SODA Foundation Overview

A NEW HOME FOR DATA AND STORAGE COLLABORATION

A LITTLE BIT OF HISTORY

SODA Foundation is a transformation of the Linux Foundation OpenSDS technical project. OpenSDS was initiated at the end of 2016 by Dell EMC, Fujitsu, Hitachi, Huawei, Intel, and Vodafone to address the challenges of storage integration in the cloud era.

As OpenSDS grew, there were increasing requests for it to expand its scope to support cloud native, edge, IoT, big data, machine learning, AIOps, DataOps, and more.

With the majority support of OpenSDS members, we are transforming the project into SODA Foundation. SODA Foundation is the home for open source data and storage projects.
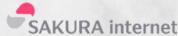

FIRST OPENSDS EVENT @OPEN SOURCE SUMMIT 2017, PRAGUE

SODA FOUNDATION

SODA Foundation is an open source project that aims to foster an ecosystem of open source data and storage software for data autonomy. SODA Foundation is organized as a Directed Fund Project under the Linux Foundation with funding from Premier Members and General Members.

The Foundation provides an open, collaborative, and neutral home for project stakeholders to coordinate their development and community investments in the SODA ecosystem.

PREMIER MEMBERS

GENERAL MEMBERS

SUPPORTERS

MISSION

To foster an ecosystem of open source data management and storage software for data autonomy
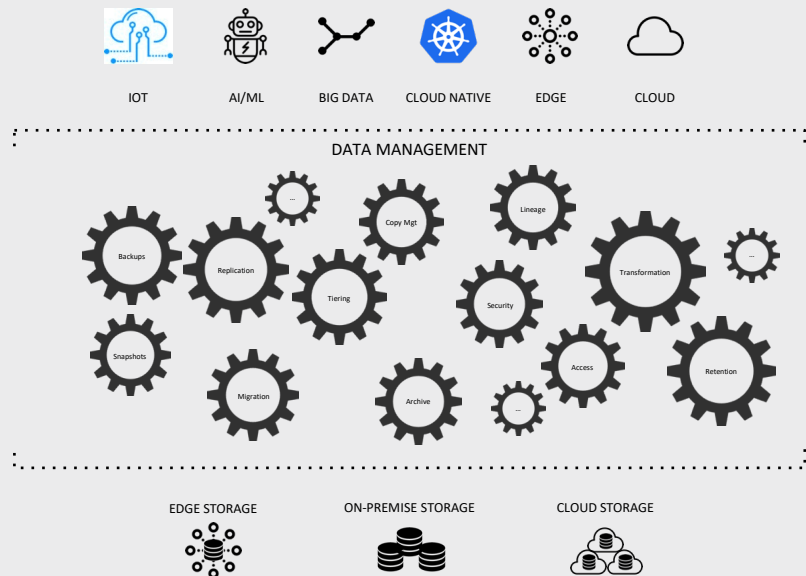
To offer a neutral forum for cross-projects collaboration and integration

To provide end users quality end-to-end solutions

NOW

FRAGMENTATION

There is no single vendor solution that can undertake the humongous task of handling all kinds of data operations from the point of data creation, to data storage, to data archive and data disposal. Challenges include virtualization, cloud native, IoT, big data, AI and machine learning, with the added complexity of data residing everywhere beyond the data center in any cloud or at the edges.
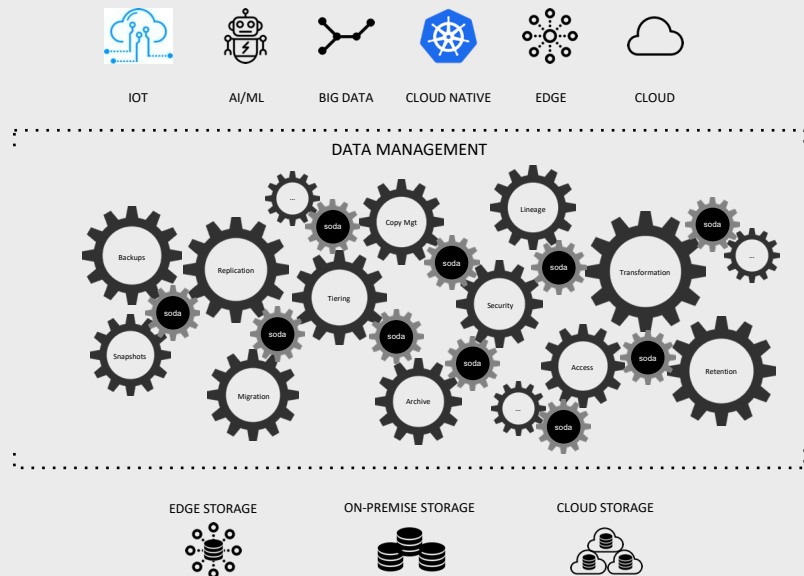
The result is that multiple vendor solutions have to be put together causing these problems:
- data silos
- non-standard interfaces
- spot solutions
- excessive data transfers
- mostly manual operations
- limited scalability

