

Neue Wege zur Speicherung, Verarbeitung und Analyse von IoT-Daten

Andreas Weininger, IBM

Andreas.Weininger@de.ibm.com

 @aweininger




Agenda

- Neue Anforderungen durch IoT
- Architektur zur Lösung
- Technologien:
 - Hochgeschwindigkeitsdatenverarbeitung
 - Datenvirtualisierung

Neue Anforderungen durch IoT



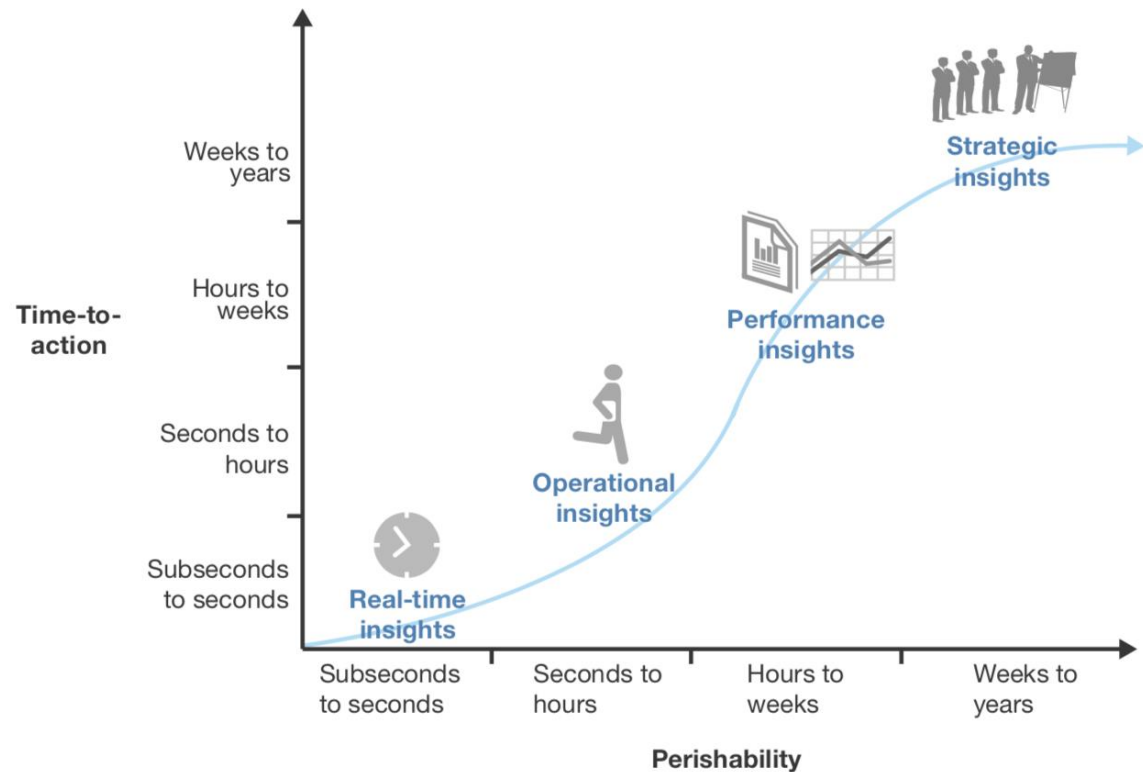
Daten sind überall und immer heterogener



Perishable Insights: Insights That Are Actionable

Source: Forrester Perishable Insights-Stop Wasting Money on Unactionable Insights Report 2016

- Being able to **act on insights in real-time** can provide maximum differentiation and value to businesses.
- To act on real time events, businesses need an Application Architecture that can ingest, analyze and act upon events as they happen in sub-second to seconds response time.
- Machine Learning combined with Reactive Programming frame works enable insights to be actioned automatically and stay current so business can react to events in sub-second timer frames



