Towards Autonomic Science Infrastructure: Architecture, Limitations, and Open Issues

Raj Kettimuthu, Zhengchun Liu, Ian Foster, Pete Beckman, Argonne National Laboratory

Alex Sim, Kesheng Wu, Lawrence Berkeley National Laboratory

Wei-keng Liao, Qiao Kang, Ankit Agrawal, Alok Choudhary, Northwestern University

1920 telephony

Situational awareness system

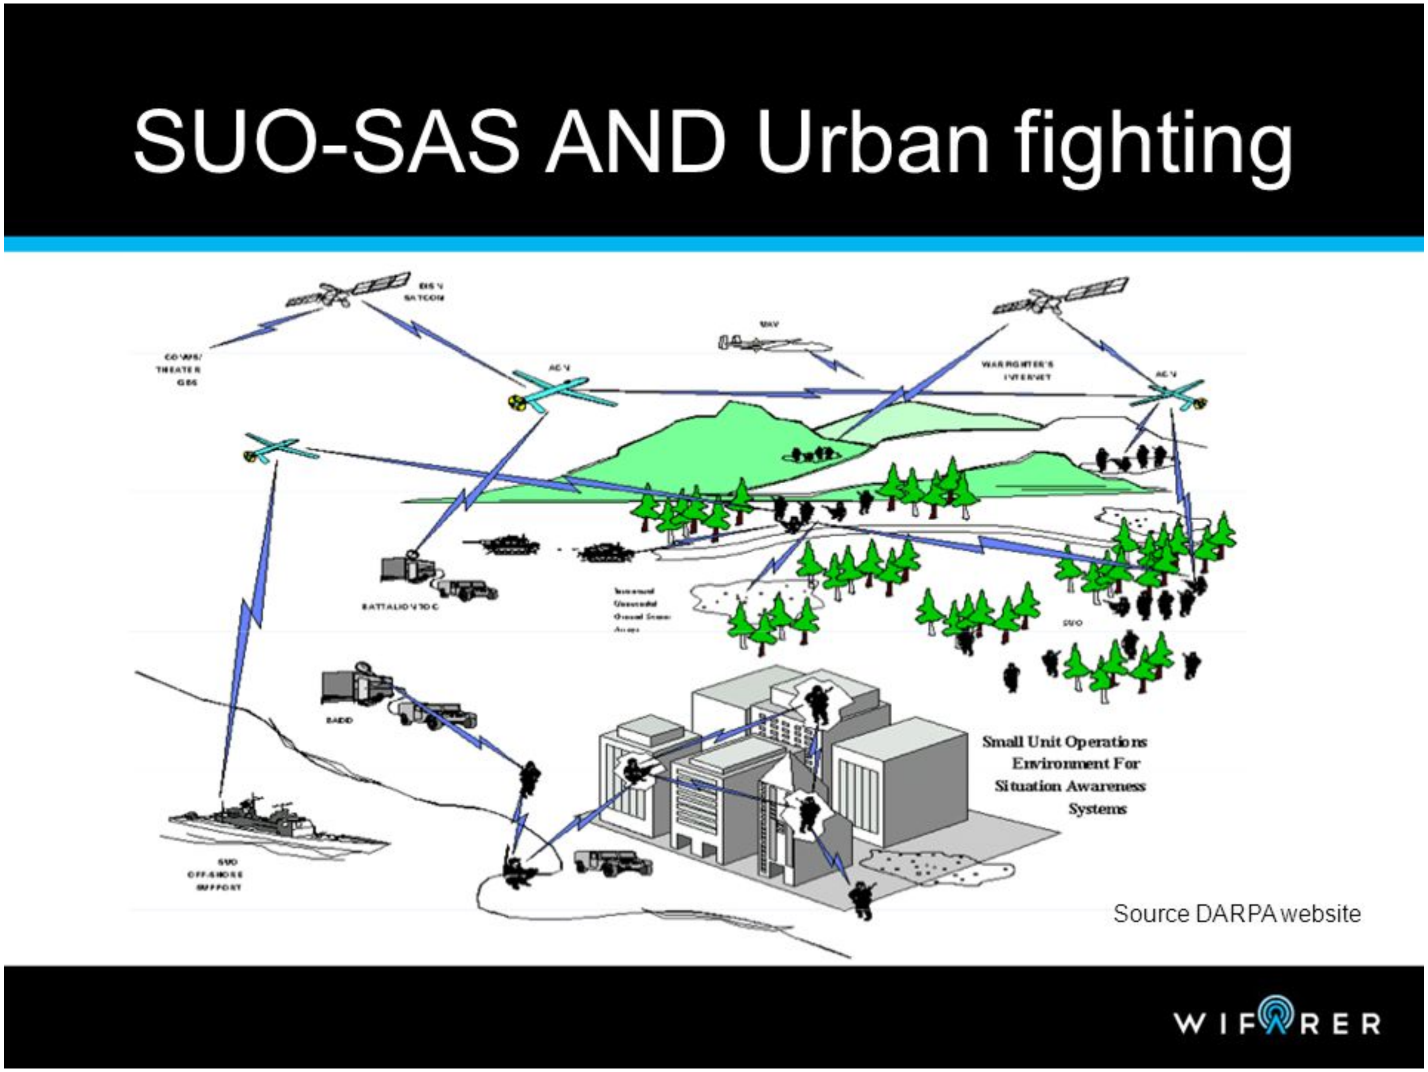

DASADA project

Dynamic Assembly for System Adaptability, Dependability and Assurance (DASADA)

The software revolution requires dynamic gauges

But gauges made it happen

And an ability to dynamically update and use models

John Salasin, PhD
Information Systems Office

IBM

- Complex computing systems should have autonomic properties
- Independently take care of the regular maintenance and optimization tasks
- Reduce workload on the system administrators
- Distilled four properties of a self-managing (i.e. autonomic) system:
  - Self-configuration: adapting to dynamically changing environments
  - Self-optimization: tuning resources and balancing workloads to maximize use of IT resources
  - Self-healing: discovering, diagnosing, and acting to prevent disruptions
  - Self-protecting: anticipating, detecting, identifying, and protecting against attacks

MAPE-K

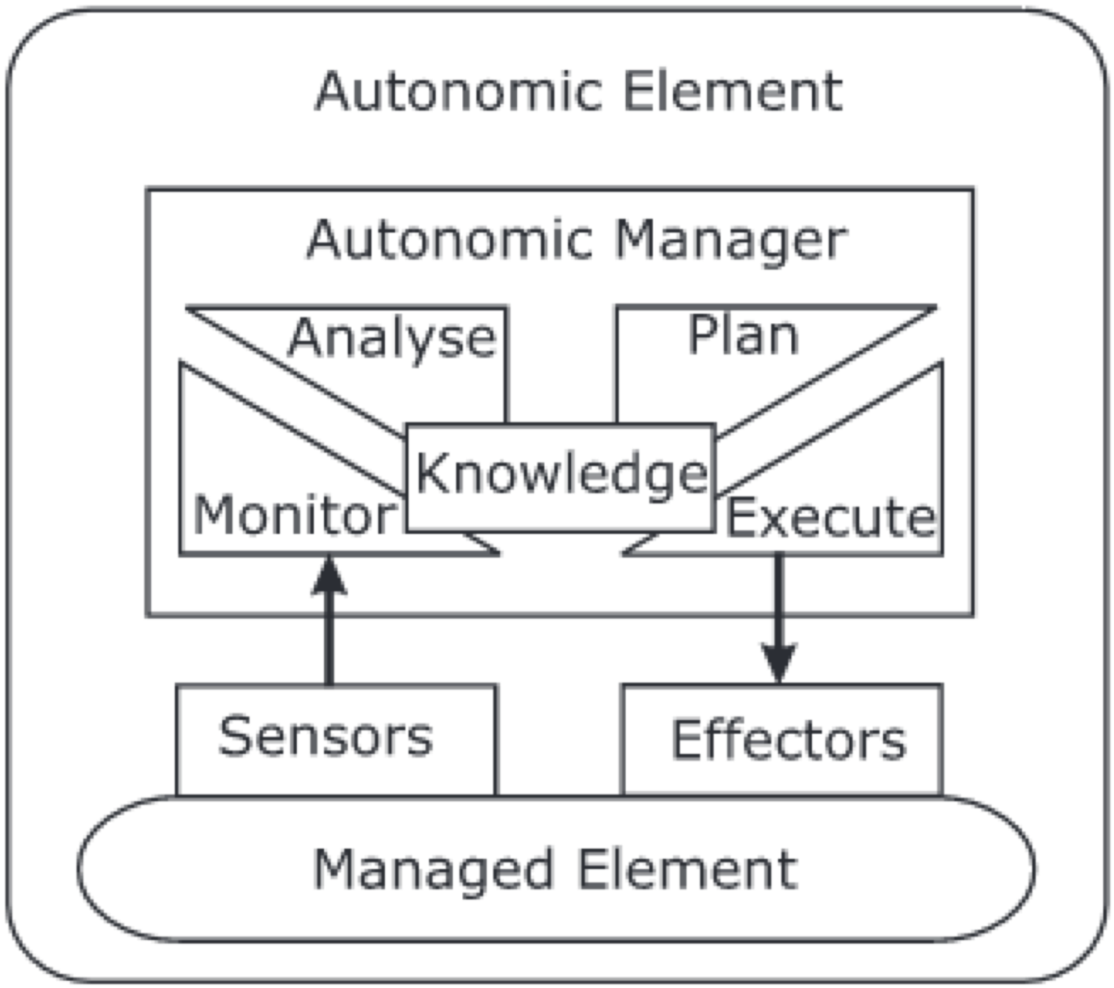
Autonomic Element
Autonomic Manager
Analyse
Plan
Knowledge
Monitor
Execute
Sensors
Effectors
Managed Element

