Karlsruhe Institute of Technology

The Smart Data Innovation Lab Research Platform

2nd Big Data All Hands Meeting, 11. October 2017, KIT

Parinaz Ameri, Frank Polgart, Nico Schlitter, Jan Erik Sundermann

STEINBUCH CENTRE FOR COMPUTING - SCC

Outline

- Introduction Smart Data Innovation Lab
- Overview SDIL Platform
- Concepts for Operation of Platform

Smart Data Innovation Lab

- SDI-X funded by BMBF 09/2015 – 08/2018
  - Joint platform to strengthen cooperation of industry and research
  - Self organizing data innovation communities
  - Cooperation in a framework of well defined projects
  - SDIL strategy board approves project proposals
  - Project members get access to SDIL platform hosted by KIT @ SCC

- SDSC-BW funded by BW 10/2014 – 09/1017
  - Educating SME about data analytics
  - Evaluating potential of existing data

SDIL Concept – Benefits for Partners

- Research partners:
  - Proof analytics concepts against real use cases and data
  - Easy access to powerful state-of-the-art technologies
- Industry partners:
  - Lower threshold to experiment with analytics and big-data methods
  - Easy access to cutting-edge research and technologies
  - Leverage data analytics for tangible business advantage
- IT providers:
  - Showcase latest software and hardware technologies
  - Test and improve products for real use cases and workloads

