Operational Data Analytics

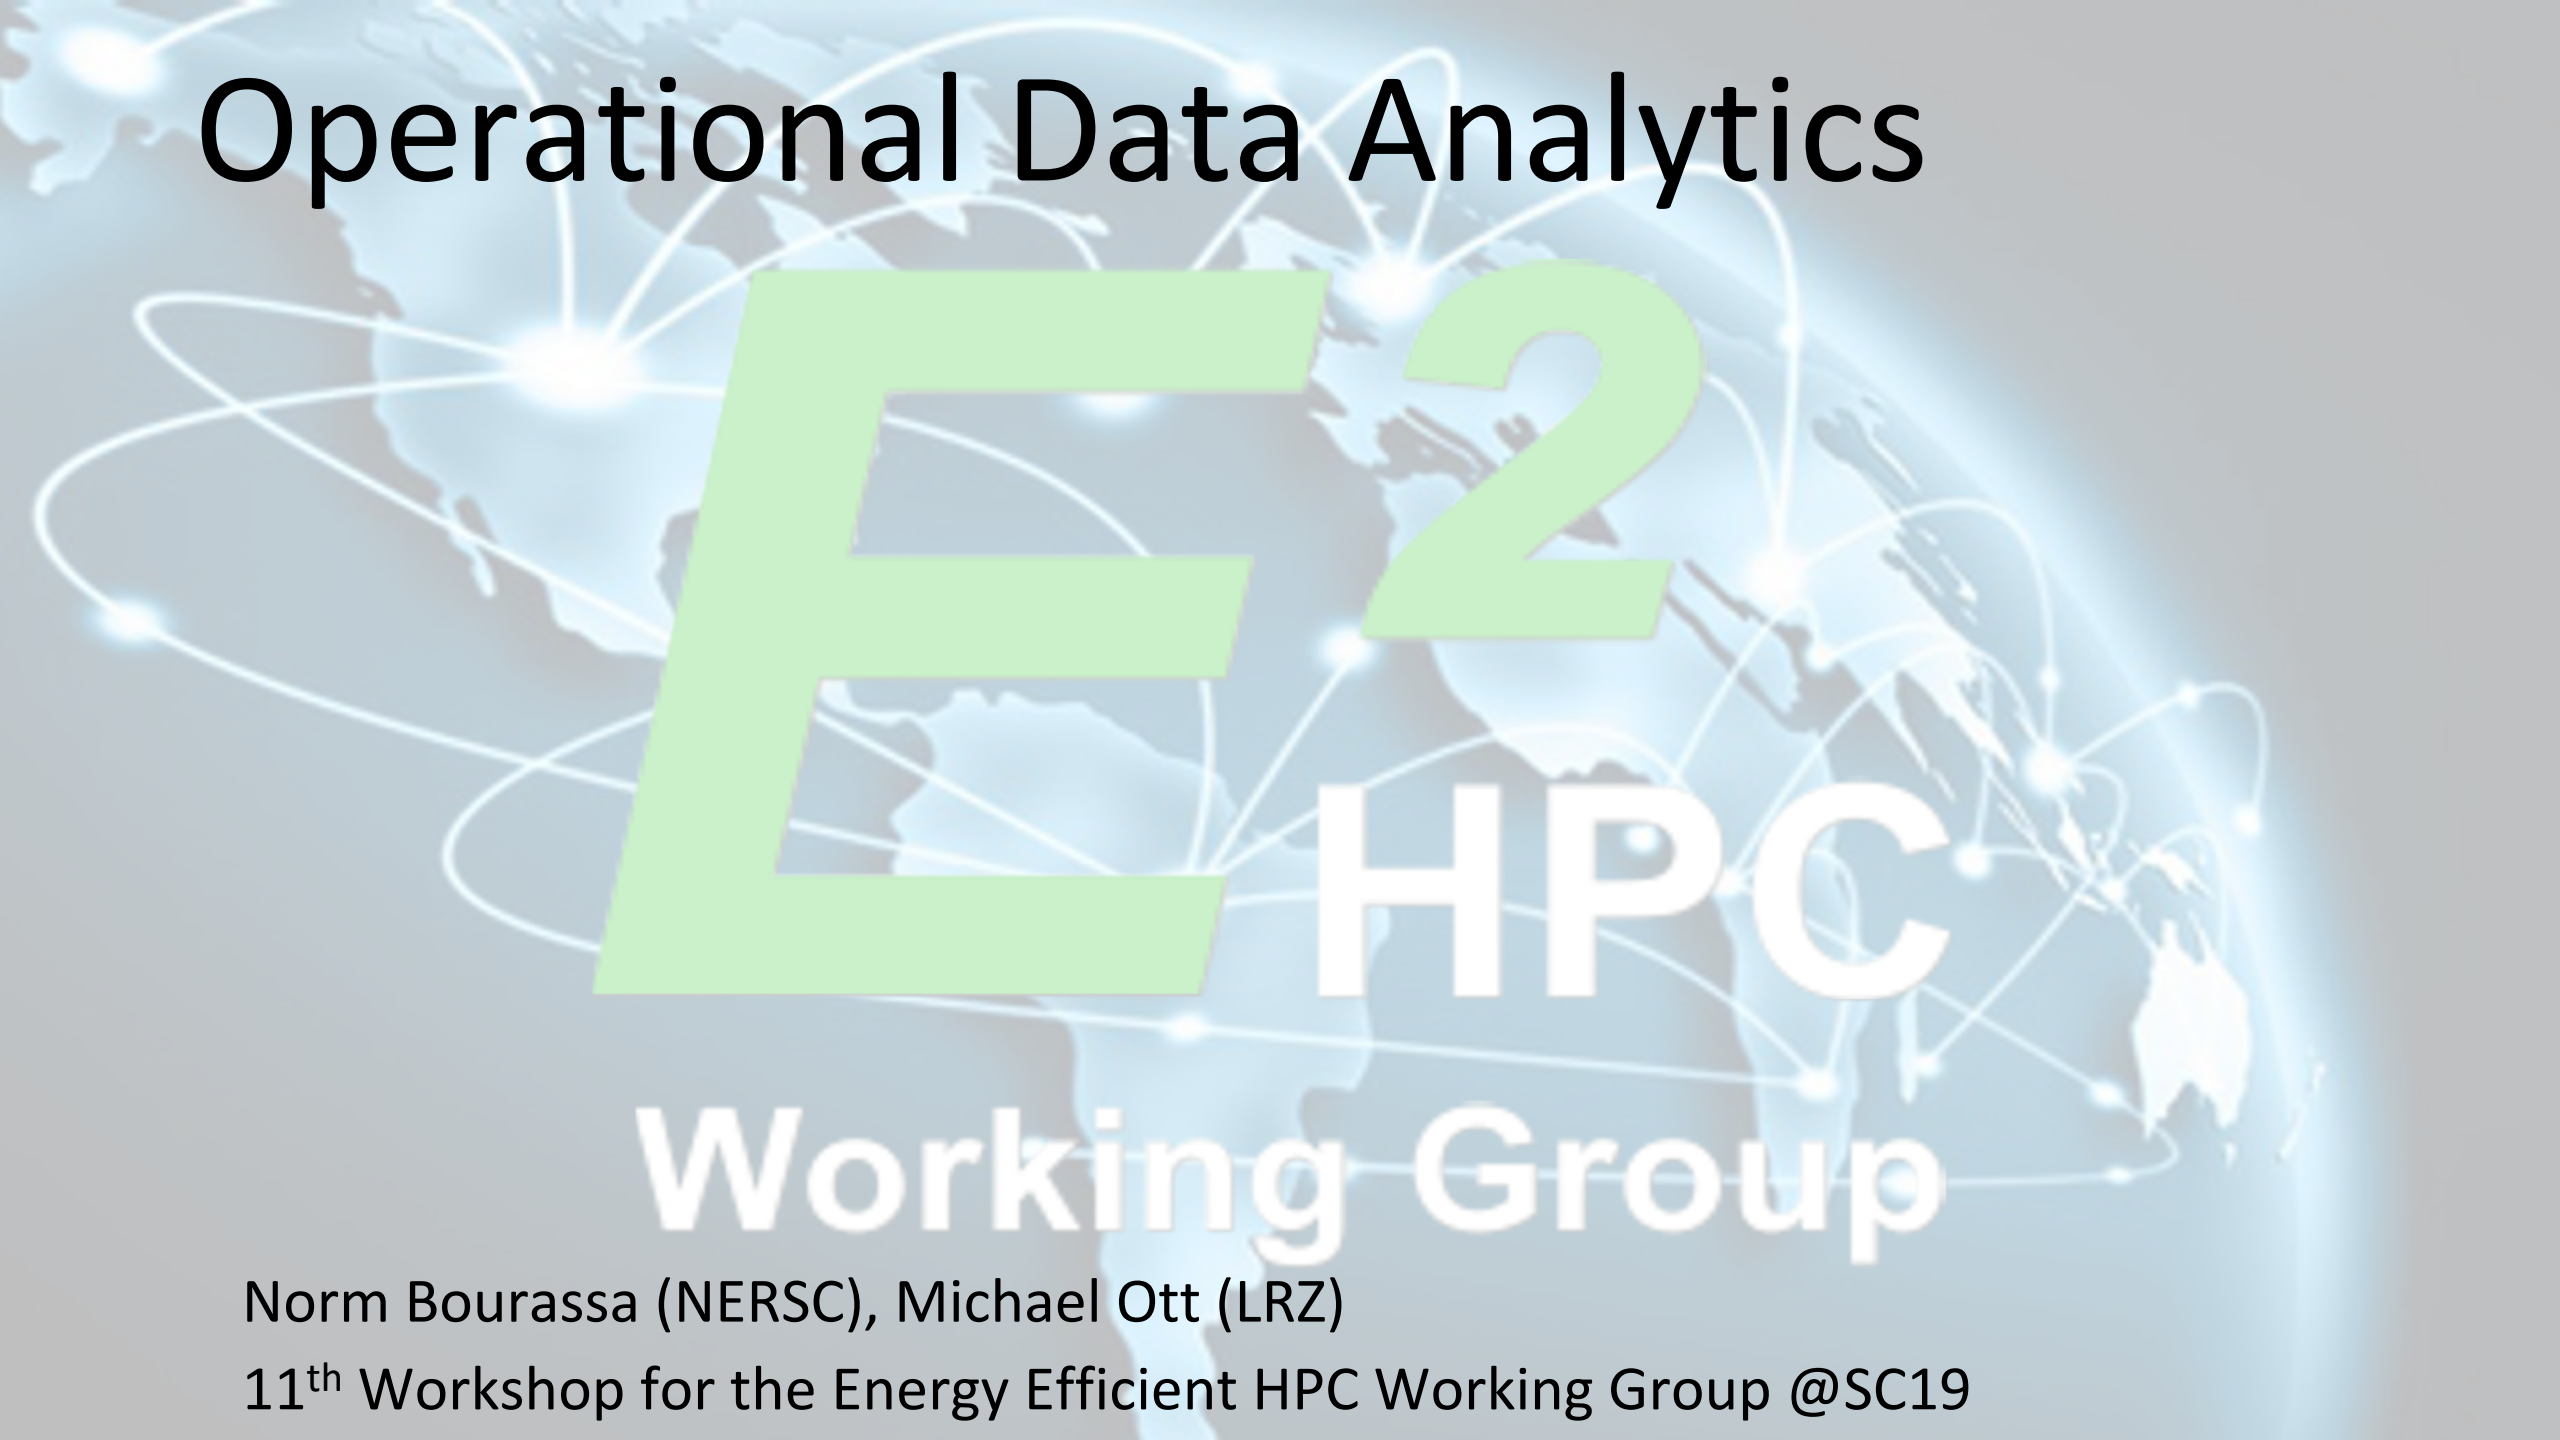

Norm Bourassa (NERSC), Michael Ott (LRZ)

11th Workshop for the Energy Efficient HPC Working Group @SC19

Mission

- First early adaptor sites are starting to deploy data collection and analysis systems that integrate HPC platform data with facility data
- As a team, we want to learn and understand:
  - What tools are the sites using for this?
  - How are they using the collected data?
  - What are the lessons learned?
- Provide guidance for the community to deploy similar systems