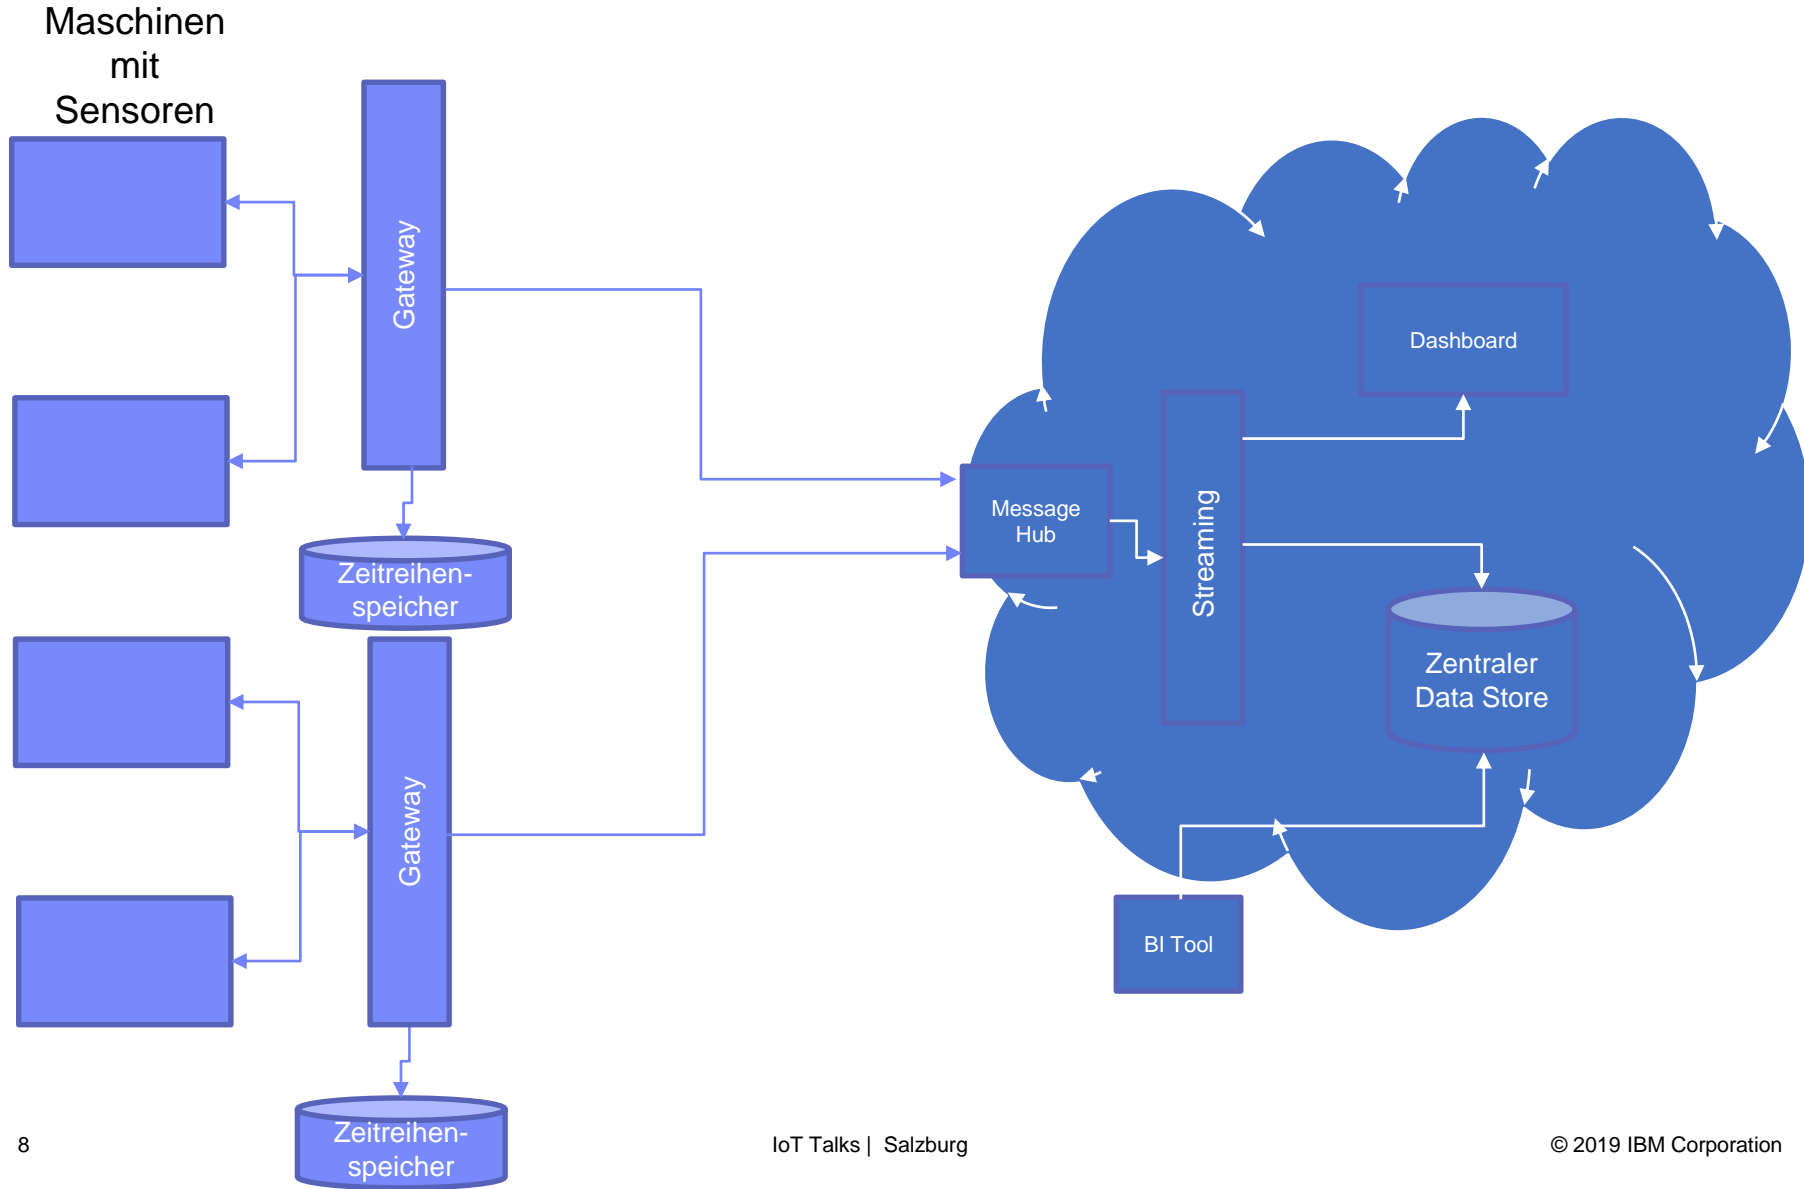
Was hat sich geändert?

- IoT bietet neue Anforderungen an die Speicherung, Verarbeitung und Analyse von Daten:
 - Extrem große Datenmengen grosse Datenmengen fallen in sehr kurzer Zeit an
 - Maximale Zeitspanne zwischen Erzeugung der Daten und ihrer (evtl. komplexen) Auswertung eventuell gering
 - Von Sensoren generierte Datenmenge kann die verfügbare Bandbreite der Anbindung zum zentralen Auswertungssystem überschreiten.