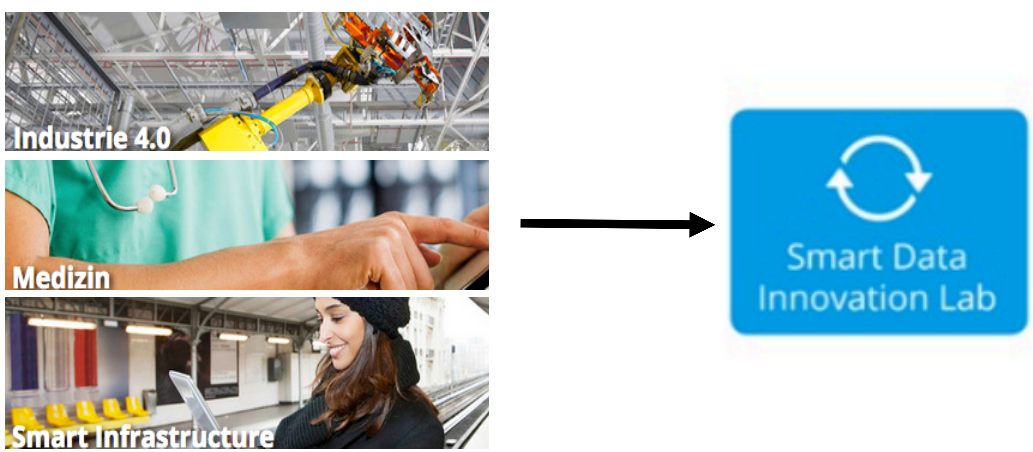

SDIL Partners

Core Partners

Research Partners

Project Partners

Associations

SDI-X Structure

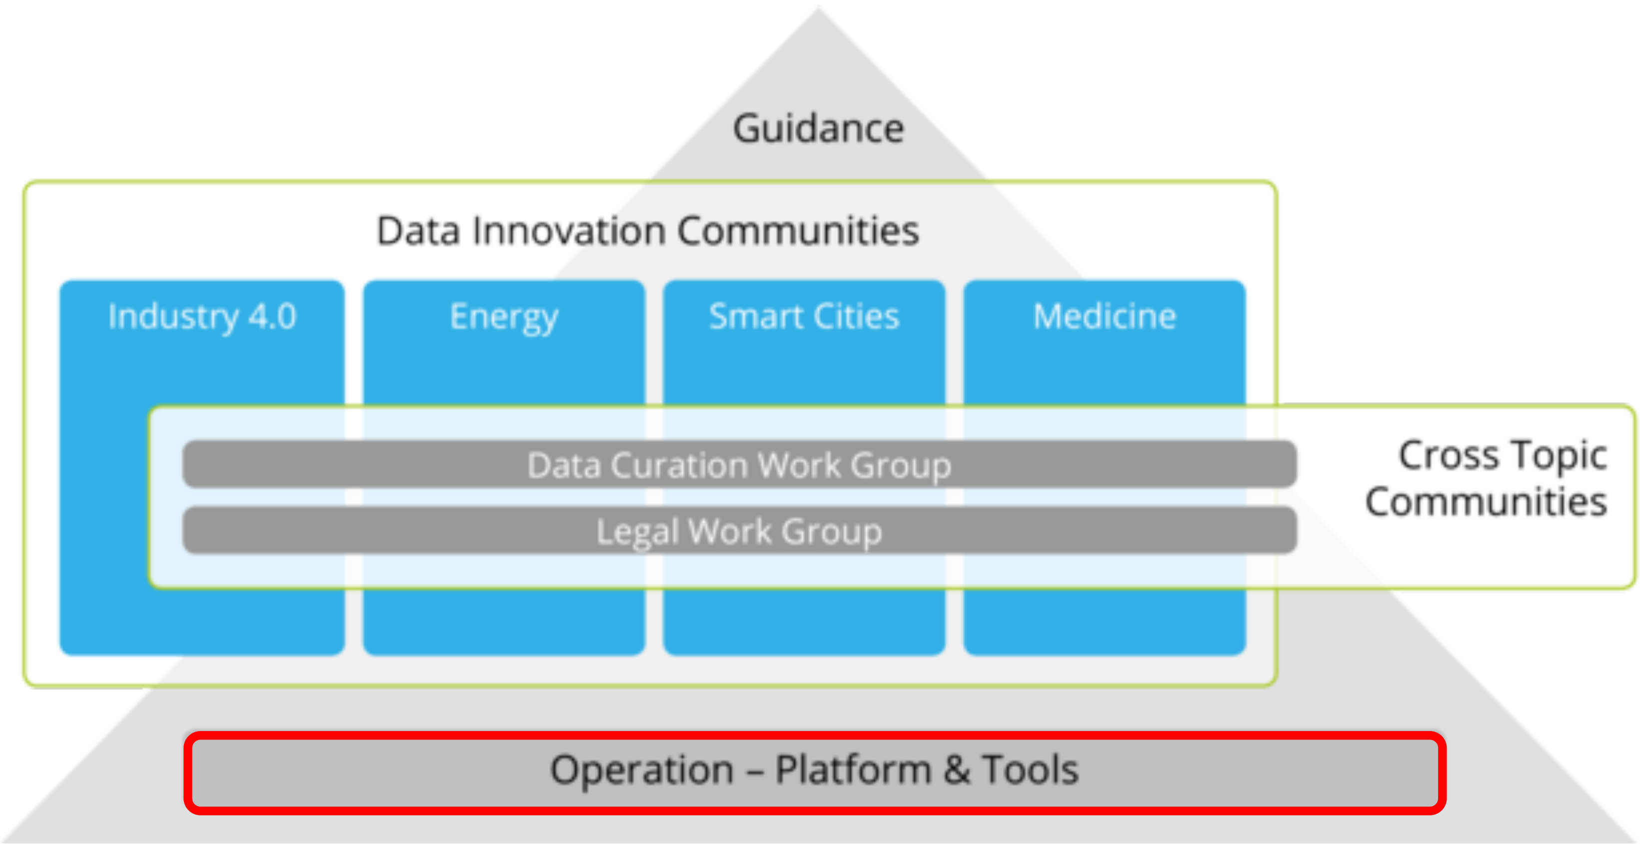

Data-Intensive Science @ SCC

- Operation of **GridKa**
  - German Tier-1 in WLCG for an international community
- Operation of the **Large-Scale Data Facility**
  - Multi-disciplinary data centre for climate research, systems biology, energy research, etc. in BaWü
- **Joint R&D&I** with scientific communities
  - Generic data management research
  - Data Life Cycle Labs in Helmholtz Programm SBD
- **Innovation driver** for SMEs, big industry und start-ups
- **Active role** in national and international projects & initiatives