SODA

ONE DATA FRAMEWORK

SODA is a single data framework connecting disparate solutions into seamless end to end solutions. This framework is open source, allowing any developer, vendor, or end user to build and extend upon it.

The value propositions of this framework are:
- data mobility
- standardized interfaces
- interoperable solutions
- optimized transfers
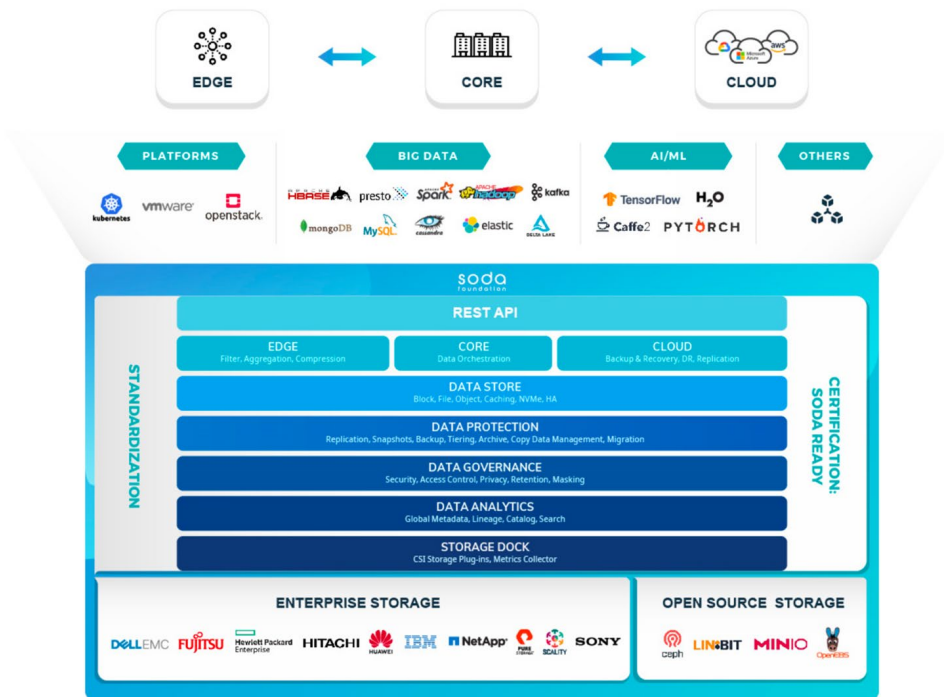- autonomous operations
- high scalability

THE SODA FRAMEWORK

SODA is an open source unified autonomous framework for data mobility from edge to core to cloud.

HIGHLIGHTS

- ⑩ open source - data and storage projects
- ⑩ standardization – data and storage management
- ⑩ ecosystem – hardware, software, solutions, services
- ⑩ certification – ecosystem components, developers, operators

SODA PROGRAMS (planning)

SODA		
PROJECTS	**SODA INCUBATOR**: Program to help projects grow	
DEVELOPERS	**SODA BOOTCAMP**: Program to help developers experience, learn and contribute to projects	
VENDORS	**SODA FOUNDRY**: Program to help vendors build and showcase SODA-based solutions in SODA Marketplace	
USERS	**SODA LAB**: Program to help users learn, test and use SODA by developing a POC solution based on user needs	
COMMUNITY	**SODA AMBASSADOR**: Program to empower members with resources and tools, with incentives to promote SODA	

EVENTS

SODA Forums
Conference Talks
Sponsorships
Meetups
Webinars

9

Governance

GOVERNANCE

SODA Foundation is overseen by the Governing Board and supported by the Technical Steering Committee (TSC), End User Advisory Committee (EUAC) and Outreach Committee (OC).

GOVERNING BOARD

Steven Tan, Chair
Futurewei

VP & CTO Cloud Solution, Storage at Futurewei where he is responsible for open source strategy and collaboration.

