



1

BIG DATA EUROPE


Empowering Communities
with Data Technologies

Collaboration Opportunities in Big Data Platform, Data Spaces & Semantic Interoperability

Sören Auer, auer@cs.uni-bonn.de

PICASSO EU-US Collaboration, Minneapolis, June 19, 2017





PICASSO ICT MEETING

Three possible streams for collaboration:

- ⊙ A Big Data Platform for societal good
- ⊙ Establishing data sharing and data value chains with the Industrial Data Space
- ⊙ Semantic Domain Models (vocabularies, ontologies) for establishing a common understanding of the data



Big Data Europe Platform

Empowering Communities with Data Technologies
Platform release



BIG DATA EUROPE


Empowering Communities
with Data Technologies



Smart Data Analytics

June 19 2017

Big Data Europe Platform



Platform Goals

- ⊙ Opensource
- ⊙ Ease of Use
- ⊙ Support a variety of use cases
- ⊙ Embrace emerging Big Data Technologies
- ⊙ Simple integration with custom components



Infrastructure

Hadoop On-Premise
Cloudera, Hortonworks, MAPR, Pivotal, IBM InfoSphere, spice, jethro

Hadoop in the Cloud
Amazon, Microsoft Azure, Google Cloud Platform, IBM InfoSphere, CAEENA, aliscale, Databricks, Qubole, xplinty

Spark
databricks, GridGain, TACHYON, XELUX

Cluster Services
Amazon, Docker, Mesosphere, Core OS, StackIO

Analytics

Analyst Platforms
Palantir, AYASDI, Quid, engma, Dope Analytics, KNOX LABS

Analytics Platforms
Microsoft, guavus, Duckaroo, Interana

Data Science Platforms
Continuum, DataRobot, Alping, MOOSE, DataSift, NOVENA, sense, what, ACORTEX

Visualization
Tableau, Google Cloud, Roambi, Looker, Qlik, CHARTIO

Applications

Sales & Marketing
RADIUS, Gainsight, bloomreach, Zeta, livefyre, blueyonder, Okta, Lattice, persado, infer, sense, AVISO, ACTIONIQ, GRANITING, ENCAGIO

Customer Service
MEDALLIA, ATTENTIFY, STERILIA, NGDATA, WISIO, DataGrip, IMPACTU, Rube PROTECTS

Human Capital
gild, Connectix, textio, entelo, hiQ, Everlaw, Brevia, FISHBOWL

Legal
RAVEL, SPICEDATA, Everlaw, Brevia, FISHBOWL

NoSQL Databases
Amazon DynamoDB, Microsoft Azure, mongoDB, KROSPIKE, Sequoia, redislabs, InfluxDB

NewSQL Databases
SAP, Clustrix, Pivotal, Oracle, memsql, paradigM4, nuodb, MariaDB, VOLTDB, citusdata, doopdb, Trafalgar, Cockroach LABS

BI Platforms
Power BI, Amazon, Domo, Micro Analytics, GoodData, birst, plotforce, lockero, tableausoft

Statistical Computing
sas, SPSS, MATLAB

Log Analytics
splunk, sumologic, HIBONO, loggly

Social Analytics
NETBASE, DATASIFT, track, bitly, synthosio, boozice, ampereach

Ad Optimization
MediaMath, Integral, OpenX, theTradeDesk, Livestorm, Uppler, DAPP

Security
CYLANCE, CounterTrack, ARLA1, Recorded Future, FORTISCALE, sift science, Xplore, feedzai, escnfy

Vertical AI Applications
Facebook, Clara, KASISTA, lumata

Graph Databases
neo4j, Orient DB, EDB

MPP Databases
Teradata, Vertica, Alteryx, Kognitio, Dremio

Cloud EDW
Amazon Redshift, Microsoft Azure, Pivotal, Snowflake, AWS Glue, Infoworks

Data Transformation
Alteryx, Trifacta, tamar, Alteryx, Pivotal, Snowflake, Dell Alation

Data Integration
MuleSoft, snapLogic, Bedrock Data

Real-Time
Amazon, Metamarkets, ScanFluent, Convoconver, DataArtisans

Machine Learning
Amazon, H2O, SKYTRIE, DataRobot, DataCamp, DataKind, checkr, IBM, Xplore

Speech & NLP
Narrative Science, Capital, Nuance, Samantrix, flintfix, IBM, Xplore

Horizontal AI
IBM Watson, Cortana, Viv, Namana, Clarifai

Publisher Tools
Outbrain, mixpanel, Chartbeat, yieldbot, Yieldmo

Govt/ Regulation
Socrata, OPENGOV, EN, Enigma, mark43, OpenDataSoft

Finance
Affirm, LendingClub, OnDeck, Kreditech, Kabbage, INSIKT, UORO, Dataminr, Lenddo, KENSHO, ALOYA, ISENTIUM, Quantopian

Management / Monitoring
New Relic, 4F-DYNAMICS, Amazon, Oclifio, Splunk, PCCOHO, Aveset

Security
Tanium, Illumio, DataGravity, CyberCloud, VECTRA, Sift, Trustlook

Storage
Amazon, Microsoft Azure, Pivotal, Qumulo, CONCURRENT

App Dev
Apigee, ORSK, Typesafe, Convoconver

Crowd-sourcing
WorkFusion

Search
HP, Oracle, Lucidworks, elastic, Thinkful, MAANA, swifttype, Algolia, Sift, DataKind

Data Services
LKO, OPERA, DataCamp, DataKind

For Business Analysts
Original Logic, DearStory, CIRRO, Import.io

SMB / Commerce
Google Analytics, Bluecore, sumai, granify, retention, custora

Education/ Learning
Knewton, Clever, Cleara, KnowU

Life Sciences
Counsyl, Recombi, Flatiron, Zyrus, HealthTap, Zephyr, Metabio, Zephyr, OVI, Gingerio, Immisrbic, Glow, entho, AICure, Ambrico

Industries
eHamony, RetailNext, Stitch Fix, WorkFusion, Tachyus, Seeq, FarmLogs, PowerD, collect, Stat, Boxever

Cross-Infrastructure/Analytics

Amazon, Google, Microsoft, IBM, SAP, SAS, HP, VMware, Talend, TIBCO, Teradata, Oracle, NetApp

Framework
Hadoop, YARN, Spark, HESOD, Flink, WCDAP

Query / Data Flow
SLAMDATA, DRILL, Hive, Pig, Tez, Mahout, Hama, Hbase, Hdfs, Hcatalog, HiveServer2, HiveMetastore, HiveTransport, HiveServer2, HiveMetastore, HiveTransport

Data Access
Cassandra, Hbase, MongoDB, CouchDB, Riak, Aerospike, ScyllaDB, Aerospike, ScyllaDB

Coordination
Talend, Apache Zookeeper, Apache Ambari

Real-Time
Storm, Spark, Flink, Tachyon, Druid

Stat Tools
Scala, SciPy

Machine Learning
mlib, Apache, SENG, MAE lib, FeatureFu, DIMSUM, Aerolve, Caffe, CNTK, DL4J, VELES, WESA, DIMSUM

Search
Elasticsearch, Solr, Recore

Security
Apache Ranger, Apache Ambari

Visualization
Tableau, Qlik, PowerBI, Looker, Domo, Alteryx, SAP, Oracle, IBM, Microsoft, Amazon, Google, Facebook, Twitter, LinkedIn, YouTube, Instagram, Snapchat, WhatsApp, Telegram, Messenger, Slack, Discord, Zoom, Microsoft Teams, Google Meet, Webex, Cisco Webex, Cisco Jabber, Cisco Duo, Cisco Duo Mobile, Cisco Duo for iOS, Cisco Duo for Android, Cisco Duo for Windows, Cisco Duo for Mac OS, Cisco Duo for Linux, Cisco Duo for Unix, Cisco Duo for Solaris, Cisco Duo for AIX, Cisco Duo for HP-UX, Cisco Duo for IRIX, Cisco Duo for Tru64, Cisco Duo for OpenVMS, Cisco Duo for OS/390, Cisco Duo for z/OS, Cisco Duo for z/OS VSE, Cisco Duo for z/OS VSE/2, Cisco Duo for z/OS VSE/3, Cisco Duo for z/OS VSE/4, Cisco Duo for z/OS VSE/5, Cisco Duo for z/OS VSE/6, Cisco Duo for z/OS VSE/7, Cisco Duo for z/OS VSE/8, Cisco Duo for z/OS VSE/9, Cisco Duo for z/OS VSE/10, Cisco Duo for z/OS VSE/11, Cisco Duo for z/OS VSE/12, Cisco Duo for z/OS VSE/13, Cisco Duo for z/OS VSE/14, Cisco Duo for z/OS VSE/15, Cisco Duo for z/OS VSE/16, Cisco Duo for z/OS VSE/17, Cisco Duo for z/OS VSE/18, Cisco Duo for z/OS VSE/19, Cisco Duo for z/OS VSE/20

