Towards a Hybrid Cloud & Edge Orchestrator for Mining Exascale Distributed Streams

Herodotos Herodotou

Cyprus University of Technology

herodotos.herodotou@cut.ac.cy

Acknowledgments

Nicolas Kourtellis

Telefonica Research

Maciej Grzenda

Politechnika Warszawska — Warsaw University of Technology

Piotr Wawrzyniak

Lodz University of Technology

Albert Bifet

The University of Waikato — Te Whare Wānanga o Waikato

Telecom Paris — IP Paris

Data Intensive Computing Research Lab

Short Bio: Herodotos Herodotou

Duke University
MSc/PhD Computer Science

Microsoft Research

Cyprus University of Technology
Assistant Professor

2012 2014 2022

Microsoft

aster data big data. fast insights.

Yahoo! Research

ACM SIGMOD

Jim Gray Doctoral Dissertation Award Honorable Mention

Data Intensive Computing Research Lab

Research Areas

- Large-scale data processing systems (e.g., MapReduce, Spark)
- Centralized and distributed database systems
- Cloud computing (compute, storage, and networking)
- Data-driven applications (maritime, tourism, social computing)

Research Team

- Supervise: 1 postdoc & 3 PhD students
- Co-supervise: 2 postdocs & 1 PhD student

Equipment

- 15-node local private cluster

Our Data-driven World

Massive Data

Entertainment
- Internet images
- Movie & music clips

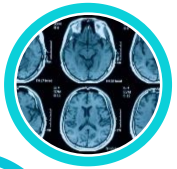

Medicine
- MRI & CT scans
- Patient records

Science
- Genomics
- Astronomy
- Weather data

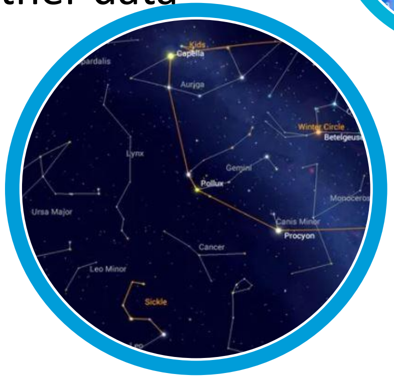

Business
- Product sales
- Stock market
- Customer data

Humanities
- Social interactions
- Historical documents

Analysis

Insights

Ad placement

Scientific breakthroughs

Business process efficiencies

Personalized recommendations

Improved healthcare

Fraud detection

Evolution of Big Data Analytics Systems

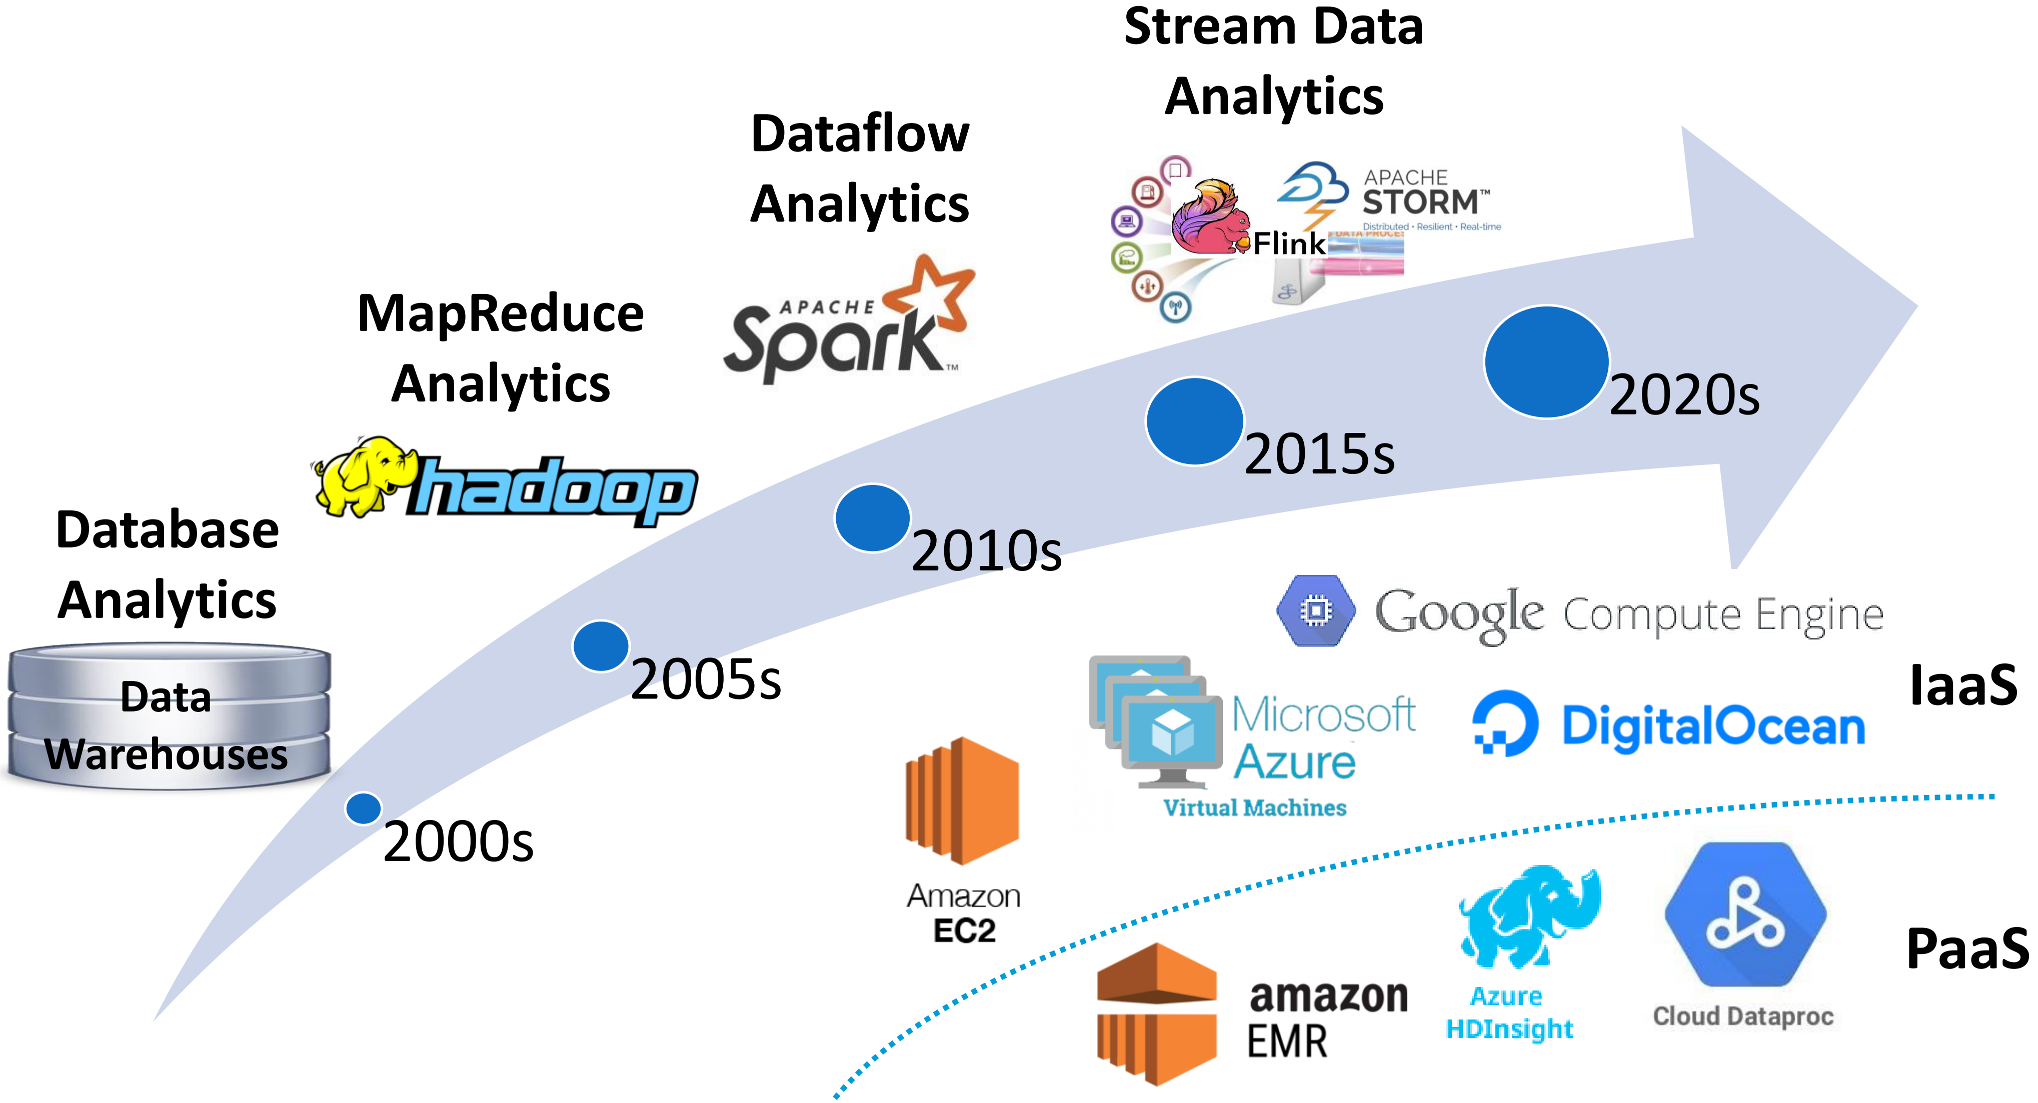

Current Landscape

- 10s of billions of devices
- 4/5 Vs: Volume, Velocity, Variety, Veracity
- Need for data modeling to extract Value (5th V)

However

- Cloud heterogeneity & management overhead
- Vendor lock-in
- 4Vs already stress current, non-scalable infra