Yuji Yazawa, Vice-Chair
Toyota Motor Corp

Principal Engineer / Group Manager. Responsible for researching & designing hybrid cloud infrastructure for connected vehicles

Xin Zhong, Treasurer
China Unicom

CTO, China Unicom. He is in charge of team that is working for public, private and multi-cloud management products.

Rakesh Jain
IBM

STSM and Researcher - Cloud, IoT, Analytics at IBM Research. Working on setting the strategic vision and technical architecture.

Shinya Hamano
Fujitsu

Manager, Fujitsu. In charge of the planning and development of new IT platform services in the fields of storage, cloud-native, AI, and IoT and open source technologies.

Kei Kusunoki
NTT Communications

Technology Development, Storage Engineer - NTT Communications. He has worked on the storage service/solution for the telecom carrier.

Anjaneya Chagam
Intel

Principal Engineer and Chief SDS Architect, Intel. He is responsible for developing software-defined storage strategy, architecture, and technology initiatives.
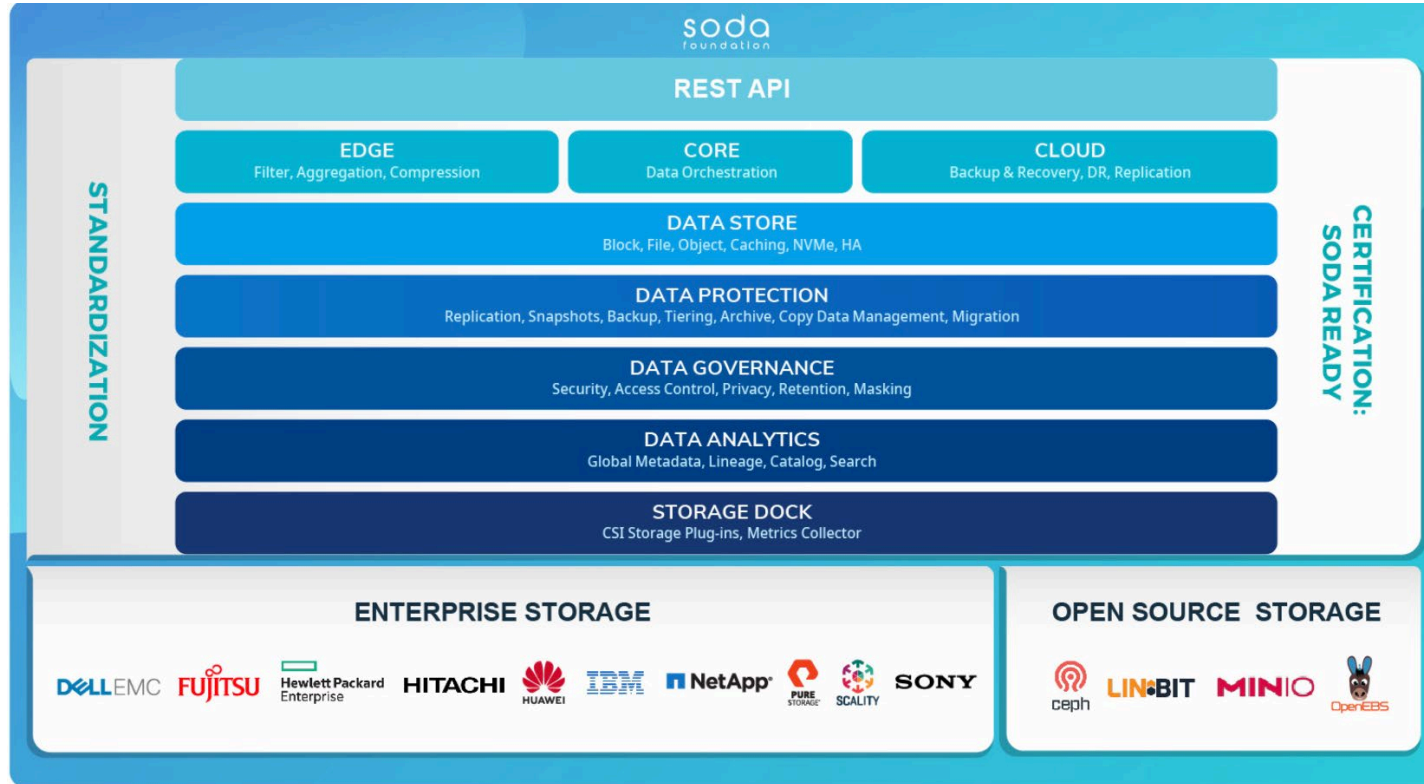
Yoshiya Eto, Advisor
Fujitsu, Linux Foundation Fellow