SDIL Platform at SCC

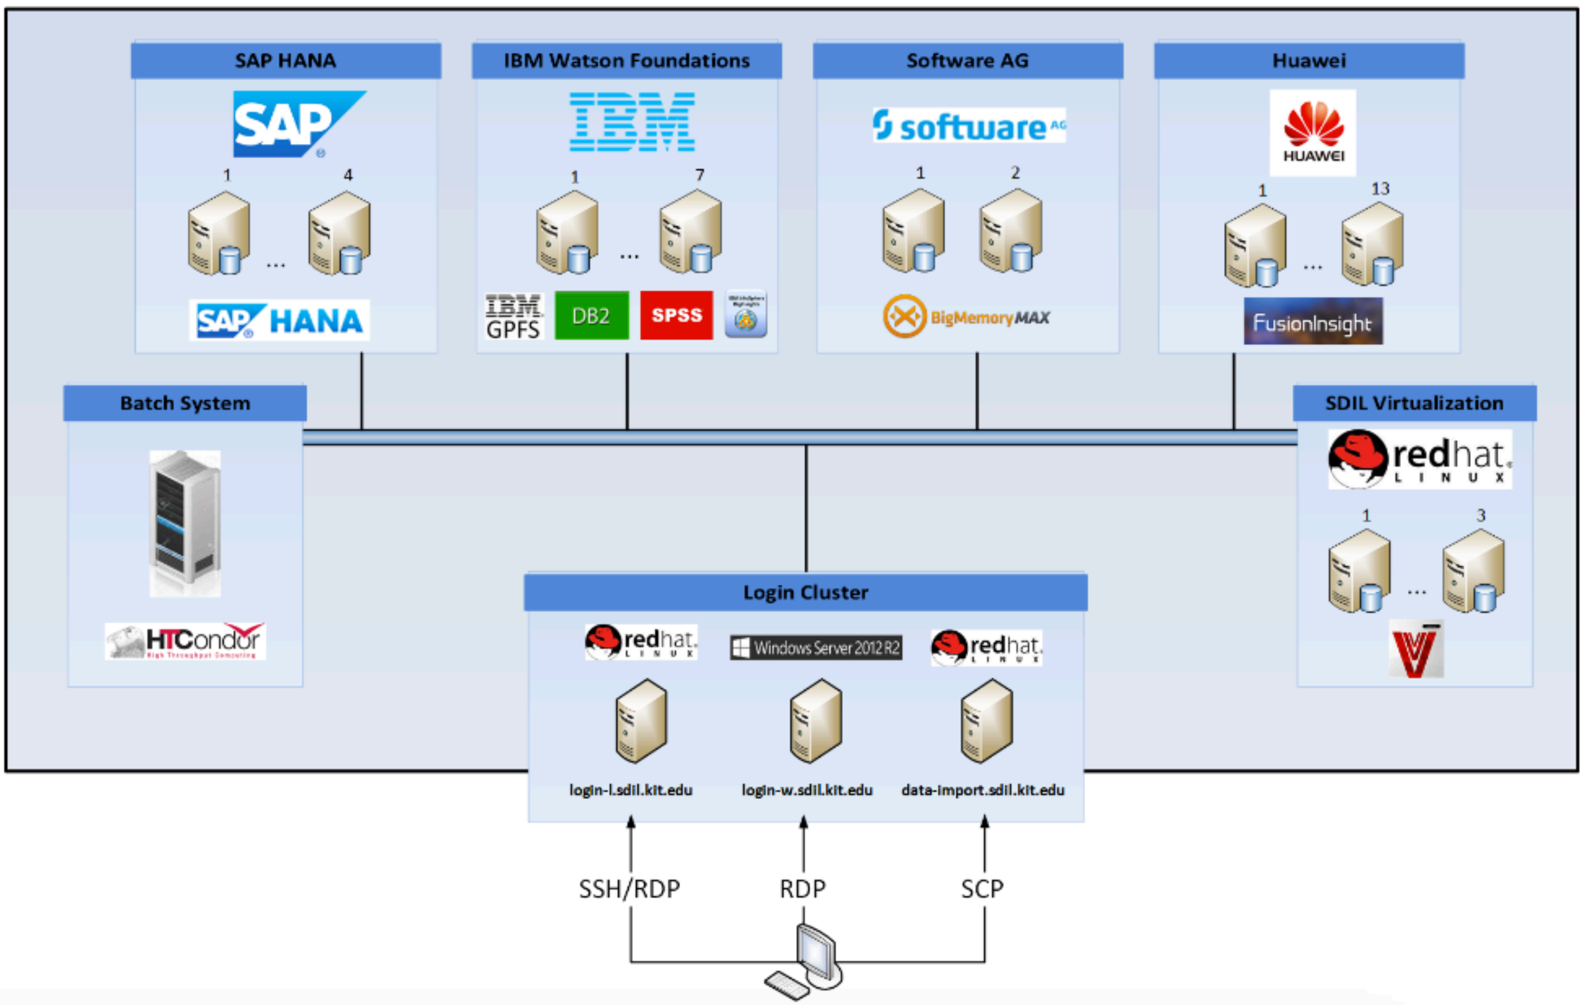

SDIL Platform Resources

	Nodes	CPU-cores	RAM	Disk Space	Network	Software
SAP HANA	4	4 x 80 = 320	4 x 1 TB	4 x 20TB	10 Gbit/s	SAP HANA Studio Predictive Analysis Library
IBM Watson Foundation Power 8	7	7 x 20 = 140	4 TB	300 TB	40 Gbit/s	IBM InfoSphere BigInsights • Hadoop • Spark • Flink IBM SPSS Modeler IBM SPSS Analytic Server DB2 with BLU Acceleration
Software AG Terracotta	2	8	256 GB	on request		Big Memory Max
Huawei	13	356	5 TB	362 TB	10 Gbit/s	Fusion Insight Platform
HTCondor		32 x 4 = 128	1 TB		1 Gbit/s	RapidMiner, Python, R
Virtualization		3 x 12 = 36	576		10 Gbit/s	Red Hat Enterprice Virtualization

SAP HANA

- SAP HANA in-memory databases operating entirely in RAM
- Persisted frequently on attached disks (H/W raid)
- Database servers setup consists of two hosts connected using Intel QuickPath Interconnects (QPI):
  - 80 cores (160 with hyper-threading)
  - 1TB RAM per database server
  - Non-uniform memory access (NUMA)

Huawei – FusionInsight

- Huawei FusionInsight:
  - Enterprise class big data platform for batch and real-time analytics
  - Includes Hadoop, Spark,
- Hardware provided by Huawei:
  - 13 servers with 365 cores
  - 5TB RAM
  - 10Gbit/s Ethernet Network
  - 362TB storage space

IBM's Watson Foundation POWER cluster

POWER8

7 nodes
140 cores
2.800 virtual systems

4 TB RAM

40 Gb/s network

260 disks with
>300 TB space

IBM Watson Foundations Software

IBM InfoSphere BigInsights

Enterprise-grade Big Data

Model-based Predictive Analytics

Semantic Text Analysis

Cognitive Computing

IBM Watson Foundation Technologies in SDIL

Karlsruhe Institute of Technology

- Open Platform with Apache Hadoop, MapReduce und Spark for distributed big data analytics
- SPSS Modeler und Server for statistical data analysis
- SPSS Analytic Server:
  - Extension of SPSS Modeler to Hadoop / Spark
- DB2 database server with in-memory BLU acceleration