S2CE: Stream AI + Cloud + Edge

Motivational Scenario 1: Telco Cloud

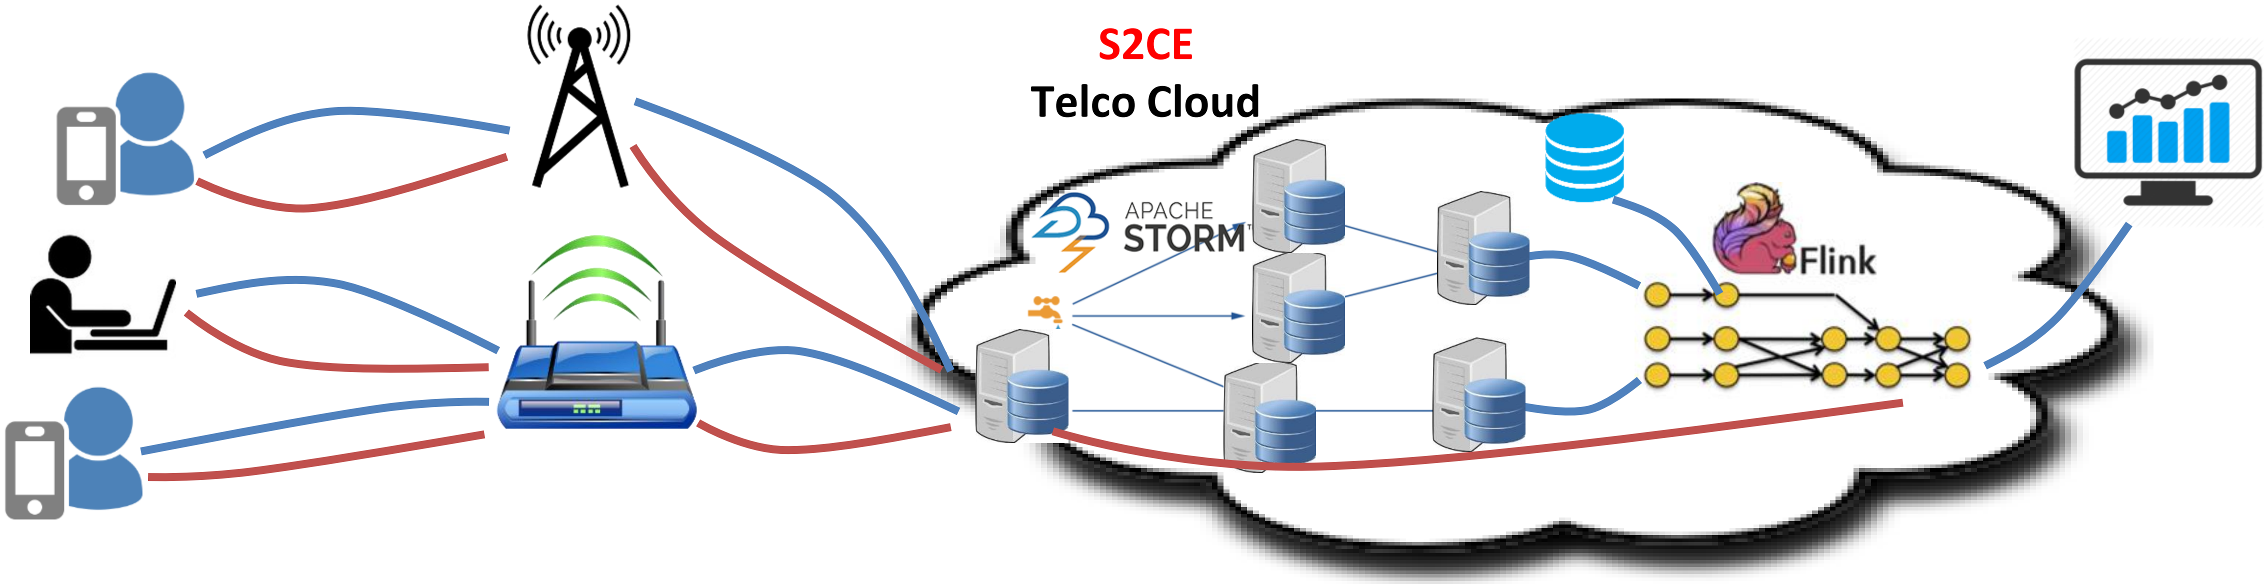

Upstream	
Downstream	
S2CE edge node	O

Motivational Scenario 2: Smart Bus & City

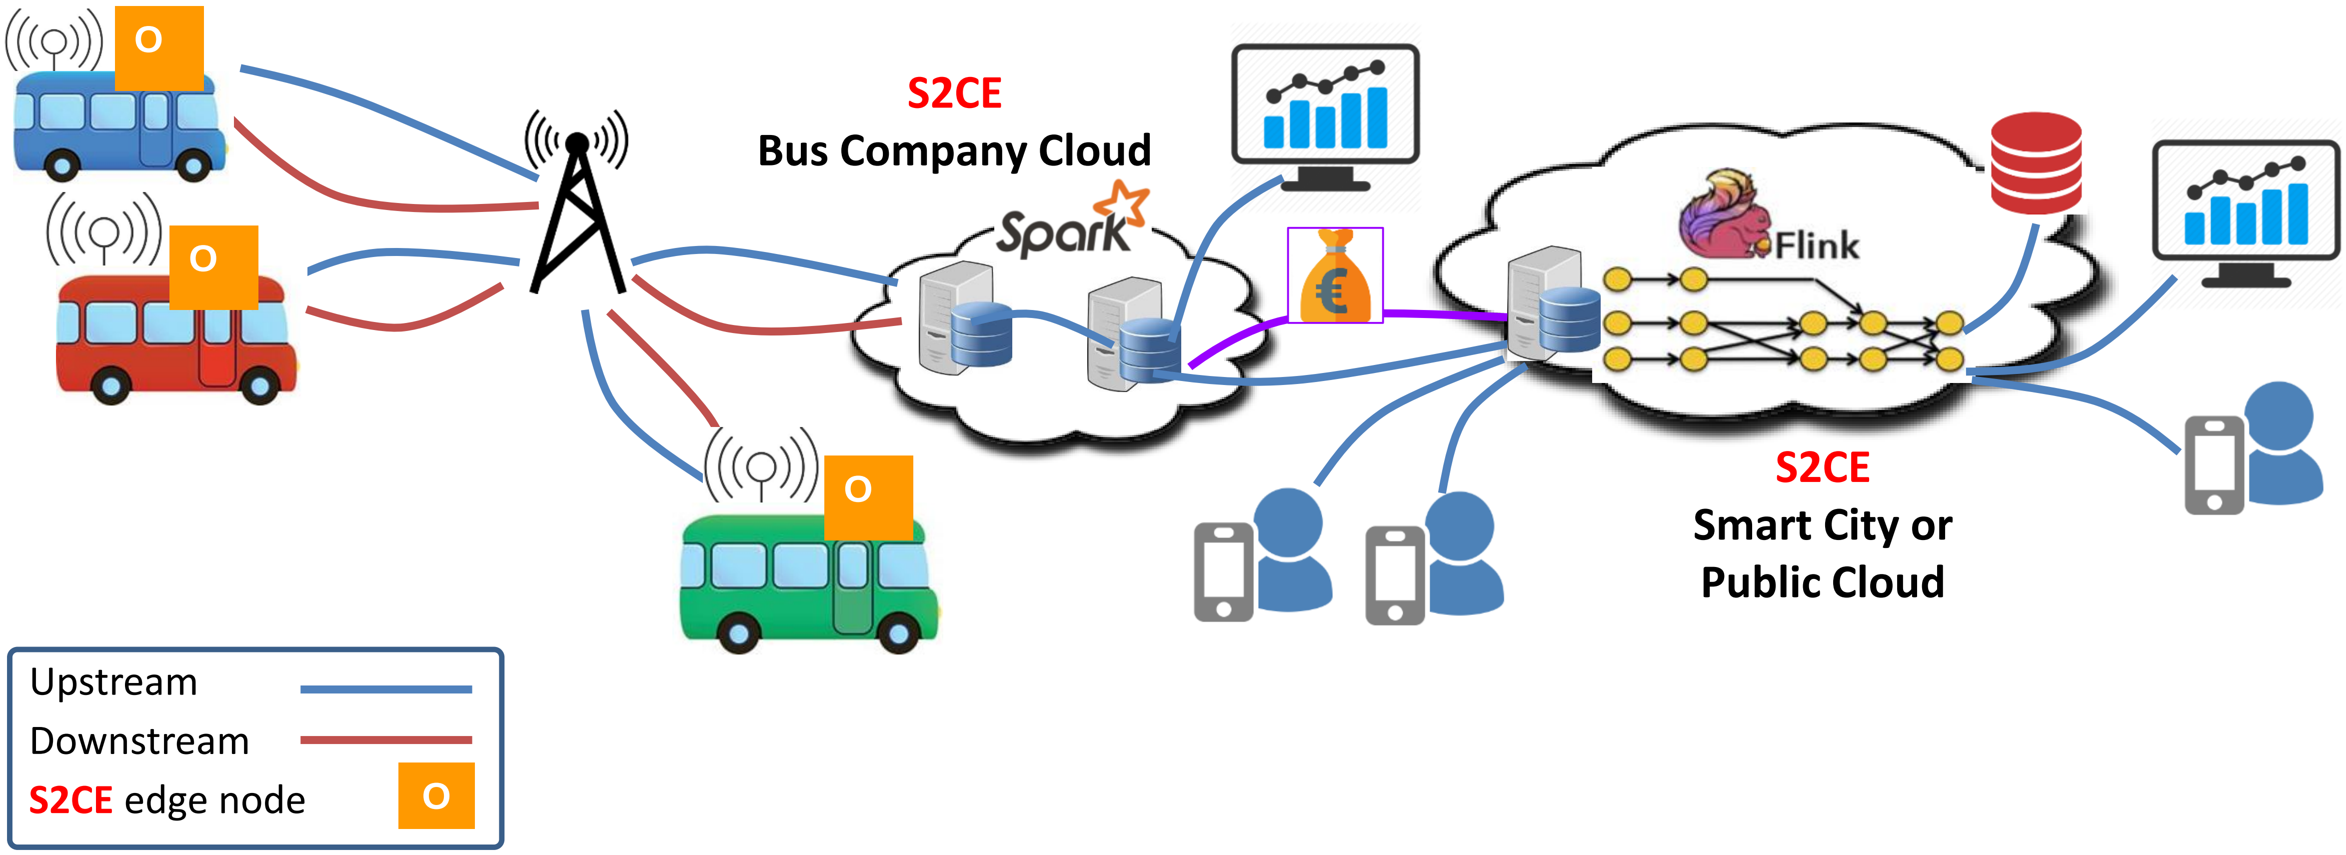

State of the Art

- Big data stream processing systems
- Cloud resource management and tuning
- Distributed stream processing at the edge
- Machine and deep learning over data streams
- Data transformation techniques

Big Data Stream Processing Systems

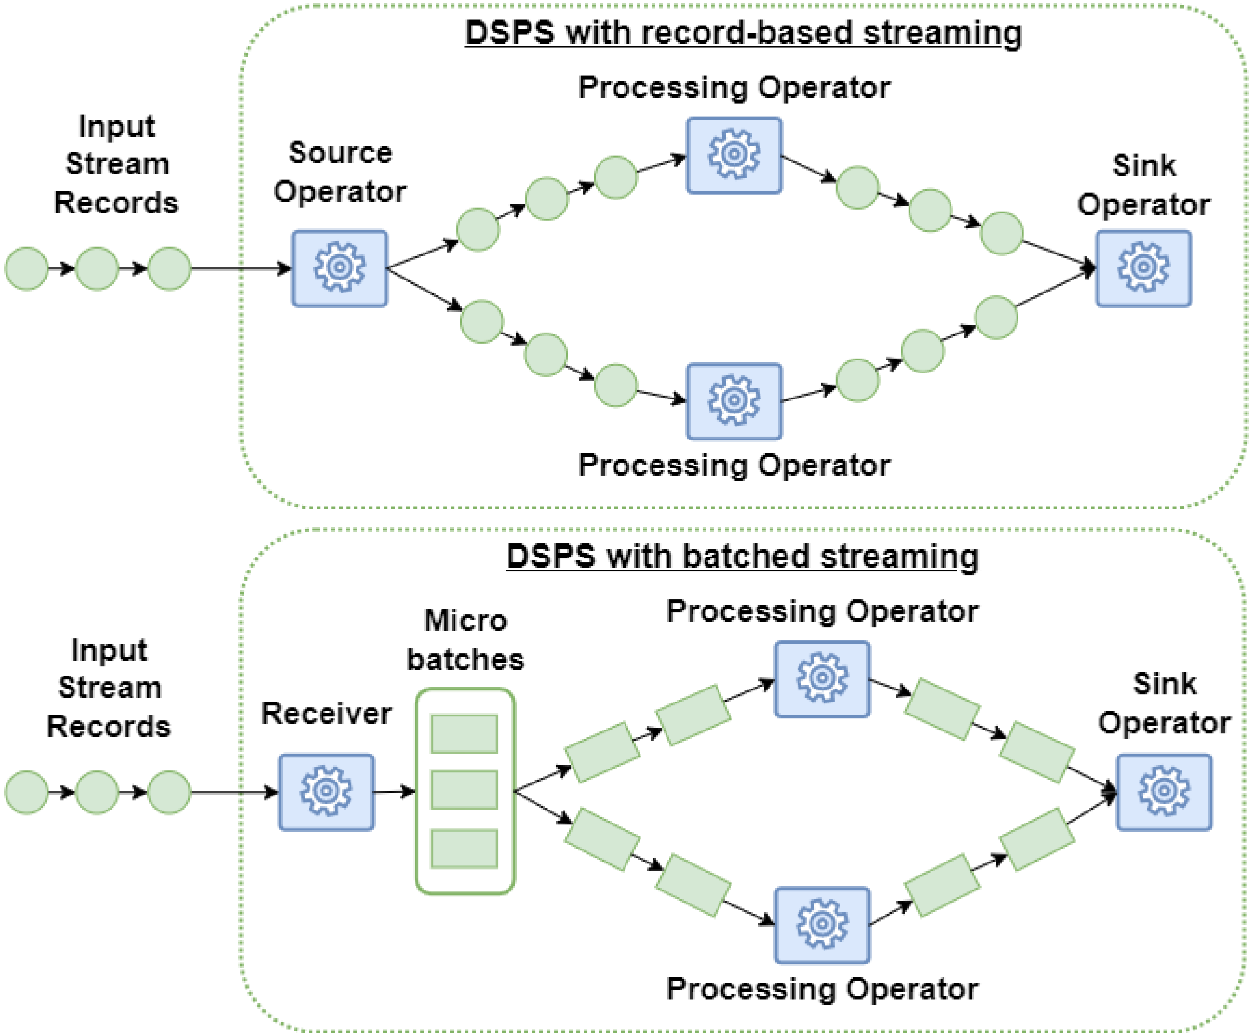

Flink ML

Big Data Stream Processing Systems – Cloud

Google Cloud Dataflow

Azure

amazon Kinesis

Cloud Resource Management and Tuning

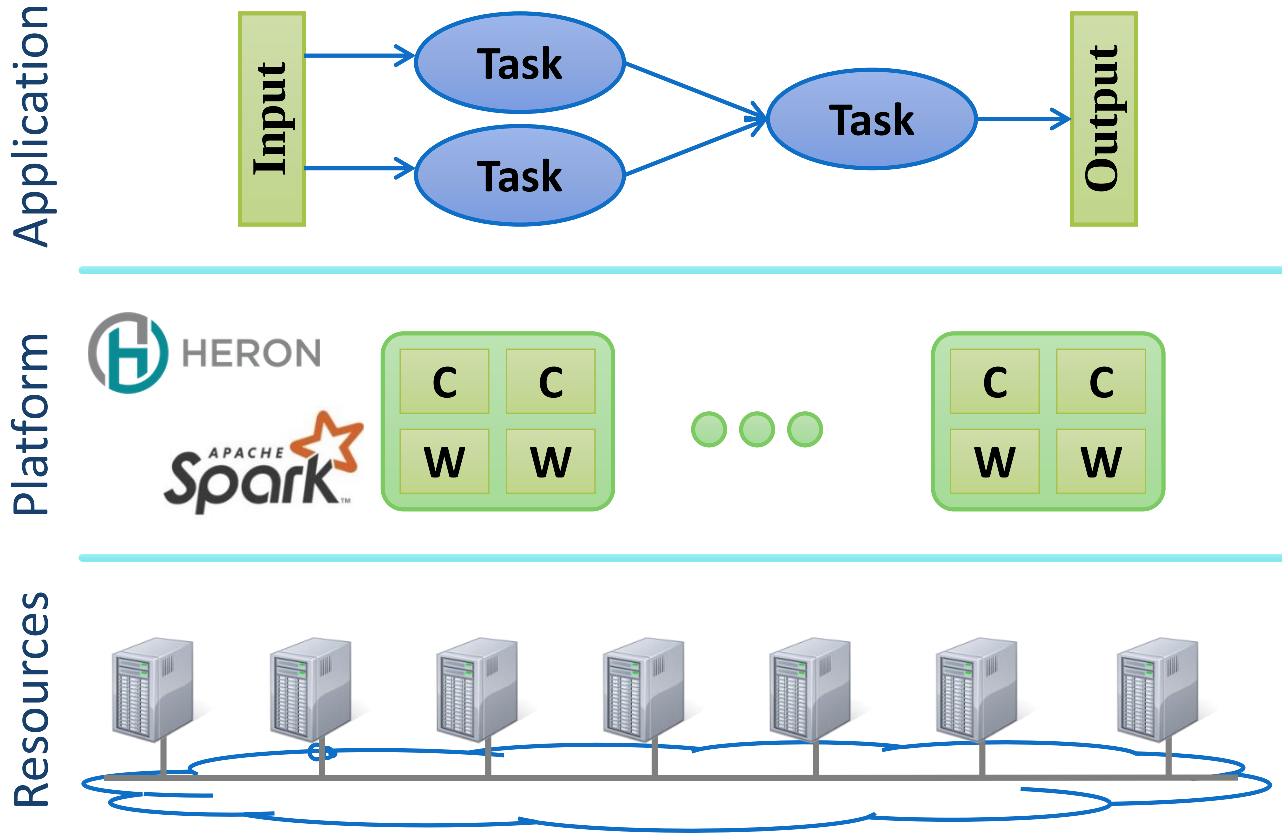

Decisions:

- Task parallelism
- Micro-batch size
- …

- Number of cores
- Memory settings
- …

- Number of nodes
- VM/Container specs
- …

Cloud Resource Management and Tuning – Approaches

Cost Modeling	Use cost models & statistics to find optimal settings
Simulation-based	Use simulator to estimate application performance
Experiment-driven	Execute application with different settings iteratively
Machine Learning	Use machine learning to model application performance
Adaptive	Change configurations while application is running

Distributed Stream Processing at the Edge

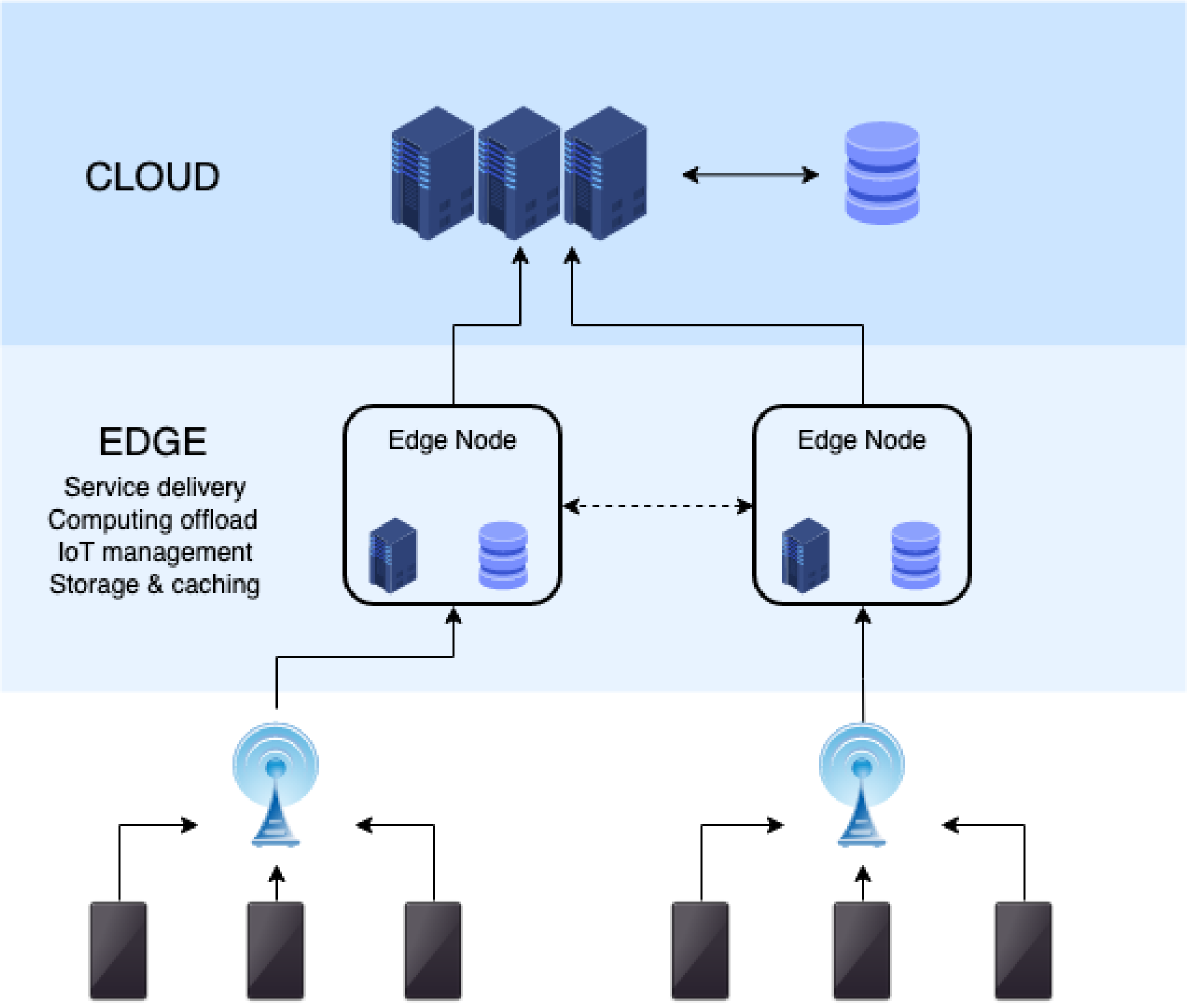

- **Benefits of edge:**
  - Reduce end-to-end latency and communication costs
  - Enable services to react to events locally
  - Offload processing from the cloud
- **Challenges:**
  - Computing, storage, and network resources are constrained
  - Deployment of data stream processing applications onto heterogeneous infrastructure has been proven to be NP-hard

Distributed Stream Processing at the Edge – Projects

Aspect	EdgeX Foundry	Azure IoT Edge	Apache Edgent	CORD	Akraino Edge Stack
Interface	Restful API	Web service	API	API or XOS-GUI	N/A
OS support	Various	Various	Various	Ubuntu	Linux
Programming framework	Not provided	Java, .NET, C, Python, etc.	Java	Shell, Python	N/A
Applications	IoT	Unrestricted	IoT	Unrestricted	Unrestricted
Deployment	Dynamic	Dynamic	Static	Dynamic	Dynamic
Target user	General users	General users	General users	Network ops	Network ops
Virtualization	Container	Container	JVM	VM & Container	VM & Container
Limitation	Lack of programming interface	Azure services chargeable	Limited to data analysis	Unable to be offline	Unable to be offline
Scalability	Scalable	Scalable	Not scalable	Scalable	Scalable
Mobility	Not supported	Not supported	Not supported	Support	Support

Liu et al. A survey on edge computing systems and tools. IEEE

Machine and Deep Learning over Data Streams

Existing libraries/systems:

- Massive Online Analysis (MOA)
  - algorithms for streaming classification, clustering, and change detection
- Vowpal Wabbit
  - based on the perceptron algorithm with a focus on reinforcement learning
- Jubatus
  - tight coupling between the ML library and the underlying custom-built DSPS
- Apache SAMOA
  - distributed computation of several ML algorithms over four DSPSs

Machine and Deep Learning over Data Streams

Challenges:

- Data Availability
  - a train, test and predict approach is not applicable for stream data; models are susceptible to changes
- Real-Time Streaming
  - need to reduce time to train models dramatically
- Concept Drift
  - models must adapt to patterns evolving over time by detecting changes quickly
- IID Random Variables
  - statistically independent variables cannot be guaranteed for the overall population

Data Transformation Techniques

- **Data Preprocessing**
  - filtering, format conversion, and multiplexing/demultiplexing
- **Dimensionality reduction**
  - statistical inference methods using hashing projections
  - different subspace tracking methods
- **Stream sampling**
  - allows one-pass algorithms for analysing big data streams
  - uniform or biased sampling along with reduction of the problem space
- **Synthetic data stream generator**
  - infer underlying statistical distributions of the real data
  - do not work for streams with concept drifts
  - protecting privacy and confidentiality is hard

Data Intensive Computing
Research Lab

State of the Art: A summary

✘ means no support
! means partial support
✓ means good/full support

Features/Capabilities	Apache Storm	Apache Samza	Apache Spark	Apache Flink	Apache Apex	Apache Beam	Google CD	MS Azure ML	AWS Kinesis	MOA	Vowpal Wabbit	Jubatus	Apache SAMOA	Desired Platform
Stream integration components	✓	✓	✓	✓	✓	✓	!	!	!	✘	✘	!	✓	✓
Data preprocessing and fusion	!	✘	!	!	✘	✘	!	✘	!	!	!	✓	!	✓
Built-in synthetic data generator	✘	✘	!	!	!	!	✘	✘	✓	✓	✘	✘	!	✓
Stream-based machine learning	!	!	✓	✓	!	!	✓	!	!	✓	✓	✓	✓	✓
Stream-based deep learning	✘	✘	!	✘	!	✘	✘	✘	✘	✘	✘	✘	✘	✓
Resource management	✓	!	✓	✓	✓	✓	✓	!	✓	✘	✘	!	!	✓
Cloud-Edge orchestration	✘	✘	✘	✘	✘	✘	!	!	!	✘	✘	✘	✘	✓
Distributed Platform	✓	✓	✓	✓	✓	✓	✓	✓	✓	✘	✓	✓	✓	✓
Open license (Apache preferred)	✓	✓	✓	✓	✓	✓	✘	✘	✘	!	!	!	✓	✓

Expected Industrial Challenges*

Expected Industrial Challenge	
Heterogeneity Scalability Data-in-motion and data-at-rest	Scalable Data processing
Hybrid (central+edge) big data architectures Decentralization & edge	Edge vs. Cloud infrastructure
Data/AI/predictive/prescriptive analytics Stream analytics frameworks & processing Advanced business analytics	ML/DL-based analytics
Heterogeneity Semantic interoperability Data quality Distributed trust infrastructures	Data fusion and input / output

*BDVA: https://www.bdva.eu

Data Intensive Computing
Research Lab

S2CE Design Objectives

Expected Industrial Challenge	(Objective) How does S2CE address the challenge?
Heterogeneity Scalability Data-in-motion and data-at-rest	(O1) Handling diverse types of cloud computing resources (O1) Distributed and parallelized dynamic analytics for real-time learning (O1) Processing data seamlessly at the same time without extra system overhead
Hybrid (central+edge) big data architectures Decentralization & edge	(O2) Optimizing an efficient mixture of central and edge resources (O2) Computing at edge for faster, more scalable, energy efficient processing
Data/AI/predictive/prescriptive analytics Stream analytics frameworks & processing Advanced business analytics	(O3) Using distributed deep and machine learning (O3) Minimal development effort, scalability, processing speed (O3) Intelligence to empower companies for accurate, instant, data-driven decisions
Heterogeneity Semantic interoperability Data quality Distributed trust infrastructures	(O4) Handling diverse data, modeling, and input/output interfaces (O4) Facilitating data and model exchange between vertical data silos (O4) Providing curation methods for data filtering, quality assessment, improvement (O4) Managing data in anonymized and decentralized fashion

