Failure Prediction and Fault Localization in Cloud Environment

Cristina Monni

Joint work with:
Leonardo Mariani, Mauro Pezzè, Oliviero Riganelli, Rui Xin

USI

UNIVERSITA' DEGLI STUDI DI MILANO BICOCCA

Cloud Computing

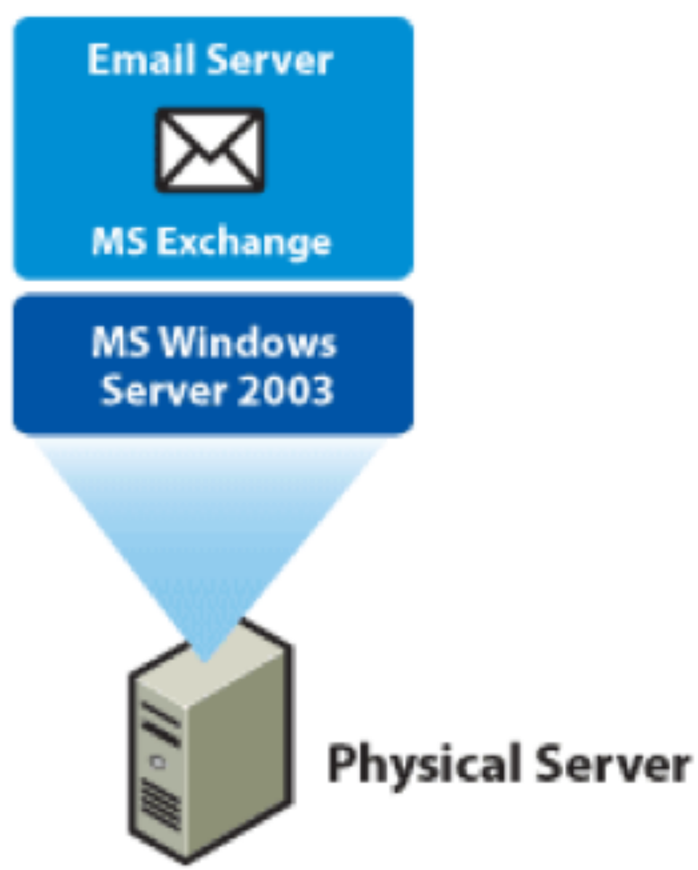

Traditional

1 Application - 1 Server

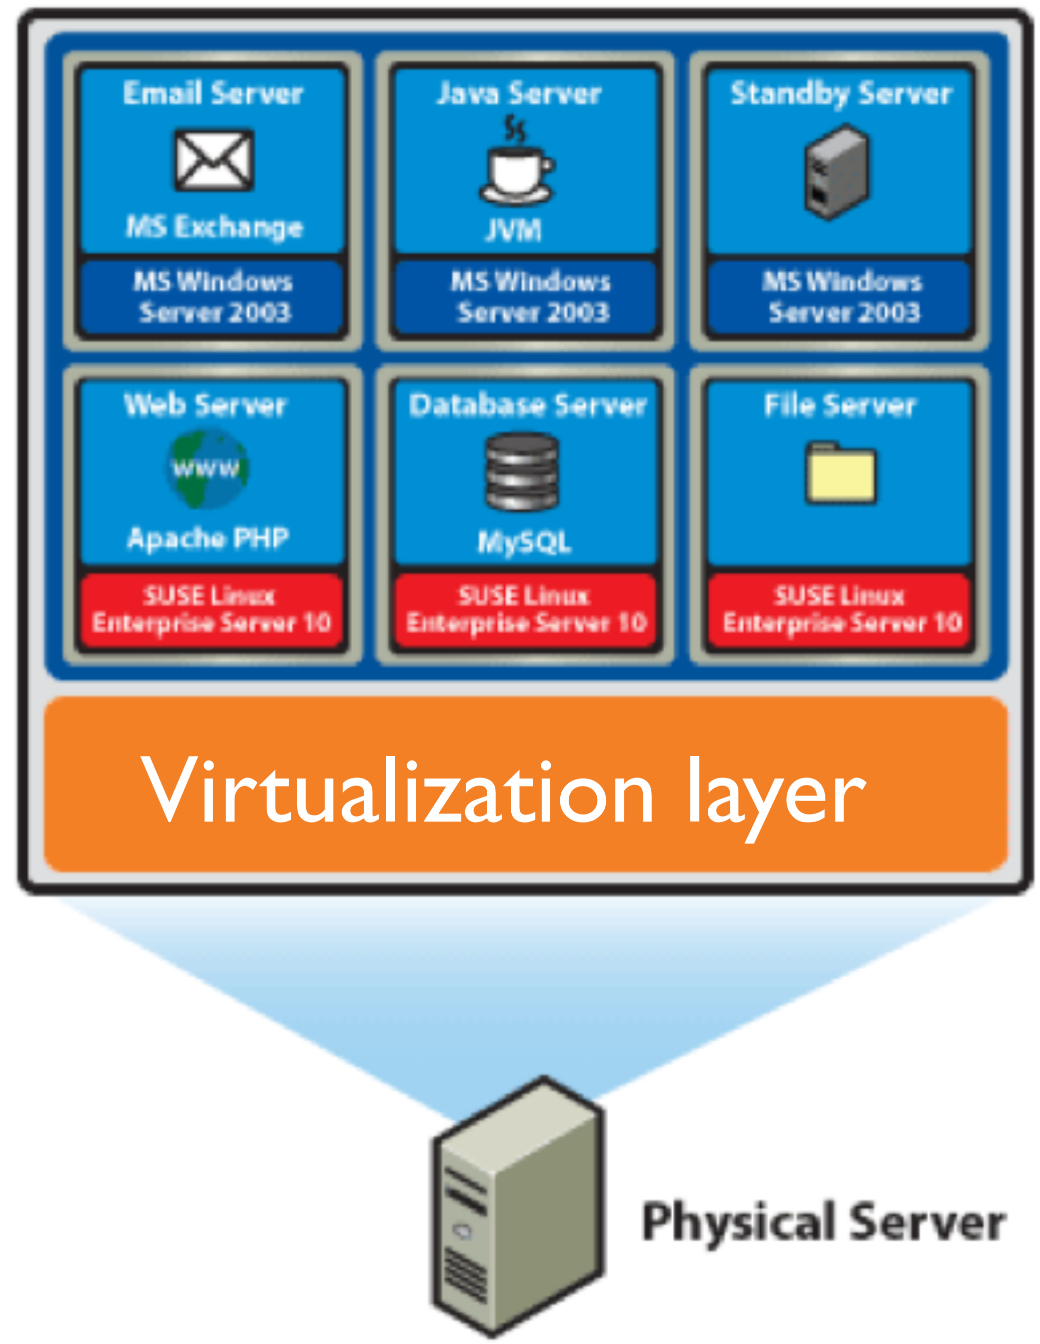

Cloud

Multiple Virtual Server

Cloud Comput

Microsoft Azure Suffers Global Outage

Michael Moore, November 19, 2014, 10:52 am

Google Suffers Compute Engine Cloud Outage

Ben Sullivan, August 17, 2015, 2:38 pm

DNS Server Error Brings Down iTunes, iCloud for 12 Hours

By Chris Preimesberger | Posted 2015-03-11 Print

are 76 Email

ıpacted sales on the iTunes Store, iBooks

AWS Suffers Another Cloud Outage

Ben Sullivan, August 10, 2015, 3:15 pm

Verizon Cloud Returns After 40-Hour Shutdown

Ben Sullivan, January 12, 2015, 8:40 am

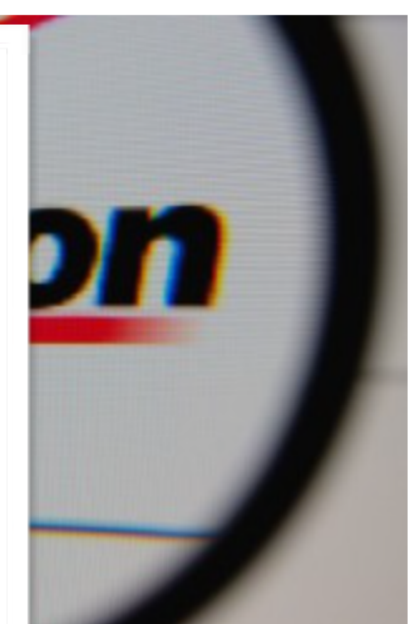

All

CIO, Deve

On jobs.tec

Twitter Suffers Worldwide Outage

Ben Sullivan, January 19, 2016, 9:11 am