Autonomic element

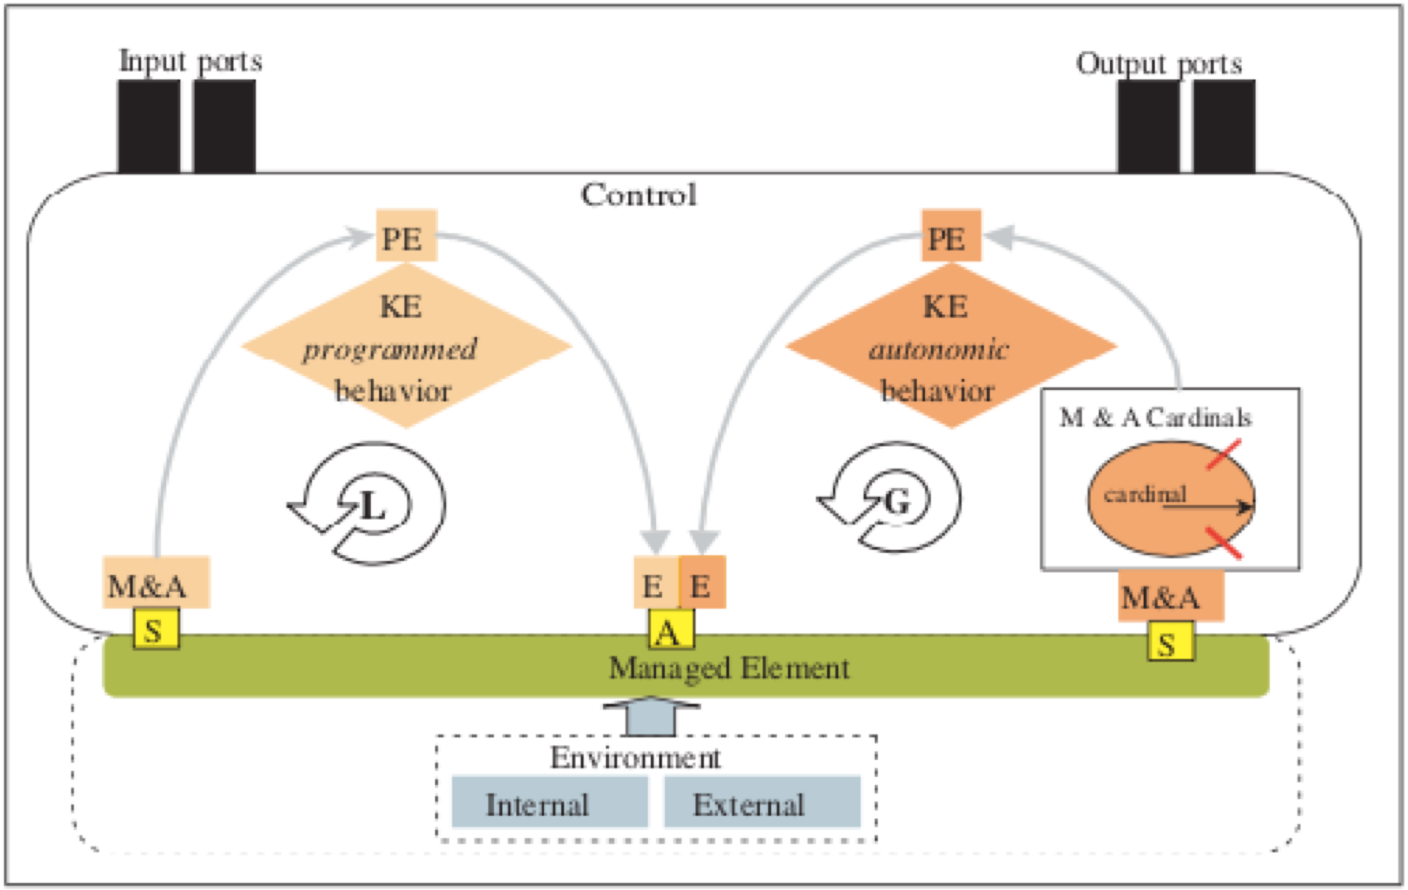

M. Parashar, S. Hariri, "Autonomic Computing: An Overview," Conference Paper in Lecture Notes in Computer Science, January 2004

Multi-tier model

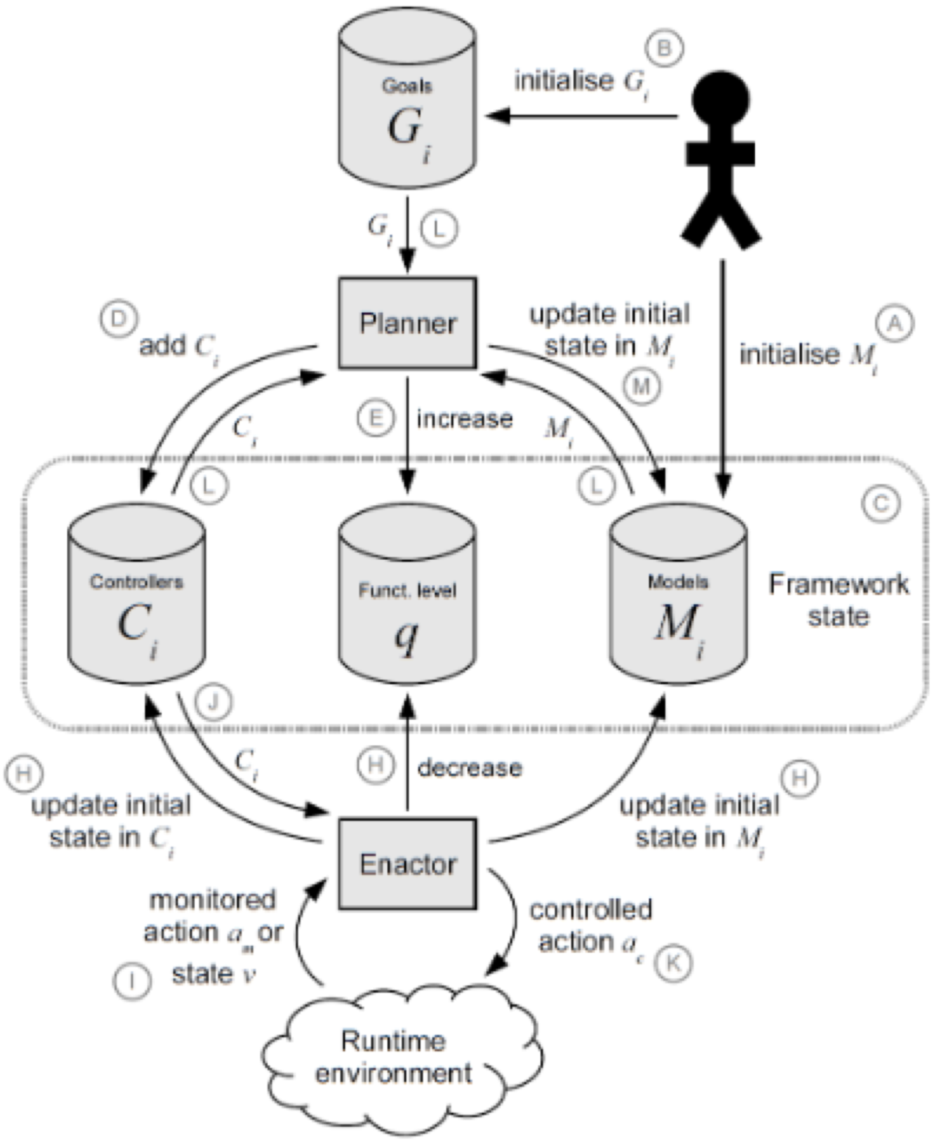

N. D'Ippolito, V. Braberman, J. Kramer, J. Magee, D. Sykes, and S. Uchitel, "Hope for the Best, Prepare for the Worst: Multi-tier Control for Adaptive Systems," ICSE 2014.