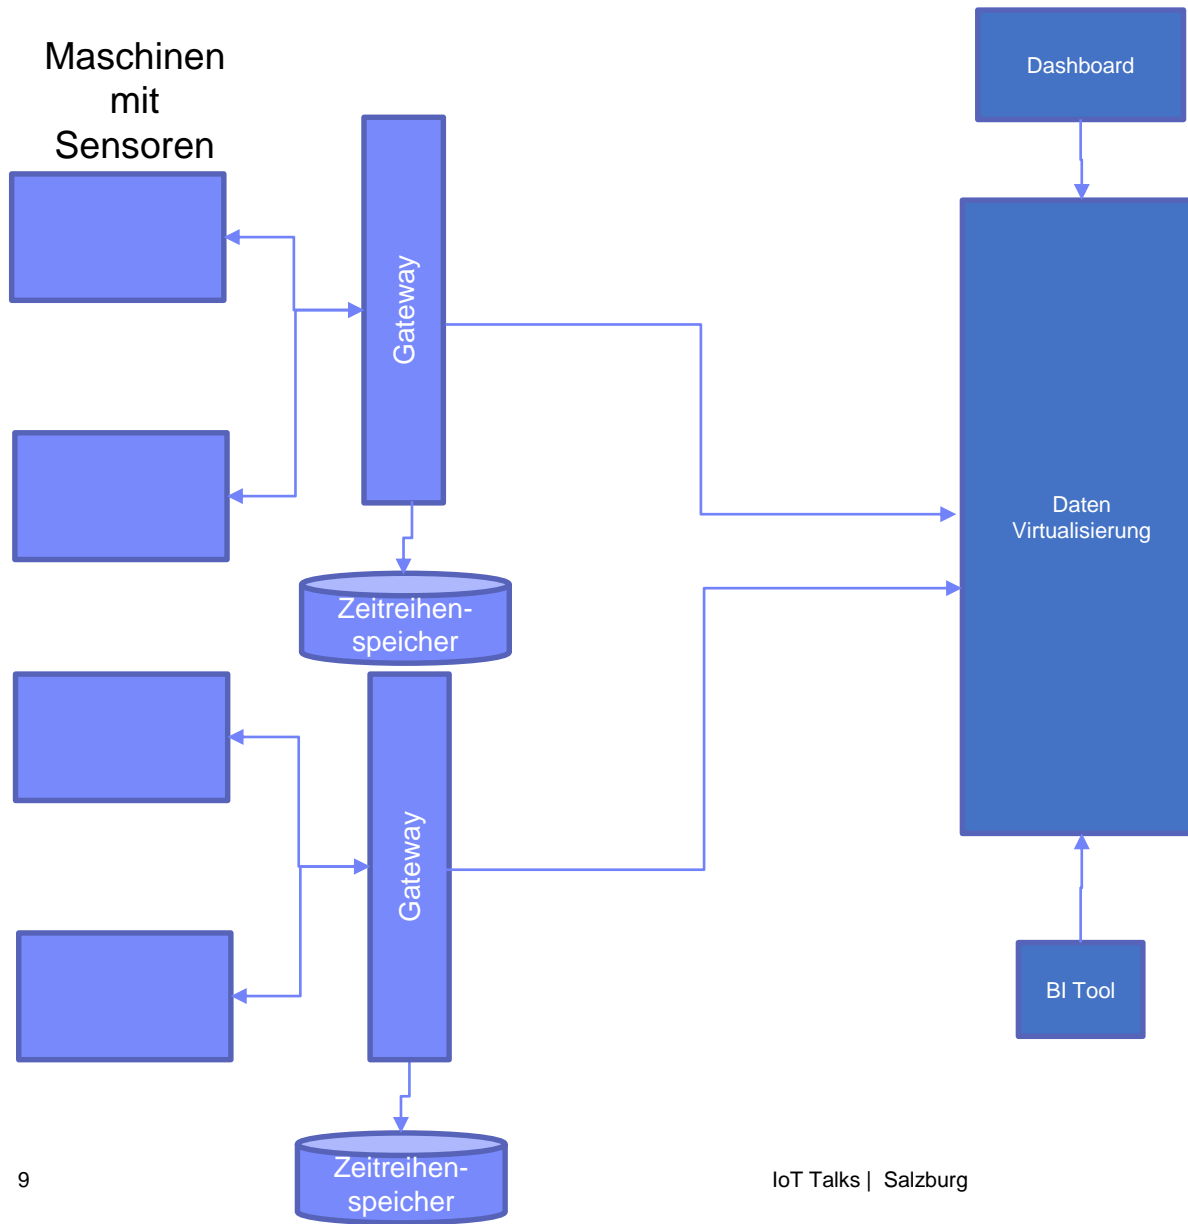
Architektur zur Lösung



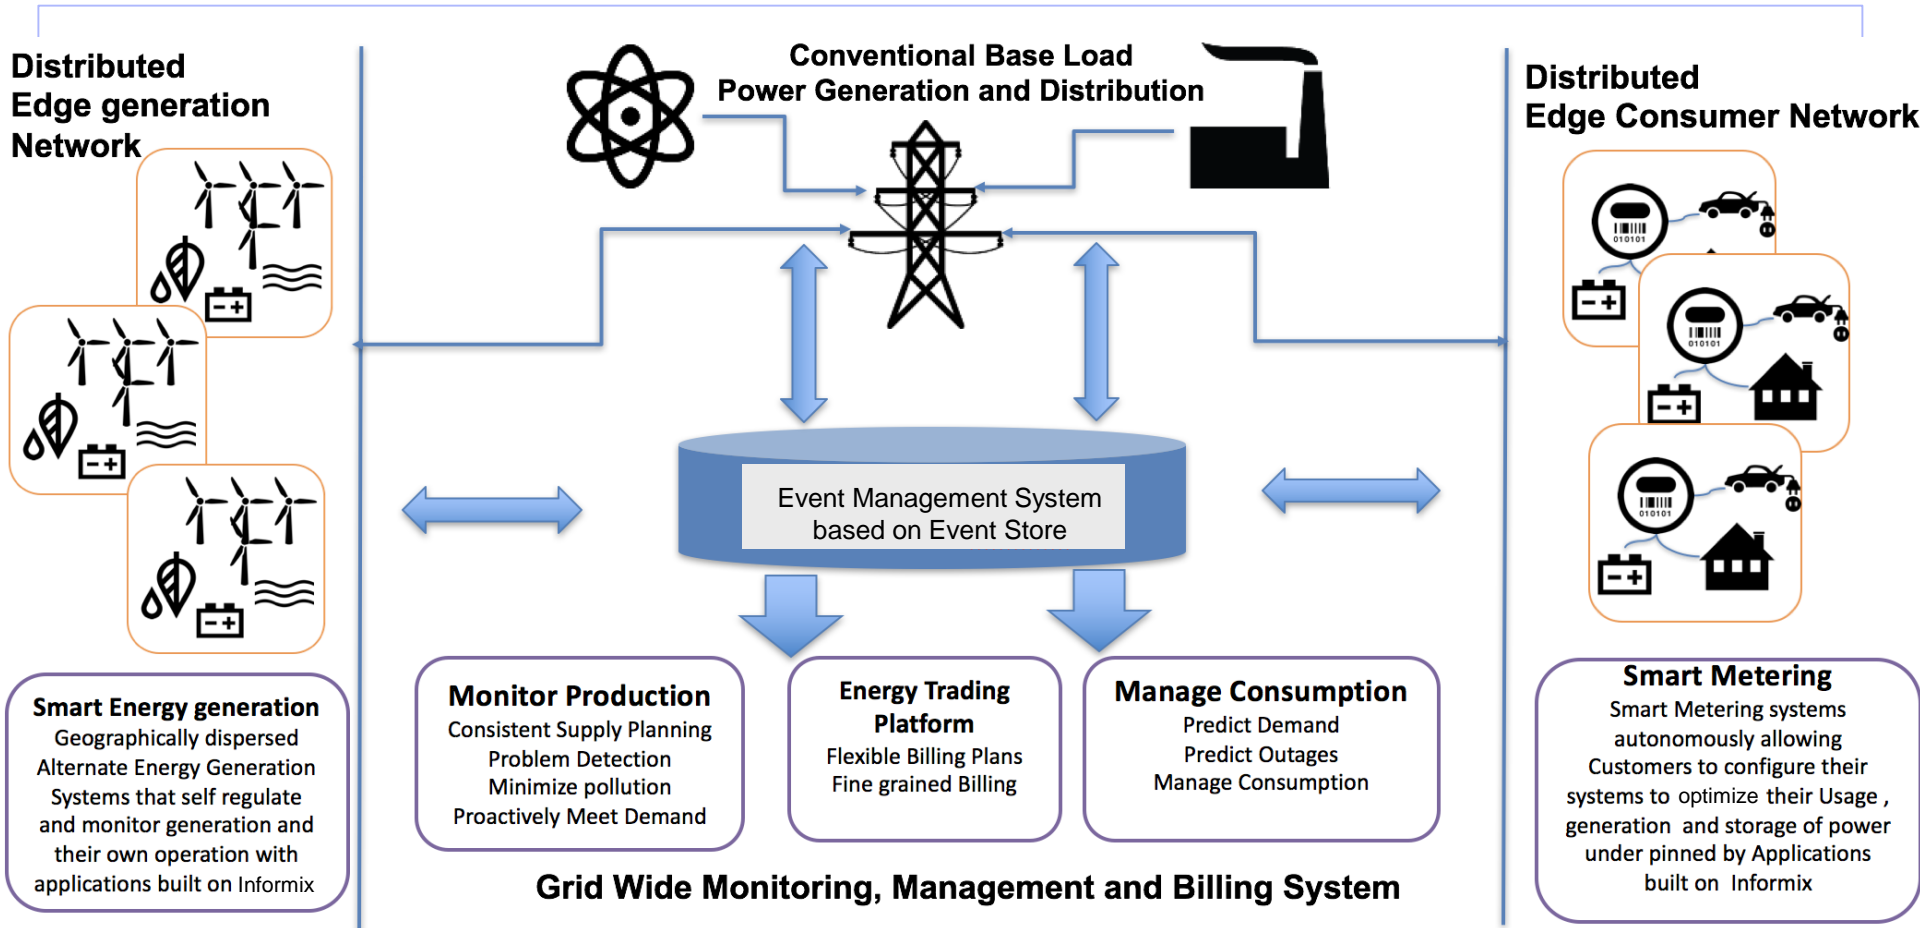
Architekturbeispiel Industrie 4.0



Architekturbeispiel Industrie 4.0



Architekturbeispiel Energie



Technologien



Hochgeschwindigkeitsdatenverarbeitung in Echtzeit



Beispiel: Connected Cars

Telematische Information:

- Geschwindigkeit, Schaltung, ABS, Airbag, Reifendruck, Lokation
- Mehr als 1500 Messwerte, einige in Millisekundenintervallen
- Wetter, locale Ereignisse
- Lokation: Verkehr
- Demographische Information, Fahrerverhalten

Möglichkeiten:

- Verständnis wie Kunden Produkt nutzen
- Wartungsplanung
- Geo-fenced Event Notification



Automotive

Beispiel: Handel

Six Billion Mobile Phone Subscribers



Retail

iBeacon-type technologies

- Brick-and-mortar
click stream analysis

Micro segmentation marketing

- Customer location and time
opportunities

IBM Db2 Event Store Characteristics

① Lightning Fast Ingest¹

- 1 Million inserts per second per node
- Ingest rate scales linearly with added nodes
- Memory-optimized indexing for fast lookups

② Real-time Analytics

- Real-time analytics over ALL ingested data
- OLAP queries faster than SparkSQL
- Integrated machine learning capabilities

③ Highly Available Atomic Transactions

- Replication for HA
- Configurable consistency
- Full insert Capabilities

④ Built for Open Hybrid Cloud

- Query elasticity through Spark
- Writes to shared storage in Parquet format
- No vendor lock-in