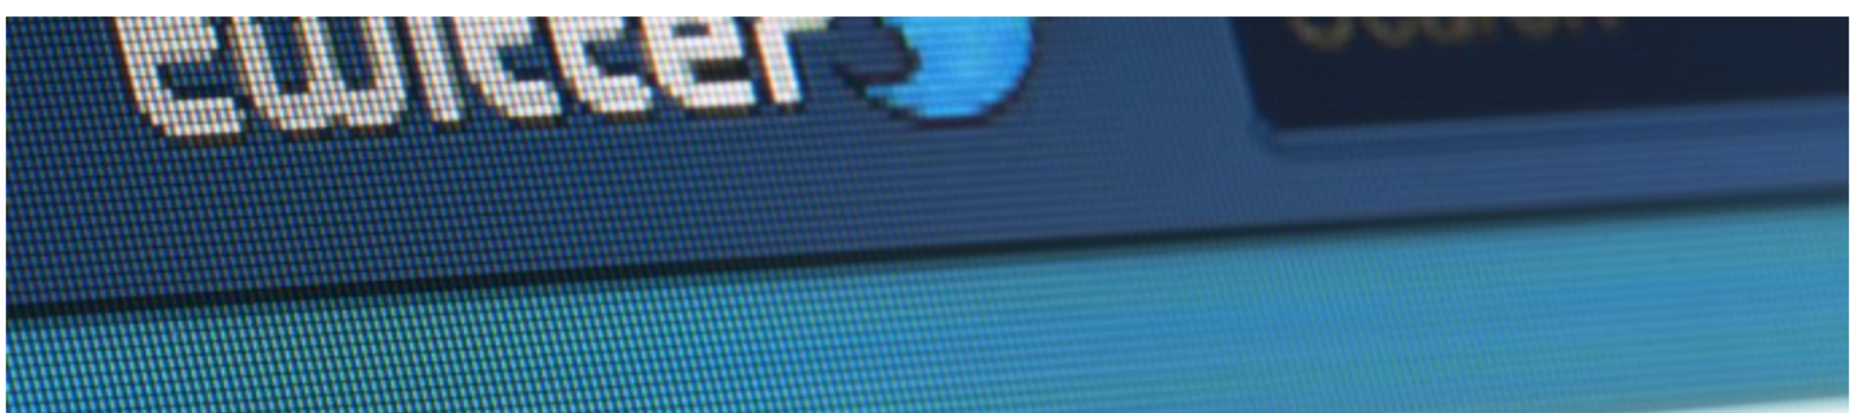

Proactive Fault Management

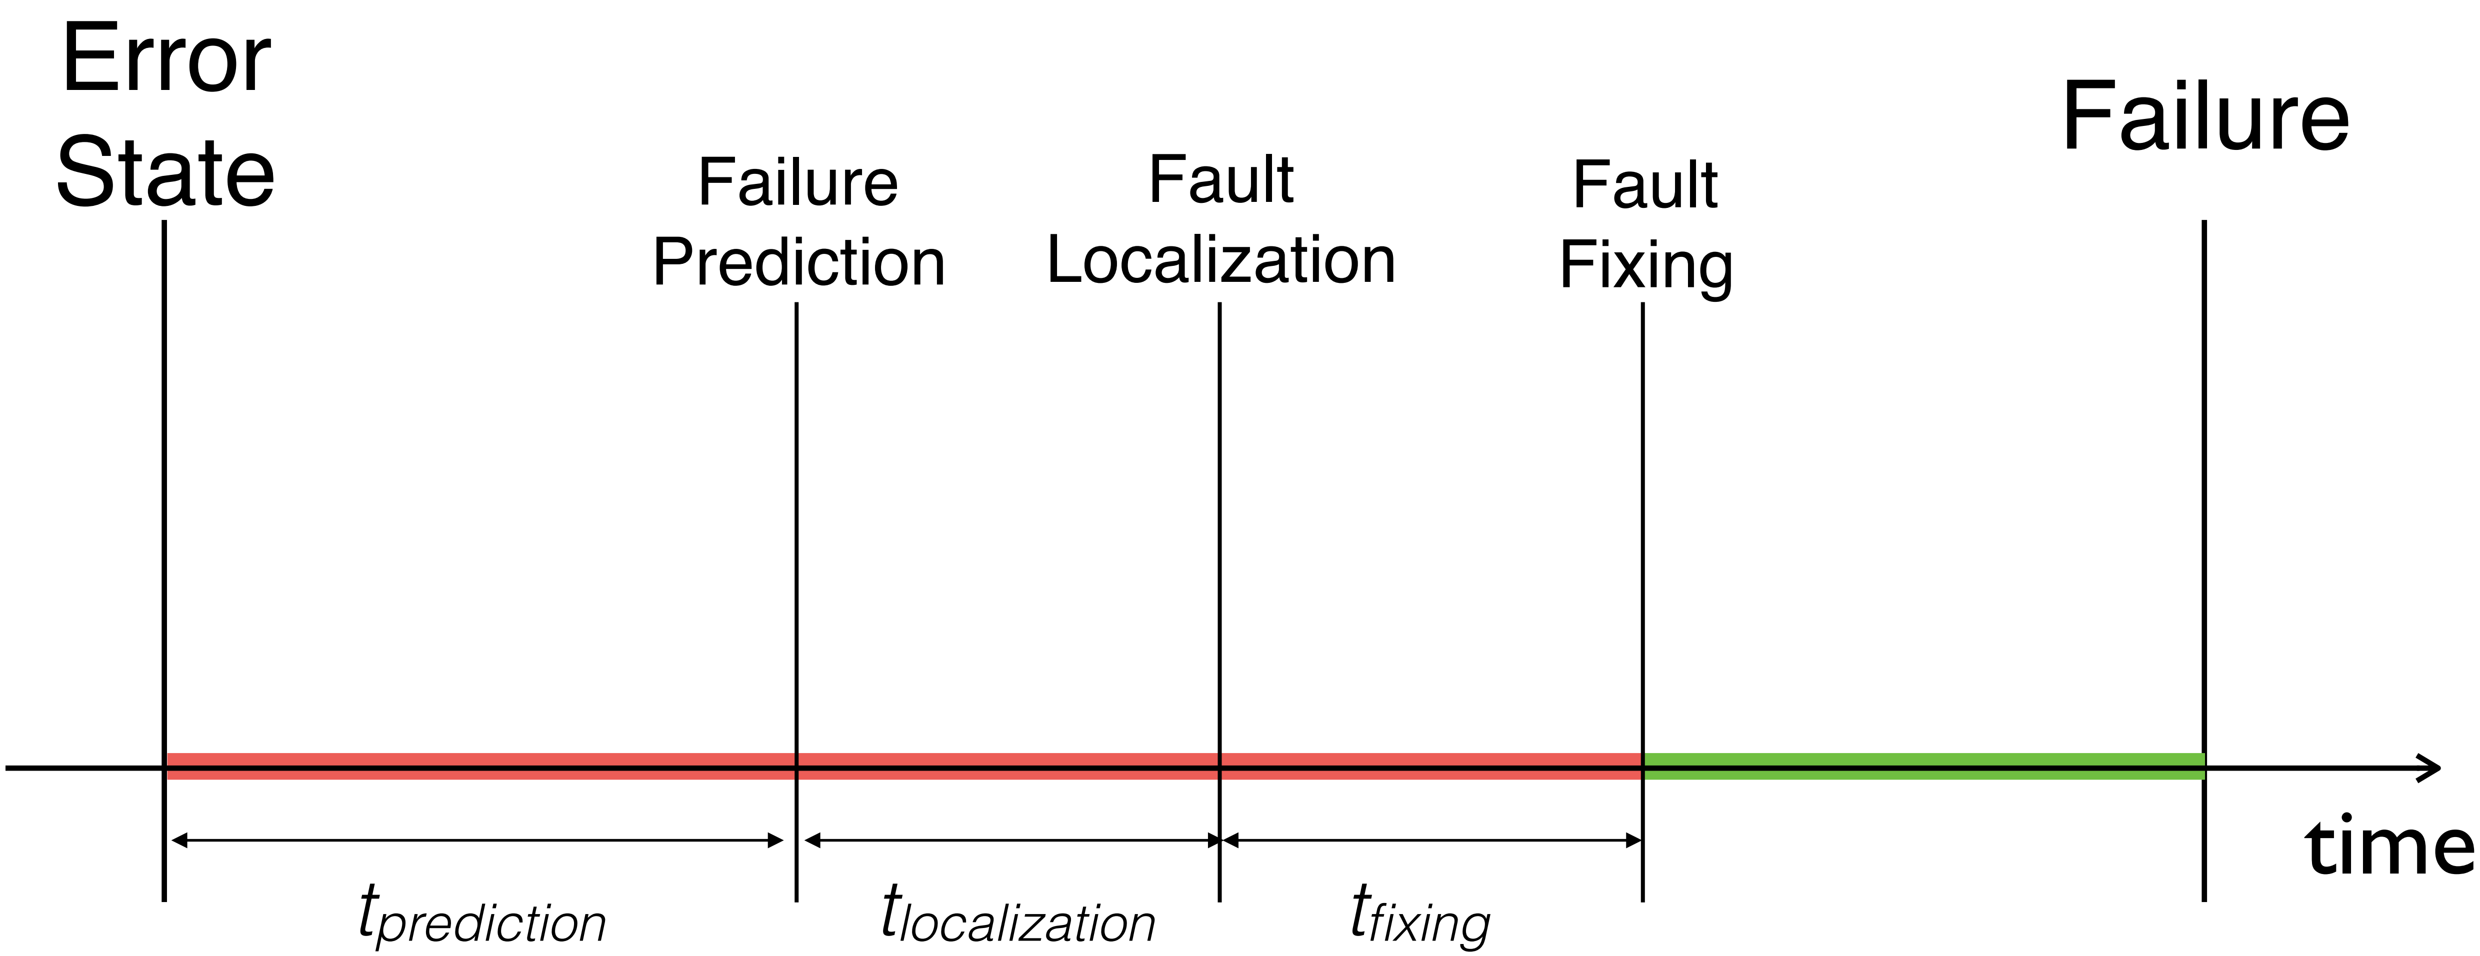

Goal 1

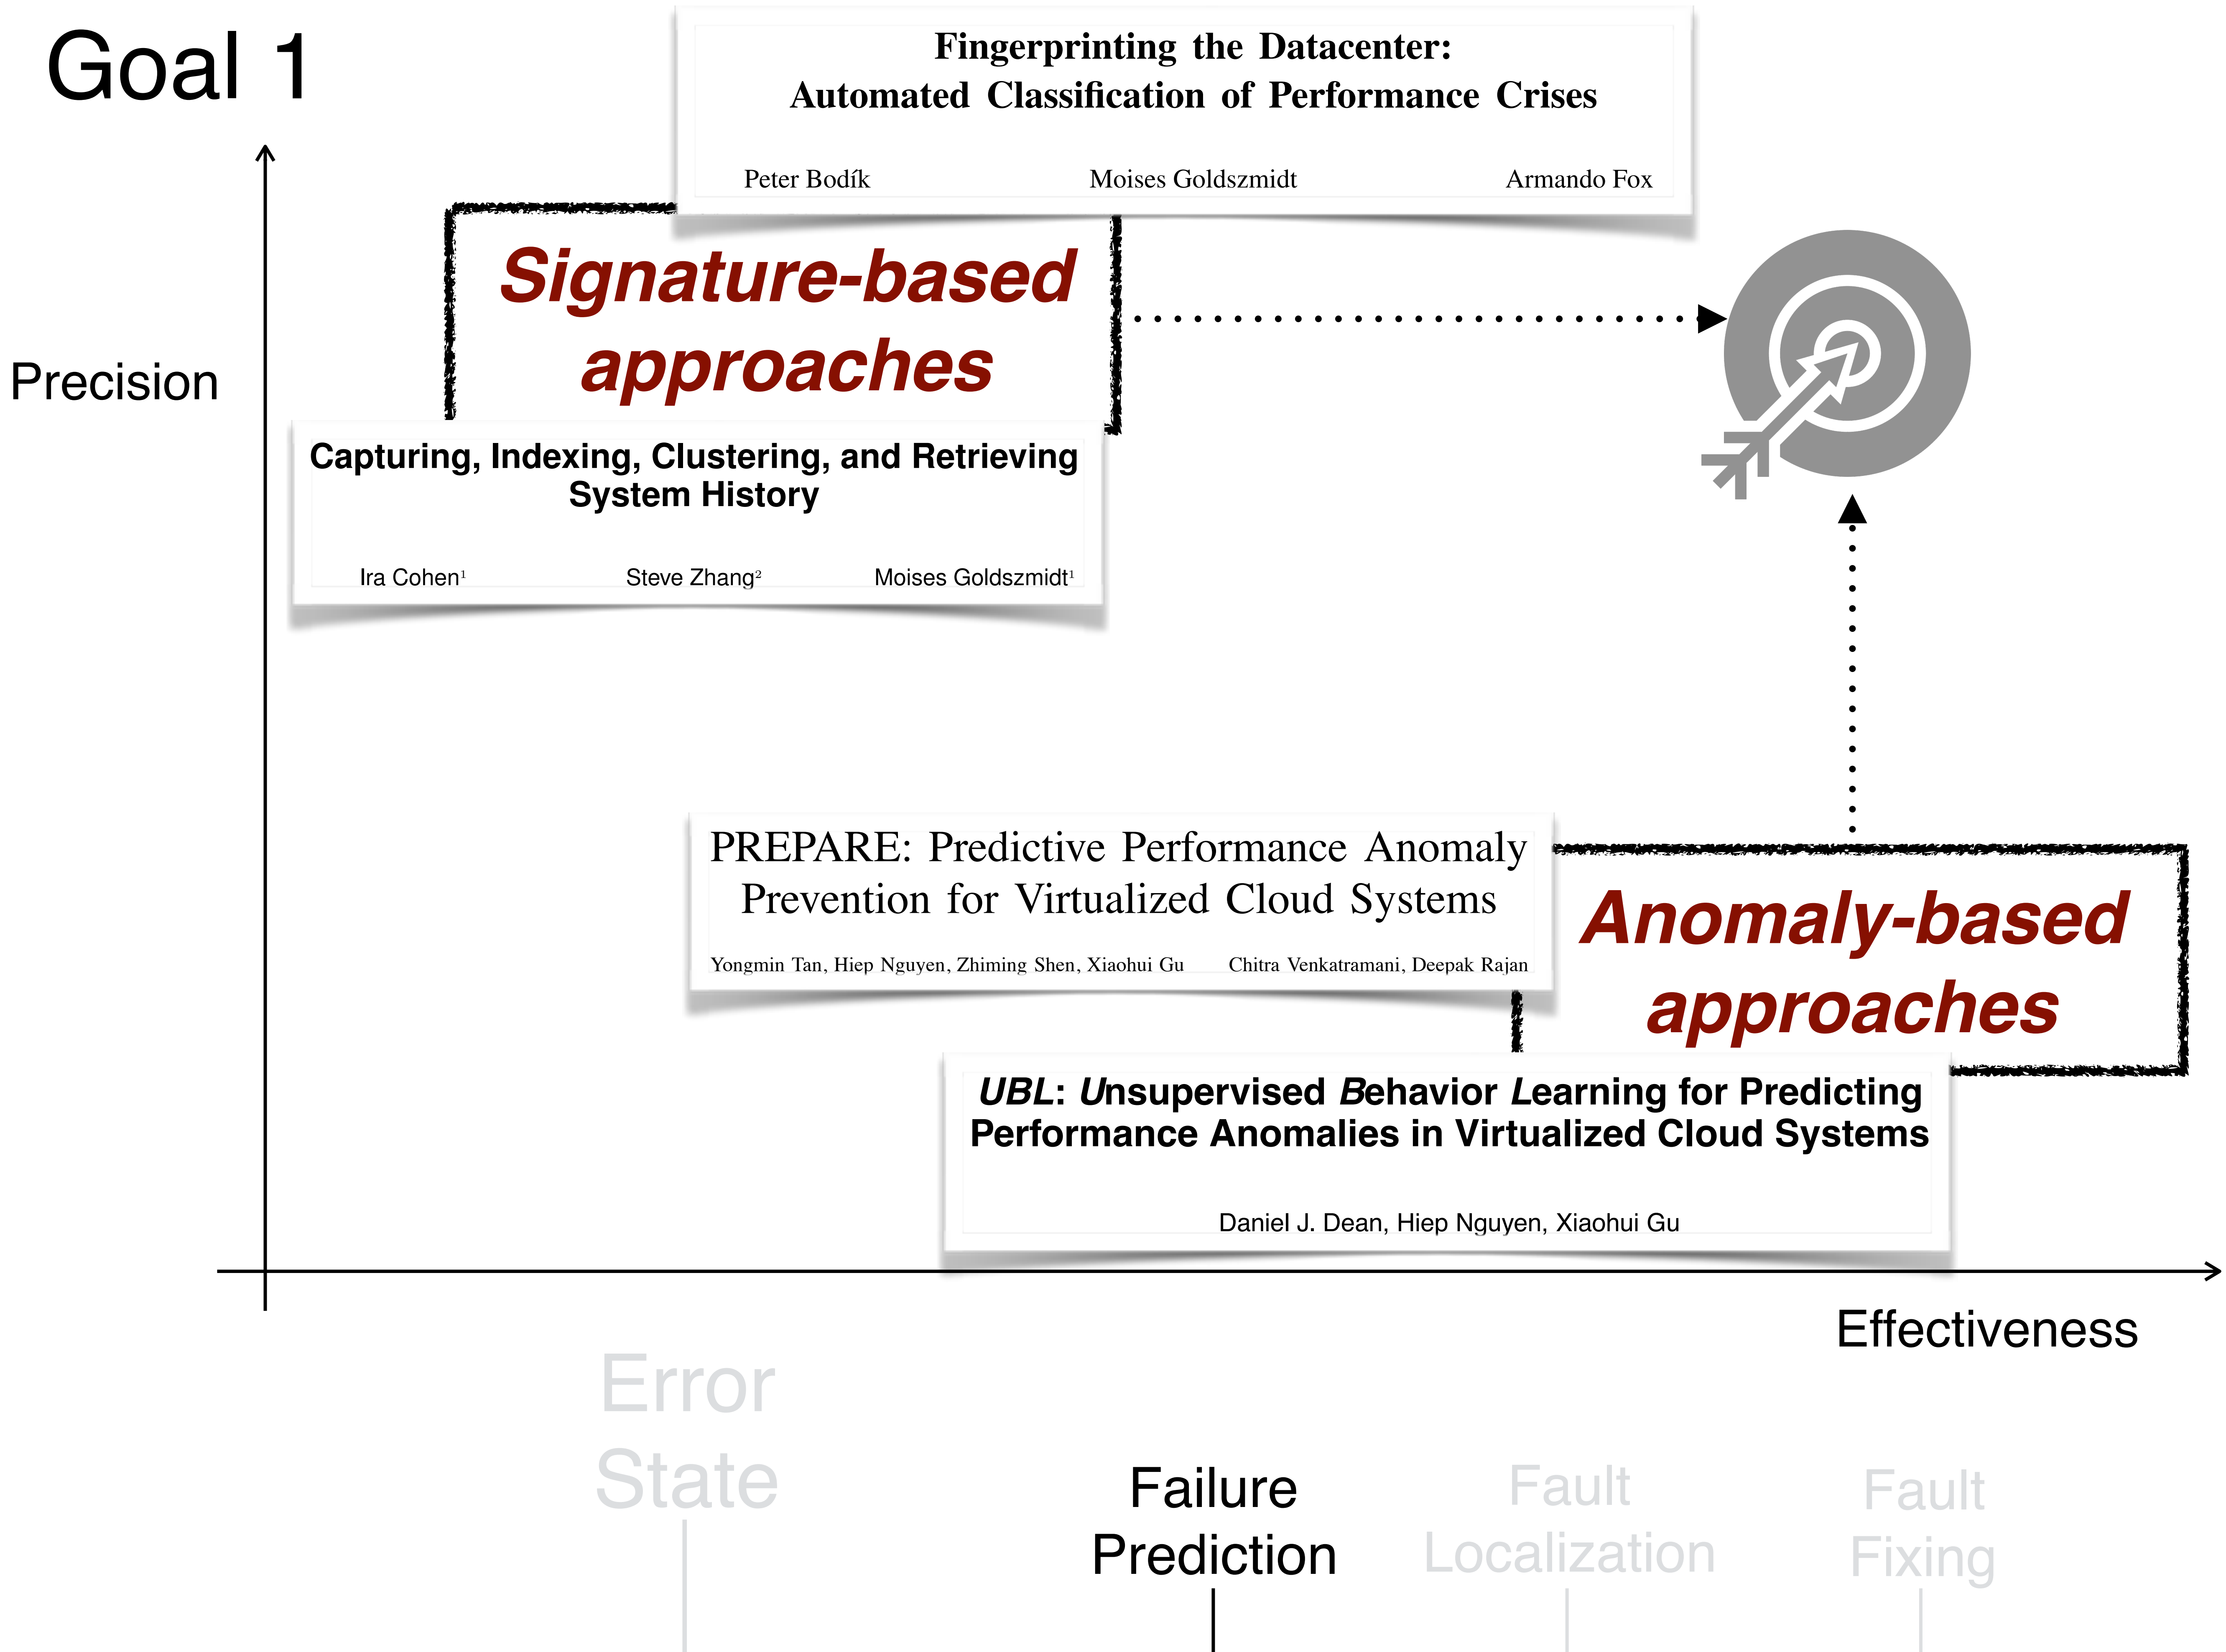
Fingerprinting the Datacenter:
Automated Classification of Performance Crises
Peter Bodík
Moises Goldszmidt
Armando Fox
Signature-based approaches
Precision
Capturing, Indexing, Clustering, and Retrieving System History
Ira Cohen[1]
Steve Zhang[2]
Moises Goldszmidt[1]
PREPARE: Predictive Performance Anomaly Prevention for Virtualized Cloud Systems
Yongmin Tan, Hiep Nguyen, Zhiming Shen, Xiaohui Gu
Chitra Venkatramani, Deepak Rajan
Anomaly-based approaches
UBL: Unsupervised Behavior Learning for Predicting Performance Anomalies in Virtualized Cloud Systems
Daniel J. Dean, Hiep Nguyen, Xiaohui Gu
Effectiveness
Error State
Failure Prediction
Fault Localization
Fault Fixing

Goal 2

Precision

Ensembles of Models for Automated Diagnosis of System Performance Problems

Steve Zhang[1], Ira Cohen, Moises Goldszmidt, Julie Symons, Armando Fox[1]