Dynamic job scheduling policies

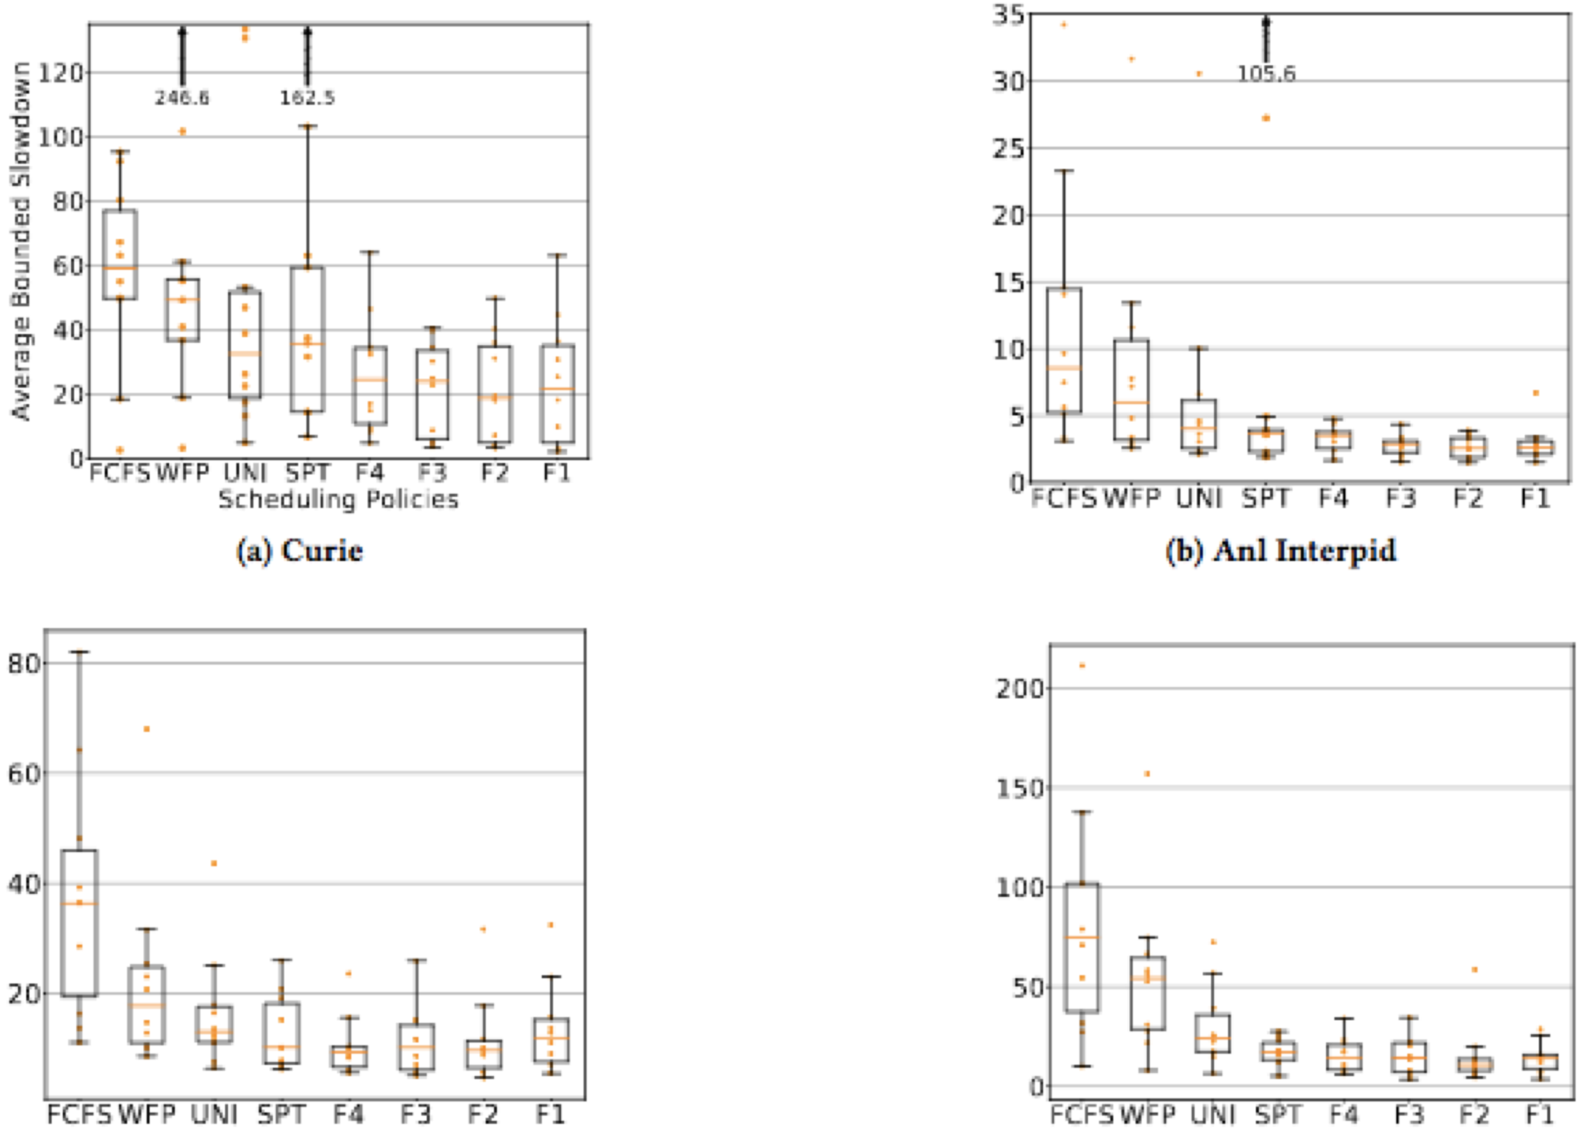

(a) Curie (b) Anl Interpid (c) SDSC Blue (d) CTC SP2

D. Carastan-Santos, R. Y. De Camargo. Obtaining Dynamic Scheduling Policies with Simulation and Machine Learning. SC'17.

Anomaly detection

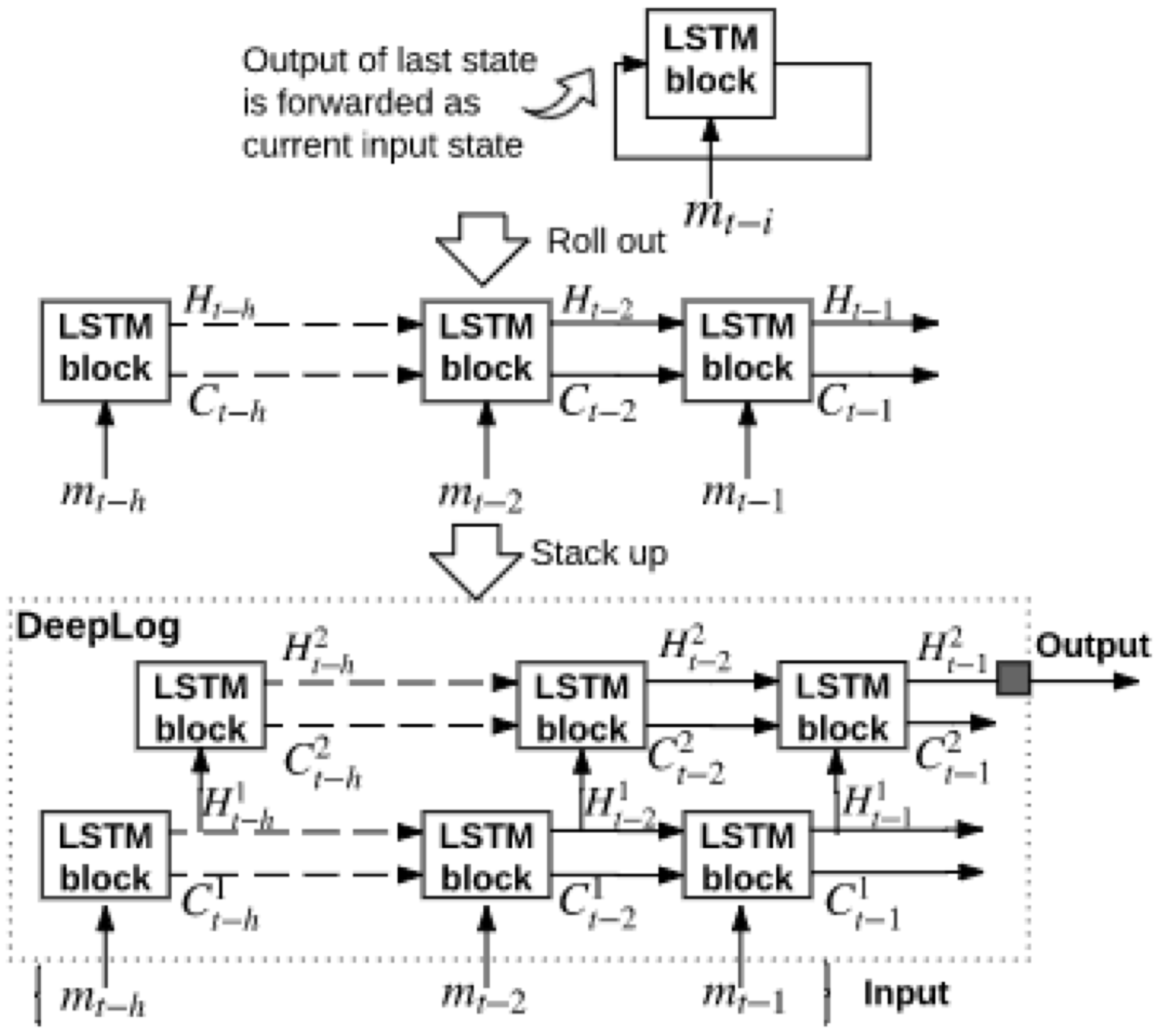

M. Du, F. Li, G. Zheng, and V. Srikumar, "DeepLog: Anomaly Detection and Diagnosis from System Logs through Deep Learning," In ACM SIGSAC Conference on Computer and Communications Security (CCS ’17).

Prefetching

(a) Clustering the address space into separate streams.

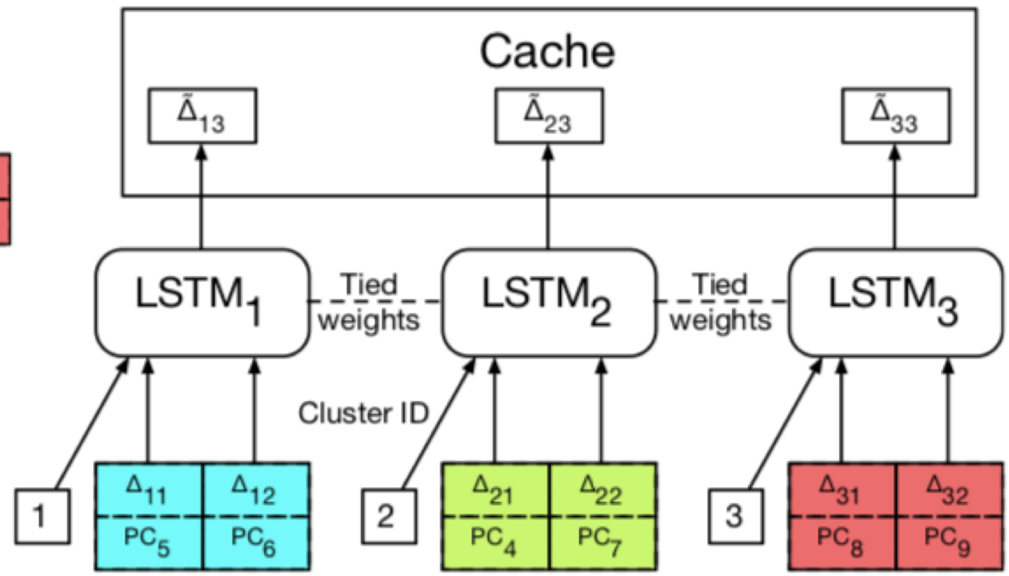

(b) The clustering + LSTM model.

M. Hashemi, K. Swersky, J.A. Smith, G. Ayers, H. Litz, J. Chang, C. Kozyrakis, P. Ranganathan, "Learning Memory Access Patterns", 2018.

Tune performance of storage systems

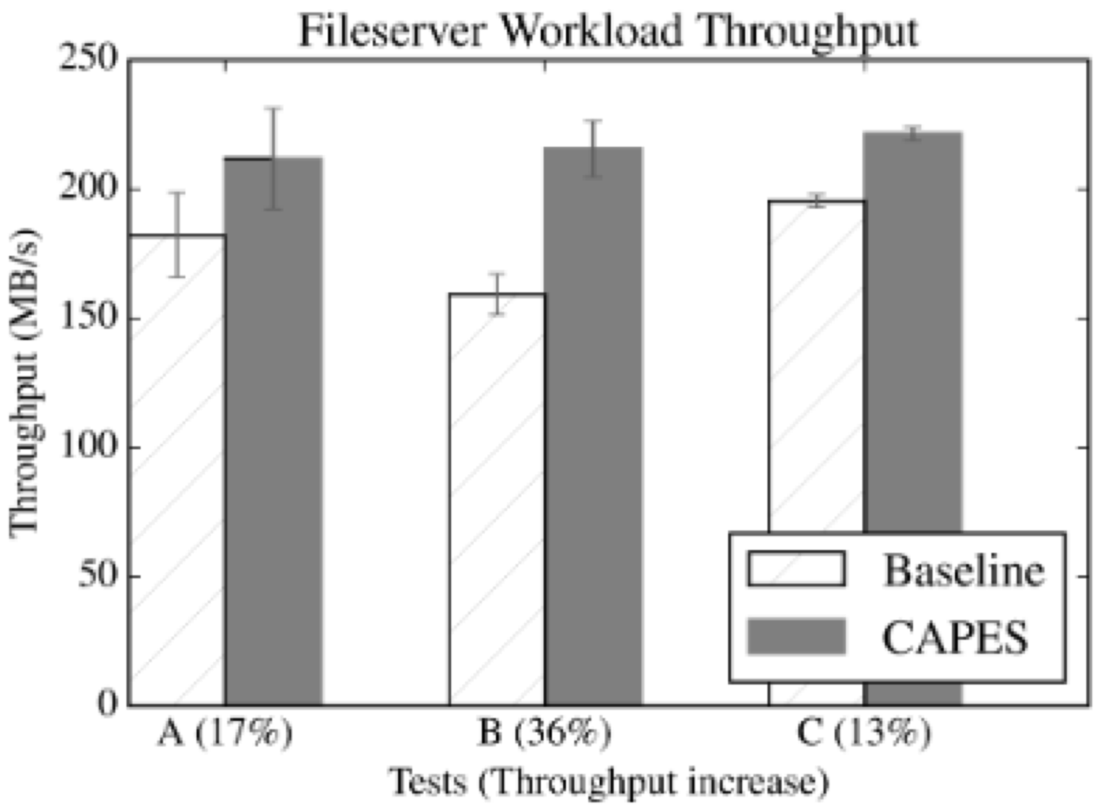

Y. Li, K. Chang, O. Bel, E. L. Miller, and D. Long, "CAPES: Unsupervised Storage Performance Tuning Using Neural Network-Based Deep Reinforcement Learning," SC17.