S2CE Architecture Overview

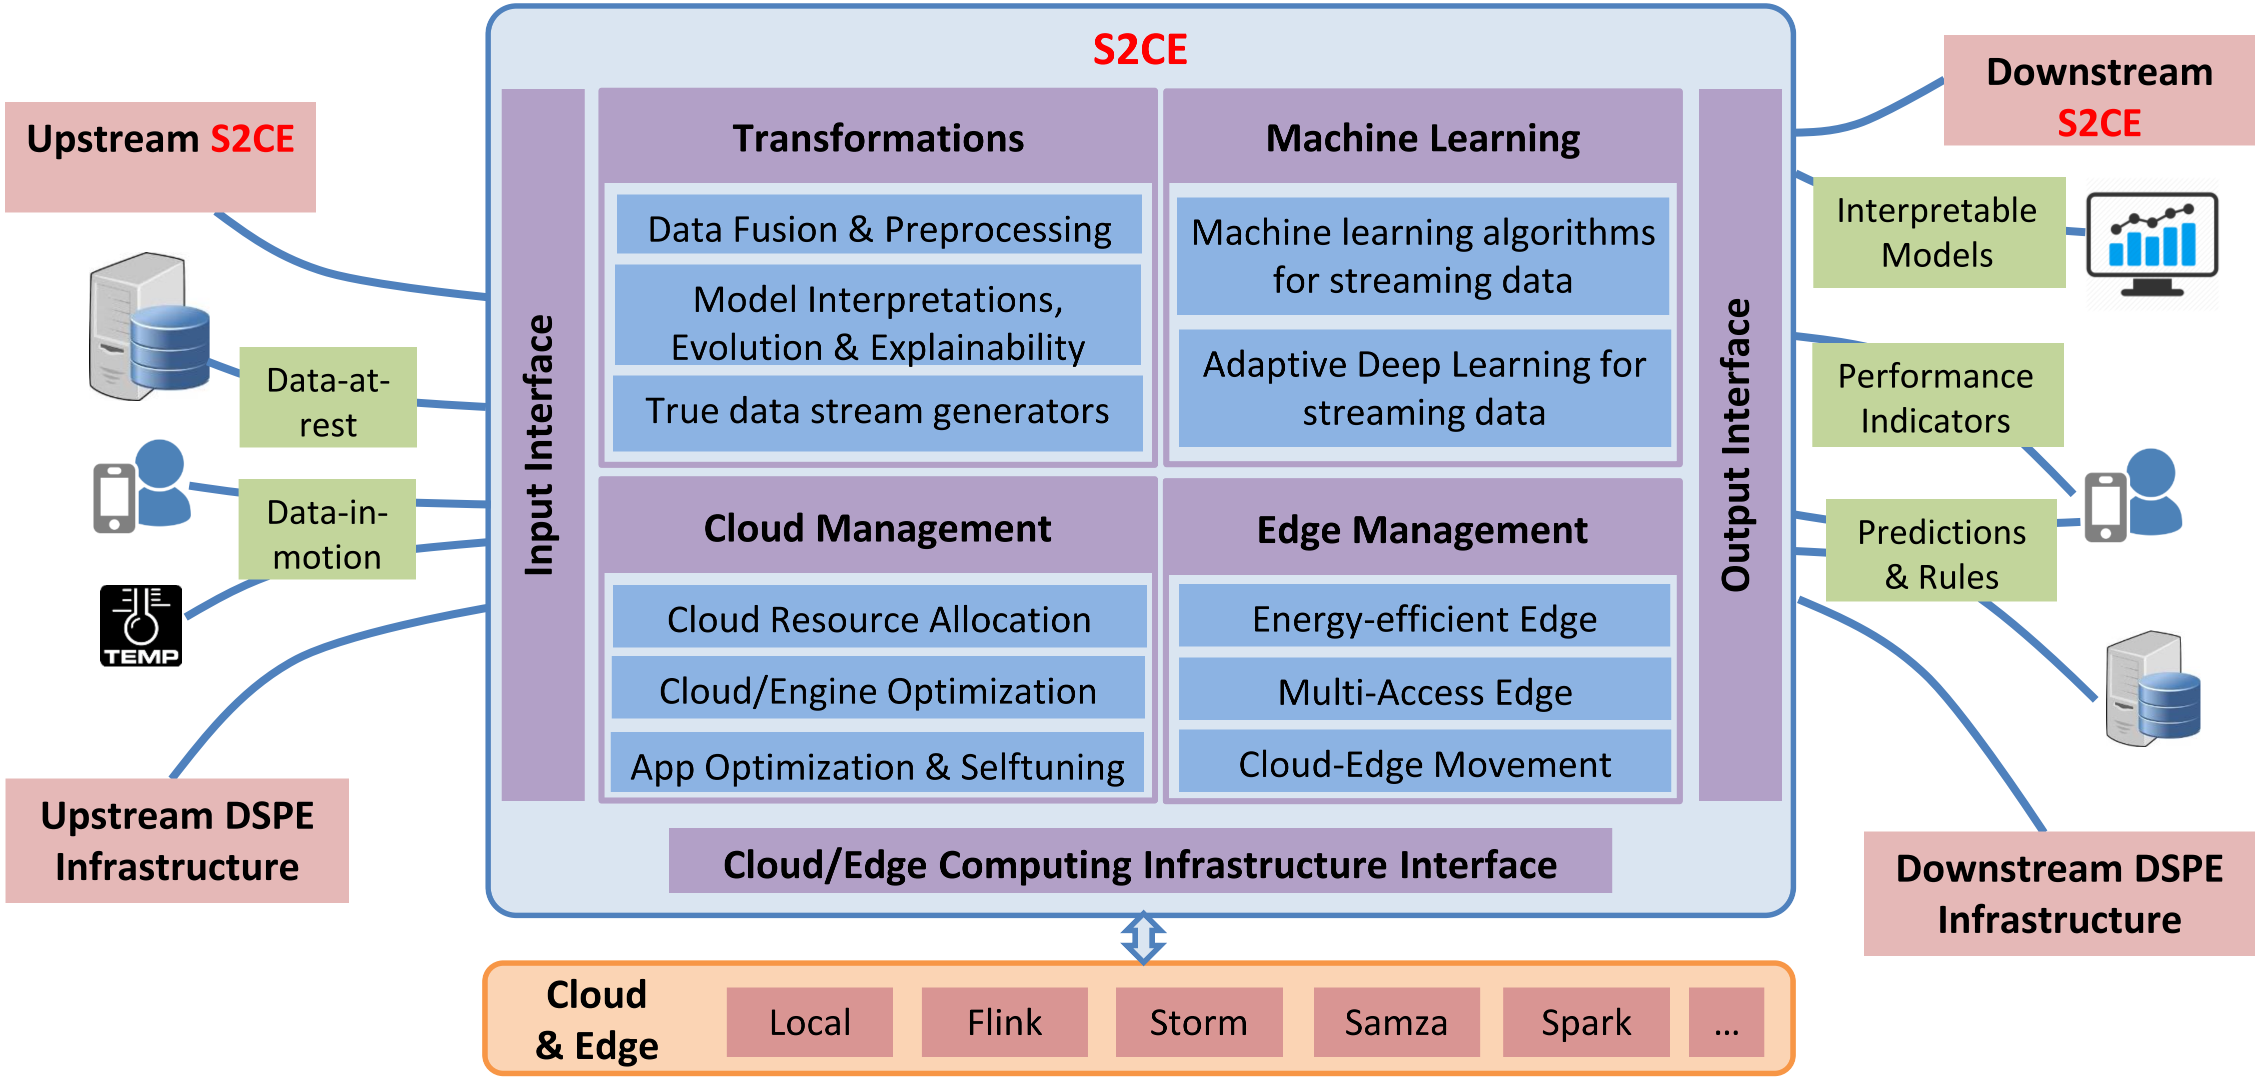

S2CE: Input Interface

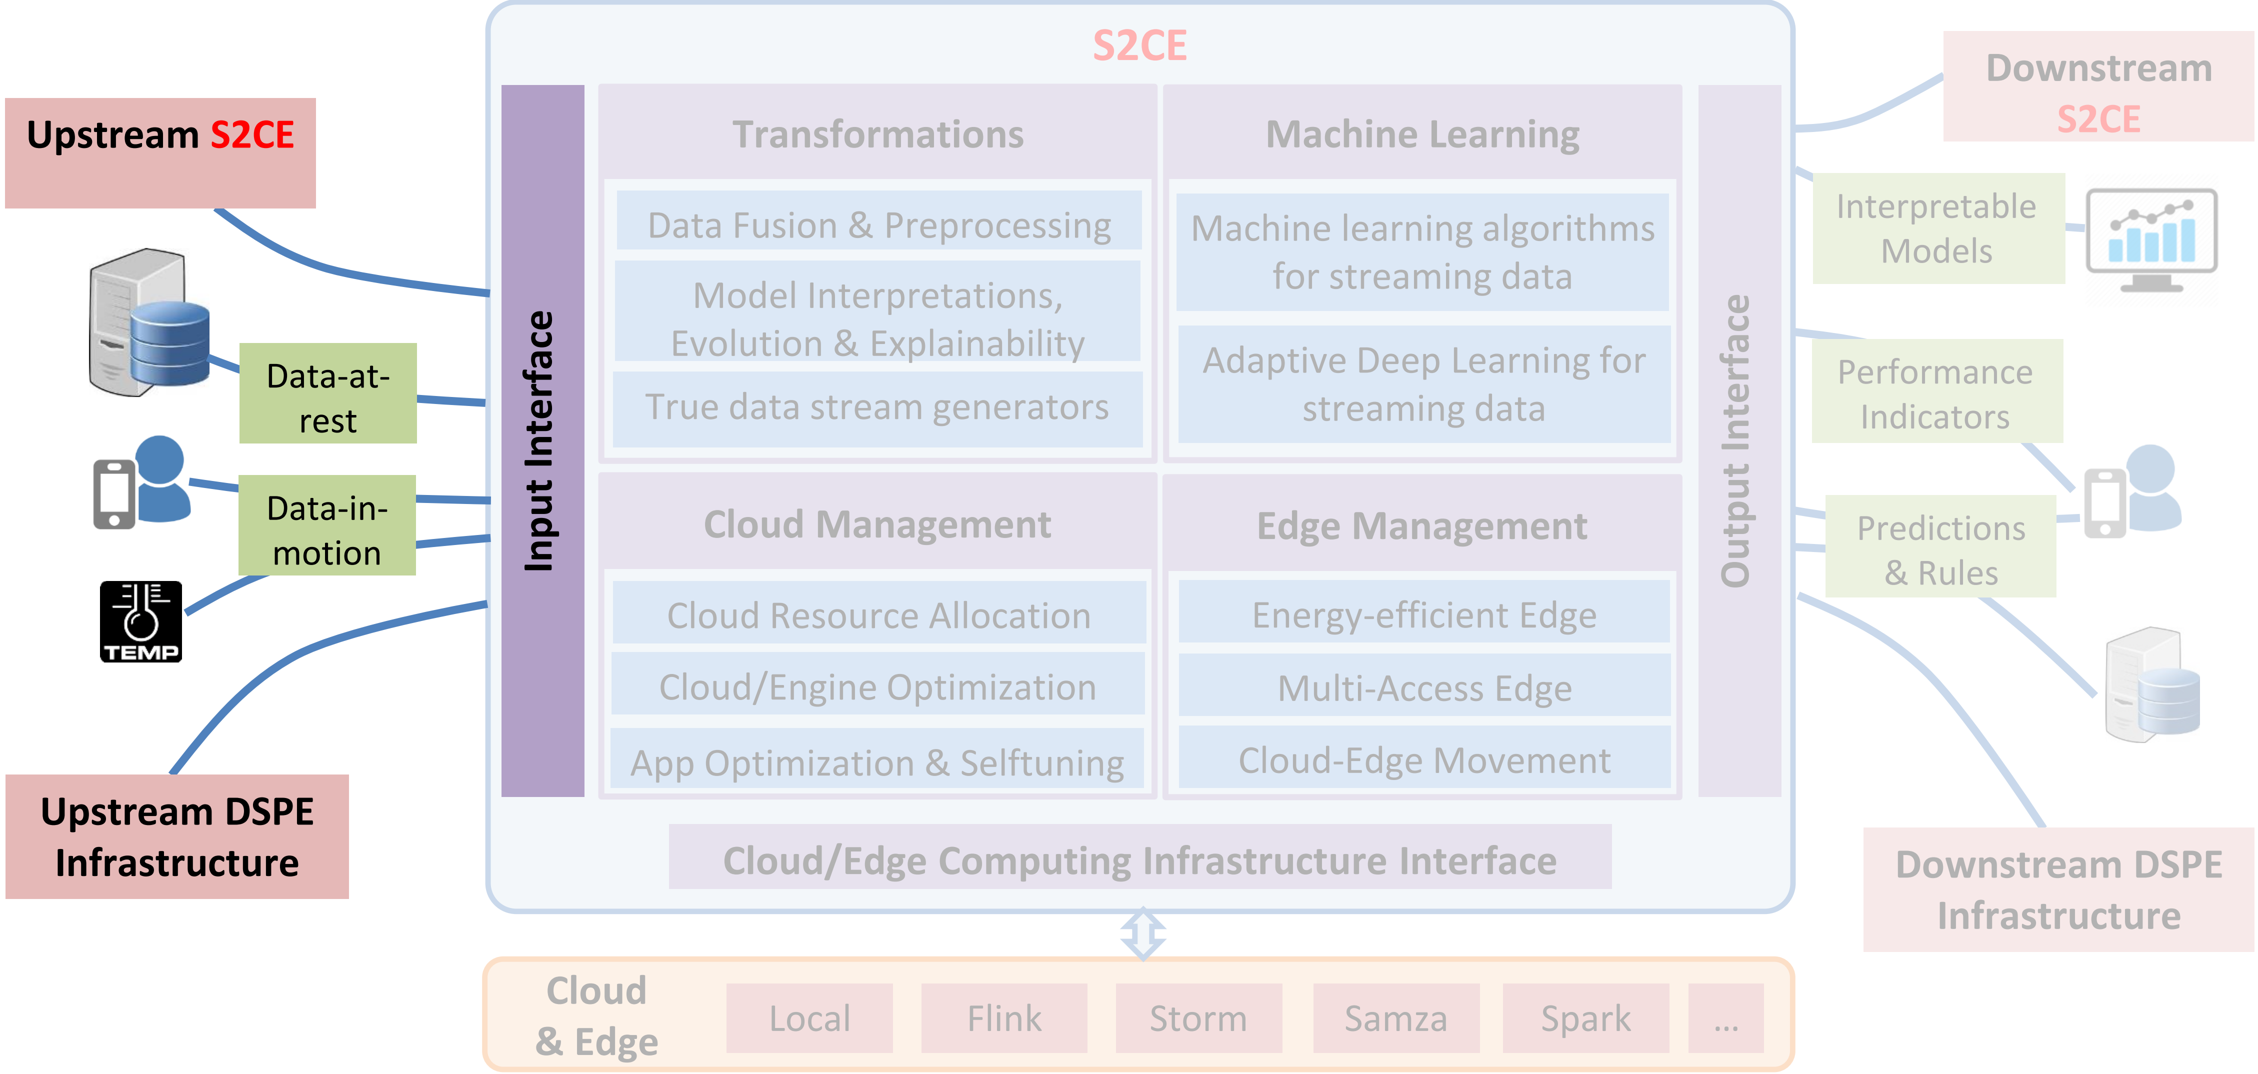

S2CE: Transformations

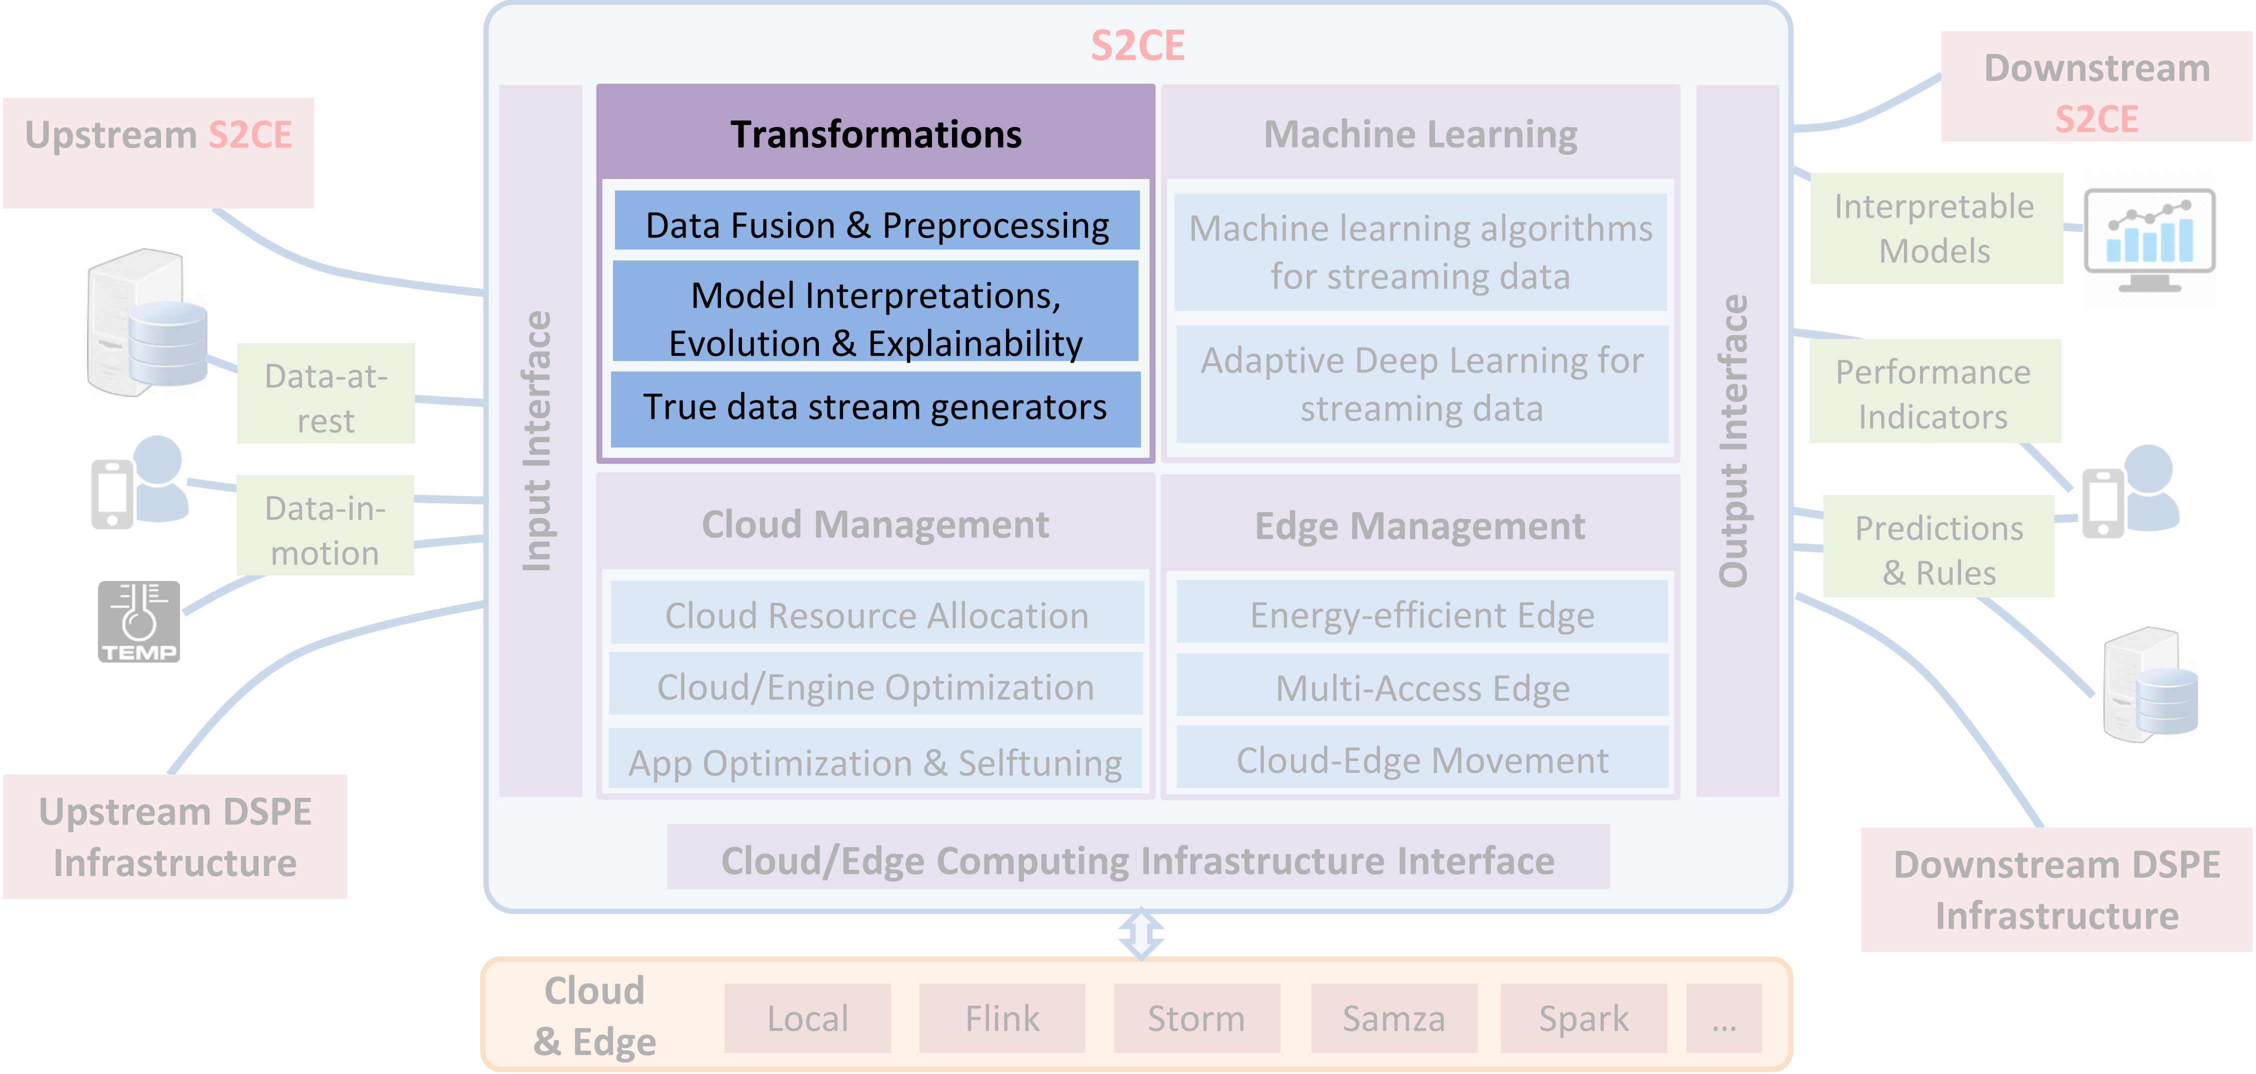

S2CE: Machine Learning