Data Sources & APIs

Health
Apple, Jawbone, Garmin, Withings, Fitbit, Validic, Human API, Knsa

IOT
Uptake, ThingWorx, InMobi, Samsara, Xata


Financial & Economic Data
Bloomberg, Dow Jones, YOLEE, PREMIER, CAPITAL IQ, Quandl, Xignite, CBHSC45, Macquarie, Sesizim, PLAD

Air / Space / Sea
Planet Labs, Spire, Airware, DroneCatcher

Location/People/Entities
Garmin, Foursquare, InsideView, Esri, Streetline, Caribb, Factual, PlaceIQ, Cision, Placemeter, Basis, Mapbox

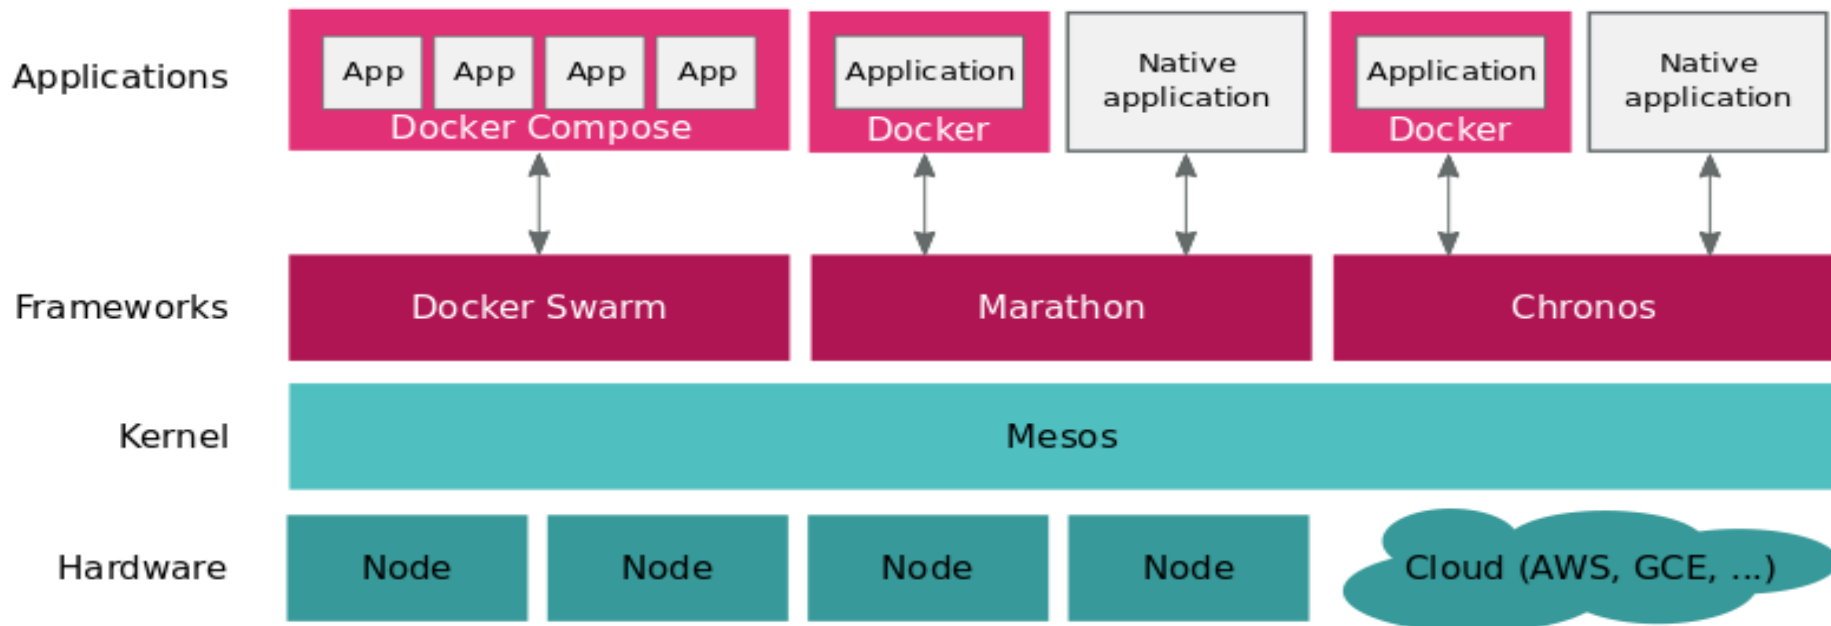
Other
Qualtrics, Panjiva, Data.gov


Incubators & Schools
GA, DataCamp, Insight, DataFlite, The Data Incubator



Platform Architecture Evolution

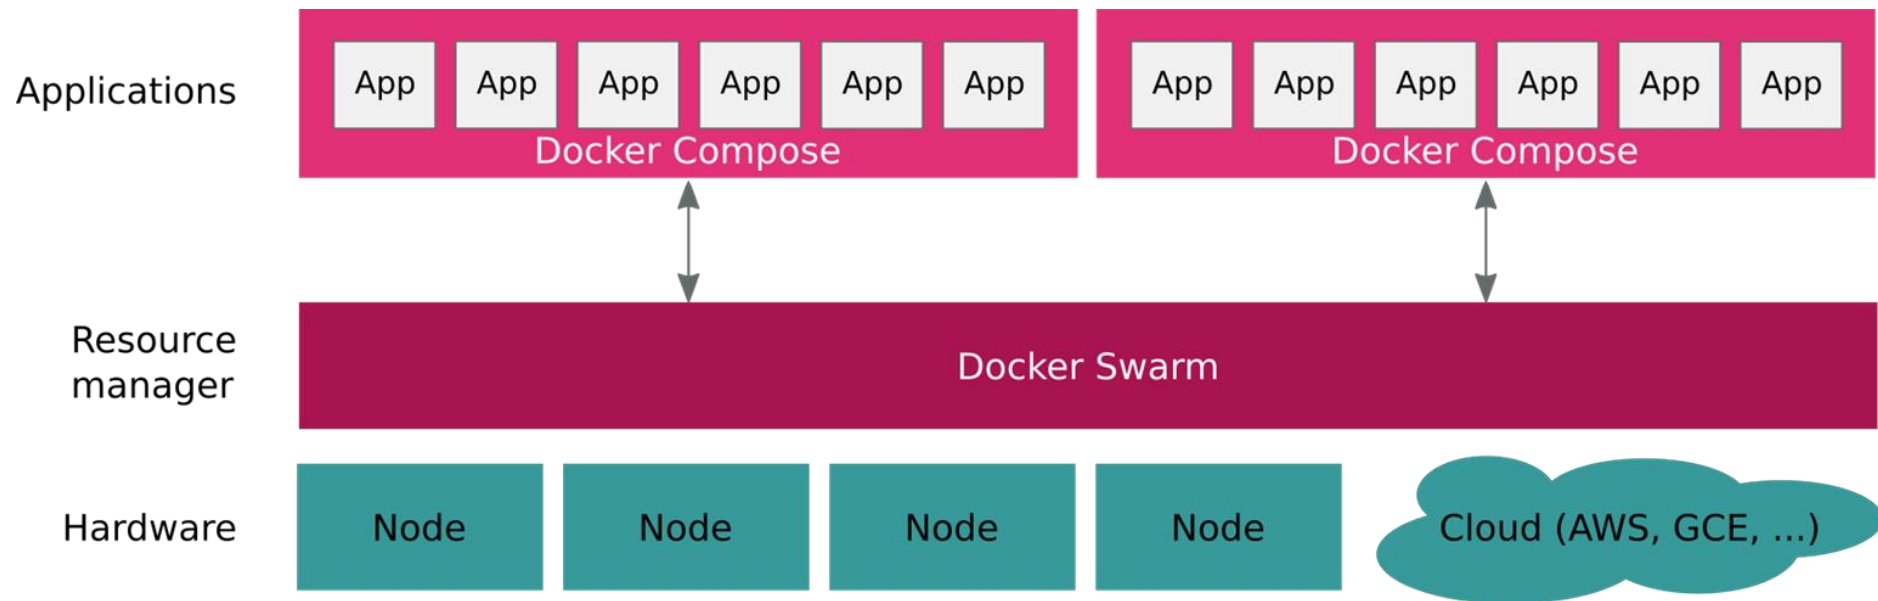
6






Platform Architecture Evolution

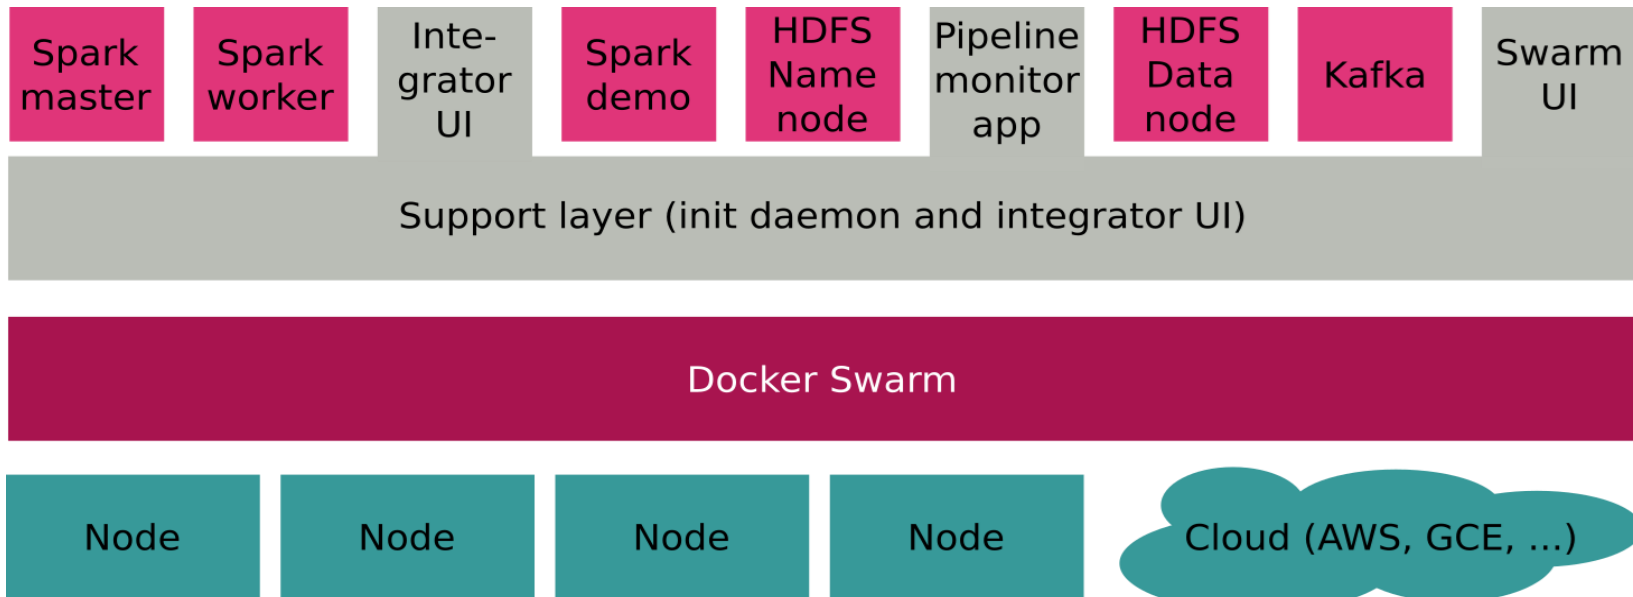
7






Platform Architecture Existing

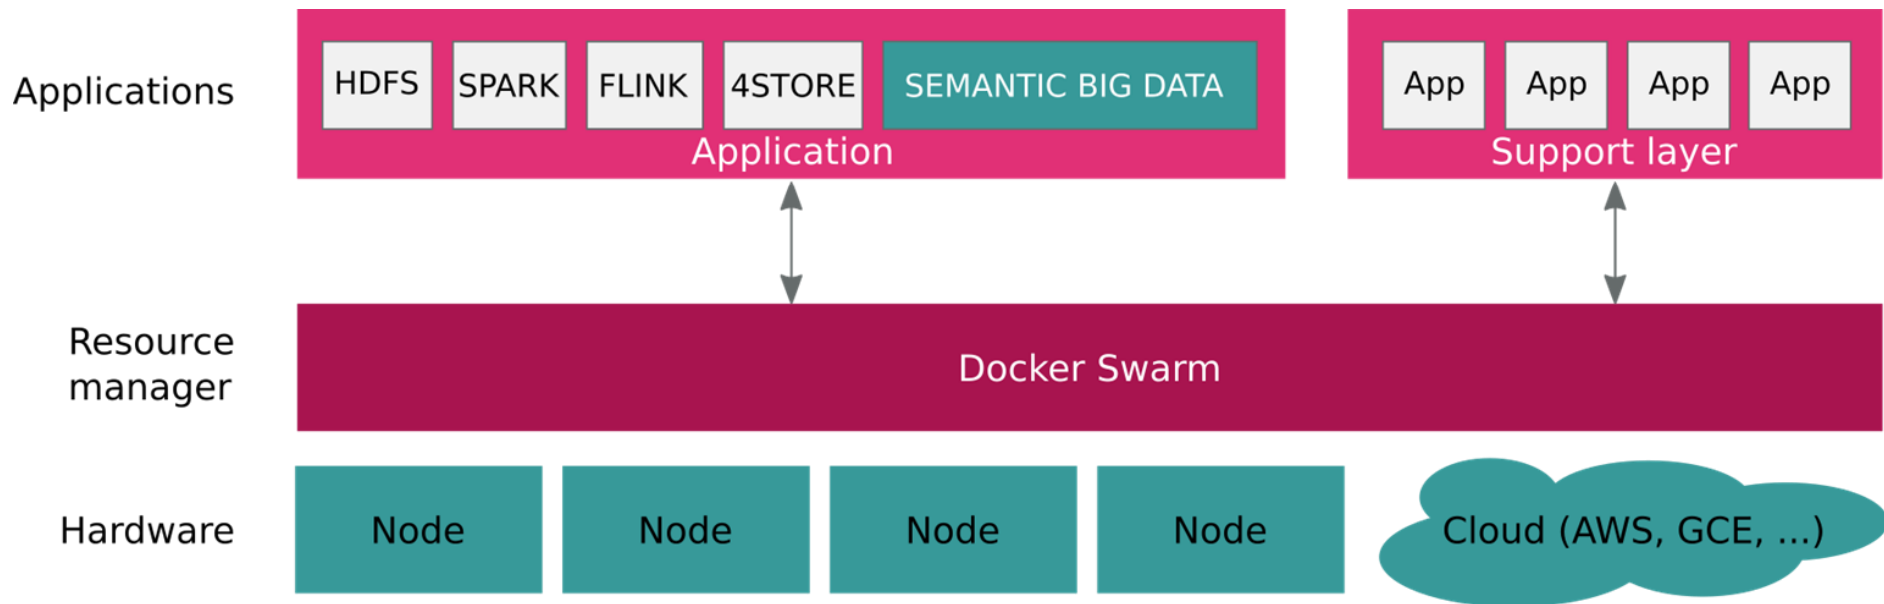
8






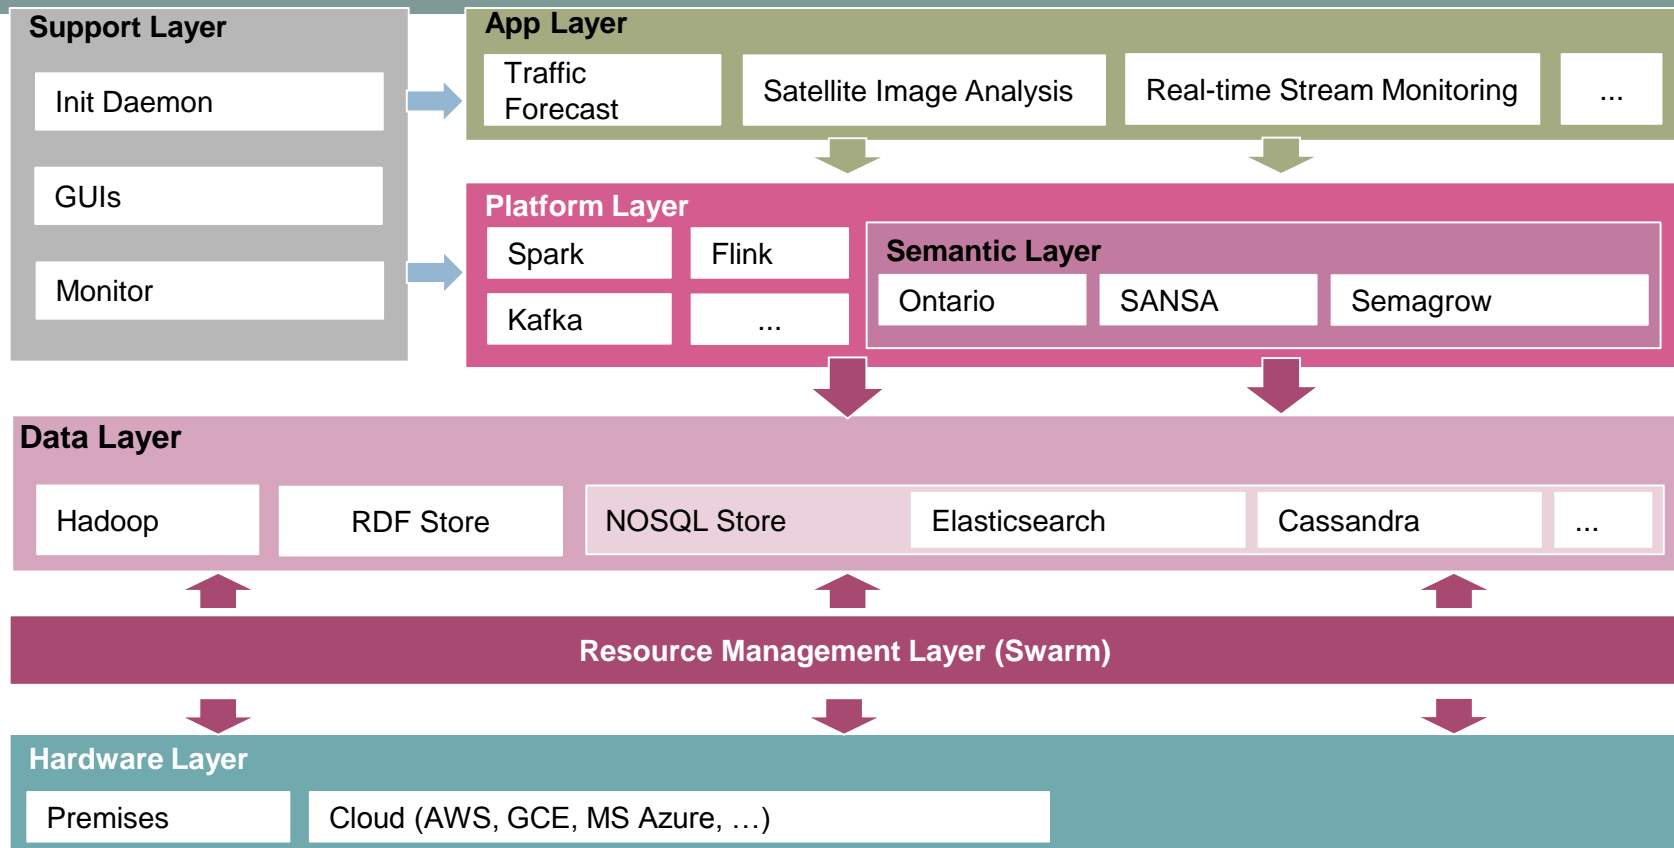