DeepMind AI Reduces Google Data Centre Cooling Bill by 40%

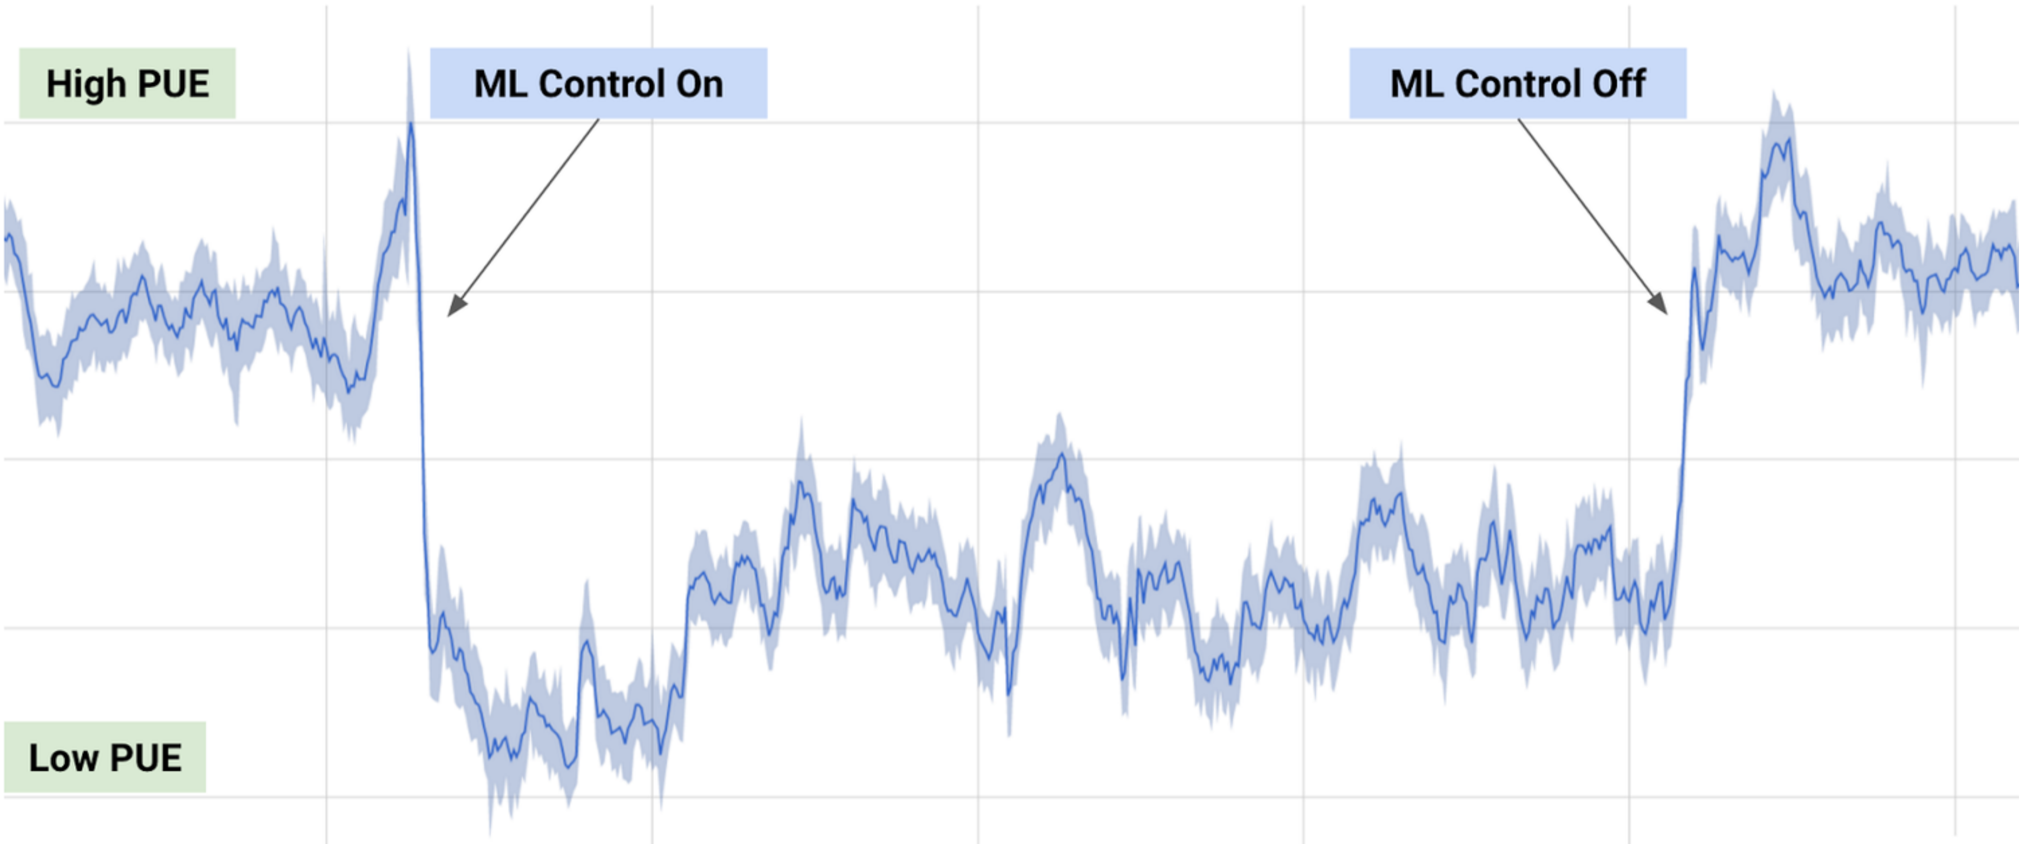

https://deepmind.com/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-40/

Self driving cars

Challenges

- ML methods have promise as a means of creating autonomous modules for computer systems
- Coexistence of several autonomous modules is required in order to handle multiple concerns - requires coordination
- As systems scale toward exascale, many resources will become increasingly constrained
- Some resources have historically been allocated explicitly, others such as N/W BW, I/O BW, and power are not
- As systems continue to evolve, we expect many such resources will need to be explicitly managed
- Autonomous management of resources and coordination to achieve higher-level goals even more important

