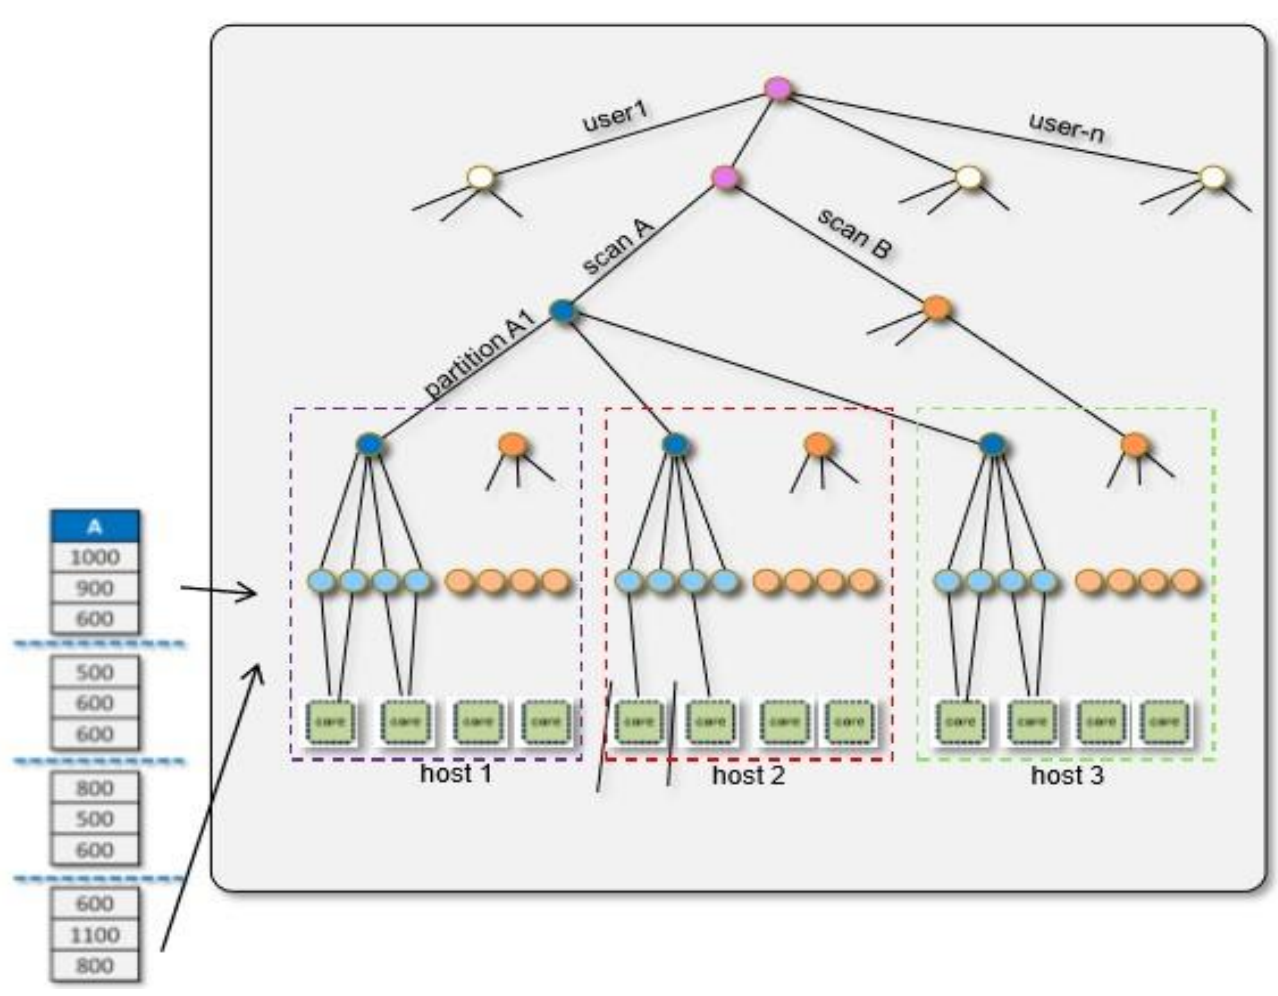

HANA Streaming

Smart data streaming extends the capabilities of the SAP HANA Platform

Stream capture

- **Capture data arriving as individual events – at potentially high speeds**
  - Hundreds of thousands or millions of events per second
  - Micro-batching and parallel processing to optimize load speeds
- **Capture events that are published from streaming sources**
  - e.g. message bus
- **Filter, transform or enrich the data on the way in**
  - Capture only the data you want, in the form you need it
- **Prioritize data**
  - Capture high value data in HANA and direct other data into Hadoop

Continuous analysis, Immediate Response

- **Monitor incoming event streams**
  - Watch for trends or patterns
  - Monitor correlations
  - Detect missing events
  - Continuously update and monitor aggregate statistics
- **Generate alerts, notifications**
- **Initiate immediate response**

HANA Streaming

Complex Event Processing extracts insight from events

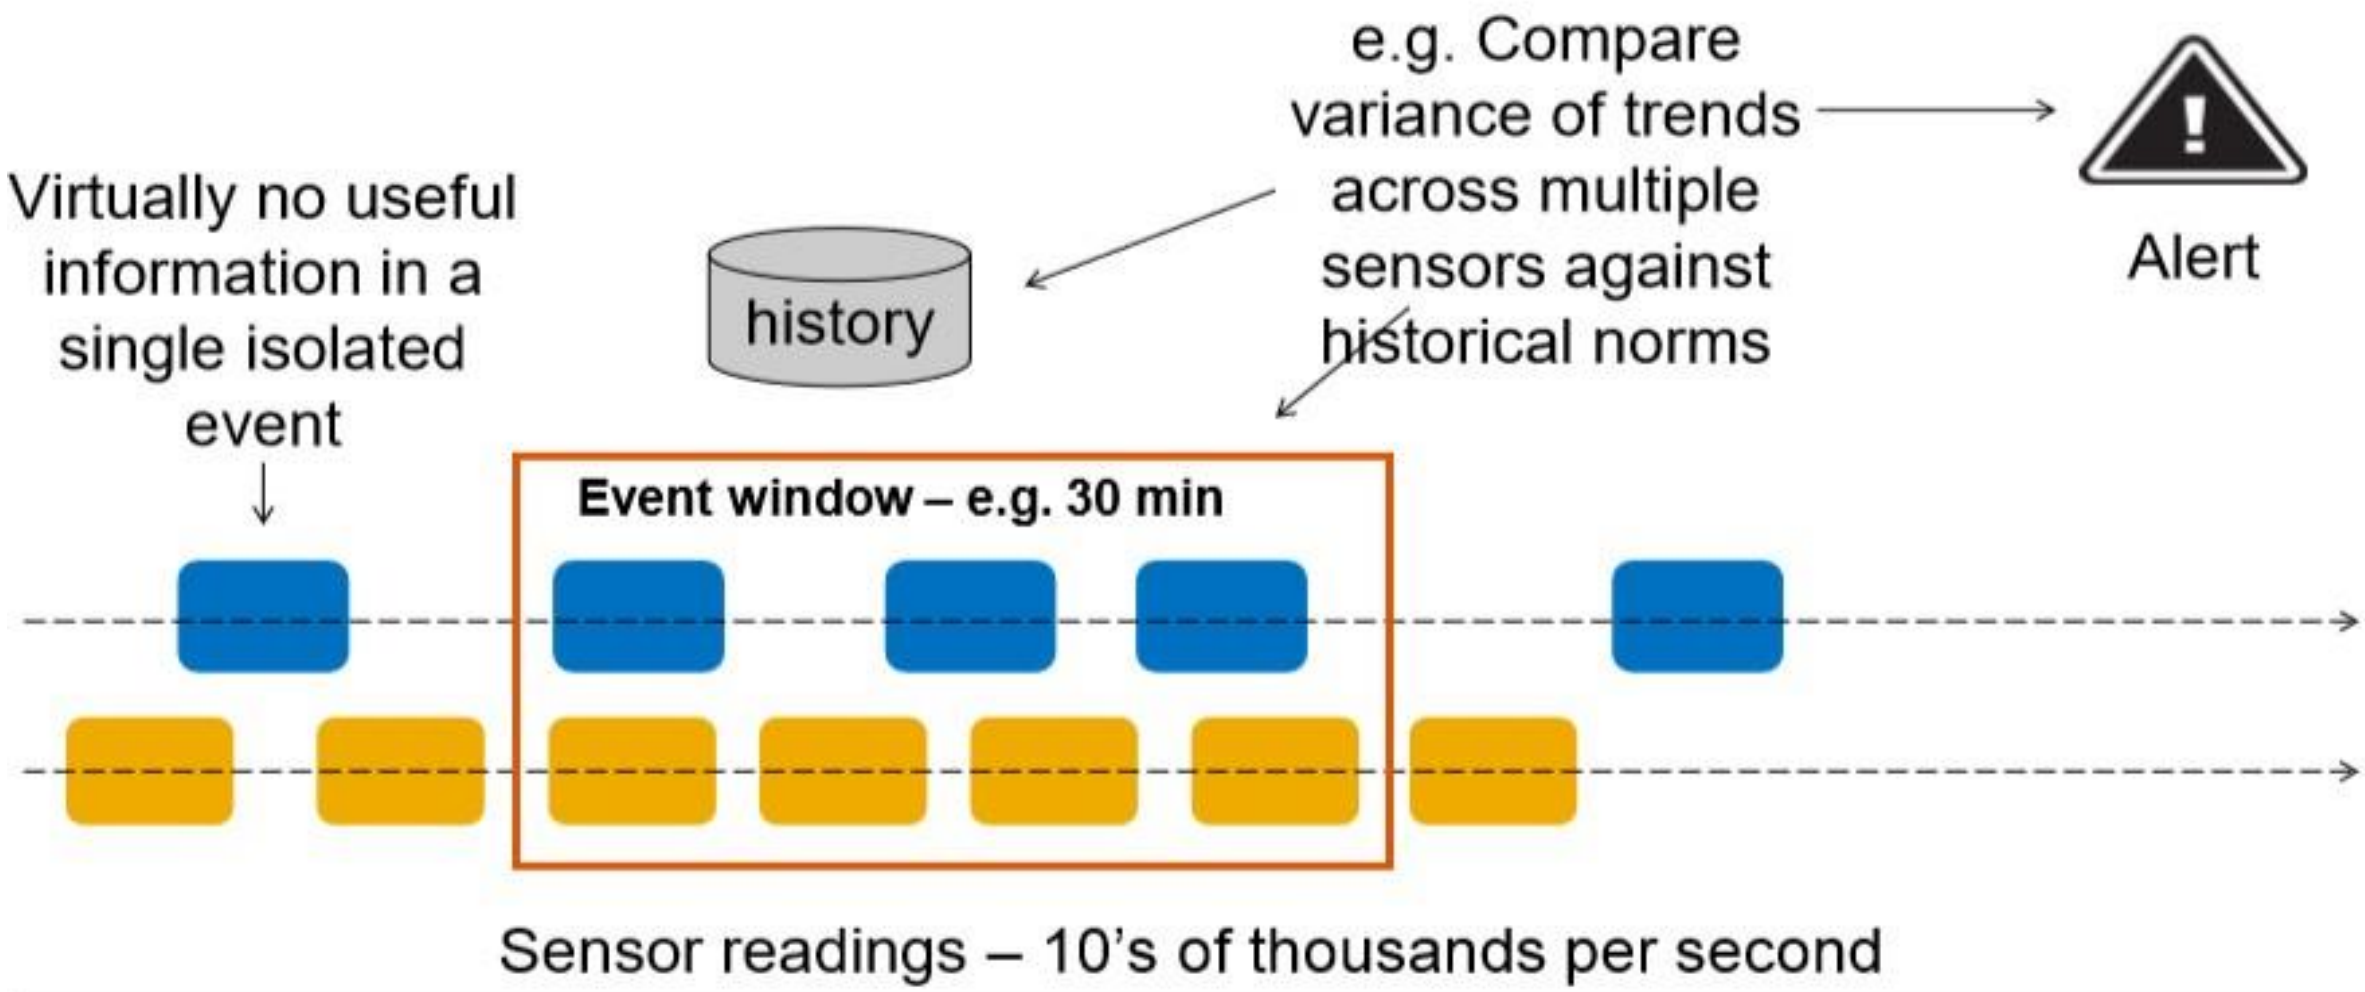

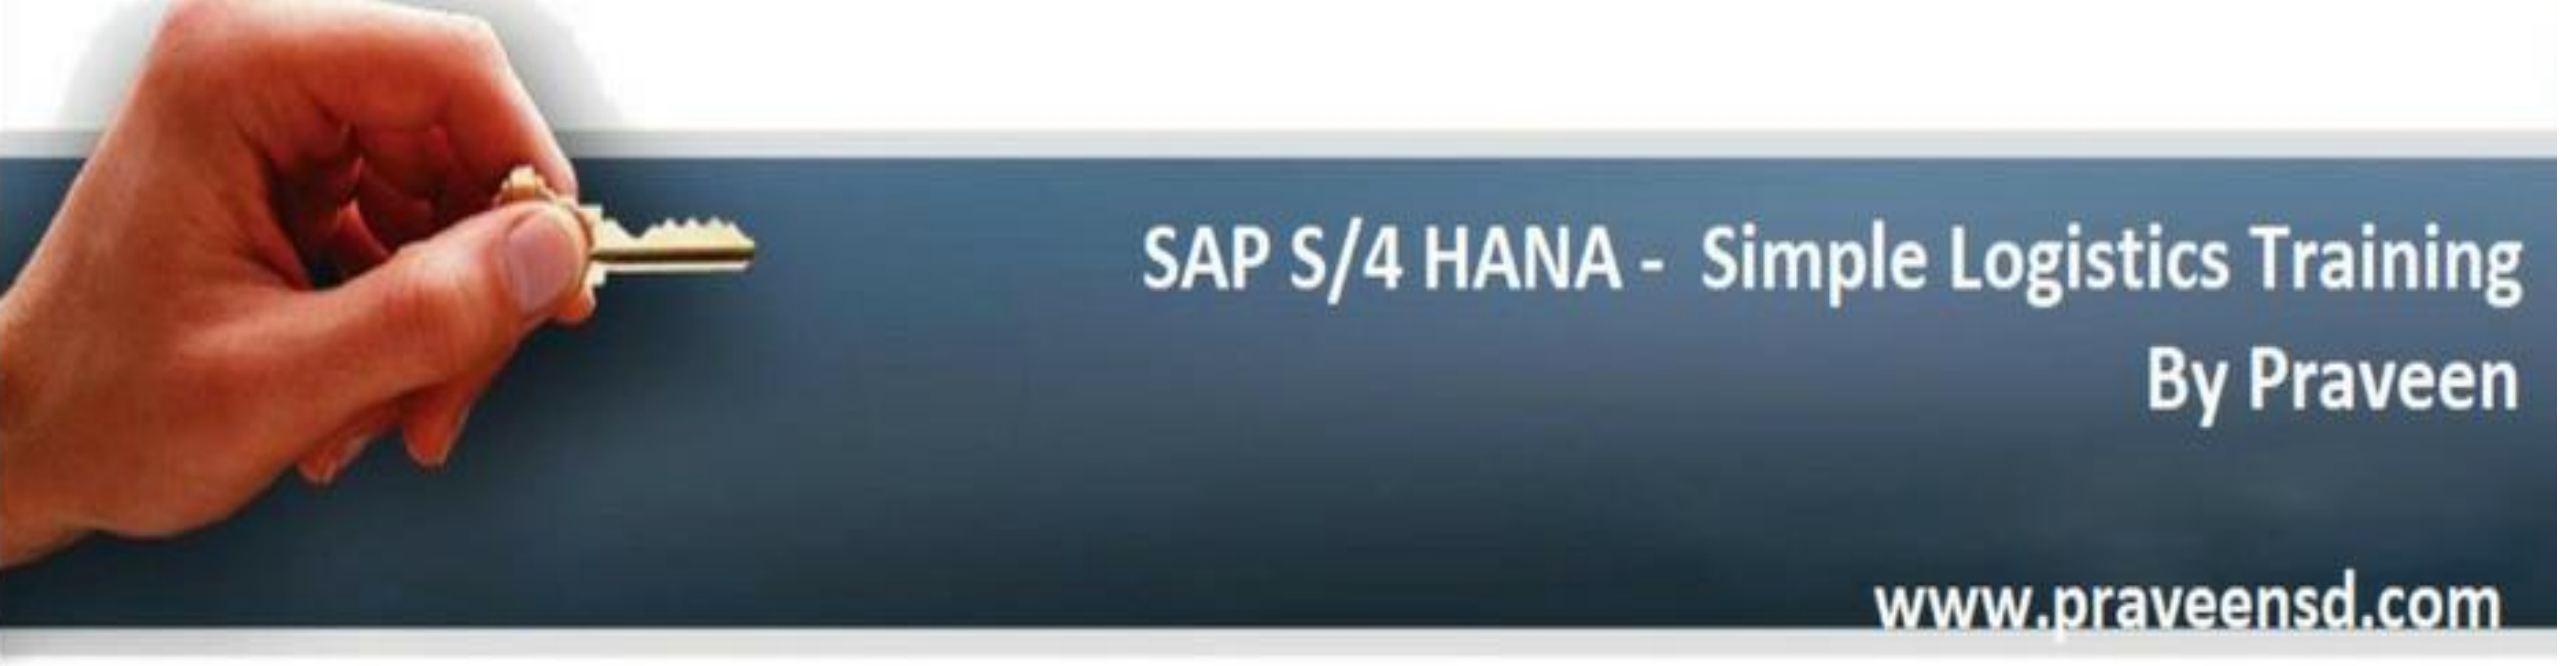
SAP S/4 HANA - Simple Logistics Training
By Praveen
www.praveensd.com

