

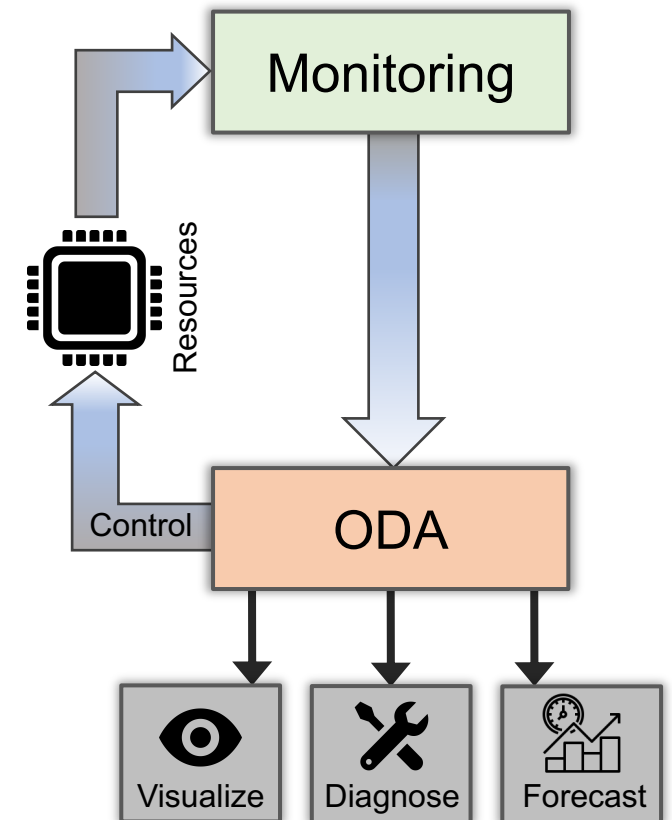


# A Conceptual Framework for HPC Operational Data Analytics

EEHPC State of the Practice Workshop 2021, 07.09.2021 | Alessio Netti,  
Woong Shin, Michael Ott, Torsten Wilde, Natalie Bates

## Introduction

- *Operational Data Analytics (ODA)* uses monitoring to extract actionable **knowledge** on system **behavior**.
  - Can improve energy efficiency and reliability.
  - More and more data centers use ODA.
- However, ODA is a **broad** and **diverse** field:
  - Predictive tuning of CPU frequencies is ODA.
  - Diagnosing infrastructure failures with ML is ODA.
  - Simulating scheduling policies is ODA.
  - Computing a data center's PUE is ODA.