Principal Architect Processor System BU, Fujitsu Ltd.
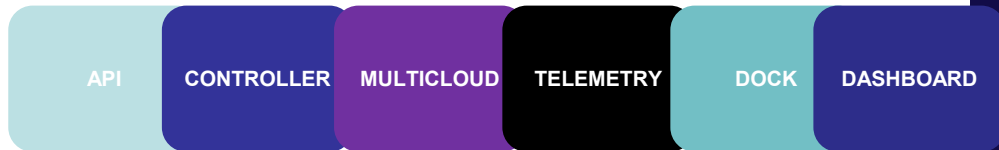
SODA Projects

SODA ARCHITECTURE

SODA PROJECTS

SODA Projects make up SODA as an open source offering

CORE PROJECTS

Core Projects are initiated by the SODA community. They form the core SODA framework and follow the SODA release roadmap. Core Projects are governed by the TSC.

| API | CONTROLLER | MULTICLOUD | TELEMETRY | DOCK | DASHBOARD |

NATIVE/Member PROJECTS

Native Projects are projects donated to SODA Foundation and are either initiated by SODA members or 3rd-party organizations. Native Projects are supported by the SODA Foundation but are governed independently. They either follow the SODA release roadmap or have their own release roadmap.

Member Projects (incubated)

Member Projects (to be incubated)

YiG

LIN STOR   ZENKO

SODA Use Cases

Use Case: Cloud Native Storage

Heterogeneous Backup, Replication, Migration, Tiering, Archiving, …
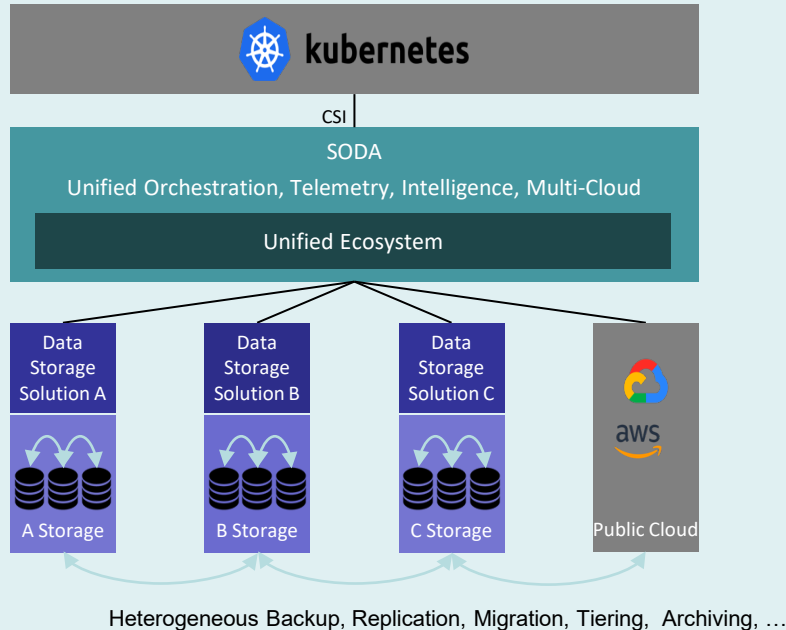
Vodafone, KPN, Yahoo Japan,