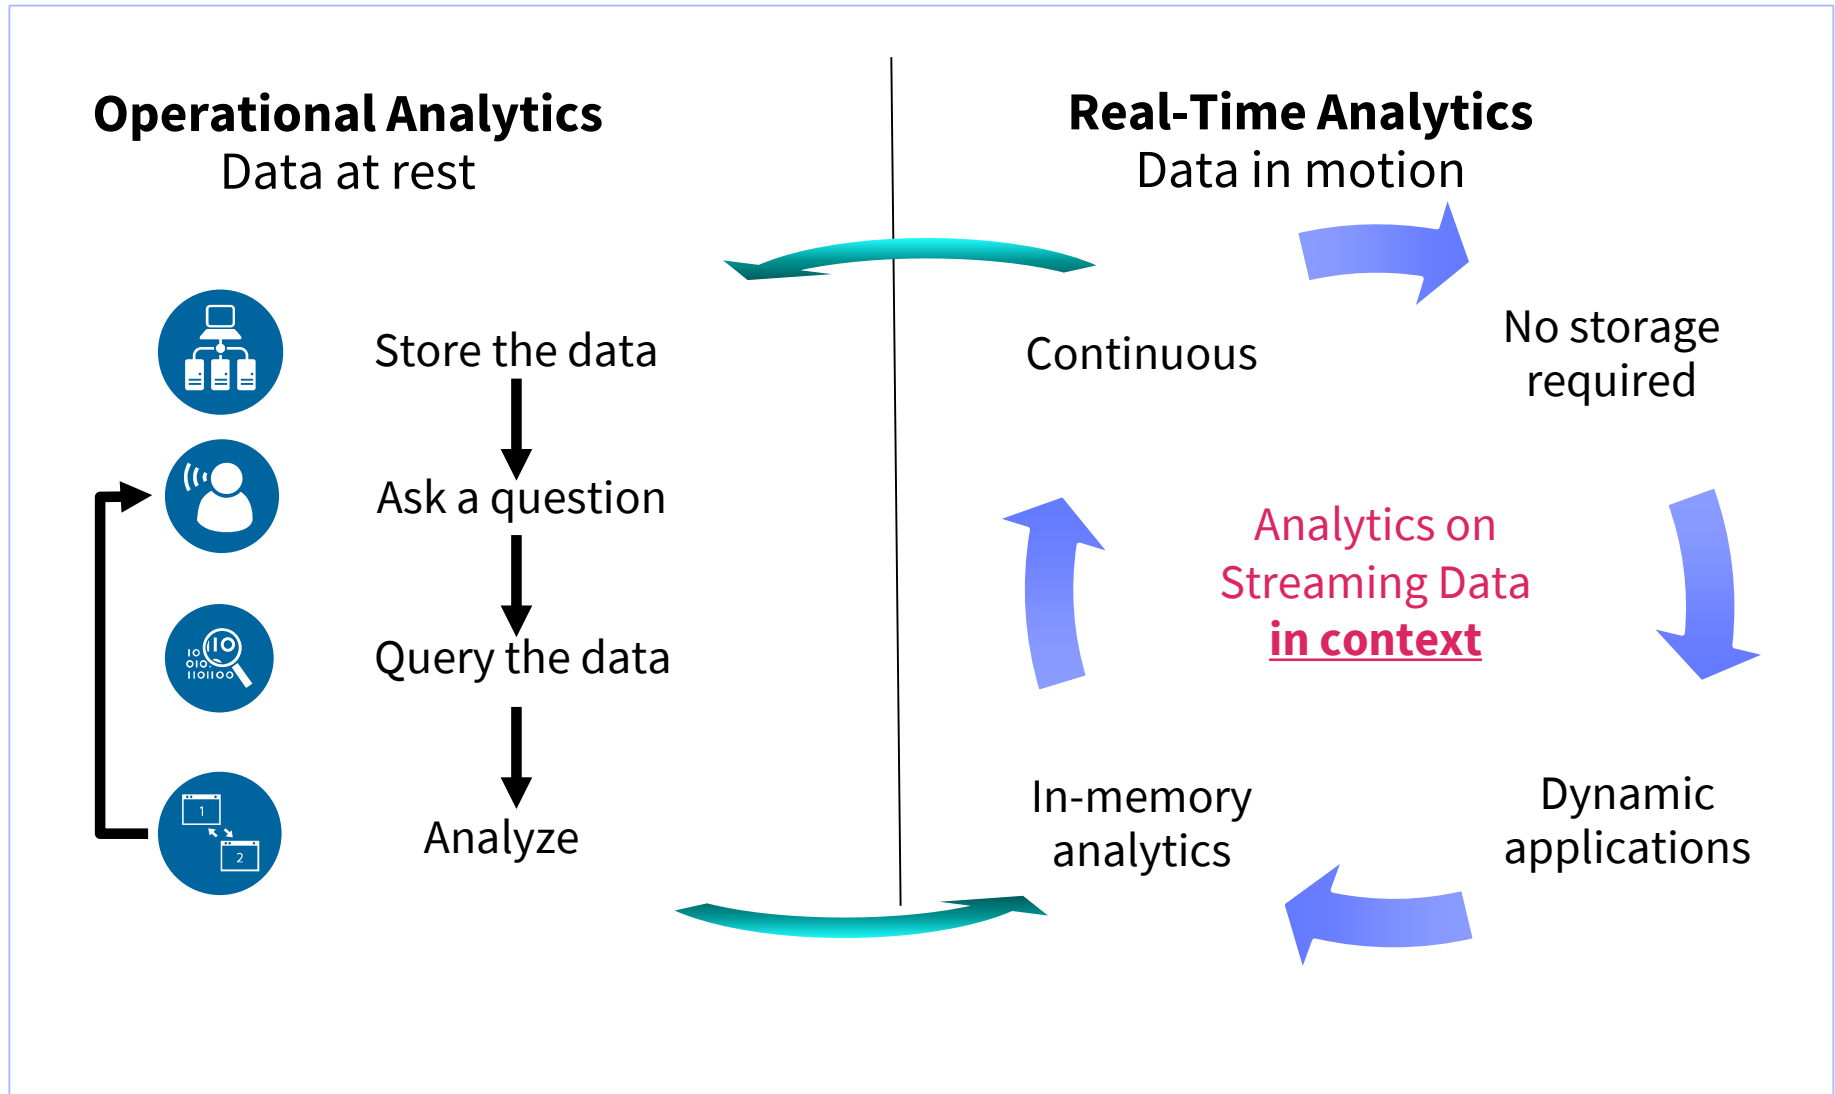
¹Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

Why Not Hadoop?