HANA IOT

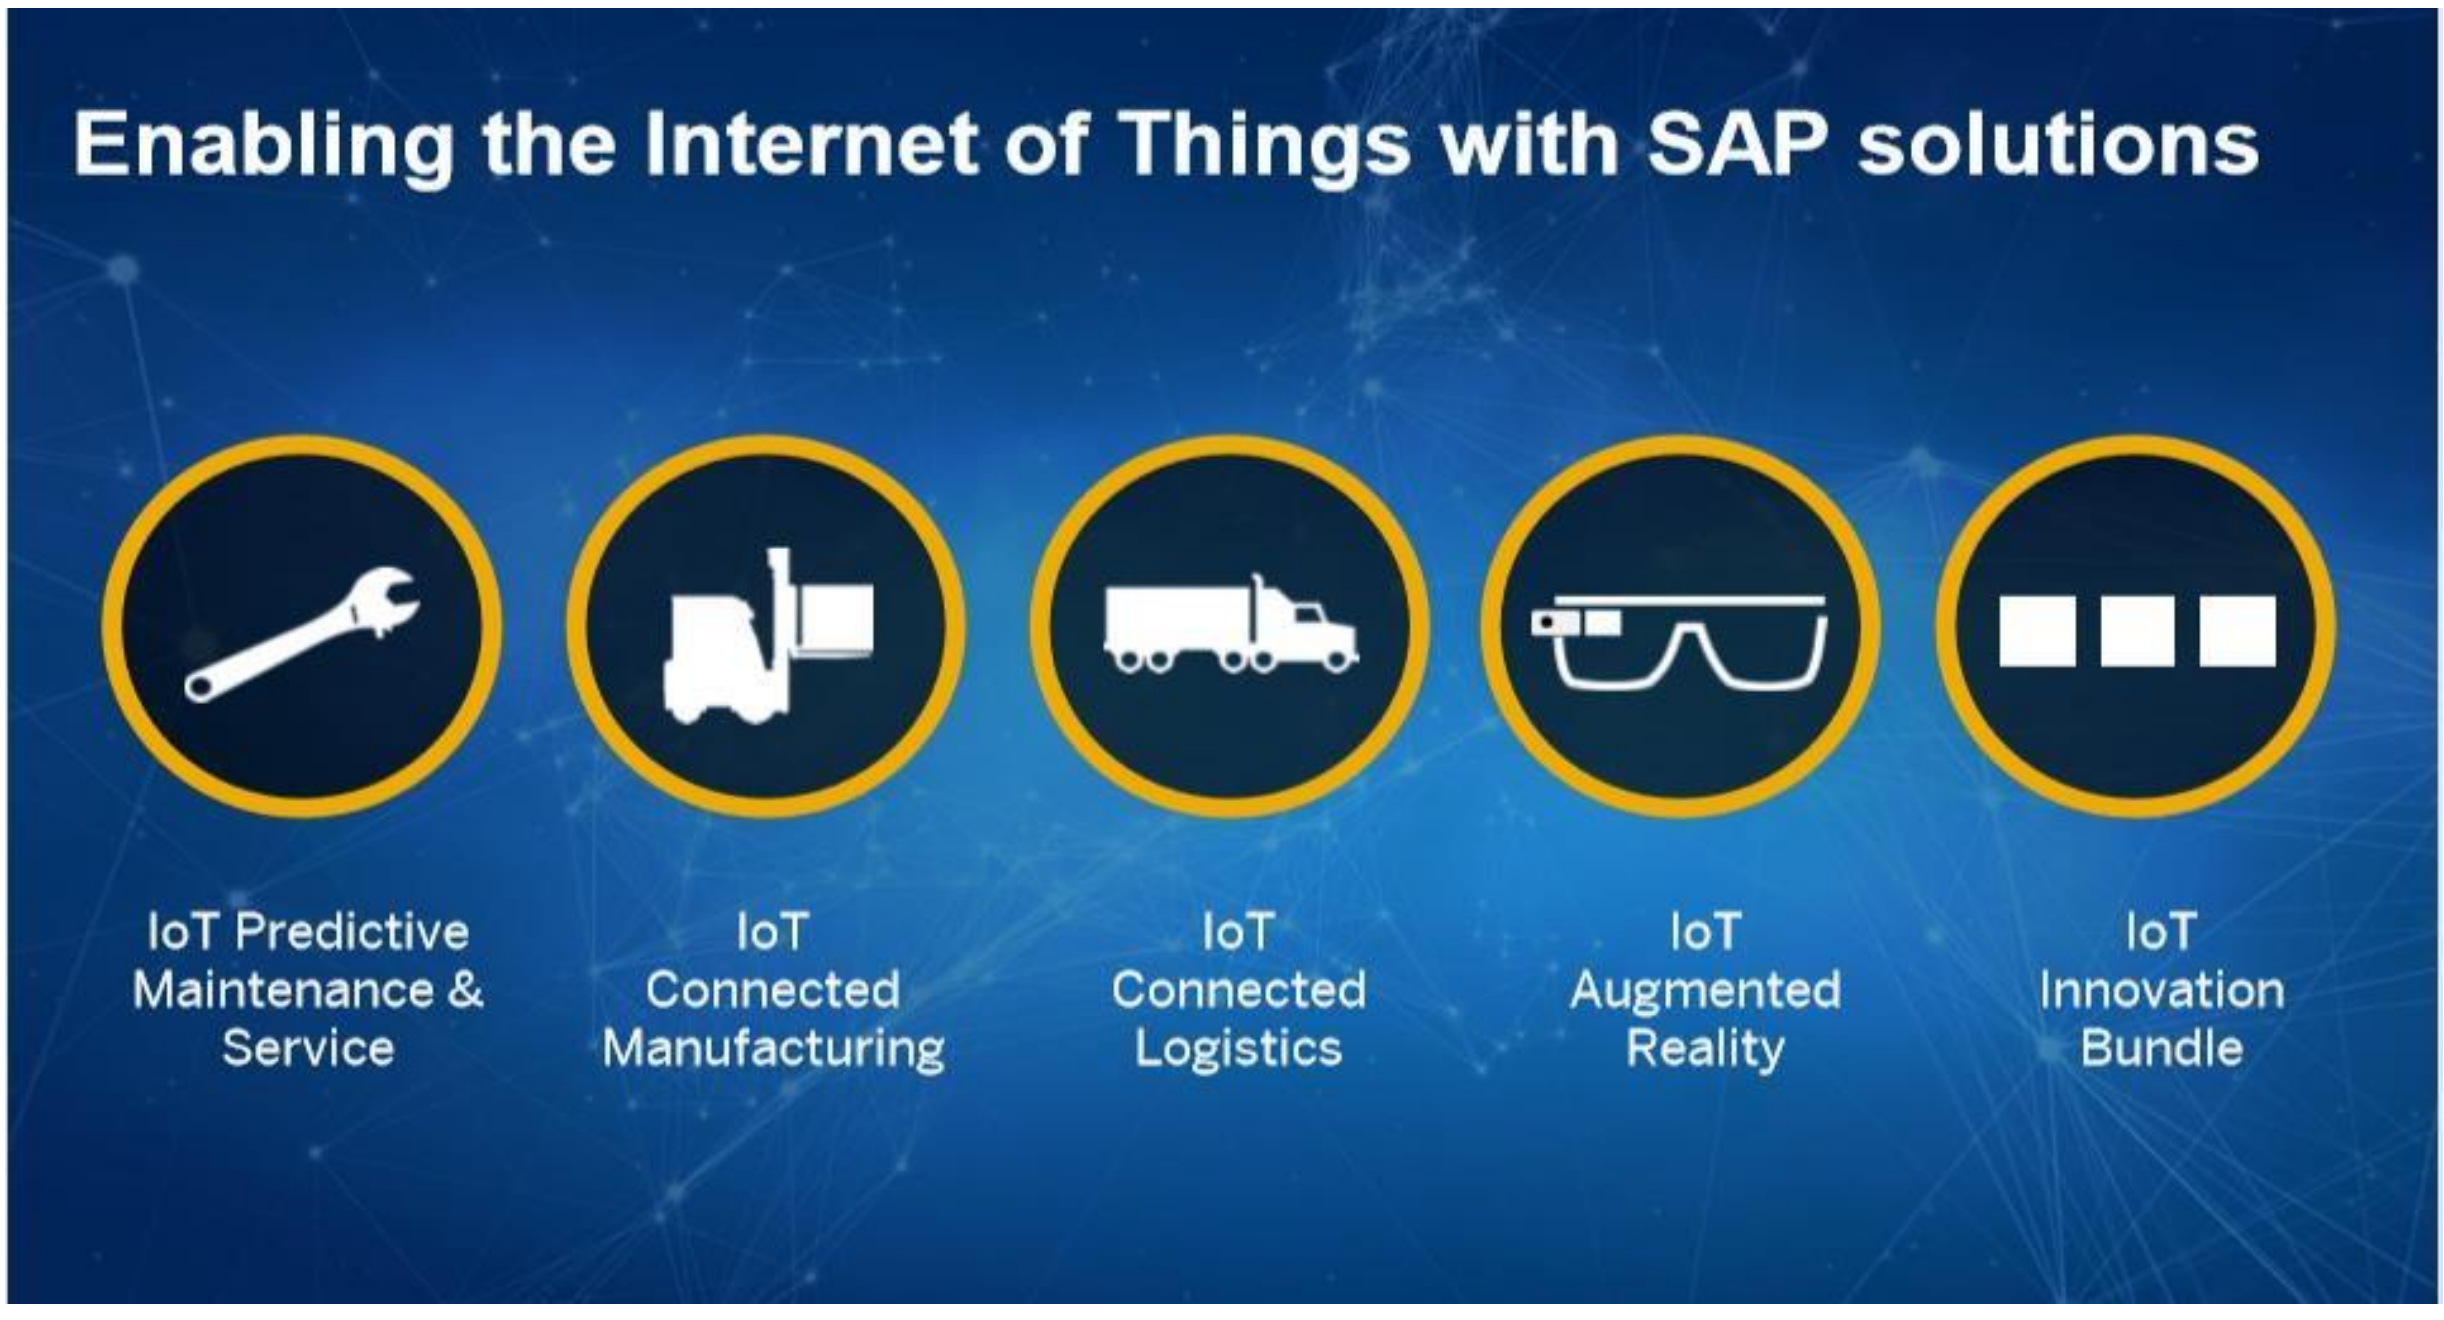
Enabling the Internet of Things with SAP solutions
IoT Predictive Maintenance & Service
IoT Connected Manufacturing
IoT Connected Logistics
IoT Augmented Reality
IoT Innovation Bundle