China Construction Bank Fintech

BACKGROUND

Implementing a cloud native strategy

ISSUES

Using storage solutions from multiple vendors

REQUIREMENTS

Must be a open, extensible

Must support public clouds

SOLUTION

Profile-based storage provisioning for Kubernetes

Flexible storage resource allocation

SODA Kubernetes Use Case
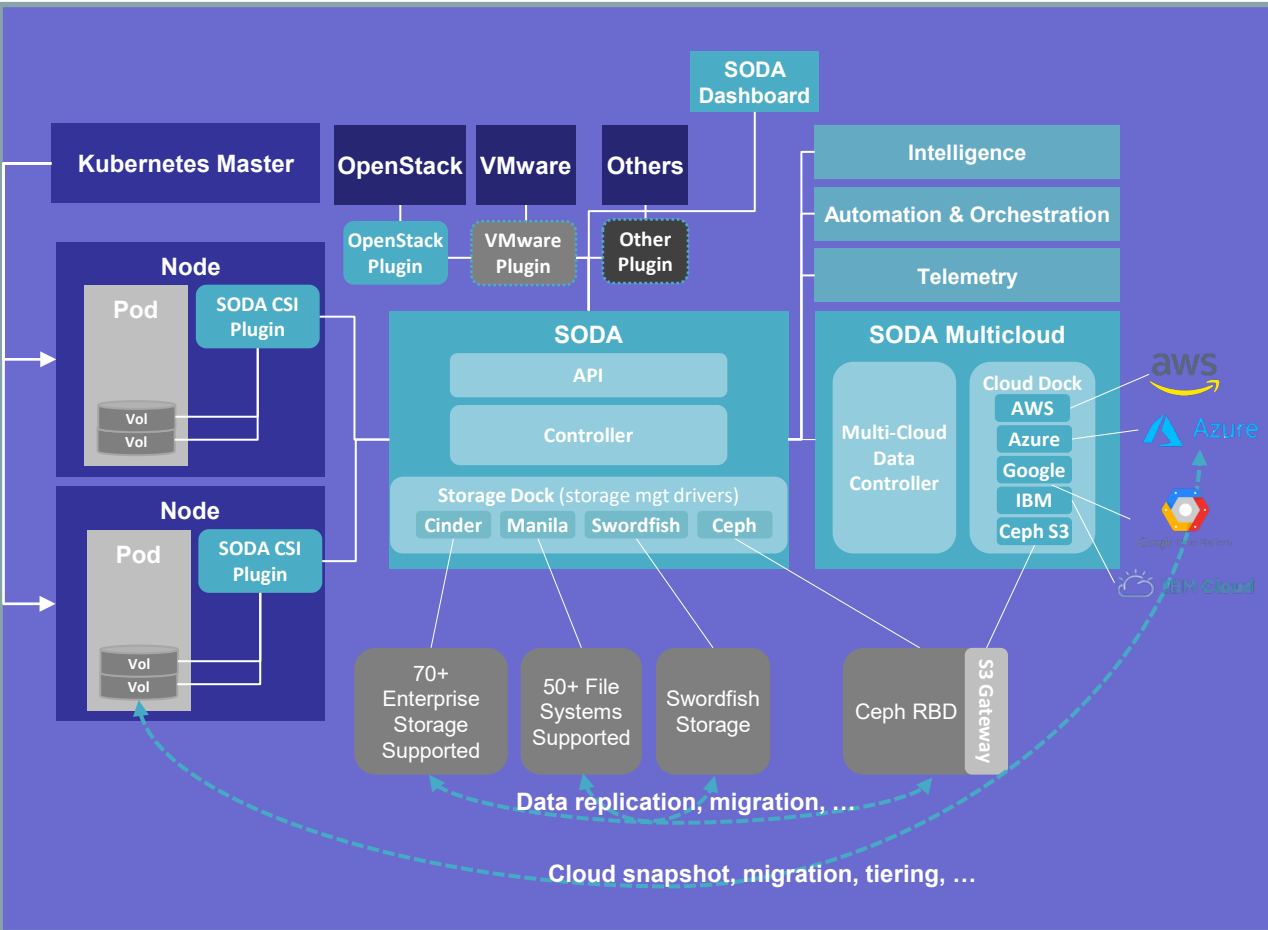
BENEFITS

- One plugin for each framework
- Add new storage with same plugin
- 100+ block and file storage supported
- Data ops across storage systems
- Common data protection, security, etc.
- Monitoring and analytics is aggregated
- Extend across data centers and clouds
- Common platform for K8S and others
- Automate workflows

Use Case: Data Lifecycle

Toyota

BACKGROUND

Car data sent to edge for edge computing, data move to primary storage for big data, to secondary storage for ML, and to archive storage