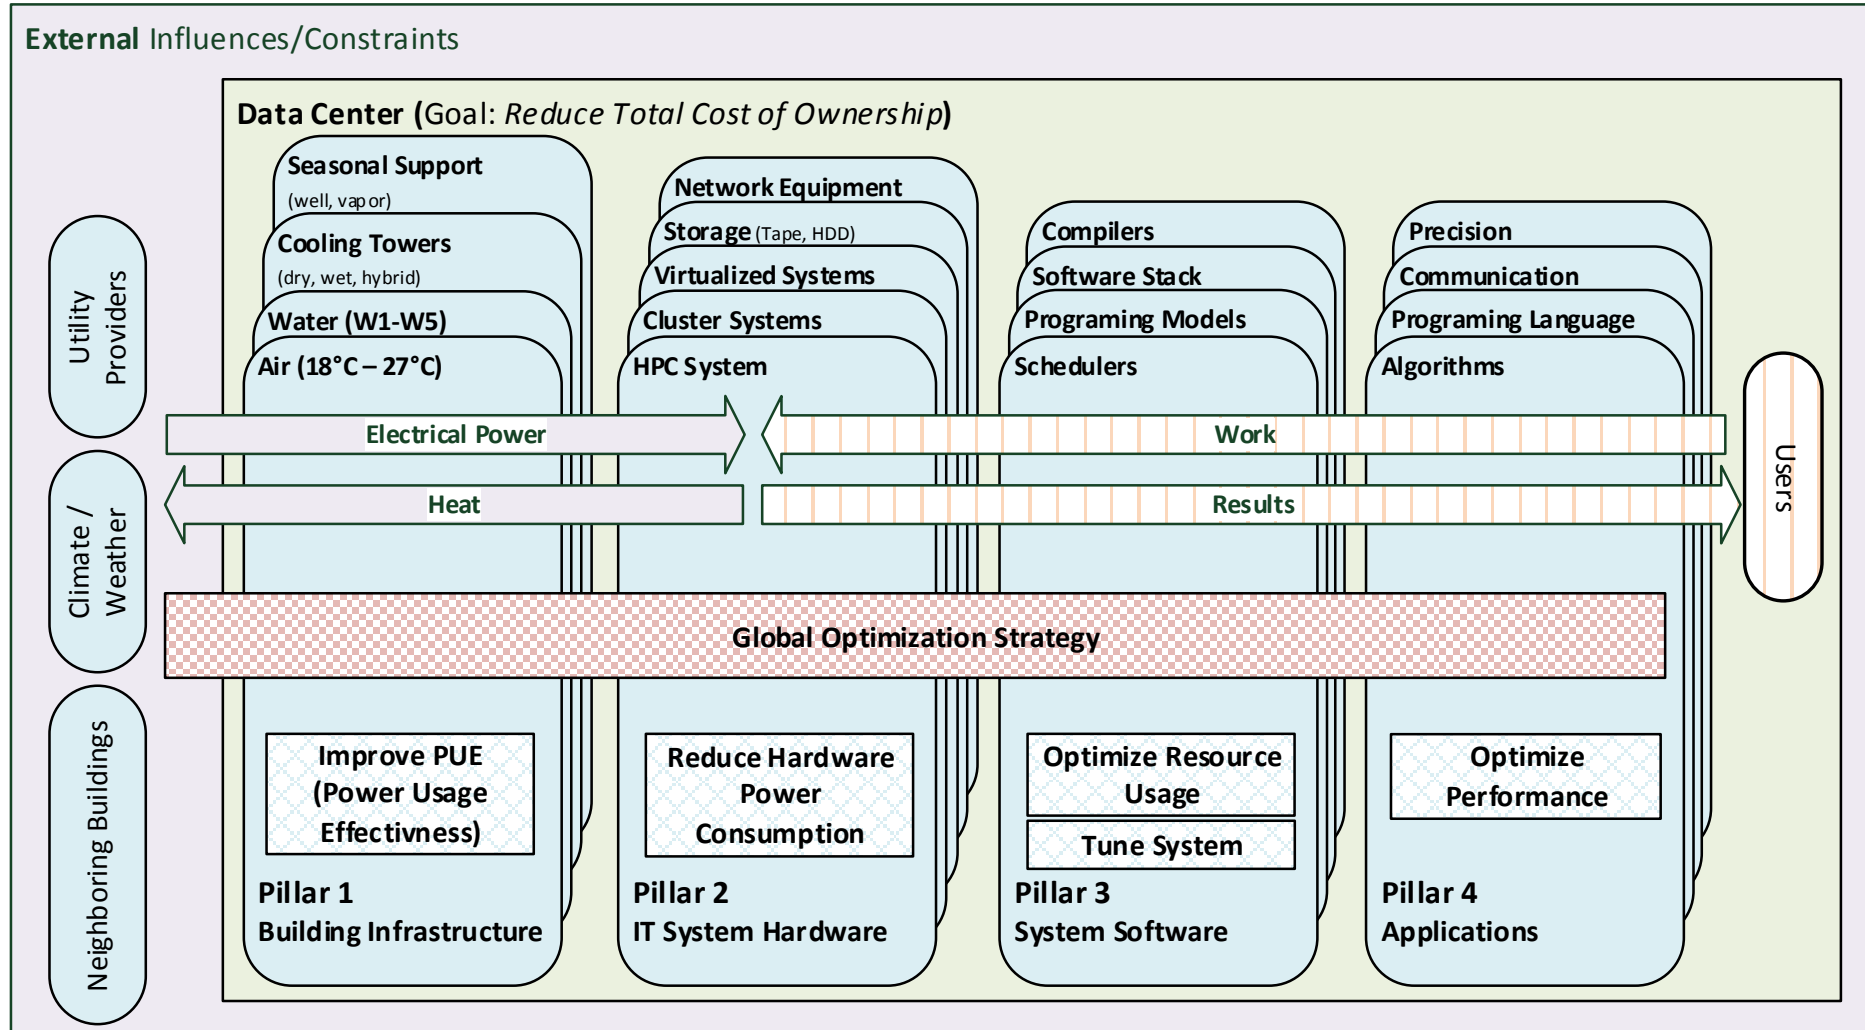
## Contributions

- There is no **common language** to reason about ODA.
  - Research gaps and opportunities are difficult to identify.
  - System design and requirements are not standardized.
  - Adoption of ODA by data centers is cumbersome.
- Our **contributions** are the following:
  - A conceptual framework to help classify ODA use cases.
  - An extensive survey and categorization of ODA literature.
  - Demonstration of the framework on state-of-the-art use cases.

## Designing a Framework for ODA

- Many possible **questions** about an ODA use case:
  - What is the functional **complexity** and data center scope?
  - How do we **decompose** it in simple, standard blocks?
  - Have other people already tackled a **similar** problem?
  - What are the deployment **requirements** and gains?
- We use two state-of-the-art frameworks as a **foundation**:
  - The “*4-Pillar Framework for Energy-Efficient HPC Data Centers*”.
  - The “*4 Types of Data Analytics Framework*”.