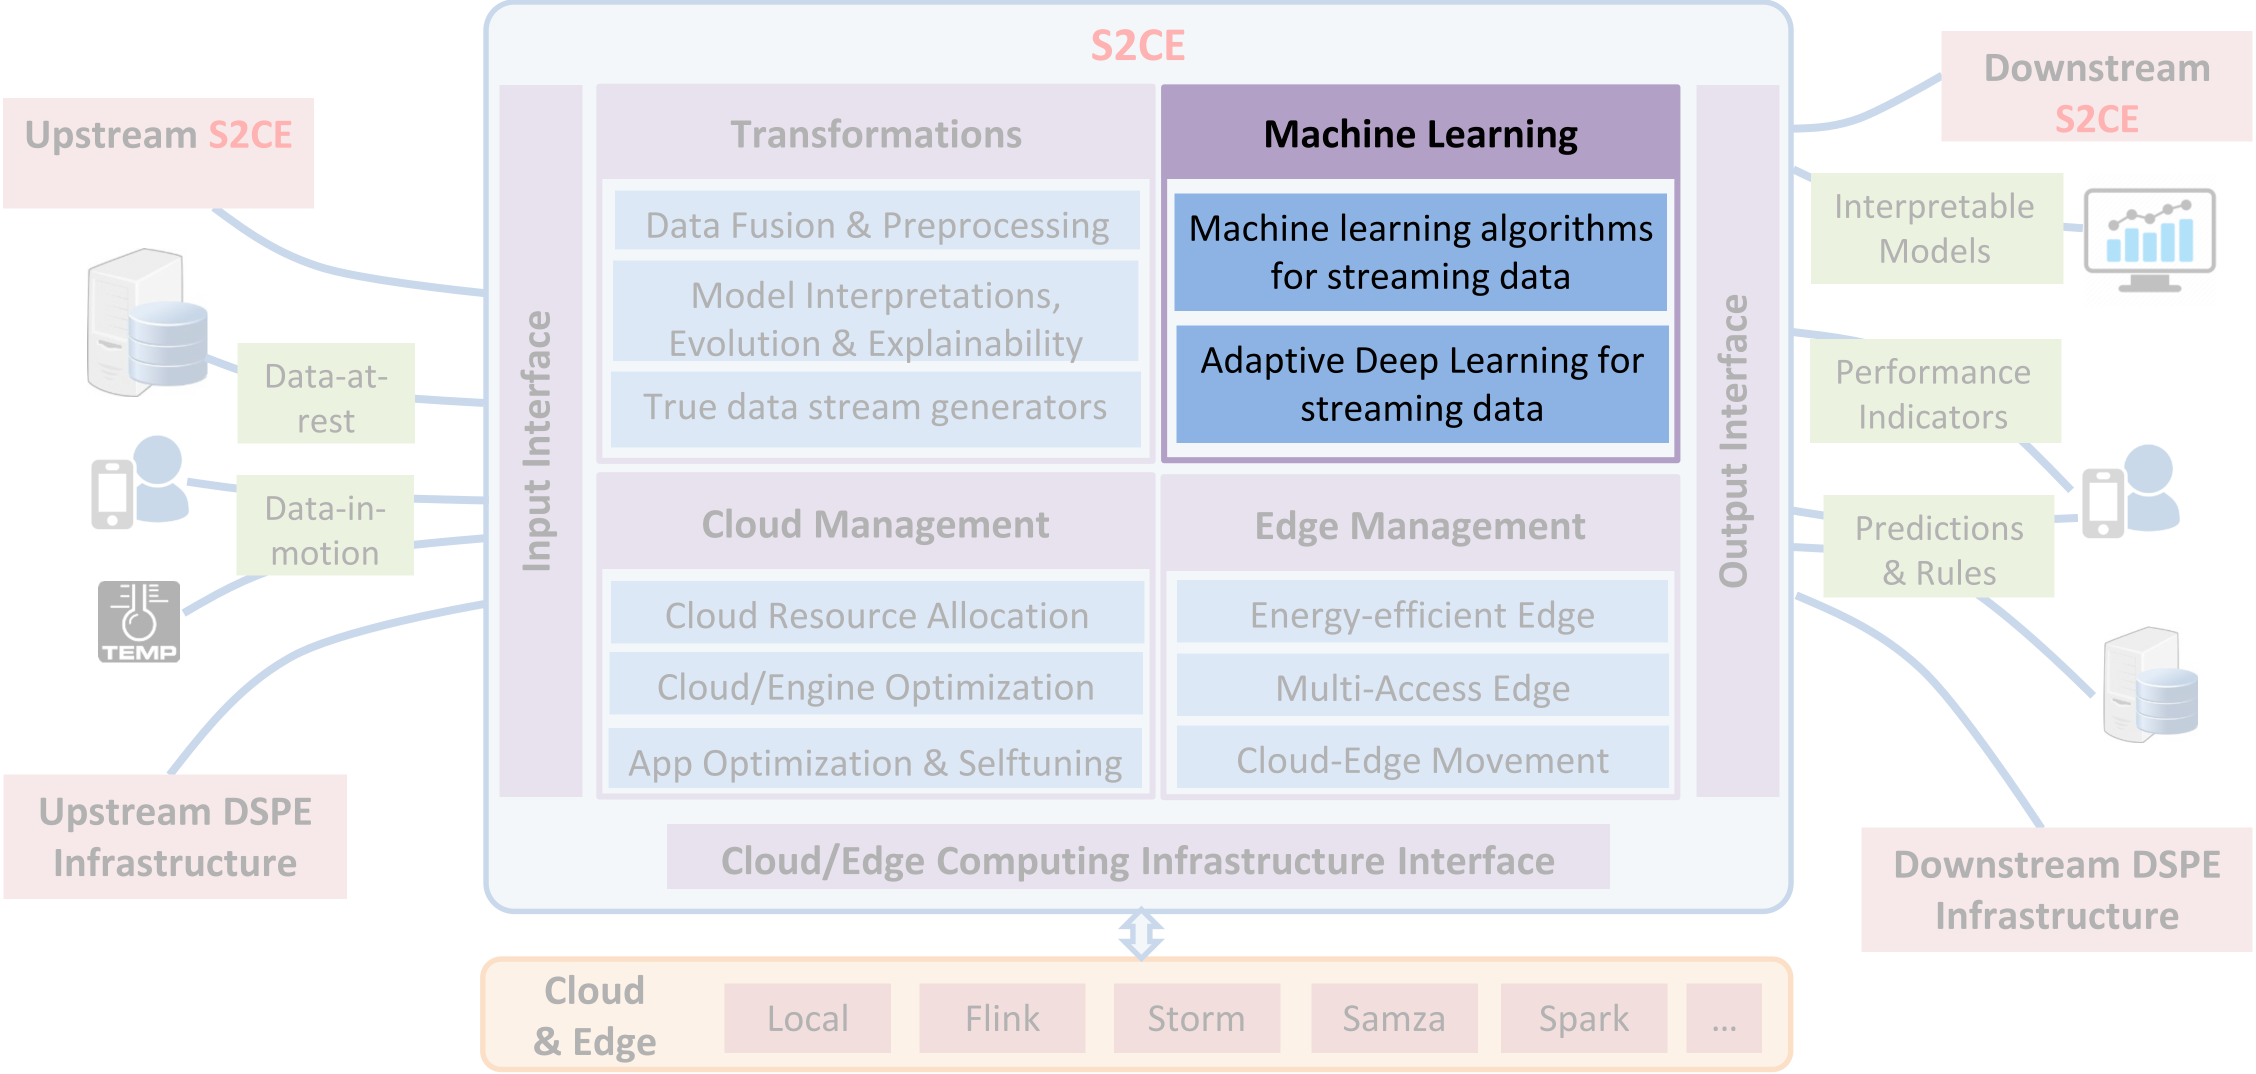

S2CE: Cloud Resource Management

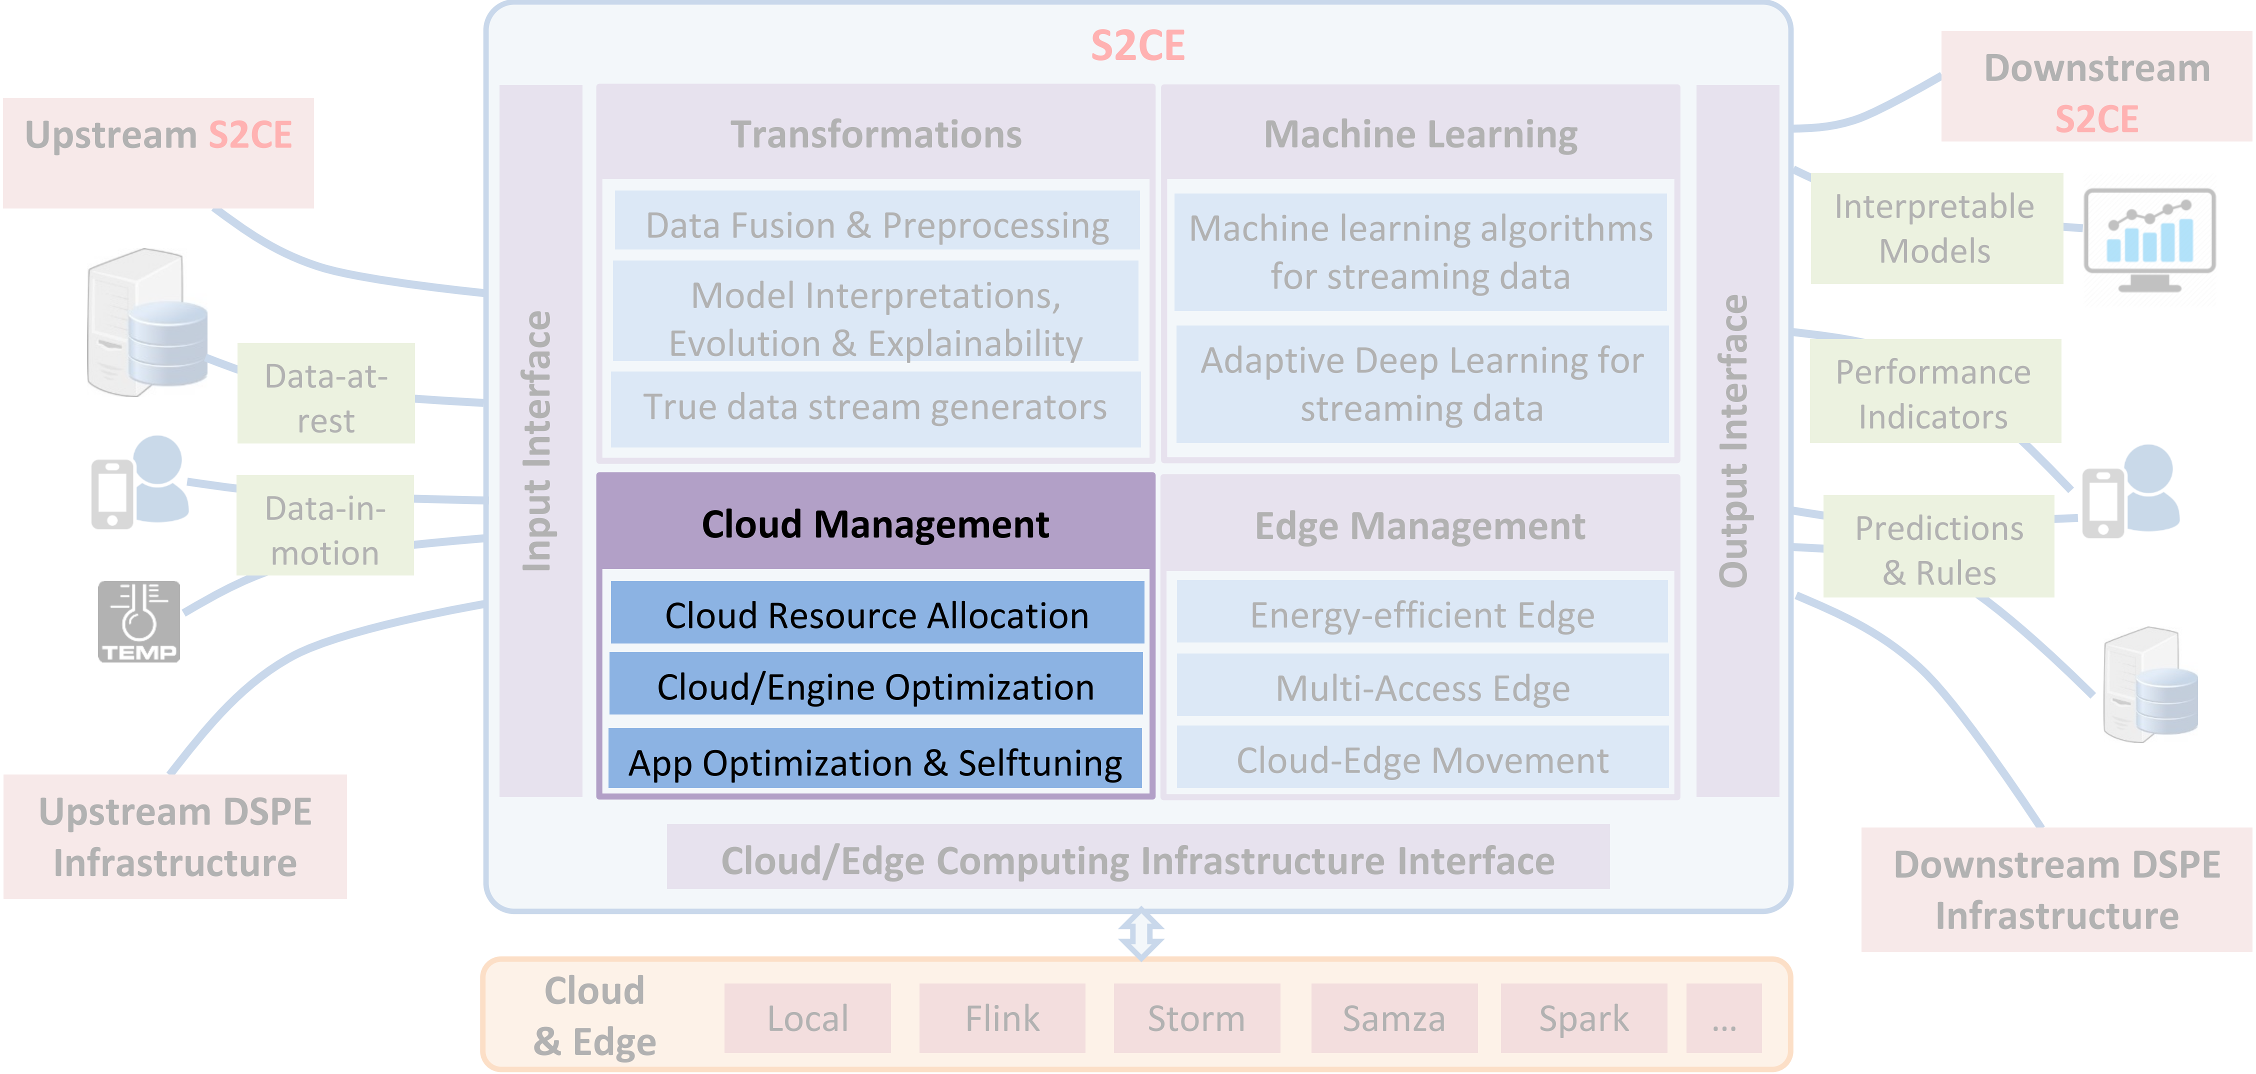

S2CE: Edge Resource Management

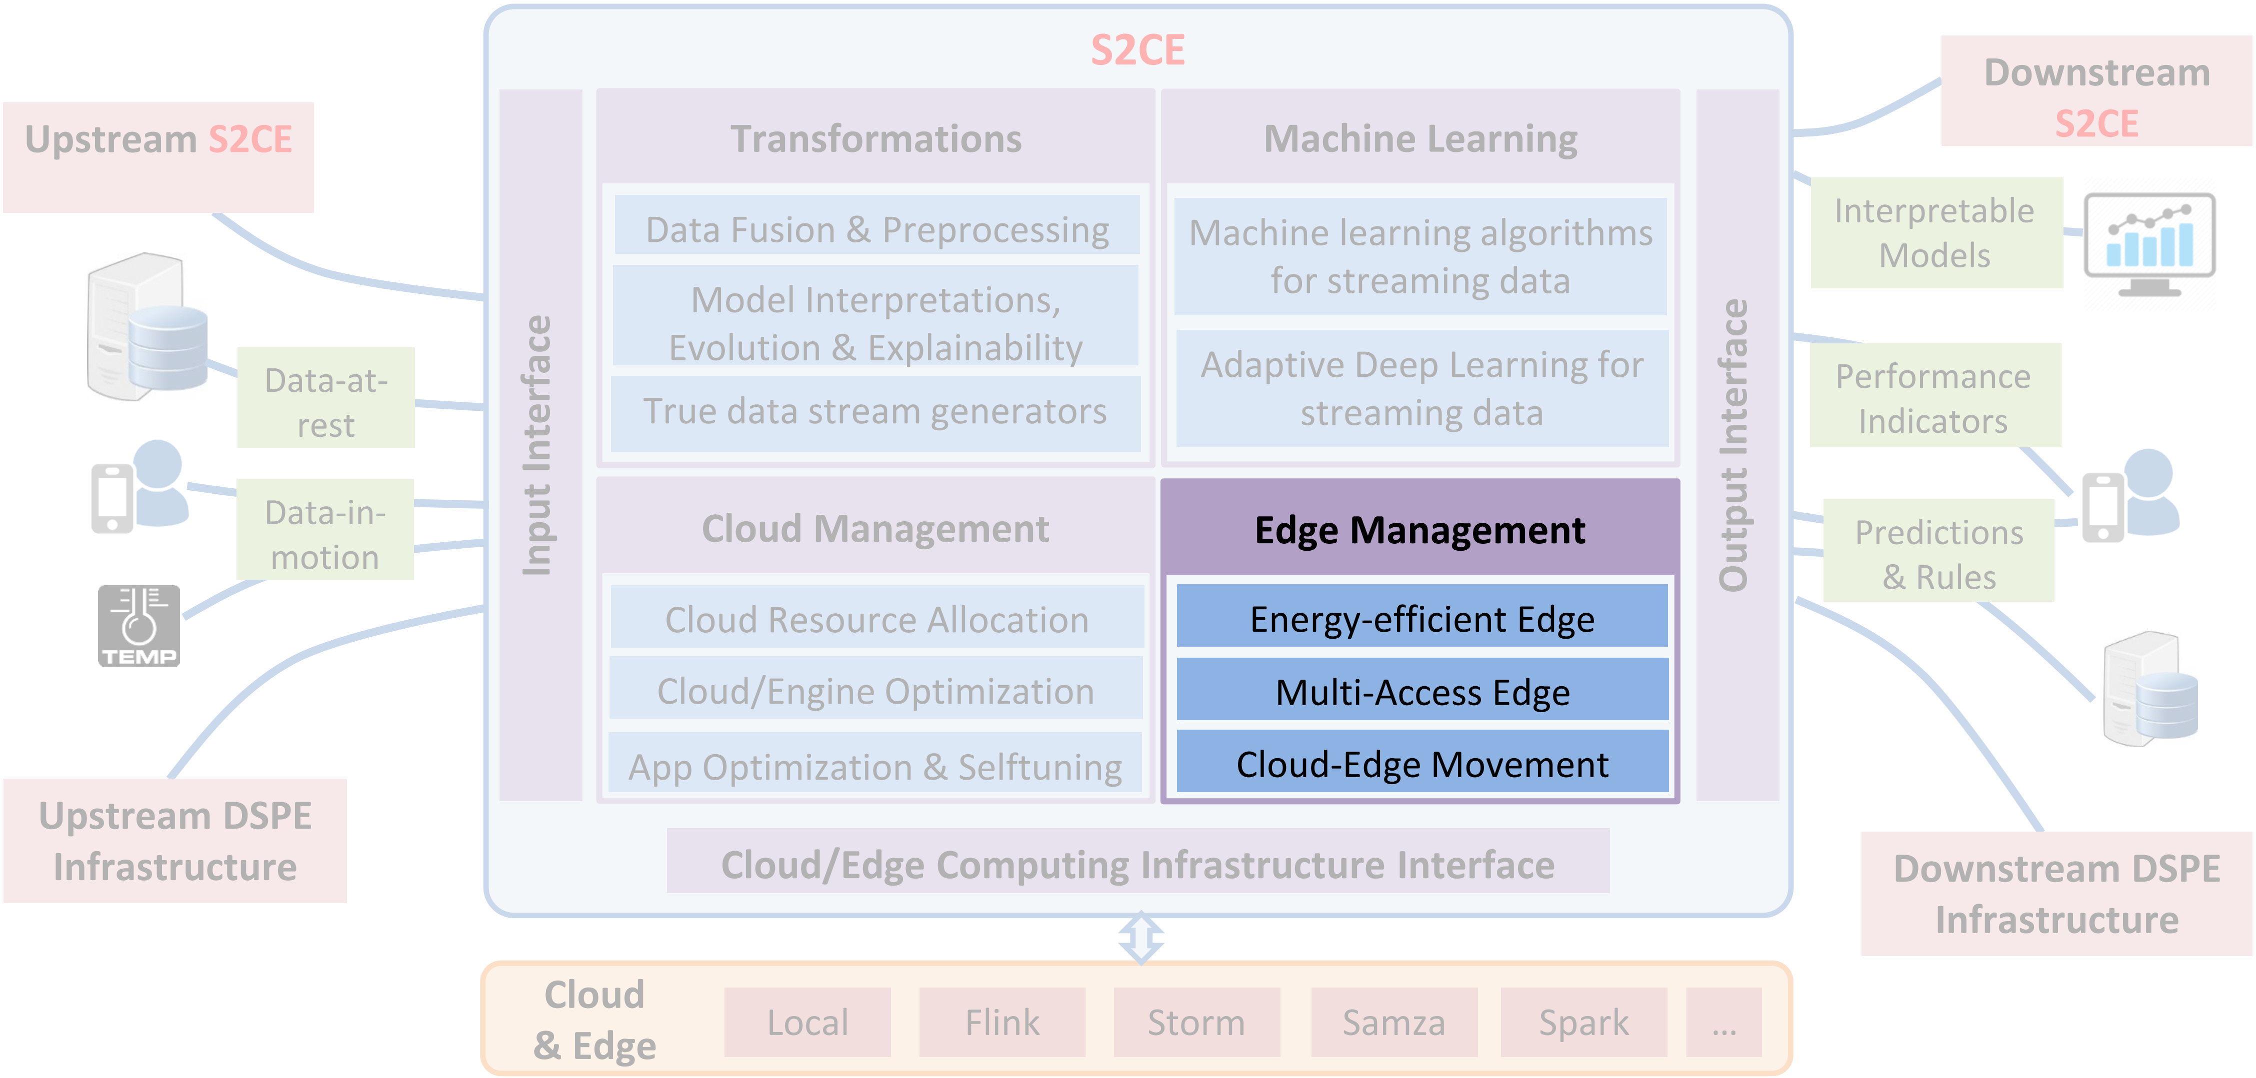

S2CE: Output Interface

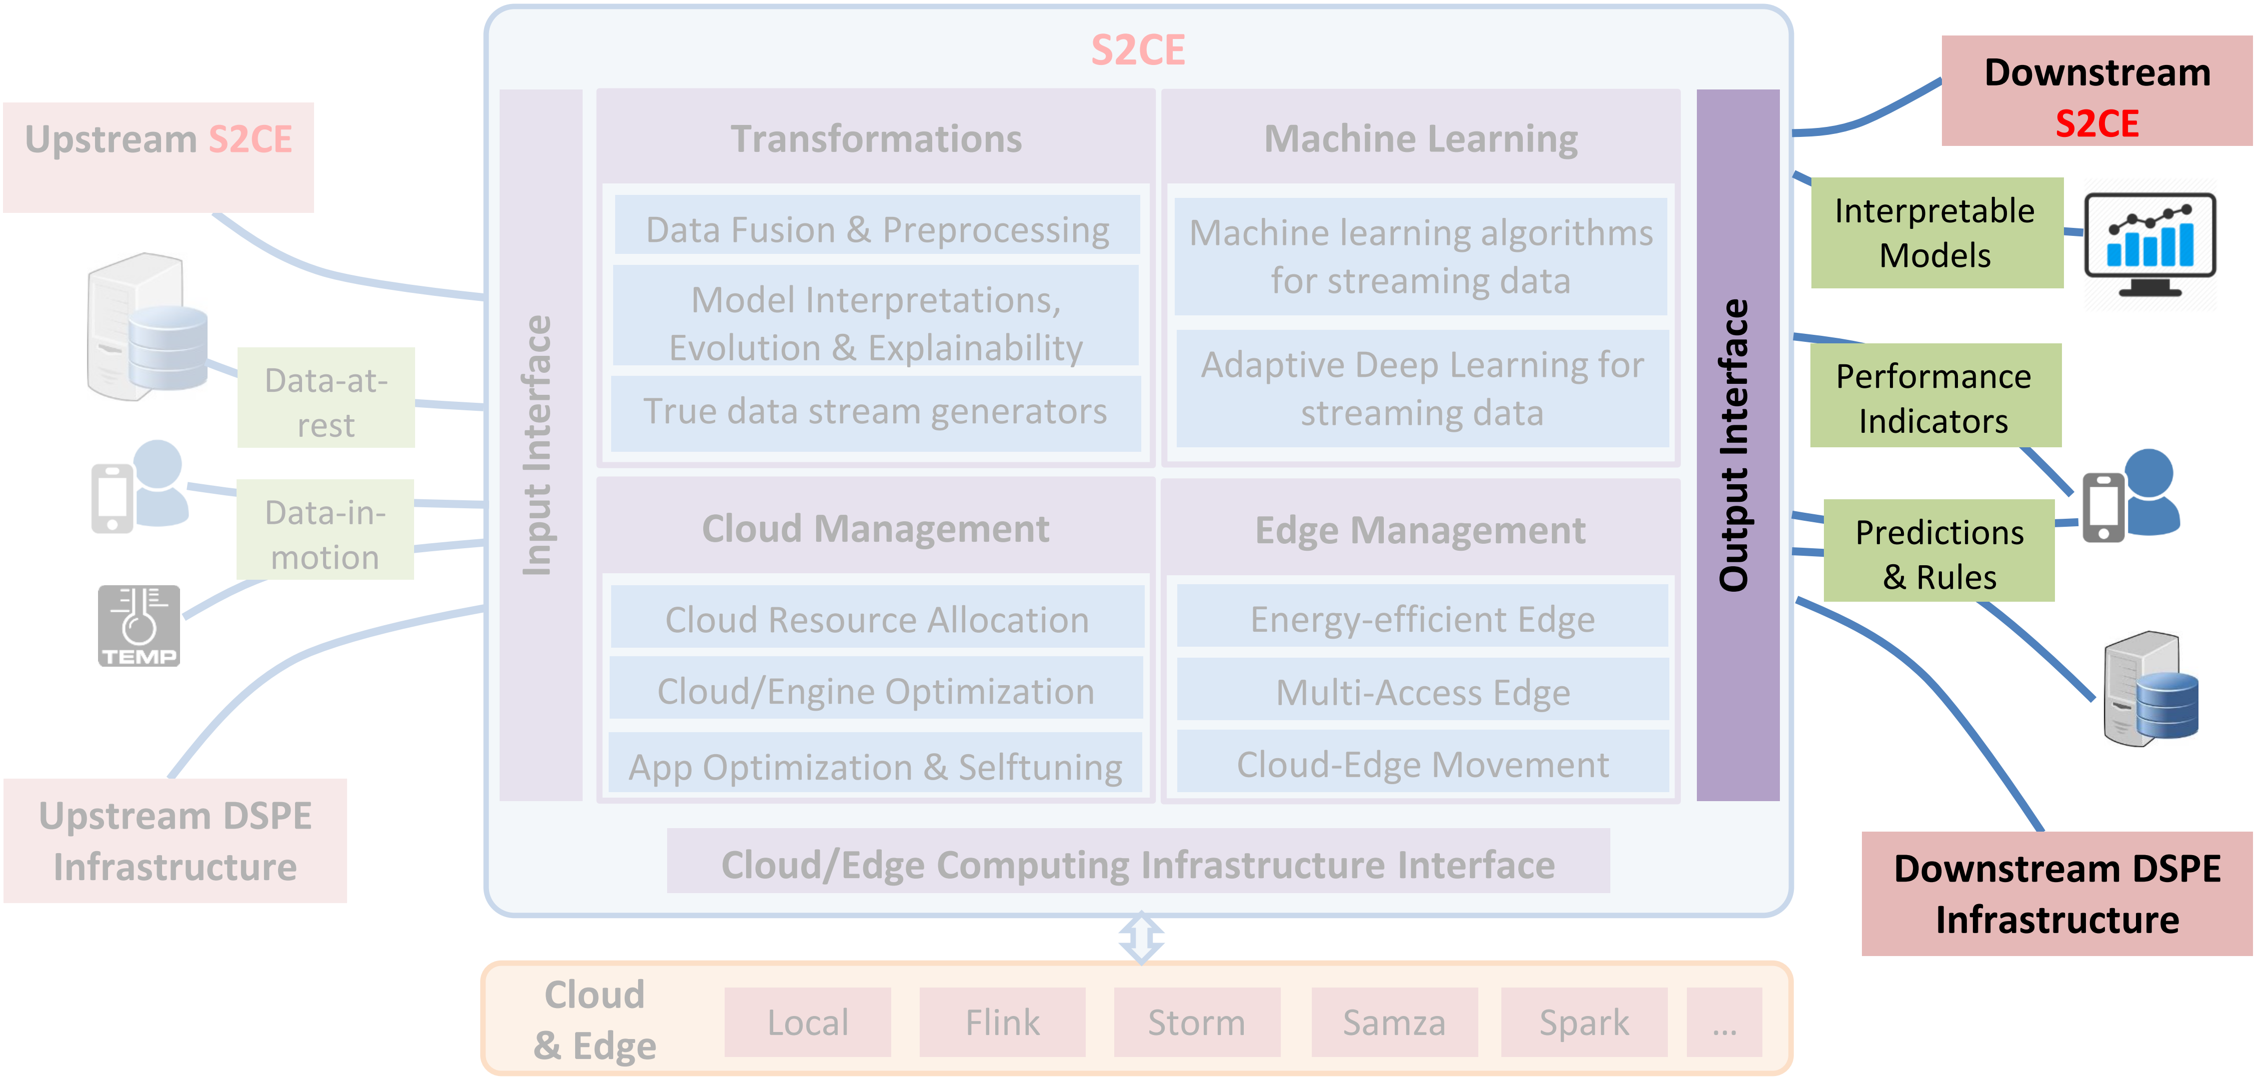