ISSUES

Need a platform for lifecycle management from edge to DC, to cloud
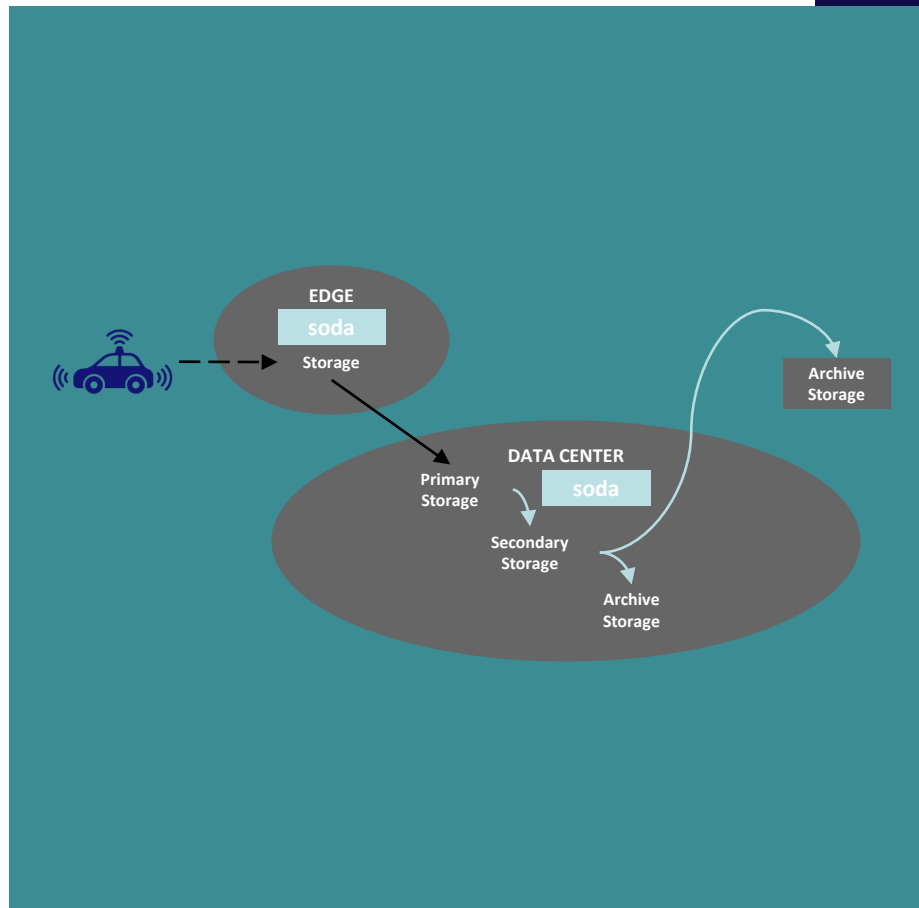
REQUIREMENTS