Autonomous science infrastructure architecture

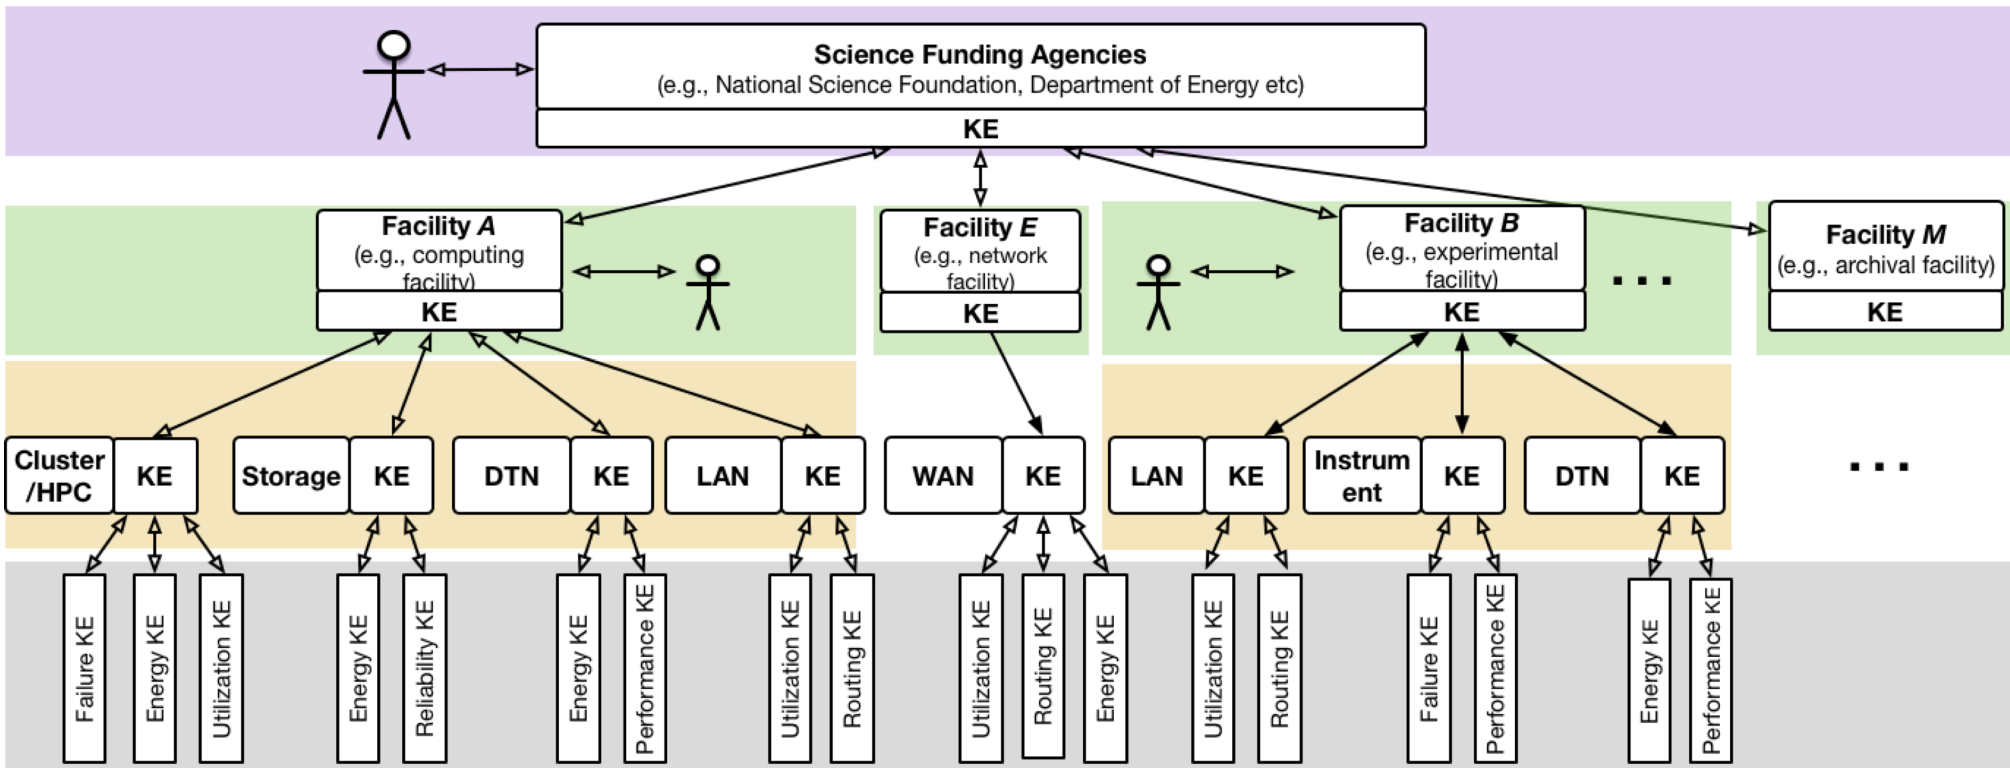

Architecture of an individual autonomous system

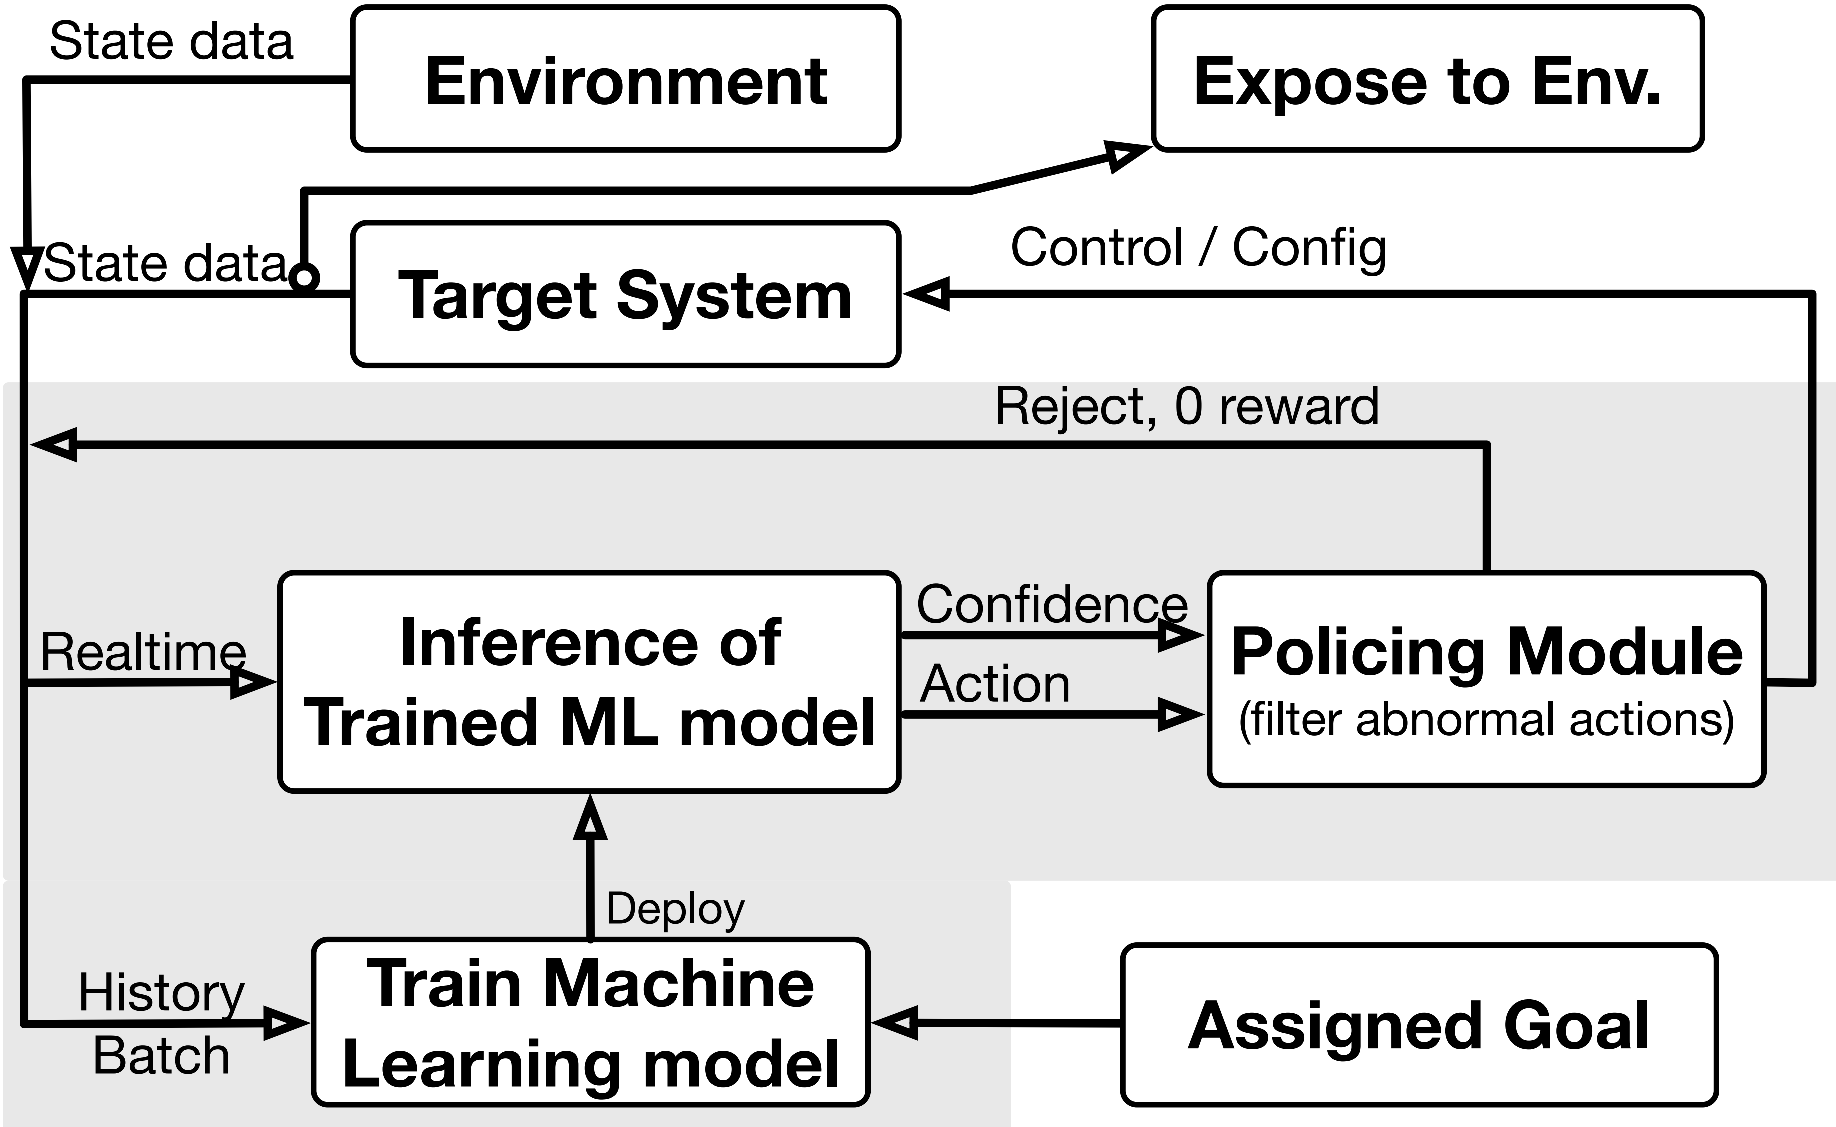

Autonomous science infrastructure architecture

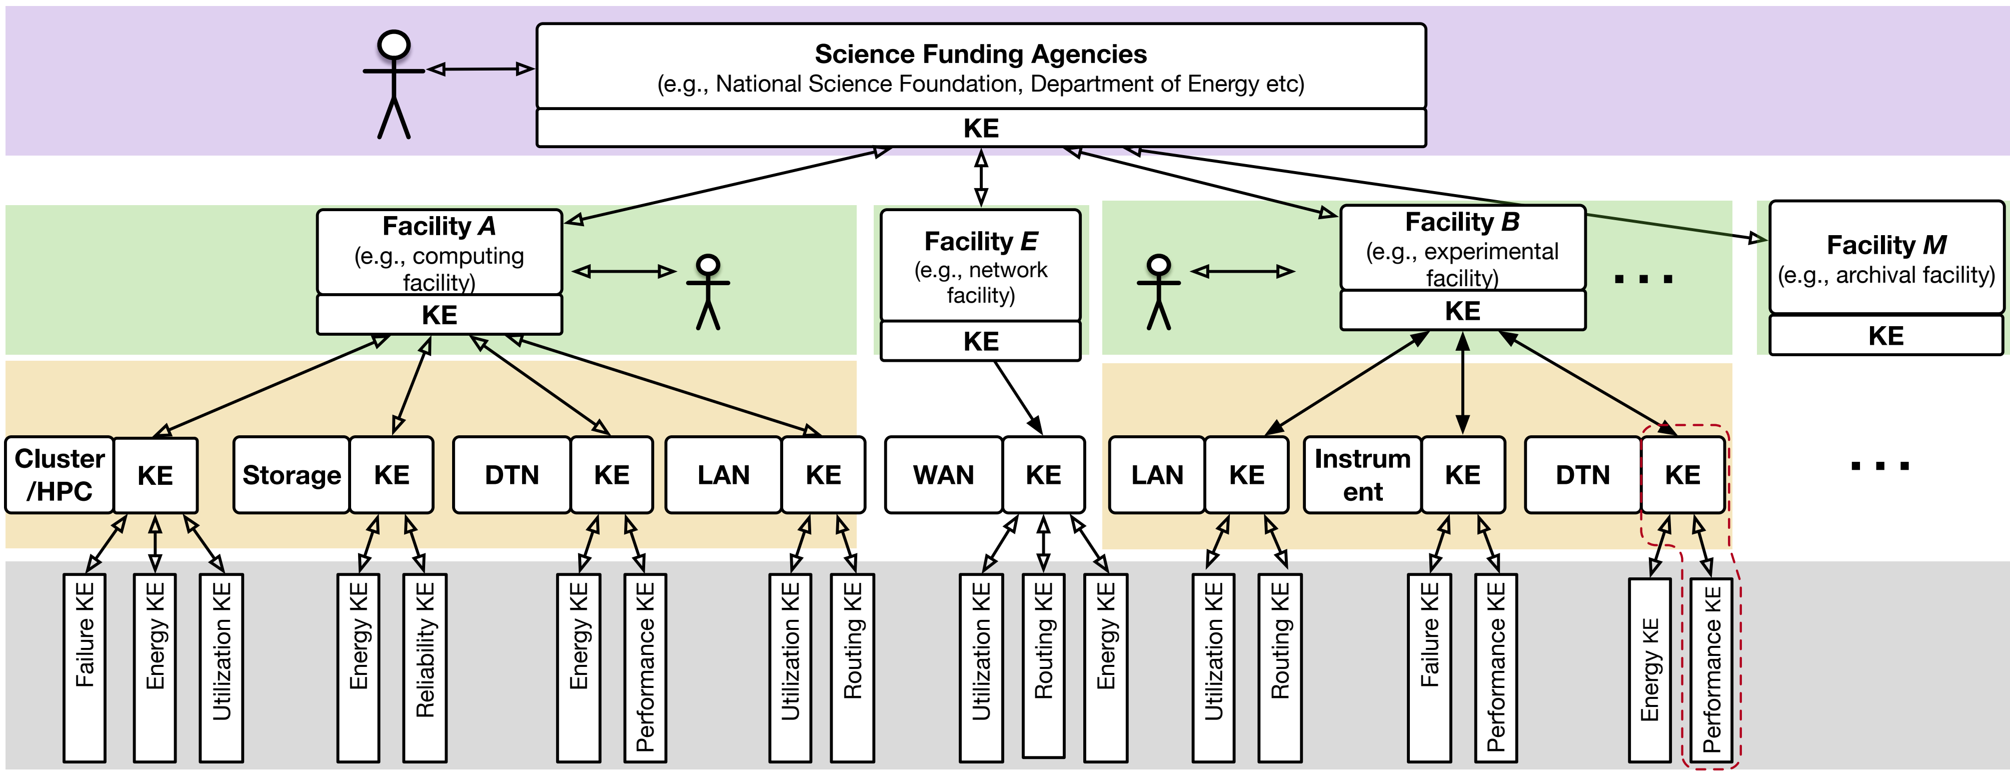

Autonomous performance module for DTN

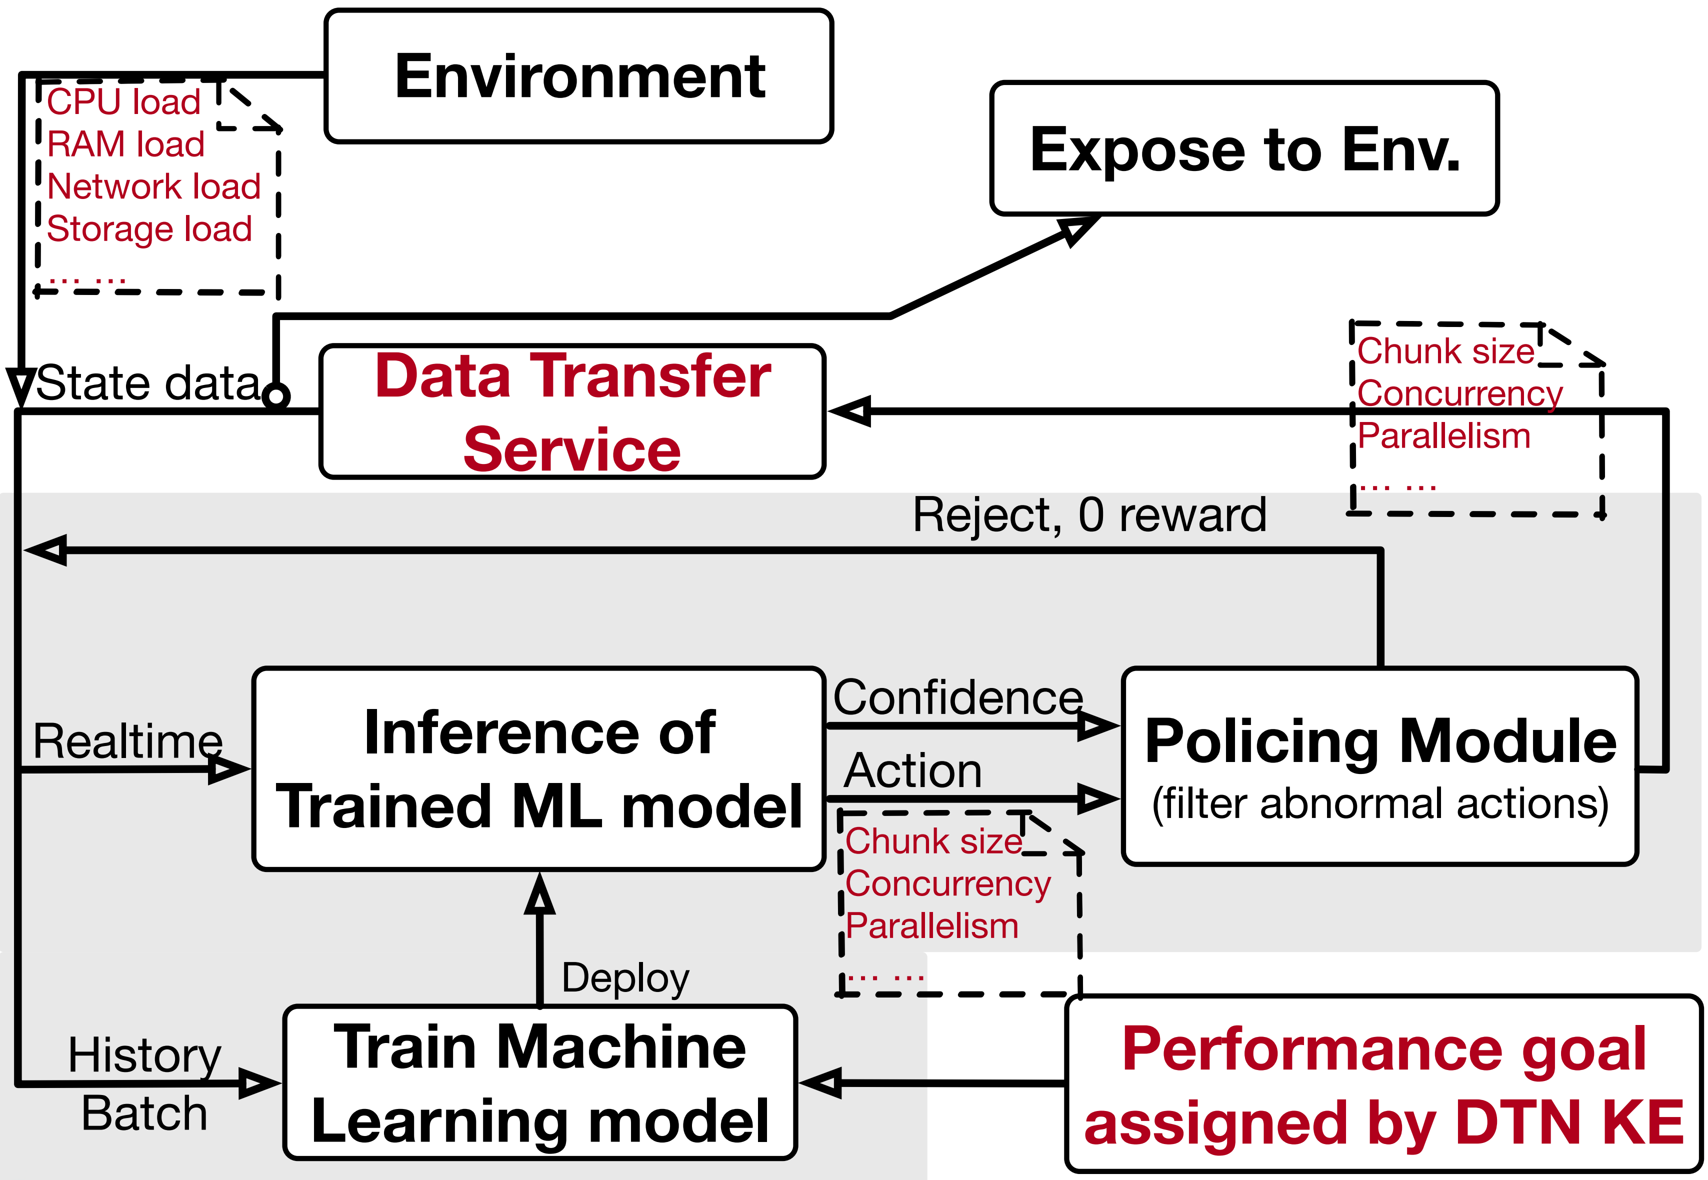

Concurrent transfers

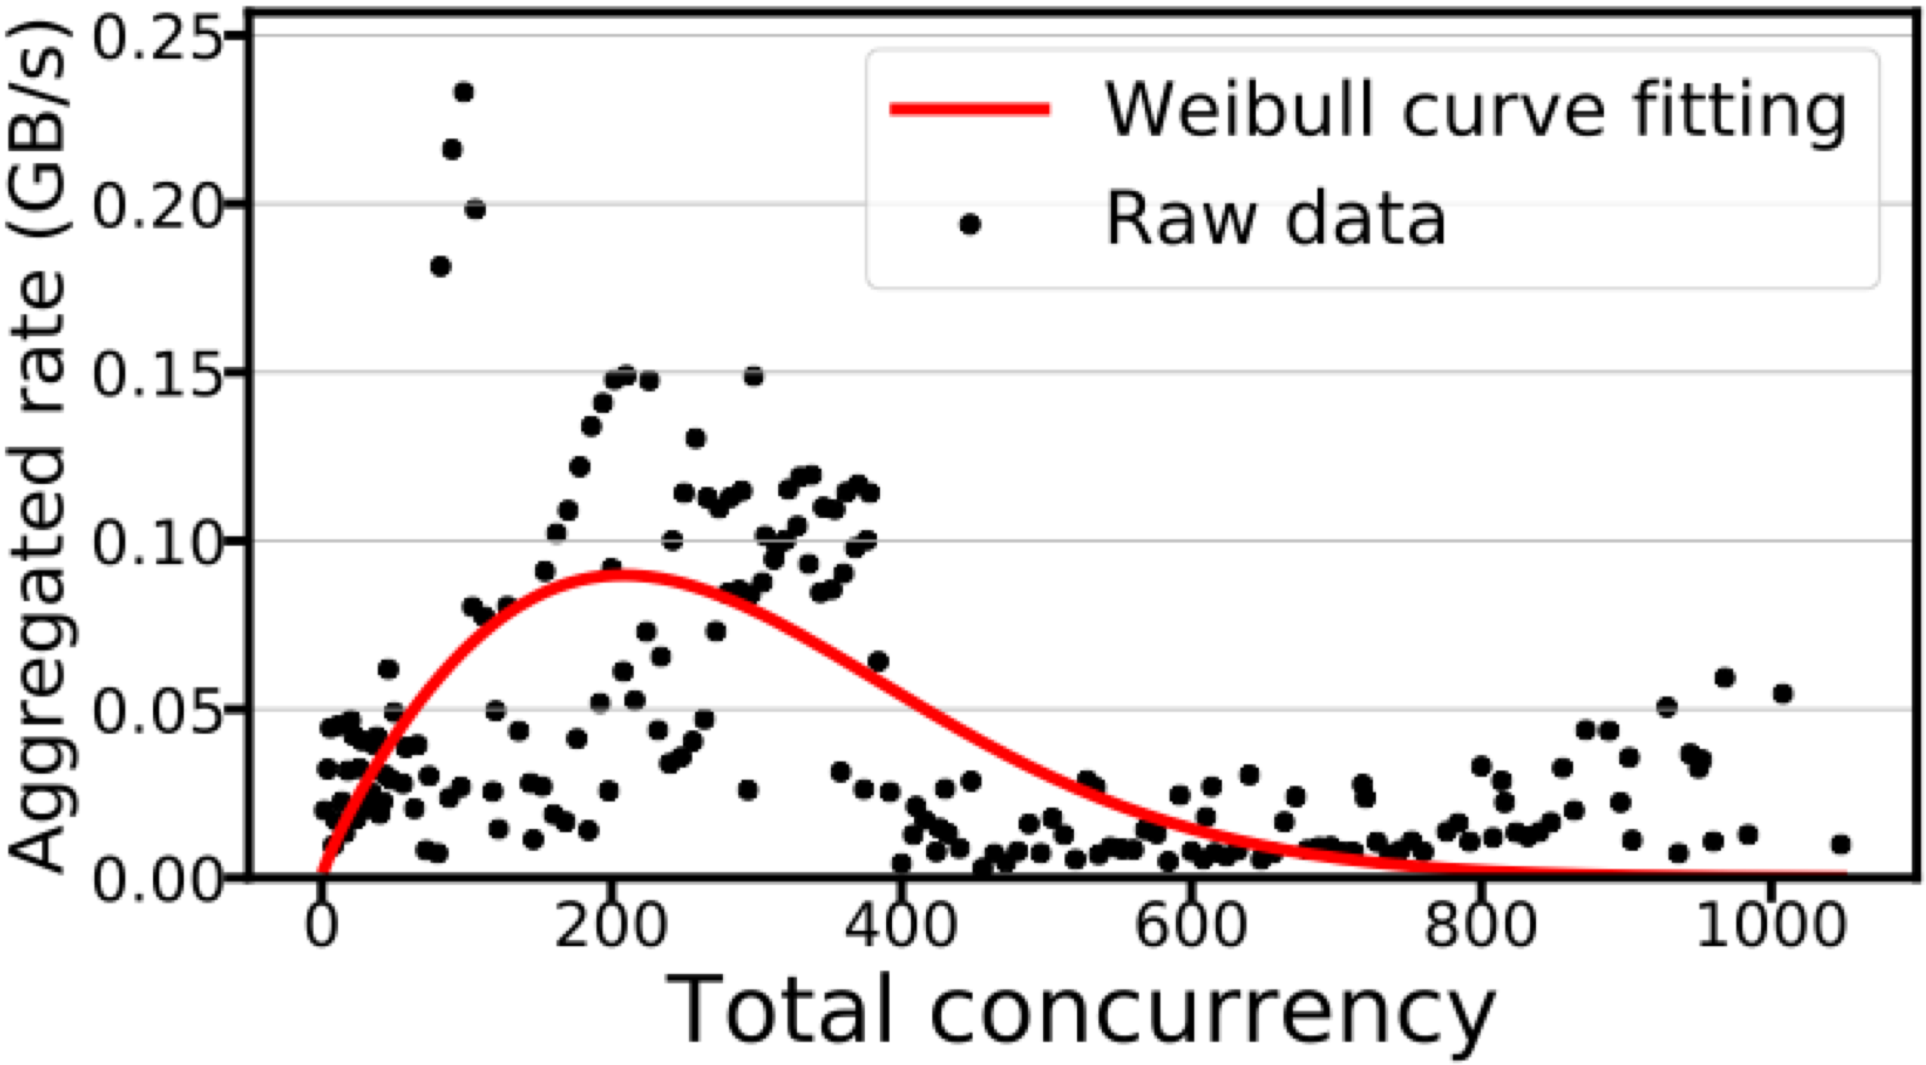

Smart data transfer node

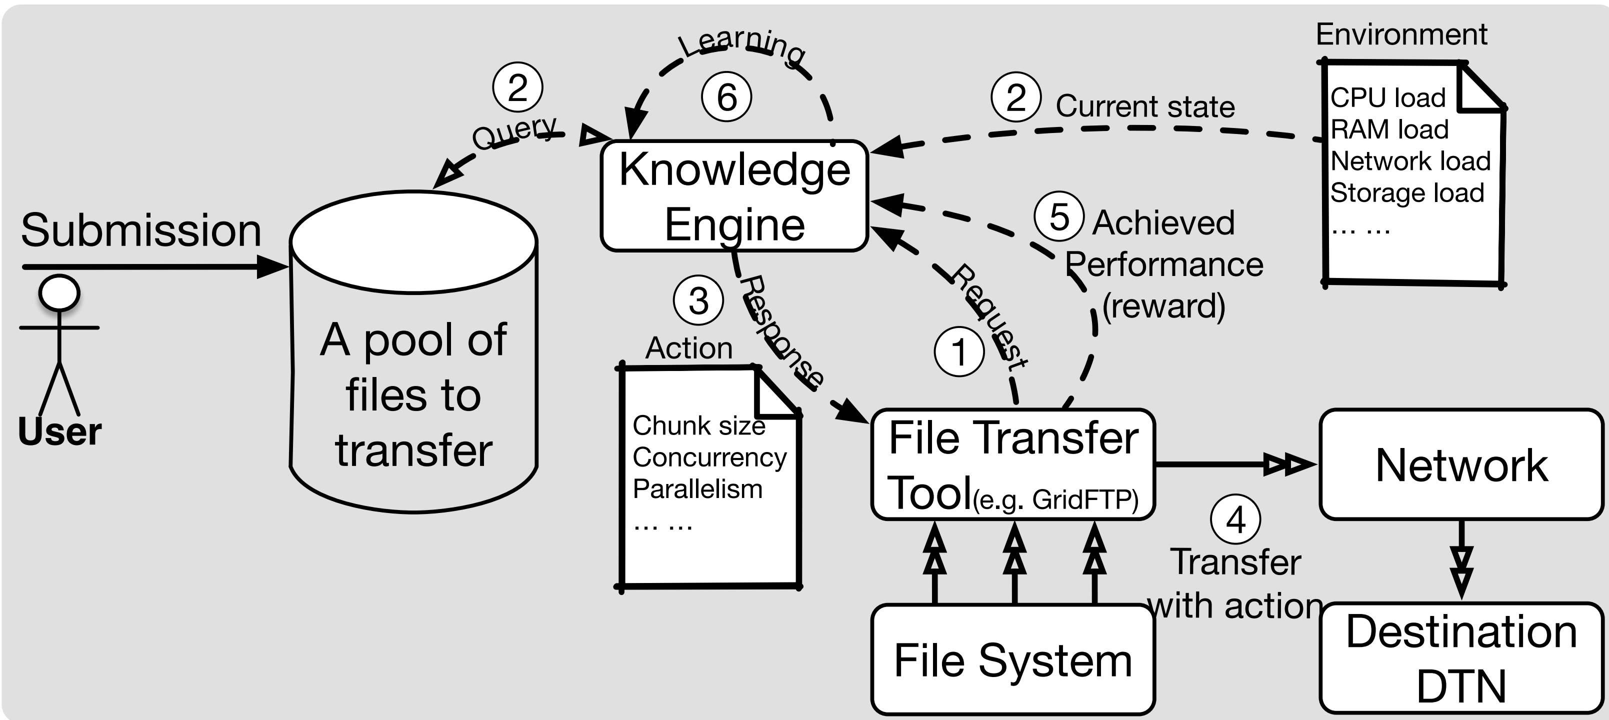

① A file transfer tool requests a file to transfer from the KE. The KE ② checks the current DTN state and ③ responds to the transfer tool with a chunk of file and corresponding optimal transfer parameters (the steering action). ④ The transfer tool transfers the associated chunk with the parameters and monitors the aggregate DTN throughput during this transfer. ⑤ Once completed, DTN's average aggregate throughput is reported to the KE as a reward for its actions. ⑥ Based on the reward (encourage or discourage), the KE updates its internal model parameters to improve its decision policy.