HANA IOT Application

Predictive Maintenance

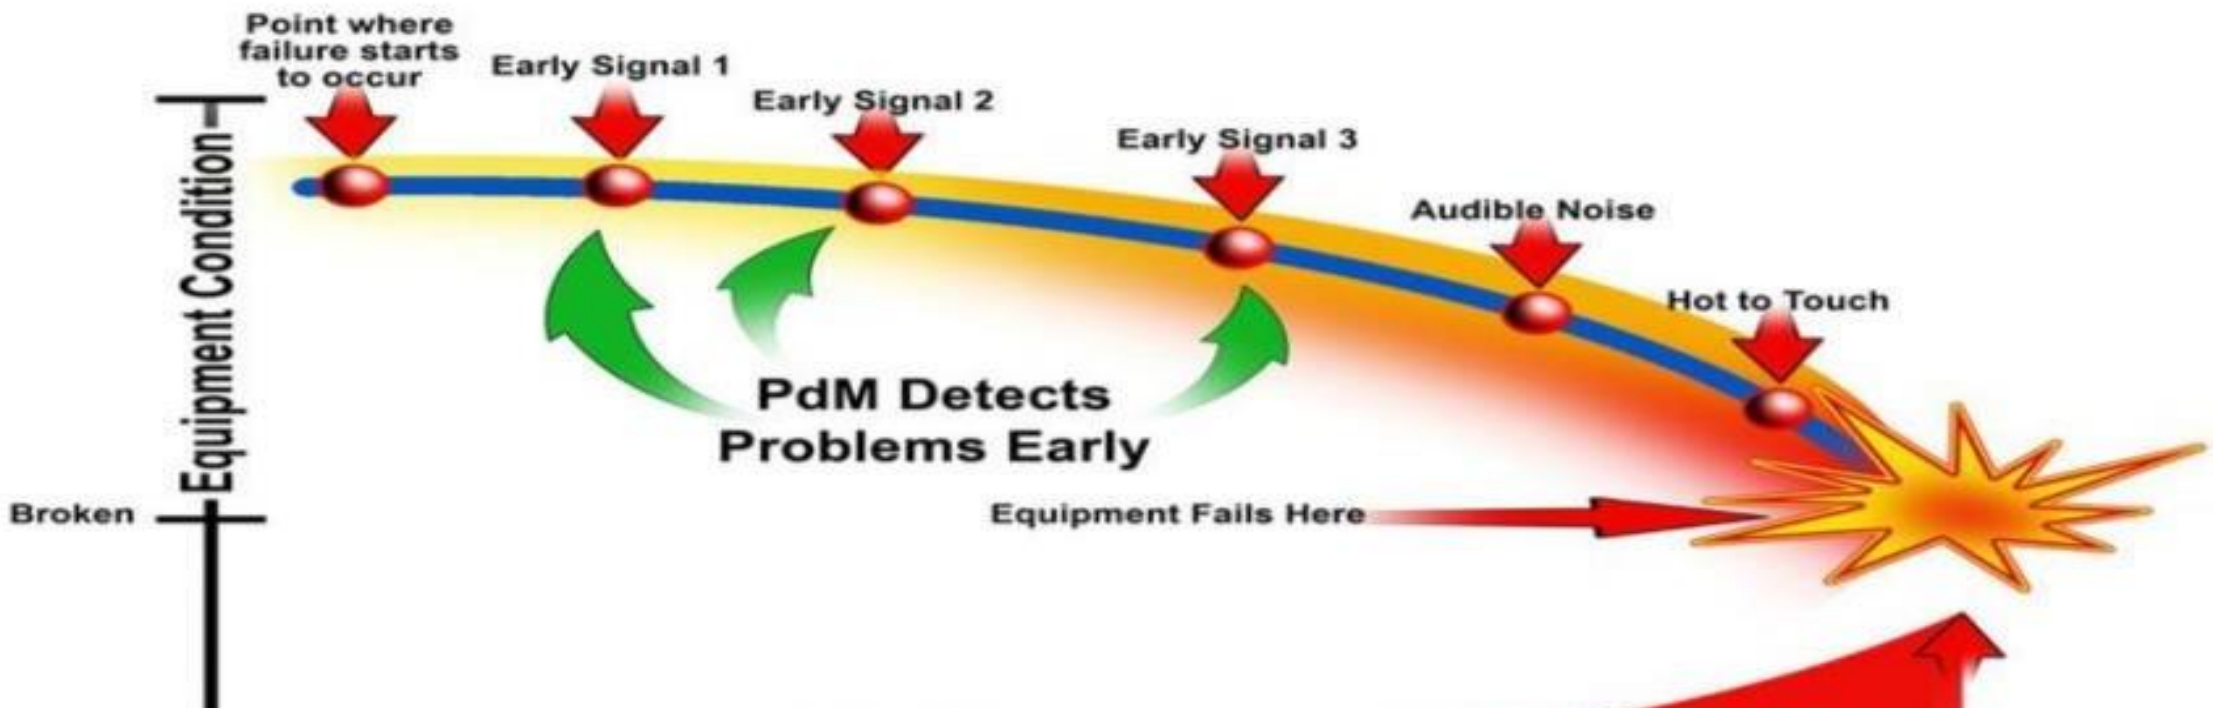

Building Blocks

Architecture

HANA Core Platform

ONE platform for simple and efficient data processing

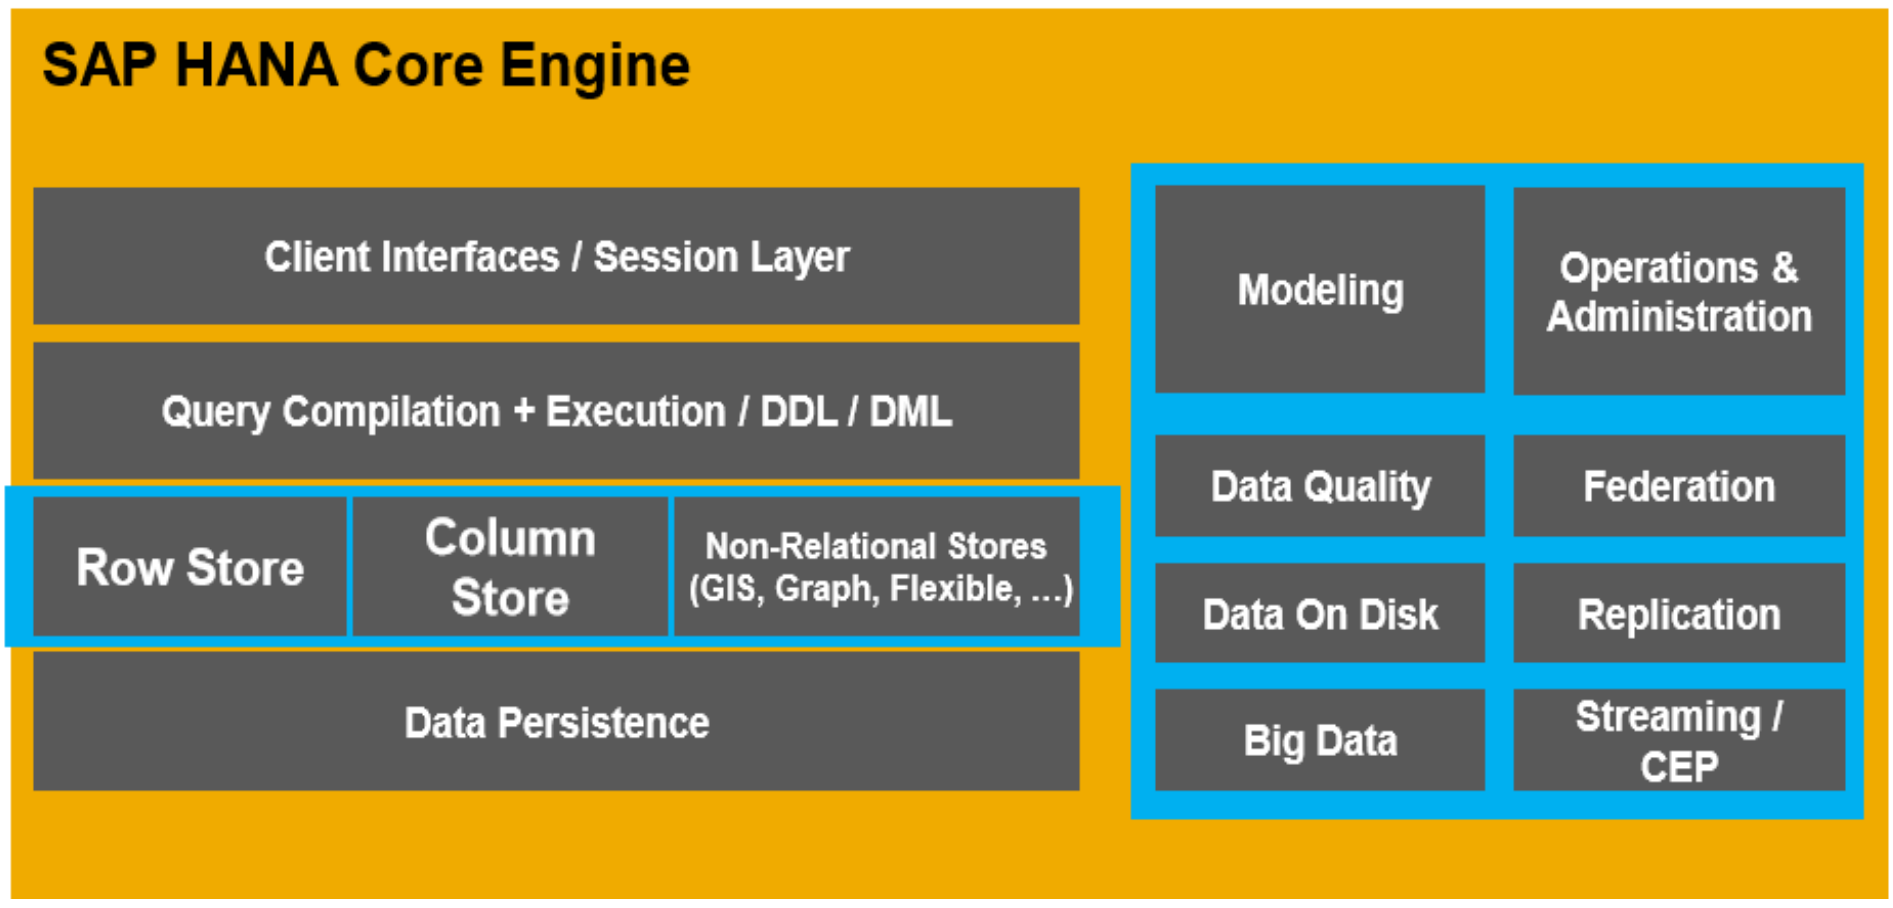

HANA Architecture

SAP HANA Server consists of:

1. Index Server
2. Pre-processor Server
3. Name Server
4. Statistics Server
5. XS Engine

Architecture Discussion

- Key points of the HANA architecture:

- In-Memory Database – the entire database is running in memory Combines OLTP, OLAP, and HW Acceleration, eliminating unnecessary complexity and latency

- Strict hardware specifications for performance designs reasons – 16GB of memory per CPU core, fixed ratio of disk data storage to 4 times RAM, log storage = 1 times RAM

- Two types of Relational Stores in HANA – Row Store (Common in previous traditional databases) and Column Store (is what HANA uses primarily)

- Persistency Layer – In Memory is volatile, so persistency later makes sure all data in memory is also stored on hard drive storage

- Not only is data stored in HANA memory, but what makes it faster is that the calculations are made in the database and only the results transfer to the application layer

Data Loading to HANA

- SAP Landscape Transformation (SLT),
- Data Services (DS),
- Smart Data Access (SDA)
- SAP HANA the major focus is real-time
- Data Services can do Near real-time and batch processing.

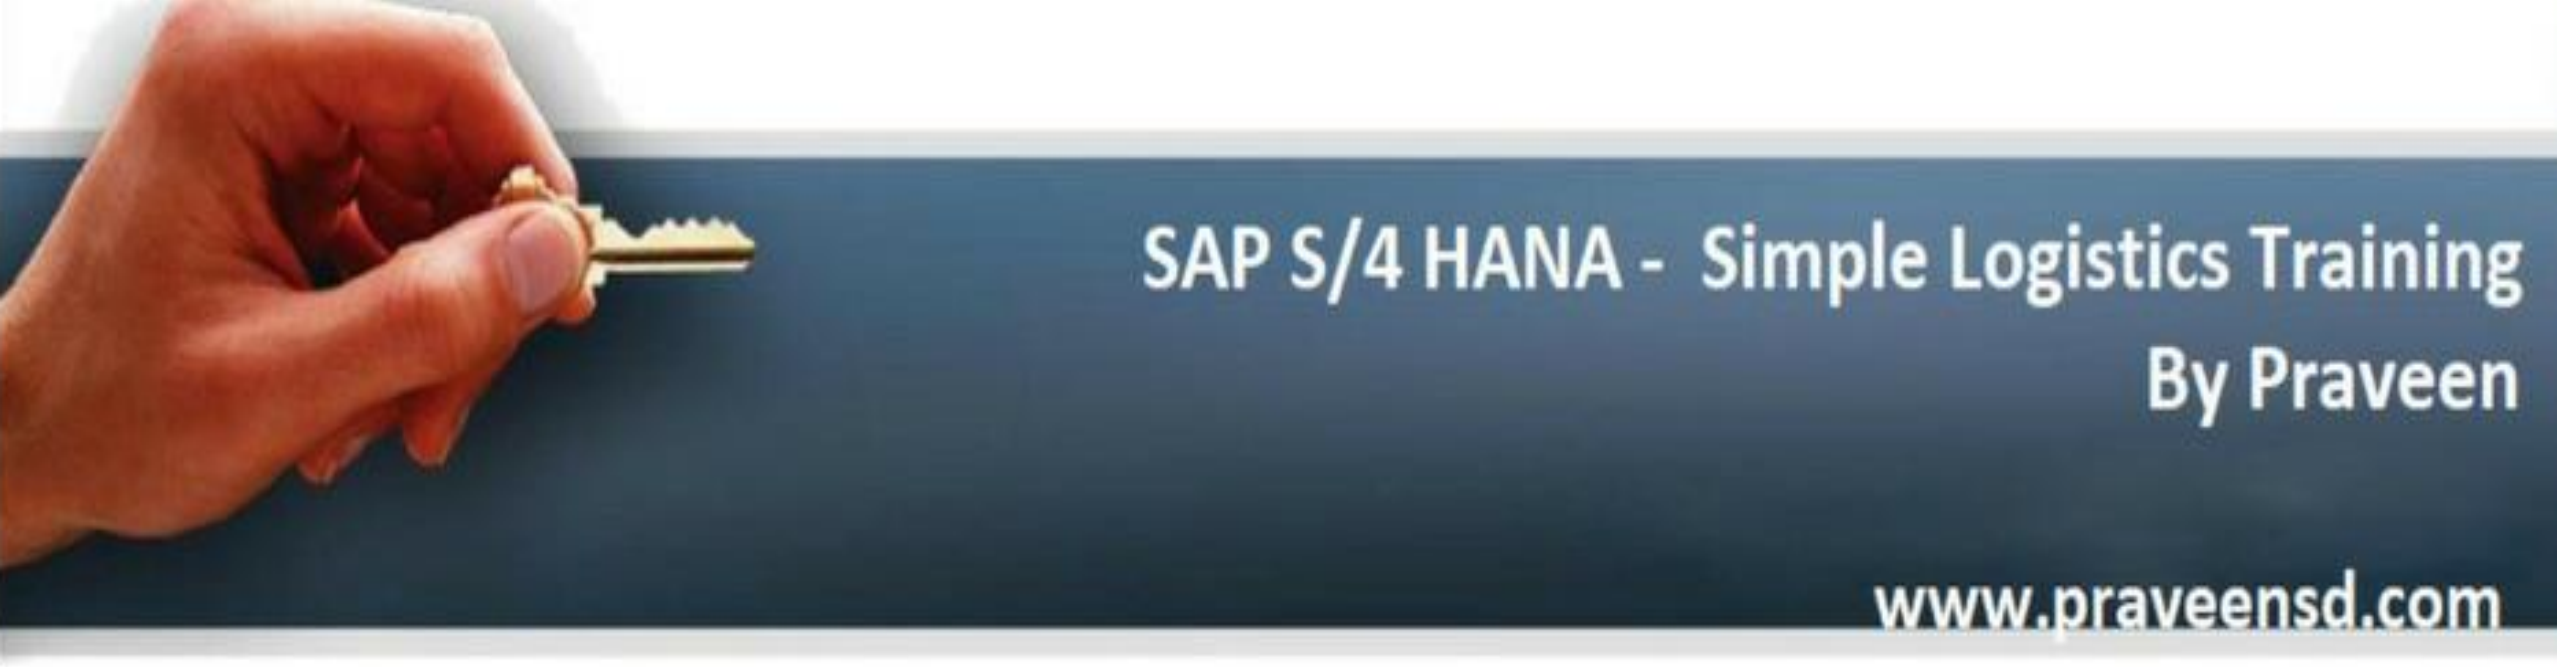

Data Provisioning For SAP HANA

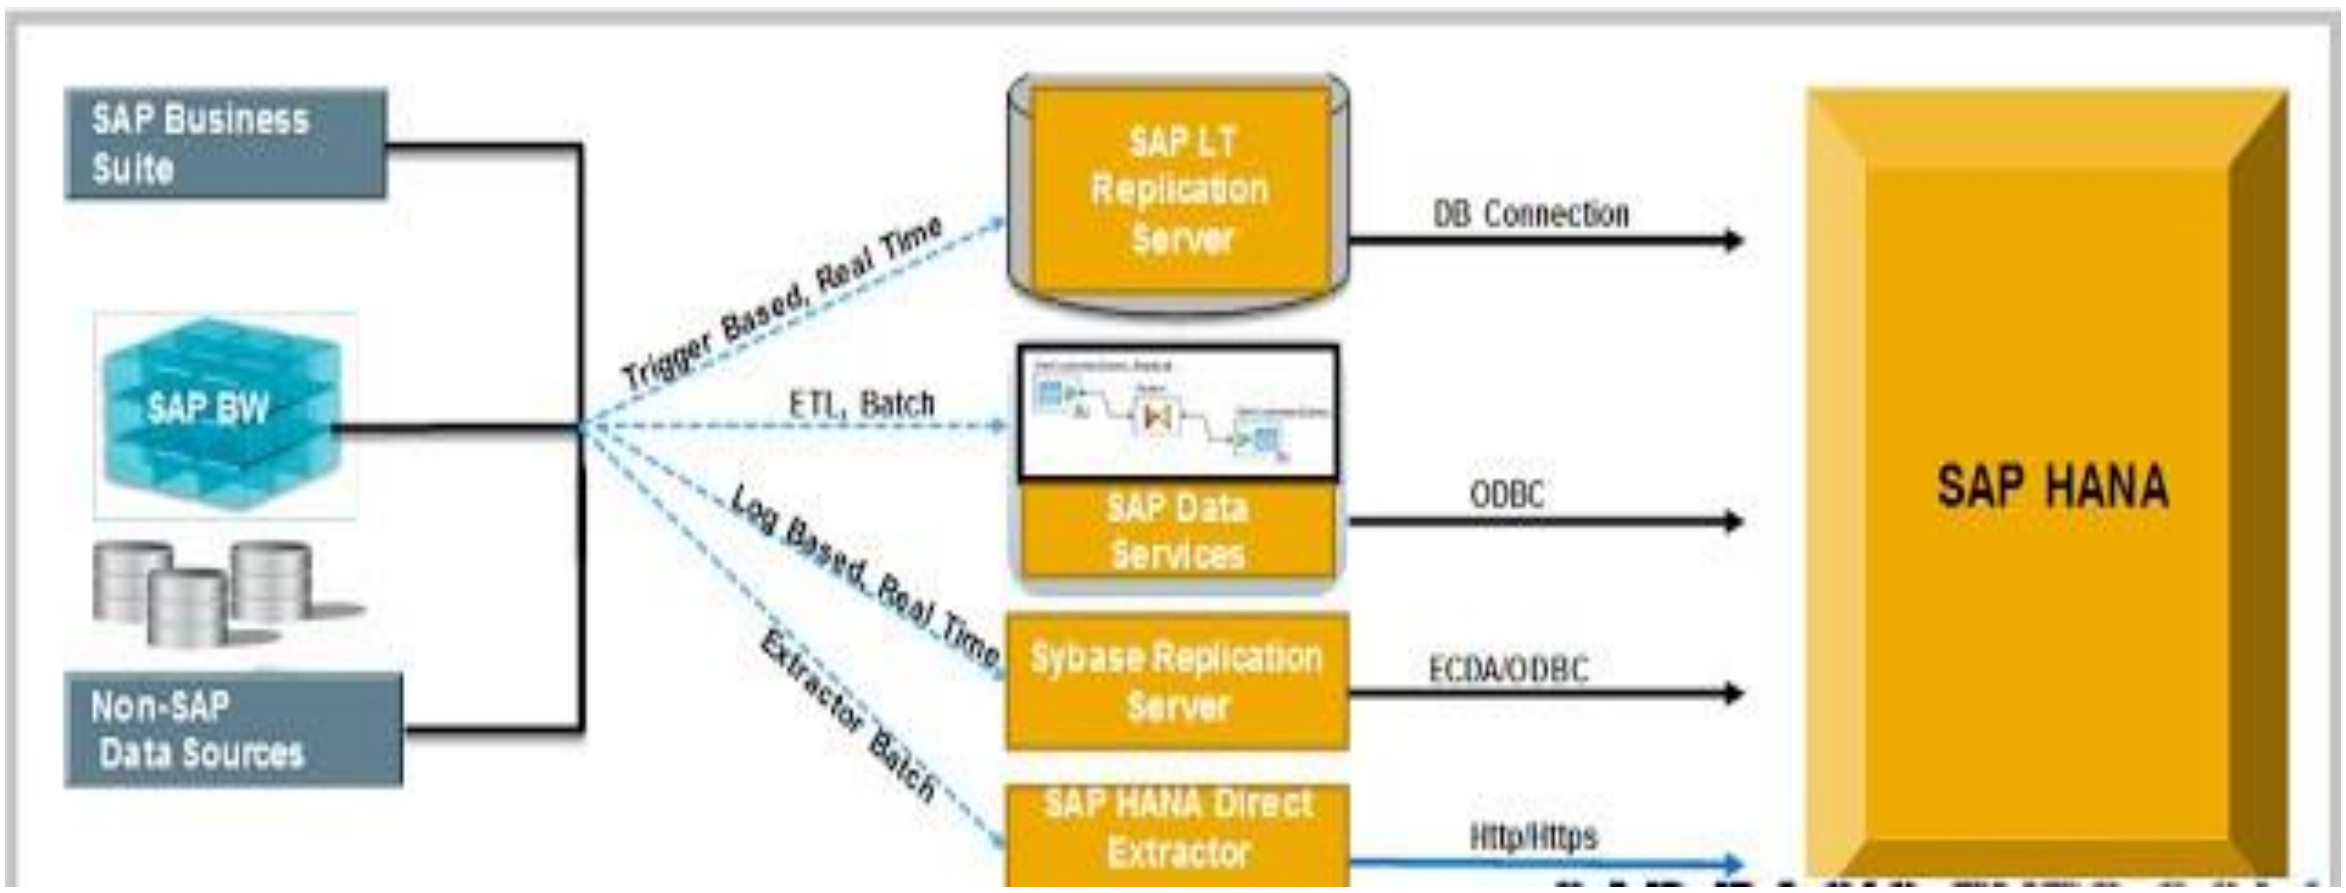

Data provisioning is the process of loading data into SAP HANA. Depending on the requirement scenario, the loading can be done once or on regular basis(real time or scheduled).

HANA Spatial- Usage

Spatial Data Types and Sources

A typical data flow

Spatial data is sourced from many systems, and different tools are used to replicate the data. Quite often, data is imported into SAP HANA via files.

