


AIOps @ Microsoft

Mathew John, Partner Director
June 2024



AIOps harnesses AI and machine learning to revolutionize operations, automating services and optimizing resources for streamlined efficiency.

The goal of any Service organization is to deliver high-quality services that meet customer needs and expectations.

But Humans!



And tools!

Developer: Tom

SomeFile.c:

```
if (condition) {  
    // Do stuff  
    <Code change by Tom>  
} else {  
    // Do other stuff  
}
```

Developer: Susan

SomeFile.c:

```
if (condition) {  
    // Do stuff  
} else {  
    // Do other stuff  
    <Code change by Susan>  
}
```

Automated Merge

SomeFile.c:

```
if (condition) {  
    // Do stuff  
    <Code change by Tom>  
  
    // Do other stuff  
    <Code change by Susan>  
}
```

And scale!



Azure is the **world's** computer

1b

Customers across enterprise + consumer segments

600+

Azure Resource types

15K+

Daily updates

160T

Terabits/sec. Marea: The highest capacity transatlantic subsea cable

79b

Events/min

76b

Active time series

240M

Queries/min for Metrics

286M

Resources monitored/min

66+

Azure regions

300+

Datacenters

175K+


Miles of fiber

190+

Network PoPs

● Regions

● Edge Network



Challenges for the Service Operator


Visibility Overcomes monitoring limitations of modern applications

Complexity Addresses the fragmented nature of today's services

Scale Manages the vast infrastructure demands of cloud services

Efficiency Analyzes and resolves a high volume of network tickets swiftly

Our Approach



Comprehensive **standardized, accurate** and **reliable**
understanding of service health



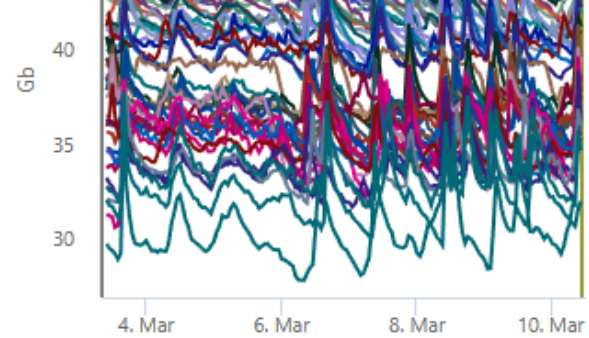
Service Health in **Near Real Time** using Service Level Indicators (**SLIs**) and Service Health Indicators (**SHIs**)



Production changes **integrated** with service health




If an incident occurs, communication
is **automated** and **timely**




- Norway West
 - South Central US STG
 - East US STG
 - Brazil US
 - East US 2 EUAP
- ▲ 1/13 ▼

All None




- East US STG
 - South Central US STG
 - East US STG
 - South Central US STG
 - Brazil US
- ▲ 1/25 ▼

All None





- Norway West / 1024
 - South Central US STG / 1024
 - East US STG / 1024
 - Brazil US / 1024
 - East US 2 EUAP / 1024
- ▲ 1/13 ▼

All None



- Norway West / 1024
 - South Central US STG / 1024
 - East US STG / 1024
 - Brazil US / 1024
 - East US 2 EUAP / 1024
- ▲ 1/13 ▼