Autonomous science infrastructure architecture

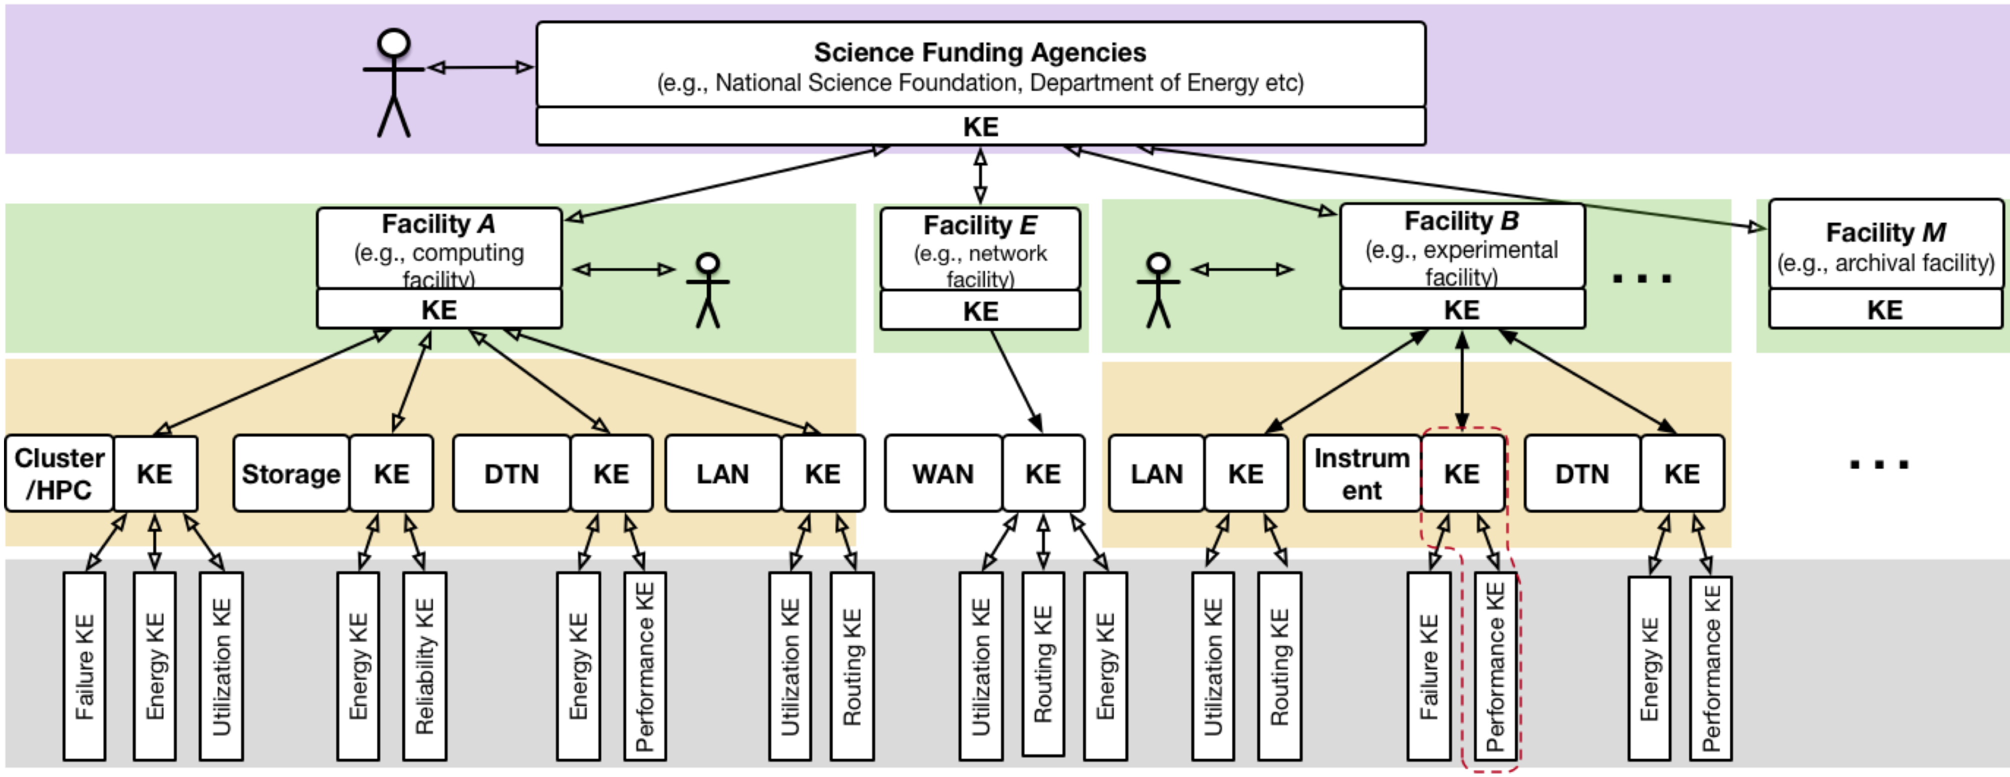

Experiment monitoring and steering

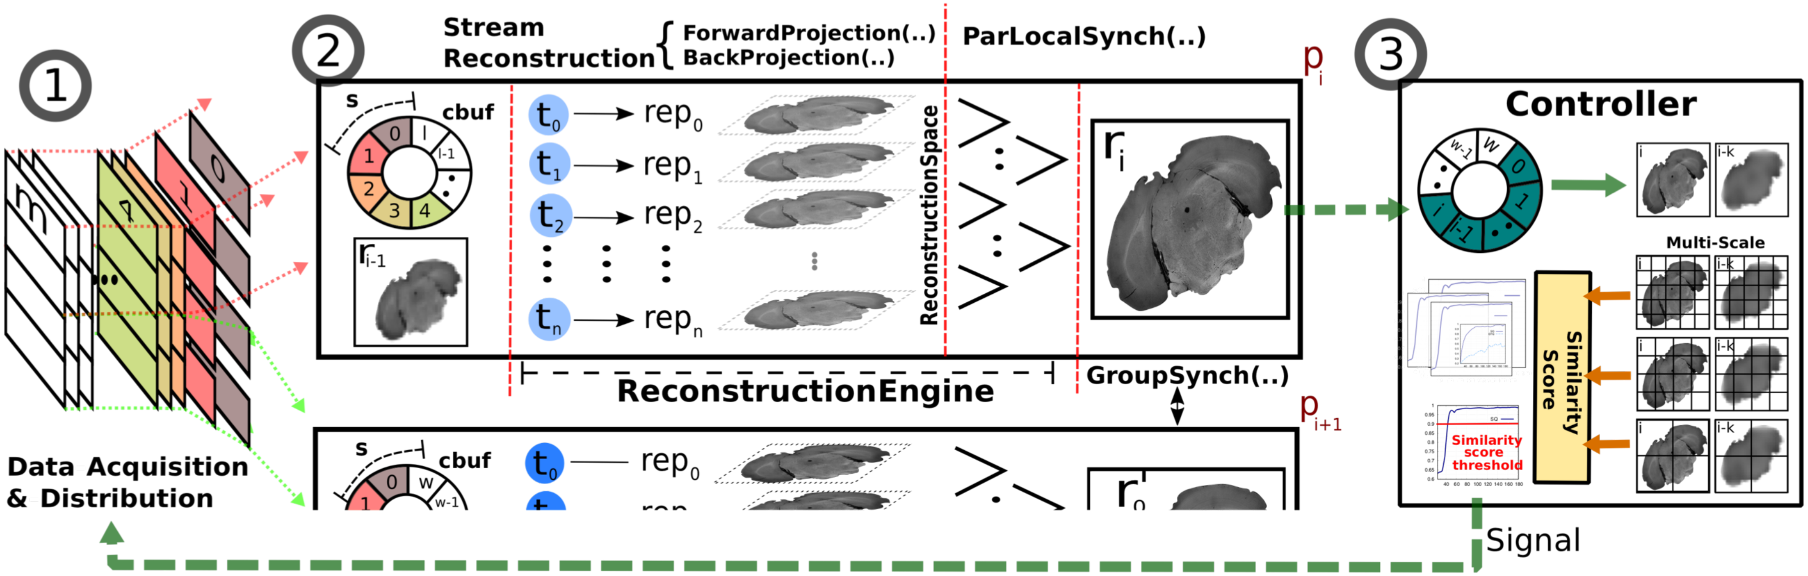

Autonomous stream processing system that allows data streamed from beamline computers to be processed in real time on a remote supercomputer with a control feedback loop used to make decisions during experimentation

– Reduce data acquisition time by 22–44% for the datasets considered in our experiments