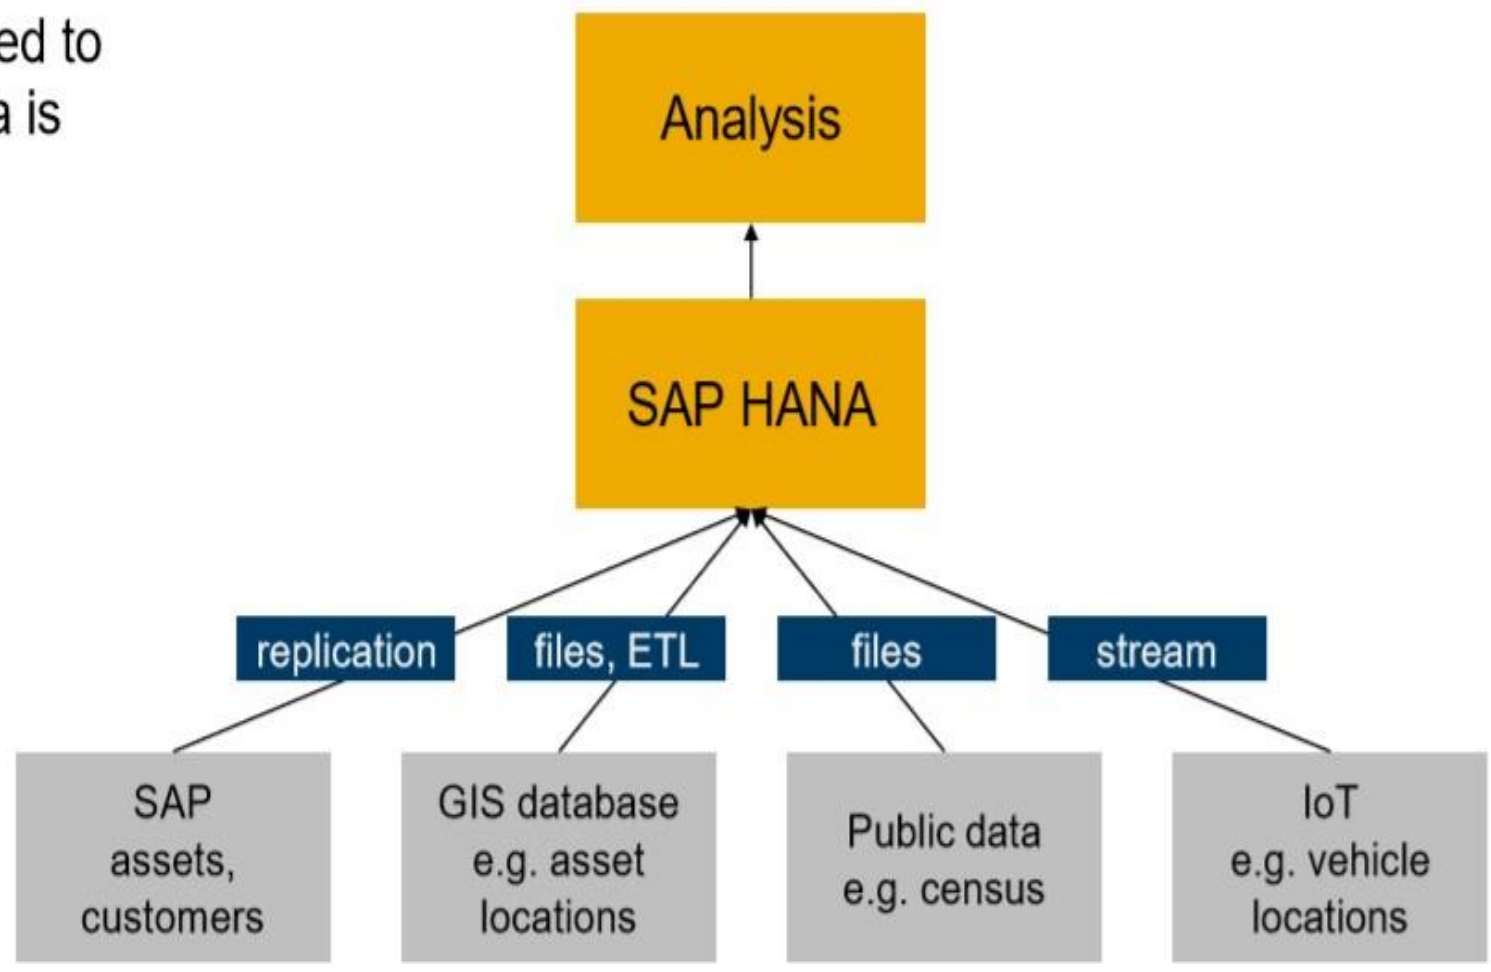

Data classification

- Hot- Data that resides in your HANA Memory
- Warm- Data that resides in your HANA Disk
- Cold - Data that resides out of your HANA

SAP S/4 HANA - Simple Logistics Training

By Praveen

www.praveensd.com

What is temperature data :

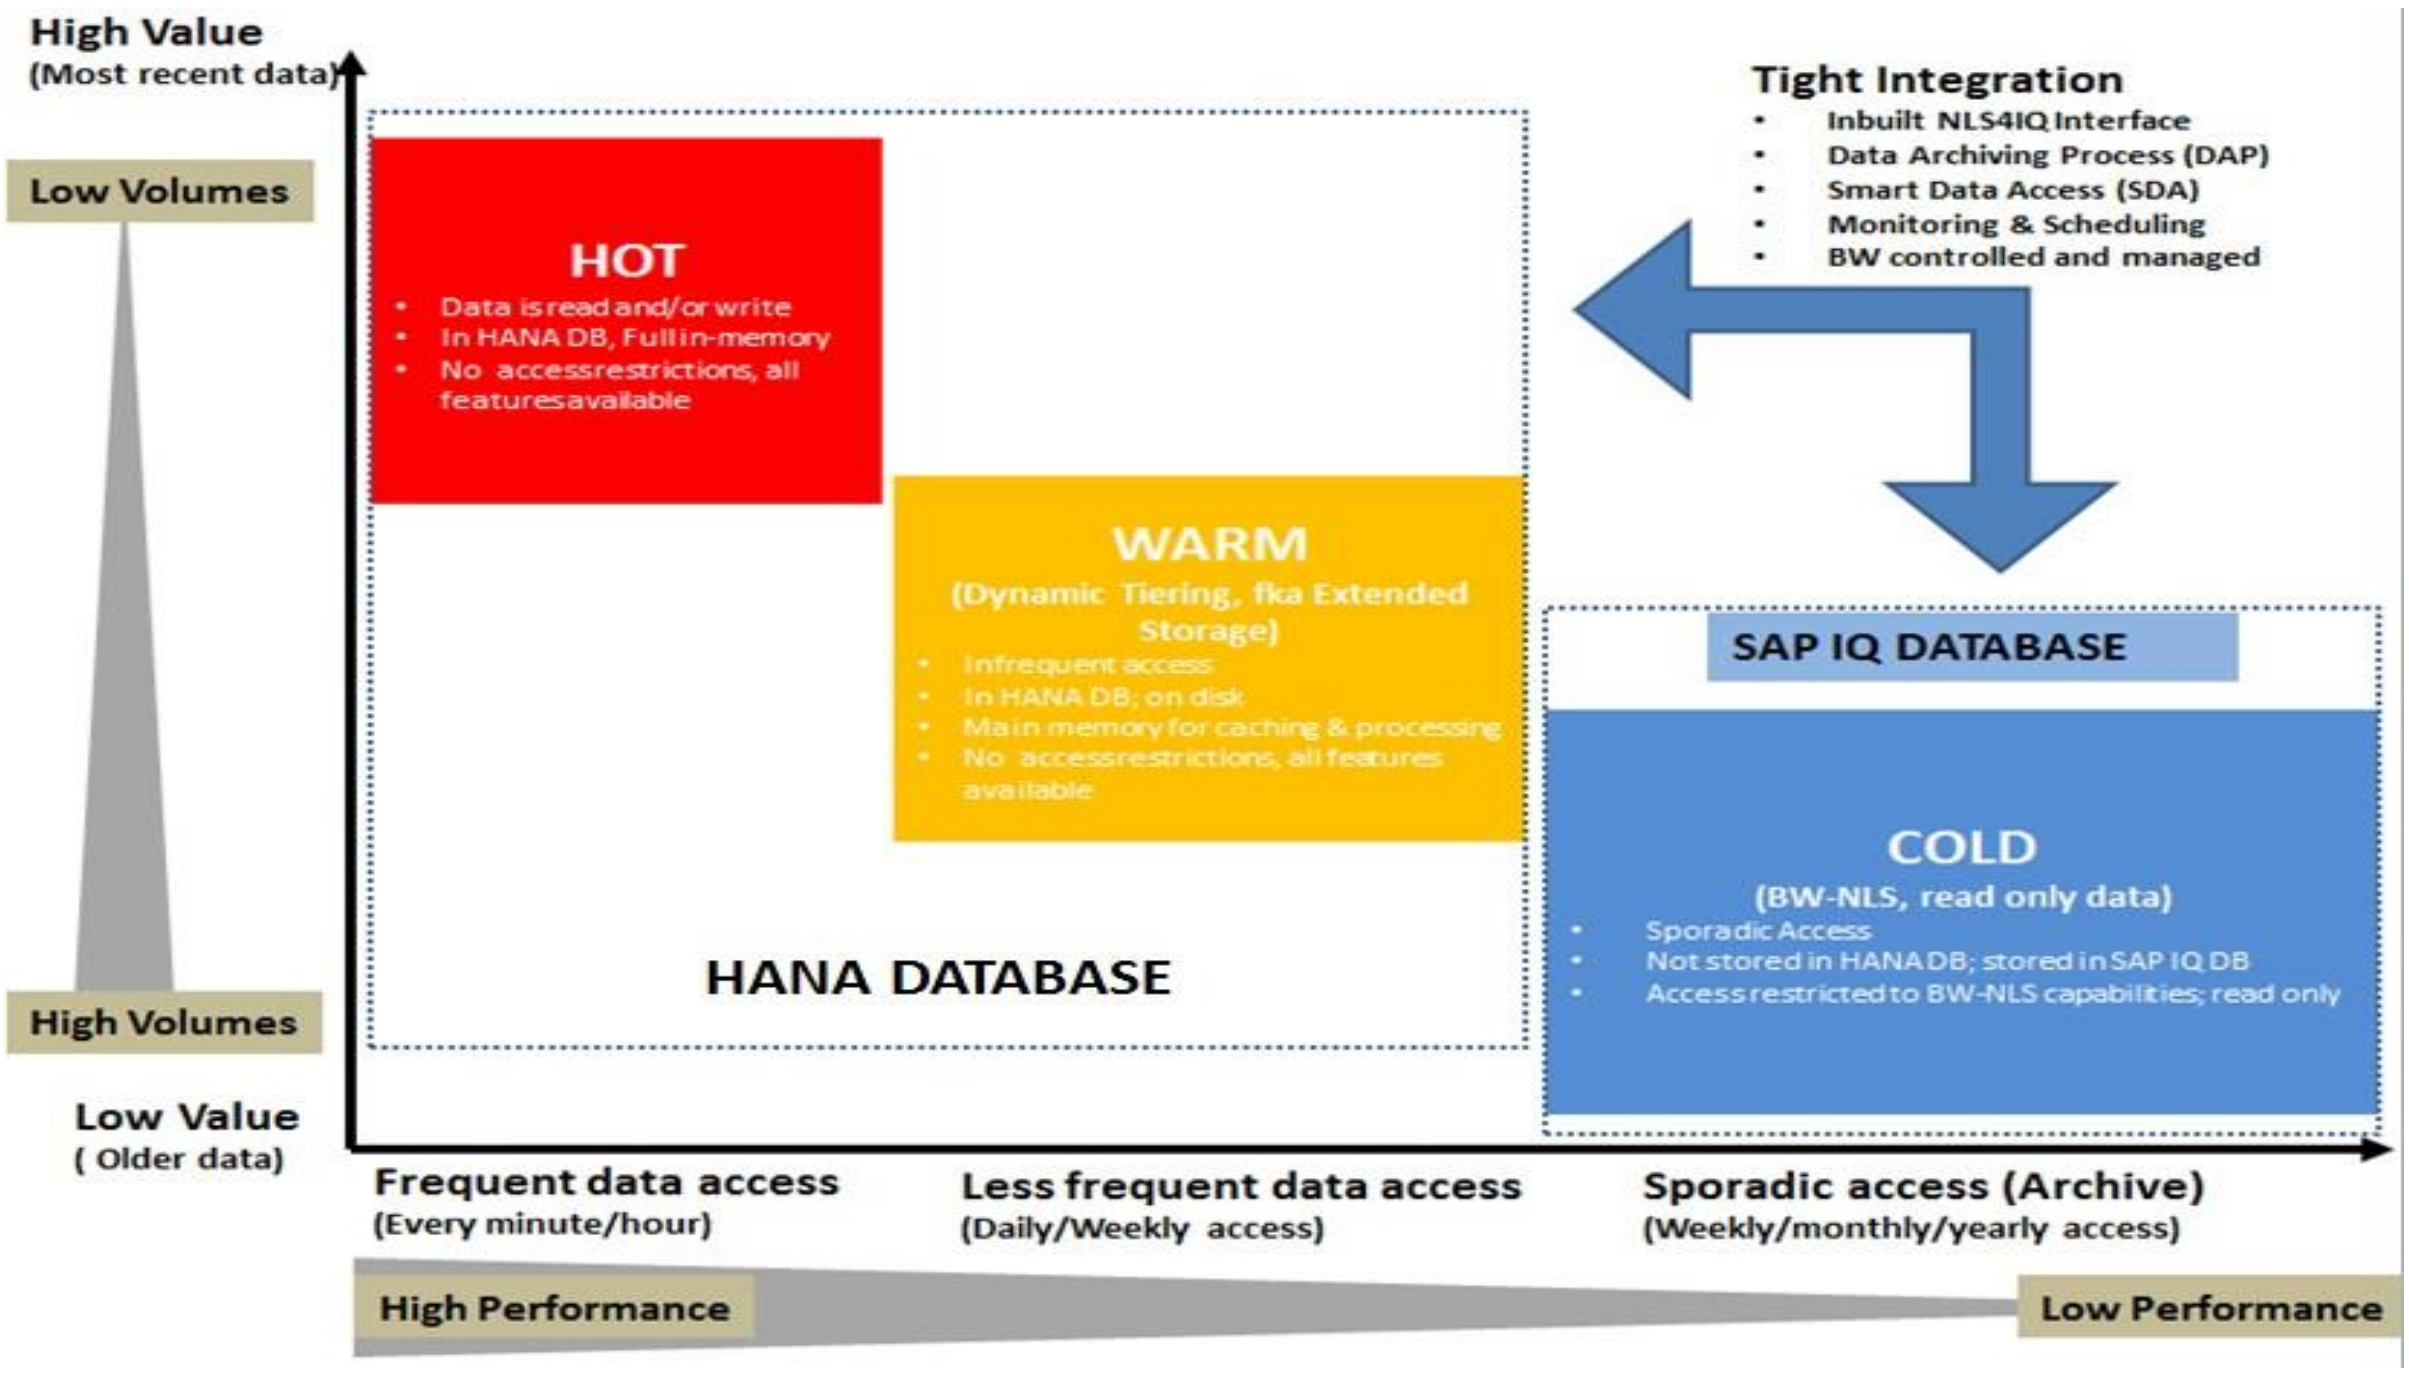

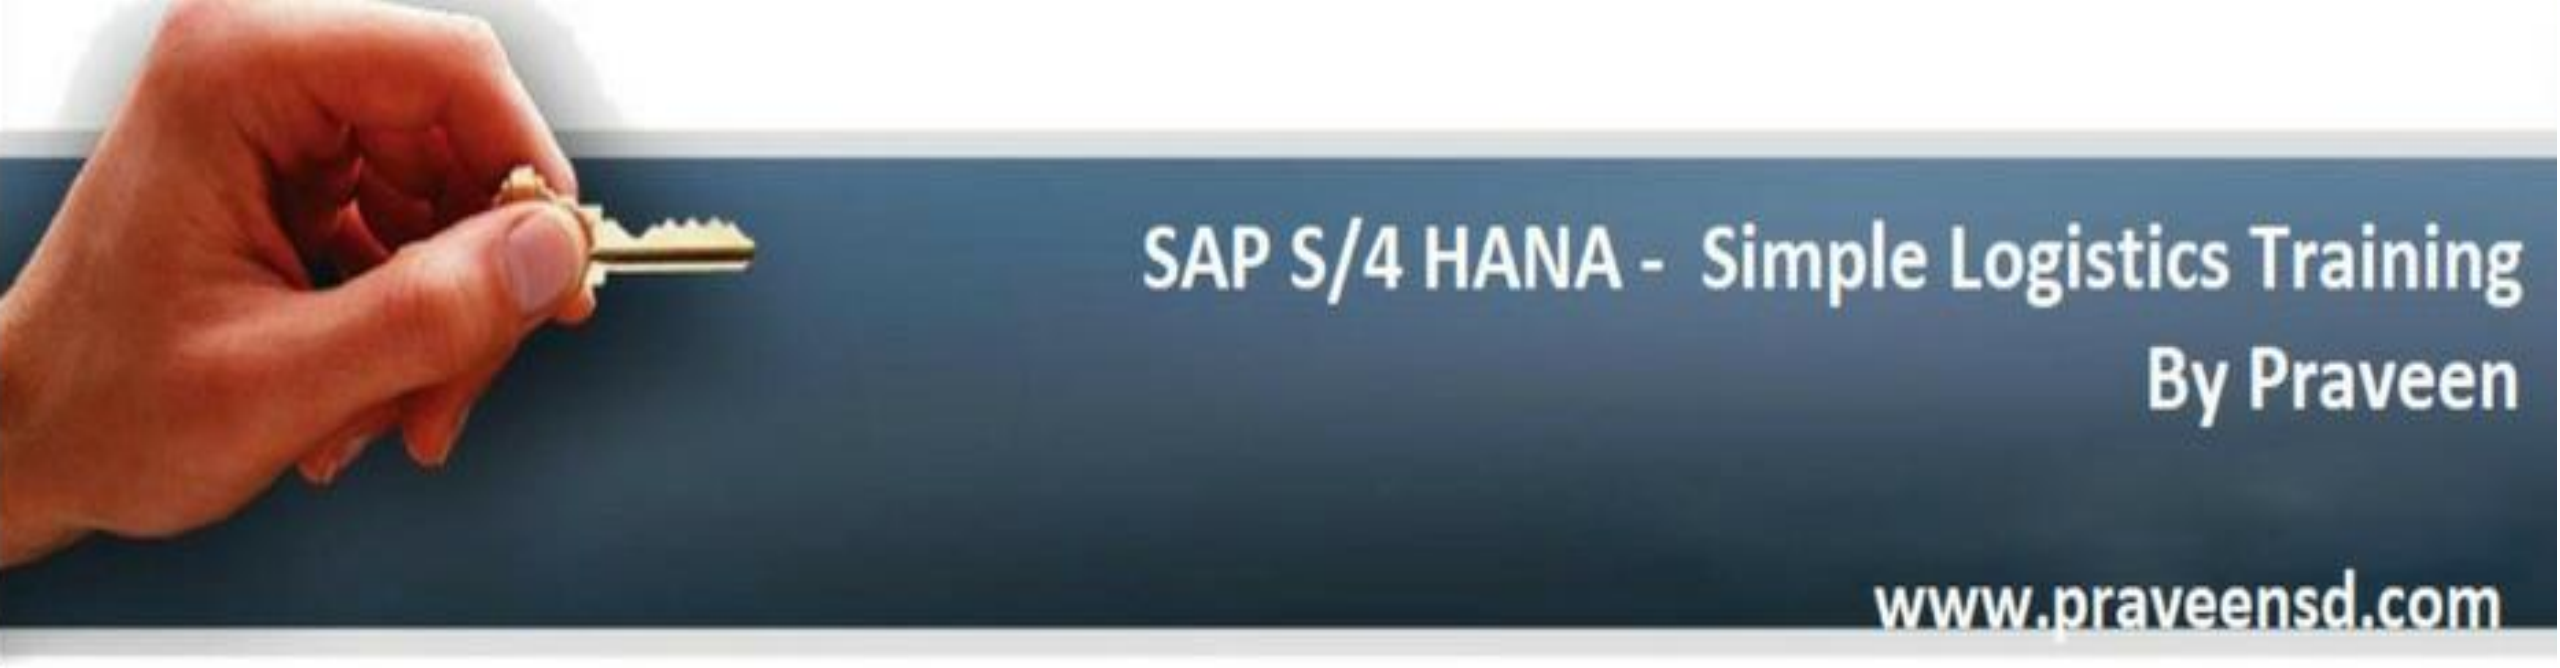

HOT DATA:

- This is the area where 100 % of your PRIMARY IMAGE DATA is in the HANA in- memory space (RAM) and is instantly available for all operations.
- In the BW world, this is typically the Info Cubes and Standard DSOs as they constitute the reporting and harmonization (EDW) areas respectively as show below. They are very frequently accessed for reporting and harmonization purposes and hence is the ideal candidates for being fully in-memory and to fully benefit from the HANA capabilities.