IBM Watson Foundations Software

Enterprise-grade Big Data

SPSS

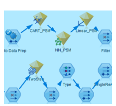

Model-based Predictive Analytics

Semantic Text Analysis

SPSS Analytics Stream Example

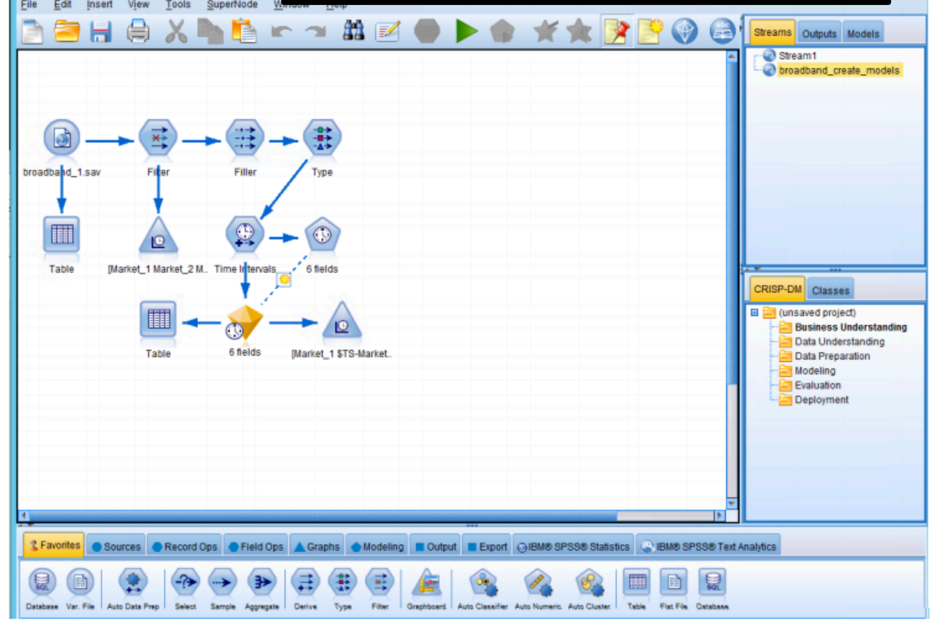

SDIL Operations – Challenges, Requirements

- Heterogeneous infrastructure:
  - Different architectures (x86_64, ppc64, ppc64le) …
  - … and operating systems (RHEL7, RHEL6, AIX, Windows)
  - Integration of different software stacks, e.g. HANA, Terracotta, Hadoop, …
- Data privacy protection, multi-tenant / user separation, permission management
- Security important → data stored in SDIL potentially sensitive
  - Strict access control of server room
  - Access only via dedicated login servers
  - Fine granular permission management for different SDIL components
- Unified storage solution to integrate different analytics products
  - Avoid migration of large amounts of data between products
- Encrypted backup of data to tape systems for disaster recovery

SDIL Operations – more details …

- 3 networks –
  1. BMC
  2. Deployment / configuration (private)
  3. Services / data (public)
- Two firewalls, strict access control (Federated LDAP), Encryption
- Need to manage (abstract) configuration of services independent of architecture / OS / compute resource
  - Foreman and Puppet for deployment and configuration
- Icinga for monitoring (integrated with puppet)
- Ticket system for operation tasks
- Weekly shifts for incident response

Storage Setup

- Unified storage solution to avoid migration of large amounts of data
- Shared file system for home and data directories
- Clustered parallel fs Spectrum Scale (formerly GPFS) on Power8
- Full Posix compliance including distributed file locking
- Heart beat / quorum mechanism to avoid “split brain” situations
- Supports HDFS over GPFS as additional protocol, i.e. direct access from hadoop / spark
- File Placement Optimizer (FPO)
  - Set of feature extensions allowing GPFS as a shared nothing architecture
  - Extended failure groups: Disks in a failure groups share common point of failure, e.g. server, switch, rack, computing center
- Redundant copies of data / metadata allow for transparent upgrade of storage cluster
- Home and project directories with snapshots