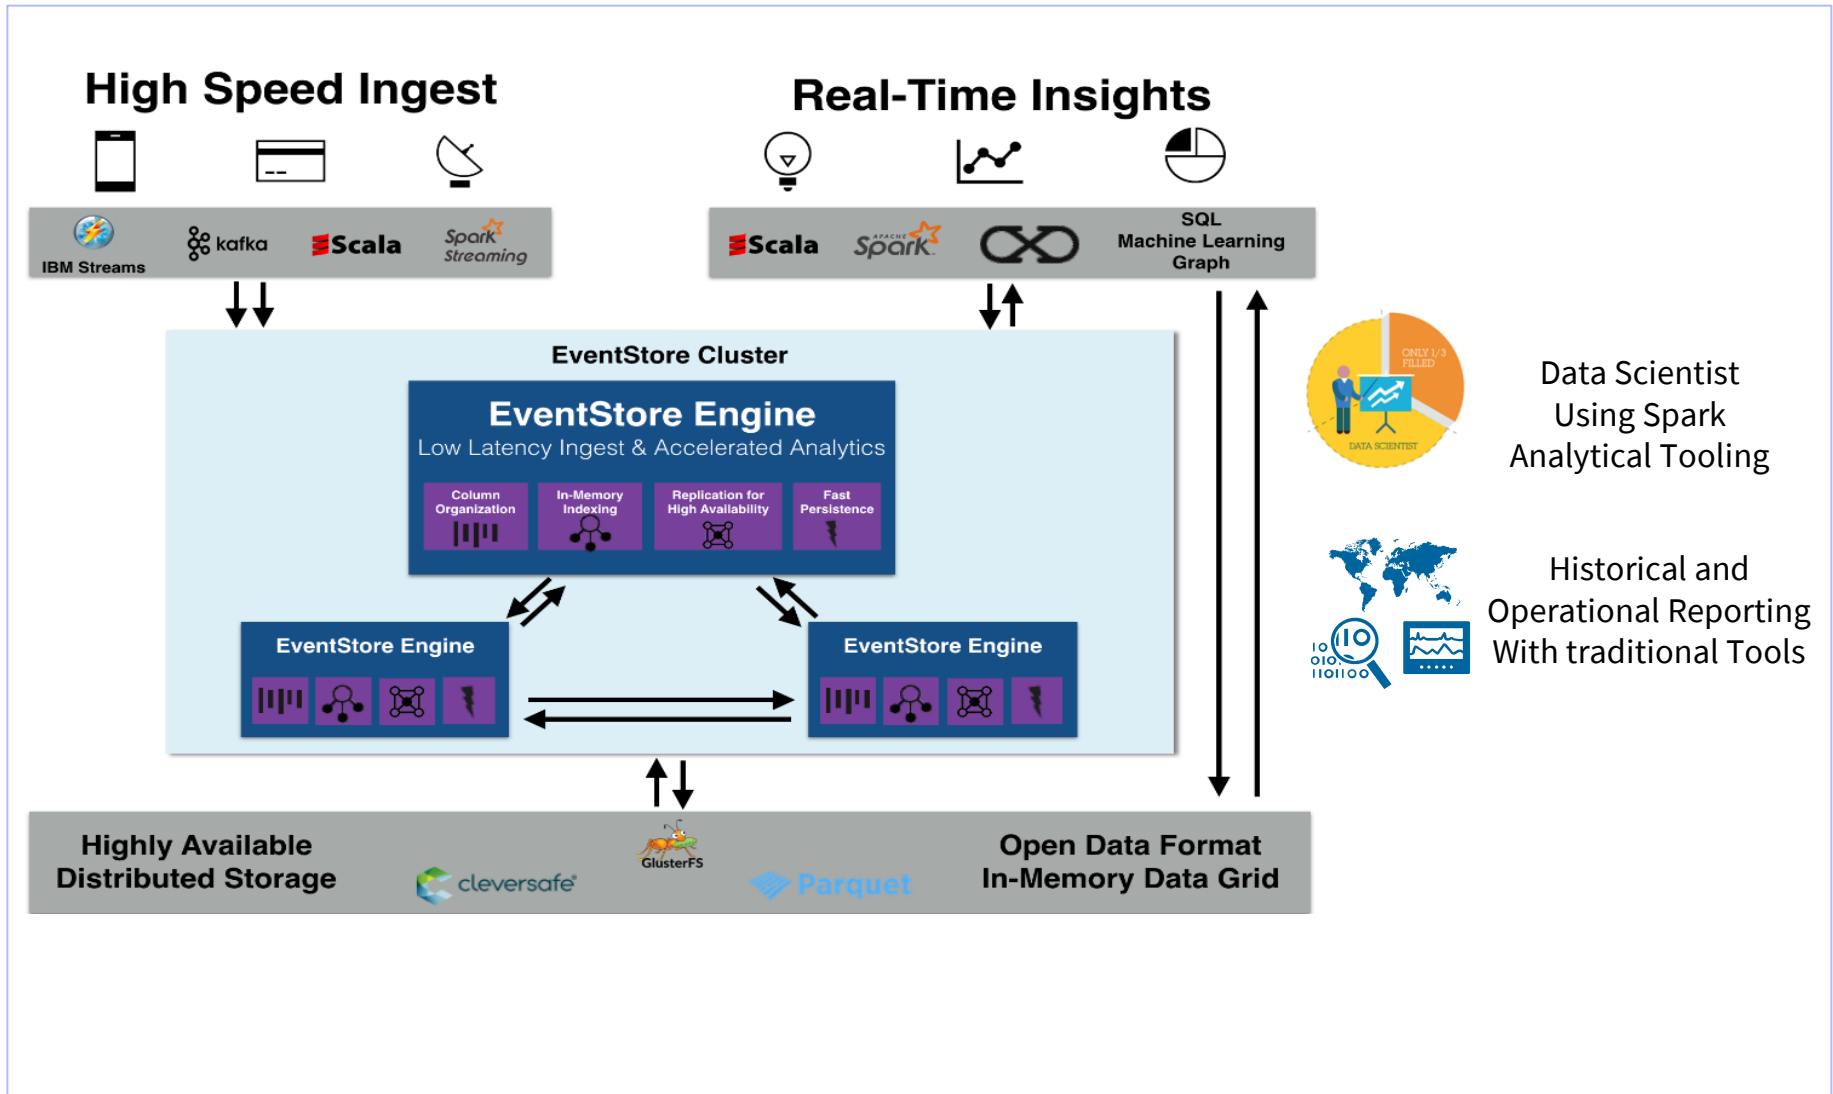
Hadoop is not designed for transactional systems
Hadoop is designed for large batch ingests

Db2 Event Store is a high-performance hybrid relational/analytics system
Db2 Event Store **continuously optimizes storage**
for high performance access