COLD DATA

- This is the area where 100 % of your PRIMARY IMAGE DATA is in a SECONDARY DATABASE (ON DISK) and the response is slightly slower than HANA but still offers reasonably fast READ ONLY access to data for reporting purposes, as if they were in one database.

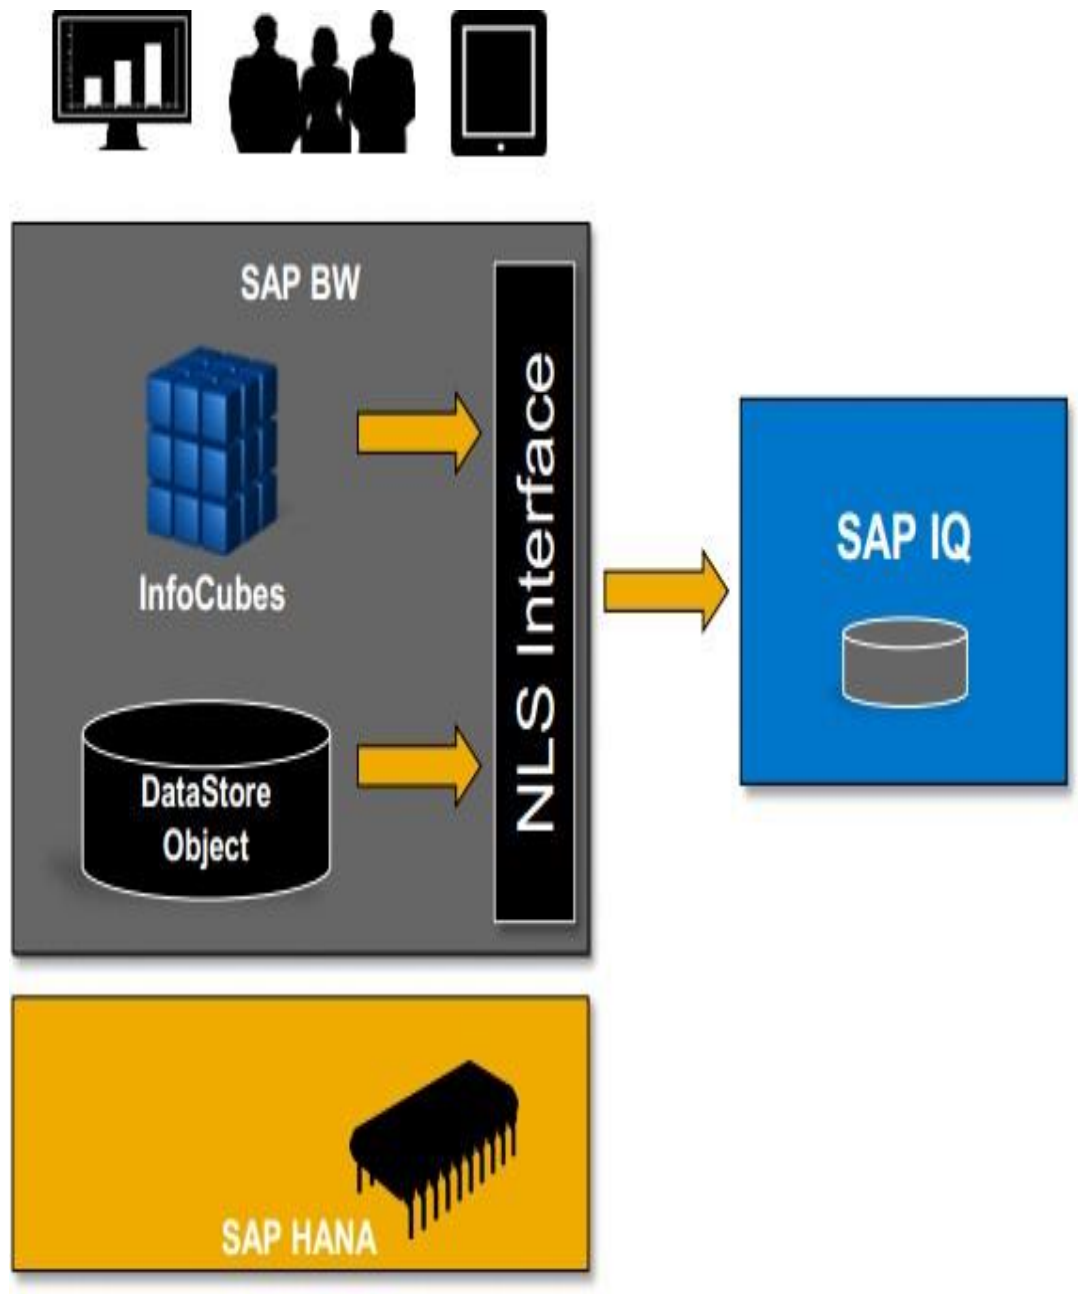

WARM DATA

- This is the area where the PRIMARY IMAGE DATA is in the DISK storage of HANA Database instance, but is always available on request. Using this you can manage your LESS RECENT DATA and LESS FREQUENT DATA more efficiently within the HANA database such that data is instantly available for READ, WRITE, UPDATE etc (all operations), but still offers the lower TCO benefit

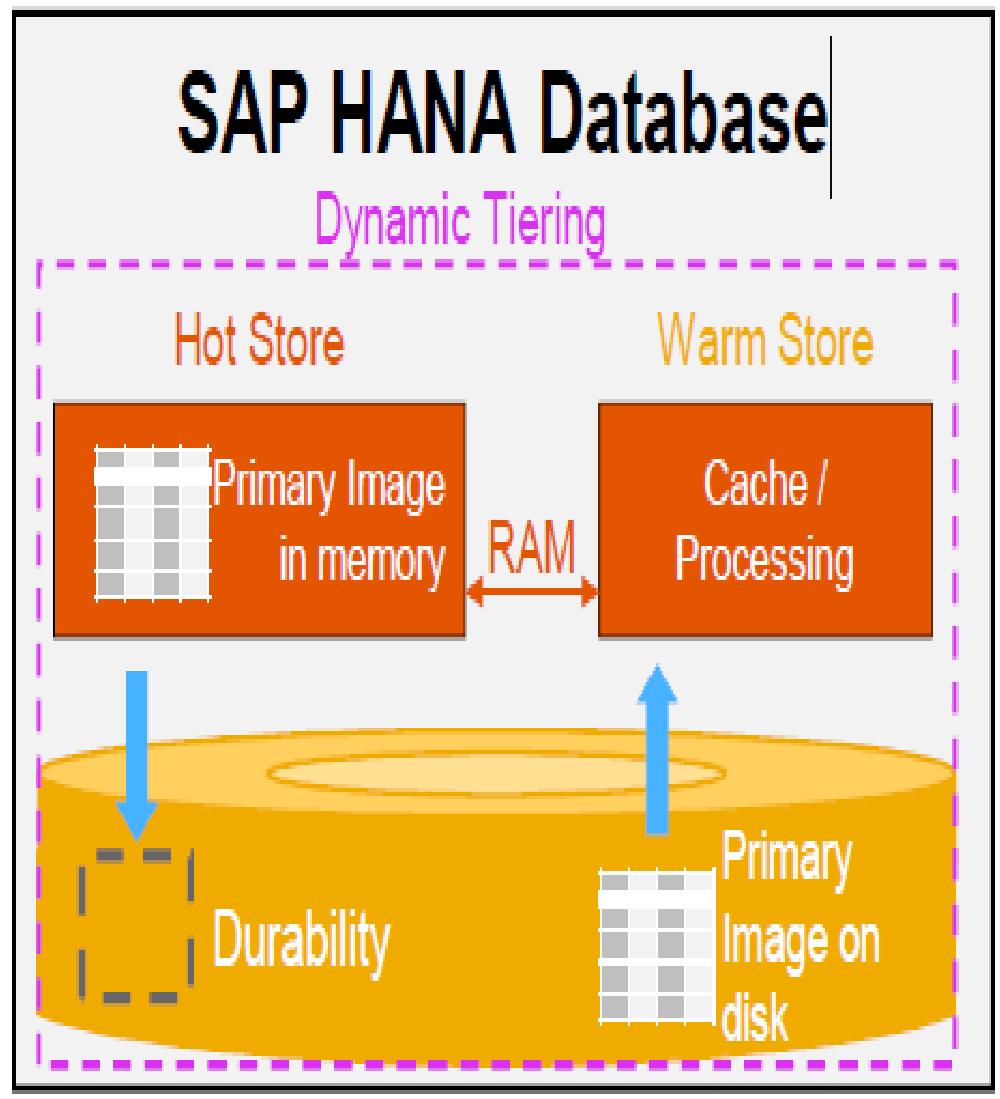

Storage of Temperature Data

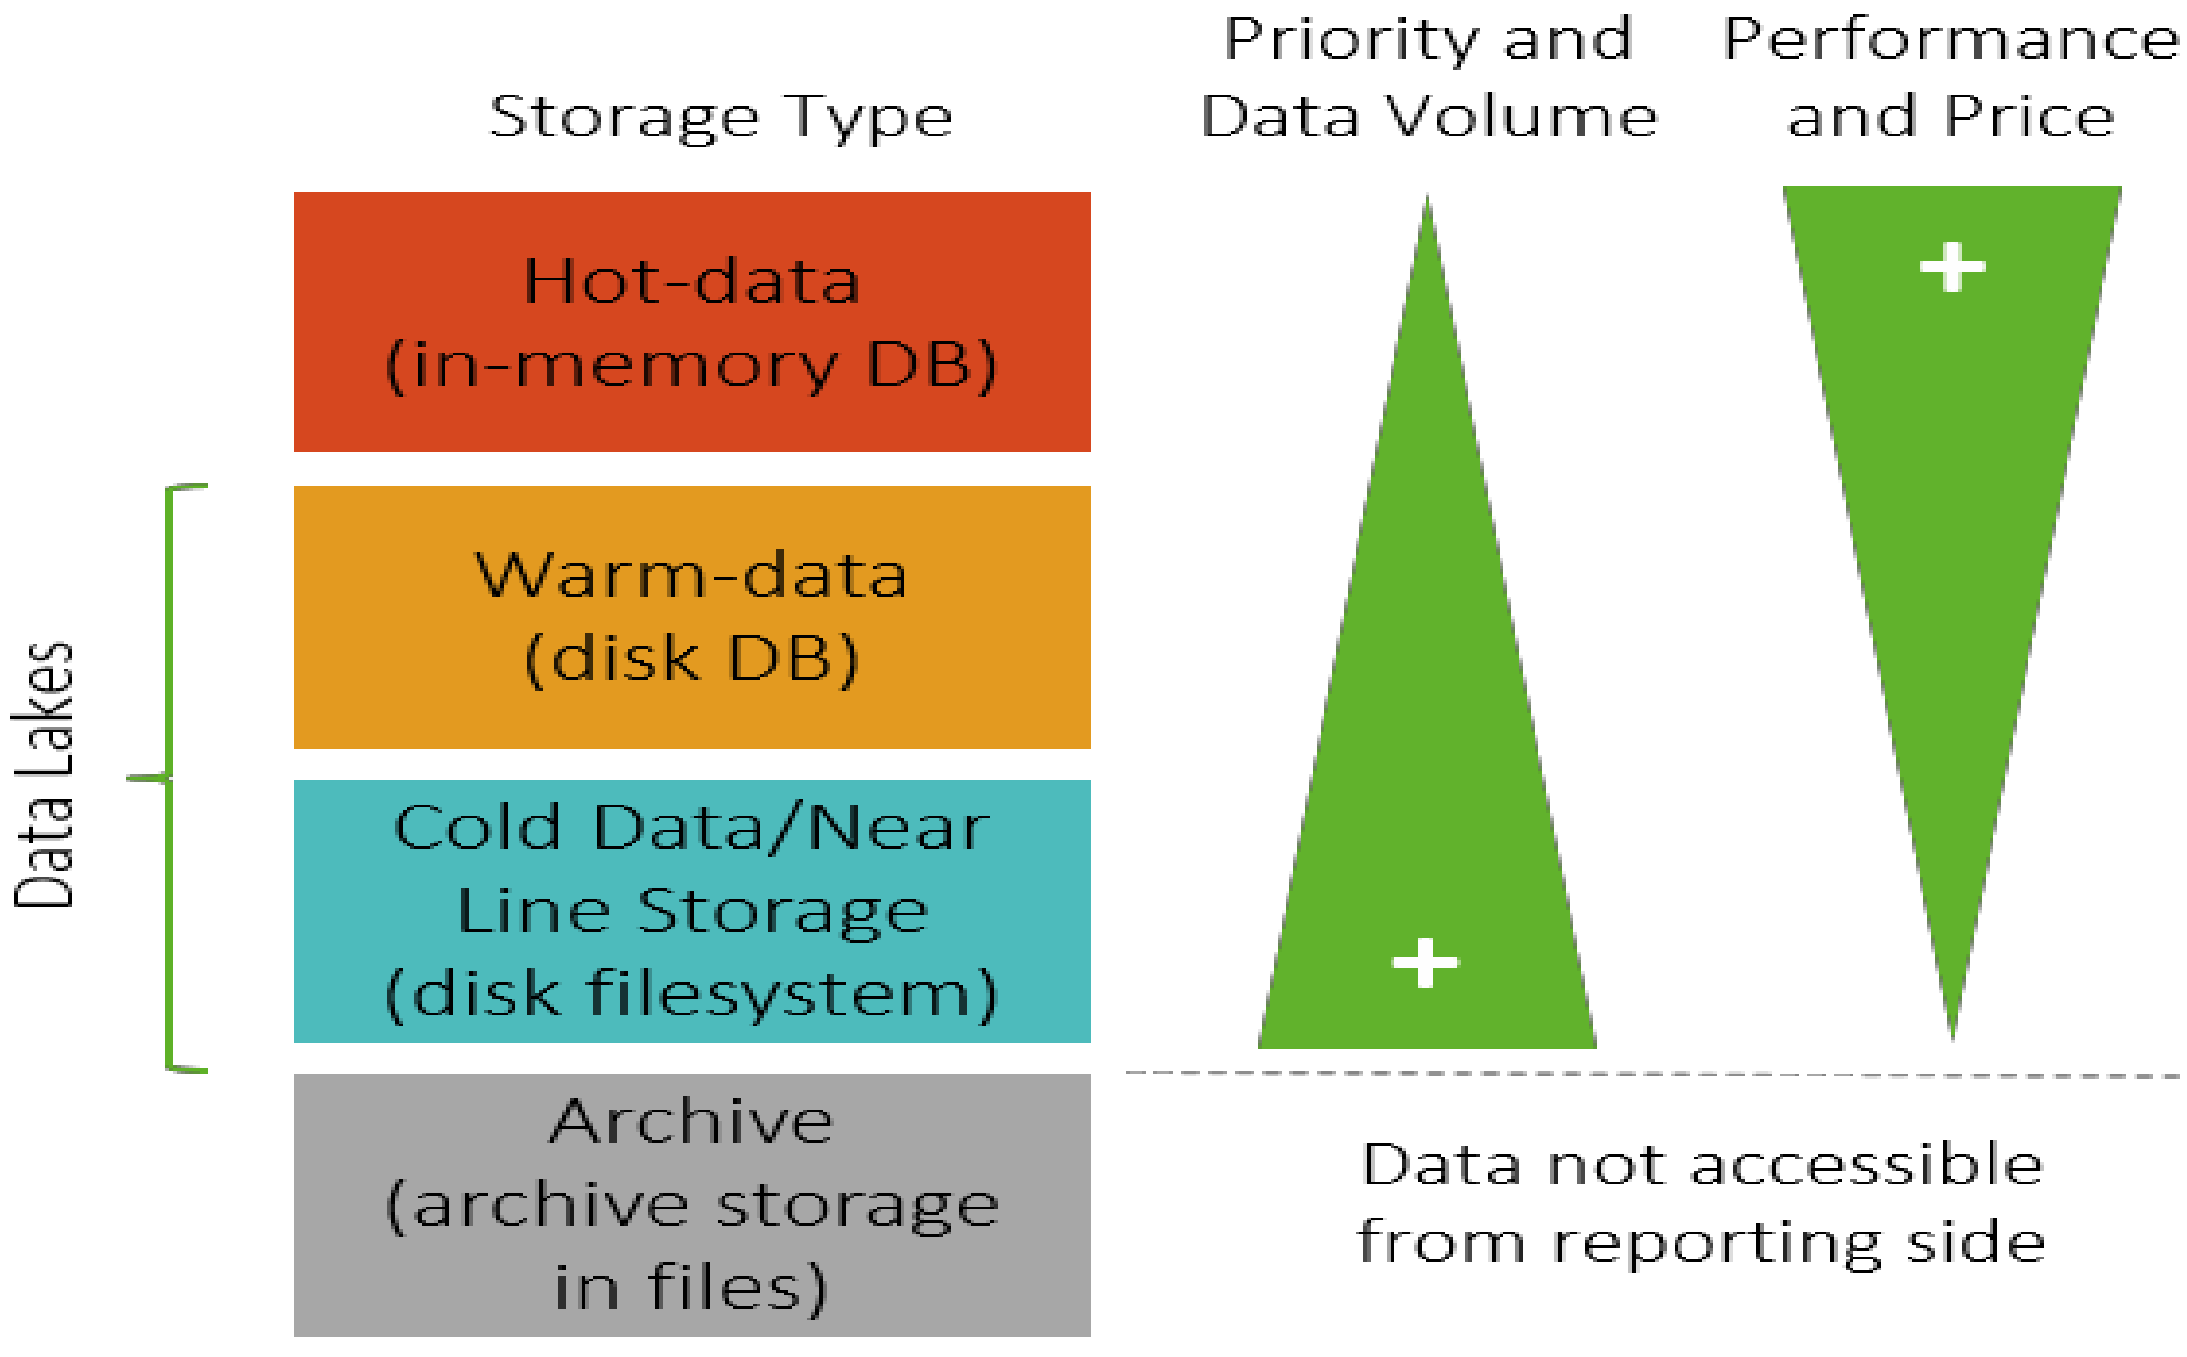

How HANA is optimized both for read and write

- Read write optimization is achieved in HANA by the concept called Delta Write

What is SOH?