All None

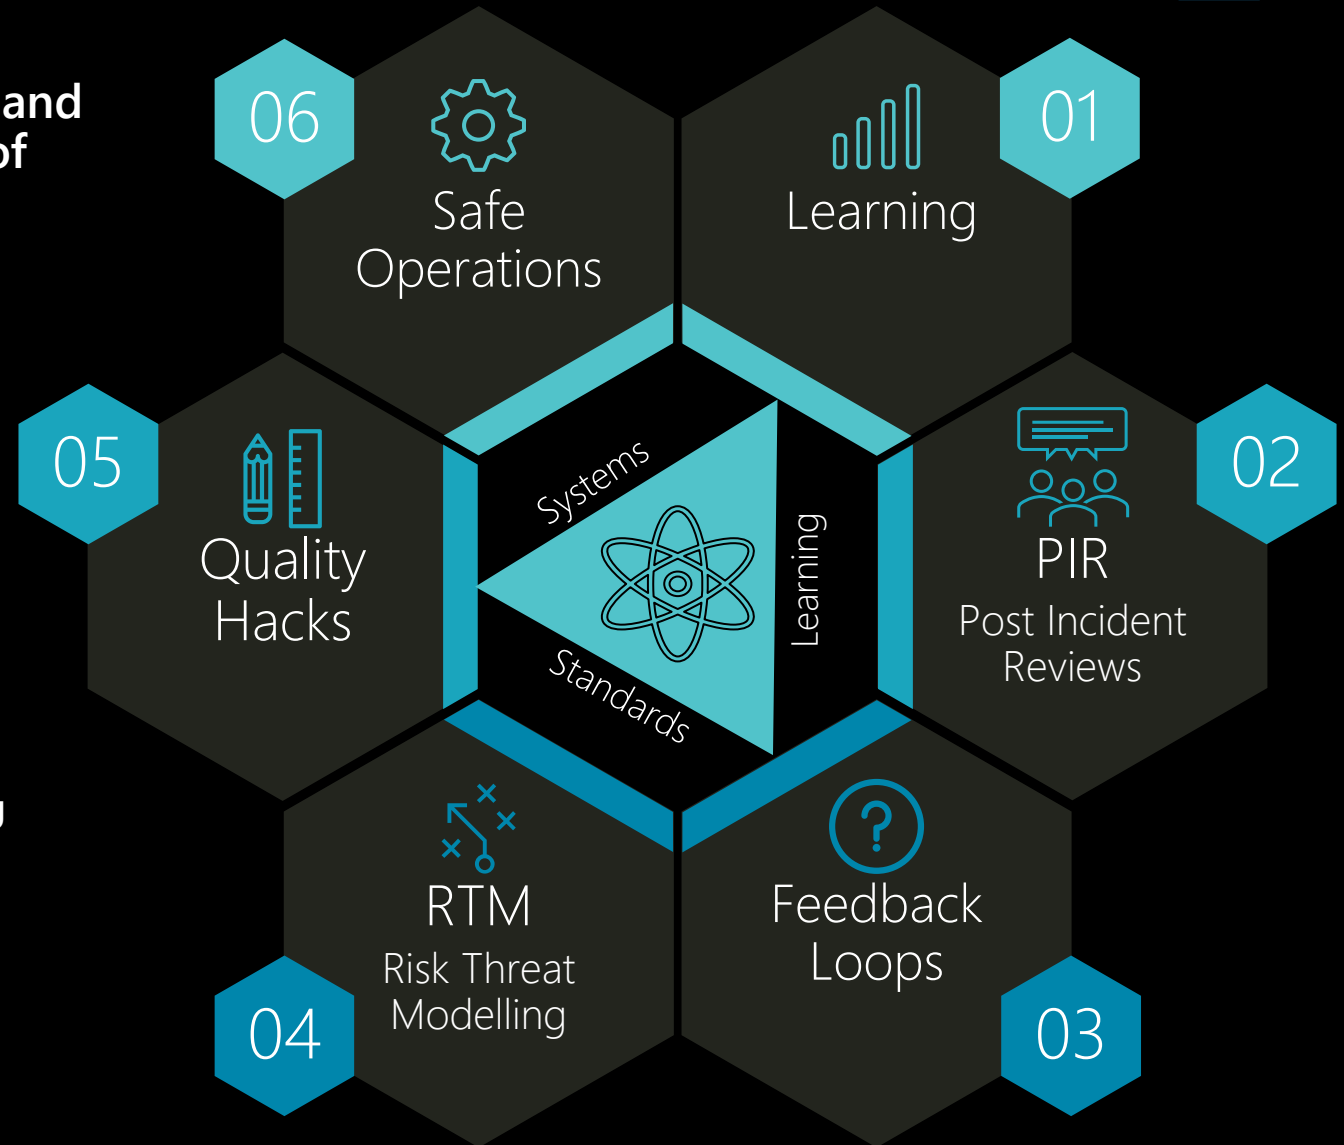


Diagnosing issues is **simple / automated**, requiring little DRI toil or manual touches

Culture of Quality


Cultivate an environment where we **Listen** and **Learn** to understand the lived experience of our DRIs and our customers.

