DMM201 – SAP IQ16
What it is and What it does

Richard Soundy SAP ESG (EMEA)

Disclaimer

This presentation outlines our general product direction and should not be relied on in making a purchase decision. This presentation is not subject to your license agreement or any other agreement with SAP. SAP has no obligation to pursue any course of business outlined in this presentation or to develop or release any functionality mentioned in this presentation. This presentation and SAP's strategy and possible future developments are subject to change and may be changed by SAP at any time for any reason without notice. This document is provided without a warranty of any kind, either express or implied, including but not limited to, the implied warranties of merchantability, fitness for a particular purpose, or non-infringement. SAP assumes no responsibility for errors or omissions in this document, except if such damages were caused by SAP intentionally or grossly negligent.

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

What is IQ16?

What Makes IQ Different?

IQ Success

Architecture Details

Where is IQ Used?

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

So, What is SAP IQ16?

What is SAP IQ?

SAP IQ is a market leading, high-speed columnar analytics server and data warehouse which is designed specifically for Data Analytics

Designed for Analytics	Unsurpassed Performance	Lowest TCO
• Columnar Database designed for high speed analytics • Eliminates the problems caused by traditional row-based systems	• Up to 1,000X faster than row-based systems • Enables high speed performance of complex queries against large datasets	• Reduce the need for aggregates and indexes • Greatly decreased ongoing maintenance costs when compared with row-based systems

All Data Sources | Large Data Volumes | Single Data Warehouse | Rapid Analytics

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

What is the SAP IQ server?

High-performance analytics platform

Inside:
Architected from the ground up for high-performance analytics

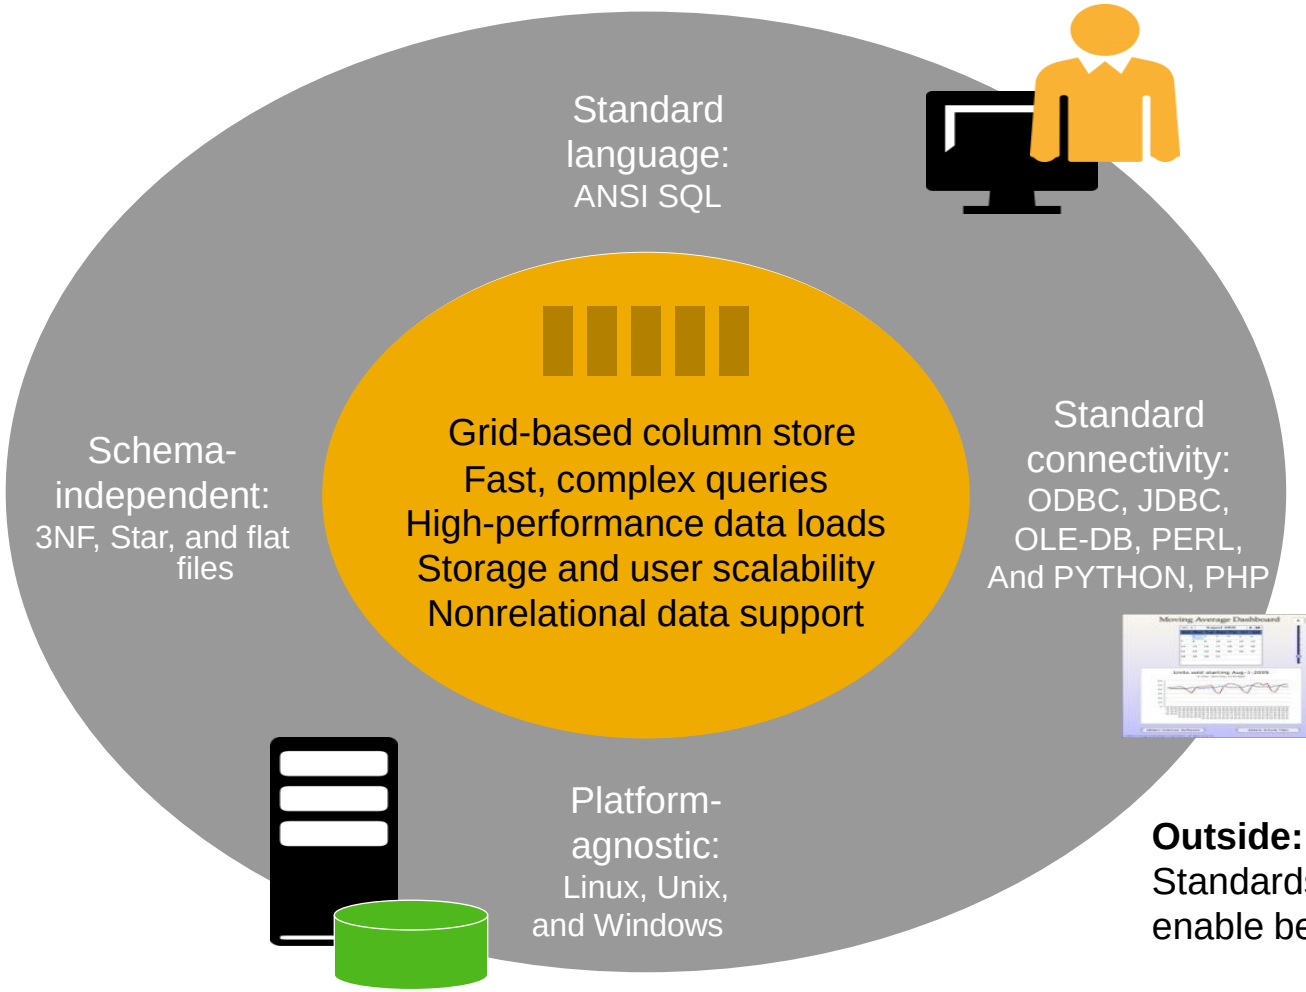

Outside:
Standards-based, open interfaces enable best-of-breed ecosystem

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

What Makes IQ Different

Why is SAP IQ Different? - 1

We chose to develop a separate database product specifically designed for data warehousing and decision support applications

This allowed design trade-offs to be made within IQ that no OLTP database could afford to make

IQ Design philosophy from day one was:

- **“Queries are King”**

OLTP database design philosophy must be:

- **“Writers are King”**

Primary design goal was to maximize the performance of many users running ad-hoc queries against vast amounts of data “Read-Optimized”

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Why is SAP IQ Different? - 2

```
Prioritized IQ Design Goals:
Multi-User Query Performance -   very high priority
Bulk Load Speed -                high priority
Single Row Insert Speed -        low priority
Single Row Update Speed -        low priority
```

Point #1:

- **Observation:** Most analytical or decision-support queries refer to only a small percentage of the columns in each referenced table
- **IQ Design Decision:** Store tables column-oriented rather than row-oriented, thereby enabling the reading of only the pages containing referenced columns

```
Table:                  "T"
Total row count:        10 billion
Number of columns:      50, "C1","C2",…"C50"
Average column width:   8 bytes
Data per row:           400 bytes
Total data volume:      3.6 Terabytes
```

What was traded for:

Example Query 1: Aggregation on all rows in largest table:

Total row count:	10 billion
Number of columns used:	10 out of 50
Average column width:	8 bytes
Useful data per row:	80 bytes out of 400
Total Data Read:	
Row-oriented storage:	3.6 Terabytes
Column-oriented storage:	0.7 Terabytes

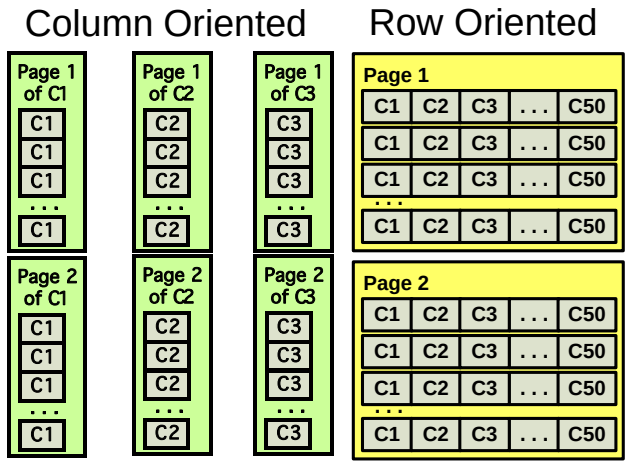

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Why is SAP IQ Different? - 3

```
Prioritized IQ Design Goals:
Multi-User Query Performance -   very high priority
Bulk Load Speed -                high priority
Single Row Insert Speed -        low priority
Single Row Update Speed -        low priority
```

Point #2:

- **Observation A:** CPU cycles have been getting cheaper much faster than I/O
- **Observation B:** Compression algorithms work more effectively when applied to data with more redundancy. There is much more redundancy in a column-store that has a single data type and domain than there is in a row-store containing multiple data types and domains
- **IQ Design Decision:** Apply compression to all pages written to disk, and use an asymmetric algorithm to reduce decompression costs at the expense of making compression more expensive

What was traded for:

Example Query 1: Aggregation on all rows in largest table:

Total row count:	10 billion
Number of columns used:	10 out of 50
Average column width:	8 bytes
Useful data per row:	80 bytes out of 400
Total Data Read:	
Row-oriented storage:	2.7 Terabytes
Column-oriented storage:	0.2 Terabytes

```
Table:                          "T"
Total row count:                10 billion
Number of columns:              50, "C1","C2",…"C50"
Average column width:           8 bytes
Data per row:                   400 bytes
Total data volume:              3.6 Terabytes

Assuming Full Asymmetric Compression
Row-Oriented Storage            2.7 TB
Column-Oriented Storage         0.2 TB
```

To illustrate the effects of compression, we will assume the following:

- "T" in a row-oriented database can be stored achieving a 25% average compression rate
- "T" stored in a column-oriented database can get a 75% average compression rate

These compression rates are not unusual, but your mileage may vary

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Why is SAP IQ Different? - 4

```
Prioritized IQ Design Goals:
Multi-User Query Performance -   very high priority
Bulk Load Speed -                high priority
Single Row Insert Speed -        low priority
Single Row Update Speed -        low priority
```

Point #3:

- **Observation:** Disk throughputs have been improving mush faster than disk average random seek time
- **IQ Design Decision A:** Reduce the number of seeks required by switching to larger database page sizes (e.g. 128K versus 4K pages) to better leverage modern disk and I/O subsystems
- **IQ Design Decision A:** Larger pages also tends to slightly improve average compression rates

What was traded for:

Example Query 1: Aggregation on all rows in largest table:

Total row count:	10 billion
Number of columns used:	10 out of 50
Average column width:	8 bytes
Useful data per row:	80 bytes out of 400
Total Data Read:	
Row-oriented storage:	966 Million I/Os
Column-oriented storage:	6 Million I/Os

```
Table:                    "T"
Total row count:          10 billion
Number of columns:        50, "C1","C2",…"C50"
Average column width:     8 bytes
Data per row:             400 bytes
Total data volume:        3.6 Terabytes

Row-Oriented Storage      4KB I/O
Column-Oriented Storage   128KB I/O
```

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Why is SAP IQ Different? – 5a

```
Prioritized IQ Design Goals:
Multi-User Query Performance -   very high priority
Bulk Load Speed -                high priority
Single Row Insert Speed -        low priority
Single Row Update Speed -        low priority
```

Point #4:

Observation A: **Many analytical queries involve multiple WHERE clause conditions on the same table**

Observation A: **Typically, OLTP database query engines can only make use of one index per table in a given query**

IQ Design Decision: **Give each row a unique number to serve as a persistent row identifier for that row**

Allow sets of rows to represented as a bitmap (i.e. a logical array of bits where if the Nth bit is on, then the row numbered N is part of the set of rows being represented).

Support indexes where each value in the index is associated with a bitmap representing the set of rows having that value

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Why is SAP IQ Different? – 5b

```
Prioritized IQ Design Goals:
Multi-User Query Performance -   very high priority
Bulk Load Speed -                high priority
Single Row Insert Speed -        low priority
Single Row Update Speed -        low priority
```

Point #4: Bitmaps

Effects of persistent row identifiers and bitmaps:

- Multiple indexes on the same table can easily be used together to help answer a query.
- Encourages the creation of many indexes, including multiple indexes of different types on the same column.
- Column stores can also be used as an index, if no better index is available for that condition.
- Rows from a table that are relevant to a given query are identified before the projection of needed columns begins, thereby reducing projection costs.
- Index-based access methods can accelerate joins, grouping, and ordering operations.
- Indexes provide meta-data as well as high performance access methods.
- No need for statistics maintenance.

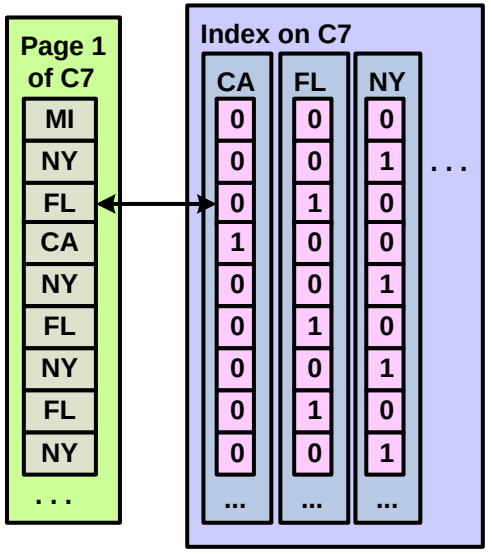

Bitmap-Based Index - Each value has an associated bitmap

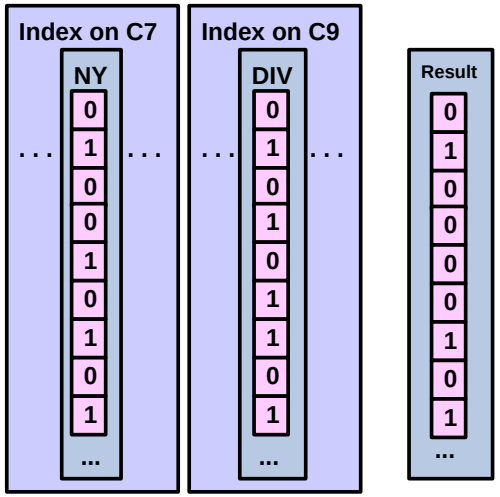

Using multiple bitmap-based indexes to identify the relevant set of rows for:

C7 = "NY" AND C9 = "DIV"

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Why is SAP IQ Different? - 6

```
Prioritized IQ Design Goals:
Multi-User Query Performance -   very high priority
Bulk Load Speed -                high priority
Single Row Insert Speed -        low priority
Single Row Update Speed -        low priority
```

Point #5:

Observation: When "Queries are King," a query should never be held up waiting for a page that is being modified

IQ Design Decision: When a page is to be modified, the modifier will copy that page to create a new version, thereby permitting queries referring to that page to continue without interruption. This technique is referred to as "Page-Level Snapshot Versioning" or "Page-Copy-On-Write"

Effects of snapshot versioning:

- A single transaction ID uniquely identifies the appropriate version of each table, and the table version identifies the appropriate pages for that version
- Different versions of a table share all pages not modified in any transaction between those two versions
- Queries are always operating at isolation level 3
- Does slightly increase disk space requirements because multiple versions of pages may be briefly necessary

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Why is SAP IQ Different? - 7

```
Prioritized IQ Design Goals:
Multi-User Query Performance -   very high priority
Bulk Load Speed -                high priority
Single Row Insert Speed -        low priority
Single Row Update Speed -        low priority
```

Point #6:

Observation: Write operations such as LOAD and INSERT are assumed to be relatively infrequent, with expected intervals in the tens of seconds or larger

IQ Design Decision: Reduce locking overhead by using table-level write locks. In combination with snapshot versions (Decision #5), this keeps the number of table versions down to a tractable level

Point #7:

Observation: The adverse effects of picking a poor query plan go up dramatically as the sizes of the tables involved goes up into the 10s or 100s of billions

IQ Design Decision A: All simple local conditions (a.k.a. SARGs) will be evaluated before most optimizer decisions are made, thereby providing more accurate information for those decisions

IQ Design Decision A: Wherever possible, make the query run-time capable of adapting to situations where the original plan was less than ideal because of inaccurate information

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Why is SAP IQ Different? – 8

```
Prioritized IQ Design Goals:
Multi-User Query Performance -   very high priority
Bulk Load Speed -                high priority
Single Row Insert Speed -        low priority
Single Row Update Speed -        low priority
```

Point #8:

Observation A: Scalability under mixed workloads requires the ability to take advantage of both SMP and cluster technology

IQ Design Decision: Support asymmetric shared disk clusters, referred to as IQ Multiplex

Within the cluster there is a single “Writer Node”, thereby avoiding the need for a distributed transaction manager and minimizing the required cluster synchronization traffic.

Users on “Query Nodes” see the most recently committed version of tables, exactly as they would if they were connected to the “Writer Node” or to a standalone (non-Multiplex) configuration

Effects of Multiplex:

- Enables many IQ servers to run against a single set of multi-ported disks.
- Only limit is the number of ports on the disks, and their total I/O bandwidth.
- Linear scaling of multi-user performance is achieved by simply adding ‘Query Nodes’.
- A site can choose to isolate any given set of users from each other and from loads

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

The Result…

Collective result of all these differences in IQ:

- **Unsurpassed multi-user query performance**
- **Unsurpassed storage efficiency**
- **Requires less expensive I/O subsystems**
- **Can scale up and out for future growth**

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Product Success

SAP IQ server

Mature, industrial-strength, and analytic database-management system (DBMS)

Leadership

- Industry-leading performance and scale benchmarks
- Recognized leader in enterprise data warehouse (EDW) market by Gartner Inc. and Forrester Group
- Pioneering technology with >10 patents

Adoption

- >5,000 installations in >2,850 accounts
- >96% customer satisfaction rates – consistently

Momentum

- Double the growth rate of the data warehouse market (last four years)
- Fast-paced product releases
- Versions 15 and 15.1 in 2009, version 15.2 in 2010, and version 15.3 and 15.4 in 2011
- Version 16 in March 2013
- HANA IQ Version Lock – May 2014

Ericsson • Sungard • Nielsen • BNP Paribas • Telefonica • hmv.com • comScore • Agricultural Bank of China

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

SAP positioned in leaders quadrant for data warehouse DBMS

Gartner 2012 Data Warehouse Database Management Systems: Magic Quadrant

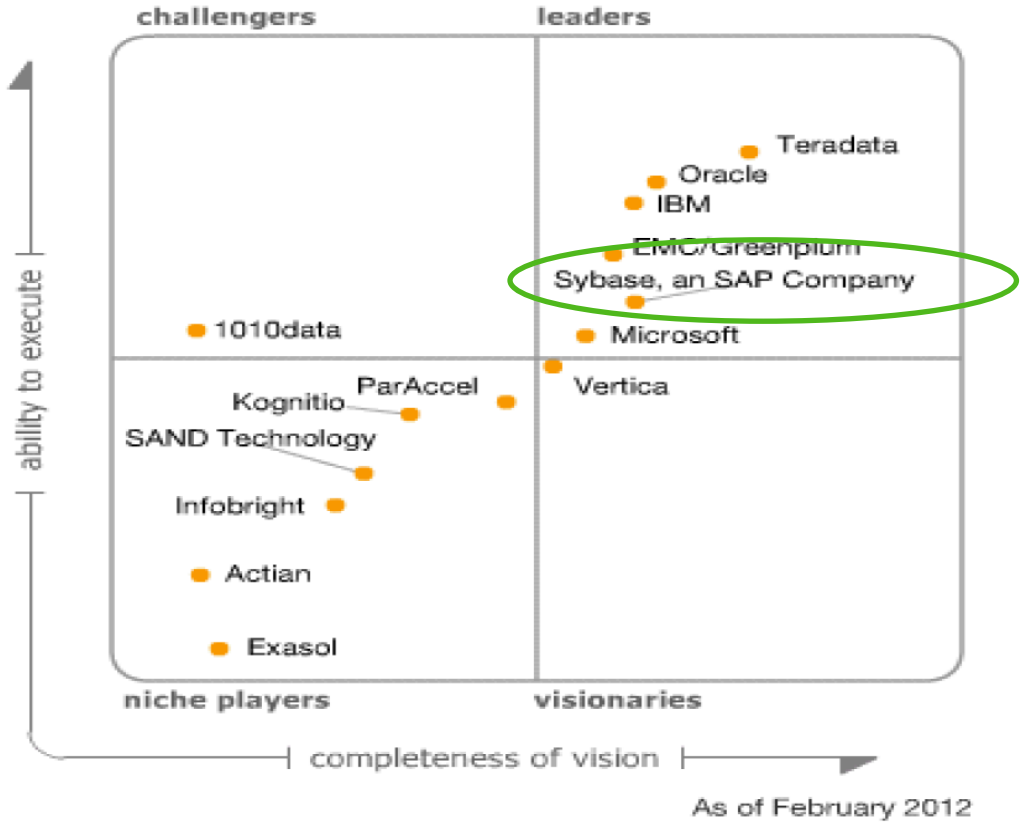

Source: Gartner
Published: February 6, 2012

This Magic Quadrant graphic was published by Gartner Inc. as part of a larger research note and should be evaluated in the context of the entire report. The Gartner report is available upon request from Sybase.