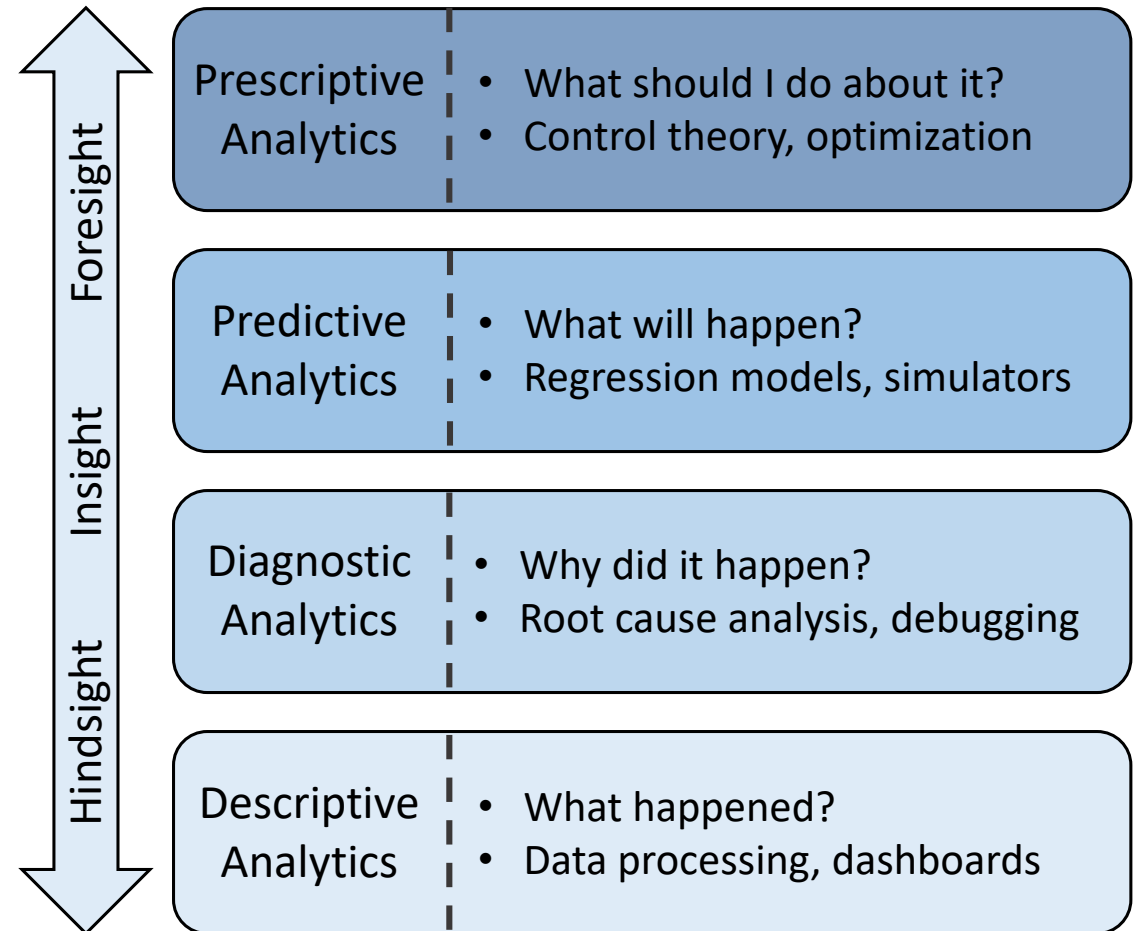
# The 4-Pillar Framework



[1] T. Wilde et al. "The 4-Pillar Framework for energy efficient HPC data centers".  
EEHPC Working Group | 07.09.2021 | Alessio Netti

