An abstract graphic on the left side of the slide, consisting of a dense, glowing blue cloud of particles or data points, with some brighter spots and a faint trail extending to the right.

Maîtriser ses données dans un monde hybride, multi-cloud: les secrets de l'innovation IBM

Dr. Robert Haas, IBM Research Zurich
rha@zurich.ibm.com

Un événement organisé par The logo for 'common ROMANDIE' features the word 'common' in a stylized, blue, outlined font, with 'ROMANDIE' in a smaller, solid blue font below it.

The logo for 'MEMBER OF common EUROPE' features the words 'MEMBER OF' in a small, blue, outlined font above the word 'common' in a stylized, blue, outlined font, with 'EUROPE' in a smaller, solid blue font below it.



Future of Storage @ IBM Research

Dr. Robert Haas

Department Head, Cloud and
Computing Infrastructure
IBM Research – Zurich

rha@zurich.ibm.com
<https://www.linkedin.com/in/rohaas/> •
twitter @robert_r_haas

**AI-powered Data
Management**

**Hybrid Cloud Data
Management**

Container Data Protection

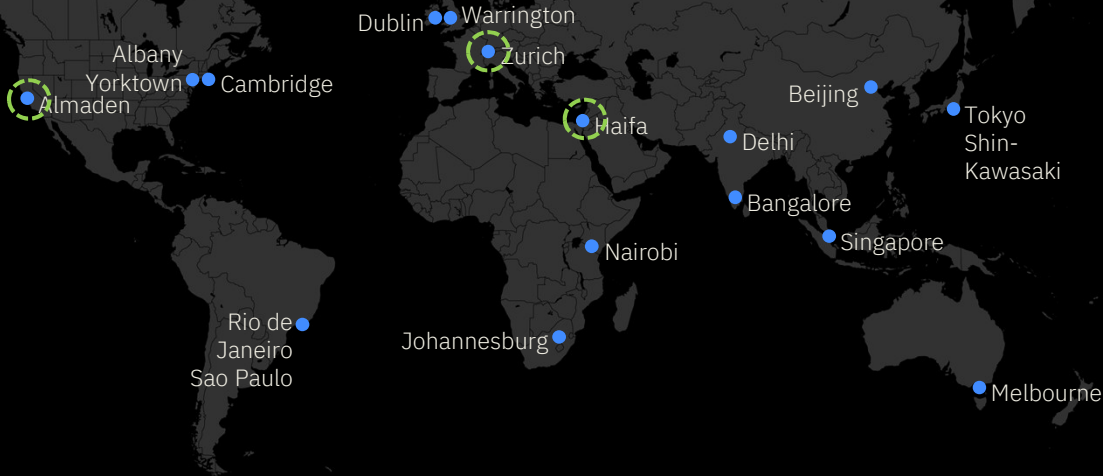
**Data Storage with Tape and
Flash**

IBM Research

3,000
Researchers

19
Locations

6
Continents



6
Nobel Laureates



10
Medals of Technology



5
National Medals of Science



6
Turing Awards

1) Tape and Flash


*Tape rocks in the cloud, and AI can hide its latency.
How we keep lowering costs of Flash while maintaining performance & reliability.*



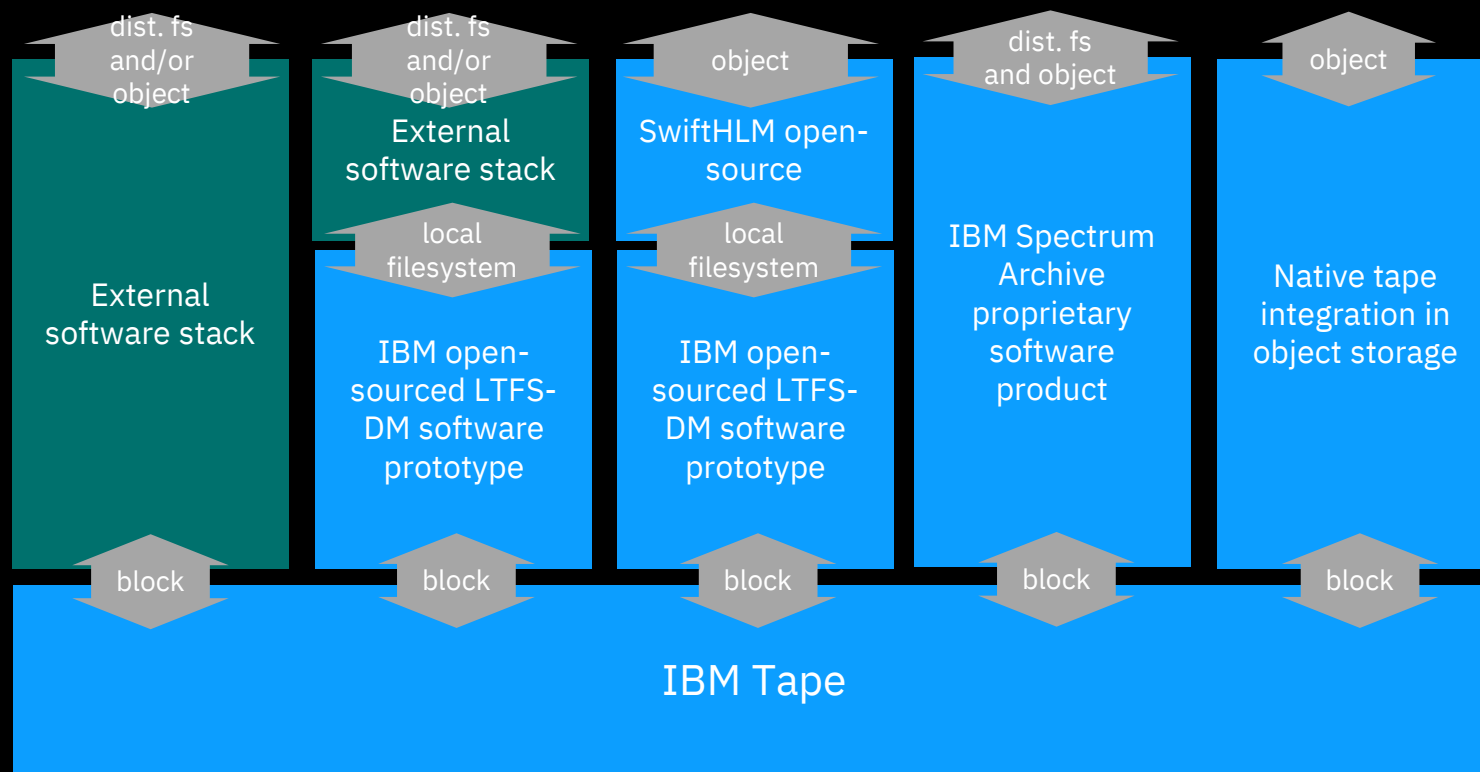
Magnetic Recording Areal Density Trends

2015: IBM-FujiFilm demonstration of 123 Gb/in² on BaFe tape

2017: IBM-Sony demonstration of 201 Gb/in² on Sputtered Tape




Tape will maintain its 5x to 10x cost advantage over other storage technologies for at least another decade.