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

EDW: Analyst Momentum

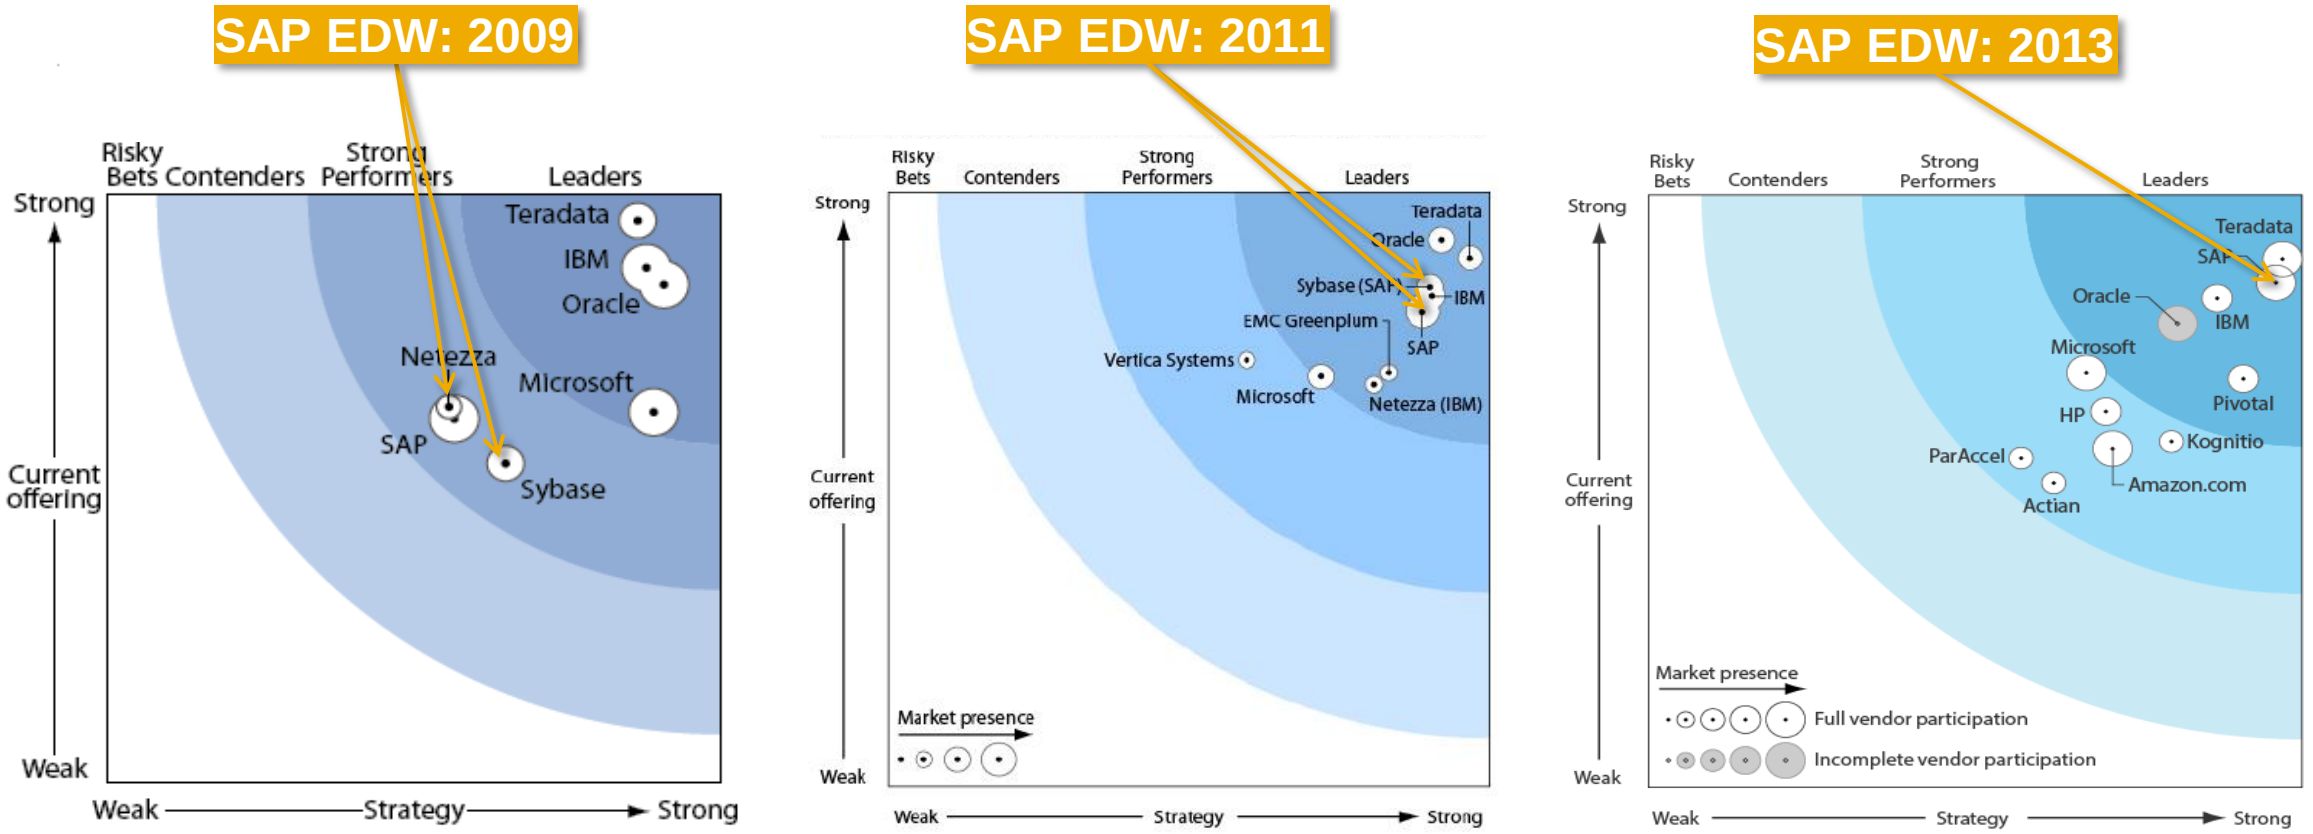

"SAP is finding its data mojo in HANA and IQ to support broader EDW use cases…."