Initial Survey: Identify Interest

- Has your site implemented data collection, aggregation and analysis for operational management of both facilities and HPC systems (including energy and power management) in a production environment on at least one large-scale system (Top500 sized system) with integration (or plans for integration) that extends from the HPC data center down through the platform to the CPU?
- Would you participate in an more in-depth questionnaire and interview?
- Survey population and responses
  - Original sites: LLNL, LBNL, LANL and NREL
  - Initial survey sent to 20 additional sites:
    - Yes (7): ORNL, LRZ, NCAR, RIKEN, CINECA, Juelich, AIST
    - Maybe (6): Kaust, NCSA, BSC, Argonne, Tokyo Tech, CEA
    - No (2): CSCS, HLRS
    - Unclear/No Response (5): NASA, TACC, JCAHPC, ECMWF, GENCI

In Depth Survey

The Team is particularly interested in understanding:

- How sites are really using their data collection systems with specific use case scenarios.
- To what extent are data streams from HPC platform and infrastructure integrated into a common platform.
- Implementation issues; such as system-to-system interface.
- Known or potential scalability constraints.

Architecture

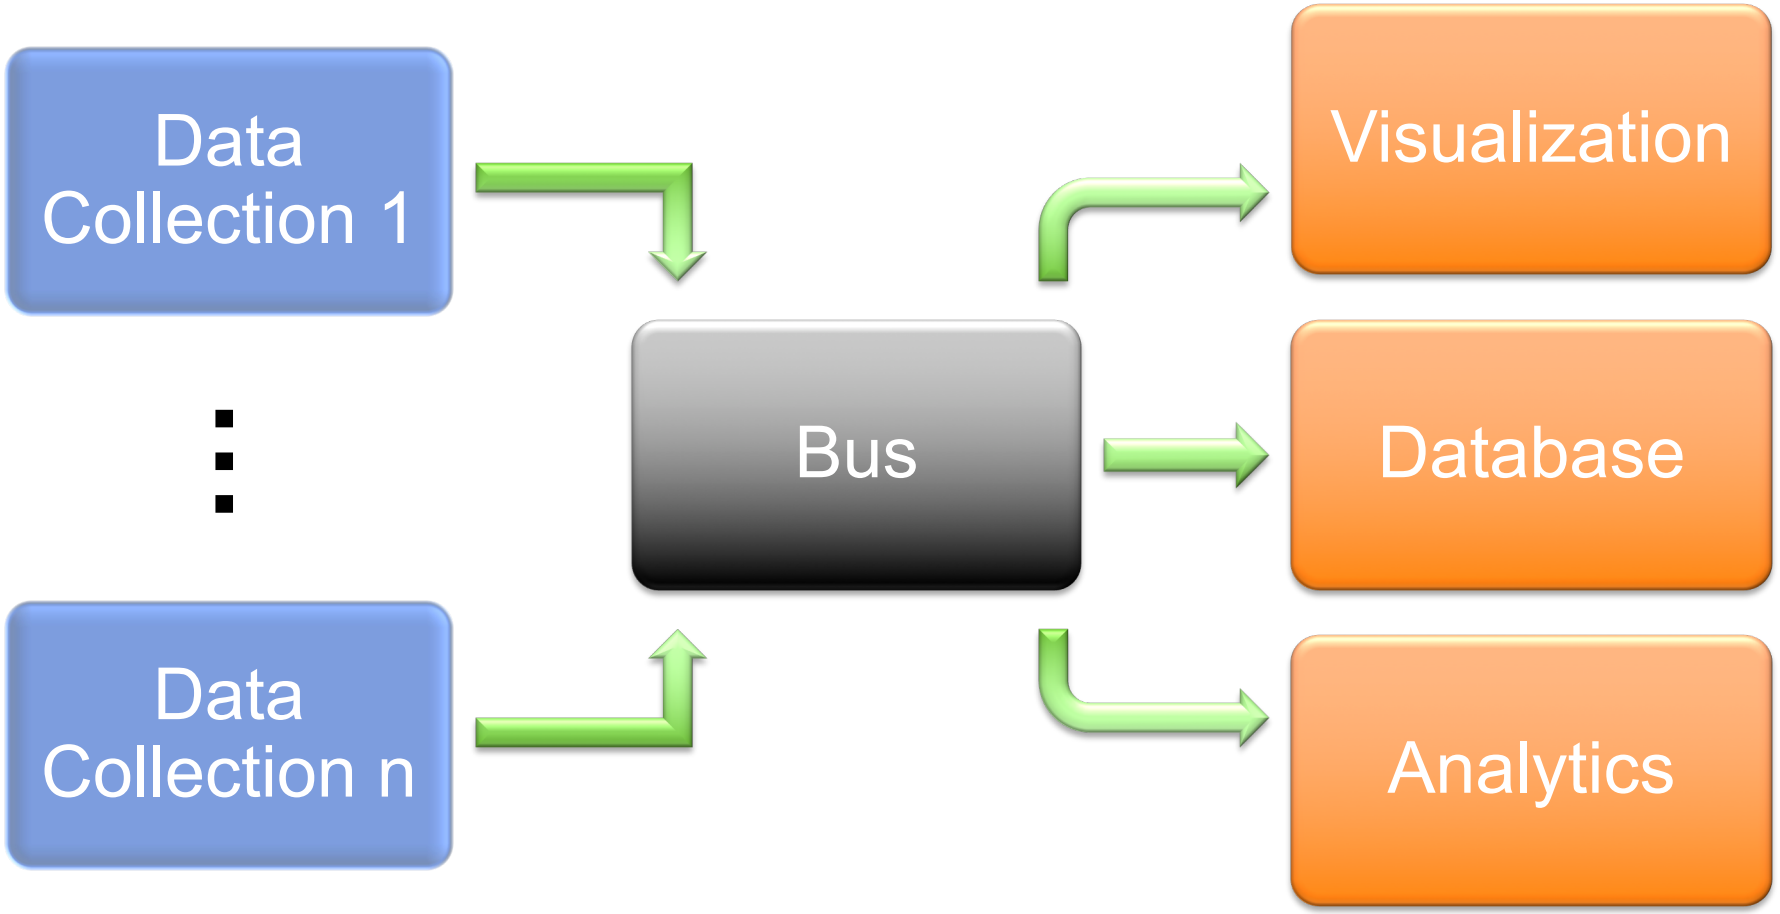

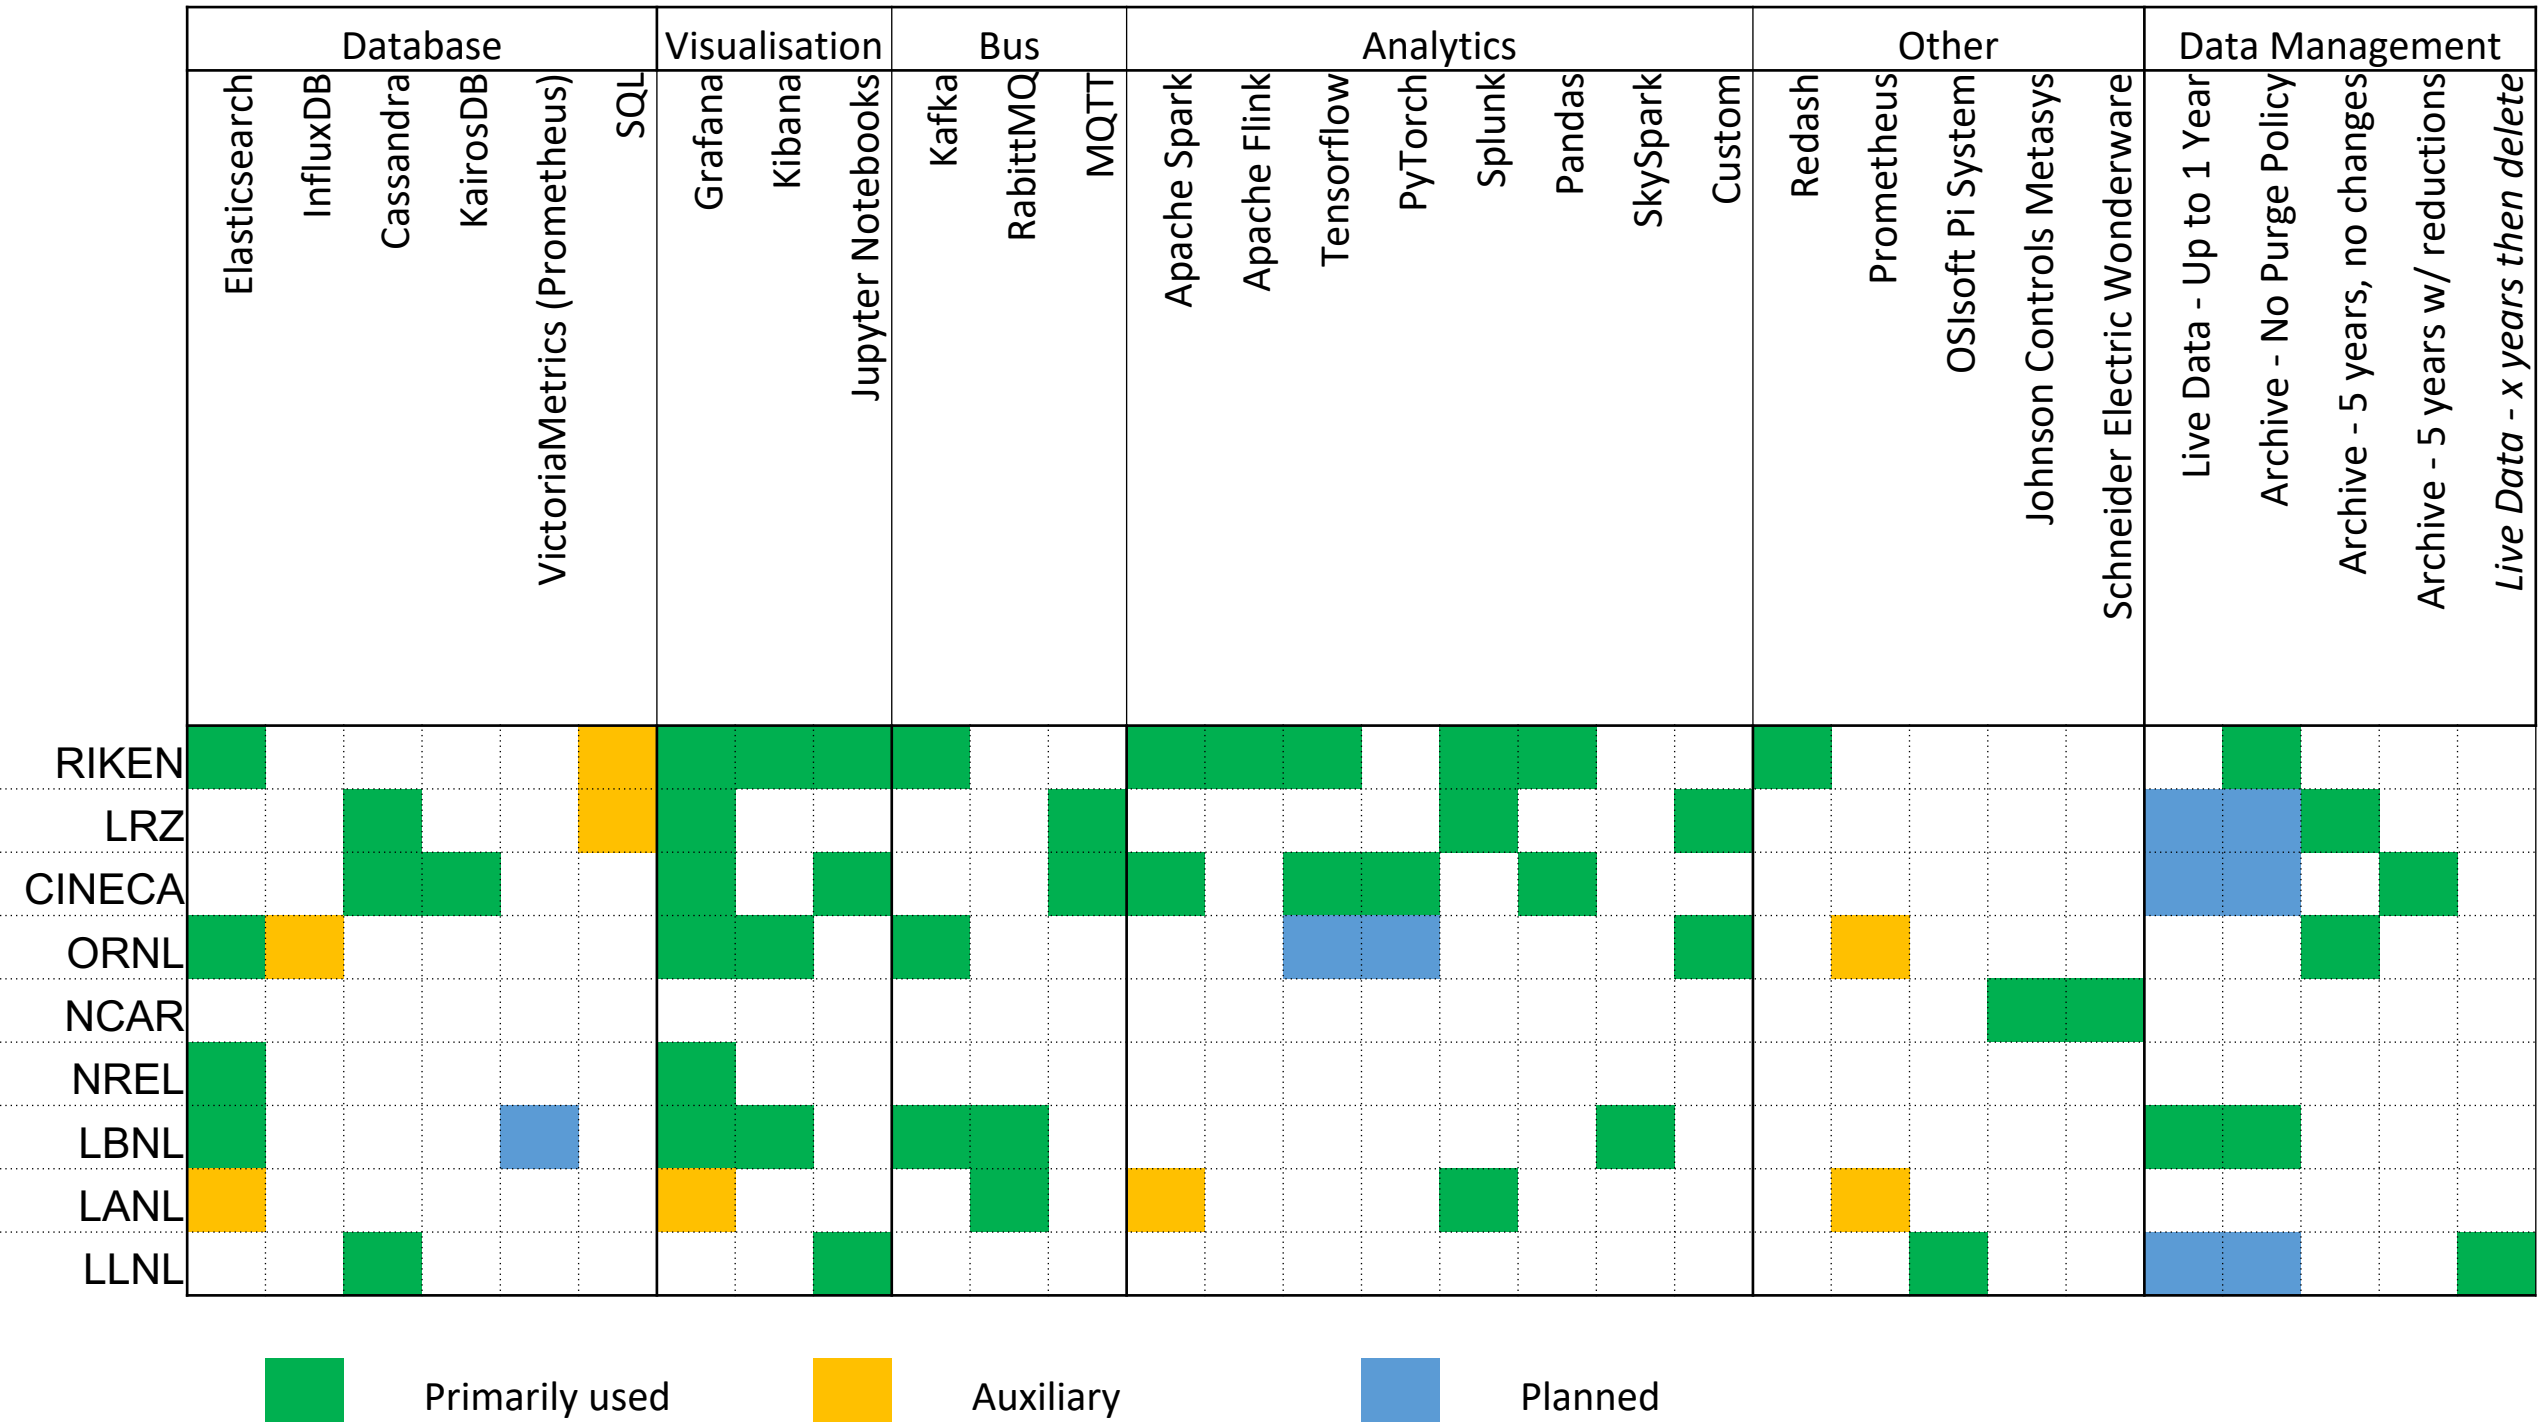
Database
Visualisation
Bus
Analytics
Other
Data Management
Elasticsearch
InfluxDB
Cassandra
KairosDB
VictoriaMetrics (Prometheus)
SQL
Grafana
Kibana
Jupyter Notebooks
Kafka
RabittMQ
MQTT
Apache Spark
Apache Flink
Tensorflow
PyTorch
Splunk
Pandas
SkySpark
Custom
Redash
Prometheus
OSIsoft Pi System
Johnson Controls Metasys
Schneider Electric Wonderware
Live Data - Up to 1 Year
Archive - No Purge Policy
Archive - 5 years, no changes
Archive - 5 years w/ reductions
Live Data - x years then delete
RIKEN
LRZ
CINECA
ORNL
NCAR
NREL
LBNL
LANL
LLNL
Primarily used
Auxiliary
Planned