AI-driven Predictive Caching/Tiering

1. Training based on metadata, access traces



2. Access prediction for caching/tiering




AI predictive caching/tiering:

Exploits metadata similarity in addition to access traces




Past access patterns applicable for caching and tiering new data


Evaluation



Cost and Performance Improvements

Data retrieval time (SLA) improved by 30% - 55% (for cache sizes 5% - 10%) compared to LRU






Flash Fundamentals – Cell types

SLC


1 bit per cell



100'000 P/E cycles

MLC

2 bits per cell



3'000 – 10'000 P/E cycles

TLC


3 bits per cell



3'000 – 5'000 P/E cycles

QLC


4 bits per cell



~1000 P/E cycles

PLC

5 bits per cell




? P/E cycles



The rise of QLC

50%

of the NAND Flash capacity shipped will be QLC by 2025



Sources: 1) Western Digital Data. 2) Forward Insights, NAND Quarterly Insights Q2/19, May 2019

Flash Challenges



- New NAND flash generation every 12-18 month
=> 4 generations of 3D NAND in ~5 years
- Availability of a single generation only a few years or even less than a year
- Development of an ASIC controller can take more than 2 years
- QLC has roughly 5-10x less endurance than TLC. Endurance is a challenge
- In QLC read, program, and erase latencies increased by ~1.5 – 2x compared to TLC

Techniques for enabling 3D QLC NAND Flash

Reliability




Performance



HW



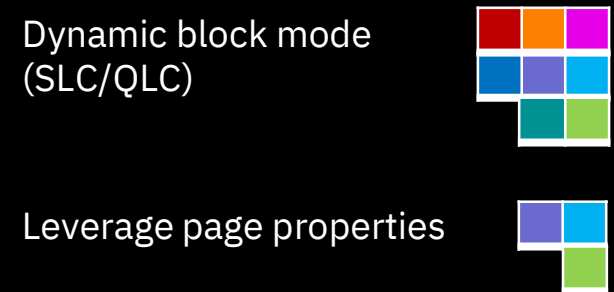
Reuse:



Change:




Invent:




Secret Sauce to Squeeze More Write Cycles from Flash

Standard approach
(wear leveling using write-cycle balancing)



Secret Sauce using health binning
and taking advantage of workload skew




2) Integrating with Kubernetes

Enable enterprise-grade storage for container environments

Storage Infrastructure In Hybrid Cloud

Composable storage services



Cloud Native Storage Services

Server Attachment

Virtual Machines

Containers

Bare Metals

Container Native Storage

Common Data Storage Services

Connect and Serve

Protect

Accelerate

Secure

Manage

Move

Storage Options

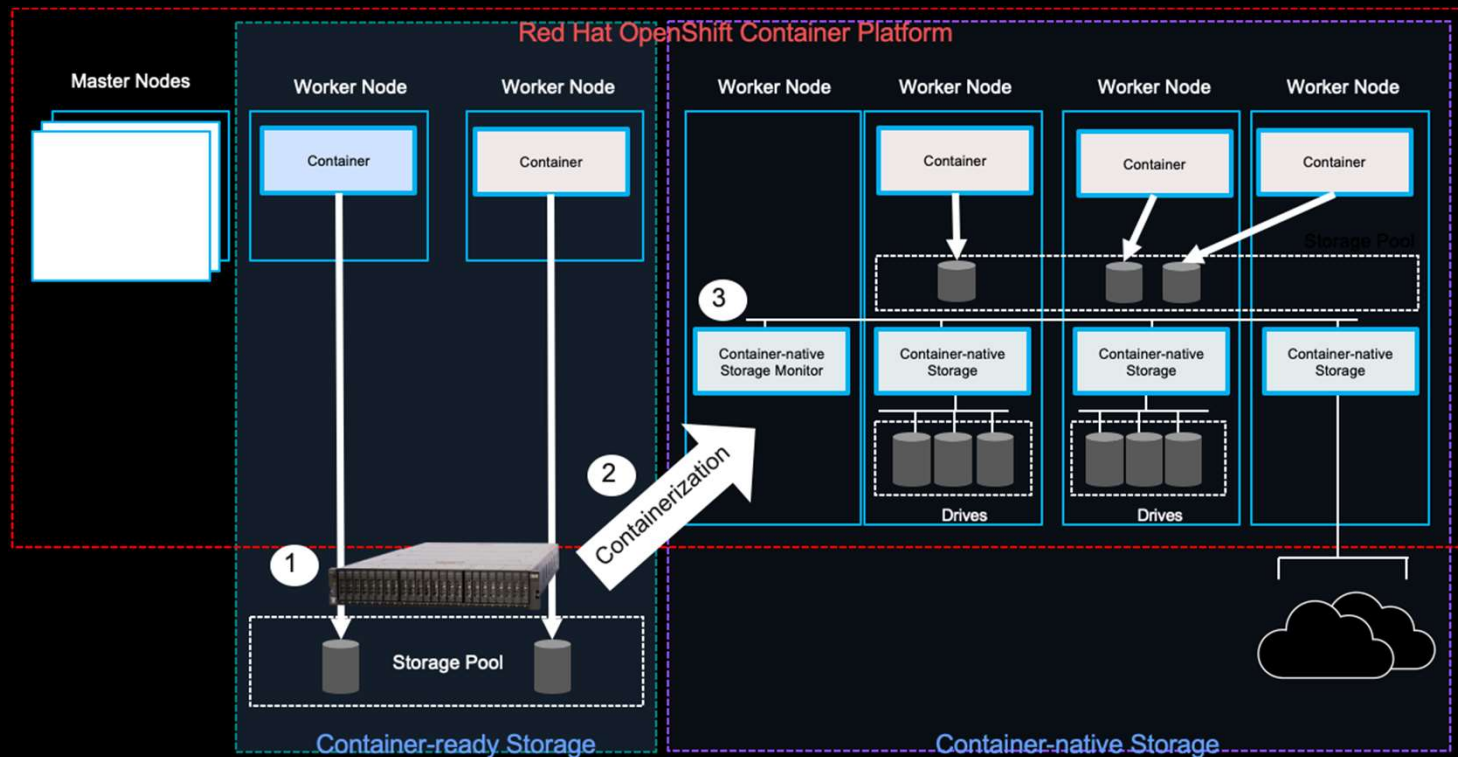
IBM Storage (Including Redhat)