Forrester Research, Dec 2013

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

SAP Sybase IQ is the leader in deployments

Sybase has more customers than all of the other specialty EDW leaders *combined*

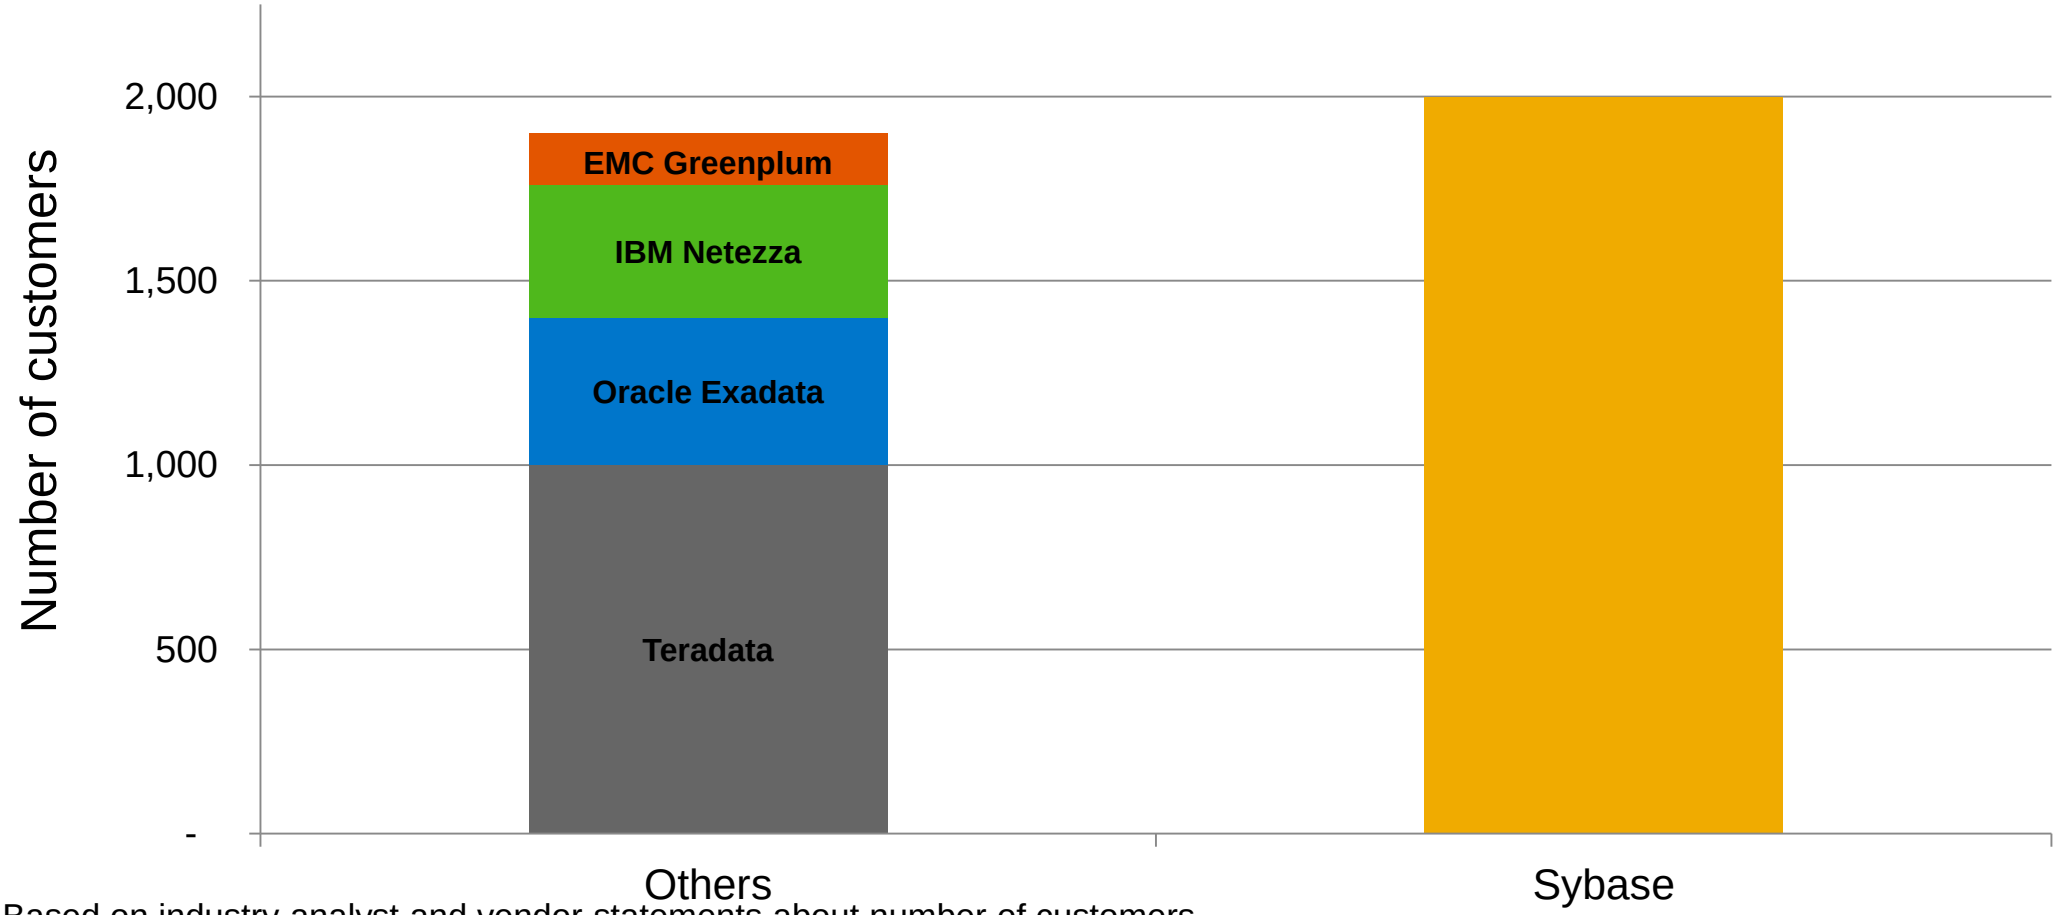

* Based on industry analyst and vendor statements about number of customers

* Vendors selected for 'Others' category based on i) 'leadership' position in Gartner MQ, ii) having a specialty Data Warehousing offering (i.e. General purpose databases are excluded)

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Architecture Details

SAP IQ 16

Innovations for extremely large databases (XLDB)

Storage Architecture

- New generation column store
- New partitioning and compression

SAP IQ BigData/XLDB Analytics

Loading Engine

- Fully parallel bulk loading
- Real-time loading into delta store

System Reliability

- Grid resiliency
- LDAP and role-based security

Query Processing

- Data affinity
- Aggressively parallel and distributed

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Multiplex grid

Architecture overview

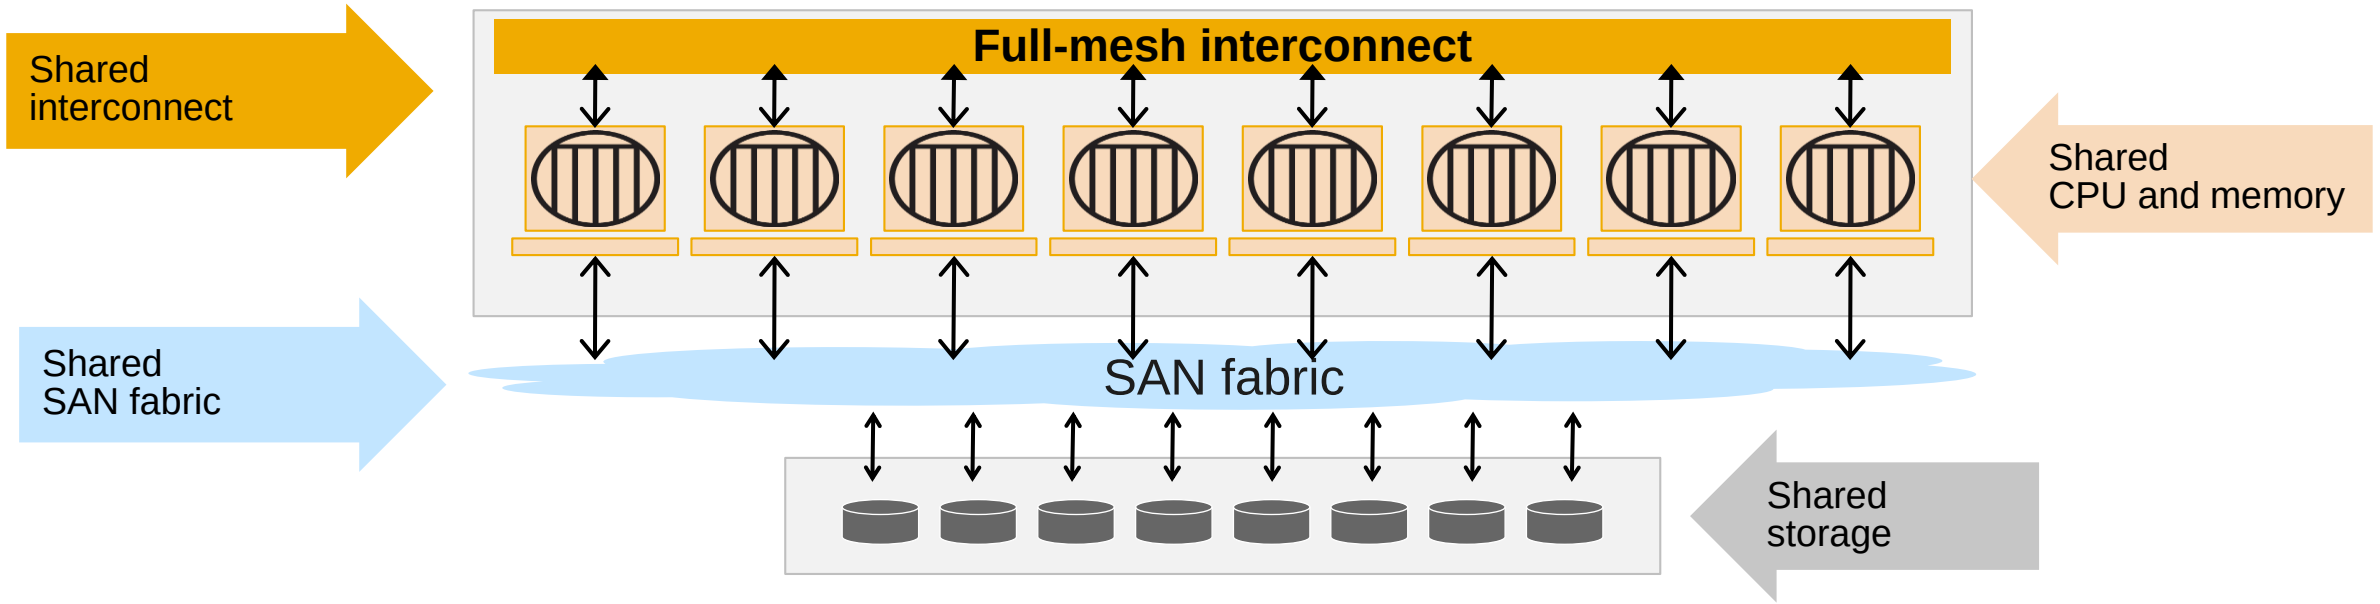

Compute nodes:

- Can be x86, IBM Power, HP Itanium, and Oracle Solaris Sparc servers
- Each node can have a different CPU core count, memory size
- One node is designated as the coordinator for meta data management
- Any node can be a write node
- Any node can be a read node

Shared storage:

- Can be SSD, FC, SATA, or a combination of all
- Shared storage can be physical SAN with FC or iSCSI
- Shared storage can be virtual SAN over DAS
- Storage can be logically grouped with compute
- Storage can be tiered for data aging

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Robust bulk load engine

Loading can be from multiple nodes:

- Load rates in excess of 250 GB/hour are common even with modest-sized hardware nodes
- Incremental loads via microbatching (change data capture)

Page-level snapshot versioning:

- No locking (only table-lock in-memory catalogs) – allows nonblocking concurrent loads and queries

Load from client machines:

- Eliminates requirement for load files to reside on database server machine

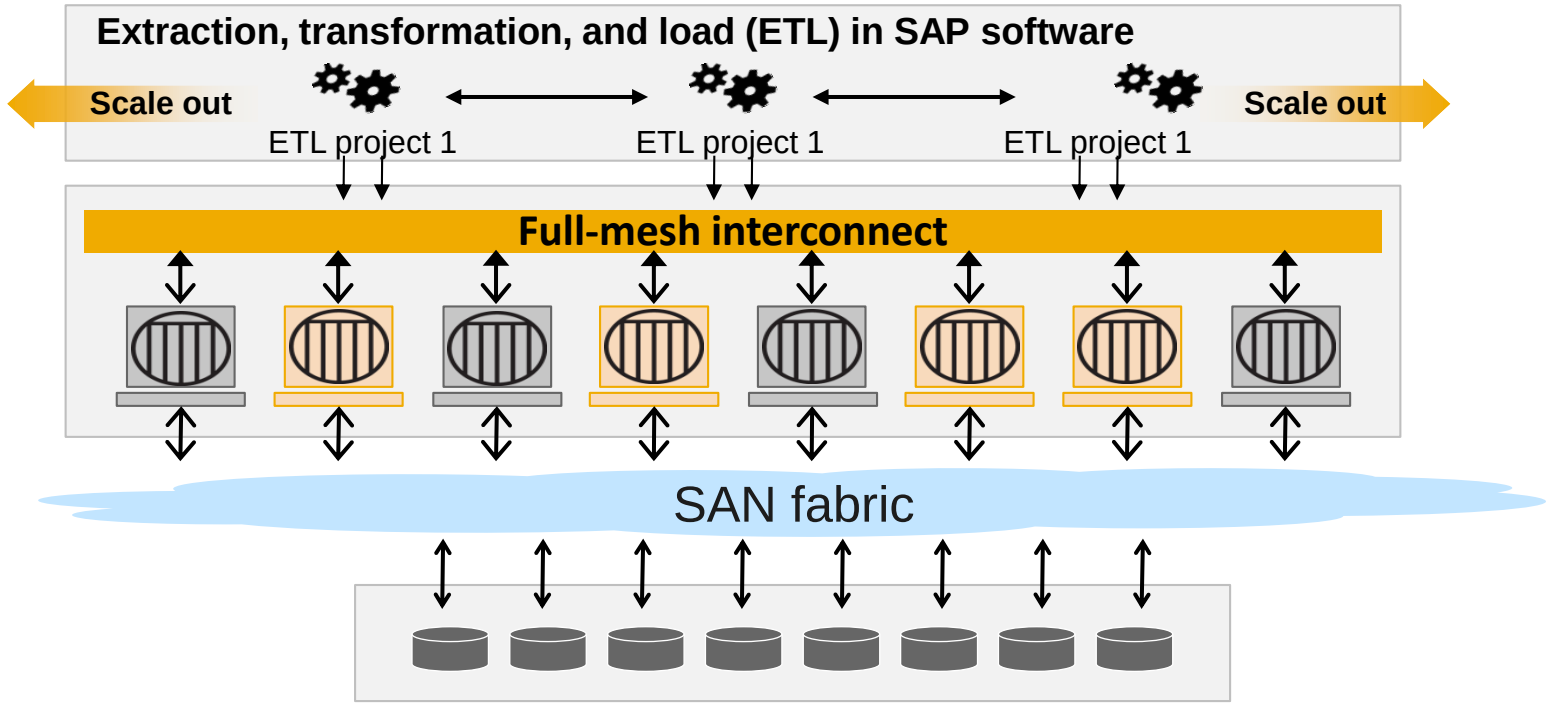

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

High velocity data loading

Row-level versioned (RLV) store

Value proposition

- Continuous analytics over operational data
- High velocity, concurrent data modifications
- Exploit large memory and core footprints

Architectural considerations

- Write optimized in-memory In-memory RLV (Row-level versioned) store
- Row level locking, and statement snapshot isolation
- Low latency micro operations
- In-memory RLV store has reduced compression, no sorting, no indexing
- Fully recoverable with dedicated transaction log
- Asynchronous data transfer from In-memory RLV store to IQ main store
- Users choose which tables are In-memory RLV tables
- Simplex only (not available on Multiplex)