-  Local Learning and Sharing Opportunities
-  Seek Diverse Perspective in PIRs
-  Create Safe Space for Feedback
-  Quality Planning through Risk Modelling
-  Dedicate Time for Quality Innovation
-  There is No Human Error



The pivot to AIOps

AIOps – Gartner's definition



Big data and ML driven IT operation automation process

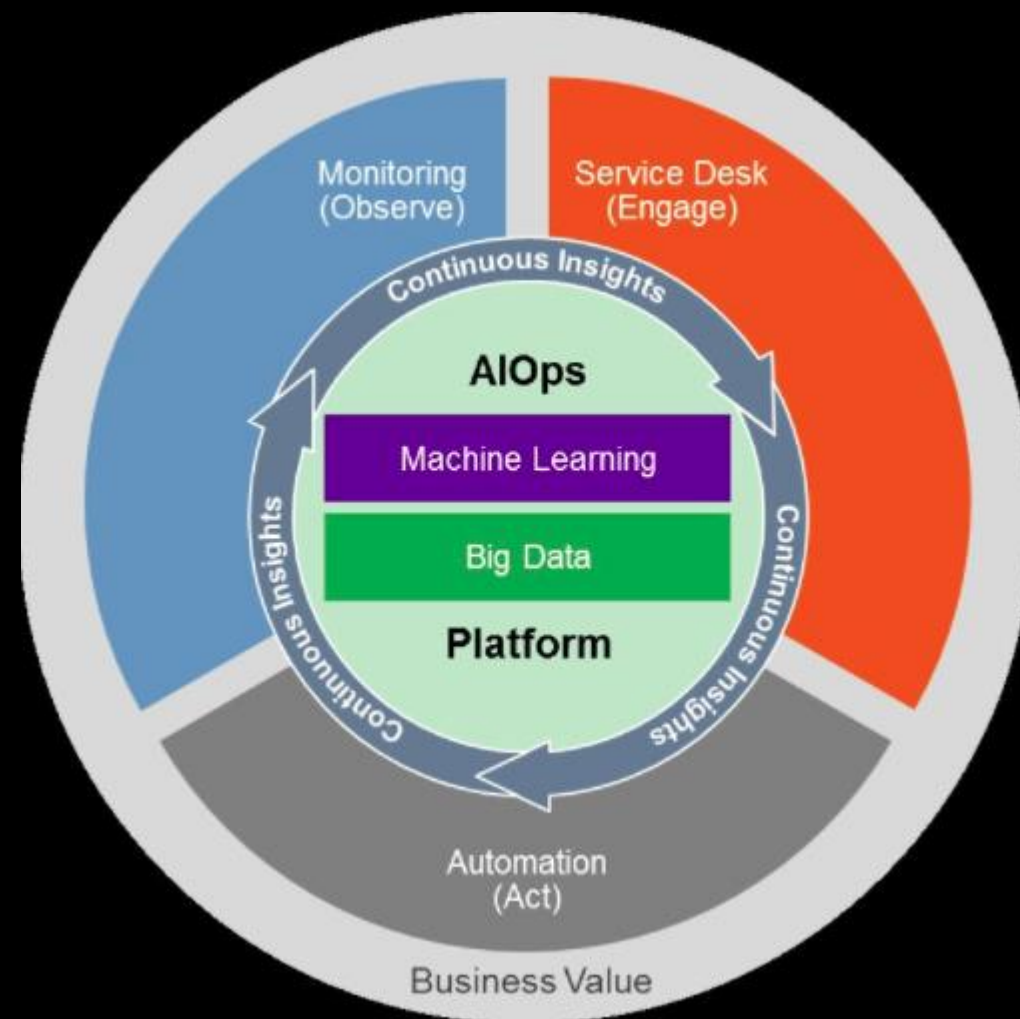