3rd party Storage

Cloud Storage

Container-ready and Container-native Storage

Composable storage services in composable hardware architectures




“Connect and Serve”:
Container Storage
Interface (CSI)

- De facto industry standard
- Simplifies consumption of storage
- Provision automation
- Snapshot enablement

Cloud Native Storage
Service as part of micro-
services mesh.

Data Protection for Kubernetes

As an **enterprise developer** I need to protect my persistent container data.



VALUE:

- Ability to protect persistent data in OpenShift environment
- Self service using native kubectl command line

Supported Environment

- OpenShift 4.1 (Kubernetes 1.13)+
- Block Storage
- Container Storage Interface (CSI) snapshots for Persistent Storage Volumes

Key Users

- Enterprise / application developer
- Kubernetes administrator

Use Case

- Backup / Restore

Key Capabilities

- Persist storage snapshots in CSI storage
- Periodic offload to vSnap (optional)
- Kubernetes aware backup scheduling automation

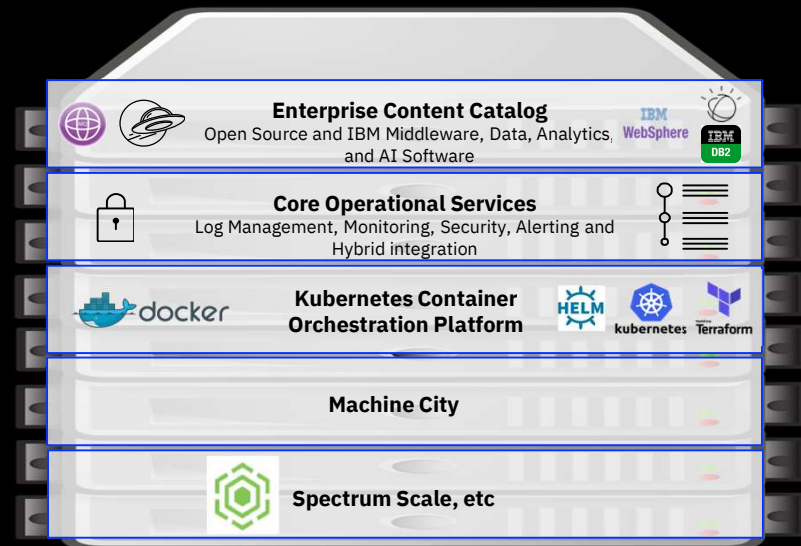
Spectrum Protect Plus

- Containers: SPP Services and SPP Data Mover
- SPP Server and SPP vSnap installed outside kubernetes/OpenShift environment (containerized in 2020)

Next Steps in Data Protection for Containers

Simplify the storage user experience for container native workloads

- **Simplify storage configuration and management**
 - Virtualization of backend storage (IBM block/file storage, open source storage, cloud storage)
 - Fully containerized, scale with Kubernetes cluster
- **Simplify data protection**
 - Continuous Snapshot (similar to Time Machine for Enterprise Storage)
 - Automatic Backup / Restore with Secondary Storage



3) Hybrid Cloud Data Management

Handle data like an asset

Data Challenge in Hybrid Cloud

Hybrid multi-cloud environments (combining premise/legacy systems, private clouds, public clouds) increase need for overall data integration and orchestration

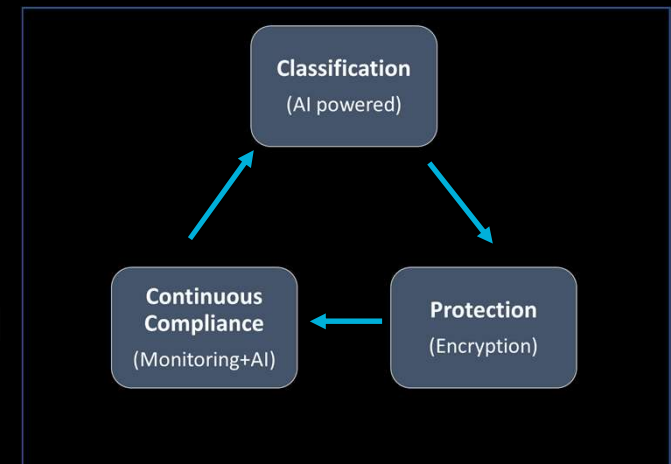
Example of a smart sales recommender system leveraging latest advances in data and AI: requires customer data from disparate sources such as:

- SaaS sales application

- On-premise order management and financial transaction systems

- Cloud-based marketing application

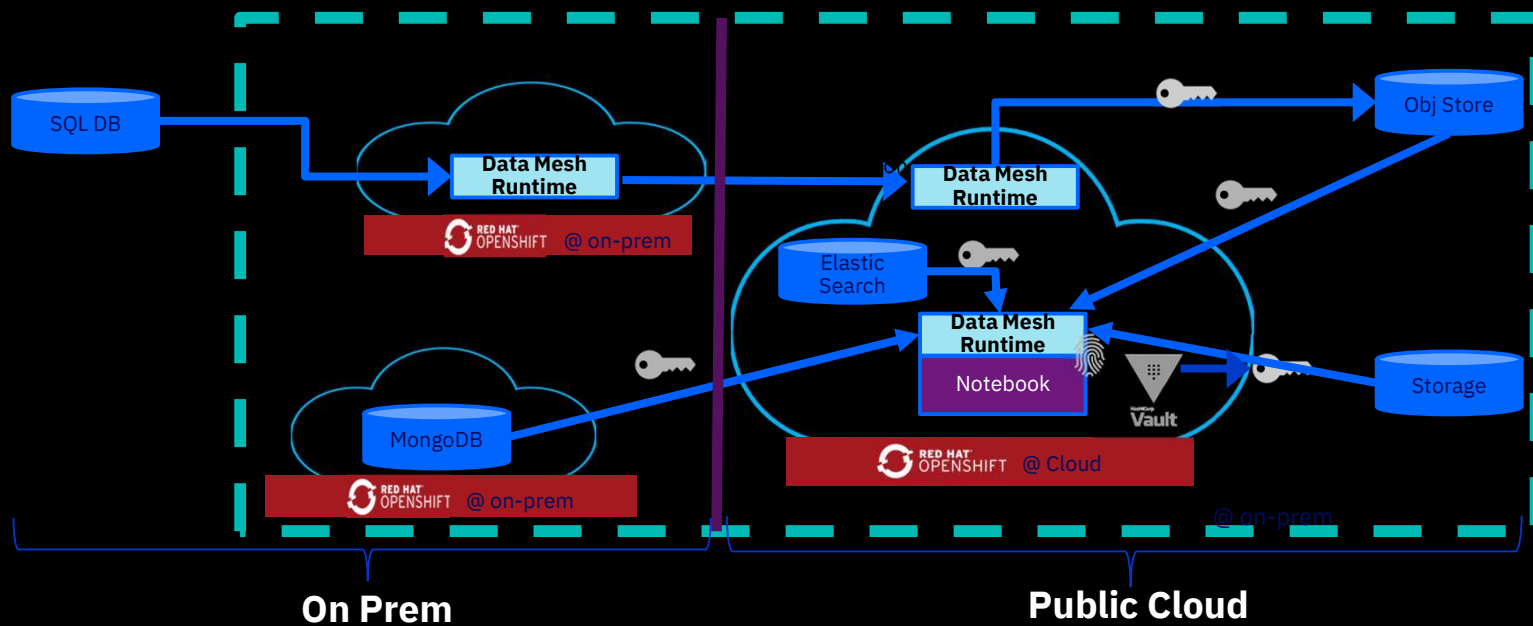
Must ensure data is treated like a proper asset when created, accessed, moved, processed, disposed



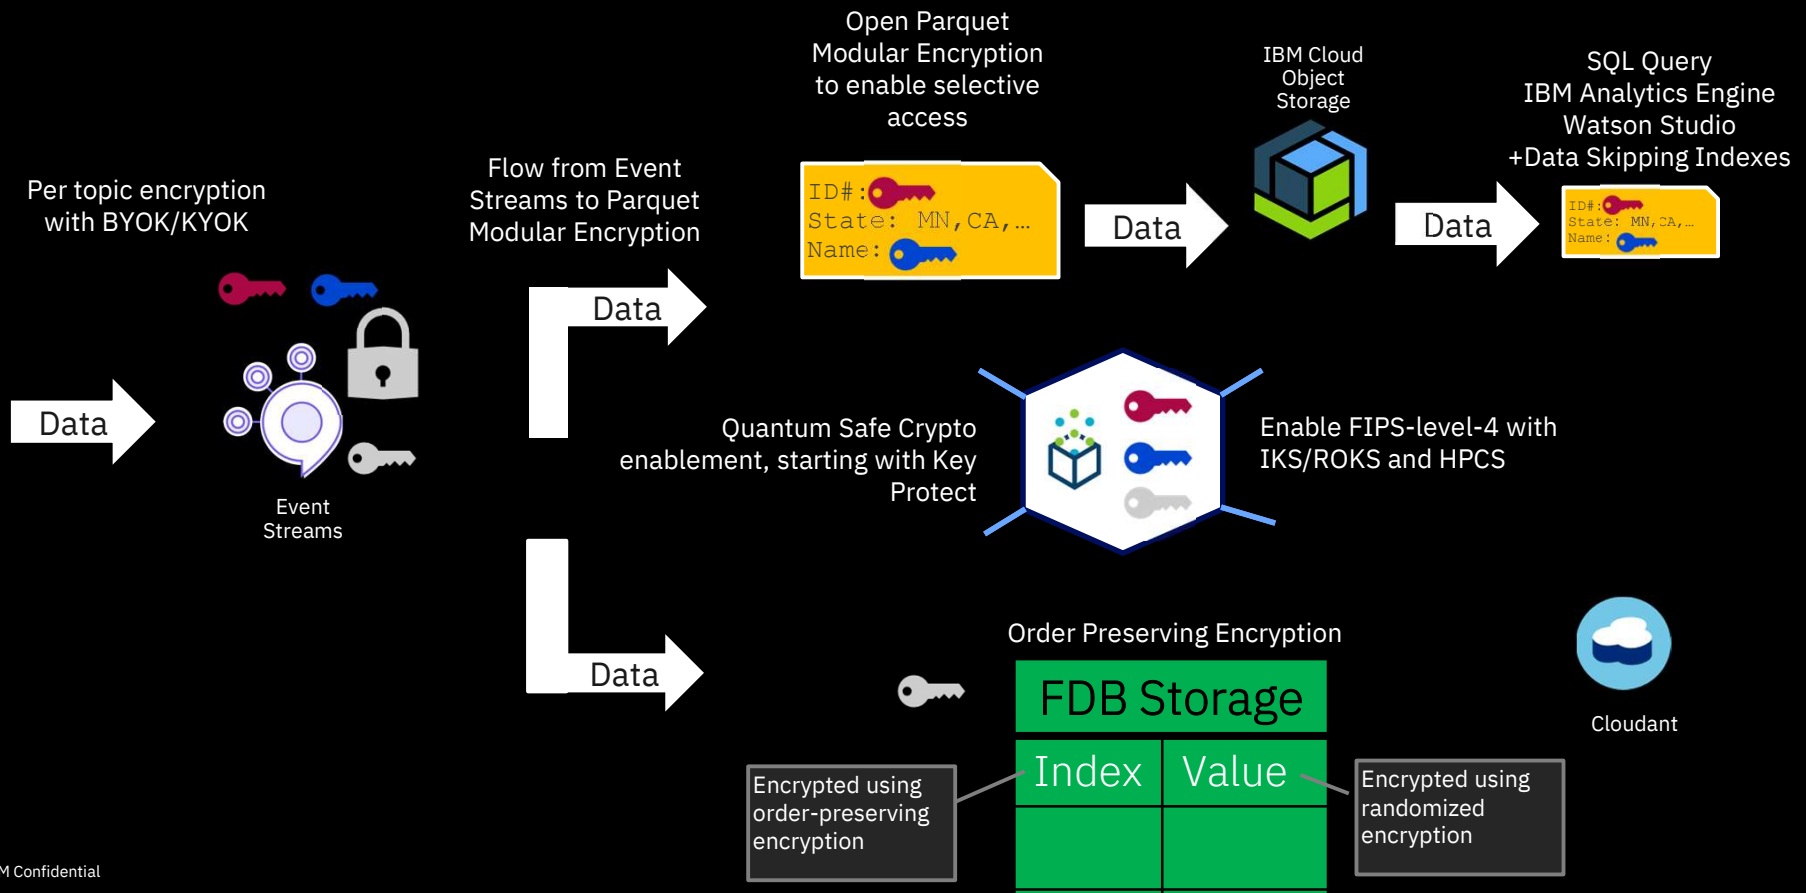
Data Management Platform

Data mesh handles

- Uniform developer experience to data
- Identification and authorization
- Governance enforcement