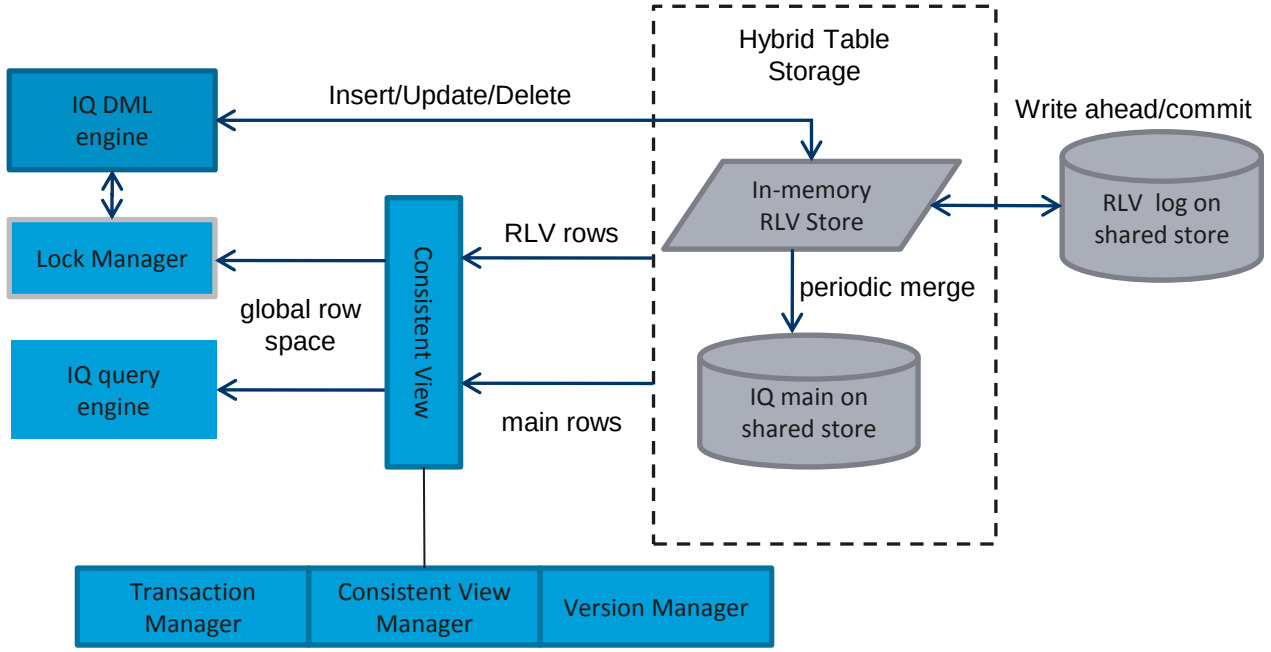

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Continuous real-time loads

Using SAP Replication Server

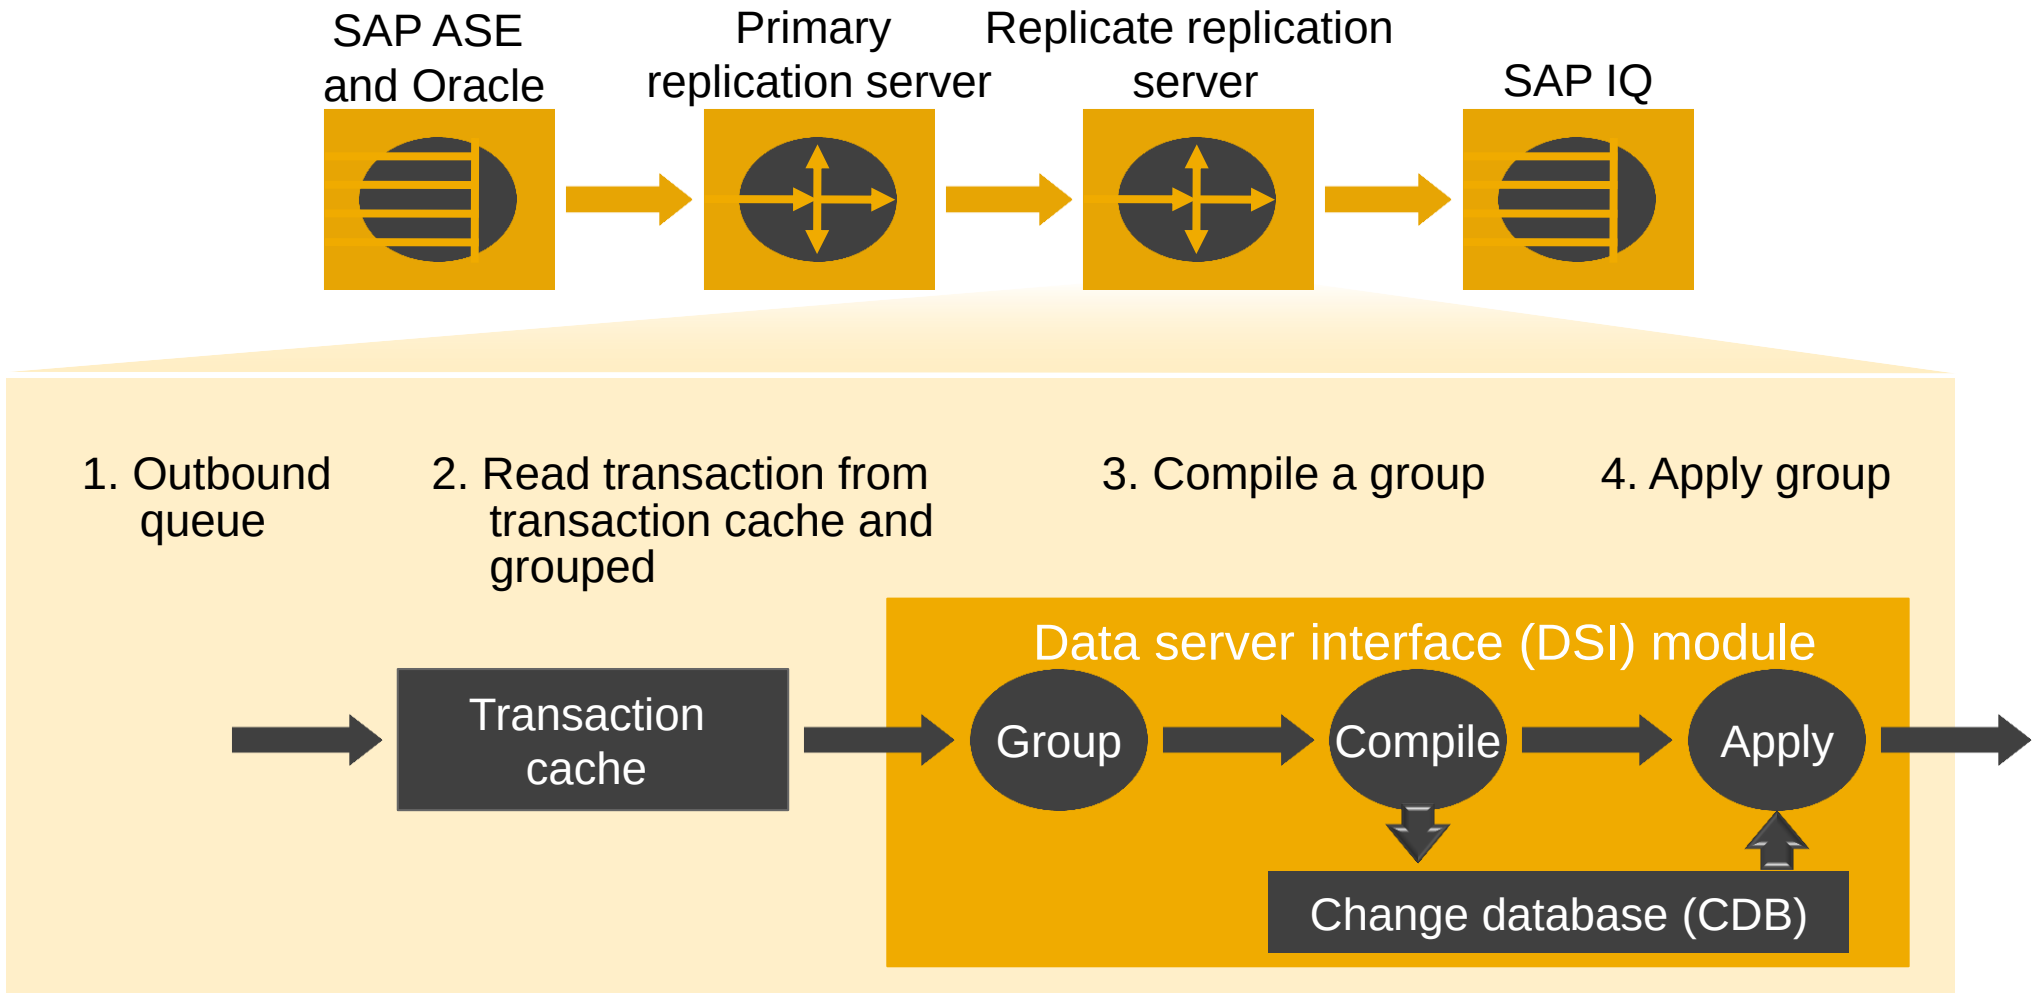

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Column store processor

Leading data compression technology

- **New generation column store architecture:**
  - Supports a variable number of cells per page
  - Supports various page formats within a column
  - High performance access paths with richer metadata
  - Insert/update/delete of variable length data processed efficiently
  - LZW compression with implicit dictionary

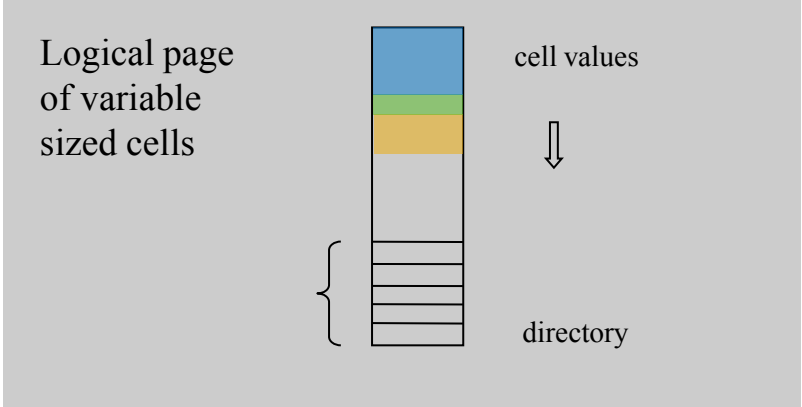

Superior compression achieved for customers!

SAP IQ: data compression in real world	Raw data loaded	SAP IQ: compress	Data explosion in traditional RDBMS
Telefonica	70 TB	15 TB	210 TB to 490 TB
Health insurance review agency	27 TB	12 TB	81 TB to 189 TB
Samsung card	15 TB	7 TB	45 TB to 105 TB

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Query engine scale out

Distributed query processing

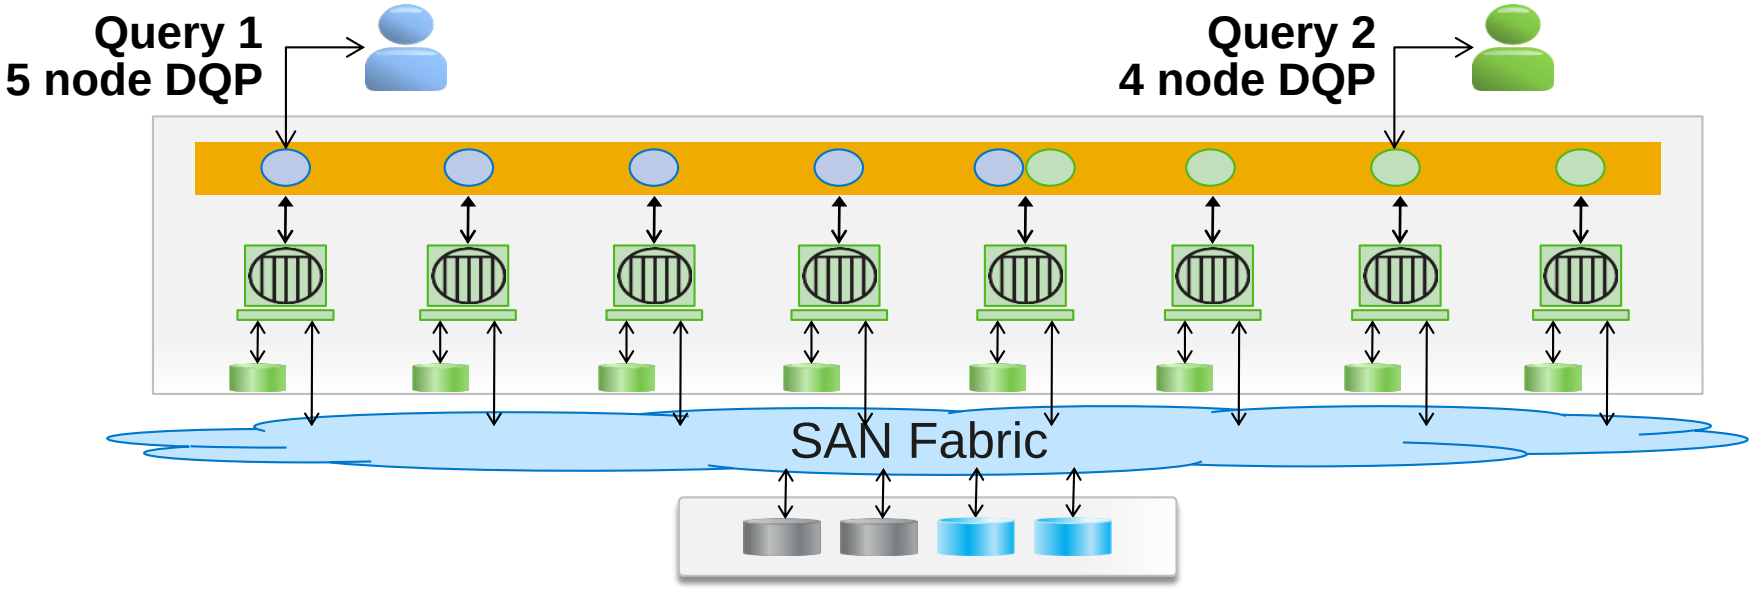

Massively parallel processing

- **Leader node:** Receives and initiates queries, including user-defined functions (UDFs)
  - Any node can be a leader, one leader per query, many concurrent leaders possible
  - Leader node may satisfy query within itself
- **Worker node:** Nodes picking up work units from leader
  - Many worker nodes per query, same worker node can serve multiple queries
  - Worker nodes are enlisted only if leader cannot satisfy query on its own

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Problems with complex analytics

Data to logic

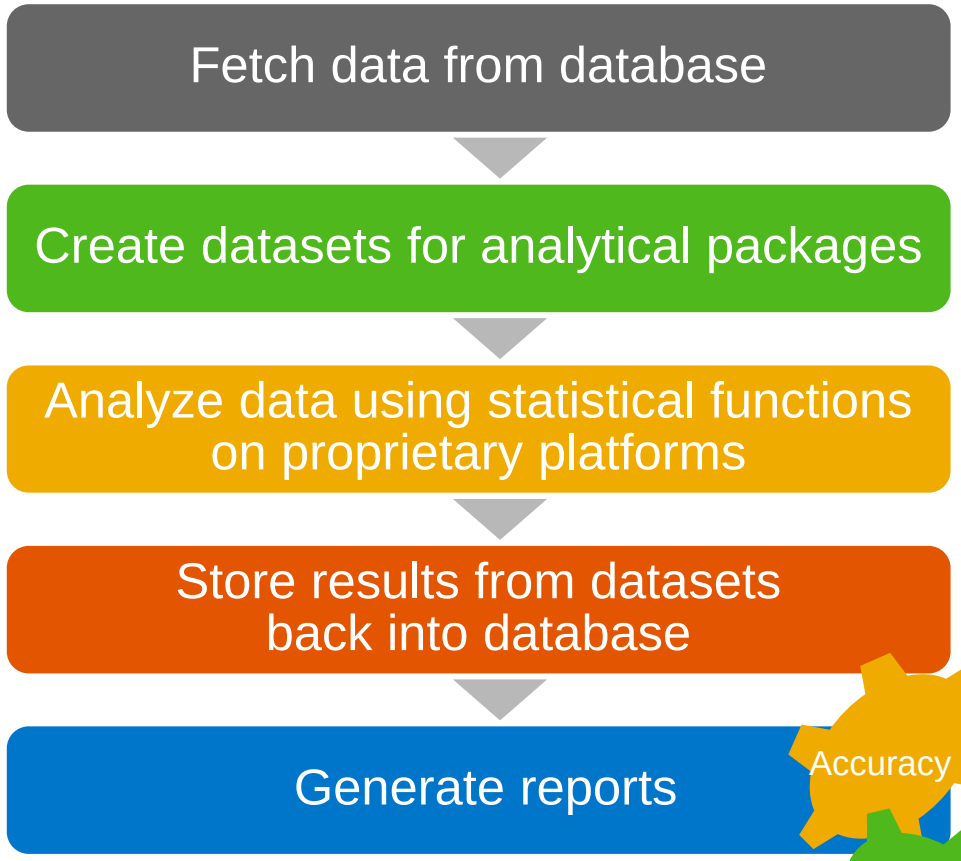

Time-consuming process –
could run into memory constraints with large data sets

Proprietary platforms make it very difficult to embed in applications

Another time-consuming process that could slow down the delivery of results to end users