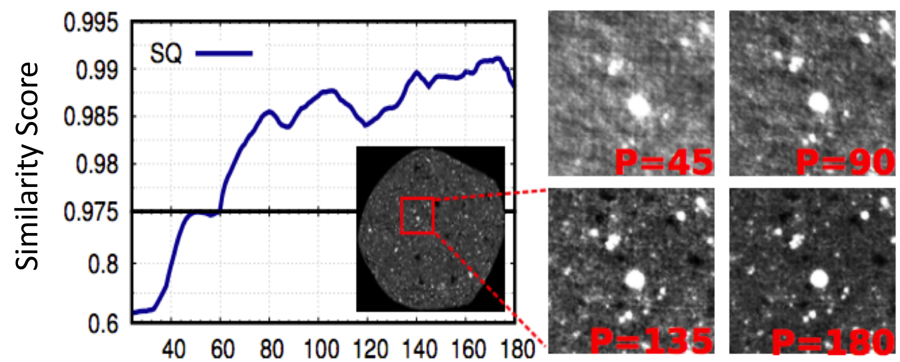

Image Quality with respect to Streamed Projections

Autonomous power module for DTN

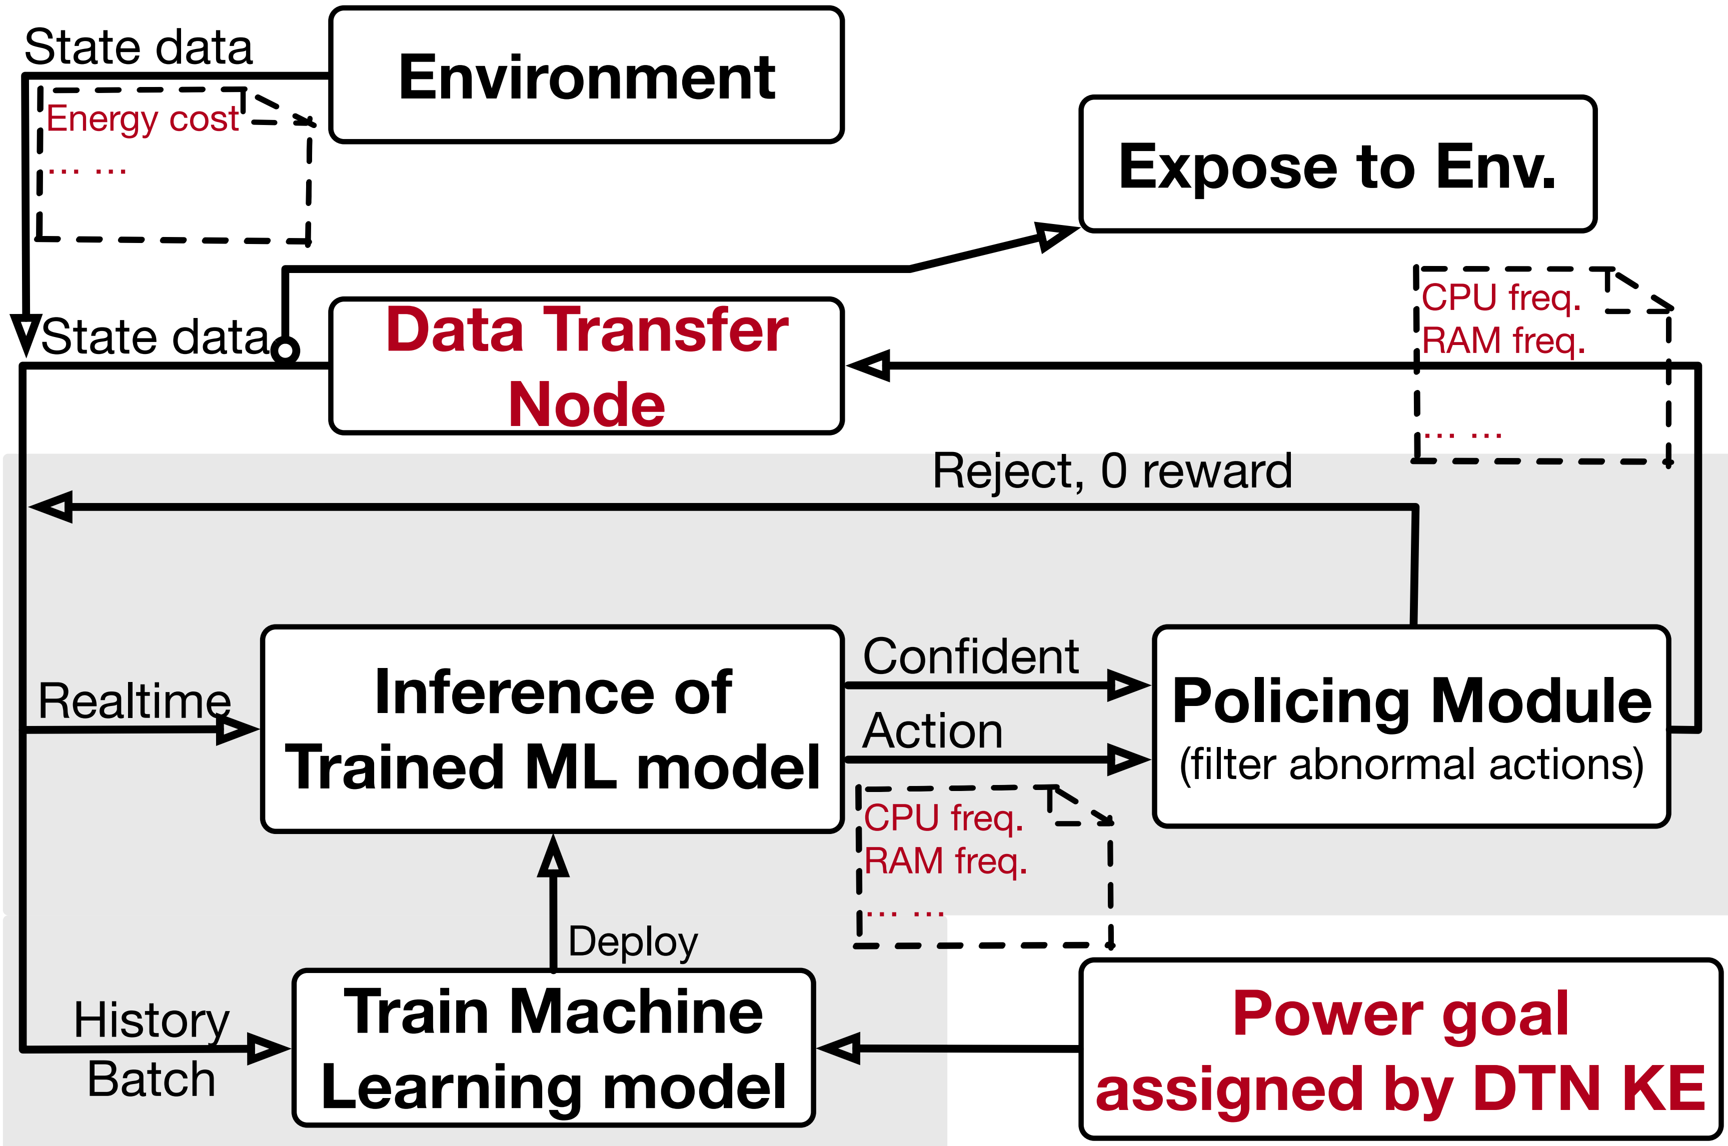

Autonomous DTN

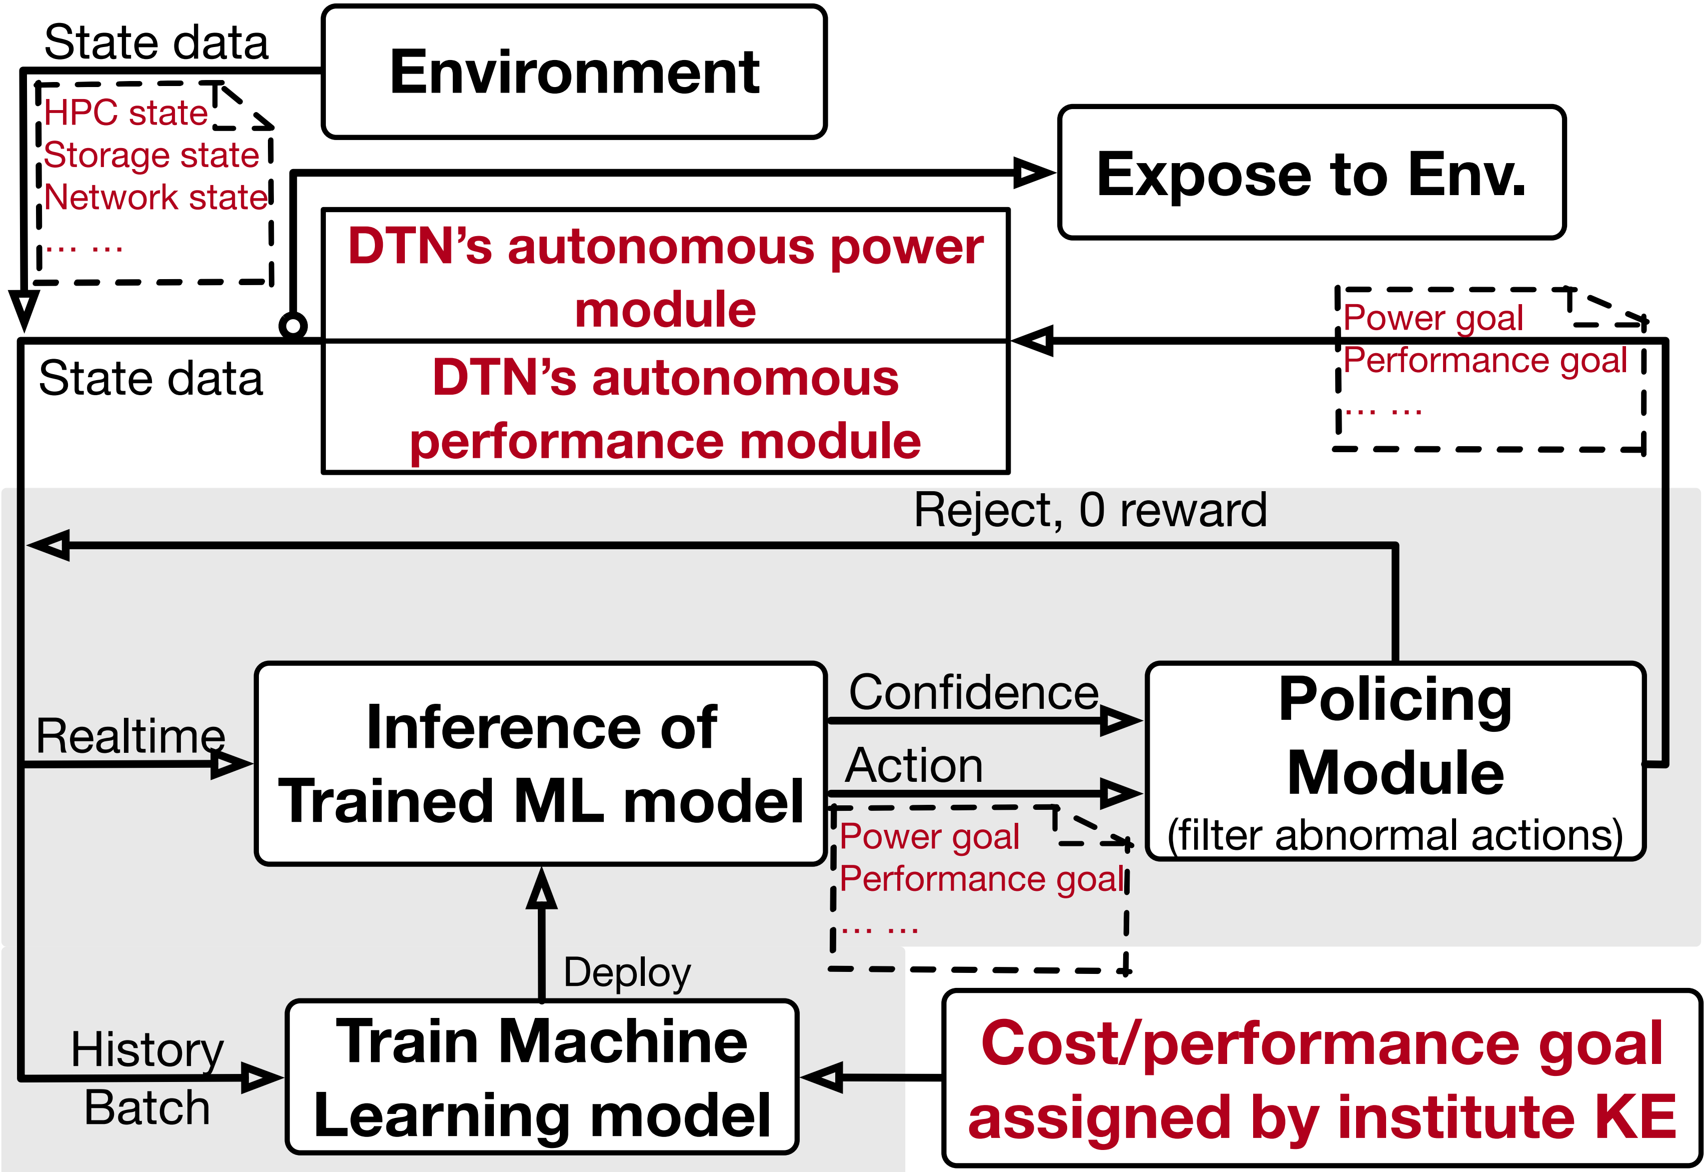

Questions