Enhancing Cloud Data Security



**Public cloud deployment of financial
services workloads is slowed by security
and compliance challenges**

16%

IBM Cloud: The most open and secure public cloud for business



Open innovation

- API services that are cloud delivered applications
- Kubernetes on IBM Cloud™: 1,000-plus clients, 19,000-plus clusters in production
- Major contributor to cloud-native open source work: Istio, Knative, Razeed and more



Security leadership

- Highest compliance for data encryption
- Configurable so that even IBM cannot see your data
- Edge-to-cloud threat management with security integration from IBM



Enterprise grade

- #1 VMware public cloud, with 2,000 clients
- Cloud migration for IBM Power® AIX®, IBM i, IBM Z®, SAP and mission-critical applications
- Broadest portfolio of compute instances, including Power and x86

Highest level of encryption
FIPS 140-2 Level 4

Isolation for cloud native
ROKS and containers on bare metal

No data egress charges with Cloud Databases
No vendor lock in and lower TCO

No-cost bandwidth between regions
Significantly lower TCO

Enhanced availability SLAs
HA: 99.99%, Non-HA: 99.9%

Higher SLA payouts versus market
25% of monthly at 60 minutes

Audit transparency to bare metal
Traceable serial number compliance

Full control to bare-metal level
Full admin control of compute

World's first financial services-ready public cloud with Bank of America



Good Design
Award for VPC



Good Design
Award for IBM API Connect®



Customer Choice
Award for Cloud IaaS



Stratus Award
for User Experience

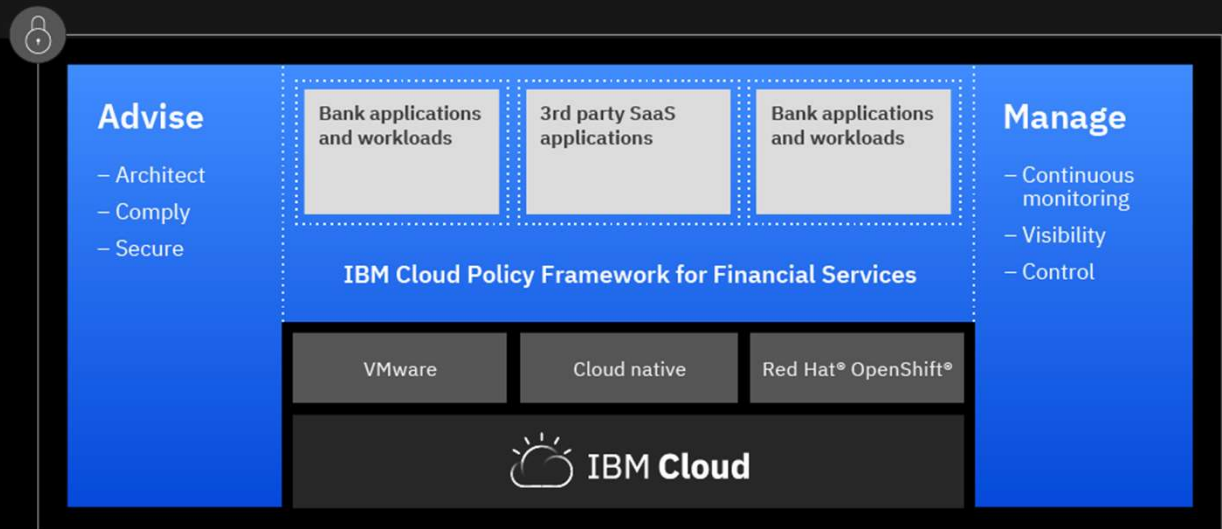
IBM Cloud for Financial Services

In collaboration with Bank of America, IBM Cloud is making public cloud financial services-ready by delivering the benefits and flexibility of a public cloud in a secure environment, enabling financial institutions, ISVs and SaaS providers to host apps and workloads in the cloud with confidence and trust.

Transformation drivers:

1. Modernization, transformation, and re-platforming (49%)
2. Business demanding more agility and/or speed from IT (49%)
3. (tie) Datacenter requiring upgrade, expansion, or replacement (45%)
3. (tie) Need functionality or services found only in cloud offerings (45%)

16% of financial services workloads are deployed in the public cloud




Continuous innovation through our global Research arm
Clients can participate through the IBM Research Cloud Innovation Lab

4) AI for Data Discovery, AI for Anomaly Detection & Prediction

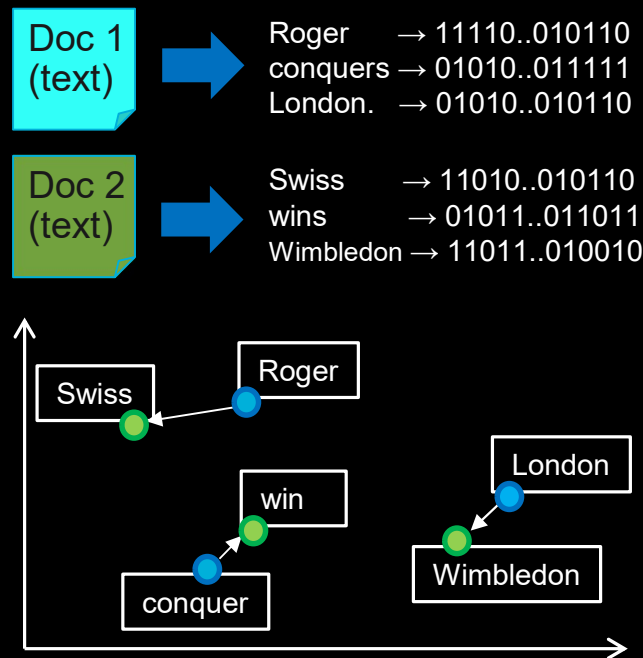
Make your data speak

Semantic Similarity Search & Visualization

1. User stores data




2. Spectrum Discover produces and stores embeddings to semantic vectors (metadata enrichment in SD)




3a. User performs semantic search

- Find documents similar to input document
- Find documents matching user query

3b. User visually explores stored data




Top-K Search Accuracy and Speed



- Dataset: 20News – 20,000 articles about 20 different topics
- Optimized version of Word Mover Distance (IBM-WMD) offers huge speed up and high accuracy