Use Cases

- Adoption and use cases has been most extensive with either infrastructure or HPC system data
- There are use cases for integrated infrastructure and HPC system, but they are limited and very early adoption
- Most use cases fall in one of these categories:
  - Cooling/infrastructure optimization
  - Power management
  - Strategic planning
  - Application performance/power tuning

Use Case: Infrastructure Optimization (NCAR)

Scenario	CDU pumps failing due to variable frequency drive (VFD) failures
Data Used	Infrastructure
Analysis	Correlation of CDU monitoring data with power quality meter showed brown out conditions in the same time frame
Action	Move CDU pumps to UPS-backed power
Value Add	Increased reliability, avoided downtime of HPC system

Use Case: Application Tuning (LRZ)

Scenario	Reduce CPU frequency to lower power consumption
Data Used	HPC Platform
Analysis	Memory bound application can run at lower frequency without hurting performance but lowering their energy-to-solution
Action	Use performance counters to assess application characteristics and DVFS to underclock memory bound applications
Value Add	Savings estimated at 1.8M € for SuperMUC Phase1 and Phase2

Use Case:
Cooling Optimization (ORNL)

Scenario	Optimize chilled water temperature and flow for energy efficiency
Data Used	Infrastructure & HPC Platform
Analysis	Correlate water temperatures & flow rates with CPU/GPU temperatures
Action	Optimize temperature and flow for energy efficiency, maintain safe & un-throttled operational conditions of HPC system
Value Add	Increased energy efficiency, lower OPEX

Use Case:
Power Management (RIKEN)

Scenario	Ensure that contractual power limit with utility is not exceeded
Data Used	Infrastructure & HPC Platform
Analysis	Monitor datacenter power draw, compare with utility rate limit and HPC system application job schedule
Action	Kill high-power jobs to reduce total power draw
Value Add	Avoid paying fine, lowering OPEX

Use Case:
Cooling Optimization (LBNL)

Scenario	Optimize computer room air handling units (AHU)
Data Used	Infrastructure & HPC Platform
Analysis	Use actual HPC air demand to control AHU instead of rule of thumb
Action	Aggregate individual server air demand (air temperatures, fan speed) using Prometheus client and feed this into AHU controls
Value Add	Increased energy efficiency, lower OPEX

Conclusion

- Sites are collecting and integrating data from HPC platform and infrastructure for analysis and visualization purposes.
- Scalable open source tools are becoming more available for implementing such pipelines.
- Analysis currently is mostly manual by visual inspection, use of AI/ML is still an active research activity.
- ODA can pay for itself by lowering OPEX.

Still more survey questions

- If you had to do it over again, what would you do differently?
- What did you want to achieve when you first began?
- What have been some of your biggest challenges or obstacles?
- Do you have any lessons learned?
- What questions do you have for your peers?