Platform Architecture Existing


9





Platform Architecture Alternate View






BDE Supported Frameworks

11


Search/indexing	Data processing
Apache Solr	Apache Spark
Data acquisition	Apache Flink
Apache Flume	Semantic Components
Message passing	Strabon
Apache Kafka	Sextant
Data storage	GeoTriples
Hue	Silk
Apache Cassandra	SEMAGROW
ScyllaDB	LIMES
Apache Hive	4Store
Postgis	OpenLink Virtuoso




Platform features

12

- ⊙ BDE Development Environment
 - Stack builder
 - Workflow builder
 - Instructions to add custom components to the BDE stack
- ⊙ Administrator Interface
 - SwarmUI
- ⊙ UI Integrator
 - Workflow monitor
 - Integrated web interface






What BDE Provides ?

13

- ⊙ Platform Installation Instructions
- ⊙ Usage Instructions
 - Creating a stack
 - Creating a workflow
 - Monitoring the Stack
 - Integration of Custom Components



Platform installation

14



Manual installation guide



Using Docker Machine


- On local machine (VirtualBox)
- In cloud (AWS, DigitalOcean, Azure)
- Bare metal



Screencasts


<https://www.big-data-europe.eu/platform/>


<https://github.com/big-data-europe>



Deploying a Big Data Stack



15

- ◎ Stack Builder 
- ◎ Stack
 - Collection of communicating components to solve a specific problem
- ◎ Described in Docker Compose
 - Component configuration
 - Application topology



Creation of WorkFlows

16

- ◎ Pipeline Builder 
 - Allows creation of dependencies among different applications
- ◎ WorkFlow Monitor 
 - Monitoring of pipeline-workflow using



Integrating Custom Components

17

⊙ Instructions

- Orchestrator required for initialization process (init_daemon)
 - ❖ Components may depend on each other
 - ❖ Components may require manual intervention
- User Interface Integration
 - ❖ Standard Interfaces from components
 - ❖ Combine and align the interfaces