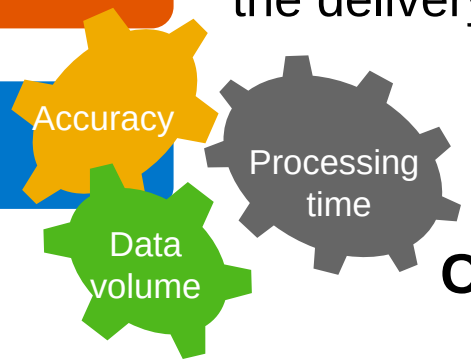

Compromise on at least one key area

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

In-database analytics in SAP IQ

No compromise for complex analytics:

- Basic to advanced analytical functions available to SQL directly from engine of SAP IQ
- Data never leaves the database until results are materialized
- Analytics code and models must be shareable yet must allow ad hoc analysis
- Analytics code and models must be applicable to the latest data set
- Standards-based access; concept extensibility is compulsory
- Performance and scalability is a given
- Average developer must be able to build in database analytical models

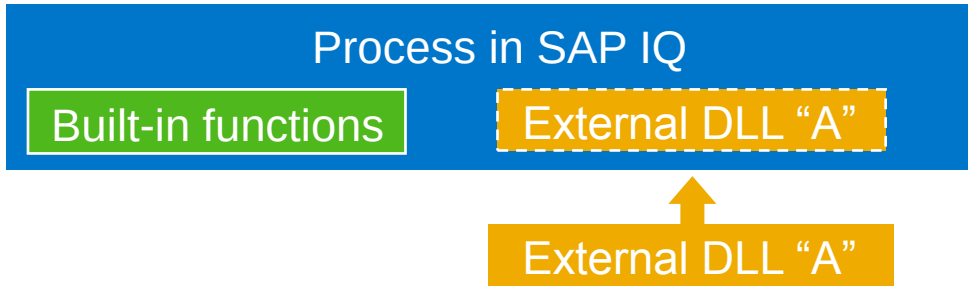

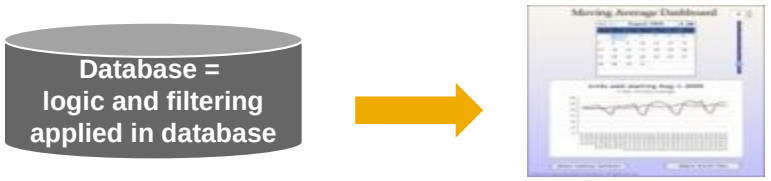

Analytics simplified: Logic to data = fast and efficient

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Query federation

With other databases

Ability to analyze data in SAP IQ as well as in remote RDBMS (only one copy of data in SAP IQ and a small set of data in remote systems)

Remote servers could expose materialized views (current data is queried to provide integrated view of the data marts

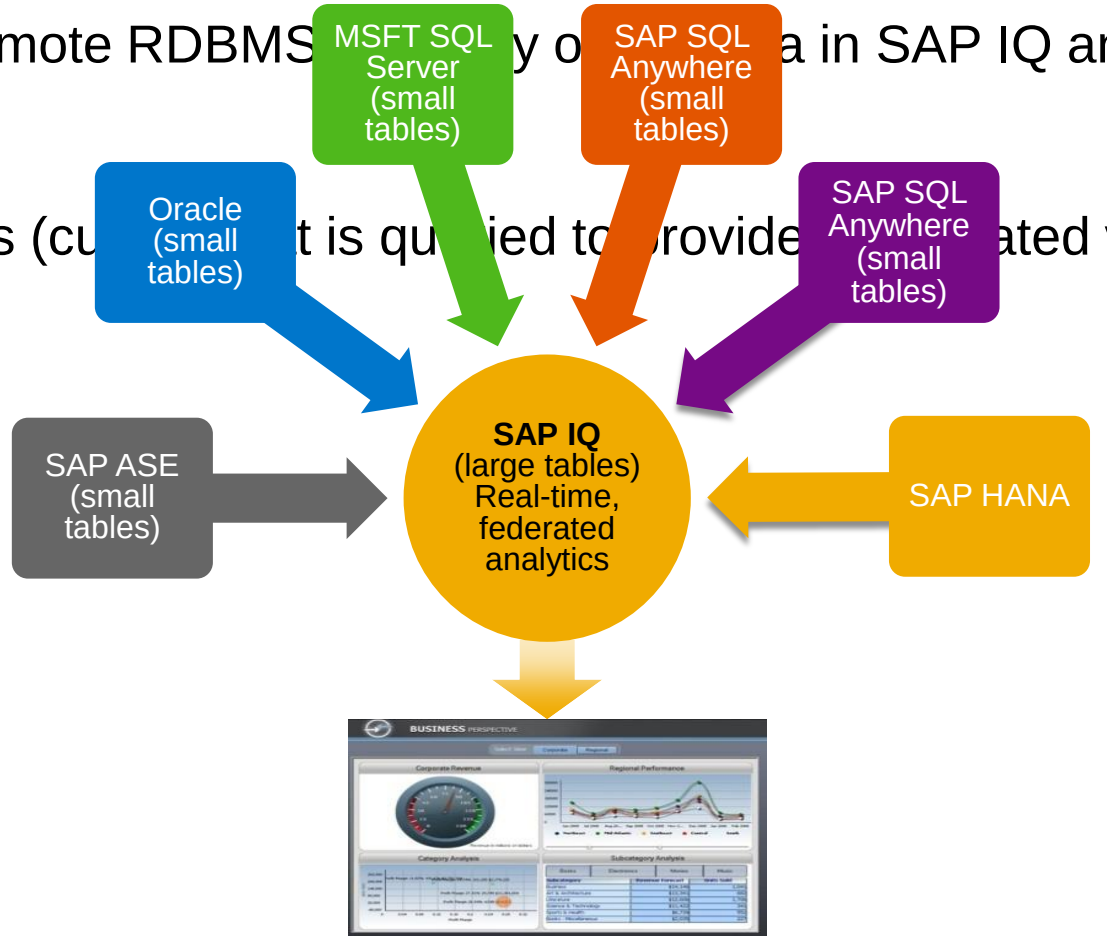

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Federation

With external file systems (Hadoop distributed file system)

1. 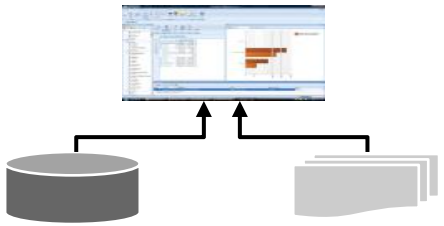

Client-side federation: Join data from SAP IQ and Hadoop at a client-application level

2. 

Load Hadoop data into column store of SAP IQ: Extract, transform, and load data from Hadoop distributed file system (HDFS) into schemas of SAP IQ

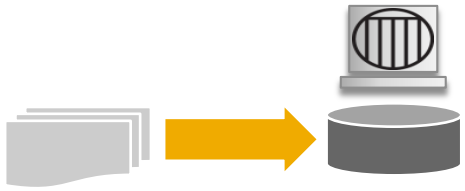

Join HDFS data with data of SAP IQ on the fly: Fetch and join subsets of HDFS data on demand, using SQL queries from SAP IQ (***data federation*** technique)

3. 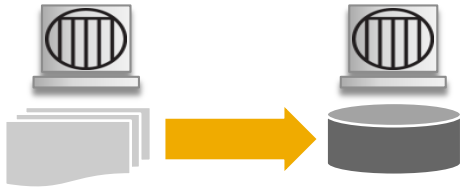

Combine results of Hadoop MapReduce (MR) jobs with SAP IQ data on the fly: Initiate and join results of MR jobs on demand using SQL queries from data in SAP IQ (***query federation*** technique)

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

So Where is IQ used

And where should you use it

SAP IQ server

Mature, industrial-strength, and analytic database-management system (DBMS)

Leadership

- Industry-leading performance and scale benchmarks
- Recognized leader in enterprise data warehouse (EDW) market by Gartner Inc. and Forrester Group
- Pioneering technology with >10 patents

Adoption

- >5,000 installations in >2,850 accounts
- >96% customer satisfaction rates – consistently

Momentum

- Double the growth rate of the data warehouse market (last four years)
- Fast-paced product releases
- Versions 15 and 15.1 in 2009, version 15.2 in 2010, and version 15.3 and 15.4 in 2011
- Version 16 in March 2013
- HANA IQ Version Lock – May 2014

Ericsson • Sungard • Nielsen • BNP Paribas • Telefonica • hmv.com • comScore • Agricultural Bank of China

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Financial Services Customers Succeed with IQ for Mission-Critical Applications

	Customer	IQ Deployment
	US Securities & Exchange Commission	With Sybase IQ, the SEC built a flexible, scalable and reliable data-rich warehouse that delivers faster load times and query results, and more comprehensive information views—all at a lower cost.
	Citigroup	With Sybase powering their data arrchive, Citigroup users can now research historical data themselves, whereas in the past IT specialists were required to do so.
	Morgan Stanley (Spain)	With Sybase IQ, organ Stanley Spain now has access to heterogeneous system-stored critical data with single, unified delivery that has dramatically reduced analysis response times.
	ICICI Bank	ICICI bank has achieved improvement in system uptime and query performance over its previous Teradata implementation, in addition to the host of other benefits of the Sybase IQ data warehouse migration.
	BNP Paribas Securities Services	Sybase IQ has improved the robustness of its online reporting system. Now the bank can compile larger amounts of data and anticipate growth in its customer base by proactively providing real-time data.