Data lifecycle based on age

Supports edge to cloud with customizable storage tiers

SOLUTION

Policy-based data lifecycle with age

Tiering across edge, data center and cloud

EDGE
soda
Storage

Archive Storage

DATA CENTER
soda
Primary Storage
Secondary Storage
Archive Storage

Use Case: Data Lake
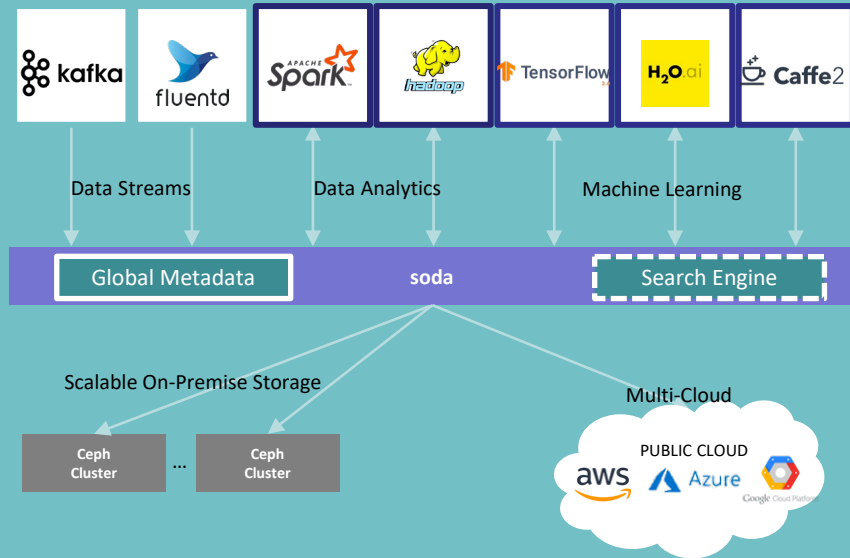
China Unicom

KEY

- A SODA Data Store project

- A collaboration between SODA Multicloud + China Unicom YIG project

- Global metadata single namespace

- Scalable on-premise storage using Ceph clusters

- Hybrid/multi-cloud storage

- Common S3 interface for data streaming, analytics, ML

- S3 Select in development

- Search Engine planned

- Status: In development, release Q2/Q3 2020

Use Case: Multi Cloud
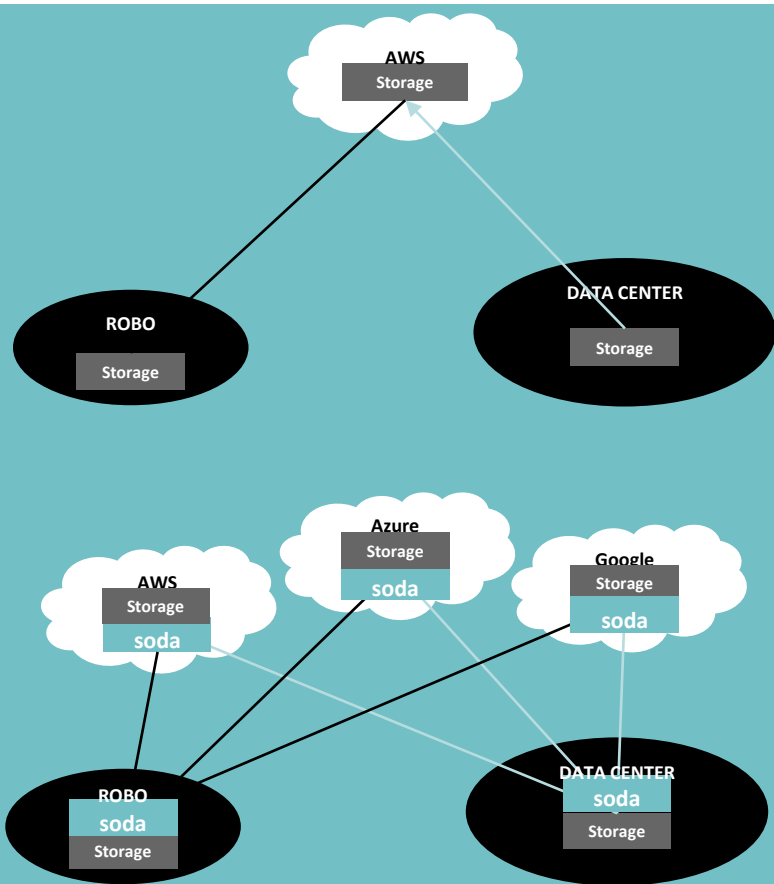
NTT Communications