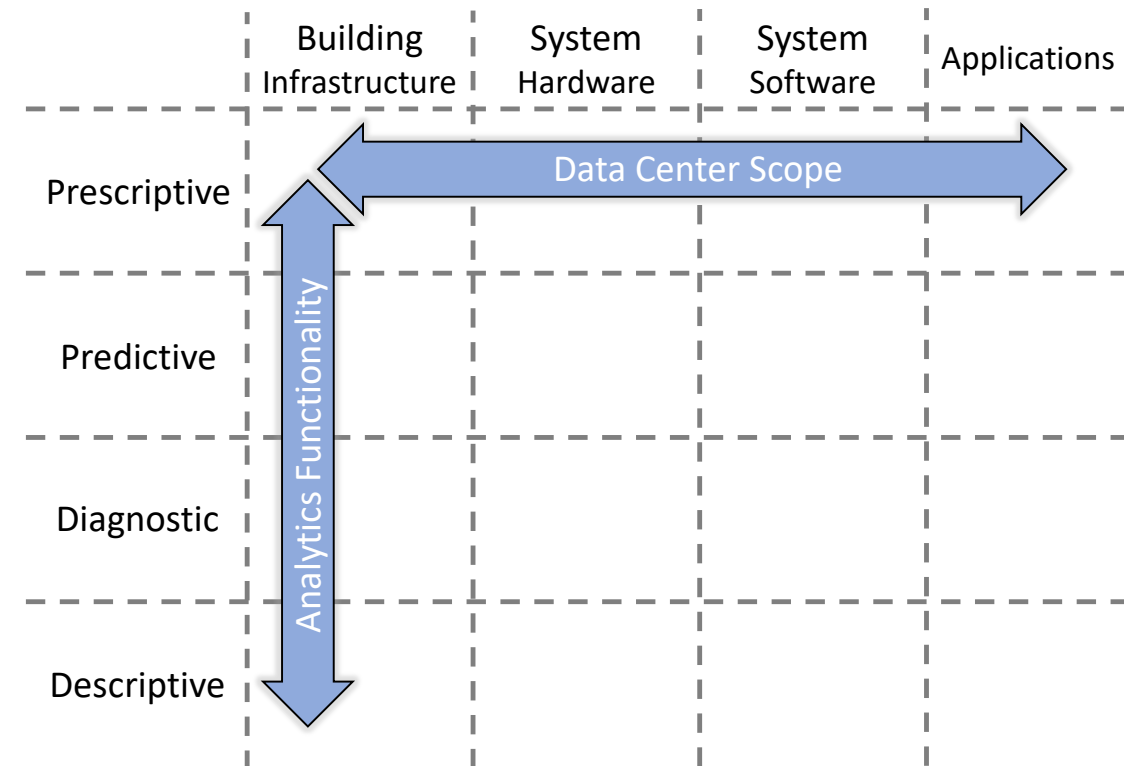
## The 4 Types of Data Analytics

- Model used by large consultancy firms to categorize **data analytics**.
- Consists of 4 *types*, which differ in the **functionality** they offer.
- Some types focus on **historical** events (*hindsight*), others on anticipating **future** ones (*foresight*).
- The types are not necessarily **ranked** by complexity.



## The 4x4 Conceptual ODA Framework

- We **combine** the 4-Pillar and the 4-Type models in a single framework.
- It consists of a 4x4 **matrix**:
  - The *pillars* in the horizontal axis describe the scope of ODA.
  - The *types* in the vertical axis describe ODA functionality.
- Any complex ODA system can be **decomposed** to fit the cells of the framework.



## Applying the ODA Framework

- We now **demonstrate** the framework's effectiveness.
- We conducted a **survey** of ODA research literature:
  - 70+ works analyzed and categorized.
  - ODA examples extracted for each category.
  - Provides an overall picture of the ODA field.
- We **applied** the framework to three state-of-the-art use cases:
  - Focus on complex ODA systems.
  - Discussion of the framework's limitations.



# Classifying ODA Research Literature

	Building Infrastructure	System Hardware	System Software	Applications
Prescriptive	<ul style="list-style-type: none"> <li>Switching between types of cooling</li> <li>Tuning cooling knobs</li> <li>Responding to anomalies</li> </ul>	<ul style="list-style-type: none"> <li>Cooling optimization at the system level</li> <li>CPU frequency tuning</li> <li>Tuning hardware knobs</li> </ul>	<ul style="list-style-type: none"> <li>Intelligent task placement</li> <li>Plan-based scheduling</li> <li>Power and KPI-aware scheduling</li> </ul>	<ul style="list-style-type: none"> <li>Auto-tuning of HPC applications</li> <li>Code improvement and recommendations</li> </ul>
Predictive	<ul style="list-style-type: none"> <li>Predict data center KPIs</li> <li>Predict cooling demand</li> <li>Models for cooling performance</li> </ul>	<ul style="list-style-type: none"> <li>Forecast sensors</li> <li>Component failure prediction</li> <li>Predict instruction mix</li> </ul>	<ul style="list-style-type: none"> <li>Simulating HPC systems and schedulers</li> <li>Predicting HPC workloads</li> </ul>	<ul style="list-style-type: none"> <li>Predicting job durations</li> <li>Predicting resource usage</li> <li>Predicting performance profiles of code regions</li> </ul>
Diagnostic	<ul style="list-style-type: none"> <li>Fingerprinting data center-level crises</li> <li>Infrastructure anomaly detection</li> <li>Stress testing</li> </ul>	<ul style="list-style-type: none"> <li>Node anomaly detection</li> <li>Root cause analysis at the system level</li> <li>Diagnose network contention issues</li> </ul>	<ul style="list-style-type: none"> <li>Detect data locality issues</li> <li>Detect software anomalies</li> <li>Diagnose OS noise</li> </ul>	<ul style="list-style-type: none"> <li>Application fingerprinting</li> <li>Identify application performance patterns</li> <li>Diagnose code-level issues</li> </ul>
Descriptive	<ul style="list-style-type: none"> <li>PUE calculation</li> <li>Facility data processing</li> <li>Facility-level dashboards</li> </ul>	<ul style="list-style-type: none"> <li>ITUE calculation</li> <li>System performance indicators</li> <li>System-level dashboards</li> </ul>	<ul style="list-style-type: none"> <li>Slowdown calculation</li> <li>Scheduler-level dashboards</li> </ul>	<ul style="list-style-type: none"> <li>Job performance models</li> <li>Job data processing</li> <li>Job-level dashboards</li> </ul>