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Healthcare Customers Succeed with Sybase IQ for Mission-Critical Applications

	Customer	IQ Deployment
	New South Wales Health	Sybase IQ solution decreases load times, accomodates increase in tables and compression for the new configuration to 73 percent of the previous storage requirement. Reduces query time across data sets.
	South Australia Health	Solution reduces lag in access to performance metrics from weeks to minutes; provides visibility across eight emergency departments, 12 metropolitan public hospitals, 17 regional hospitals and other services.
	GWQ Service Plus AG	Solution now stores and analyzes large data sets efficiently so company can more easily add new customers and conduct research that helps create new product offerings.
	Medihelp	Solution allows various groups within Medihelp, from senior management to product development and client service, to rapidly access and act on comprehensive data sets to make better decisions.
	HealthTrans	Solution improves HealthTrans's analytics capabilities, opening up a new realm of potential revenue while reducing storage expenses and maintenance costs.

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Information Provider Customers Succeed with SAP Sybase IQ for Mission-Critical Applications

	Customer	IQ Deployment
	comScore	comScore implemented a Sybase IQ data warehouse that ensures its Customer Knowledge Platform scales economically to eventually handle over 20TB of data while supporting intense data queries from customers.
	Nielsen Media Research	Sybase IQ's multiplexing capability also allows Nielsen to add new products and applications that can access the audience data warehouse in order to satisfy ever-evolving client demands.
	Experian	Offering targeted marketing solutions driven by advanced analytics, Sybase IQ helped Experian increase sales per lead by 22%; sales per hour jumped by 13%; and the number of leads generated grew by 3%.
	Alvion Technologies	Implementing Sybase IQ dramatically improved Alvion's competitive advantage by shortening data fulfillment times while still maintaining the trademark query speeds so valued by its customers.
	AdOn Network	Sybase Analytic Appliance allows AdOn Network to quickly query large amounts of data, thereby helping publishers and advertisers to maximize revenue and manage timely targeted promotions, respectively.

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Insurance Customers Succeed with SAP Sybase IQ for Mission-Critical Applications

	Customer	IQ Deployment
	Allianz Australia	In a mere eight weeks Allianz Australia implemented a completely new data warehouse and developed analytical reporting capabilities, leveraged by the suite of pre-built business intelligence templates in Sybase IWS.
	ICICI Prudential Life Insurance	With Sybase technology, ICICI Prudential can now help customers more effectively and to a greater extent because employees have immediate access to accurate, up-to-date information.
	DEVK	Sybase IQ underpins DEVK's high-performance direct mailing database through the Web, giving campaign decision-makers much more flexibility.
	Samsung Life Insurance	Sybase technology unifies and simplifies Samsung Life Insurance‘s data warehouse, providing an integrated view for its system.
	HDI Insurance	With Sybase technology, HDI Insurance can technologically support company growth with improved database transaction capabilities as well as more secure operations and increased management capacity.
	Caser Seguros	The deployment of Sybase IQ has enabled Caser to centralize information, standardize its presentation, and employ consistent means of access.

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Public Sector Customers Succeed with SAP Sybase IQ for Mission-Critical Applications

Customer	IQ Deployment
U.S. Department of Transportation	The Department developed BTS TranStats, a business intelligence solution powered by SAP Sybase IQ. TranStats contains more than 400 tables (the largest of which exceeds 250 million rows) across 100 databases with raw data storage of 2.5 terabytes compressed by Sybase IQ to about 1 terabyte.
Korea Customs Service	Active data analysis, enabled by Sybase IQ, has increased the amount of collected customs taxes, exposed more illegal foreign exchange, and enabled KCS to crack down on smuggling and money laundering significantly more effectively compared to the levels achieved prior to the implementation of the CDW. As a result, government revenues have increased.
Agencia Tributaria	Sybase IQ was chosen for its ability to store, process and access data in columns, rather than by rows as in traditional relational database architectures. This architecture translates into efficient processing, faster query responses, and dramatic storage compression rates.
Alfred Wegener Institute	AWI chose Sybase IQ to serve as the centralized data store for the enormous volume of scientific data gathered. Key considerations in this selection were Sybase IQ's column-based architecture, enabling faster, more flexible searches and its patented compression capability, reducing the amount of storage hardware required and consequently, reducing costs.
Cambridge Astronomical Survey Unit	After looking for a relational database that could store very large data volumes in a compressed, binary way, as telescopes do, CASU chose instead to implement Sybase IQ primarily for its impressive data compression capability and its rapid query response times.
European Southern Observatory	Sybase IQ enables ESO to react rapidly to changing conditions, provides to two-way information flow across multiple times zones, and addresses operational problems at the observatories. ESO's Data Flow System allows for the instantaneous database replication a transcontinental distances, linking worldwide observatories and allowing them to address operational issues.

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Telecommunications Customers Succeed with Sybase IQ for Mission-Critical Applications

	Customer	IQ Deployment
	China Telecom	The new Sybase network management system handles daily network management + expedites the analysis of significant amounts of historical data for forecasting.
	Spice Telecom	The new Sybase IQ data warehouse is readily available to the people who most require the information, helping the company manage explosive growth in a competitive market.
	Taiwan Mobile	Sybase IQ allows Taiwan Mobile to process the complex analytics required to keep its business on course as it continues to gain new customers
	Utimaco Safeware	Sybase IQ gives Utimaco a state-of-the-art data warehouse that has become the backbone of the DRS solution. It allows for massive data handling, and scalable analytics.
	Suntel	Sybase IQ Suntel is able to run business-critical reports and queries against a larger volume of historical data without negatively impacting the performance of its production database

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

SAP d-code Virtual Hands-on Workshops and SAP d-code Online

Continue your SAP d-code education after the event!

SAP d-code Virtual Hands-on Workshops

- Access hands-on workshops post-event
- Starting January 2015
- Complementary with your SAP d-code registration

http://sapdcodehandson.sap.com

SAP d-code Online

- Access replays of keynotes, Demo Jam, SAP d-code live interviews, select lecture sessions, and more!
- Hands-on replays

http://sapdcode.com/online

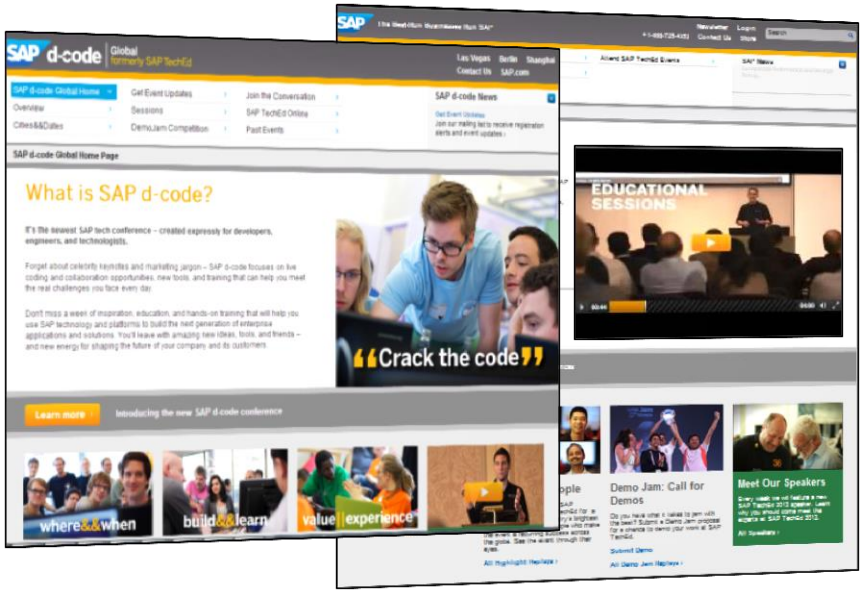

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Further Information

SAP Public Web

scn.sap.com
www.sap.com

SAP Education and Certification Opportunities

www.sap.com/education

Watch SAP d-code Online

www.sapcode.com/online

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

Feedback

Please complete your session evaluation for

DMM201

Richard Soundy (richard.soundy@sap.com)

Thanks for attending this SAP TechEd && d-code session.

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

© 2014 SAP SE or an SAP affiliate company. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP SE or an

SAP affiliate company.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP SE (or an SAP affiliate company) in Germany and other countries. Please see http://global12.sap.com/corporate-en/legal/copyright/index.epx for additional trademark information and notices.

Some software products marketed by SAP SE and its distributors contain proprietary software components of other software vendors.

National product specifications may vary.

These materials are provided by SAP SE or an SAP affiliate company for informational purposes only, without representation or warranty of any kind, and SAP SE or its affiliated companies shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP SE or SAP affiliate company products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

In particular, SAP SE or its affiliated companies have no obligation to pursue any course of business outlined in this document or any related presentation, or to develop or release any functionality mentioned therein. This document, or any related presentation, and SAP SE's or its affiliated companies' strategy and possible future developments, products, and/or platform directions and functionality are all subject to change and may be changed by SAP SE or its affiliated companies at any time for any reason without notice. The information in this document is not a commitment, promise, or legal obligation to deliver any material, code, or functionality. All forward-looking statements are subject to various risks and uncertainties that could cause actual results to differ materially from expectations. Readers are cautioned not to place undue reliance on these forward-looking statements, which speak only as of their dates, and they should not be relied upon in making purchasing decisions.