Probabilistic approaches

Scalable Near Real-Time Failure Localization of Data Center Networks

Herodotos Herodotou
Microsoft Research
herohero@microsoft.com

Bolin Ding
Microsoft Research
bolin.ding@microsoft.com

Shobana Balakrishnan
Microsoft Research
shobanab@microsoft.com

Google

Localizing packet loss

In a large complex network

Nicolas Guilbaud nguilbaud@google.com
Ross Cartlidge rossc@google.com

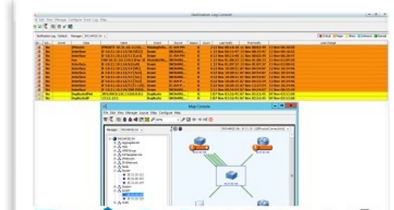

DELL EMC

Graph-based approaches

Virtual Network Diagnosis as a Service

Wenfei Wu, Guohui Wang, Aditya Akella, Anees Shaikh

Efficiency

Error State

Failure Prediction

Fault Localization

Fault Fixing

Failure

Failure Prediction and Fault Localization

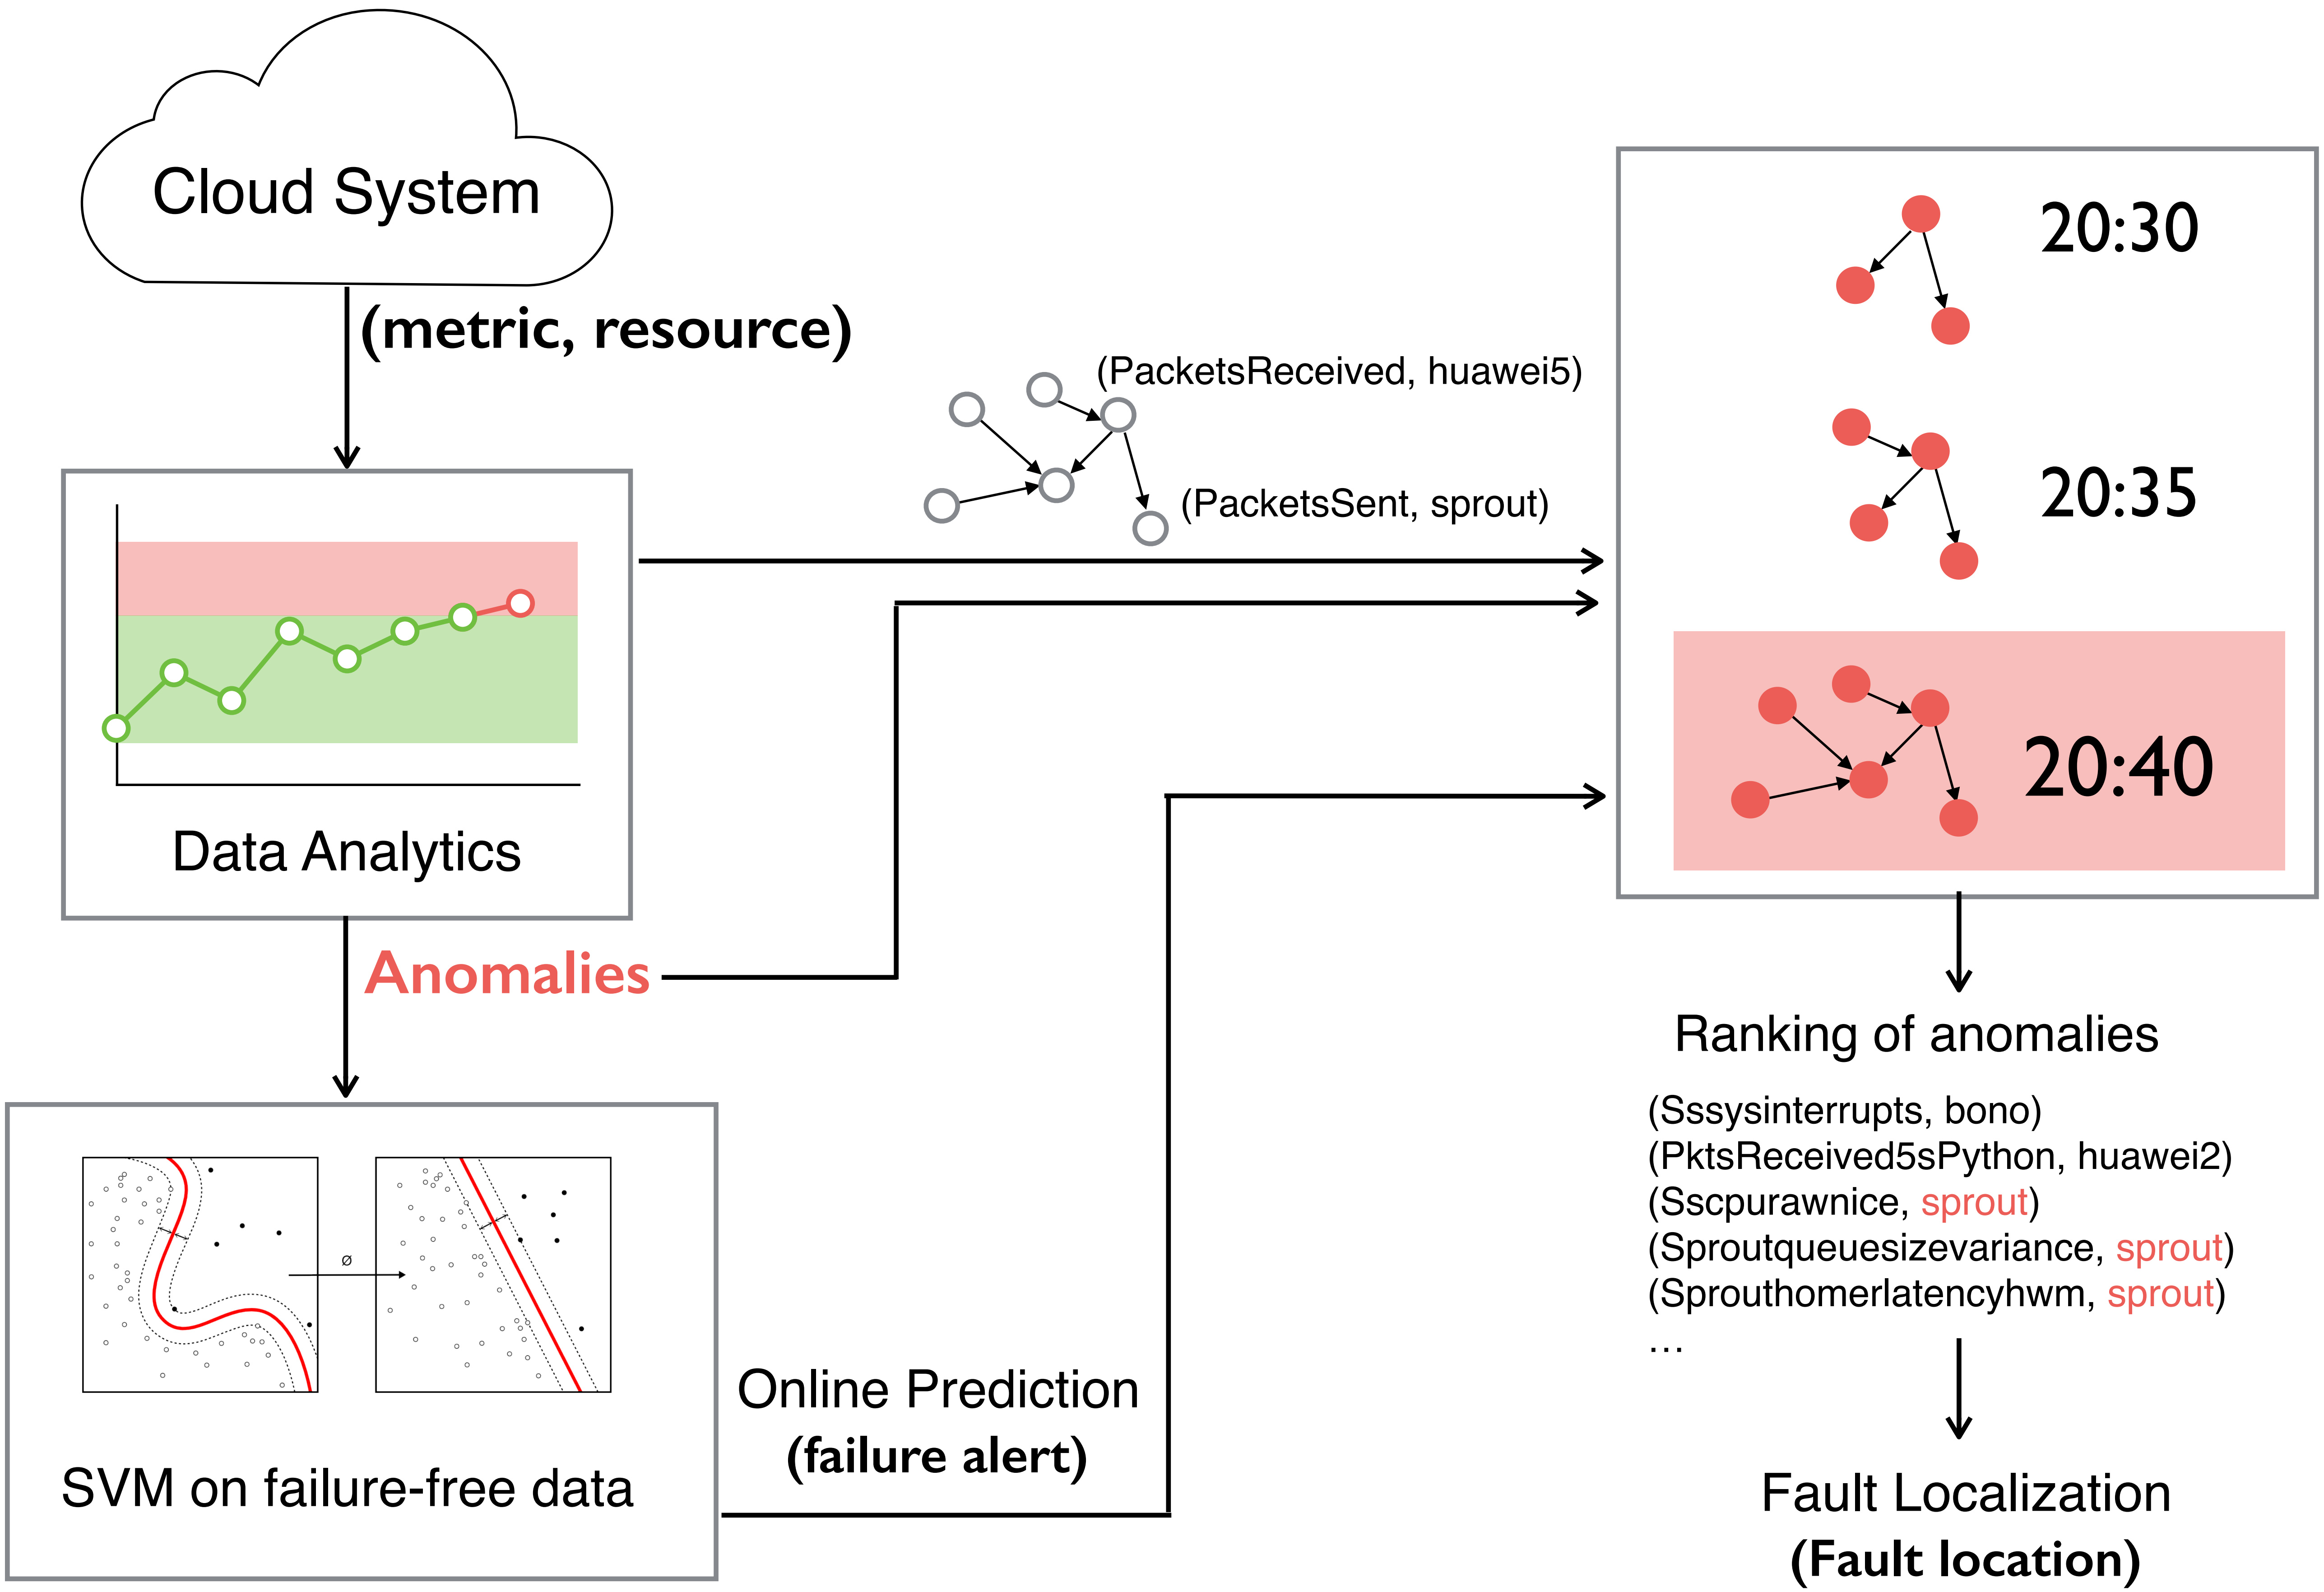

Failure Prediction and Fault Localization

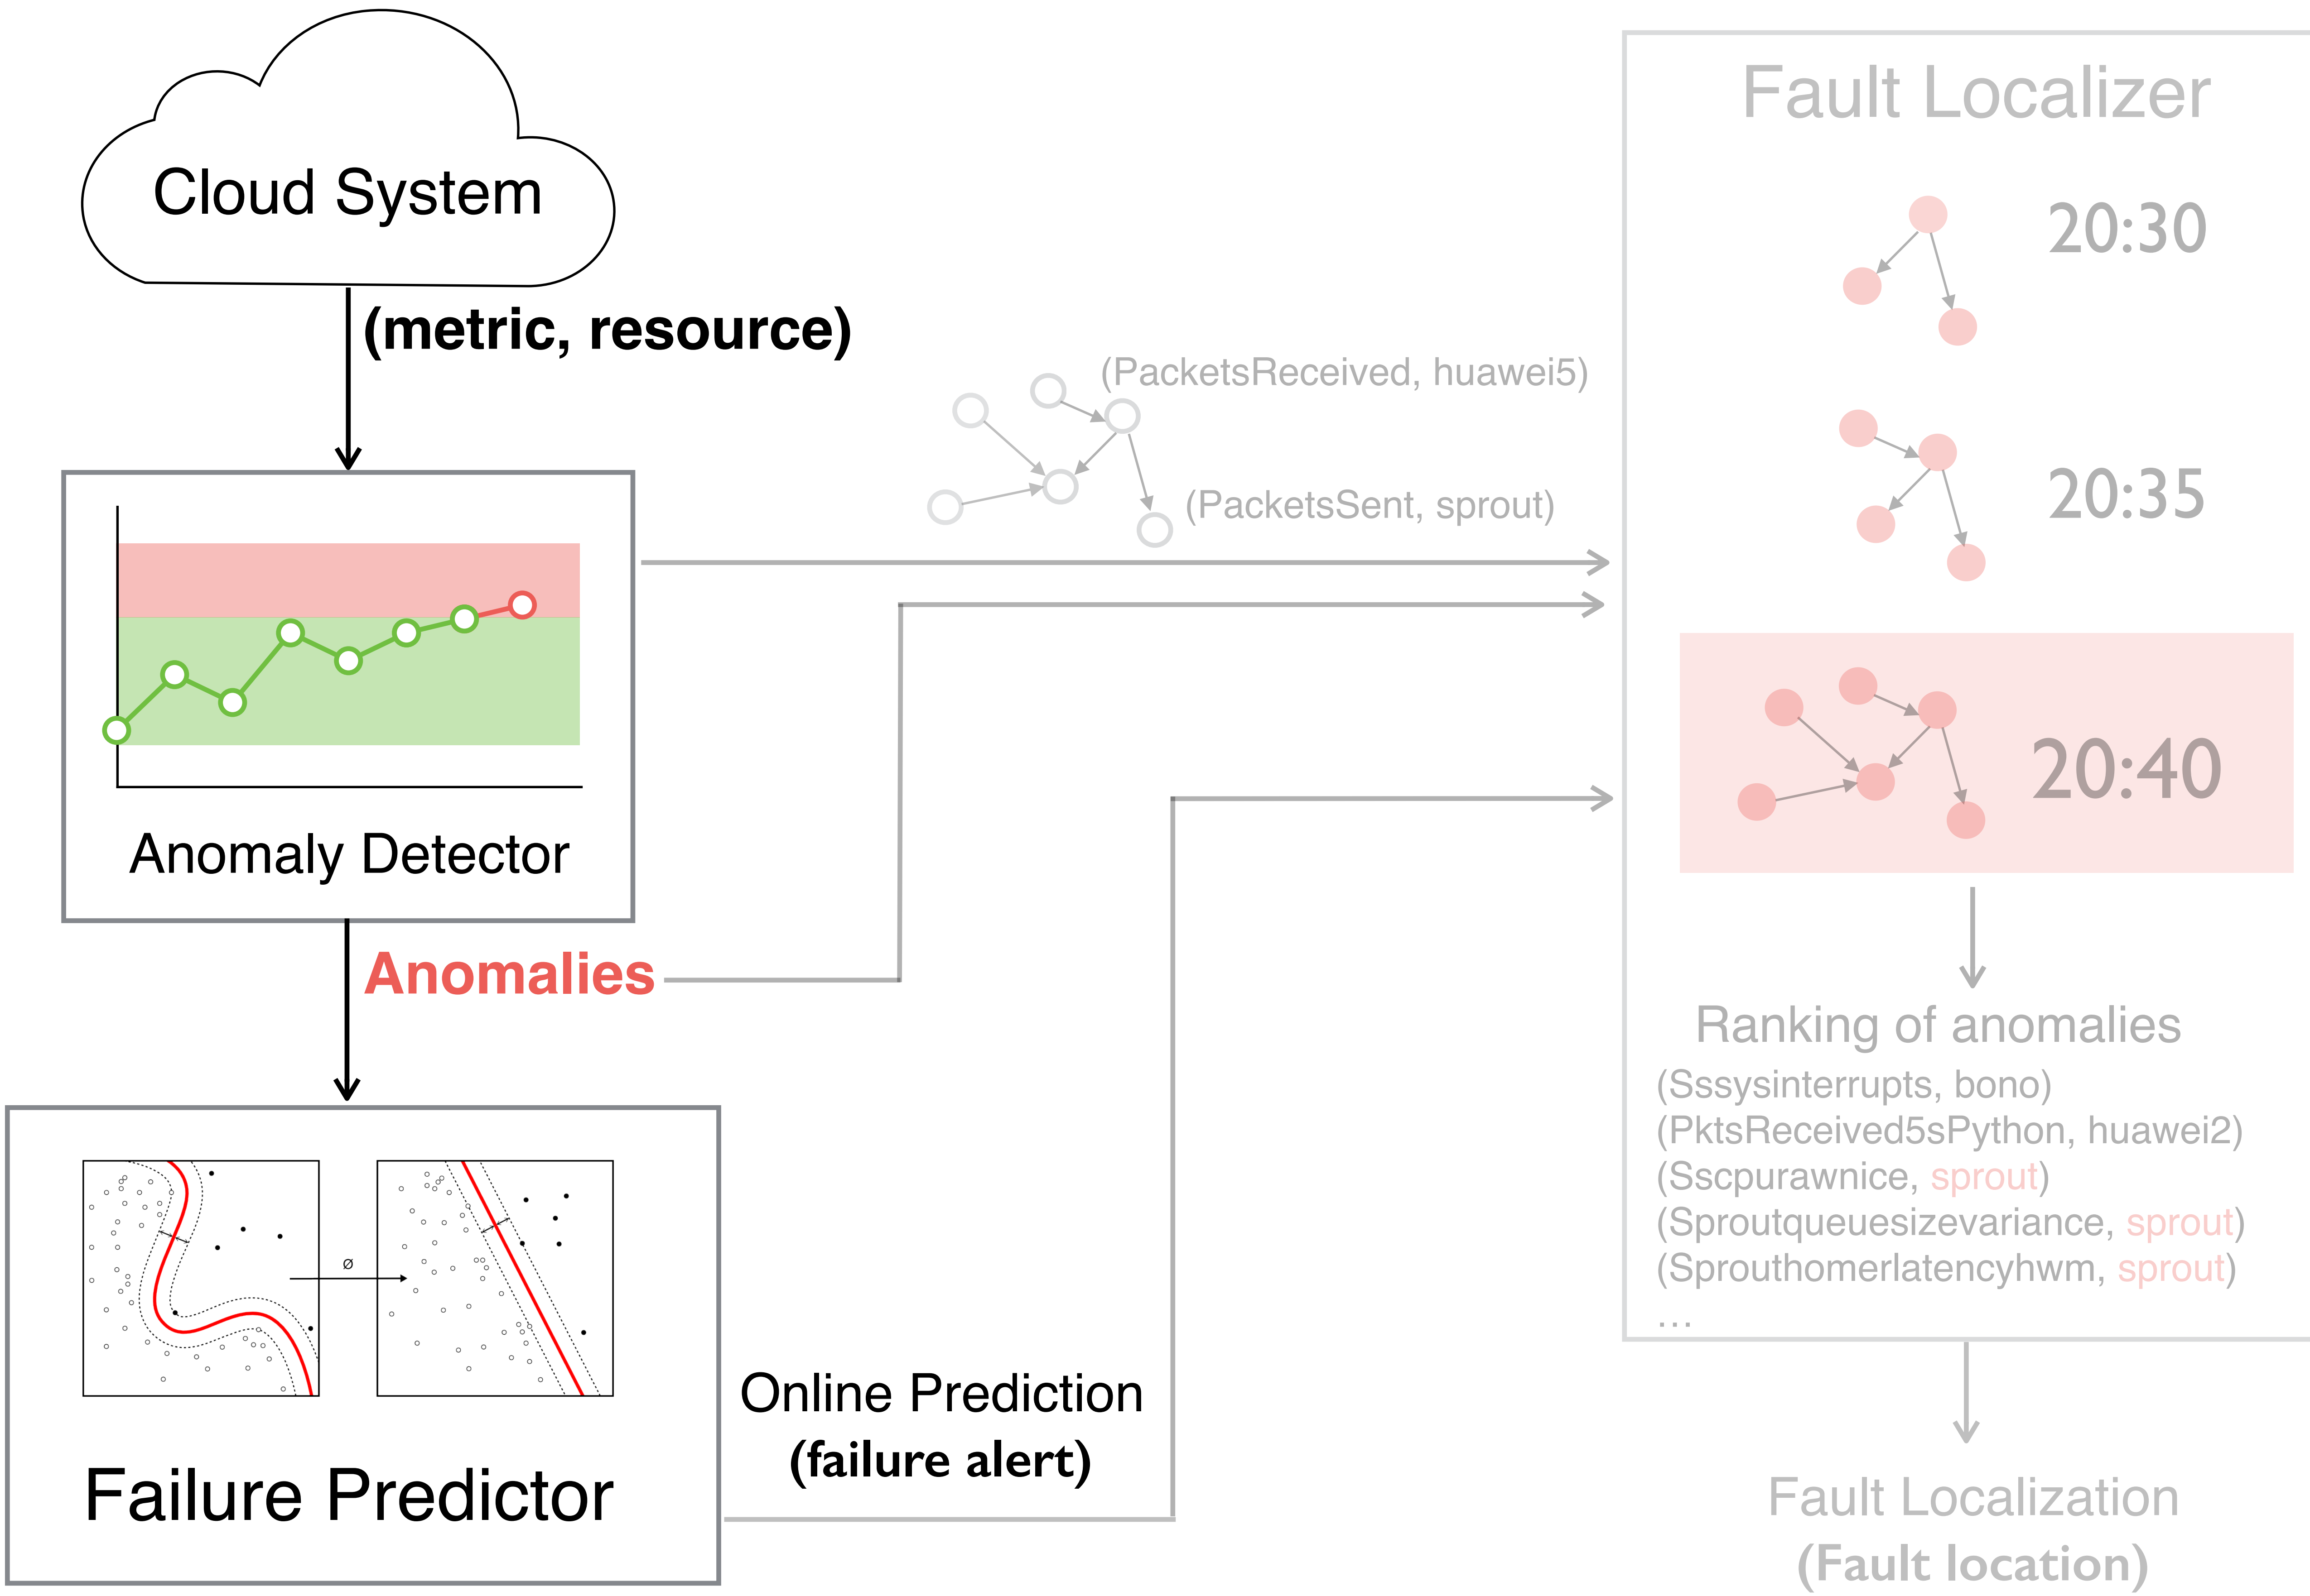

Failure Prediction

Anomaly Detector

Time Series Analysis

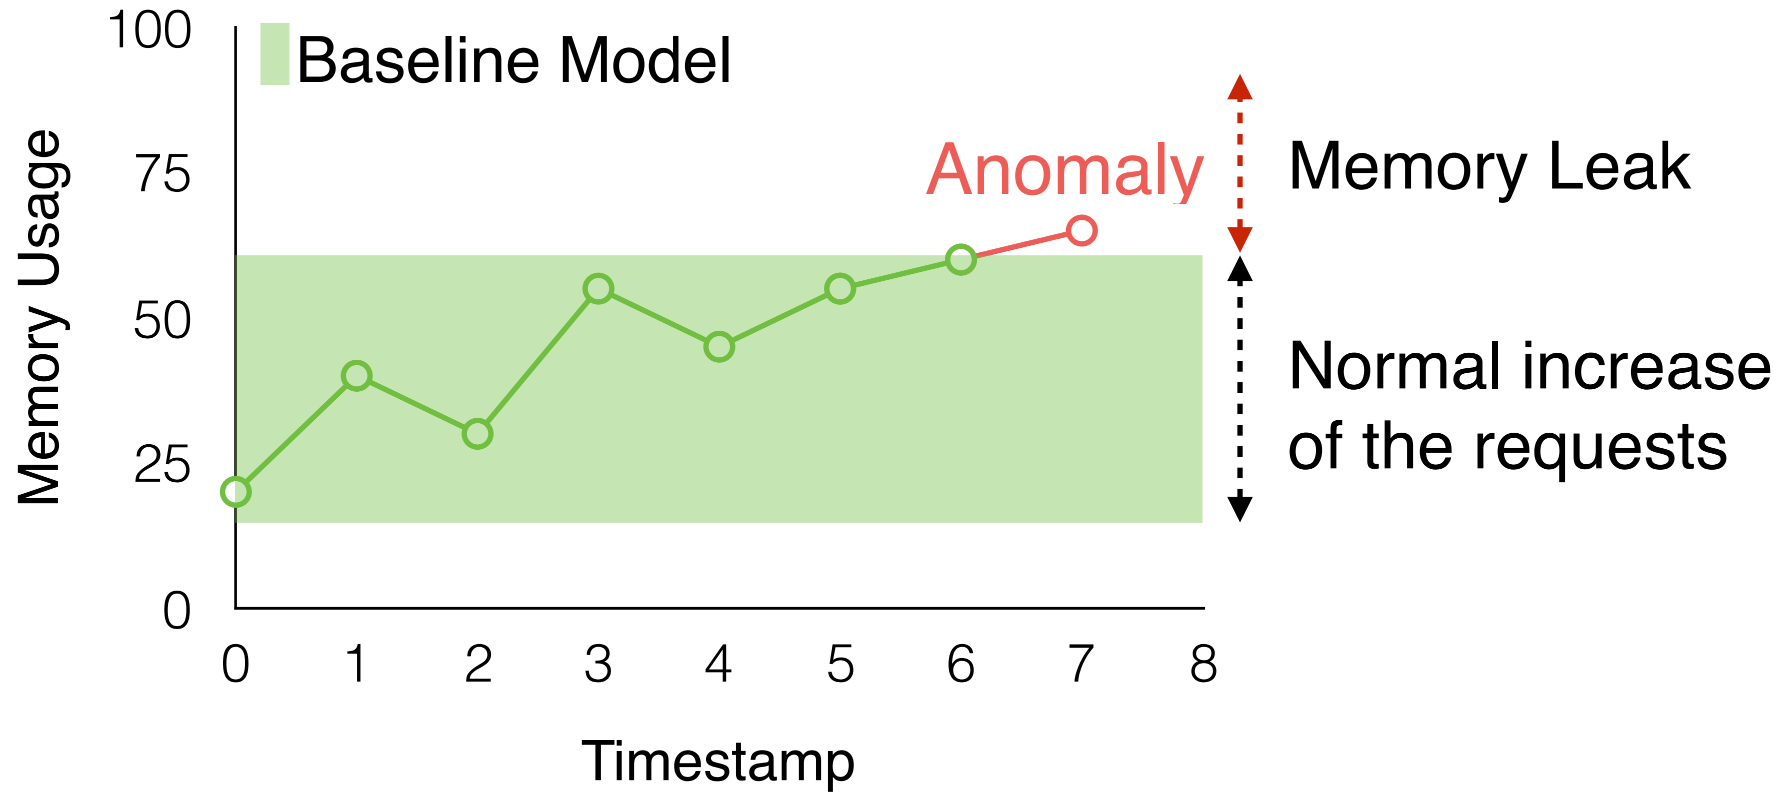

Failure Predictor

SVM with Gaussian Kernel

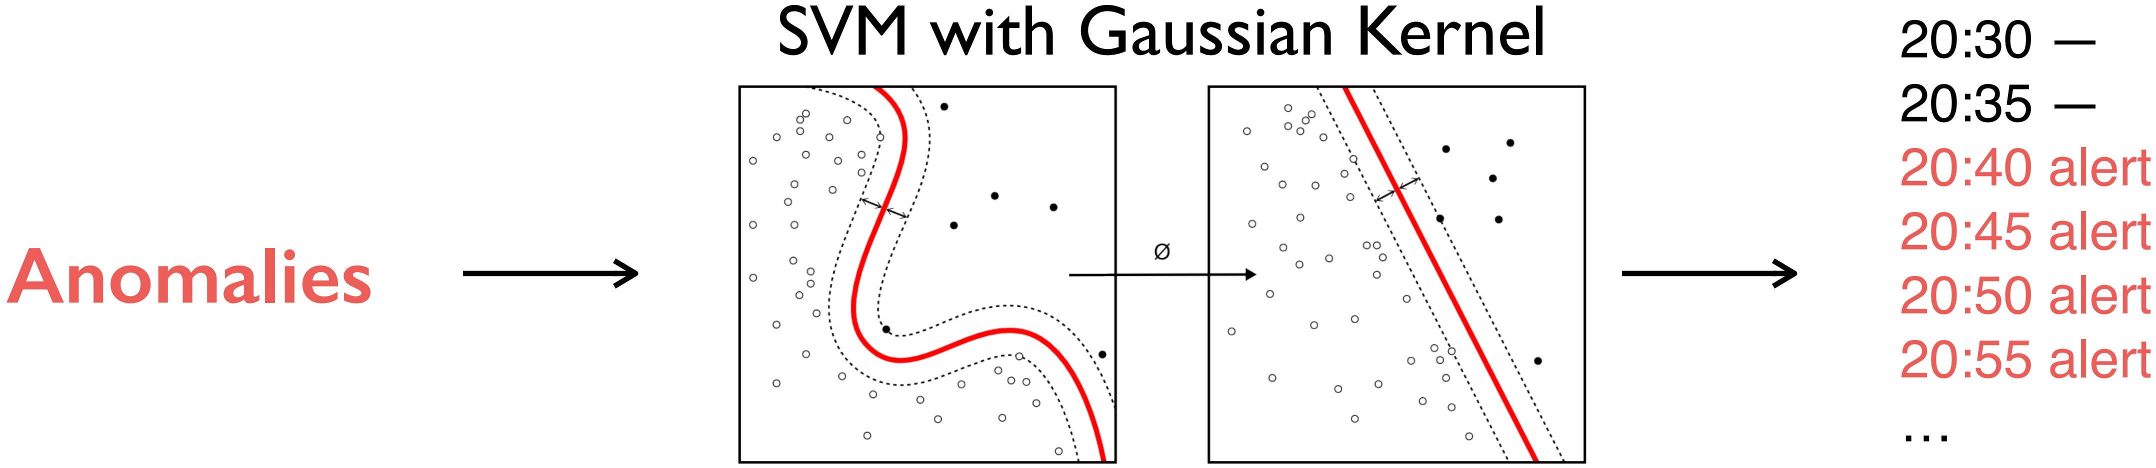

Failure Prediction and Fault Localization

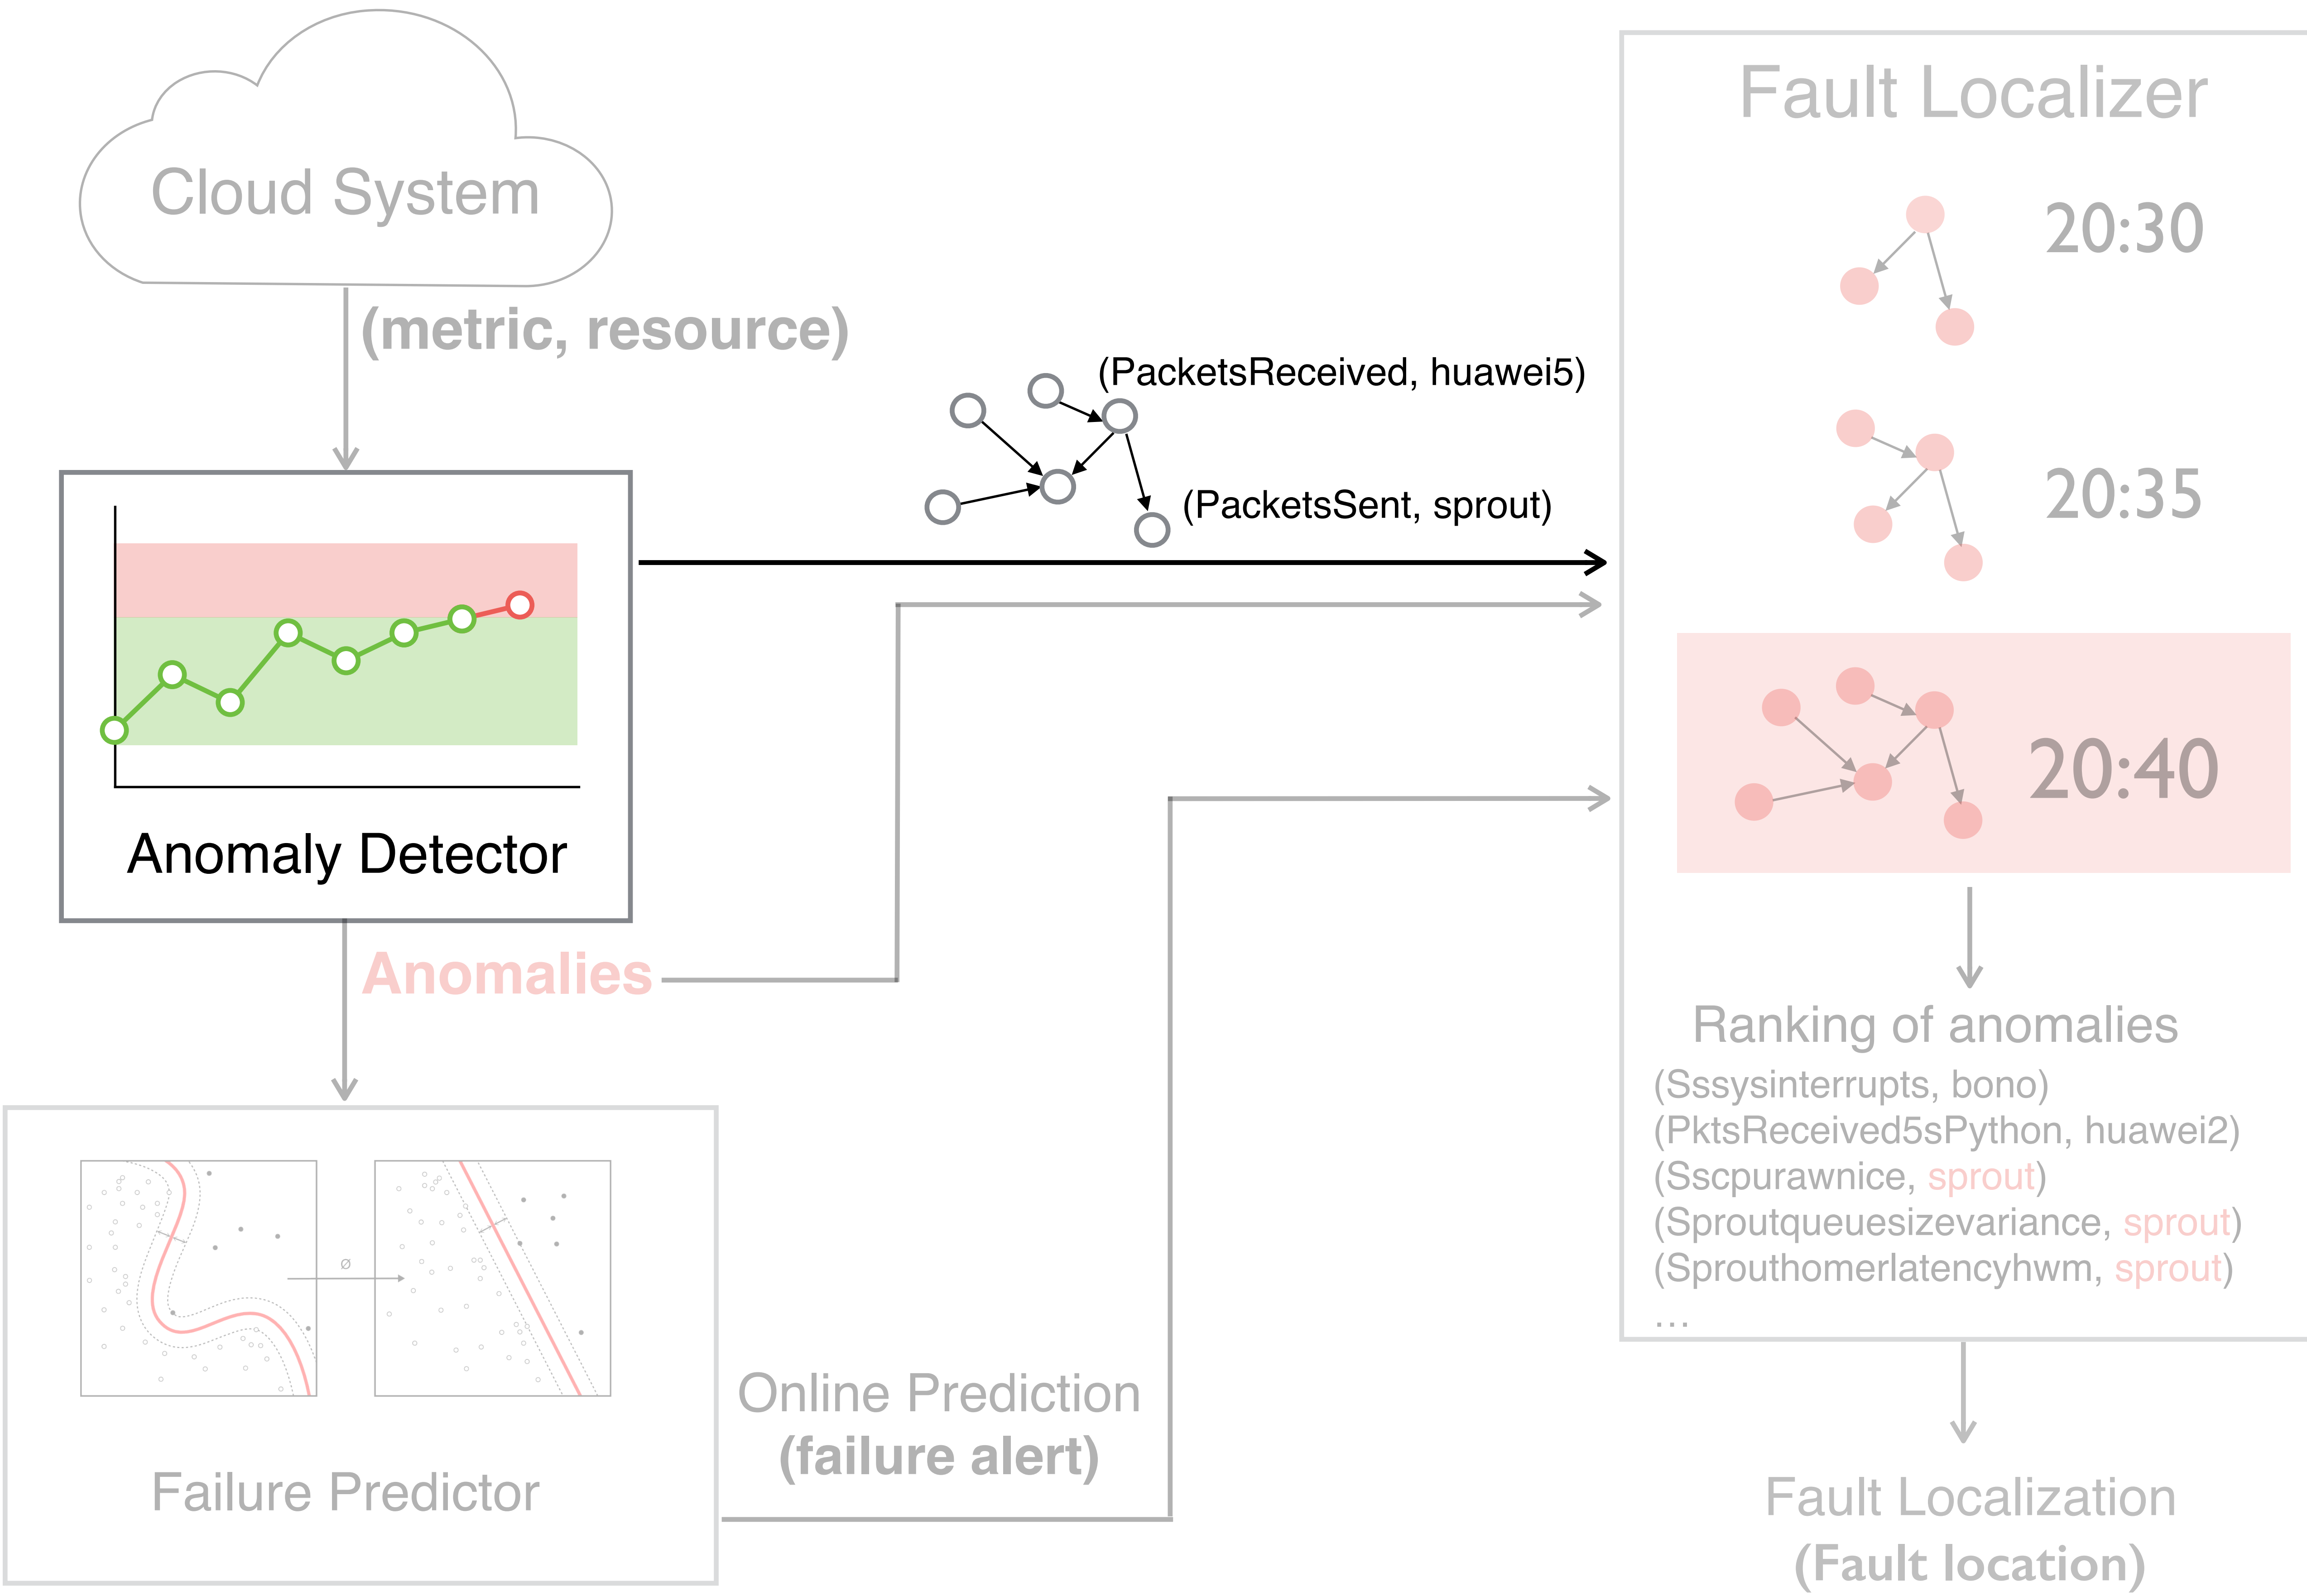

Fault Localization

Causality graph

- Built from the baseline model with machine learning
- Vertices: (metric, resource)
- Edges: correlations (Granger causality)
- Edges labels: weight probabilities

$$\text{KPI} = (\text{metric}, \text{resource})$$

vertex	KPI
1	$\langle$AvgReqQueueLength, bono$\rangle$
2	$\langle$AvgWaitTime, bono$\rangle$
3	$\langle$Bonoconnectedclients.0, bono$\rangle$
4	$\langle$Bonolatencyaverage, bono$\rangle$
5	$\langle$Bonolatencycount, bono$\rangle$
6	$\langle$Bonolatencyhwm, bono$\rangle$
7	$\langle$Bonolatencylwm, bono$\rangle$
8	$\langle$Bonolatencyvariance, bono$\rangle$
9	$\langle$Bonoqueuesizehwm, bono$\rangle$
10	$\langle$Bonoqueuesizelwm, bono$\rangle$
11	$\langle$Bonoqueuesizevariance, bono$\rangle$