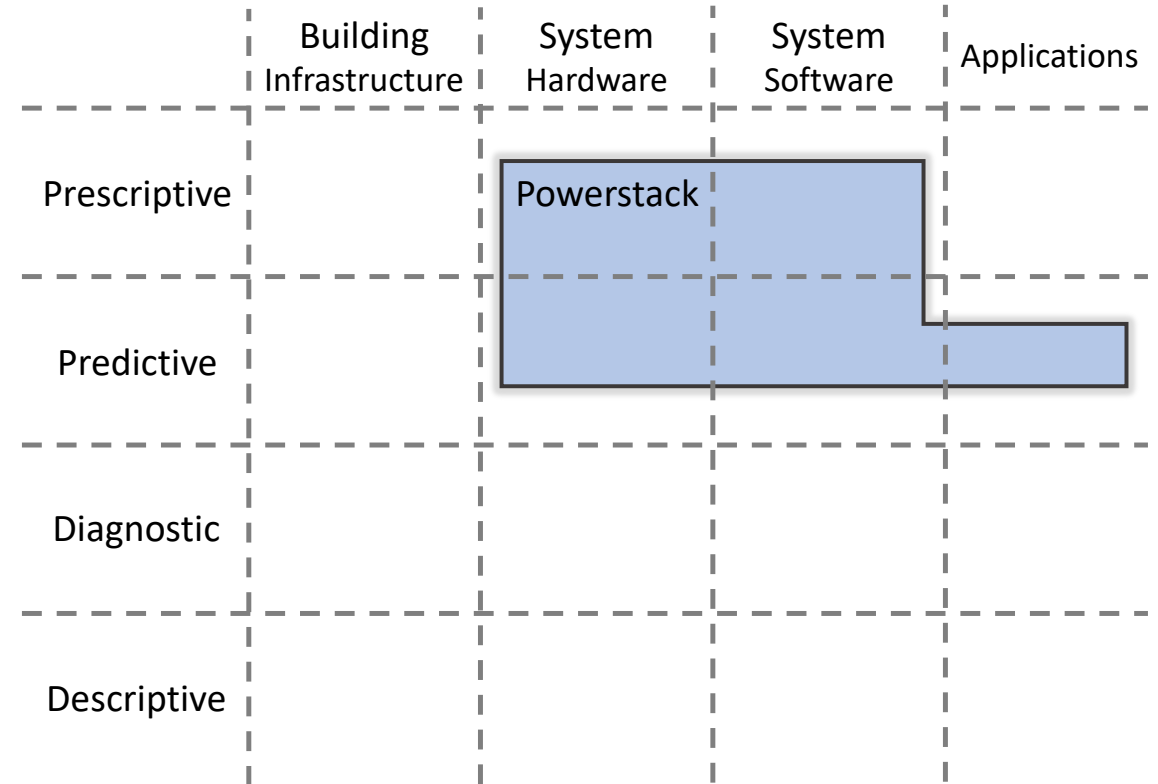
## Analytics across Multiple Types

- Infrastructure anomaly detection (*diagnostic*) and cooling set-point tuning (*prescriptive*) at **ENI** [3].
- **Better prescriptive** decisions can be made with the help of *predictive* and *diagnostic* components.
- Higher technical **complexity**.
- Requires **fusion** of heterogeneous disciplines.