IBM Watson Cluster Architecture / PowerVM

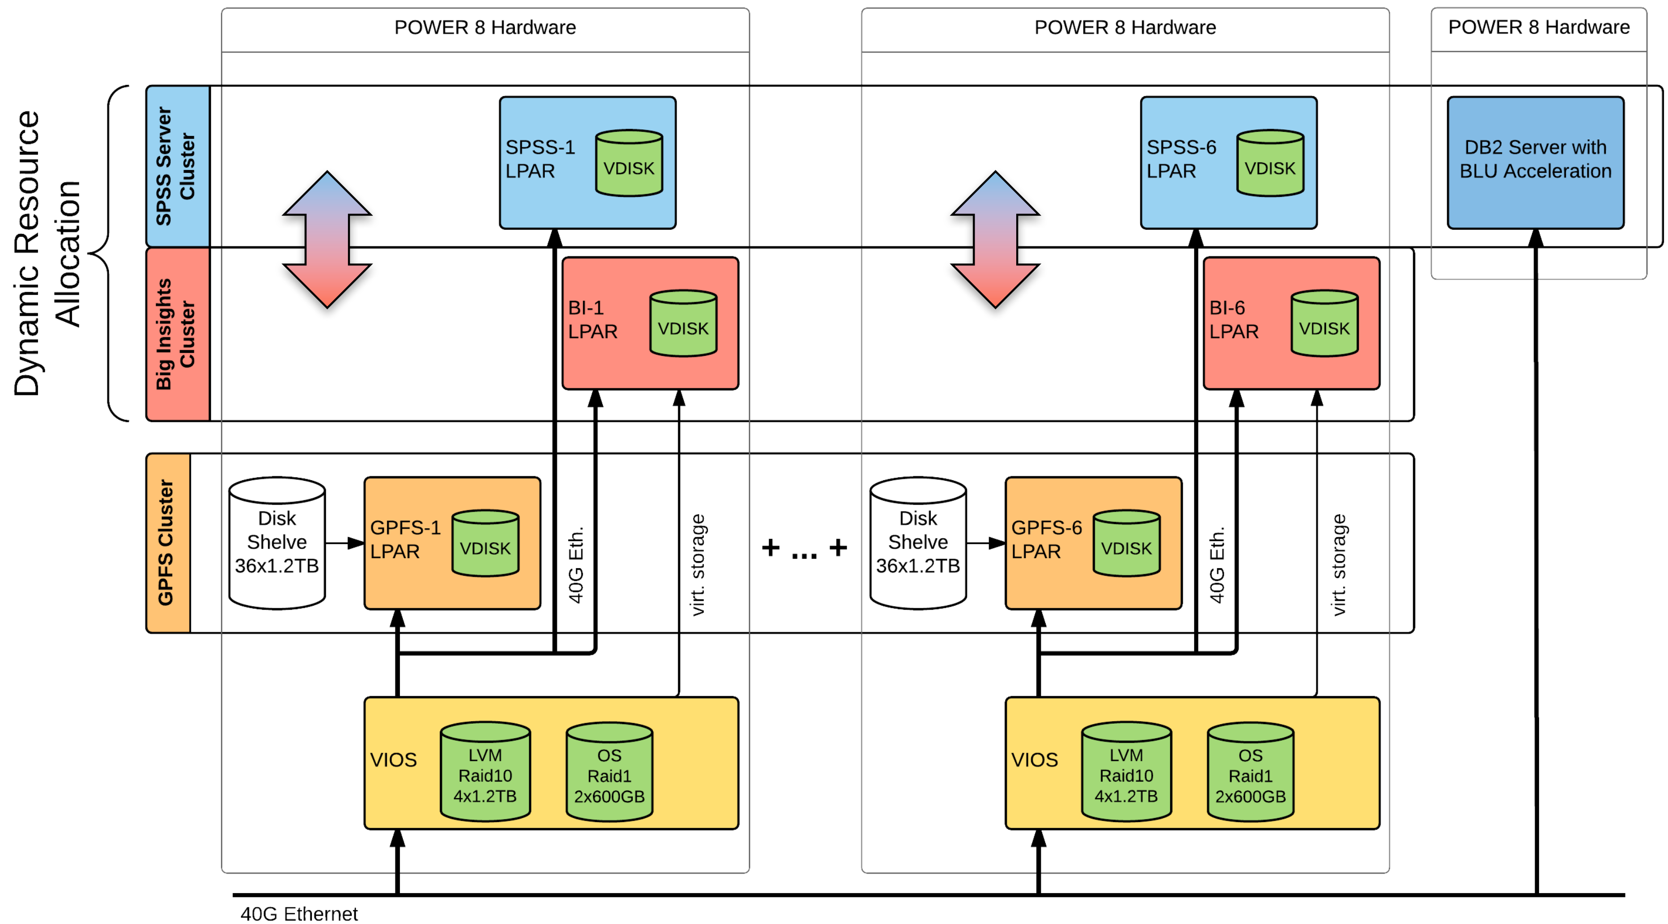

Red Hat Enterprise Virtualization

- Virtualization cluster to provide compute resources / services to users and projects “on demand”
- Using Red Hat Enterprise Virtualization (RHEV) based on KVM
- Hyper-Converged setup with GlusterFS
  - Integrates storage, compute and network in commodity hardware
  - Storage and compute on same nodes / hypervisors
  - 3-way replica of VM images in GlusterFS

Deployment / Configuration

- Using Katello / Foreman and Puppet for deployment and configuration
- Configuration of whole platform as puppet code
- Gitlab as code repository with continuous integration
- Separate development / production configuration and software stacks
  - Allows for promotion of tested setups to production
- Most SDIL machines stateless → easy redeployment any time
- Katello / Foreman integrates with virtual compute resources
  - Automatic creation / destruction of virtual machines
  - Out of the box for KVM, started to implement integration for PowerVM