Real-Time Analytics: Db2 Event Store and IBM Streams

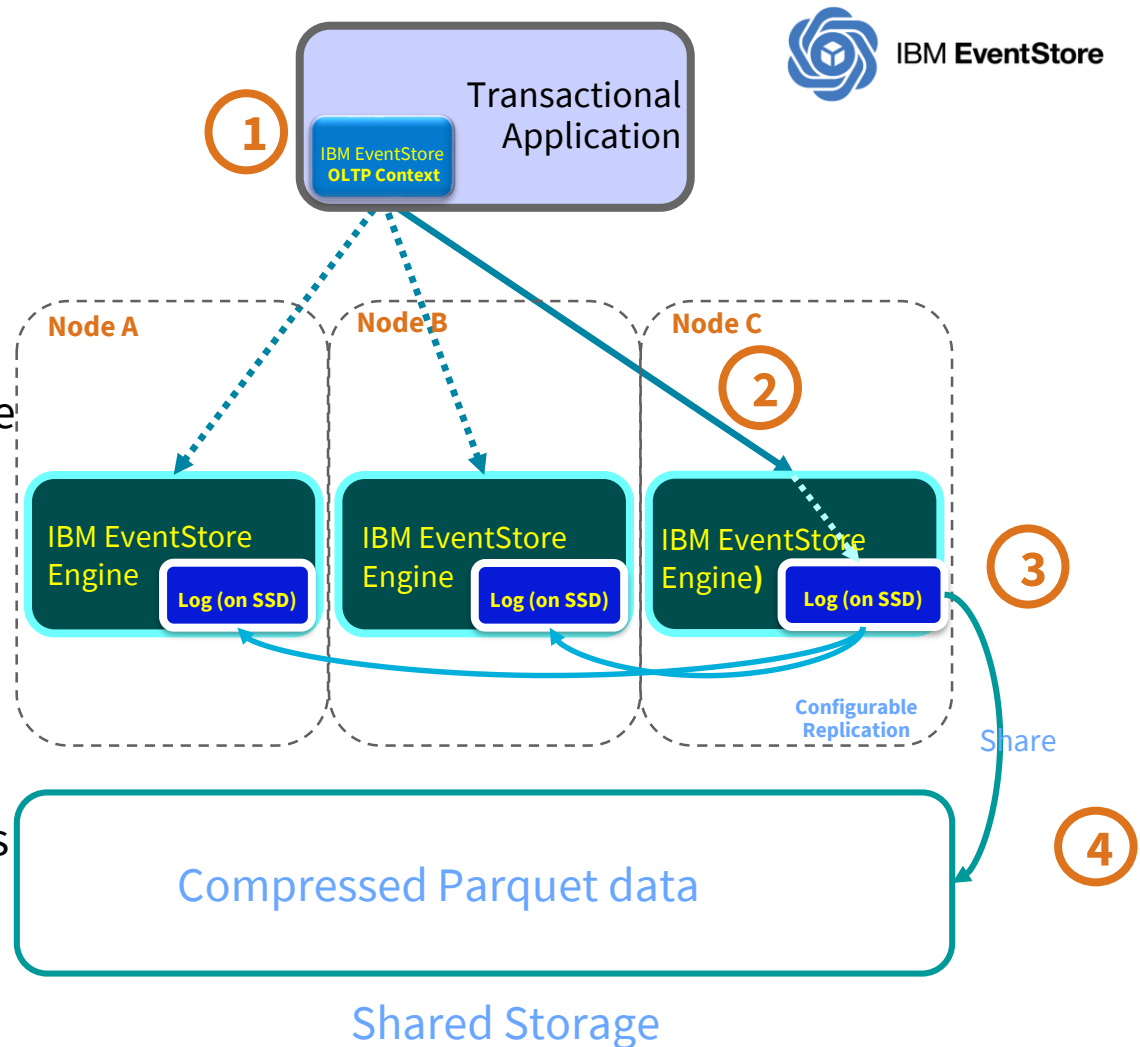


IBM Db2 Event Store System



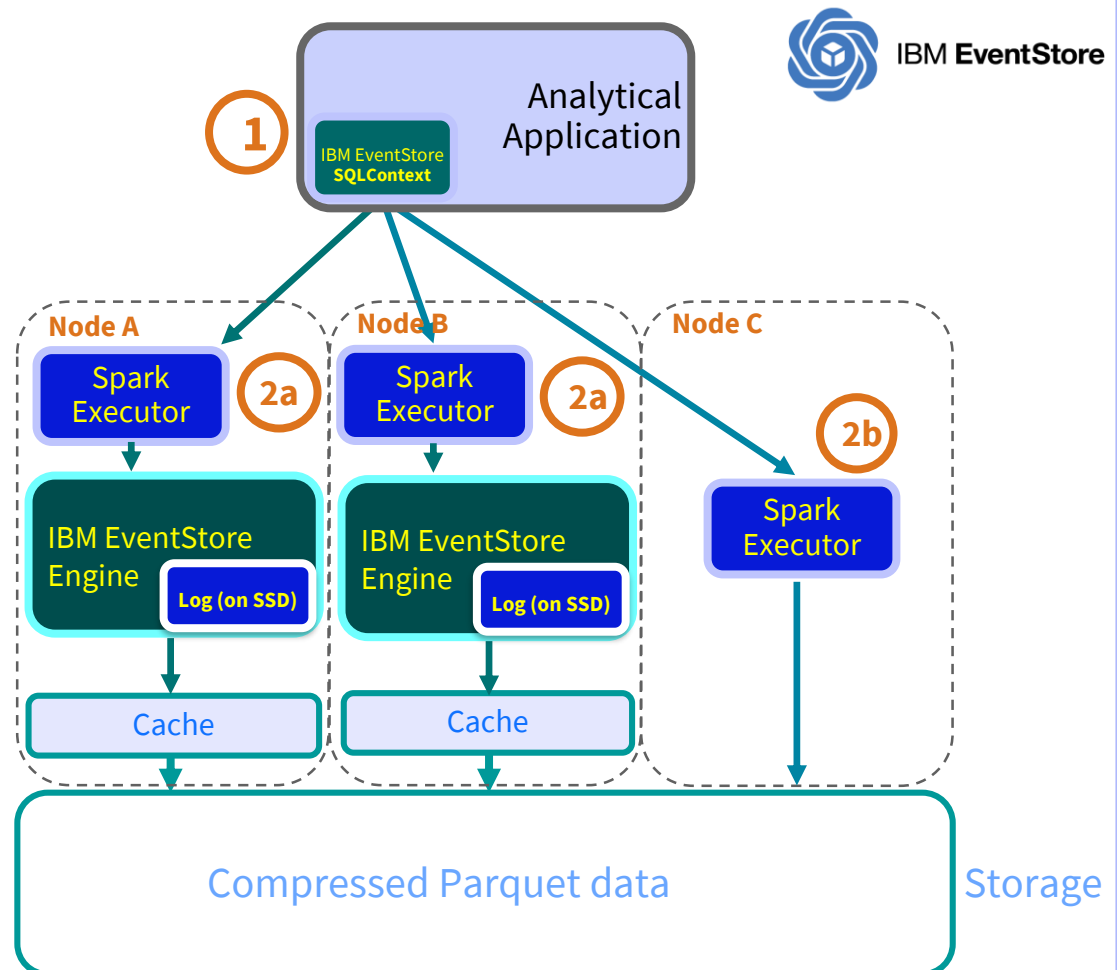
IBM EventStore Ingest

1. Ingest occurs using Scala API or streaming sink
 - Scala API accessible through Spark-friendly Context
2. Batches of rows are formed and sent asynchronously from client to the appropriate EventStore nodes
3. Rows are placed in the queryable log, replicated to replica nodes, and reply is sent to client
4. After some time, data in logs is formed into Parquet blocks and written to the storage layer



IBM EventStore Analytics


1. Analytics queries leverage the EventStore SQLContext which extends the Spark SQLContext
2. Queries are sent to either EventStore nodes, or vanilla Spark nodes depending on performance requirements and whether most recent (just inserted) data is required
 - a) Query is sent to EventStore nodes to retrieve most recent data and combine with shared data in cache or in storage layer
 - b) Query is sent to Spark node(s) to read data all but most recent data from storage