	Building Infrastructure	System Hardware	System Software	Applications
Prescriptive	ENI			
Predictive				
Diagnostic	ENI			
Descriptive				

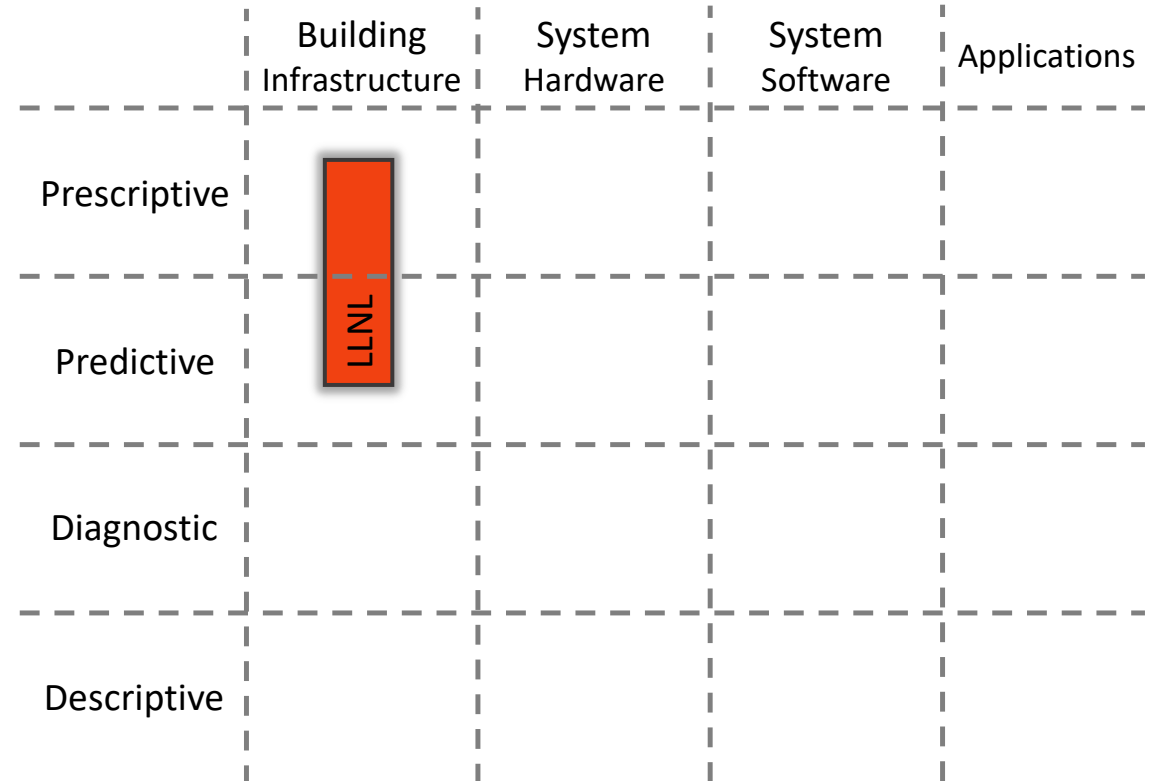
## Implementing Multi-pillar ODA

- The **Powerstack** framework for power management (*prescriptive*) using data science (*predictive*) [4].
- Most ODA systems are **closed** and cover a single pillar (or *silo*).
- Multi-pillar designs must be **holistic** and integrate many levels of scope.
- Major operational **opportunities**.



## ODA beyond Building Infrastructure

- Forecasting (*predictive*) and notifying (*prescriptive*) excessive power swings at **LLNL** [5].
- The electrical grid as an **extension** of the data center's infrastructure.
- Monitoring and control capabilities are **limited**.
- Practical implementation can be **challenging**.



## Conclusions

- Use of *Operational Data Analytics* (ODA) is becoming more and more **common** in HPC data centers.
- We propose a conceptual **framework** to classify ODA use cases according to their scope (*pillars*) and functionality (*types*).
- We aim to establish a common language to **simplify** discussion, analysis and adoption of ODA by the community at large.
- Thank you for your attention!