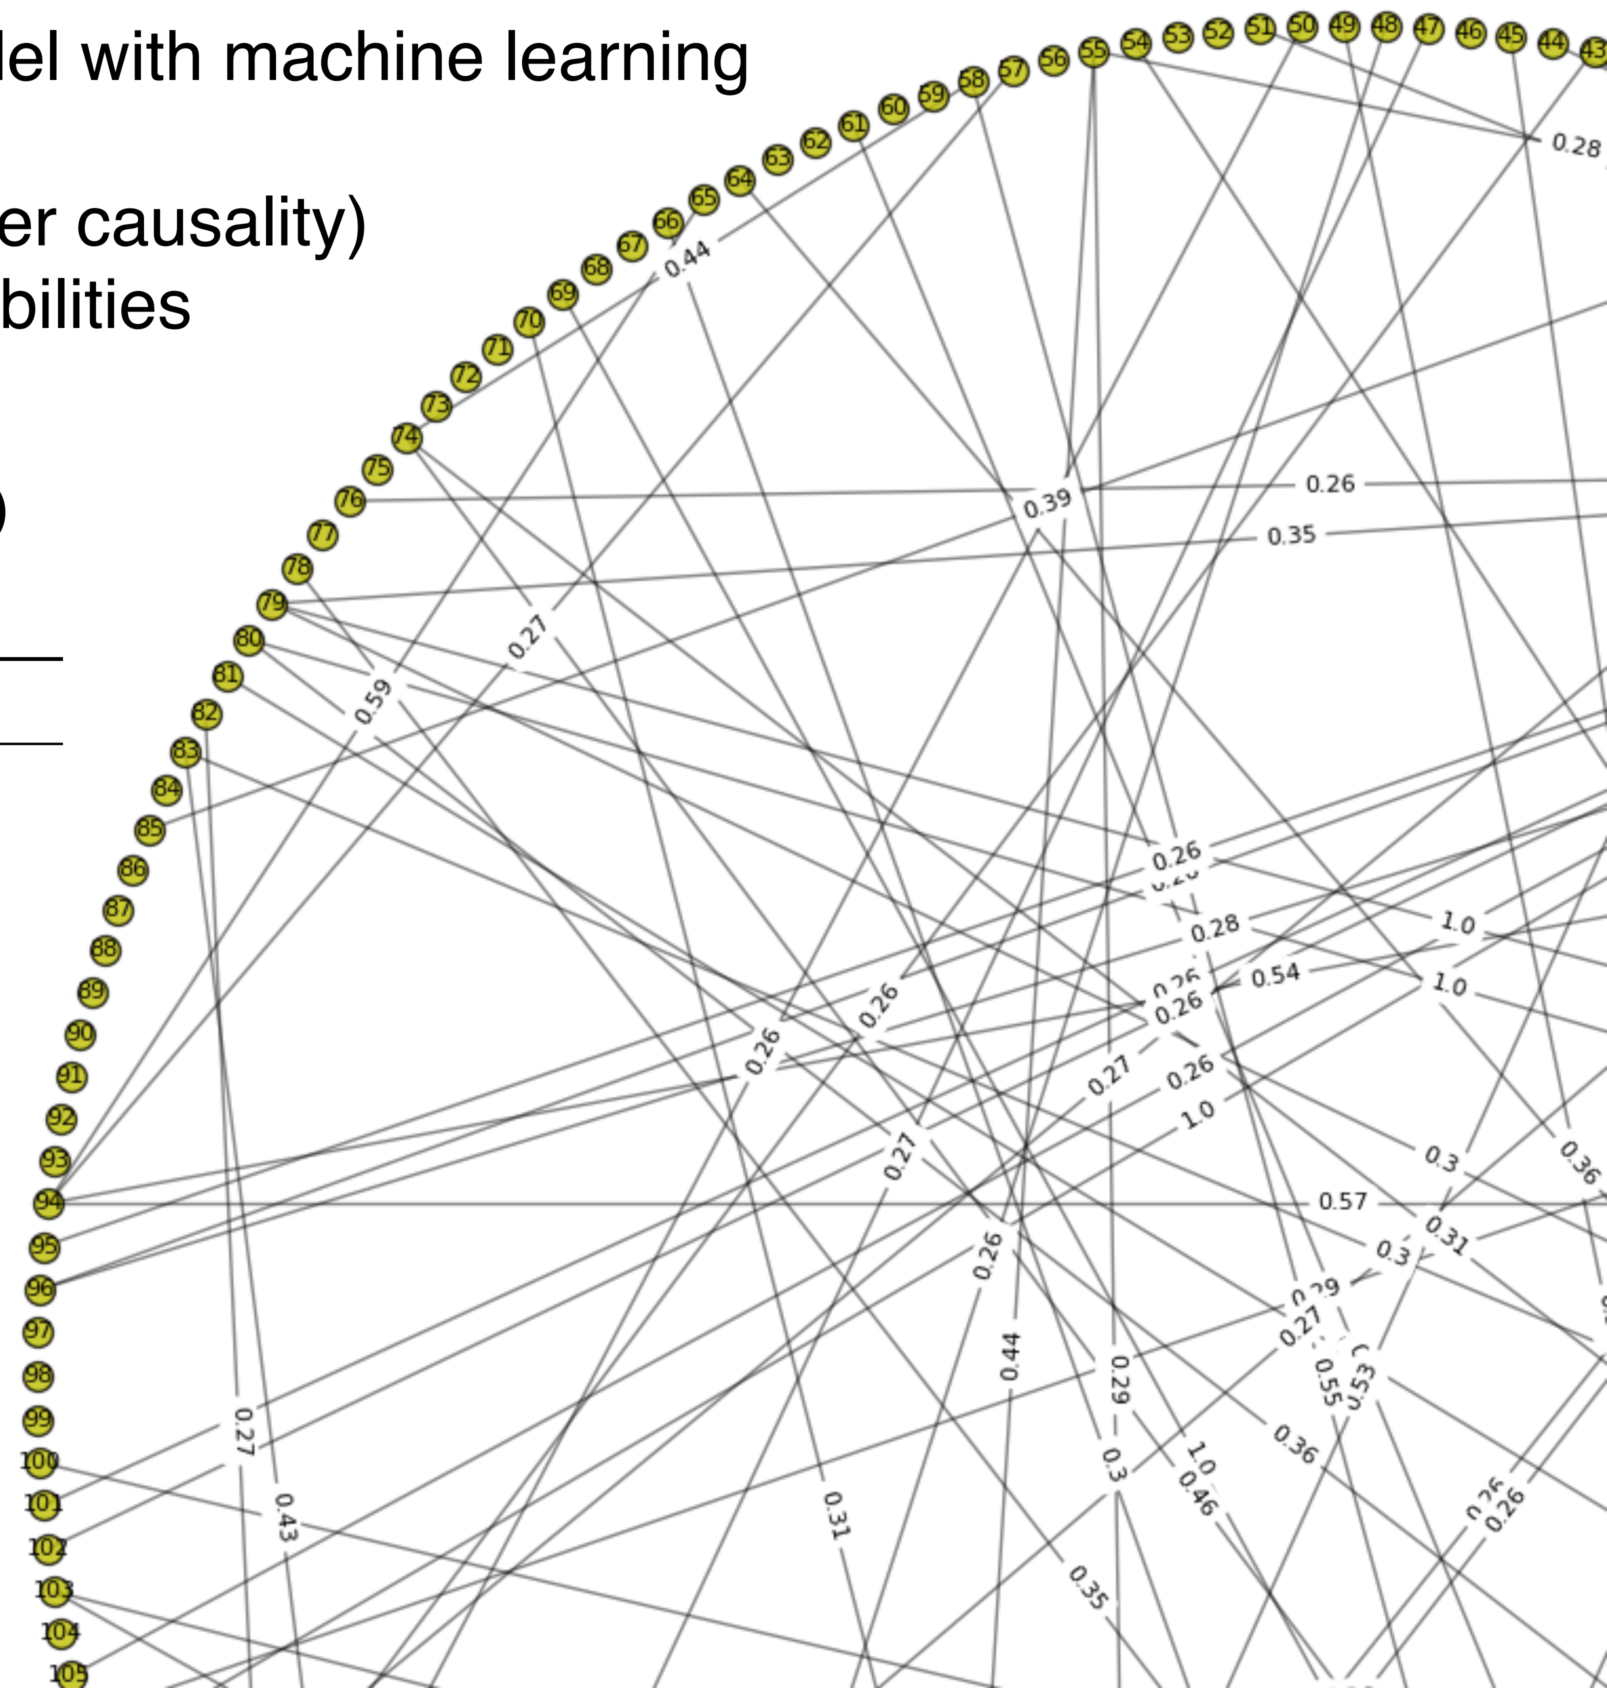

Failure Prediction and Fault Localization

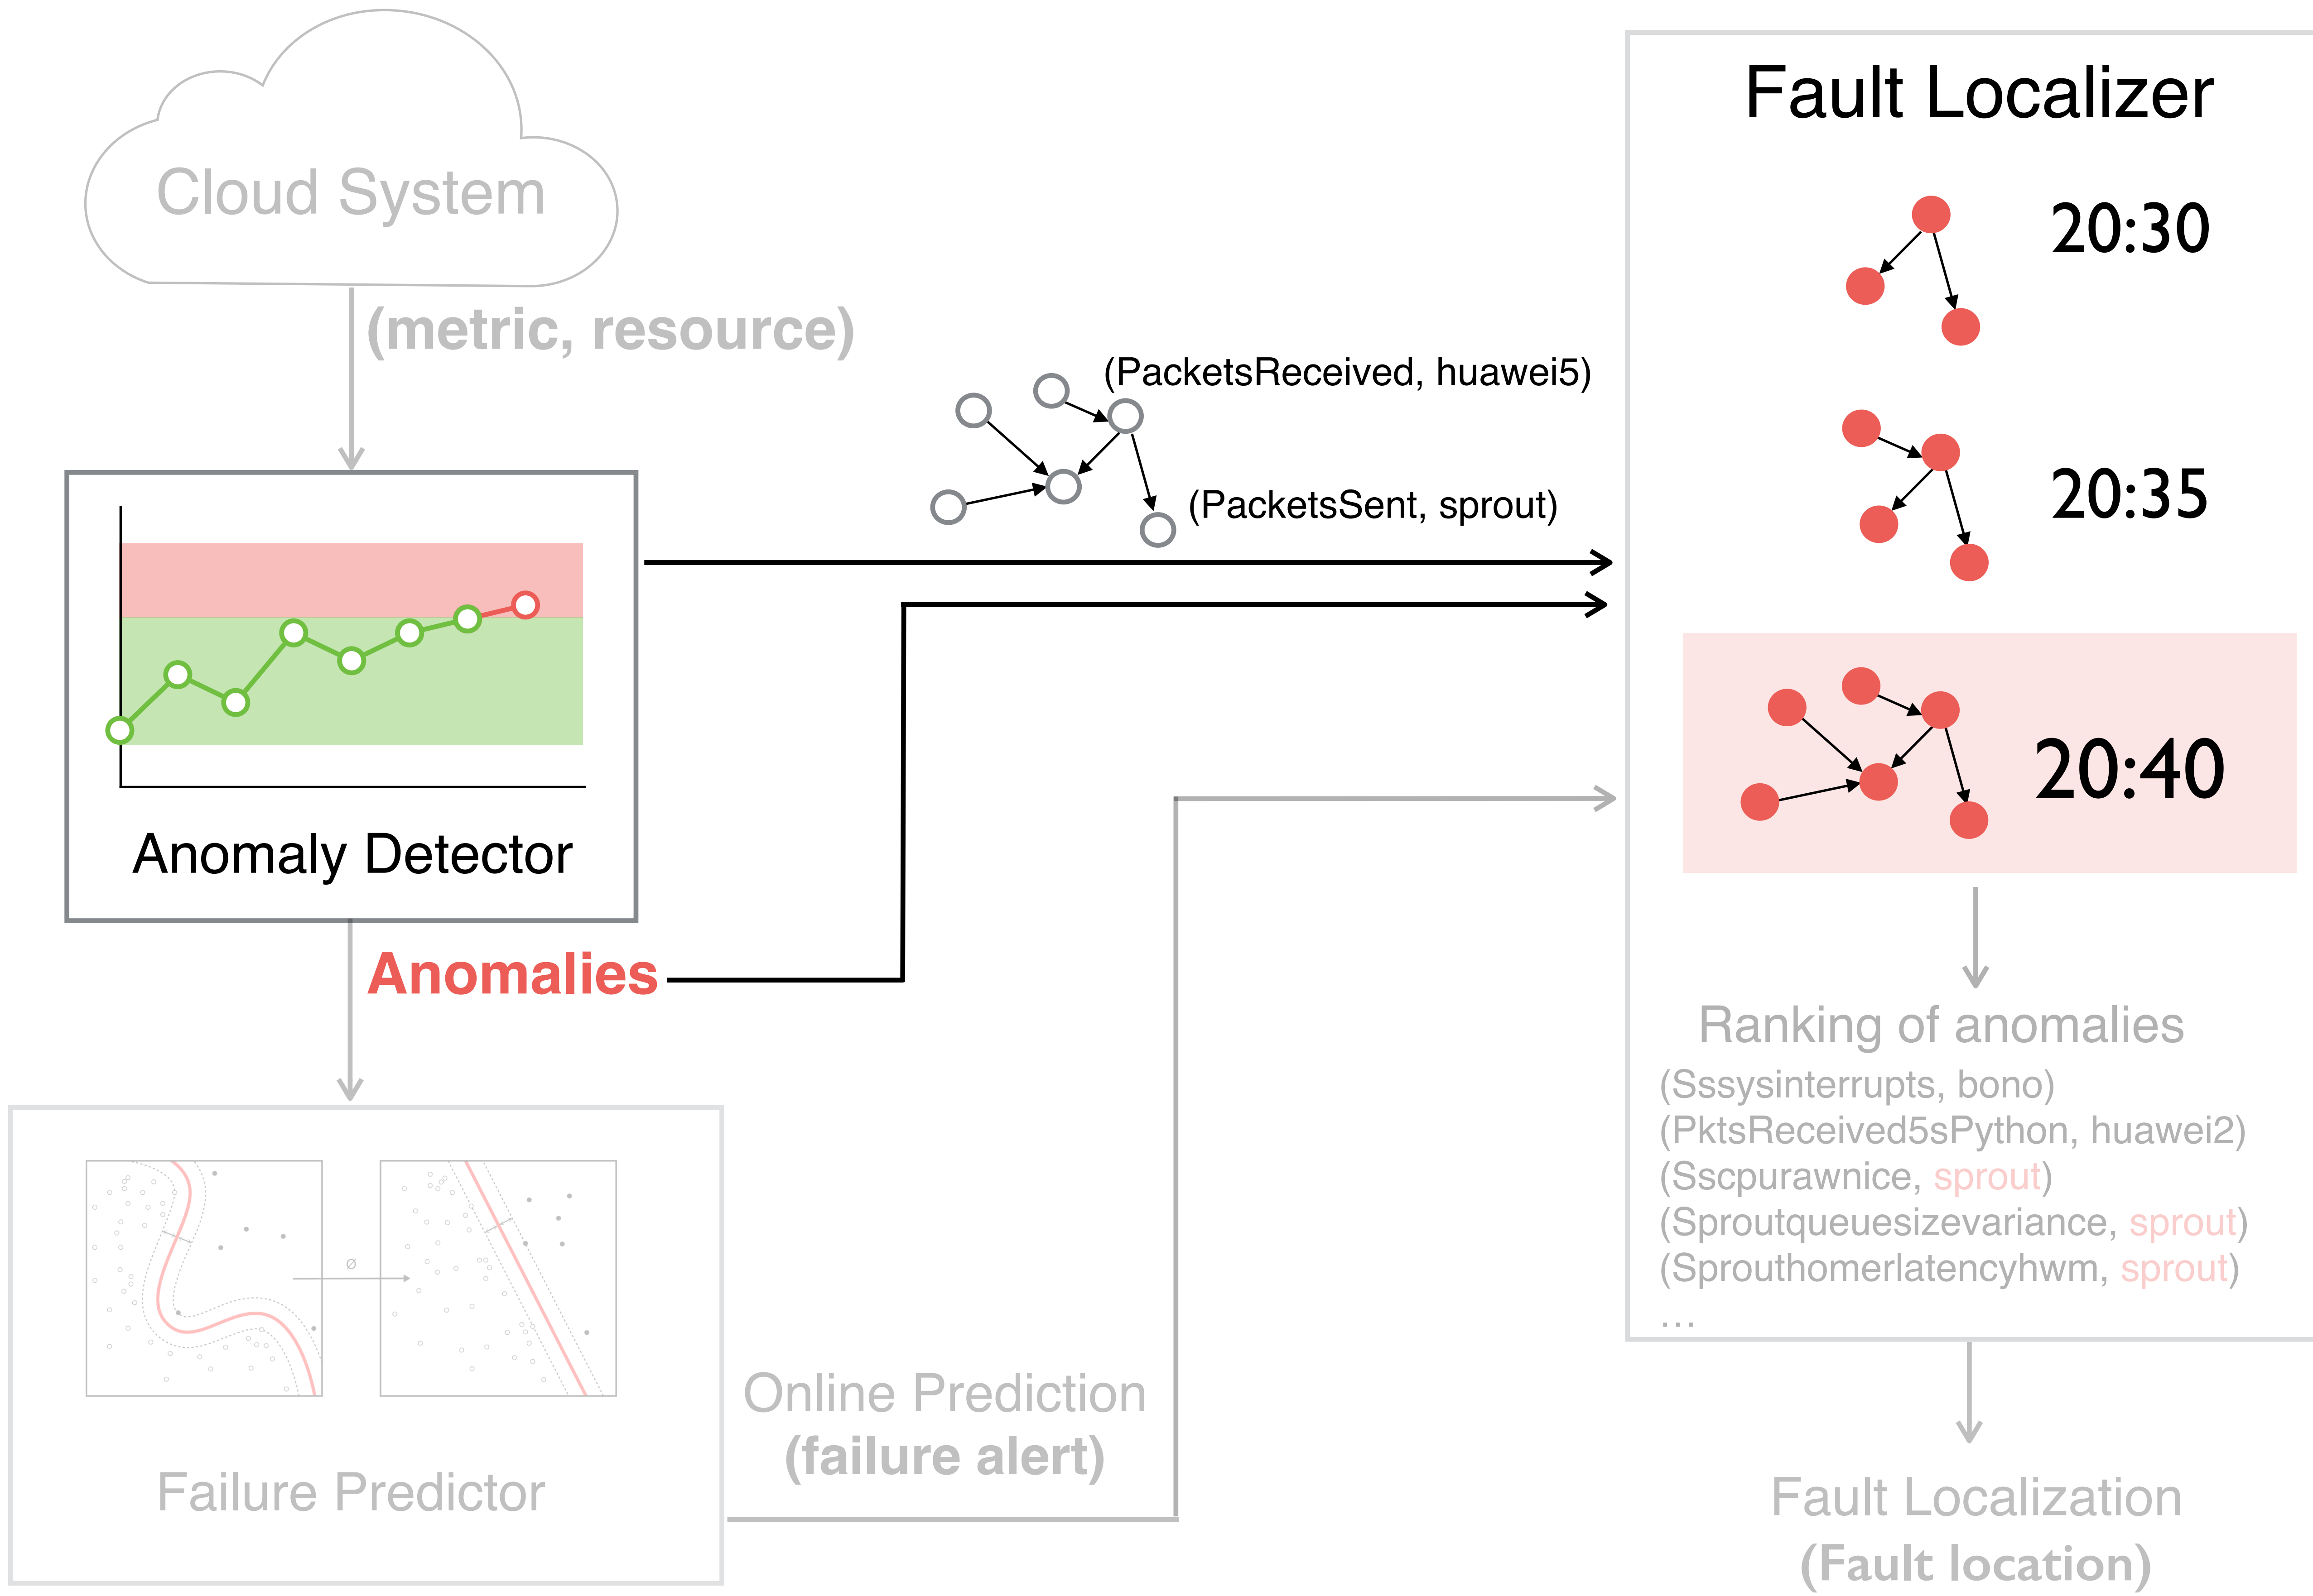

Fault Localization

Causality graph + anomalous KPIs = Evolution of Causality graph

Vertices: KPIs

Colored vertices: anomalous KPIs

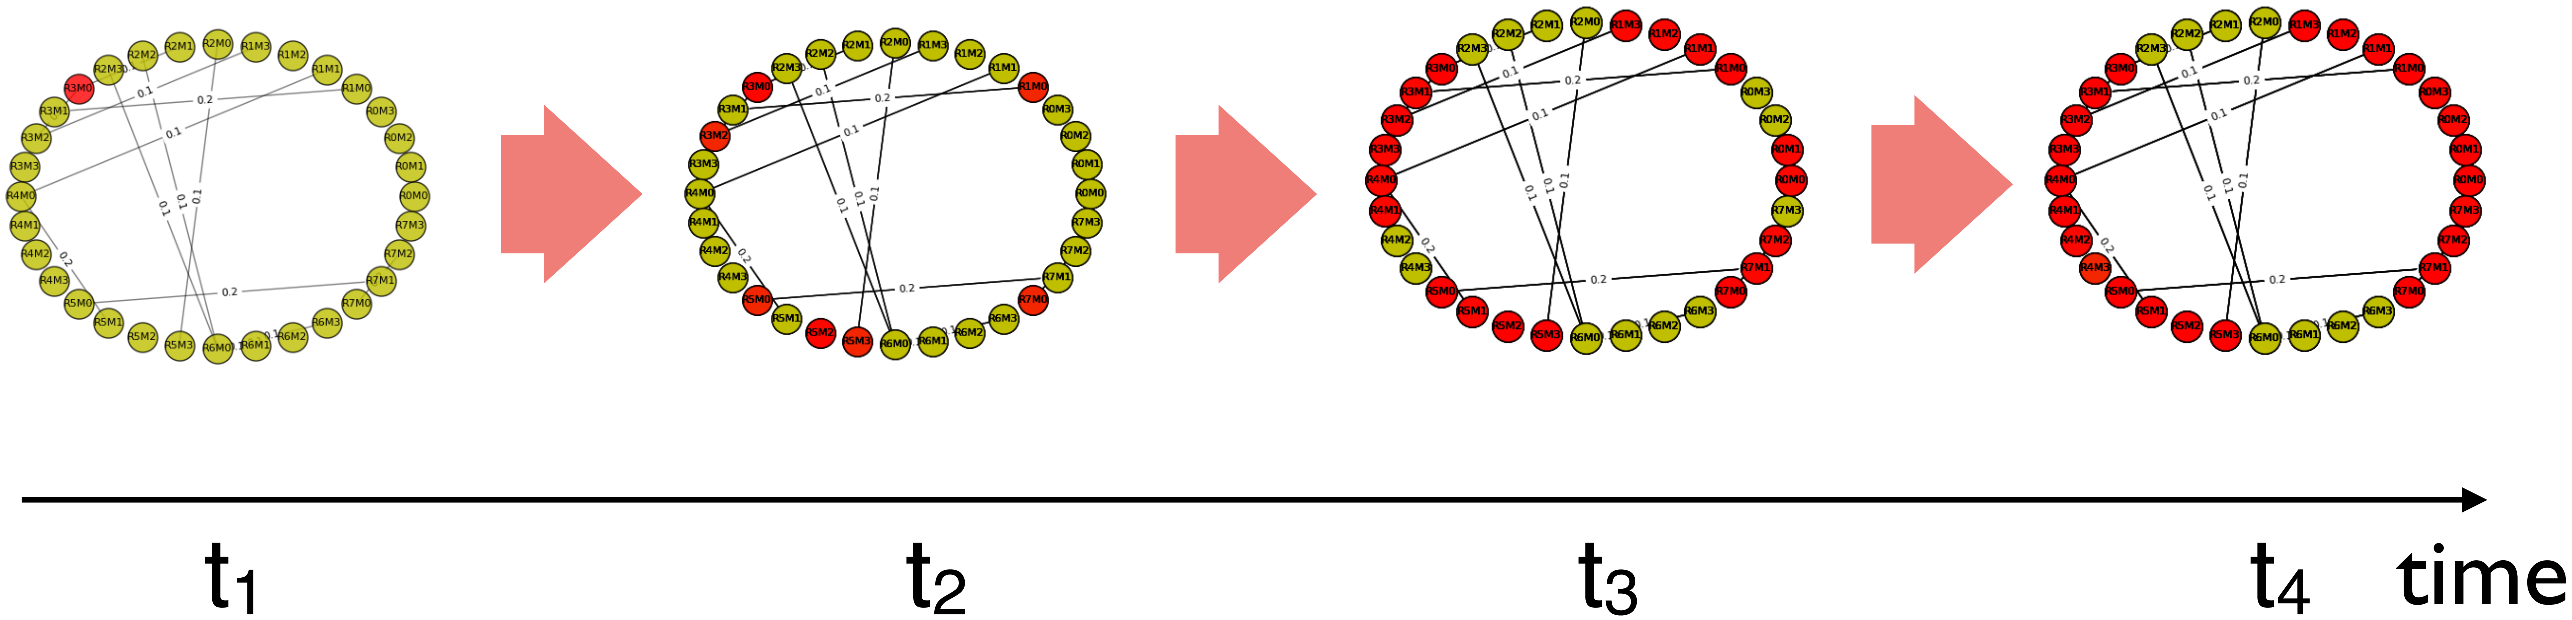

Failure Prediction and Fault Localization

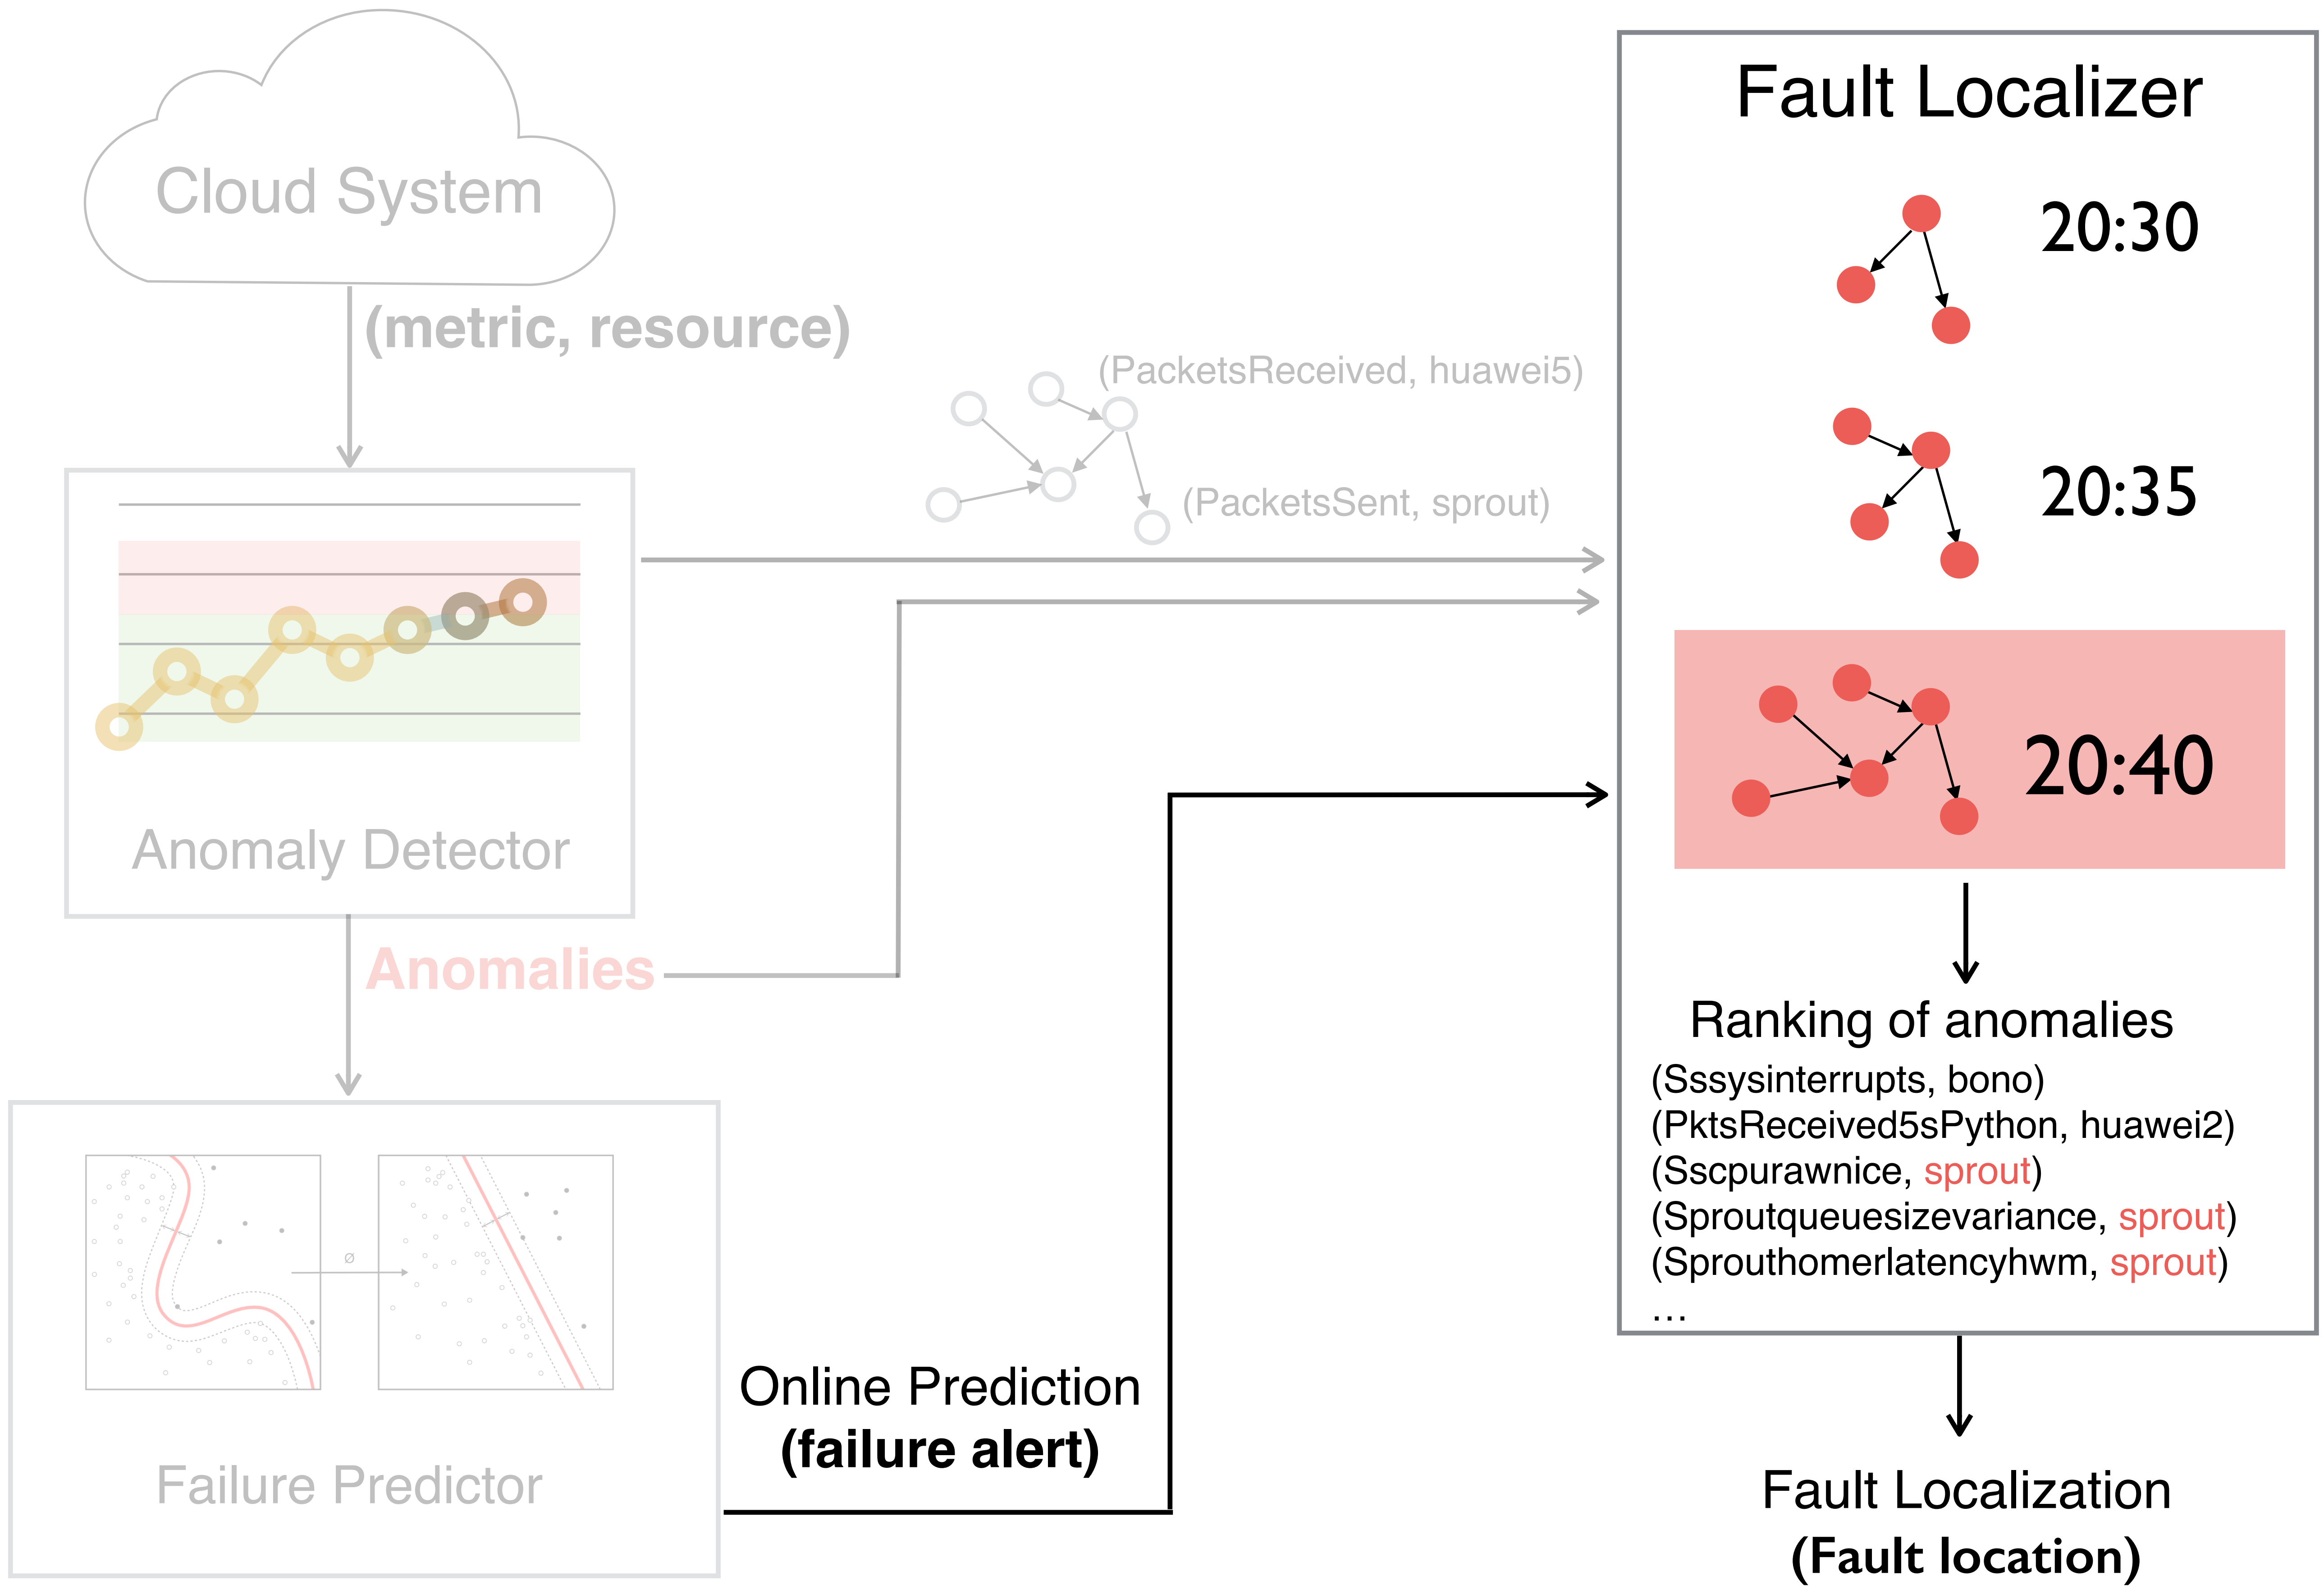

Fault Localization

Scores assigned to vertices in the evolution causality graph

Root cause → most important node in anomalous subgraph

Ranking Algorithm

PageRank

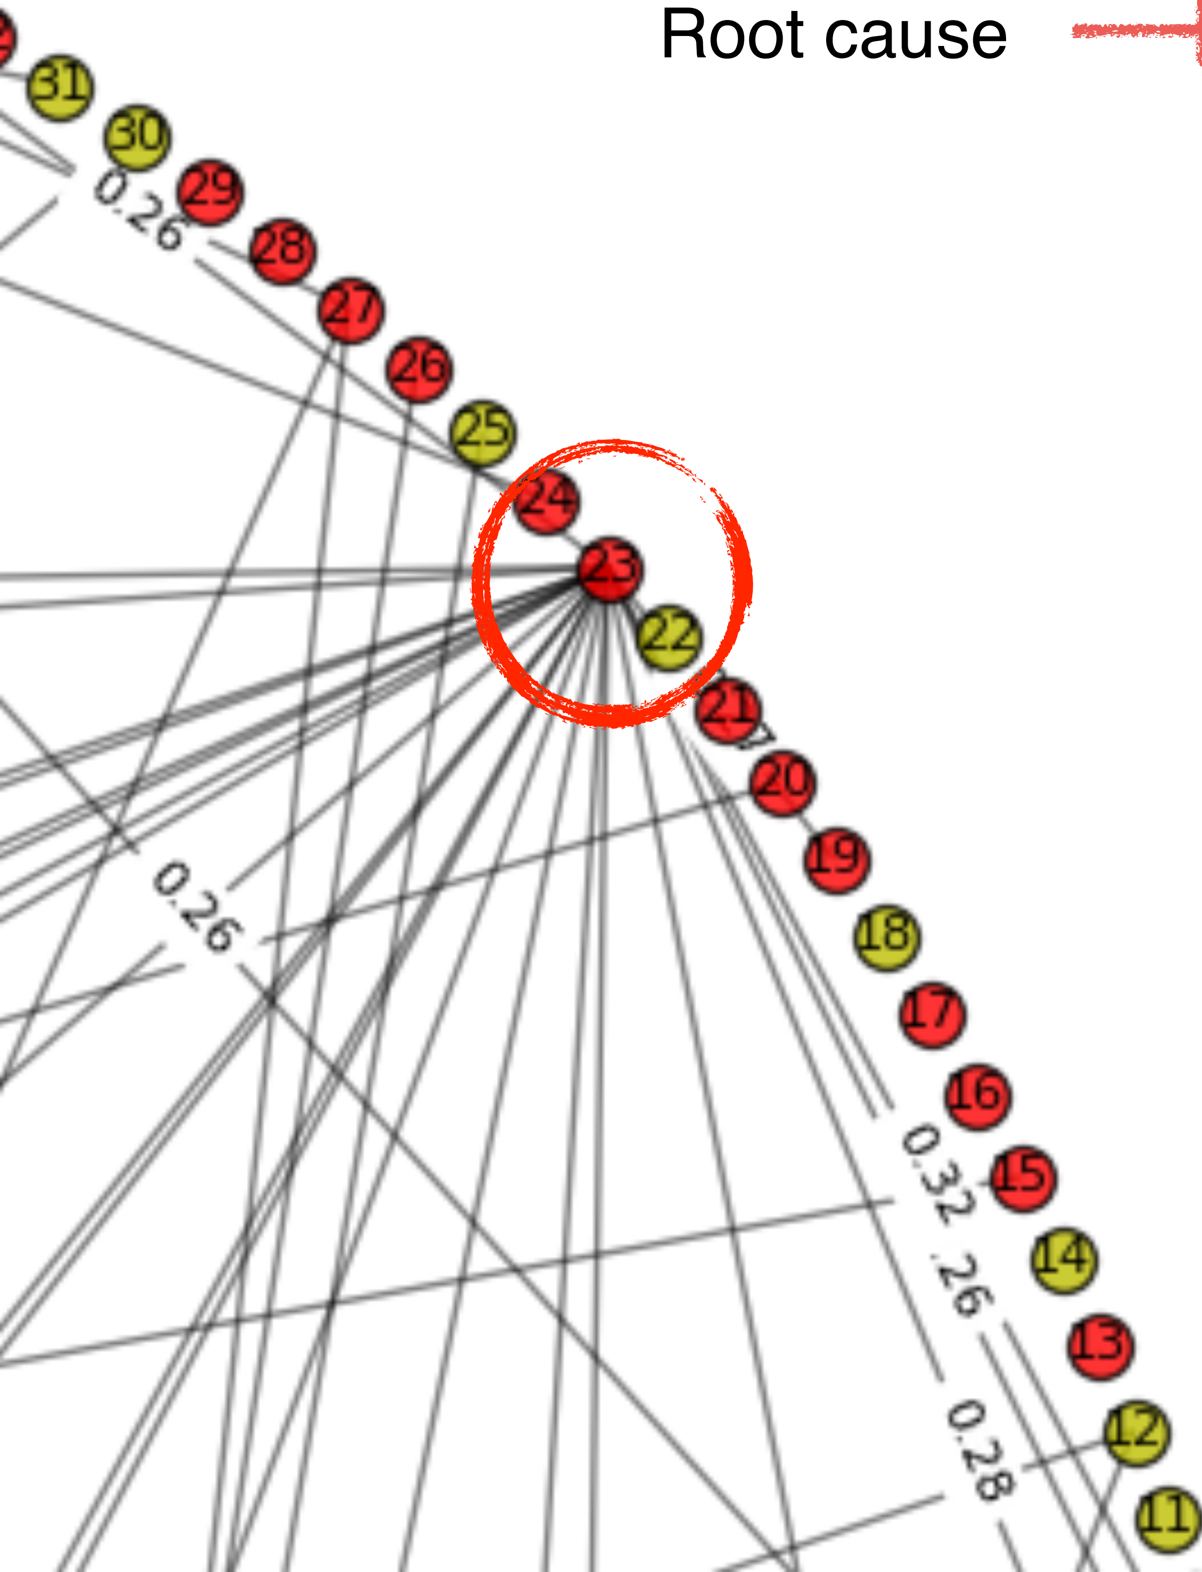

Fault Localization

Sprout

Packet loss injected in sprout
Failure alerts + vertex scores

homestead_Sscpurawnice
sprout_Sprouthomerlatencyhwm
homestead_SocketsInUse
homestead_Sscpuidle
homestead_Sscpurawsystem
…

bono_Sssysinterrupts