S2CE: Computing Infrastructure Interface

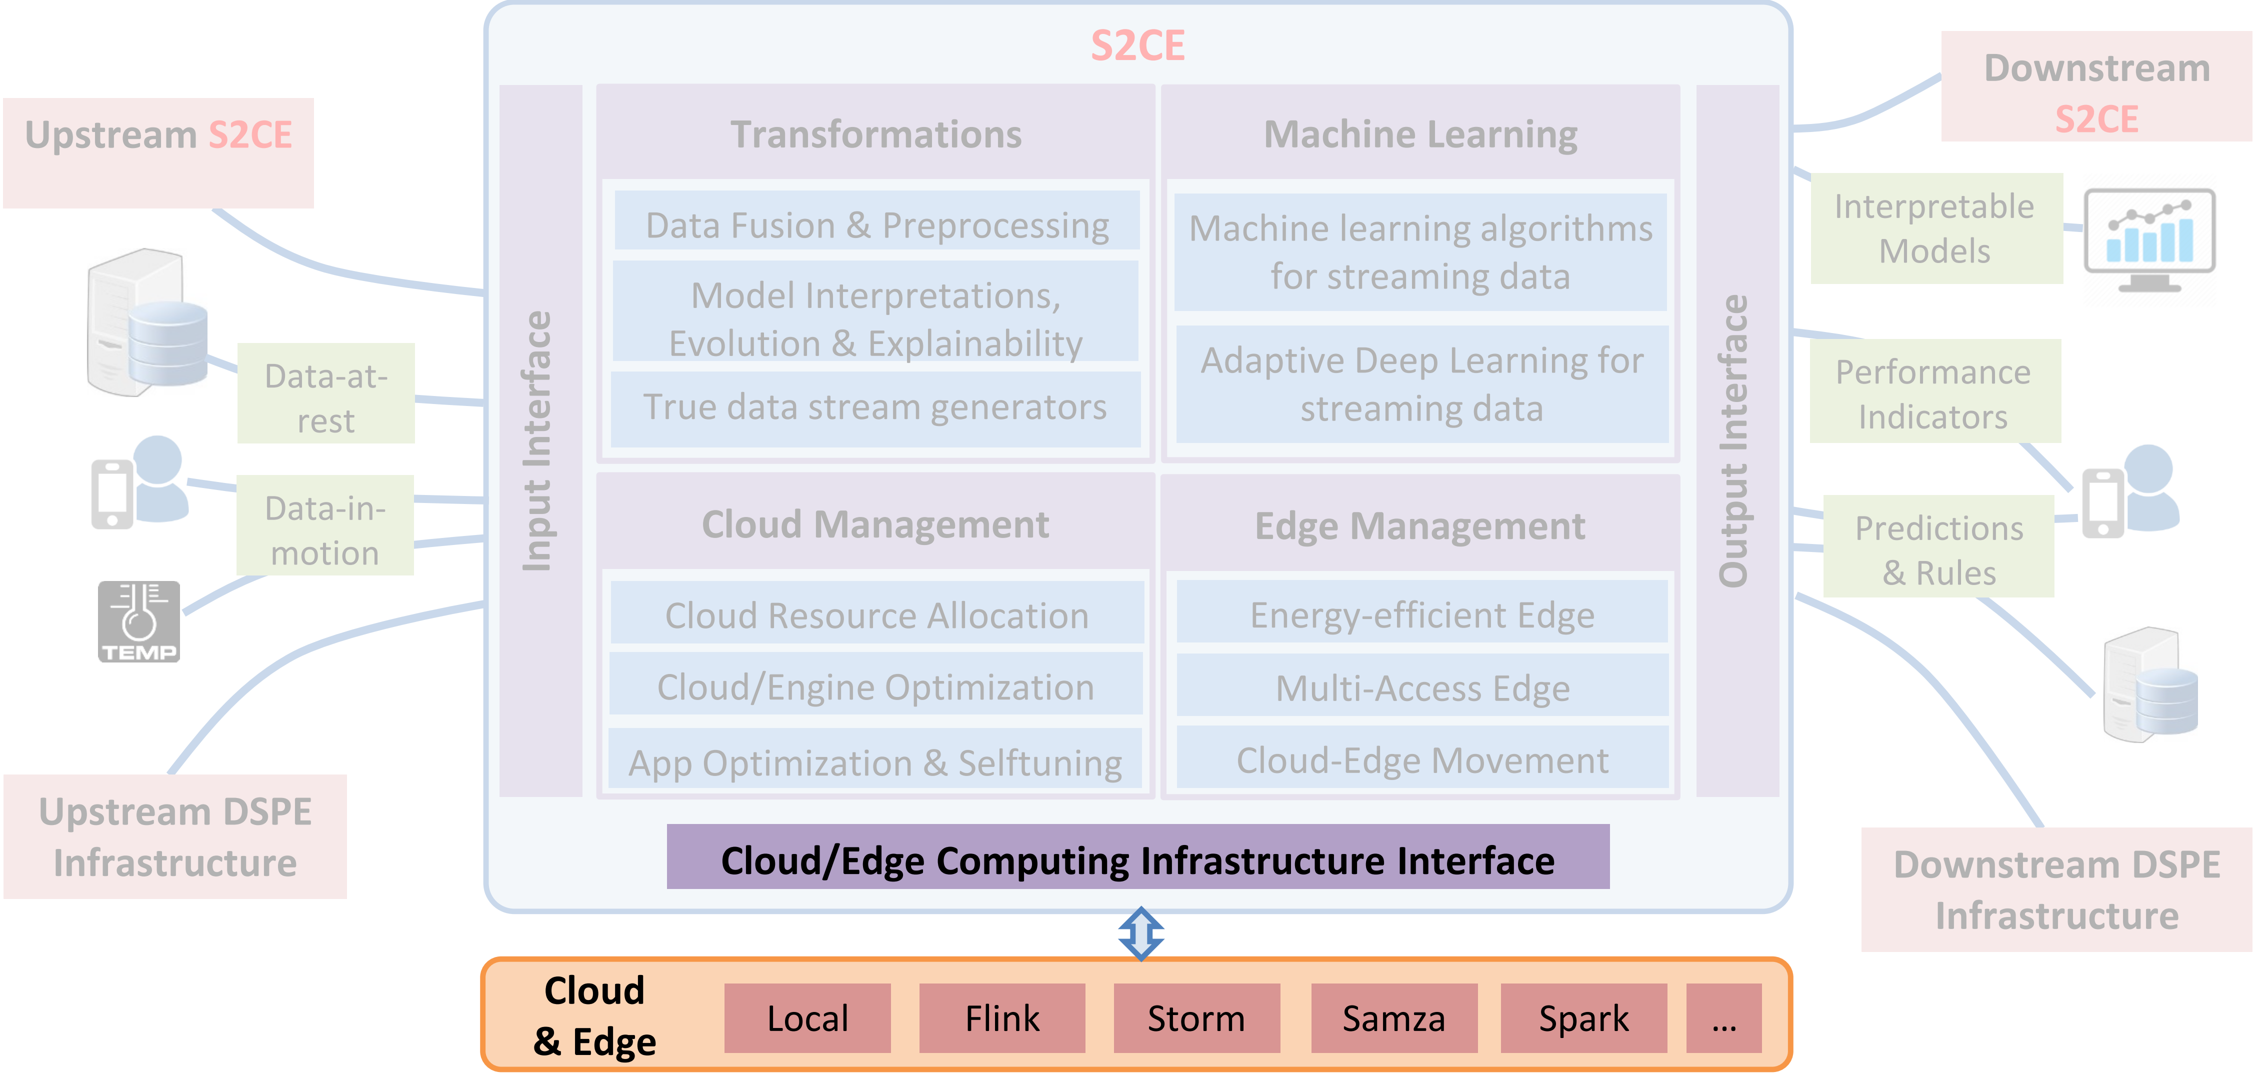

S2CE Performant Interconnections

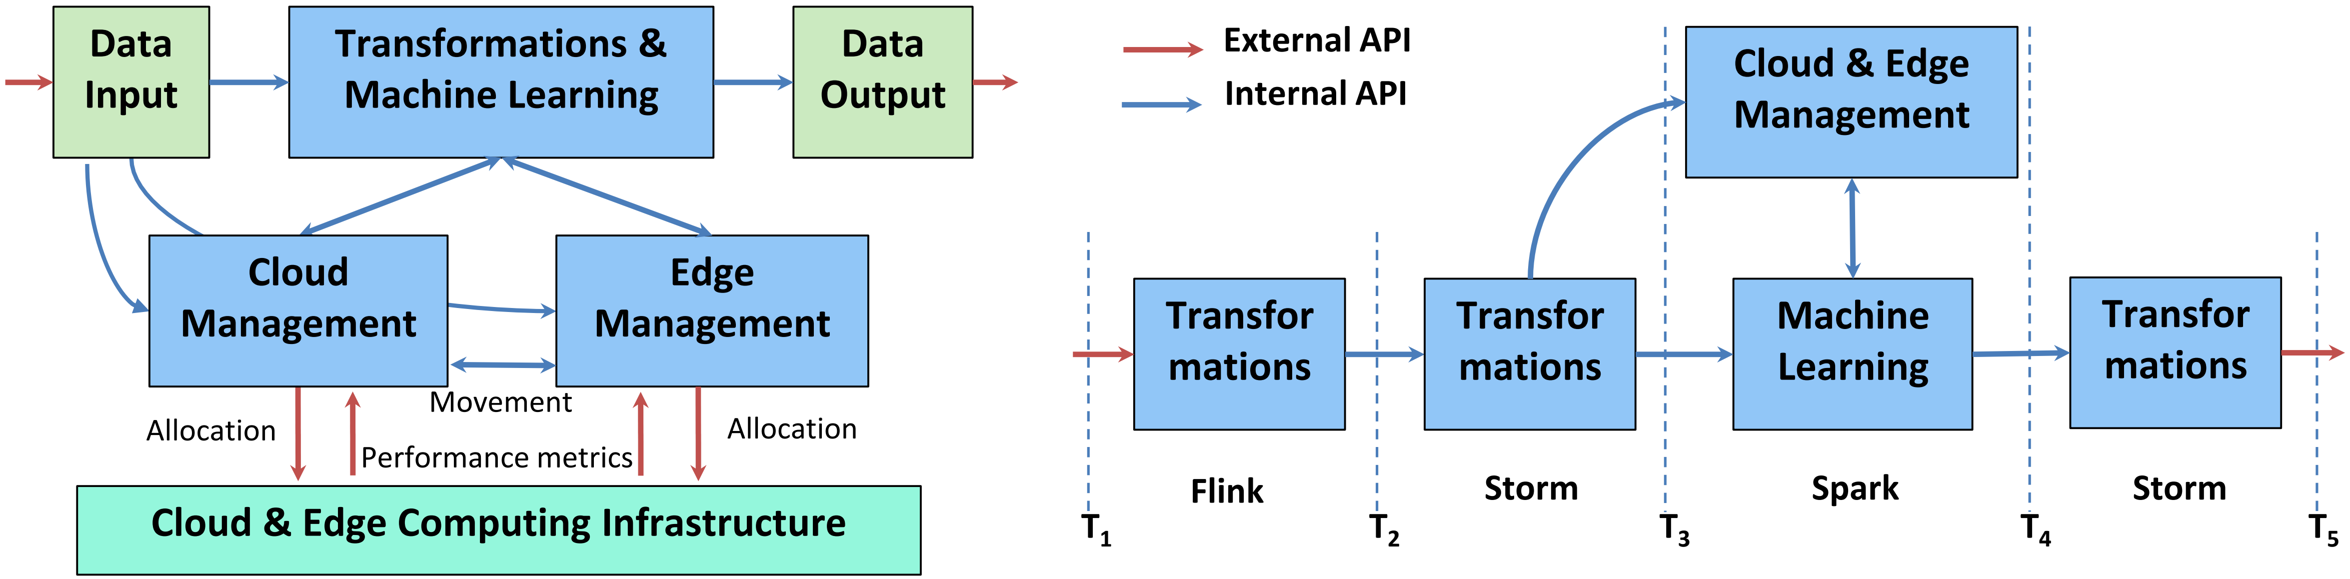

- APIs:
  - High-performant & secured
  - Scalable to support voluminous, fast & rich data