Datenvirtualisierung



Performing Analytics Today



Costly and Complex

High Latency

Does not meet Business need

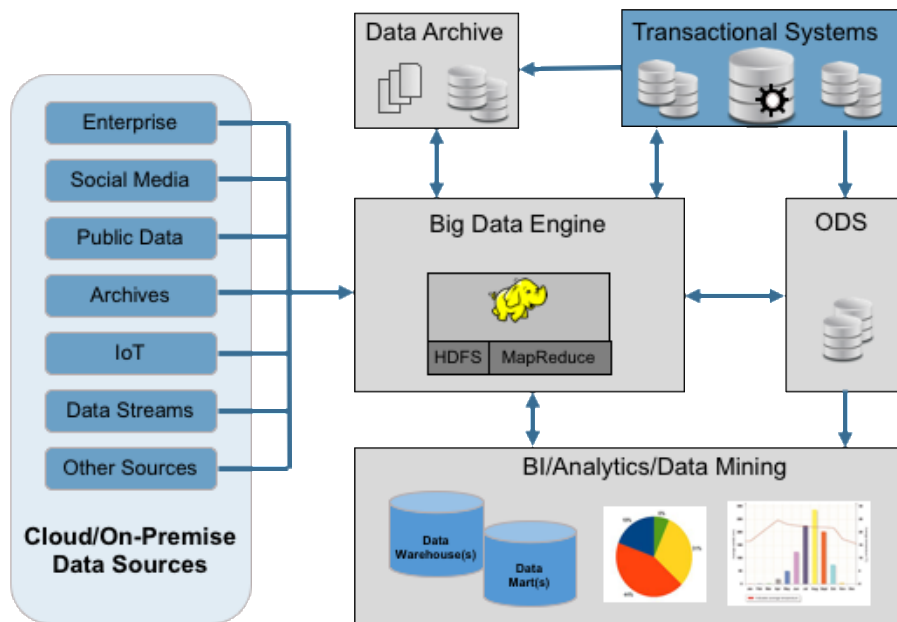
Compute resources at source not utilized

Error prone, data integrity challenges

Applications expect homogeneity

Not all data needs to be moved or copied

Resulting Data Architectures



Numerous ETLs

Unnecessary duplication,
replication

Data governance issues
accelerating

A new approach to Data Virtualization

Now in beta trial. Coming to ICP for Data in November

1 Query anything, anywhere.

Query **many heterogeneous data sources as one** across cloud, on-premise and mobile with advanced analytics using the most popular languages and tools

2

Extreme simplicity.

Automatically discover, and connect **few to many devices and data stores** into a single self balancing constellation. Avoid the complexity of centralized copies. Data only persists at the source.

3


Execution speedup.

Many times acceleration using the power of every device to compute and aggregate results.

4


Security.

Fully secure and encrypted communication and preservation of data access rights at source.



Federation vs. Queryplex Datenvirtualisierung

Classic Federation & Edge Computing




Query issued against the system

A coordinator receives the request and fans the work out to edge nodes

Edge nodes individually perform as much work as they can based on their own data. Individual results are sent back to the coordinator for final merging and remaining analytics.

Coordinator receives intermediary results from all edge nodes, merges results, and performs remaining analytics

Queryplex's Computational Mesh



Query issued against the system


A coordinator receives the request and fans the work out to edge nodes

Edge nodes self organize into a constellation where they can communicate with a small number of peers. Nodes collaborate to perform almost all analytics, not only analytics on their own data.

Coordinator receives mostly finalized results from just a fraction of nodes. Completes the final work for the query result.

Early Performance Benchmark – Query with Copy & Load

- Query includes six aggregates and grouping on 2.5 years of data, with 2,500 data points per day.
 - San Jose: Head node
 - Toronto: Edge devices
- RDBMS and Hive tests run on Intel CPU, roughly 4x faster than ARM.
- Edge/Federation and Queryplex data resides on MySQL databases inside of 36 Raspberry Pi devices, with ARM processors.
 - San Jose: Head node
 - Toronto: Edge devices



Fragen?



European IBM Informix Days 2019 im Watson IoT Center in München am 3. und 4. Juni

Vorträge vom Labor zu Informix (Zeitreihendatenbank)
Führung IoT Center

Anmeldung unter:
<https://ibm.biz/InformixDays2019>



Grazie धन्यवाद *Merci* ありがとうございます *Obrigado* 多谢
ITALIAN HINDI FRENCH JAPANESE BRAZILIAN PORTUGUESE SIMPLIFIED CHINESE

Thank You

多謝 Gracias Спасибо நன்றி ชอบคุณ *Danke* شكراً
TRADITIONAL CHINESE SPANISH RUSSIAN TAMIL THAI GERMAN ARABIC

Notices and disclaimers

- Copyright © 2018 by International Business Machines Corporation (IBM). No part of this document may be reproduced or transmitted in any form without written permission from IBM.
- **U.S. Government Users Restricted Rights — use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.**
- Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. **This document is distributed “as is” without any warranty, either express or implied. In no event shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity.** IBM products and services are warranted according to the terms and conditions of the agreements under which they are provided.
- IBM products are manufactured from new parts or new and used parts.
In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply.”
- **Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.**
- Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.
- References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.
- Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.
- It is the customer's responsibility to insure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer's business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer is in compliance with any law.

Notices and disclaimers continued

- Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. IBM does not warrant the quality of any third-party products, or the ability of any such third-party products to interoperate with IBM's products. **IBM expressly disclaims all warranties, expressed or implied, including but not limited to, the implied warranties of merchantability and fitness for a particular, purpose.**
- The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents, copyrights, trademarks or other intellectual property right.
- IBM, the IBM logo, ibm.com, Aspera®, Bluemix, Blueworks Live, CICS, Clearcase, Cognos®, DOORS®, Emptoris®, Enterprise Document Management System™, FASP®, FileNet®, Global Business Services®, Global Technology Services®, IBM ExperienceOne™, IBM SmartCloud®, IBM Social Business®, Information on Demand, ILOG, Maximo®, MQIntegrator®, MQSeries®, Netcool®, OMEGAMON, OpenPower, PureAnalytics™, PureApplication®, pureCluster™, PureCoverage®, PureData®, PureExperience®, PureFlex®, pureQuery®, pureScale®, PureSystems®, QRadar®, Rational®, Rhapsody®, Smarter Commerce®, SoDA, SPSS, Sterling Commerce®, StoredIQ, Tealeaf®, Tivoli® Trusteer®, Unica®, urban{code}®, Watson, WebSphere®, Worklight®, X-Force® and System z® Z/OS, are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at: www.ibm.com/legal/copytrade.shtml.