huawei2_PktsReceived5sPython
sprout_Sscpurawnice
sprout_Sproutqueuesizevariance
sprout_Sprouthomerlatencyhwm
…

bono_Ssiosent
huawei6_TotalNumberProcesses
huawei7_PktsReceived5sPython
sprout_SystemCpu
sprout_BusyCpu
…

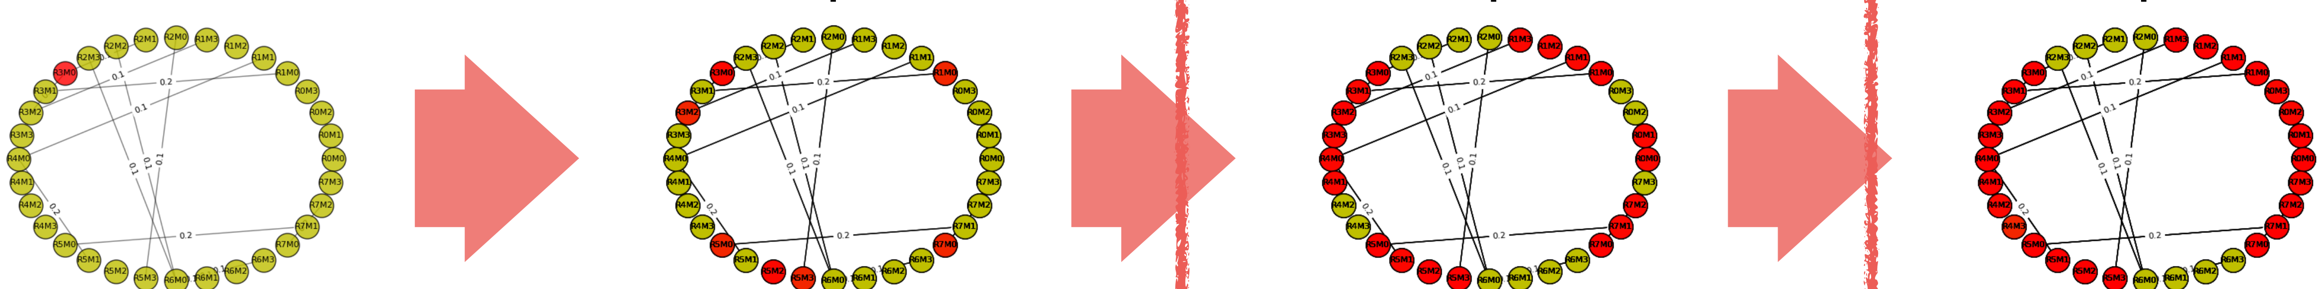

t_1	$t_2 = 10'$	$t_3 = 30'$	$t_4 = 2$ h
Normal	Normal	Alert	Failure

Case Study

Project Clearwater	IMS	162 KPIs
openstack	6 Virtual Machines	121 KPIs
KVM	8 Physical Machines	350 KPIs

- **Fault types:** Packet loss, Packet Latency, Packet Corruption
- **Workload pattern:** low traffic during week end and two peaks for day
- **# of Normal Samples:** 100
- **# of Faulty Samples:** 172

Can we accurately predict failures?

$$\text{Precision} = \frac{\text{Correctly predicted failures}}{\text{All predicted failures}}$$

$$\text{Recall} = \frac{\text{Correctly predicted failures}}{\text{All actual failures}}$$