Deployment / Resource Management

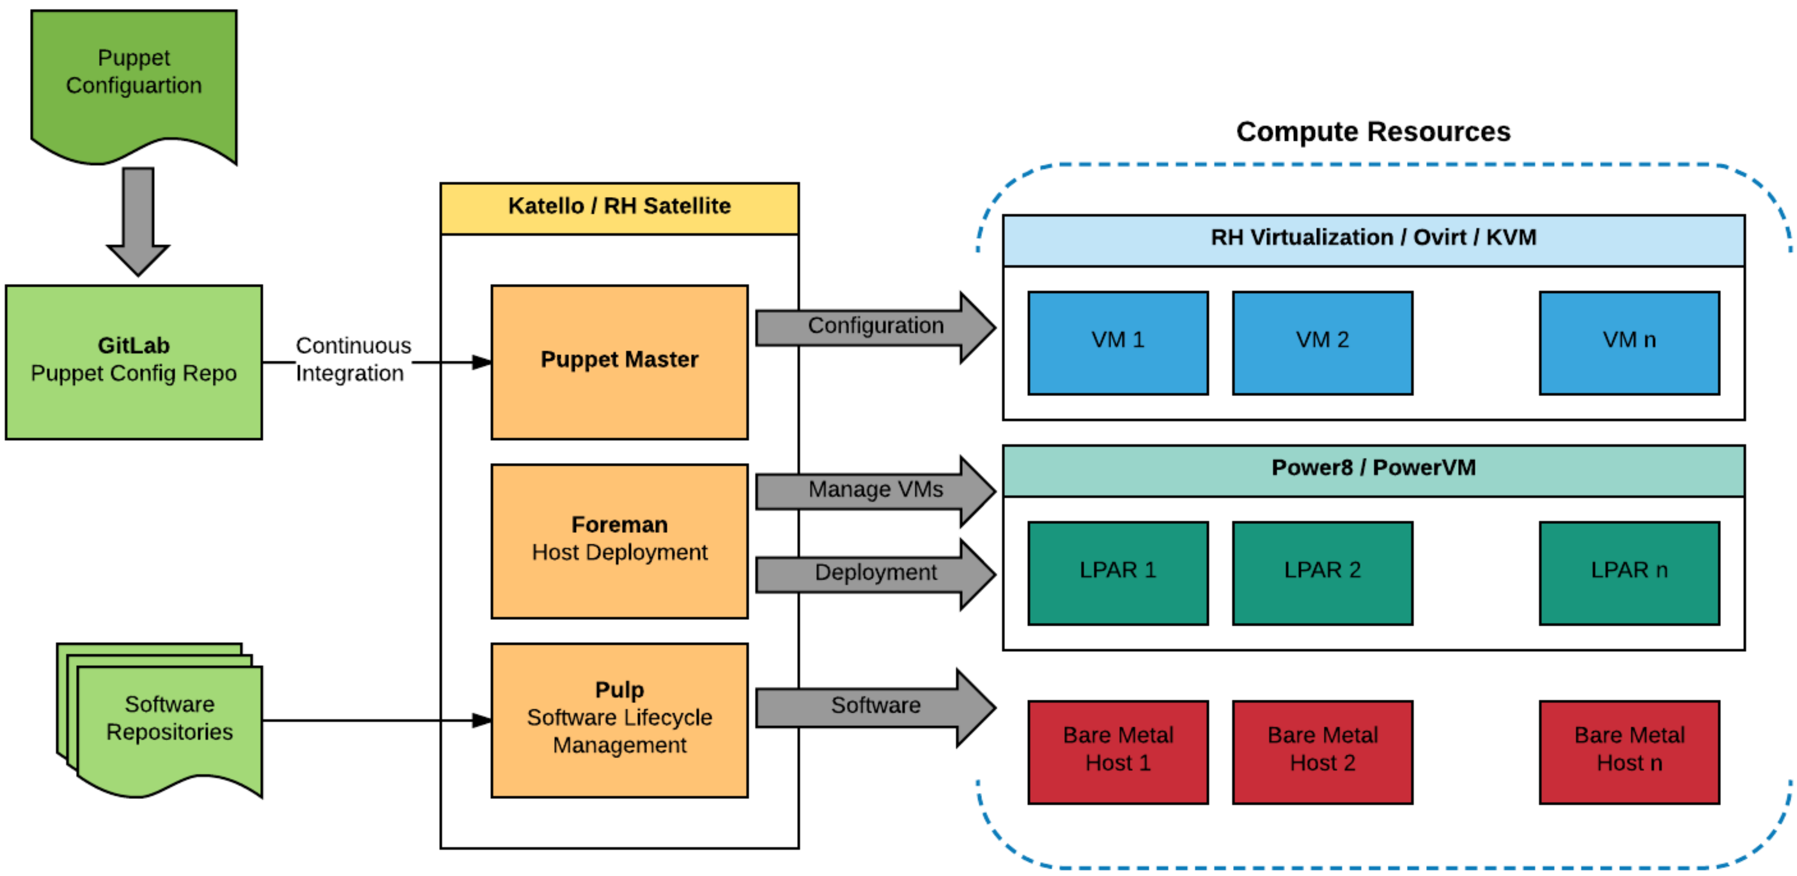

Interactive Data Analytics with Jupyter Notebooks

- Setup service to provide Jupyter notebooks via central multi-user hub
- Web application which allows to create and share documents containing code, visualizations and explanatory text
- Notebooks can serve different kernels:
  - We provide Python (2.7 and 3.3) and R by default
  - Users can install their own custom kernels
  - Users can setup their own software environments using conda
- Also provide kernel allowing users to access IBM Big Insights Spark installation via python notebooks on Power8
- Spawned notebooks with restricted resources (RAM, CPU, Time)
  - Using exiting HTCondor batch system for resource management
- Provide tools to submit notebooks non-interactively to the batch system