Adoption has increased with the uptick of digital transformation



Business value


Higher efficiency
Higher Service quality
Lower COGS



Source: Gartner

AIOps in Azure

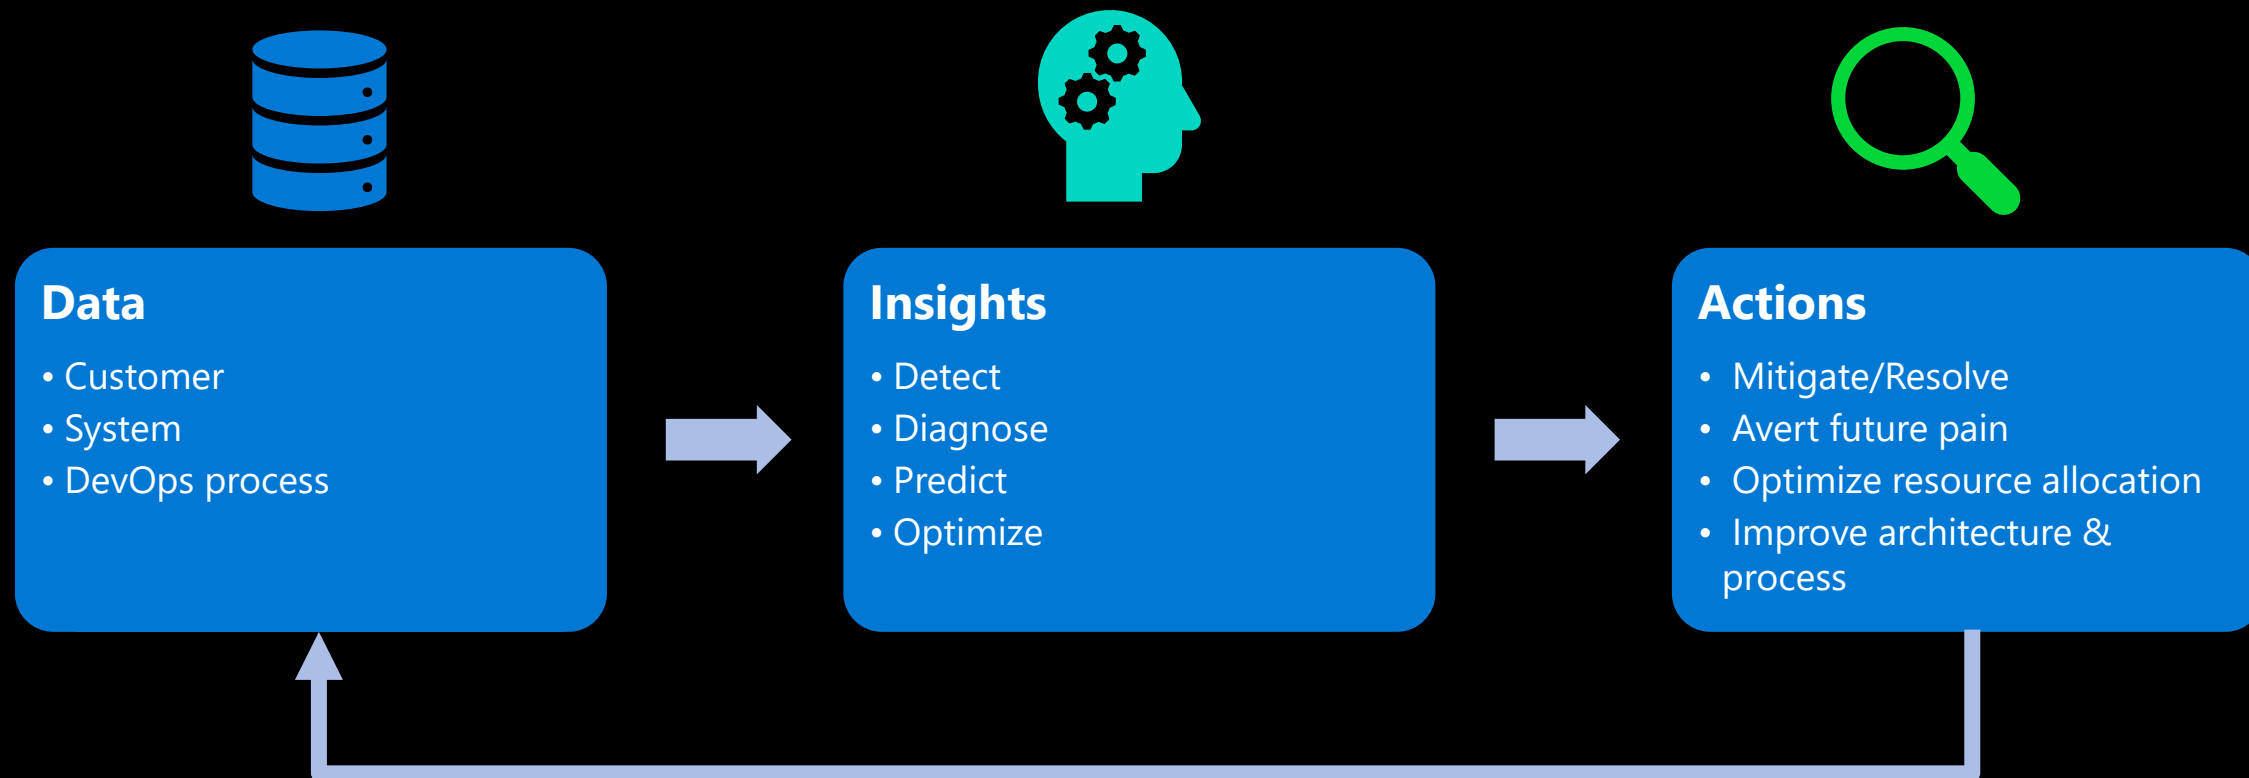
Innovating AI/ML technologies to effectively and efficiently **design, build, and operate** complex **cloud services at scale**