User Interfaces

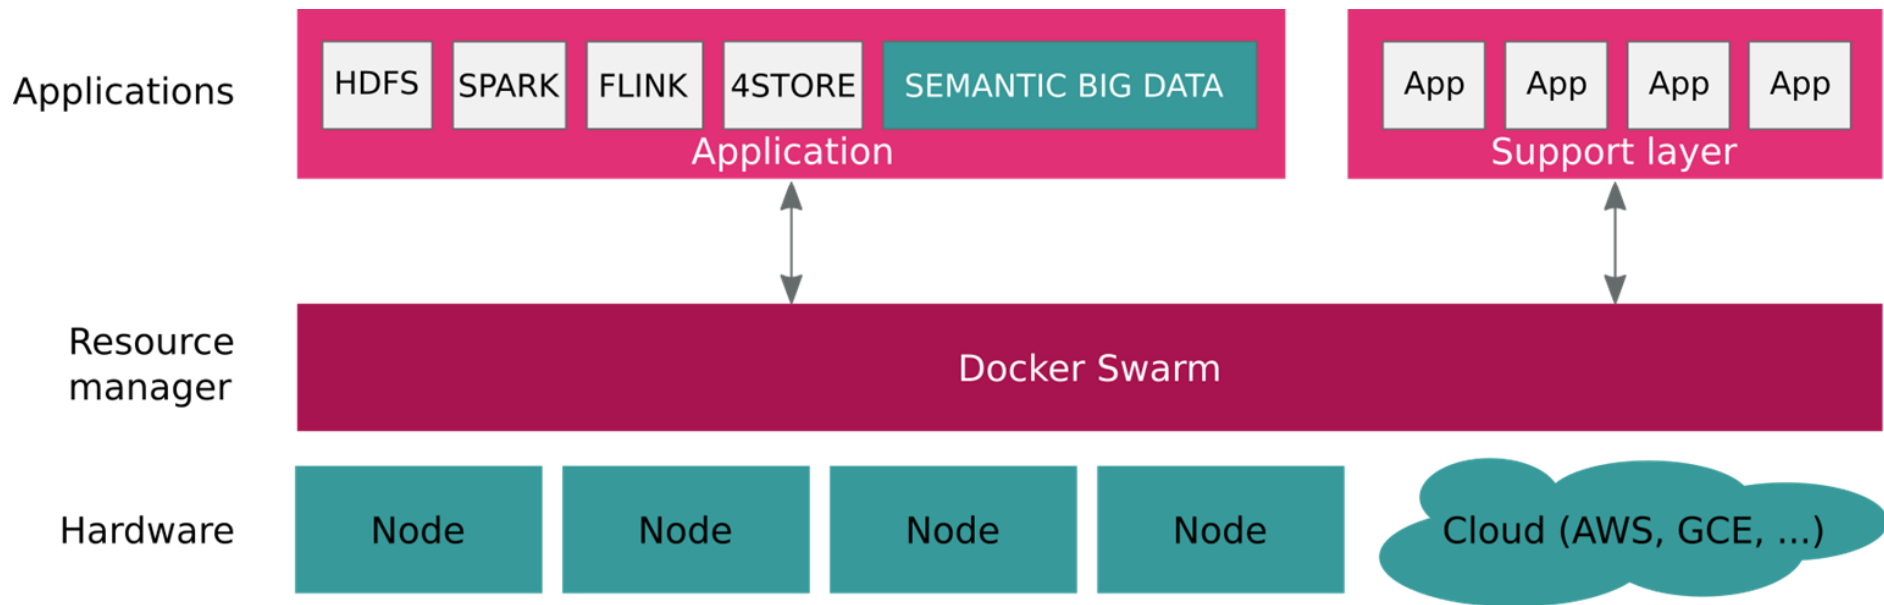
18


- ⊙ Target: Facilitate the use of the platform
 - User Interface Adaption
- ⊙ Available interfaces
 - Workflow UIs
 - ◆ Workflow Builder
 - ◆ Workflow Monitor
 - Swarm UI
 - Integrator UI



Platform Architecture

19





Pilot Show Cases

20



Health

SC1



Food & Agriculture

SC2



Energy



Transport

SC3



Climate

SC4



Social Sciences

SC5



Security

SC6

SC7

SC1 - Open PHACTS discovery platform relating to biological/medical questions

SC2 - Discovery and Linking of Viticulture-relevant information


SC3 - System monitoring in energy production units

SC4 - Short-Term traffic flow forecasting.

SC5 - Supporting data-intensive climate research

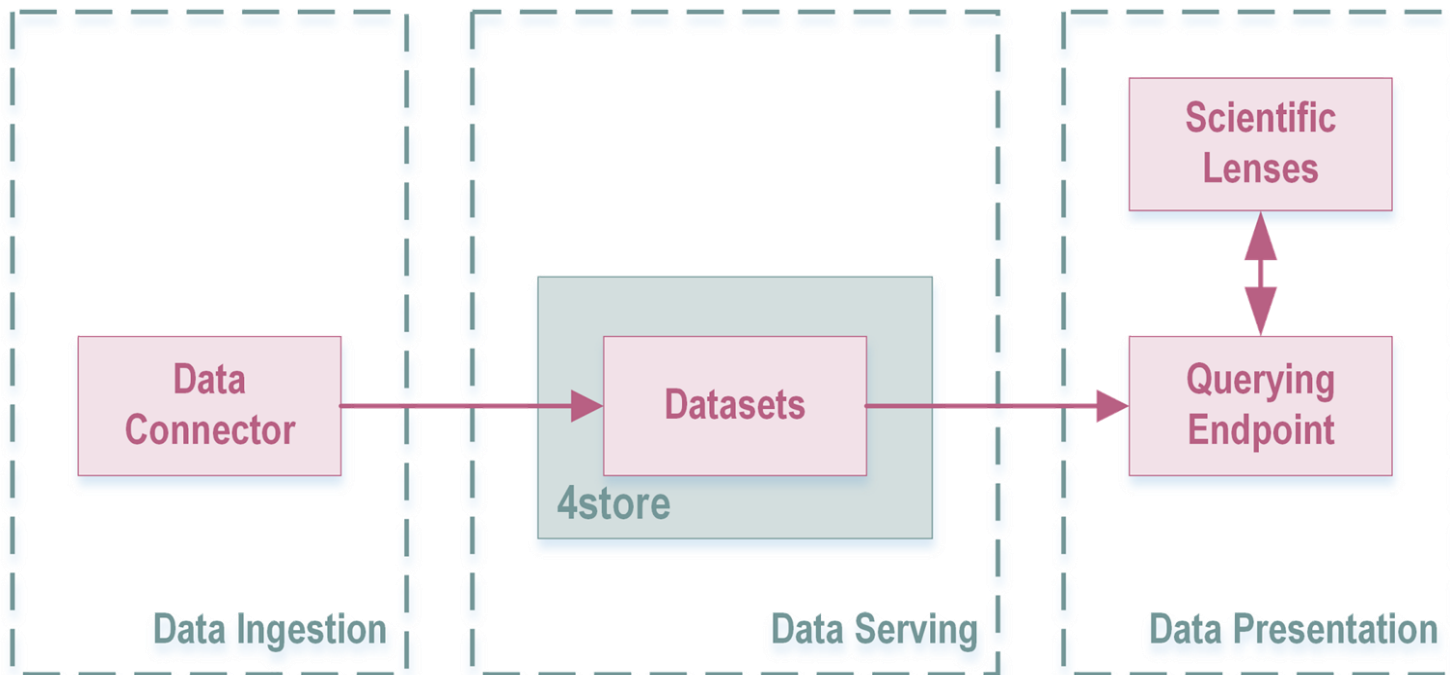
SC6 - Citizens & Researchers Budget on Municipal Level


SC7 - Ingestion of remote sensing images and social sensing data to detect and verify changes on the Earth surface for security applications