Interactive Data Analytics with Jupyter Notebooks

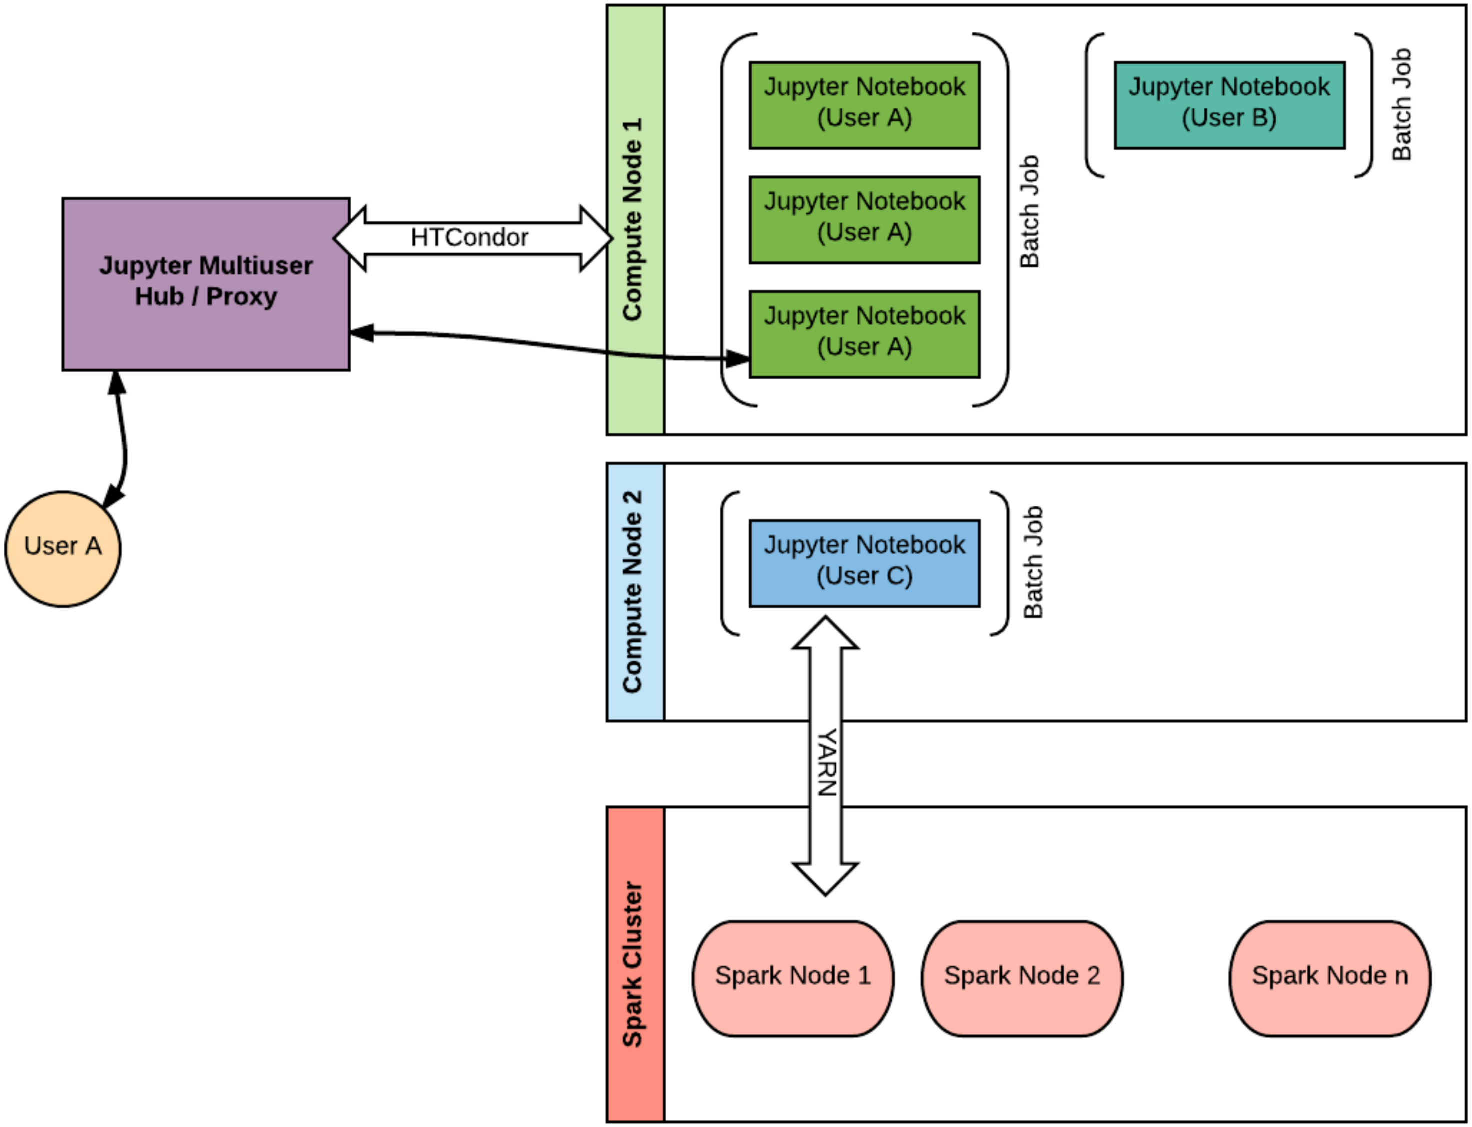

Jupyter Hub Resource Selection

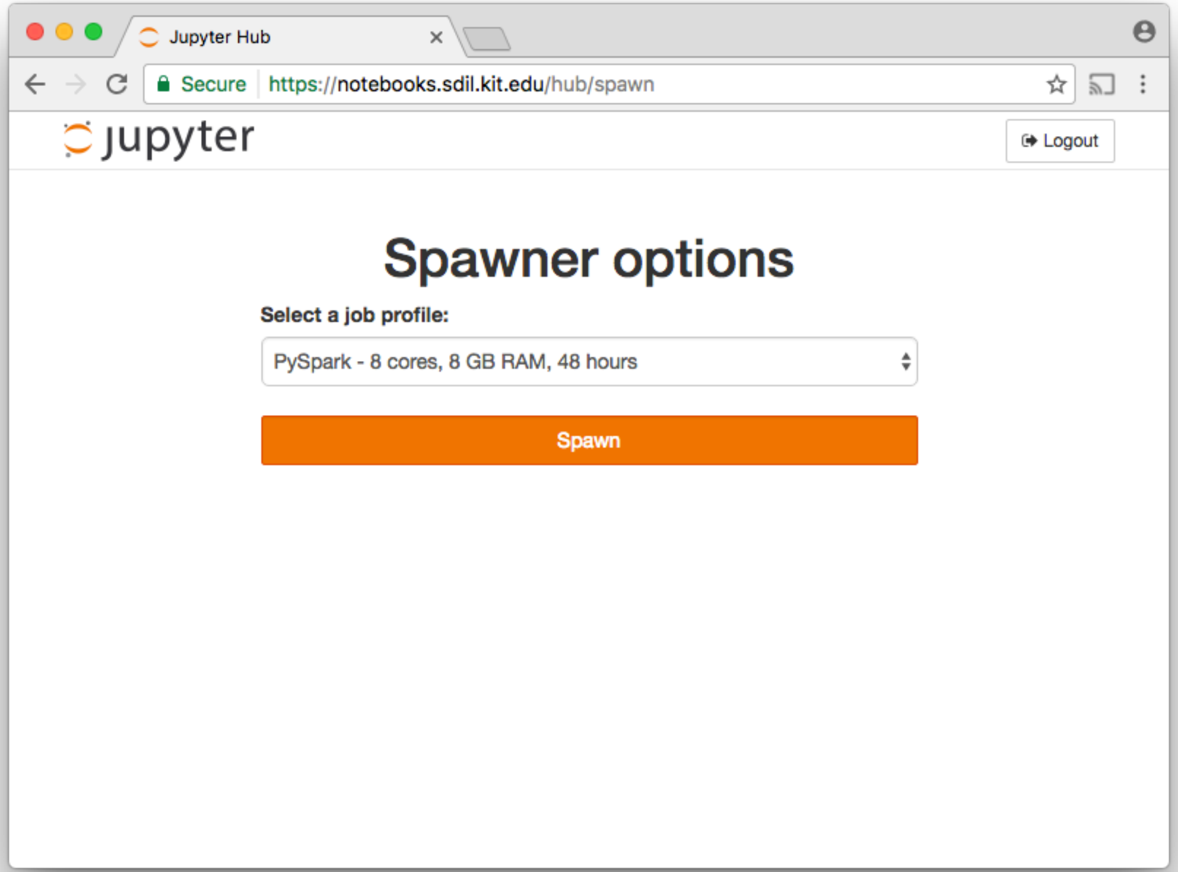

Jupyter Notebook

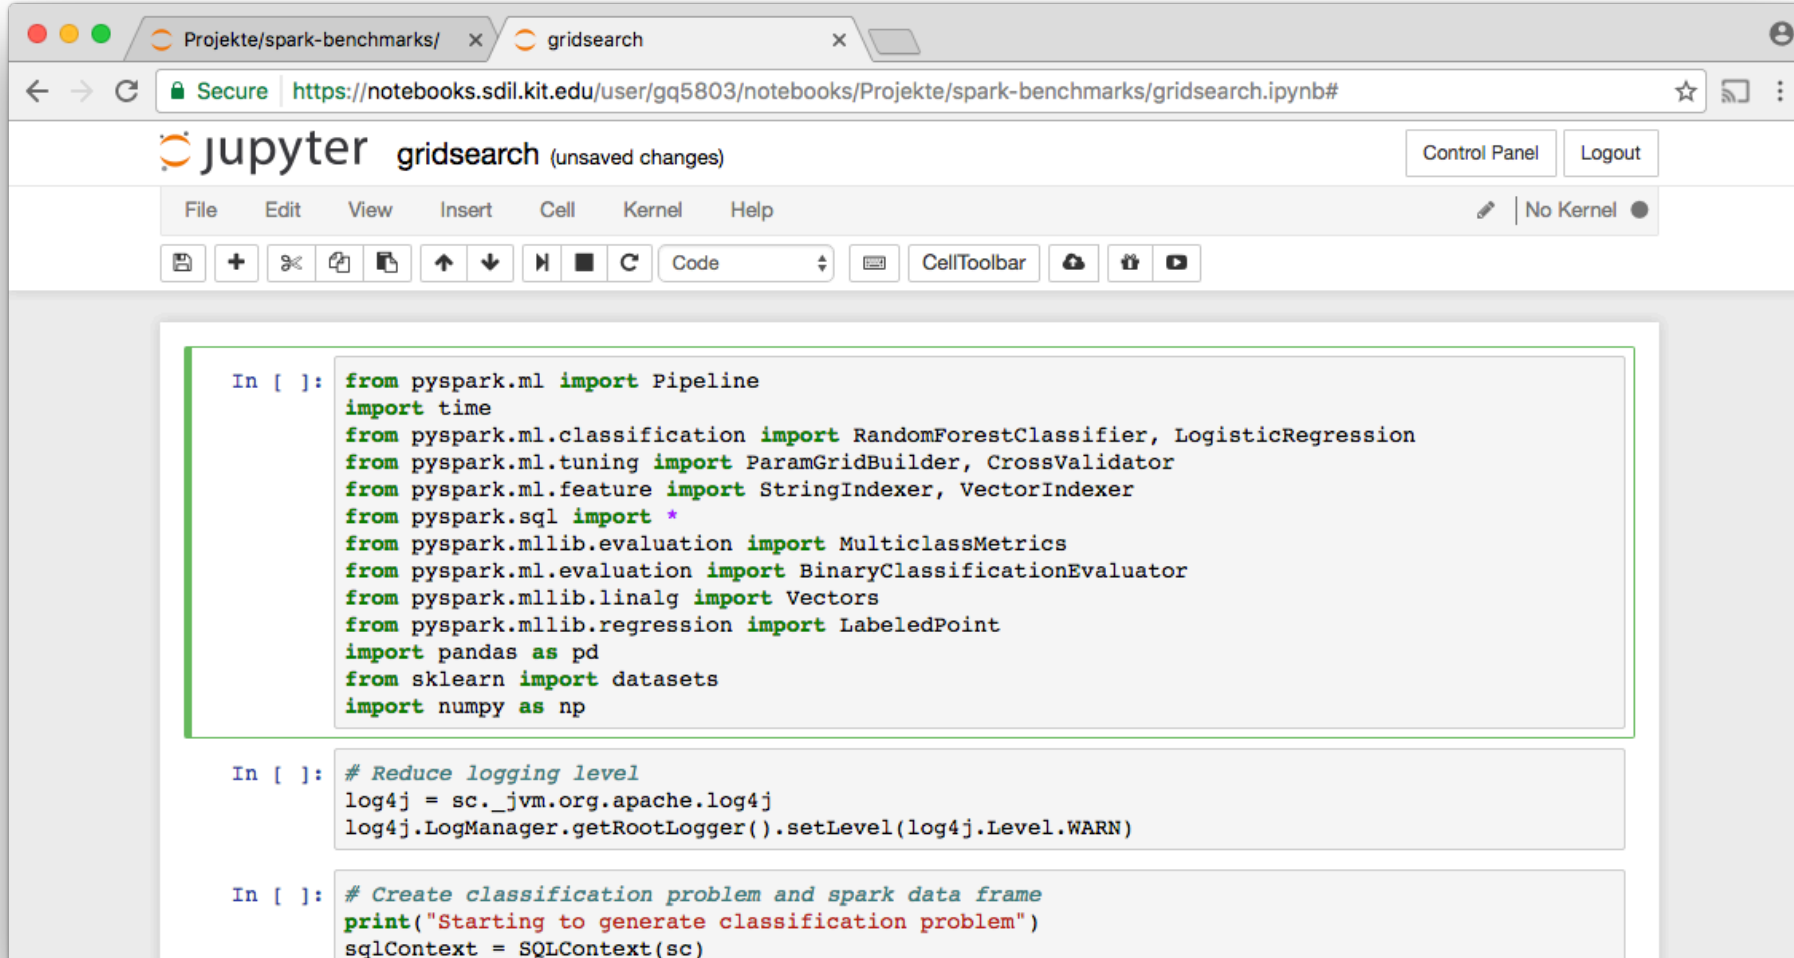

Summary

- SDIL offers wide range of data analytics technologies
- Data protection and security important
- Operation requires integration of very heterogeneous environment (software, os and hardware wise)