Data Clustering/Visualization

bbcsport dataset visualization



User: show me what files I have in my collection

- Semantically similar documents are clustered together
- Clusters are annotated with most important terms
 - Here a single label per cluster is shown but it can be multiple
- Helps user to visually explore the data

Search based on query document

Query Document:

QUERY document: Federer breezes into semi-finals

Label: Tennis

Roger Federer reached the last four of the Qatar Open with an easy 6-1 6-2 win over seventh seed Feliciano Lopez... Russian Nikolay Davydenko... meanwhile, upset French third seed Sebastien Grosjean 2-6 6-3 6-2. Fabrice Santoro ... France when he was forced to retire when 6-2 3-0 down to Albert Costa.. Ivan Ljubicic after the sixth seed beat Rafael Nadal 6-2 6-7 (3/7) 6-3.

Neighbor Documents:

NEIGHBOR document: Johansson takes Adelaide victory

Label: Tennis


... top contender at the Australian Open, which starts on 17 January. "I believe men's tennis is all about holding serve and if he's playing like that on his own serve I don't see how guys are going to break him...

NEIGHBOR document: Clijsters hope on Aussie Open

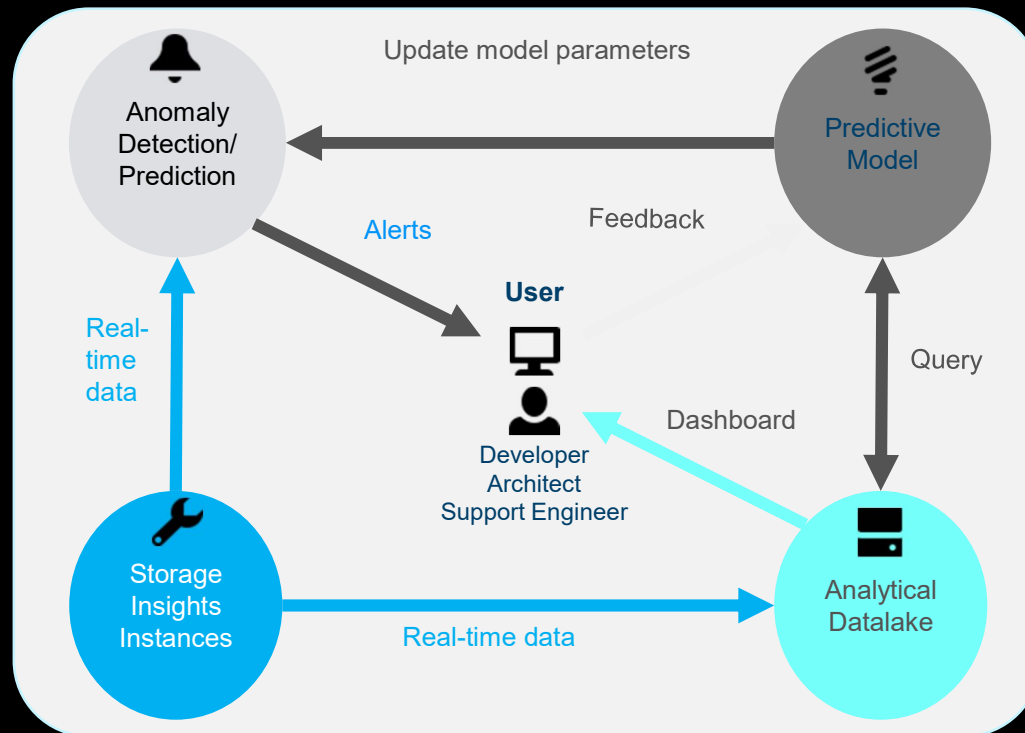
Label: Tennis

...has pulled out of January's Australian Open... Kim 17 January, Clijsters... French Open, is another absentee because of an injured left foot.