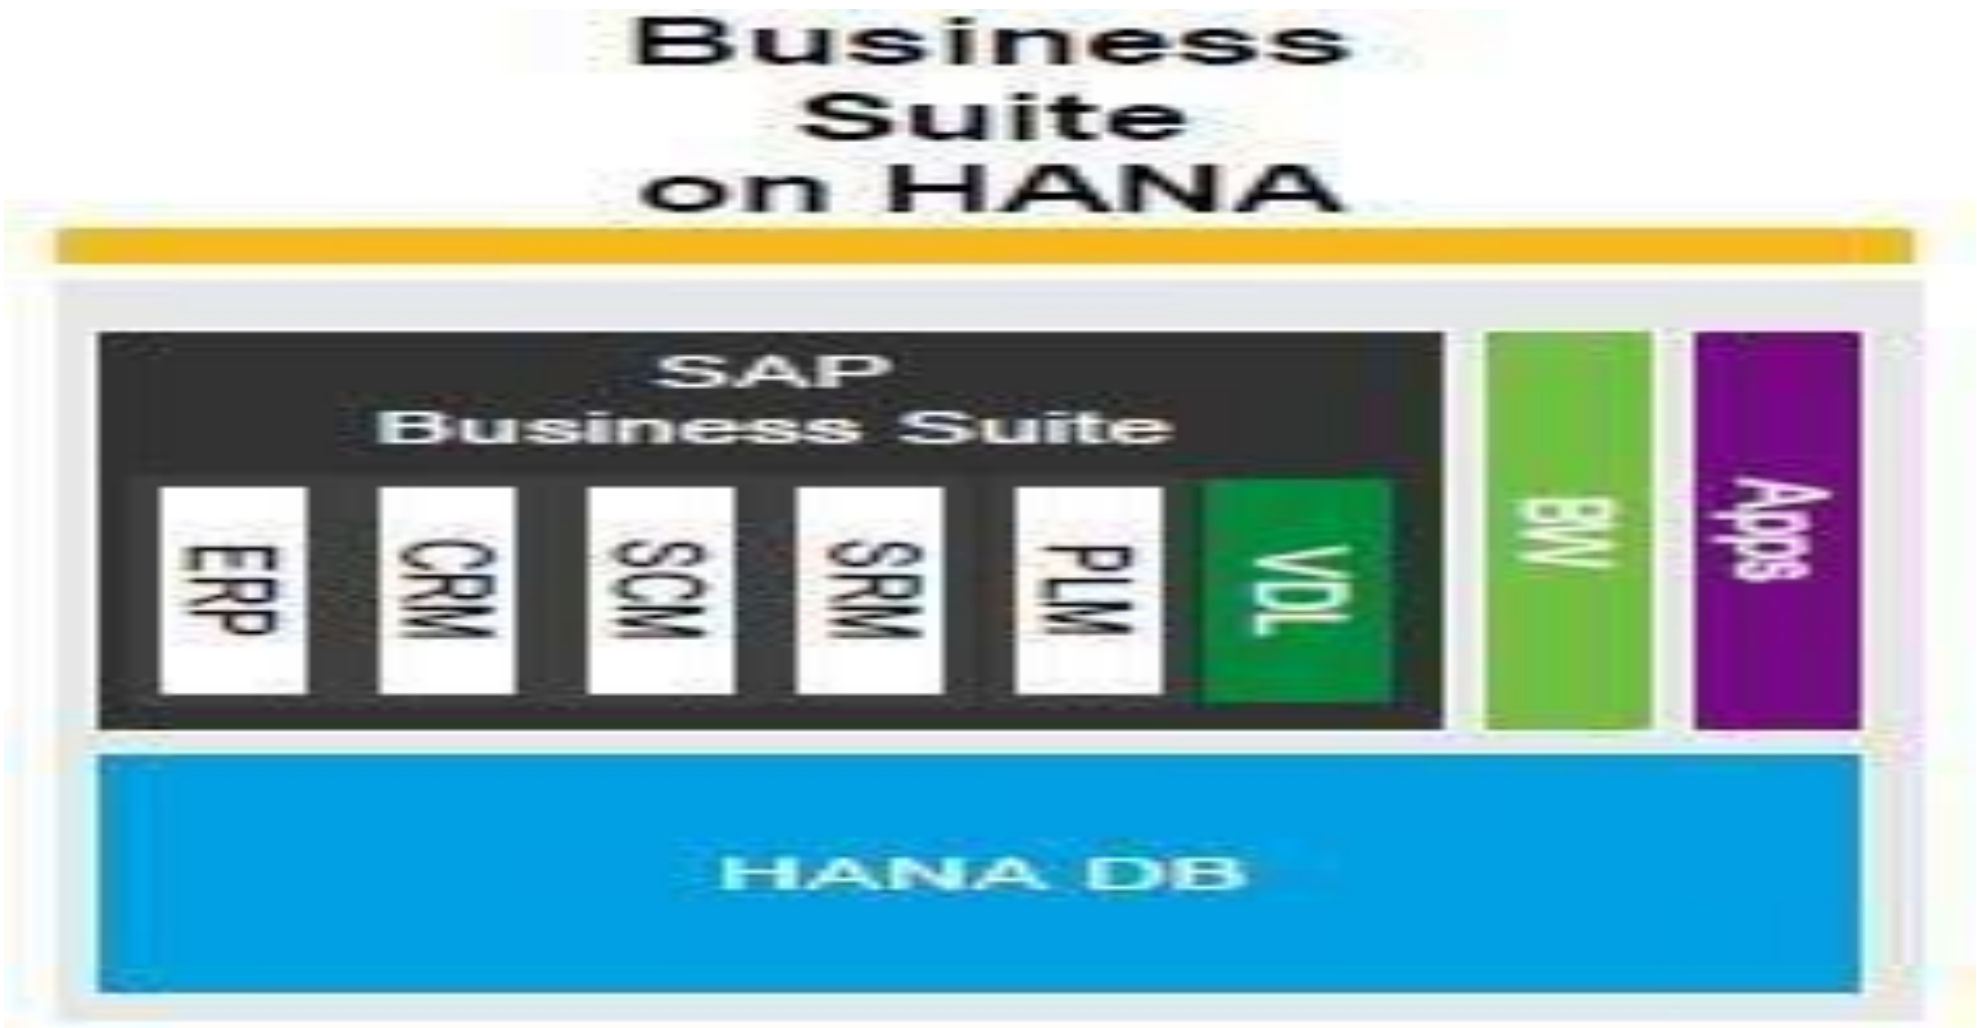

- SOH : Suite on HANA
- Here database is HANA where as application is ERP Suite

HANA live

- SAP HANA Live (previously known as SHAF – SAP HANA Analytic Foundation) is solution for real-time reporting on HANA. It is a separate package that comes with predefined SAP HANA content across the SAP Business Suite.
- SAP HANA Live provides SAP-delivered content (similar in concept like SAP BW content), in form of SAP HANA calculation views for real-time operational reporting. The calculation views spans across majority of ECC modules (FI, CO, MM, PP, SD, PS, CRM, GTS, AM and GRC). The content is represented as a VDM - virtual data model, which is based on the transactional and master data tables of the SAP Business Suite.

Currently more than 2000 views are delivered in HANA Live Package.

HANA Fuzzy search

- **Full Text Search**
- **Full Text Indexing**
- **Fuzzy Search**

- Fuzzy search is the technique of finding strings that match a pattern approximately (rather than exactly).It is a type of search that will find matches even when users misspell wordsor enter in only partial words for the search.

  A Real World Example:

  If a user types "SAP HANO " into Yahoo or Google (both of which use fuzzy matching), a list of hits is returned along with the question, "Did you mean "SAP HANA ".

Top Business Use-cases of textAnalysis:

- **Brand/ Product/ Reputation Management** Market research and social media monitoring, i.e. what people are saying about my brand or products
- **Voice of the Customer/ Customer Experience Management** Do I need to step in and offer customer service?
- How many people recommend my brand vs. advocate against it?
- **Search, Information Access, or Questions Answering** Which bloggers are negative towards USA Policies?
- Which of the hotels on India get great reviews for the room service?
- **Competitive Intelligence** What competing products are people considering and why?
- Are competitors‘ media sped generating purchase intent?

SAP S/4 HANA - Simple Logistics Training

By Praveen

www.praveensd.com

HANA modelling Introduction:

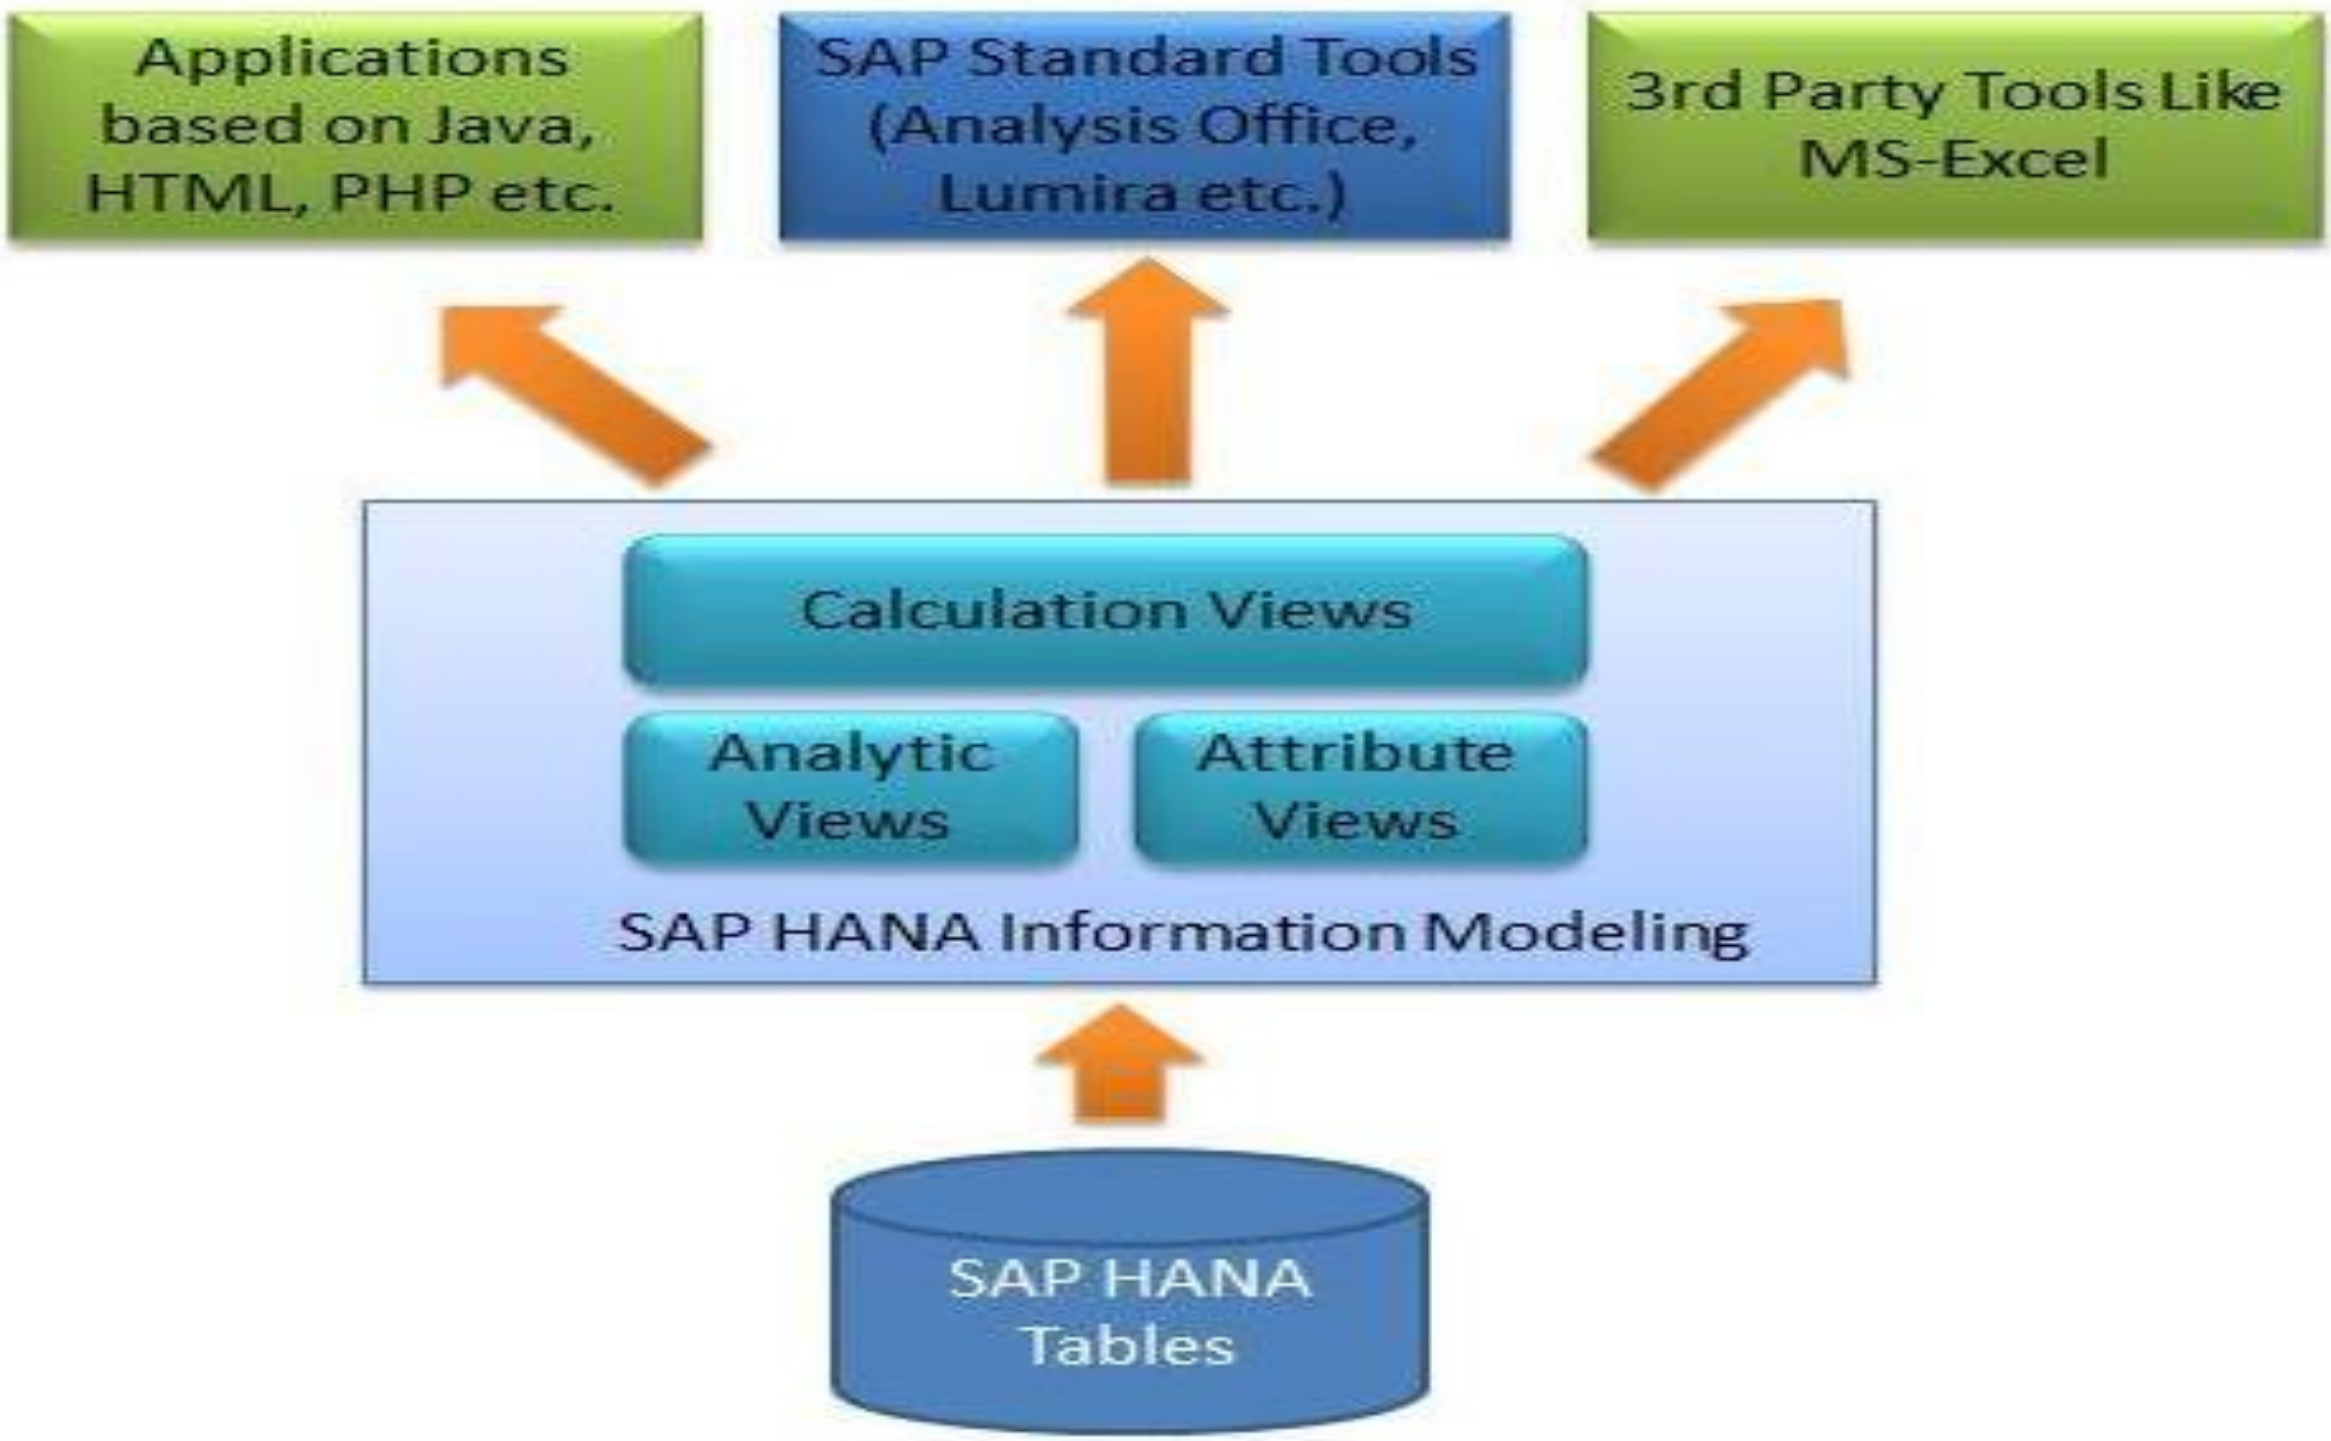

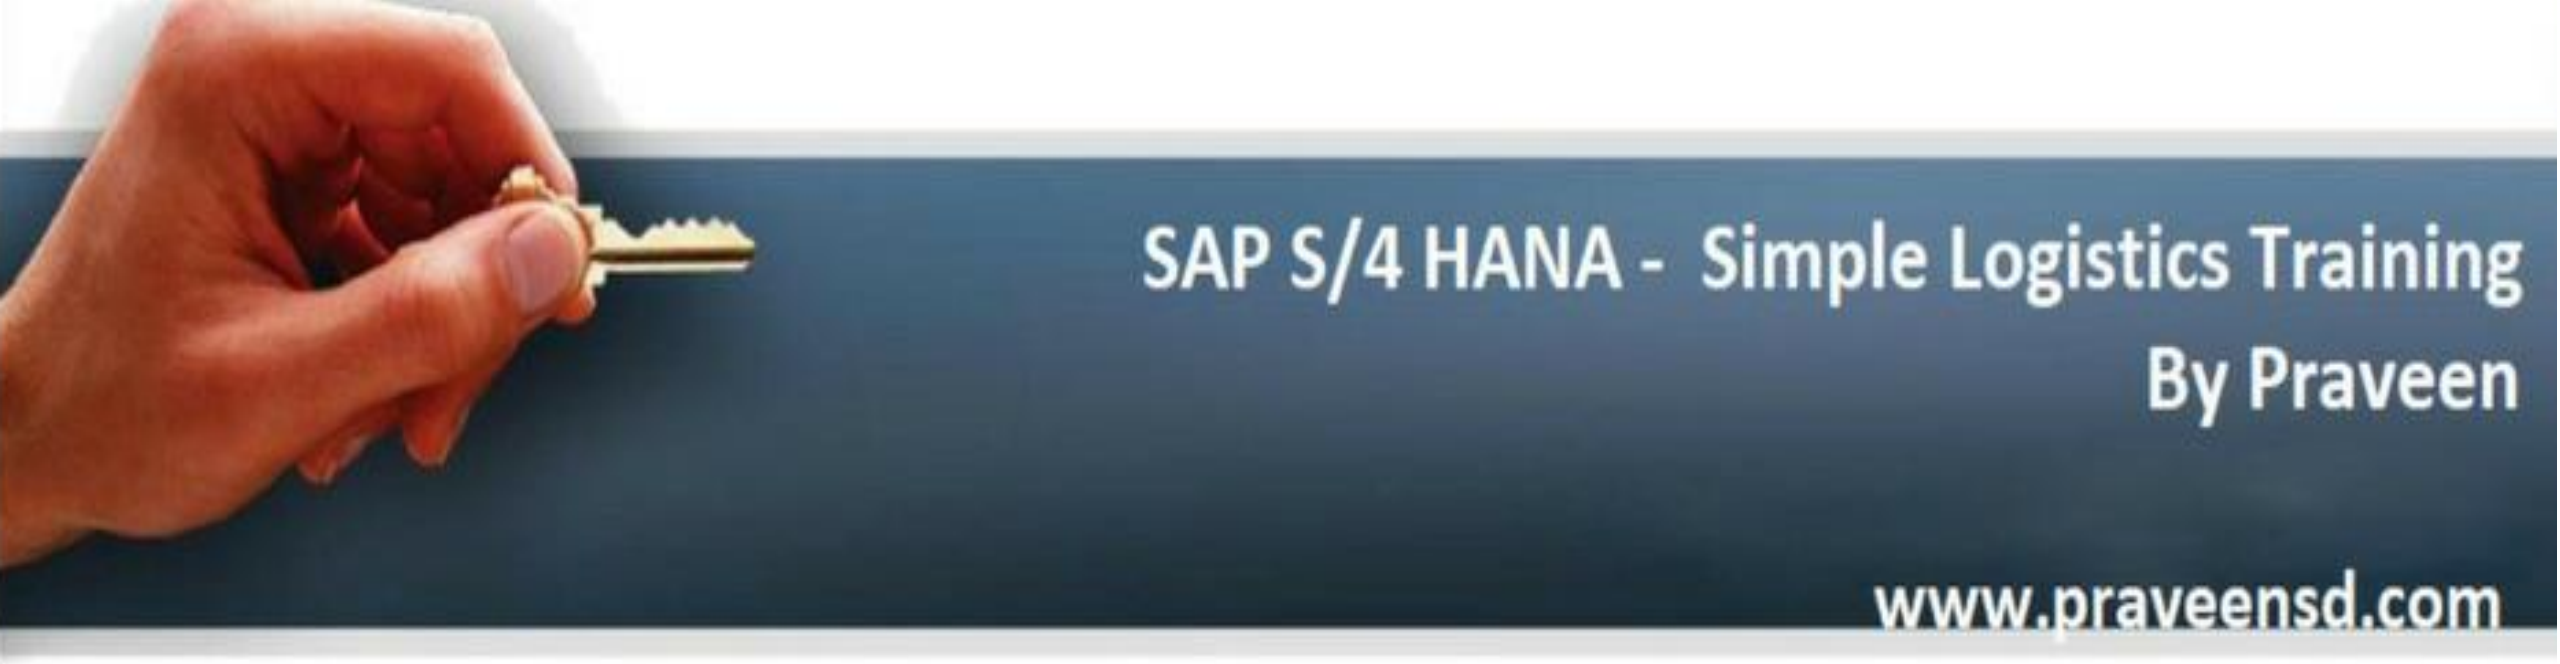

Hana Frontend Tools

SAP HANA Information Models – Client Tools Access

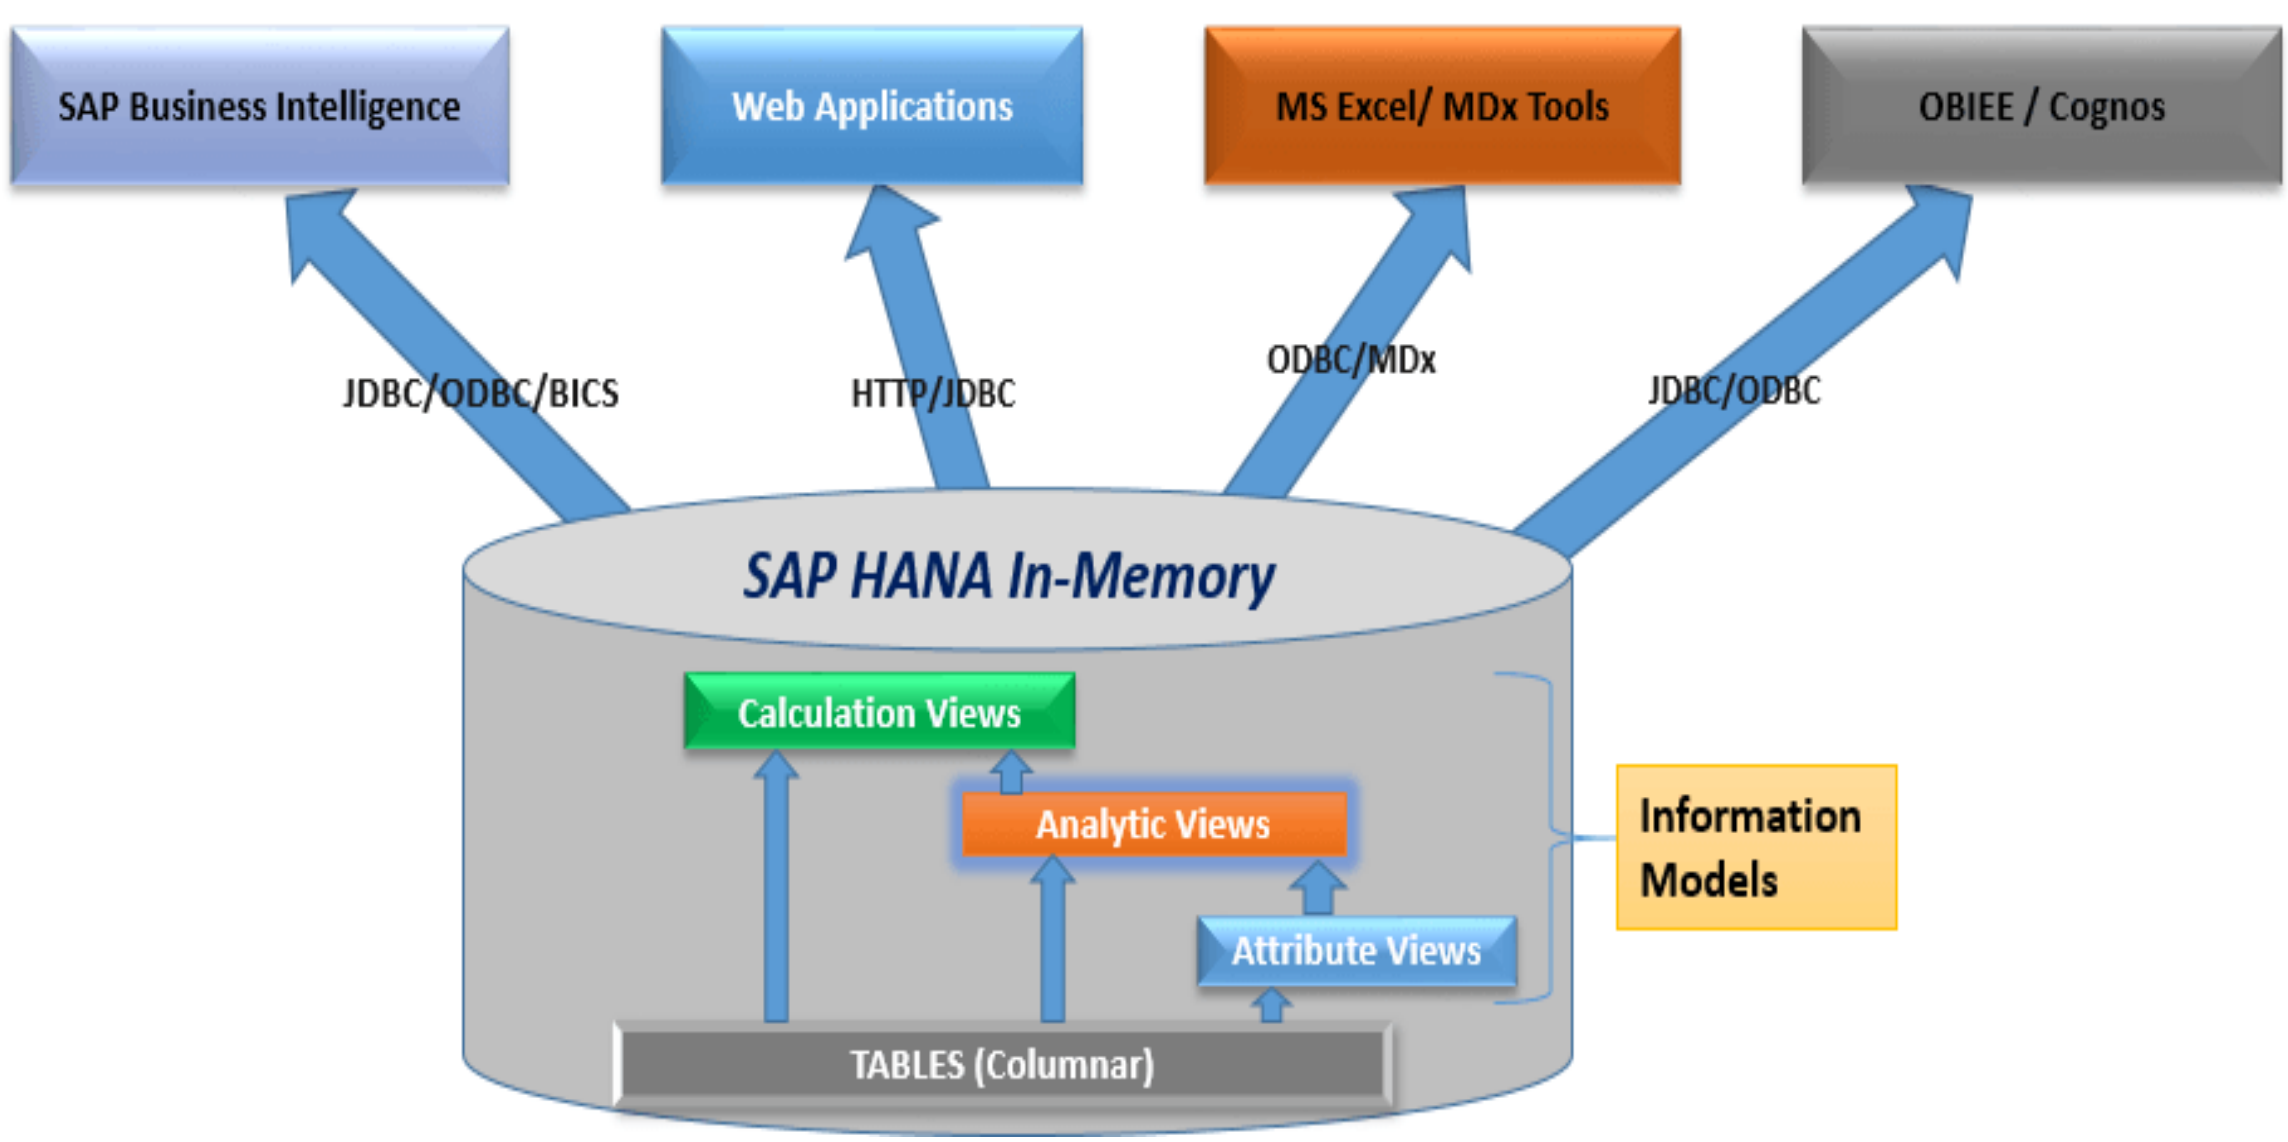

Server Sizing

Sizing HANA

- Important SAP Notes about sizing HANA:
- 1514966 – Sizing HANA
- 1704499 – System Measurement for License Audit
- 1637145 – Sizing BW on HANA
- Static + Dynamic Ram requirements determine sizing. To do this, determine uncompressed data volume to be loaded, then apply compression factor, then multiply result by two.
- Only 50% of the total RAM should be used for the in-memory database. The other 50% is needed for temporary objects (for example, intermediate results), the operating system, and application code.
- Disk size for the persistence layer is equal to 4 times RAM
- Disk size for the log files equals 1 times RAM
- CPU equals 300 SAPS per active user. Never to exceed 65% of CPU server load.
- BW on HANA has a quick sizing tool available at: https://service.sap.com/quicksizing
- HANA is available in incremental sizes: XS, S, M, L