SC1-Health

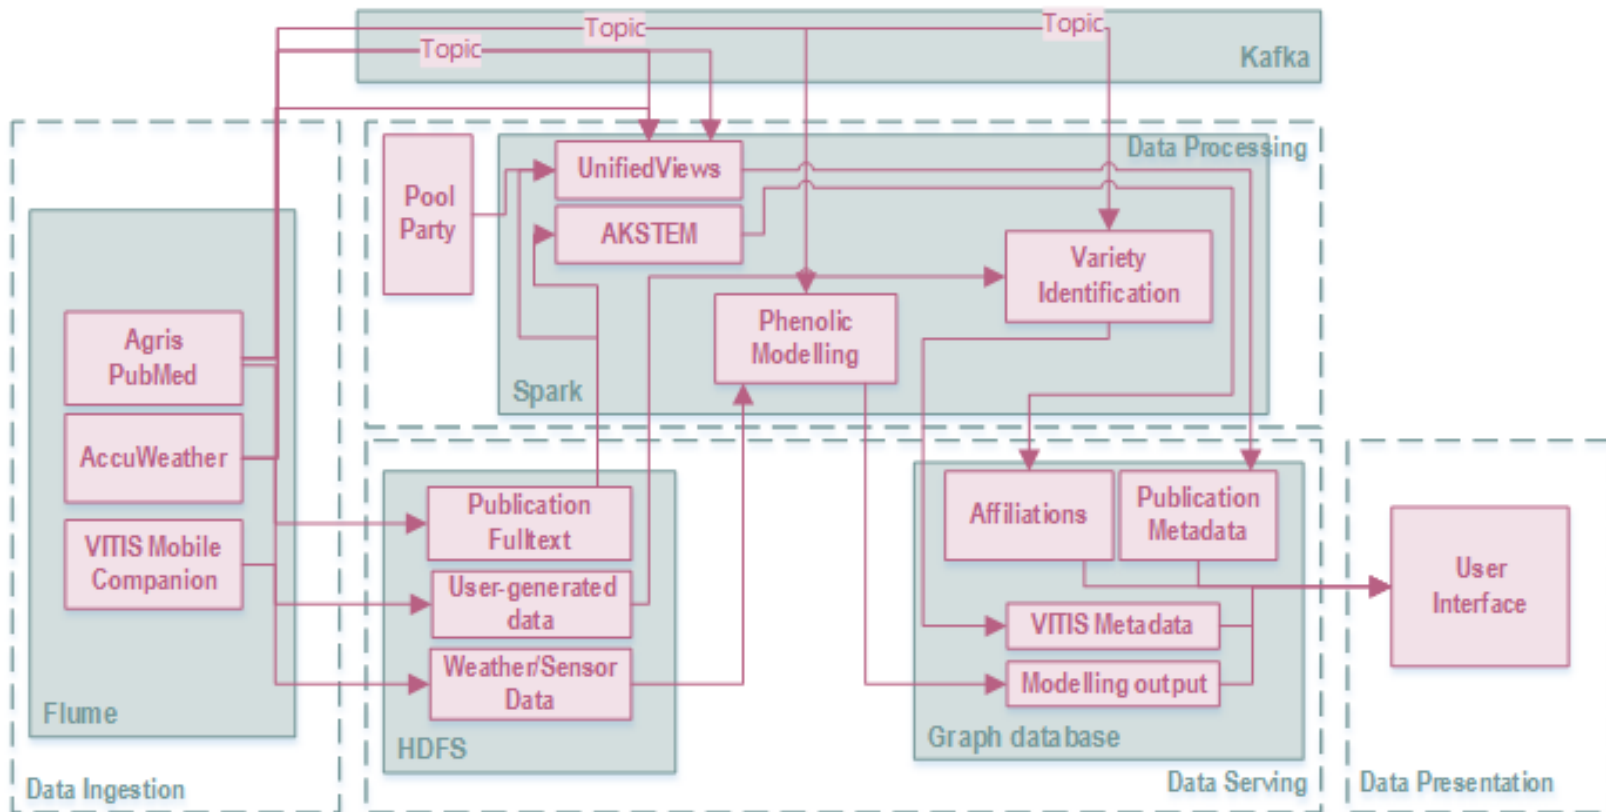
21






SC2 - Food

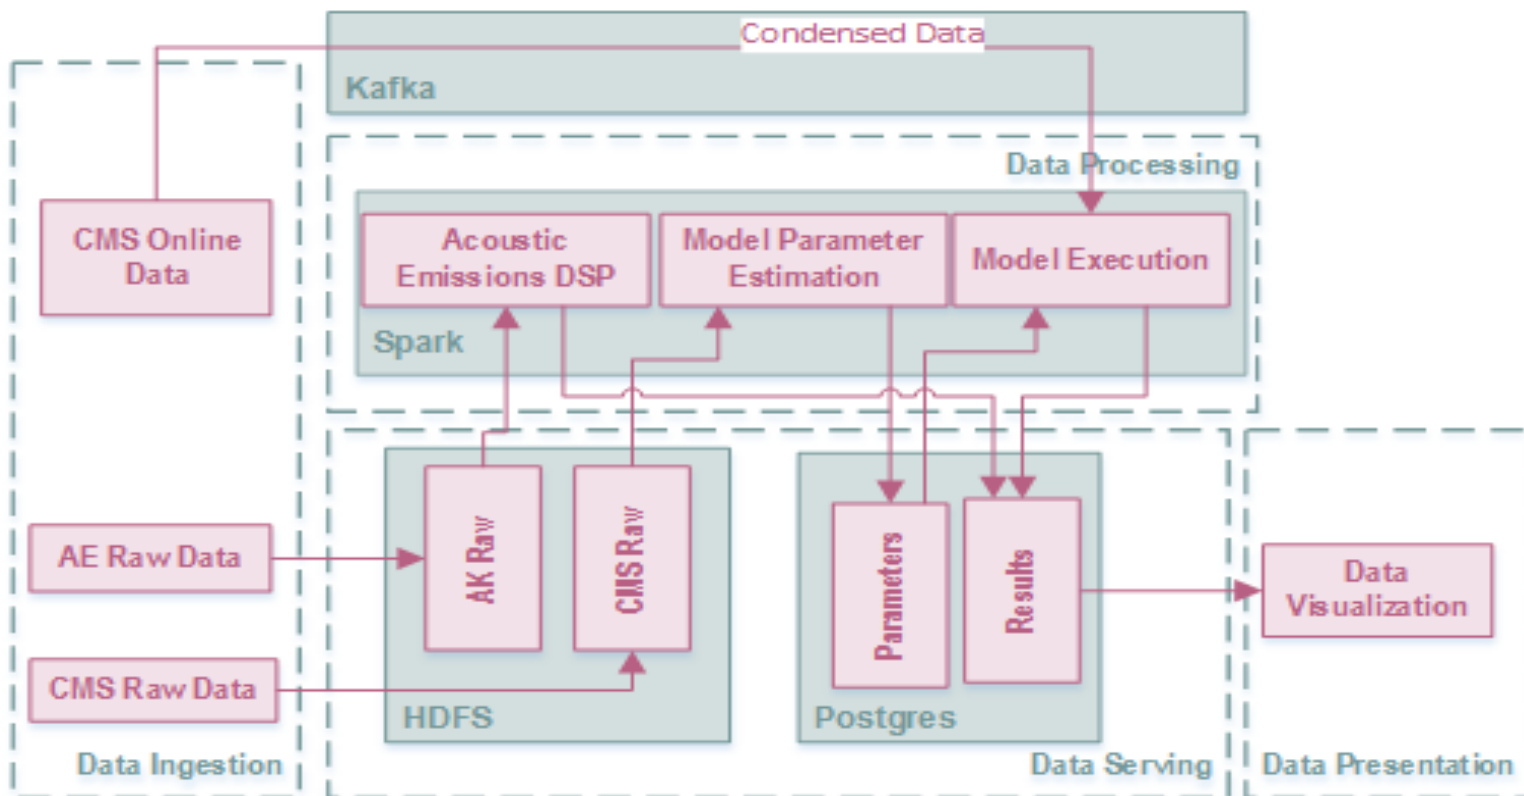
22






SC3 - Energy

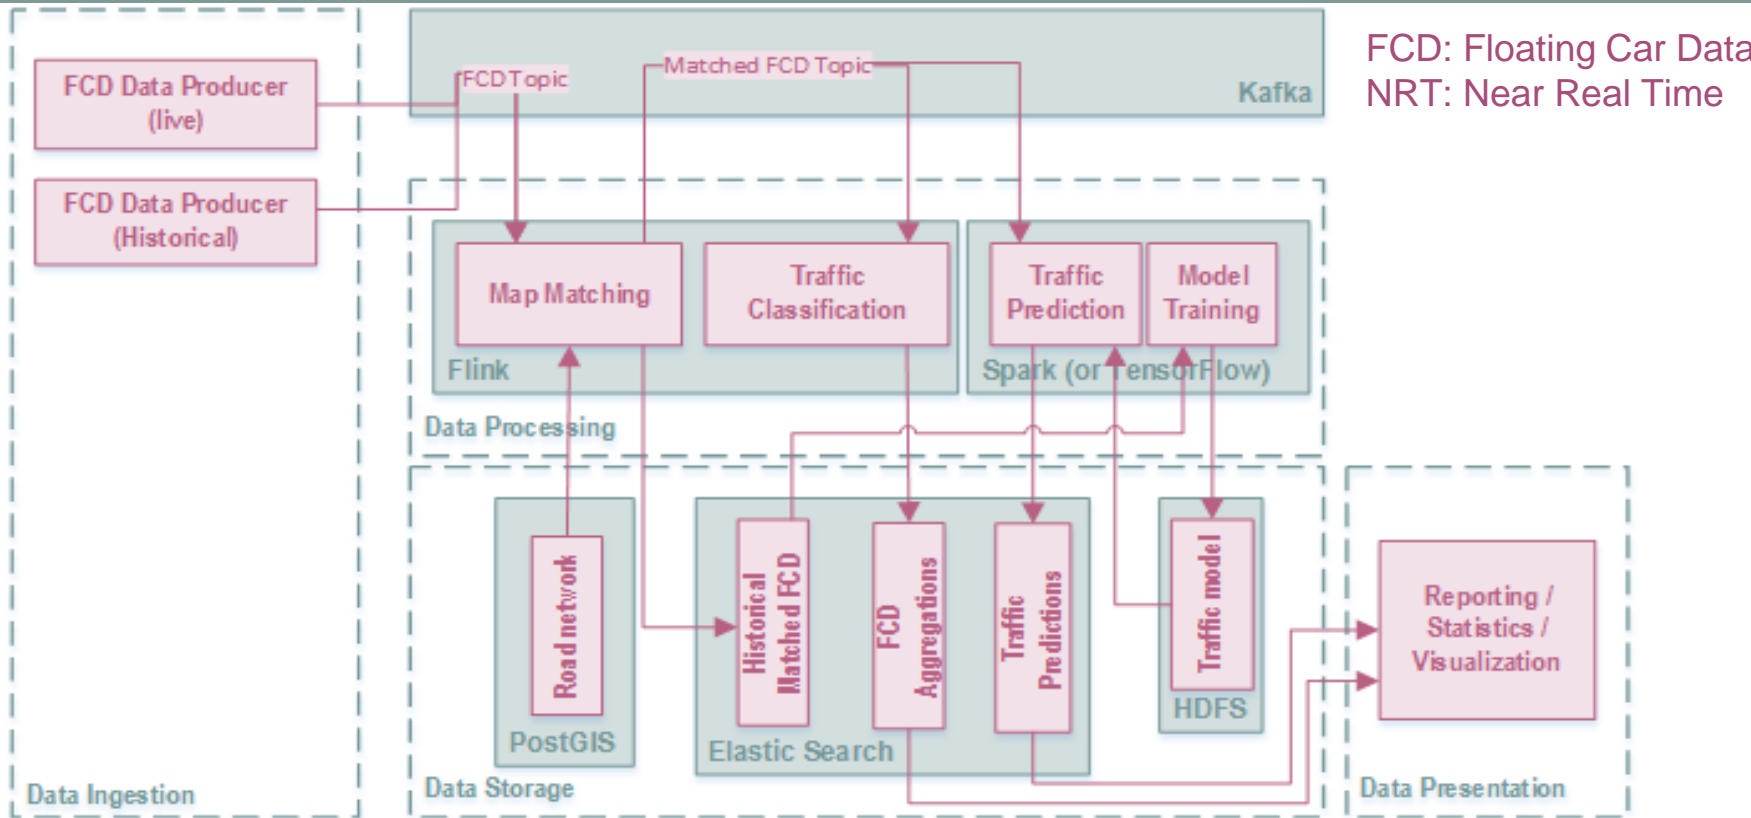
23






SC4 - Transport

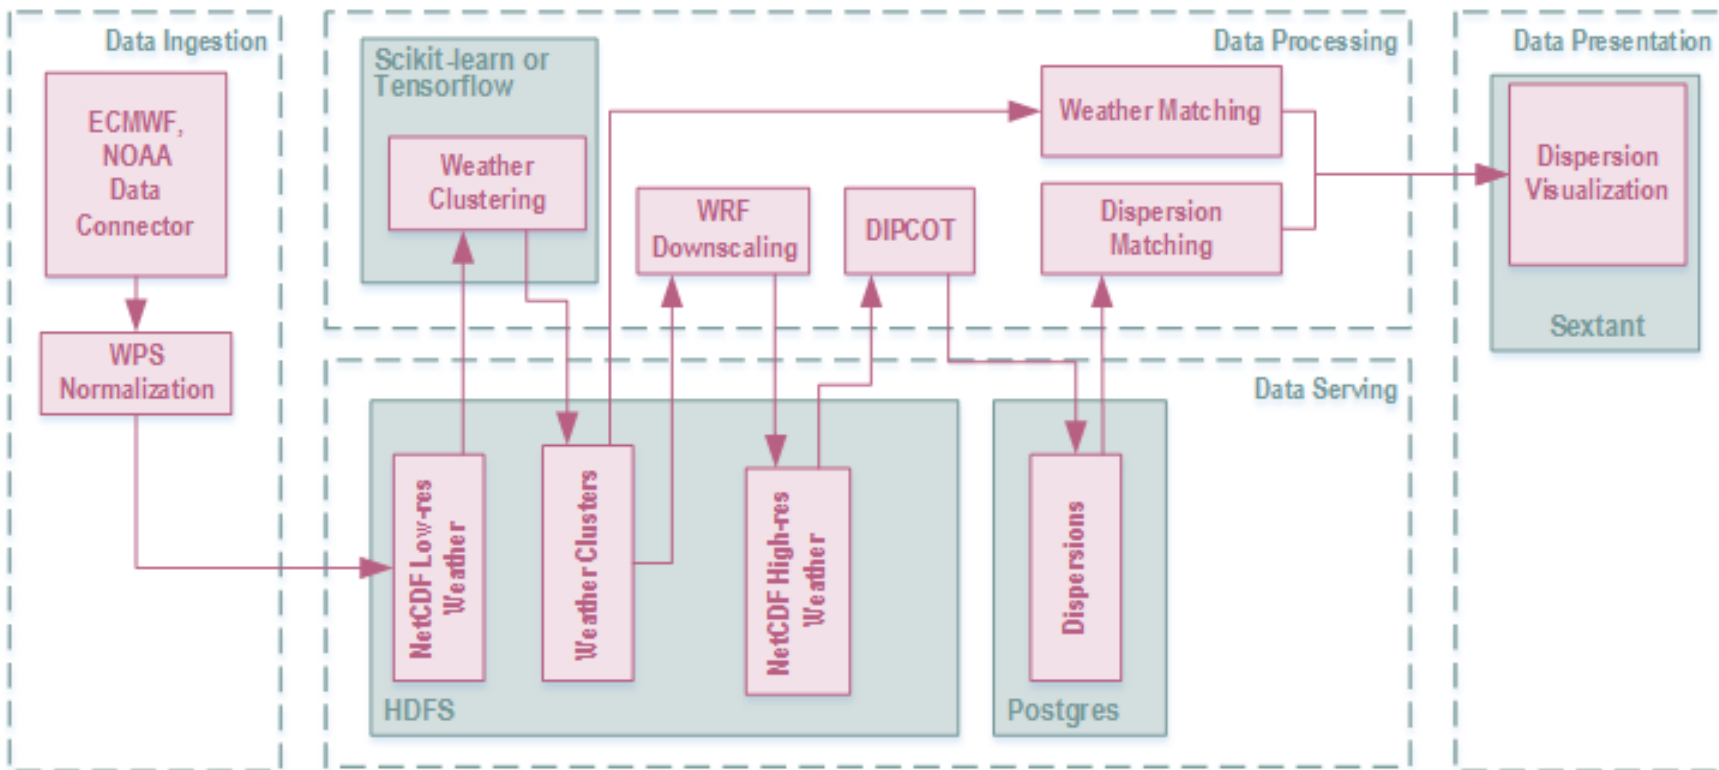
24






SC5 - Climate

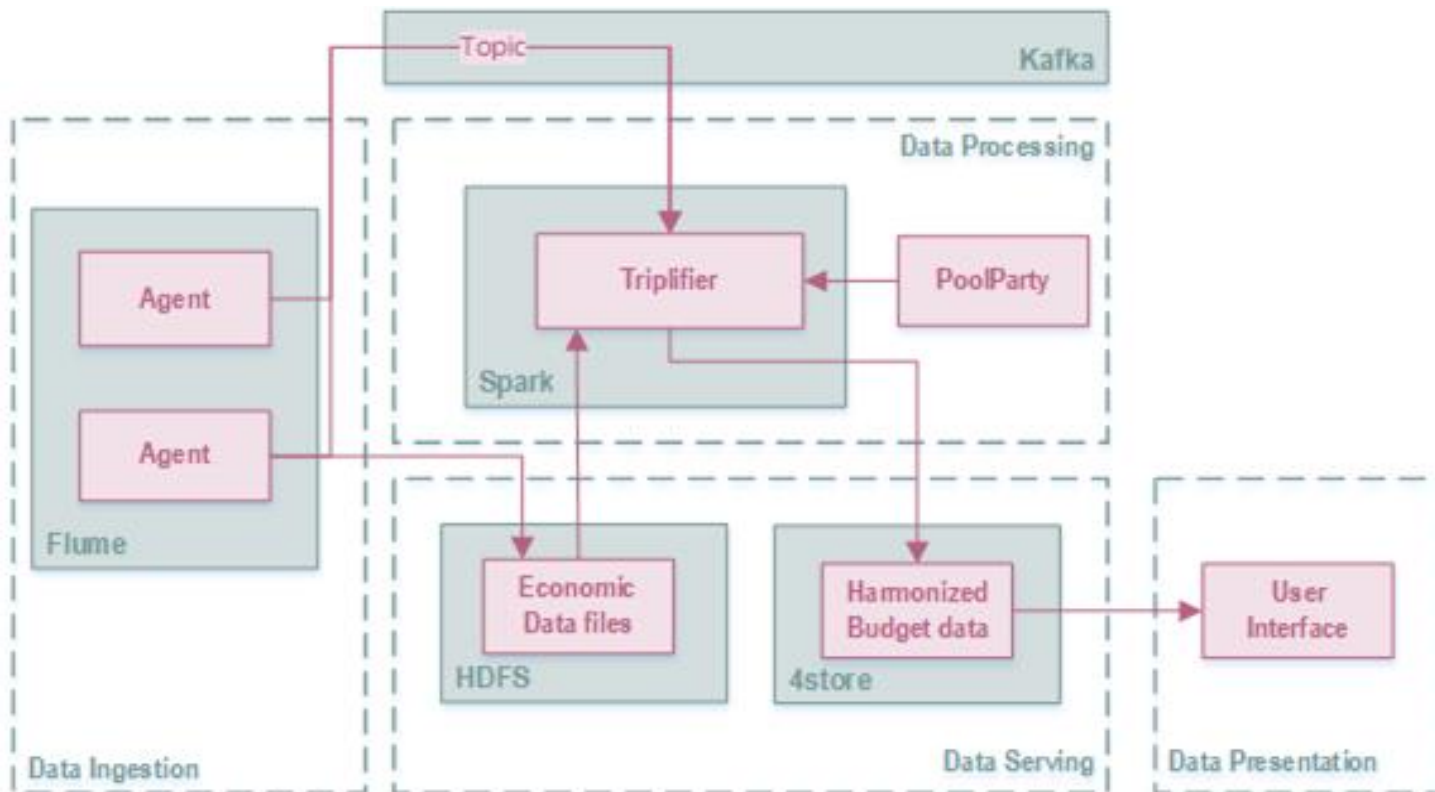
25






SC6 - Social Sciences

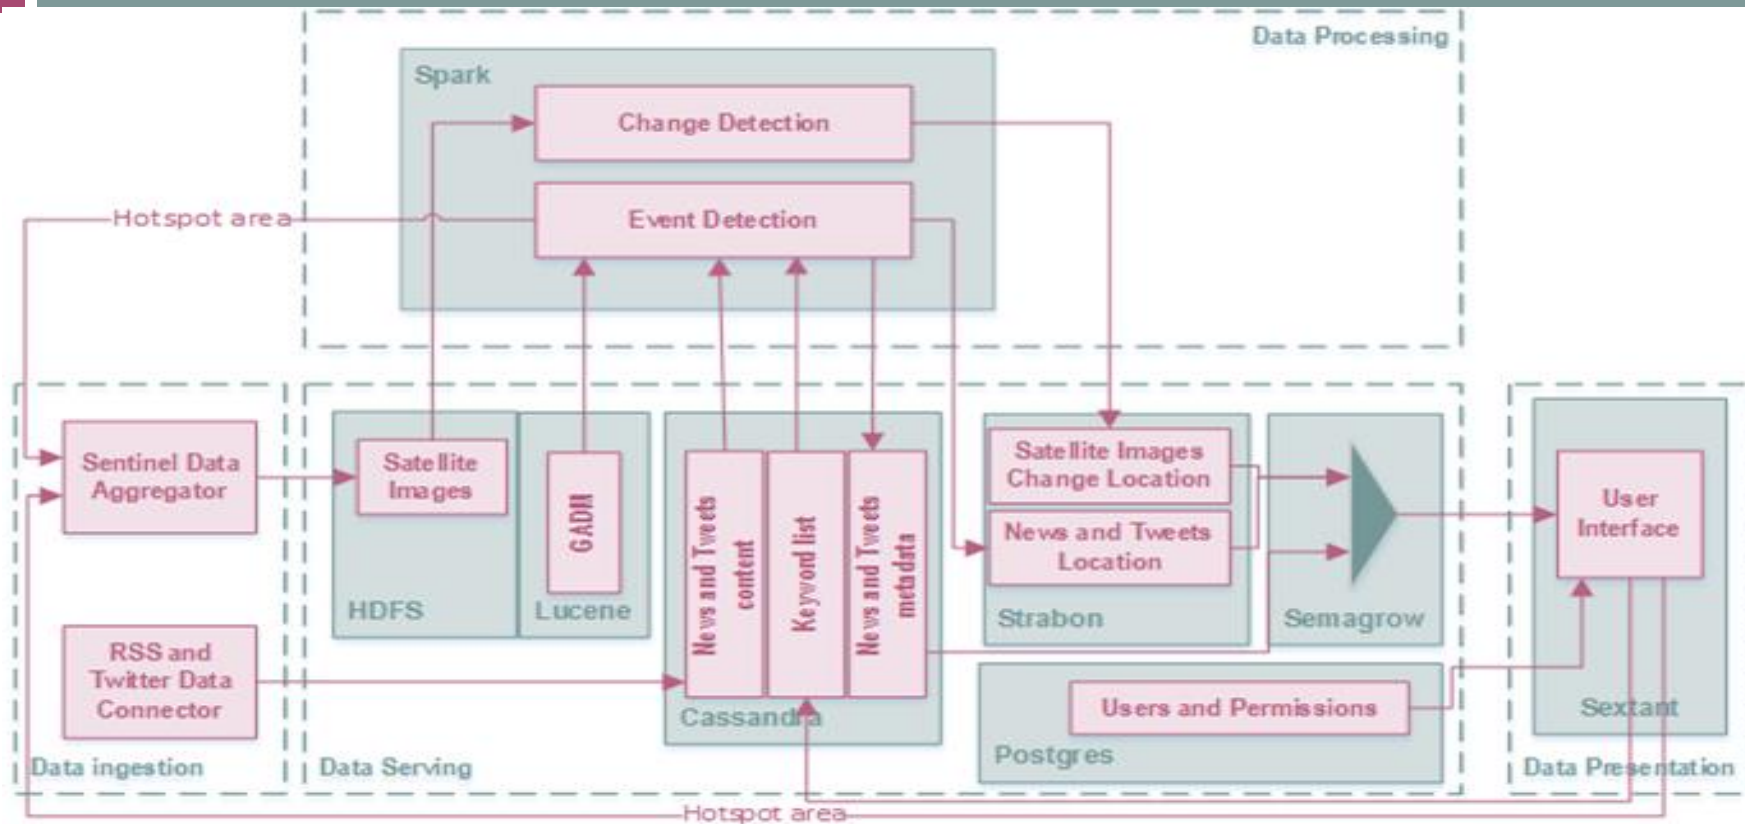
26






SC7 - Security

27






BDE vs Hadoop distributions

28


	Hortonworks	Cloudera	MapR	Bigtop	BDE
<i>File System</i>	HDFS	HDFS	NFS	HDFS	HDFS
<i>Installation</i>	Native	Native	Native	Native	lightweight virtualization