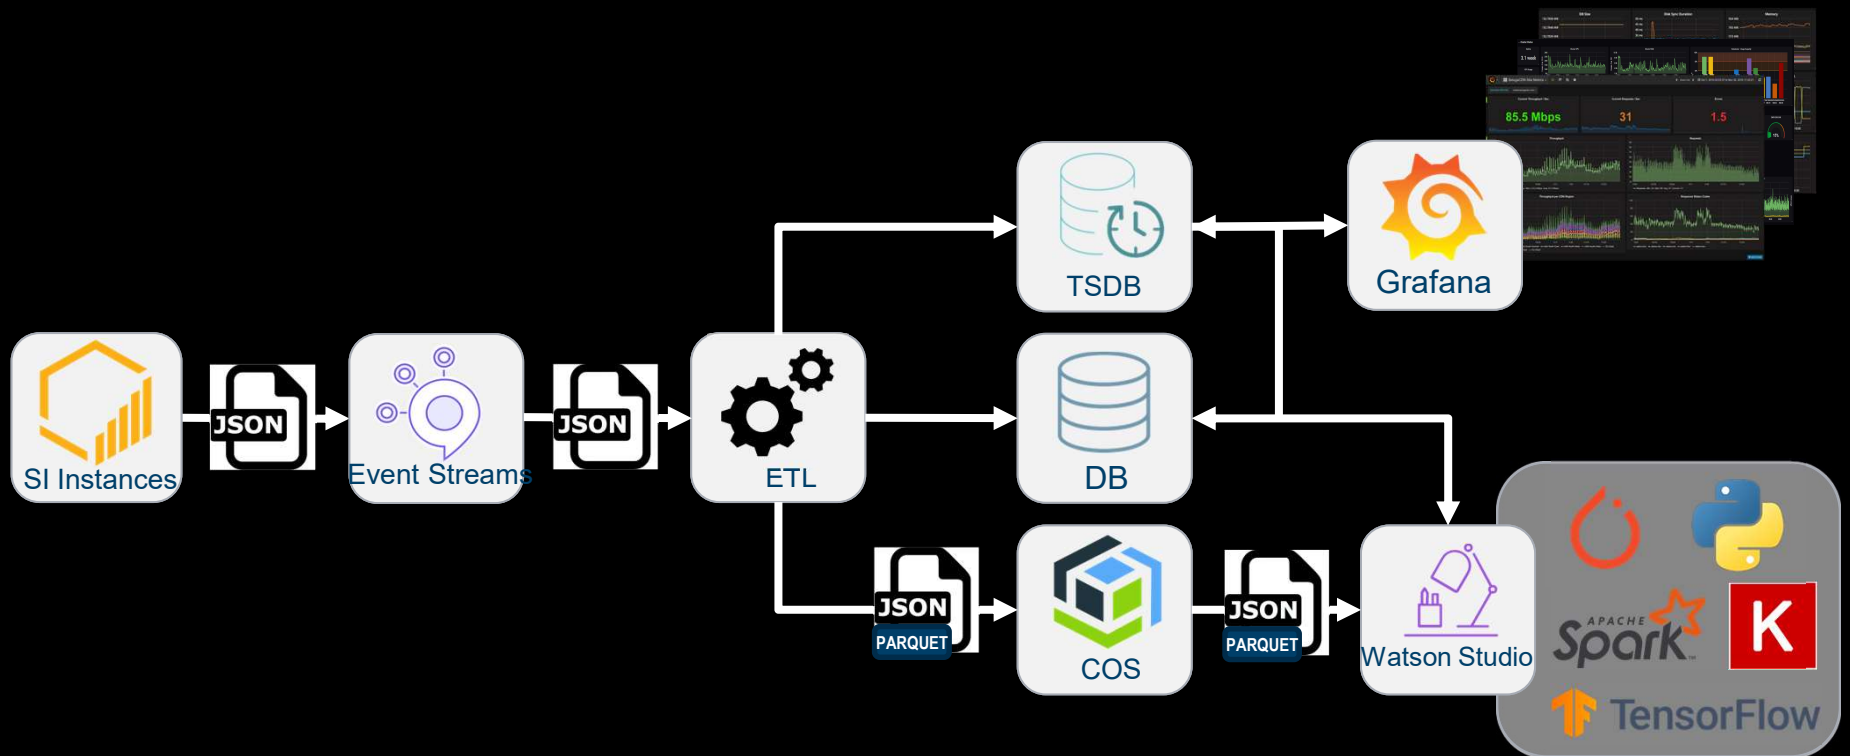
User: show me the k most similar documents to my query document