S2CE in Big Data Stream Processing Ecosystem

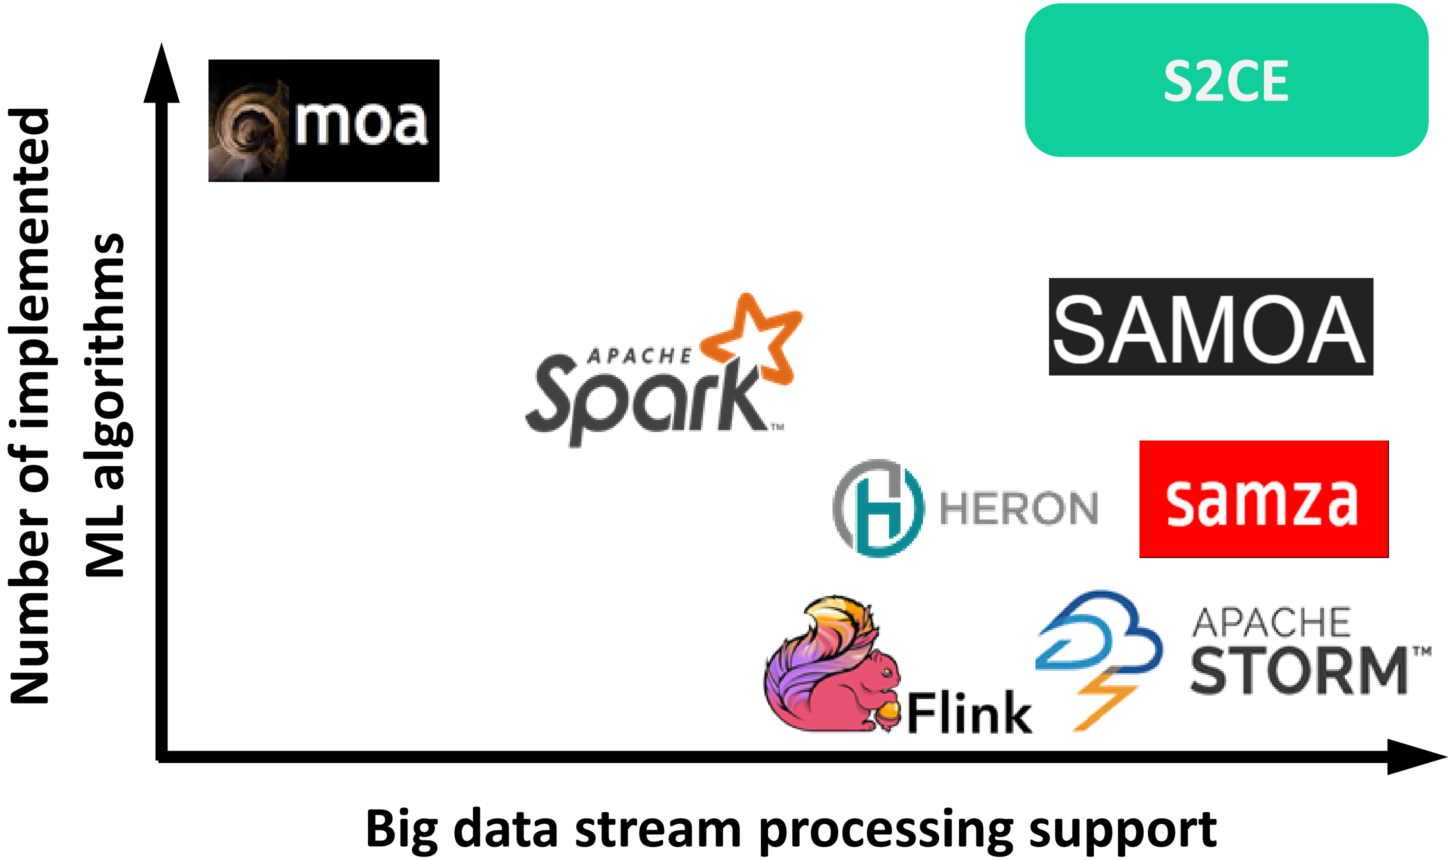

Data Intensive Computing Research Lab

Concluding Remarks: Innovation Potential

Cloud Provisioning and Orchestration

Edge Preprocessing and Movement

Data-driven Decisions based on Streams

Innovation Exchange with Apache Ecosystem’s Open Source

Unifying Efforts for a Stream ML Library

Thank you!

Herodotos Herodotou

herodotos.herodotou@cut.ac.cy

https://dicl.cut.ac.cy