<i>Flexible Modular Architecture</i>	no	no	no	no	yes
<i>High Availability</i>	Single failure recovery (yarn)	Single failure recovery (yarn)	Self healing, mult. failure rec.	Single failure recovery (yarn)	Failure recovery
<i>Cost</i>	Commercial	Commercial	Commercial	Free	Free
<i>Scaling</i>	Freemium	Freemium	Freemium	Free	Free
<i>Addition of custom components</i>	Not easy	No	No	No	Yes
<i>Integration testing</i>	yes	yes	yes	yes	--
<i>Operating systems</i>	Linux	Linux	Linux	Linux	Windows/Mac/Linux
<i>Management tool</i>	Ambari	Cloudera manager	MapR Control system	-	Docker swarm UI+ Custom



BDE vs Hadoop distributions

29

- ⊙ BDE is not built on top of existing distributions
- ⊙ Targets
 - Communities
 - Research Institutions
- ⊙ Bridges scientists and open data
- ⊙ Multi Tier research efforts towards Smart Data



Wrap Up, thanks for your Attention

Three possible streams for collaboration:

- ⊙ A Big Data Platform for societal good
- ⊙ Establishing data sharing and data value chains with the Industrial Data Space
- ⊙ Semantic Domain Models (vocabularies, ontologies) for establishing a common understanding of the data

Please get in touch: Sören Auer (coordinator Big Data Europe), auer@cs.uni-bonn.de