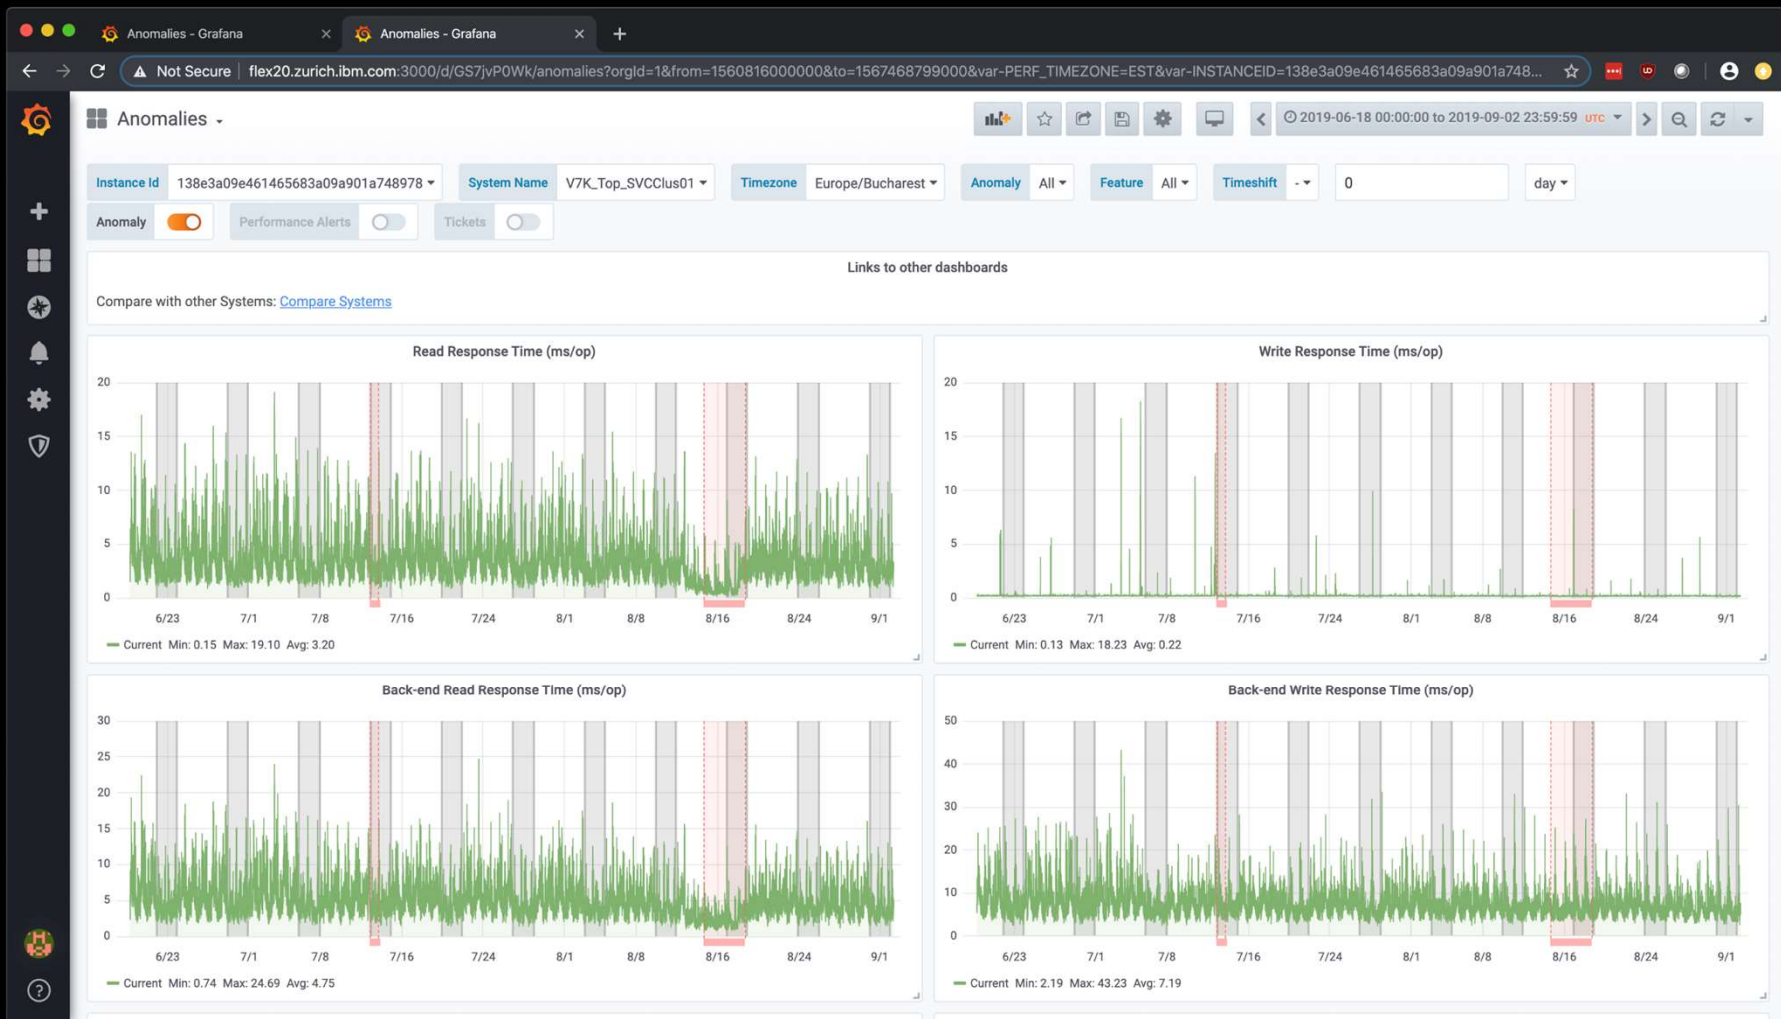


Anomaly Detection & Prediction

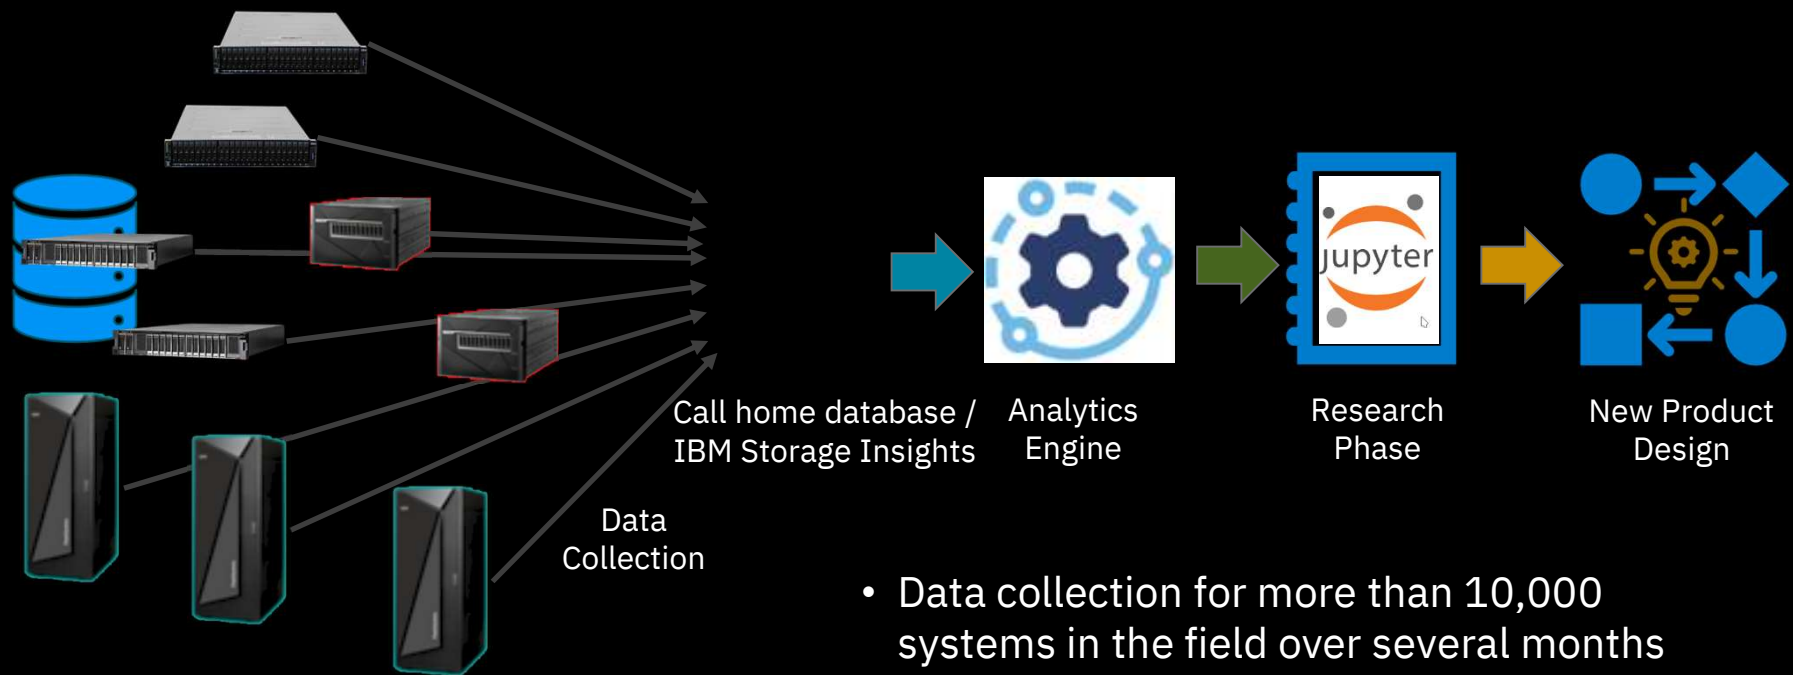


Analytical Datalake Architecture





Towards Data-driven Development



Summary

**AI-Powered Data & Storage
Management in Hybrid Cloud**

Containerized Data Protection

NVM and Reinvented Tape

Notices and Disclaimers

Copyright © 2019 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 Con't.

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.