BACKGROUND

Uses vendor specific solution for cloud snapshots and restore

ISSUES

Snapshot and restore on single vendor storage

Supports AWS only

REQUIREMENTS

Multi-vendor storage snapshot and restore

Snapshot and restore to/from multiple clouds

SOLUTION

Snapshot/restore with any storage

Policy-based multicloud data control

Year 2020 Plan

Adoption

- China Unicom
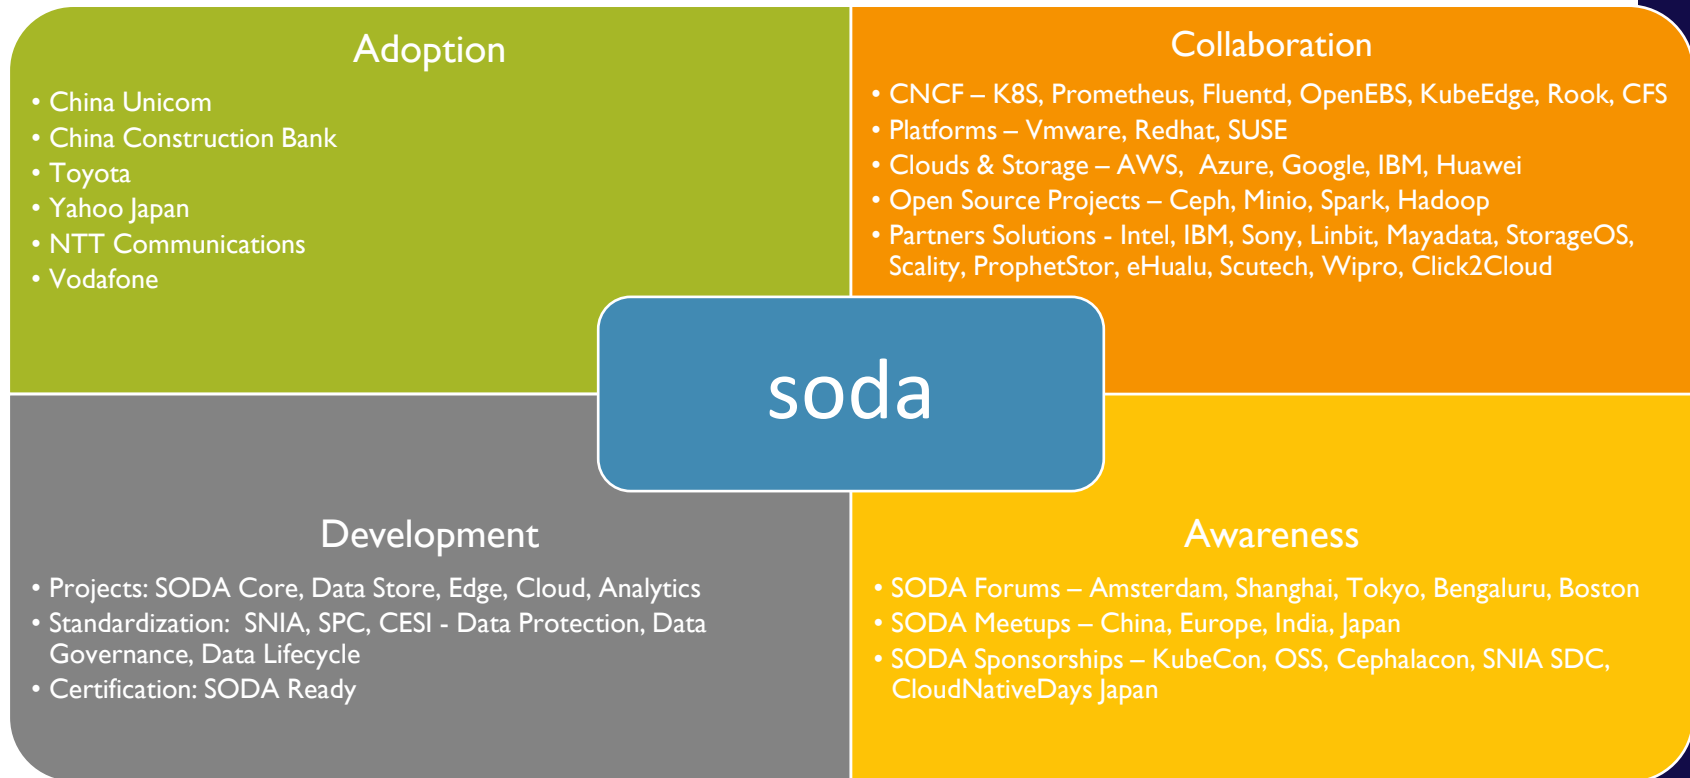- China Construction Bank
- Toyota
- Yahoo Japan
- NTT Communications
- Vodafone

Collaboration

- CNCF – K8S, Prometheus, Fluentd, OpenEBS, KubeEdge, Rook, CFS
- Platforms – Vmware, Redhat, SUSE
- Clouds & Storage – AWS, Azure, Google, IBM, Huawei
- Open Source Projects – Ceph, Minio, Spark, Hadoop
- Partners Solutions - Intel, IBM, Sony, Linbit, Mayadata, StorageOS, Scality, ProphetStor, eHualu, Scutech, Wipro, Click2Cloud

soda

Development

- Projects: SODA Core, Data Store, Edge, Cloud, Analytics
- Standardization: SNIA, SPC, CESI - Data Protection, Data Governance, Data Lifecycle
- Certification: SODA Ready

Awareness

- SODA Forums – Amsterdam, Shanghai, Tokyo, Bengaluru, Boston
- SODA Meetups – China, Europe, India, Japan
- SODA Sponsorships – KubeCon, OSS, Cephalacon, SNIA SDC, CloudNativeDays Japan

SODA Demo

Cloud Vendor agnostic Multicloud Data Management

Demo: (Object)

- Easy addition of Cloud Vendor Backend
- Basic Object Creation
- Seamless Data Mobility Across cloud vendors (migration, simple copy/paste)

SODA Northbound VMWare Plugin

One platform to provide storage to VMWare Datastore from different backends

MISSION

- to foster **an ecosystem of open source data management and storage** software for data autonomy

  - to offer a neutral forum for **cross-projects collaboration and integration**

    - to provide end users **quality end-to-end solutions**

Thank you

SODA CHINA COMMUNITY LAUNCH Dec 20, 2019

**Please take a moment
to rate this session.**

Your feedback matters to us.