Precision	Recall
98.2%	100%

- **Fault types:** Packet loss, Packet Latency, Packet Corruption
- **Workload pattern:** low traffic during week end and two peaks for day
- **# of Normal Samples:** 100
- **# of Faulty Samples:** 172

Can we accurately localize faults?

$$\text{Precision} = \frac{\text{Correctly localized faults}}{\text{All predicted failures}}$$

$$\text{Recall} = \frac{\text{Correctly localized faults}}{\text{All actual failures}}$$

Precision	Recall
97%	100%

- **Fault types:** Packet loss, Packet Latency, Packet Corruption
- **Workload pattern:** low traffic during week end and two peaks for day
- **# of Normal Samples:** 100
- **# of Faulty Samples:** 172


WHAT'S NEXT?

Self-Healing

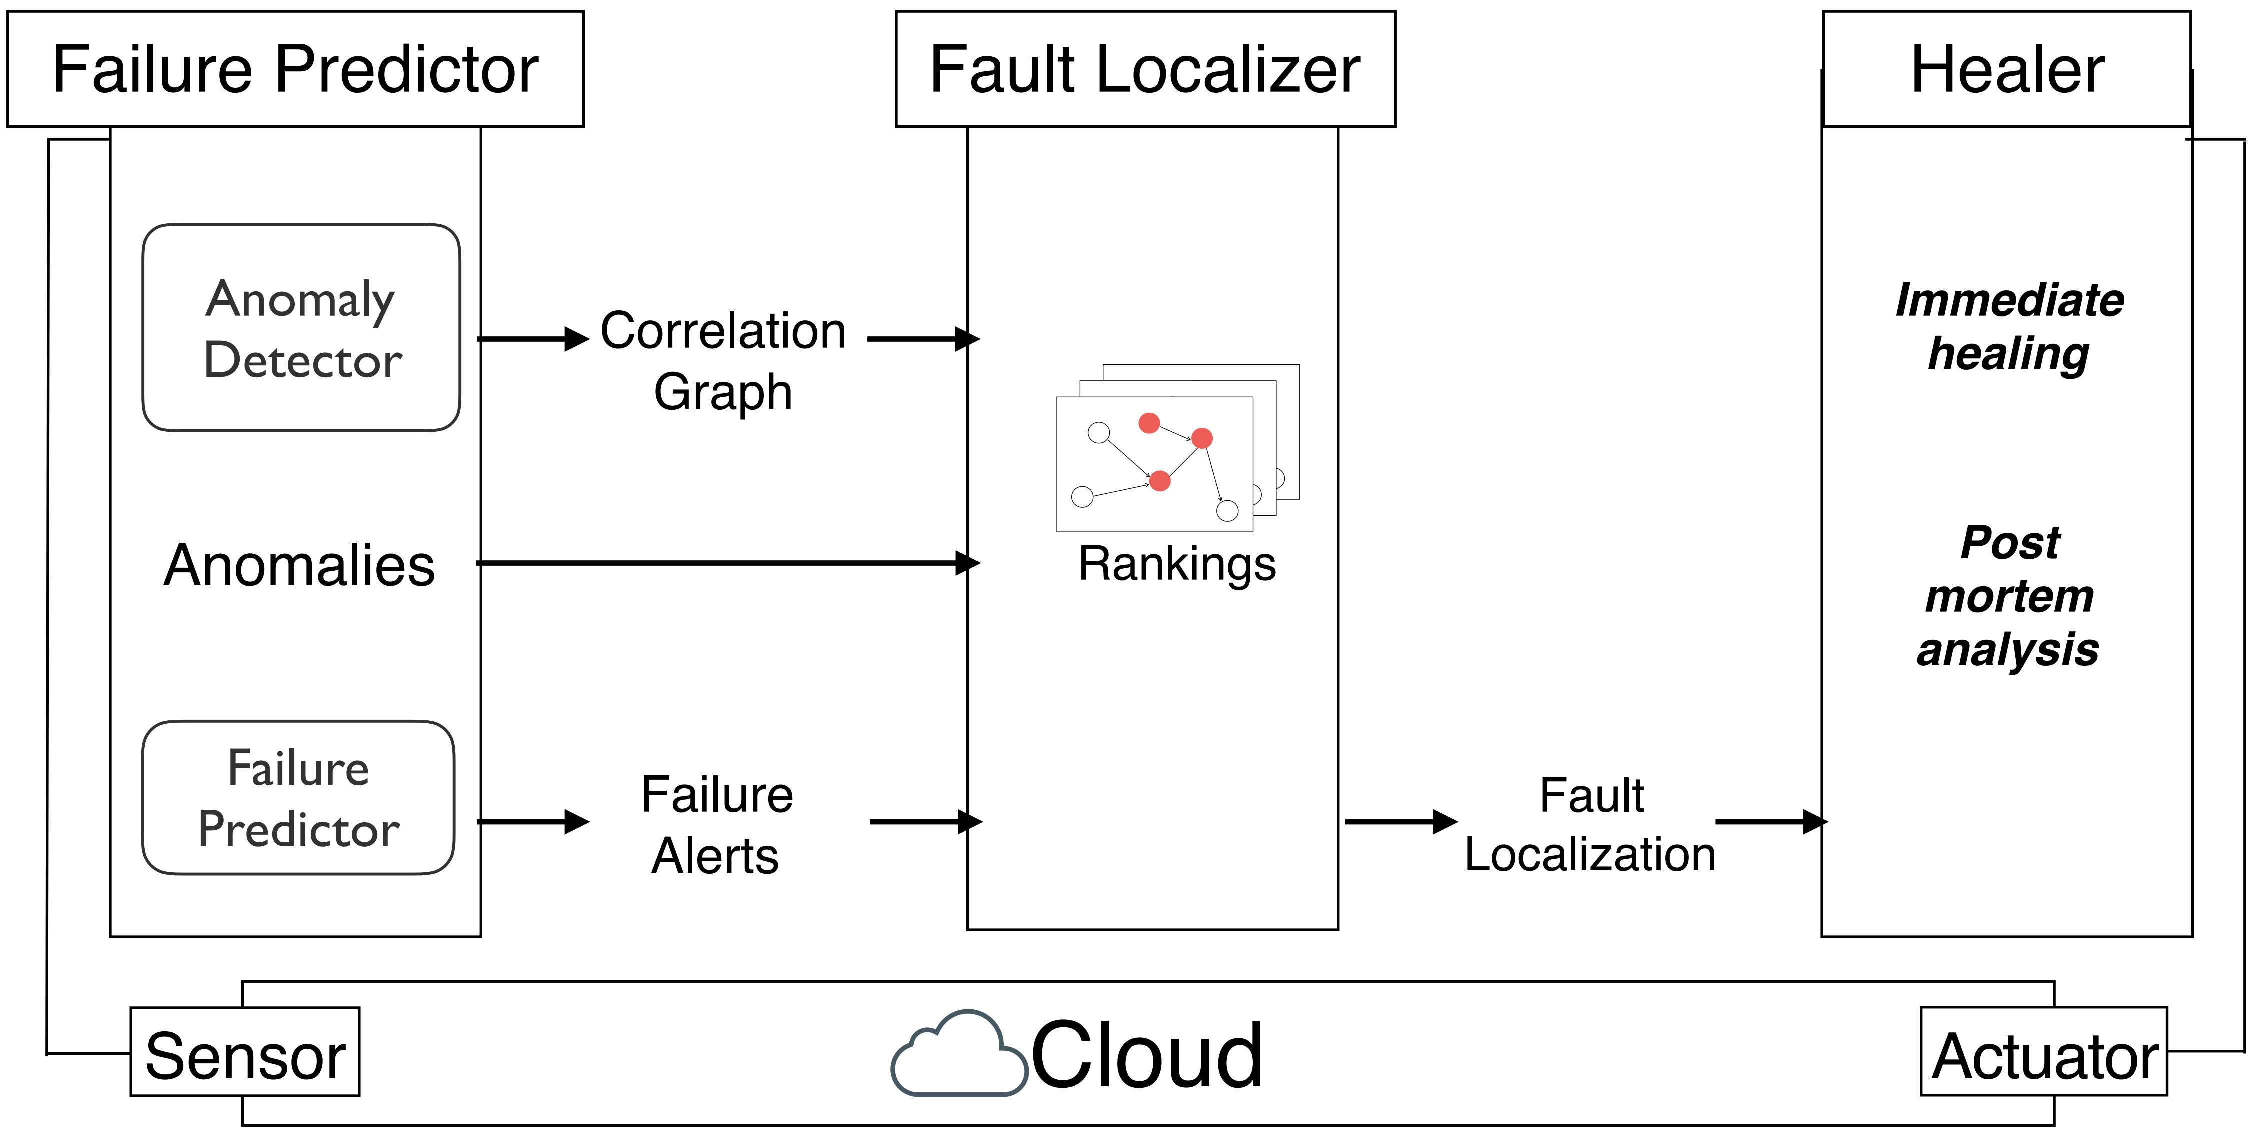

New anomaly detector based on neural network