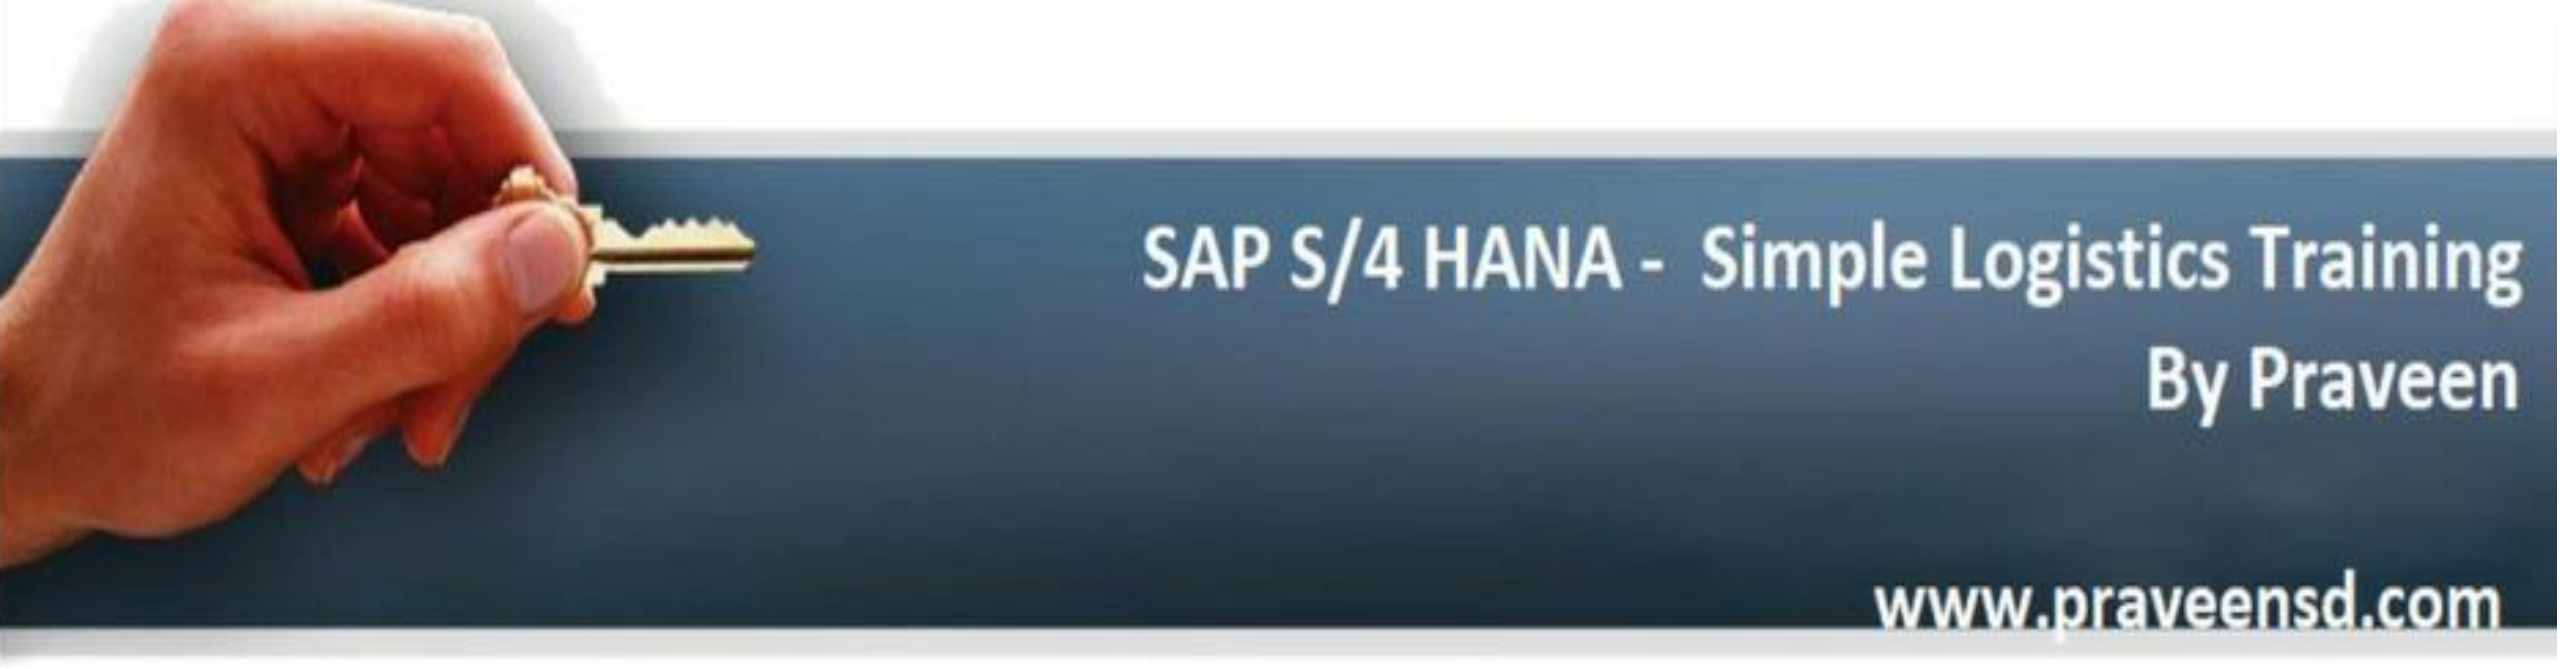

SAP HANA Database:

SAP HANA is an in-memory database:

- It is a combination of hardware and software made to process massive real time data using In-Memory computing.
- It combines row-based, column-based database technology.
- Data now resides in main-memory (RAM) and no longer on a hard disk.
- It's best suited for performing real-time analytics, and developing and deploying real-time applications.

SAP S/4 HANA - Simple Logistics Training

By Praveen

www.praveensd.com

Smart Data Access in SAP HANA:

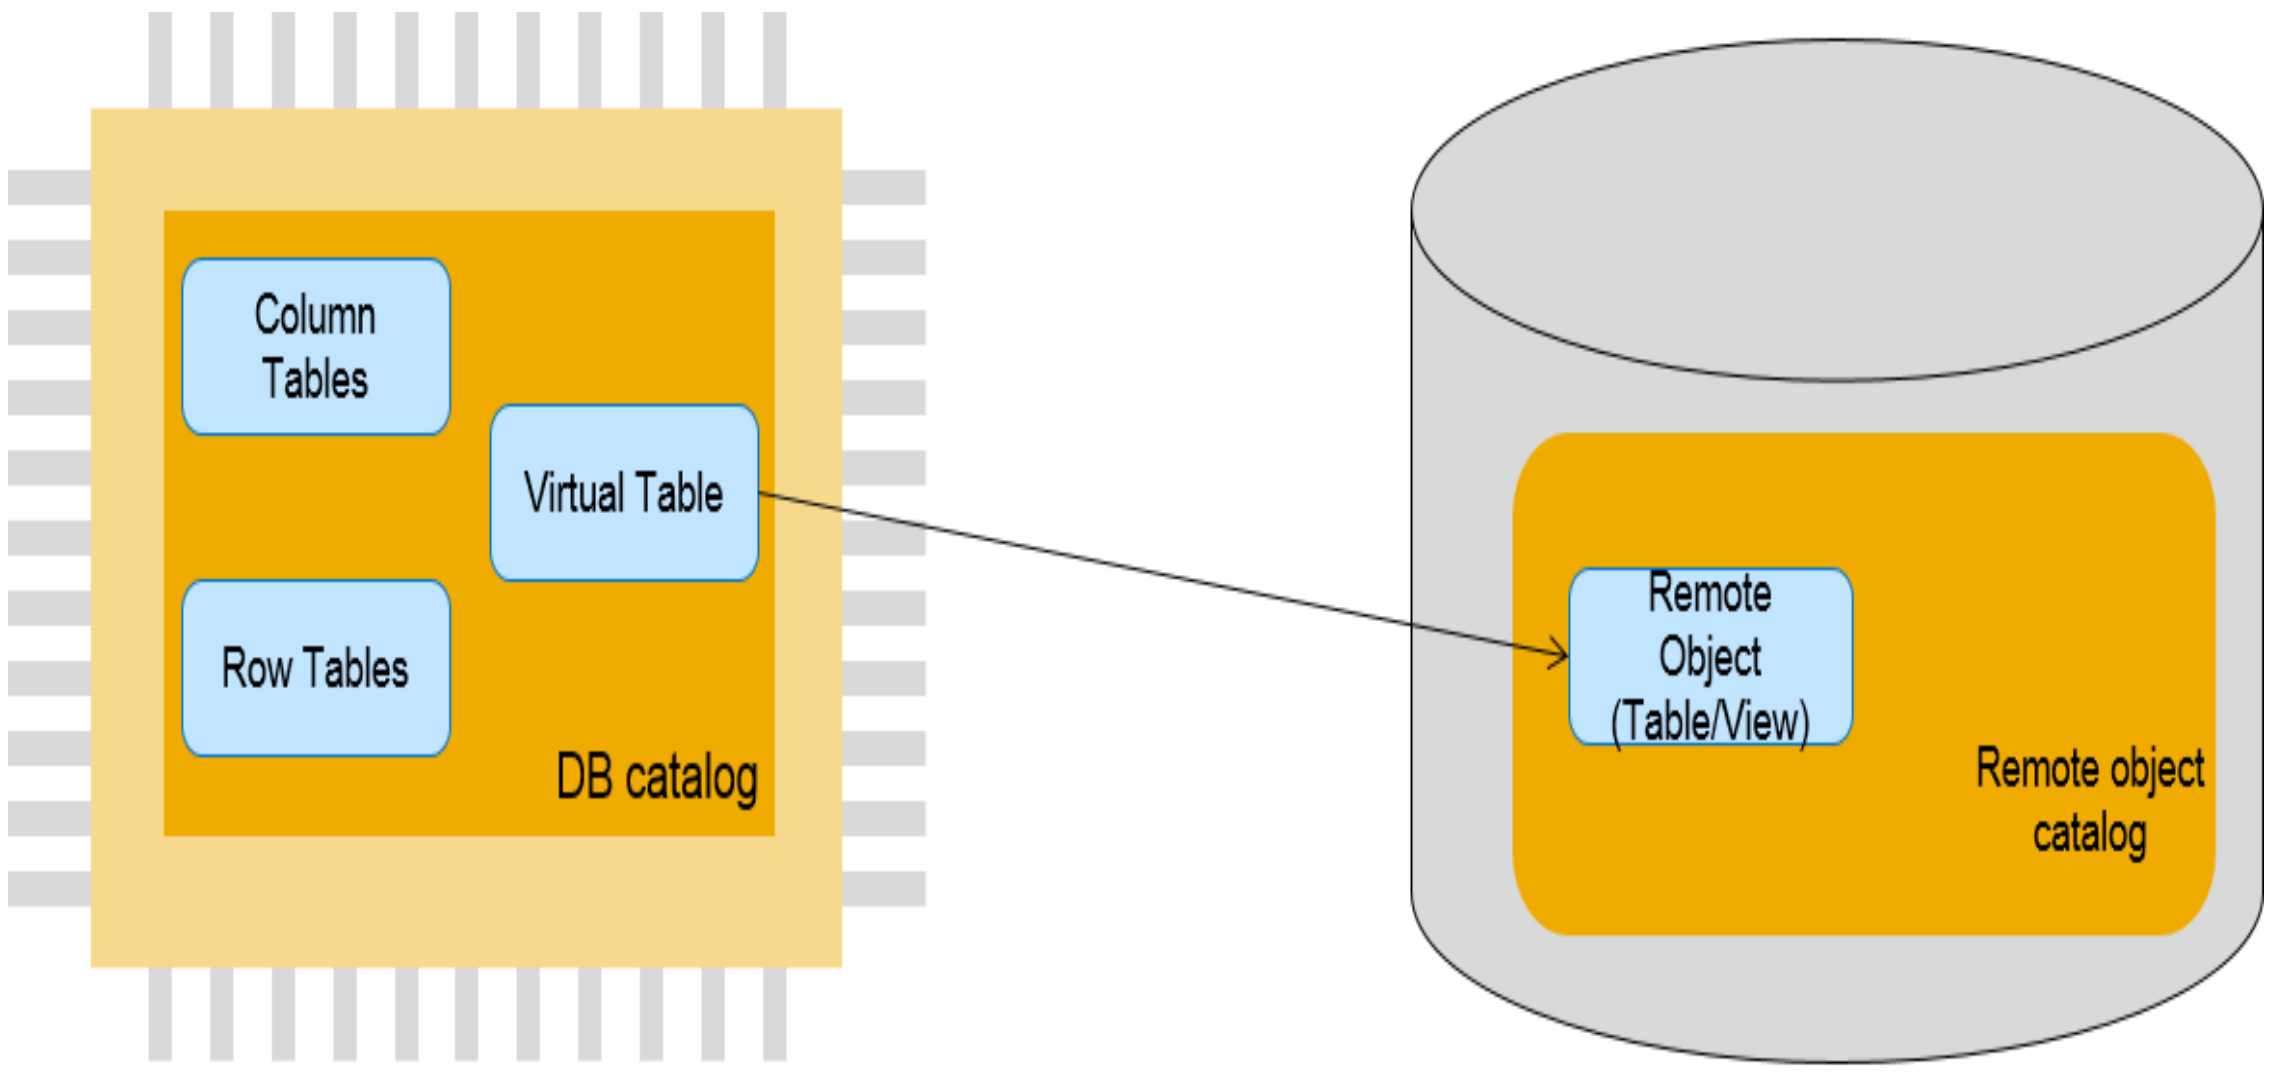

HANA Data Analysis

- Sensor data from equipment is monitored for trends or correlations that indicate a problem, alerting an operator to take immediate action before equipment damage occurs

- A commodity pricing application continuously adjusts quoted prices in response to market conditions – where delays mean either lost business or lost profit.

- IT system events are continuously monitored to watch for patterns that indicate a possible security threat

- User actions on a web site are analysed to determine the best offers to show, not just based on historical data for the user but also considering current context

Server Sizing

- SAP HANA Sizing
- Sizing is a term which is used to determine hardware requirement for SAP HANA system, such as RAM, Hard Disk and CPU, etc.
- The main important sizing component is the Memory, and the second important sizing component is CPU. The third main component is a disk, but sizing is completely dependent on Memory and CPU.
- In SAP HANA implementation, one of the critical tasks is to determine the right size of a server according to business requirement.
- SAP HANA DB differ in sizing with normal DBMS in terms of –
- Main Memory Requirement for SAP HANA (Memory sizing is determined by Metadata and Transaction data in SAP HANA)
- CPU Requirement for SAP HANA (Forecast CPU is Estimated not accurate).
- Disk Space Requirement for SAP HANA (Is calculated for data persistence and for logging data)
- The Application server CPU and application server memory remain unchanged.
- For sizing calculation SAP has provided various guidelines and method to calculate correct size.

- We can use below method- 1.Sizing using ABAP report.
- 2.Sizing using DB Script. 3.Sizing using Quick size To**ol.**

Quick sizer

Average:	Starttime:	9		Endtime:	18
Peak:	Starttime:	12		Endtime:	13

Valid from 23.02.2016 to 30.05.2016

Table 41: Concurrent Users - Standard Sizing © SAP SE

Sizing Element	Timeinter TI: Y/P/H/S	Average/P A/P	No. of con Low activi	No. of con Medium a	No. of con High activ	No. of mo Mon.	Checks fo Arch.	ID for mul ID	Fiori Apps FIORI	A short te Short text
ALM-USER	S	A								
MM-USER	S	A								
PP-USER	S	A								
PS-USER	S	A								
QM-USER		A								

Table 42: Throughput - Standard Sizing © SAP SE

Sizing Element	Timeinter TI: Y/P/H/	Average/P A/P	Sizing obj Objects	Avg. num Items	No. of cha % chg.	No. of dis % dsp.	No. of mo Mon.	Checks fo Arch.	Start of pr S.t.	End of pro E.t.	ID for mul ID	Fiori Apps FIORI	A short te Short text
ALM-PM	Y	A							9	18			
ALM-PM	P	P							12	13			
MM-IM	Y	A							9	18			
MM-IM	P	P							12	13			
MM-PUR	Y	A							9	18			
MM-PUR	P	P							12	13			
PP-CONF	Y	A							9	18			
PP-CONF	P	P							12	13			
PP-REM	Y	A							9	18			
PP-REM	P	P							12	13			
PP-SFC	Y	A							9	18			
PP-SFC	P	P							12	13			
QM-IM	Y	A							9	18			
QM-IM	P	P							12	13			

Table 43: Throughput - Project System © SAP SE

Sizing Element	Timeinter TI: Y/P/H/	Average/P A/P	Sizing obj Objects	Structural WBS	Average n Netw.	Average n Act/nw	No. of cha Changes	No. of dis Displays	No. of mo Mon.	Checks fo Arch.	Start of pr S.t.	End of pro E.t.	ID for mul ID	Fiori Apps FIORI	A short te Short text
PS P	Y	A									9	18			
PS P	P	P									12	13			

Table 44: Throughput - Material Requirements Planning © SAP SE

Sizing Element	Timeinter TI: Y/P/H/	Average/P A/P	Pl Avg. no. o Plan. orde	Average n Comp	Purchase ReorderR	Purchase Non-reord	Duration o Horiz.	Number o % chg.	% of BOM % dep.	Number o % LTS	Start of pr S.t.	End of pro E.t.	ID for mul ID	Fiori Apps FIORI	A short te Short text
MRP RUN	P	P									12	13			

Legend:	With examples	
Table ...	Table heading	
hea Sizing objects ... Cell	p	
toolti Objects	Cell	ption
descri		-in the s of the first three columns of a
A	Do not fill	cell
	Inputfield	
	Inputfield	

T Shirt Sizing

Sizing HANA Based on T-Shirt Sizing

- 2 x 10 core Westmere EX (Intel 2 socket system)
- 128 GB Main memory
- 160 GB PCIe-Flash / SSD for Log volume
- 1 TB SAS / SSD for Data volume
- 3 x 1 GB n/w or 1 x 10GB n/w (trunk), Redundant n/w infrastructure

- 2 x 10 core Westmere EX (Intel 2 or 4 sockets system)
- 256 GB Main memory
- 320 GB PCIe-Flash / SSD for Log volume
- 1 TB SAS / SSD for Data volume
- 3 x 1 GB n/w or 1 x 10GB n/w (trunk), Redundant n/w infrastructure

- 4 x 10 core Westmere EX (Intel 4 or 8 socket system)
- 512 GB Main memory
- 640 GB PCIe-Flash / SSD
- 2 TB SAS / SSD for Data volume
- 3 x 1 GB n/w or 1 x 10GB n/w (trunk), Redundant n/w infrastructure

- 8 x 10 Westmere EX core Intel (Intel 8 socket system)
- 1 TB Main memory
- 1.2 TB PCIe-Flash / SSD
- 4 TB SAS / SSD for Data volume
- 3 x 1 GB n/w or 1 x 10GB n/w (trunk), Redundant n/w infrastructure

Bare-metal sizing requirement	VMware t-shirt size	configuration
10 physical cores	XXS	20 vCPUs
20 physical cores	XS	40 vCPUs
30 physical cores	S	60 vCPUs
40 physical cores	n.A.	VMware vSphere 5.1 currently supports max. 64 vCPUs.
50+ physical cores	n.A.	