- **AI for Systems**
Building high-quality services with better reliability, performance, and efficiency
- **AI for DevOps**
Achieving high productivity in DevOps via empowering engineers with intelligent tooling
- **AI for Customers**
Improving customer satisfaction with intelligence and better user experiences

Reference: [Advancing Azure service quality with artificial intelligence: AIOps](#) published in *Azure Reliability Blog Series*

AIOps Methodologies: From Data to Actions




Our scale solution: Azure Brain

Health & AIOps system

Azure Brain



Customer Experience




Azure Services



Infrastructure devices



Critical Environment and Mechanical



Brain

Network of Intelligence



Automatic Alert correlation



Fast and actionable anomaly detection



Auto-communication



Automatic impacted service identification



Impact assessment



Root cause service identification



Efficient outage management

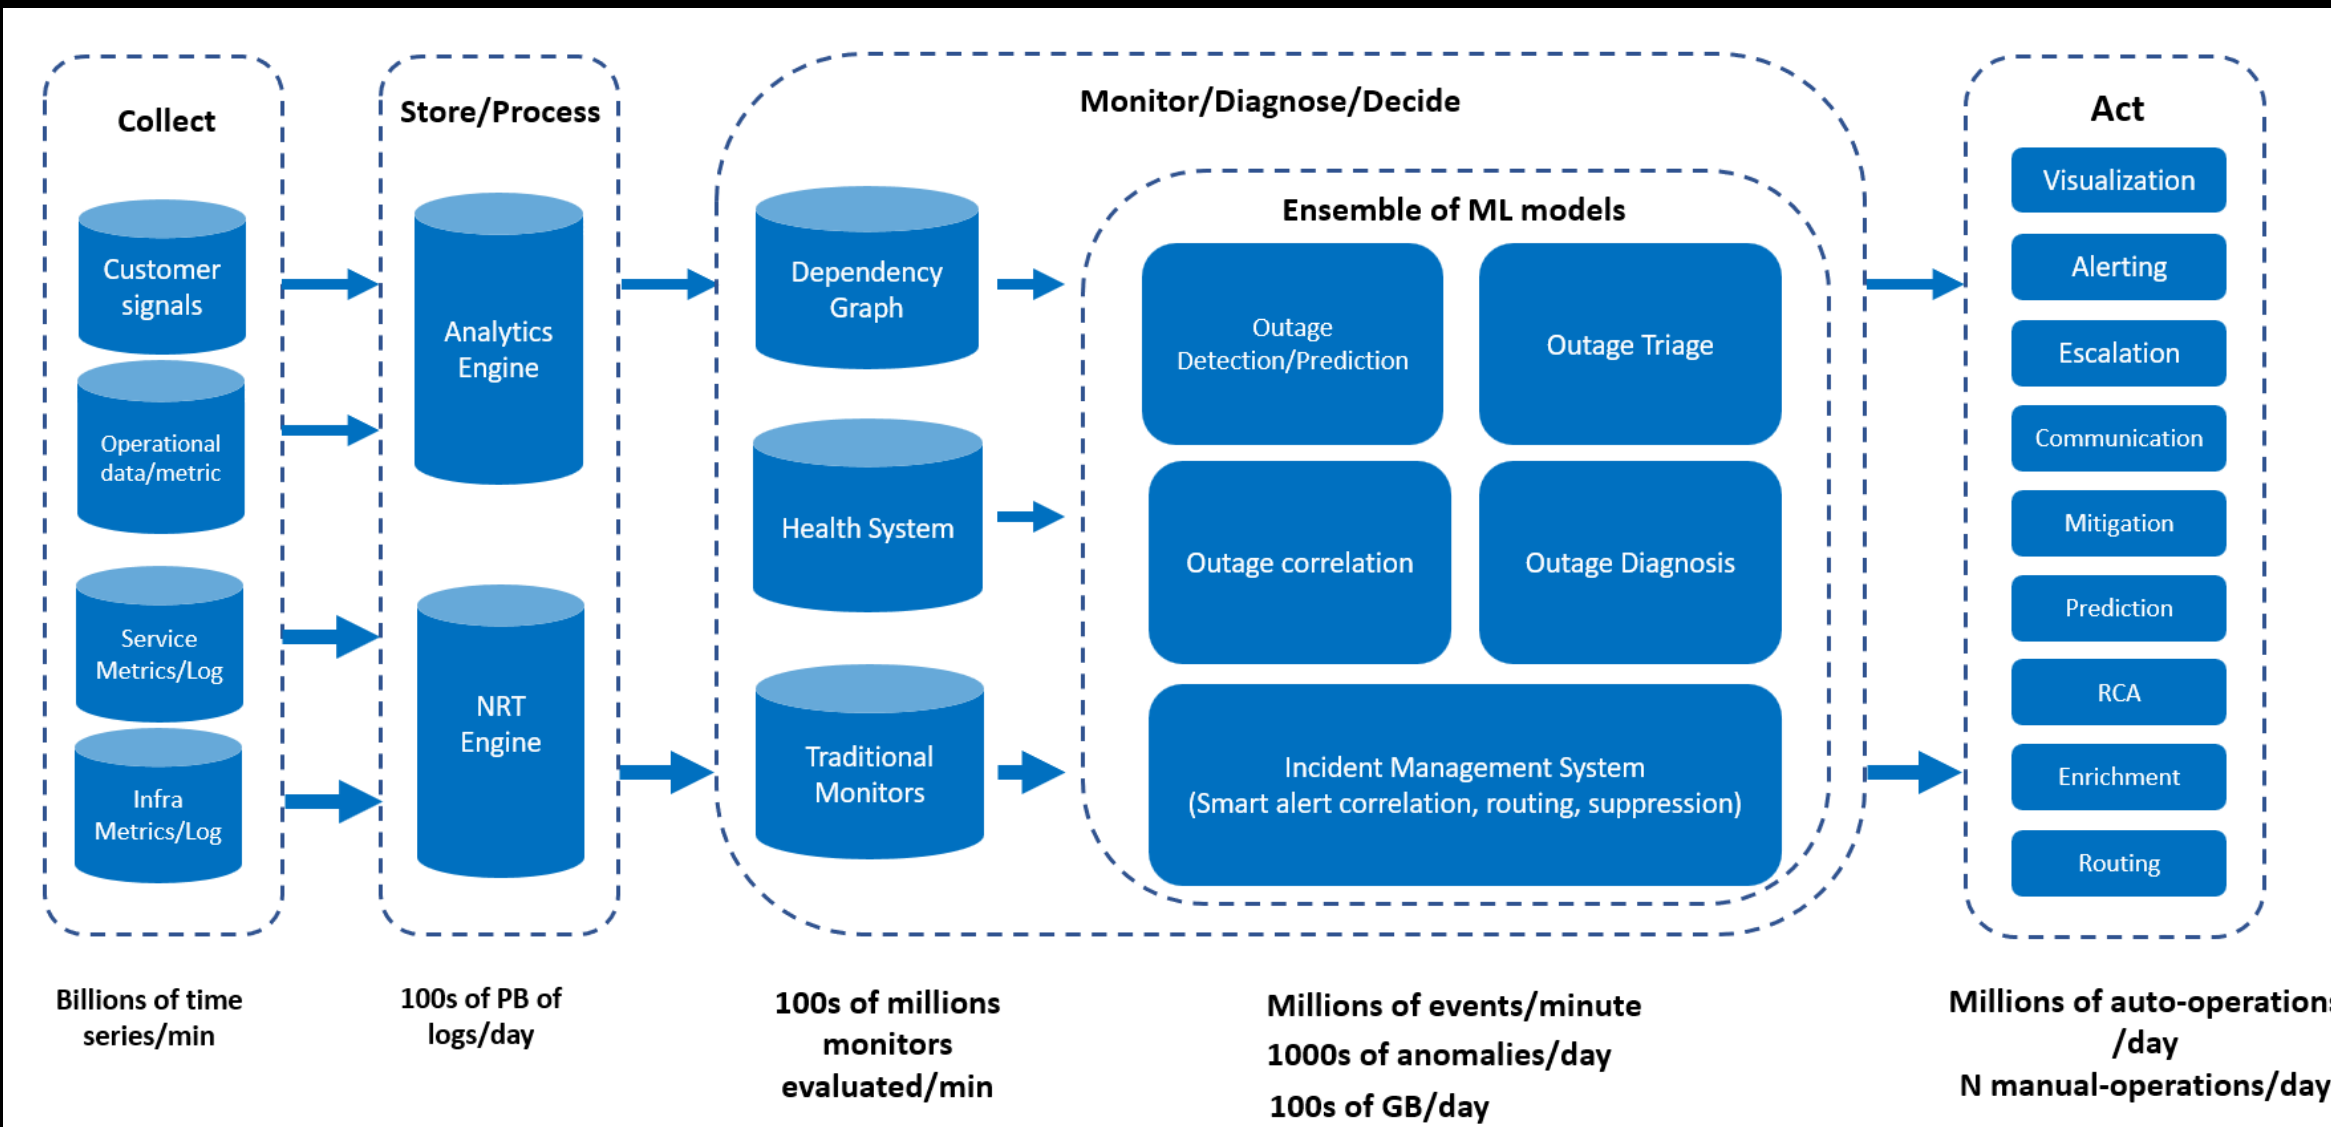


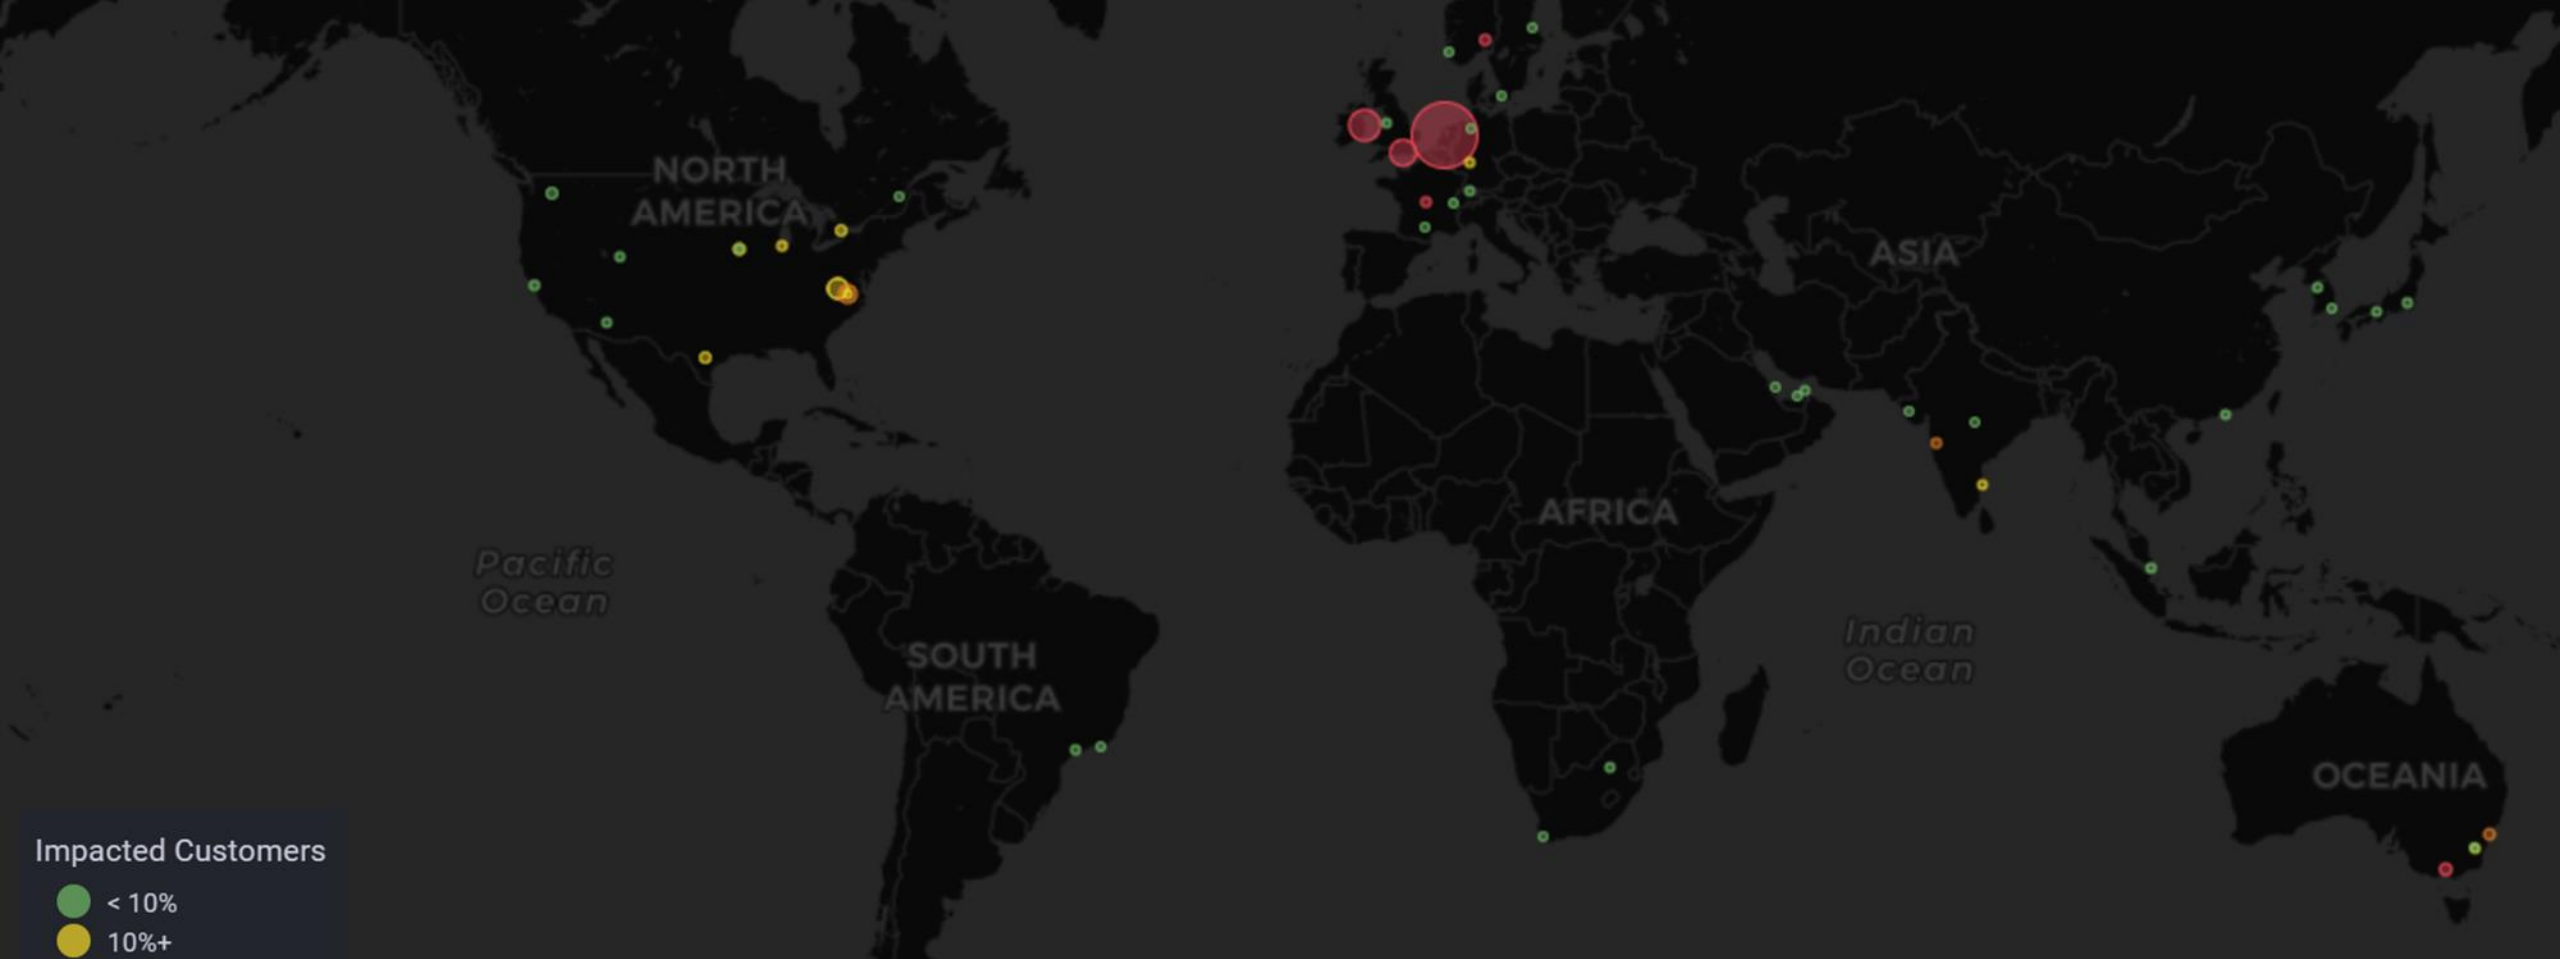
Diagnostic experiences



Auto-Mitigation

Azure Brain Pipeline





Service Health at Scale in **Near Real Time** using **SLOs/SLIs**

AI/ML challenges for AIOps


Large-volume and heterogeneous non-uniform data

Extremely imbalanced samples

Lack of canonical ground truth

AI system and human interaction

No universal intelligence for diverse scenarios



Abnormal:Normal

1:10,000

Azure Gandalf: AIOps for Infrastructure Health

Proactive prevention of issues: integrating intelligence into Azure Infrastructure

- Preventing code regressions into fleet

- Increasing host resilience

- Governance of host resource usage

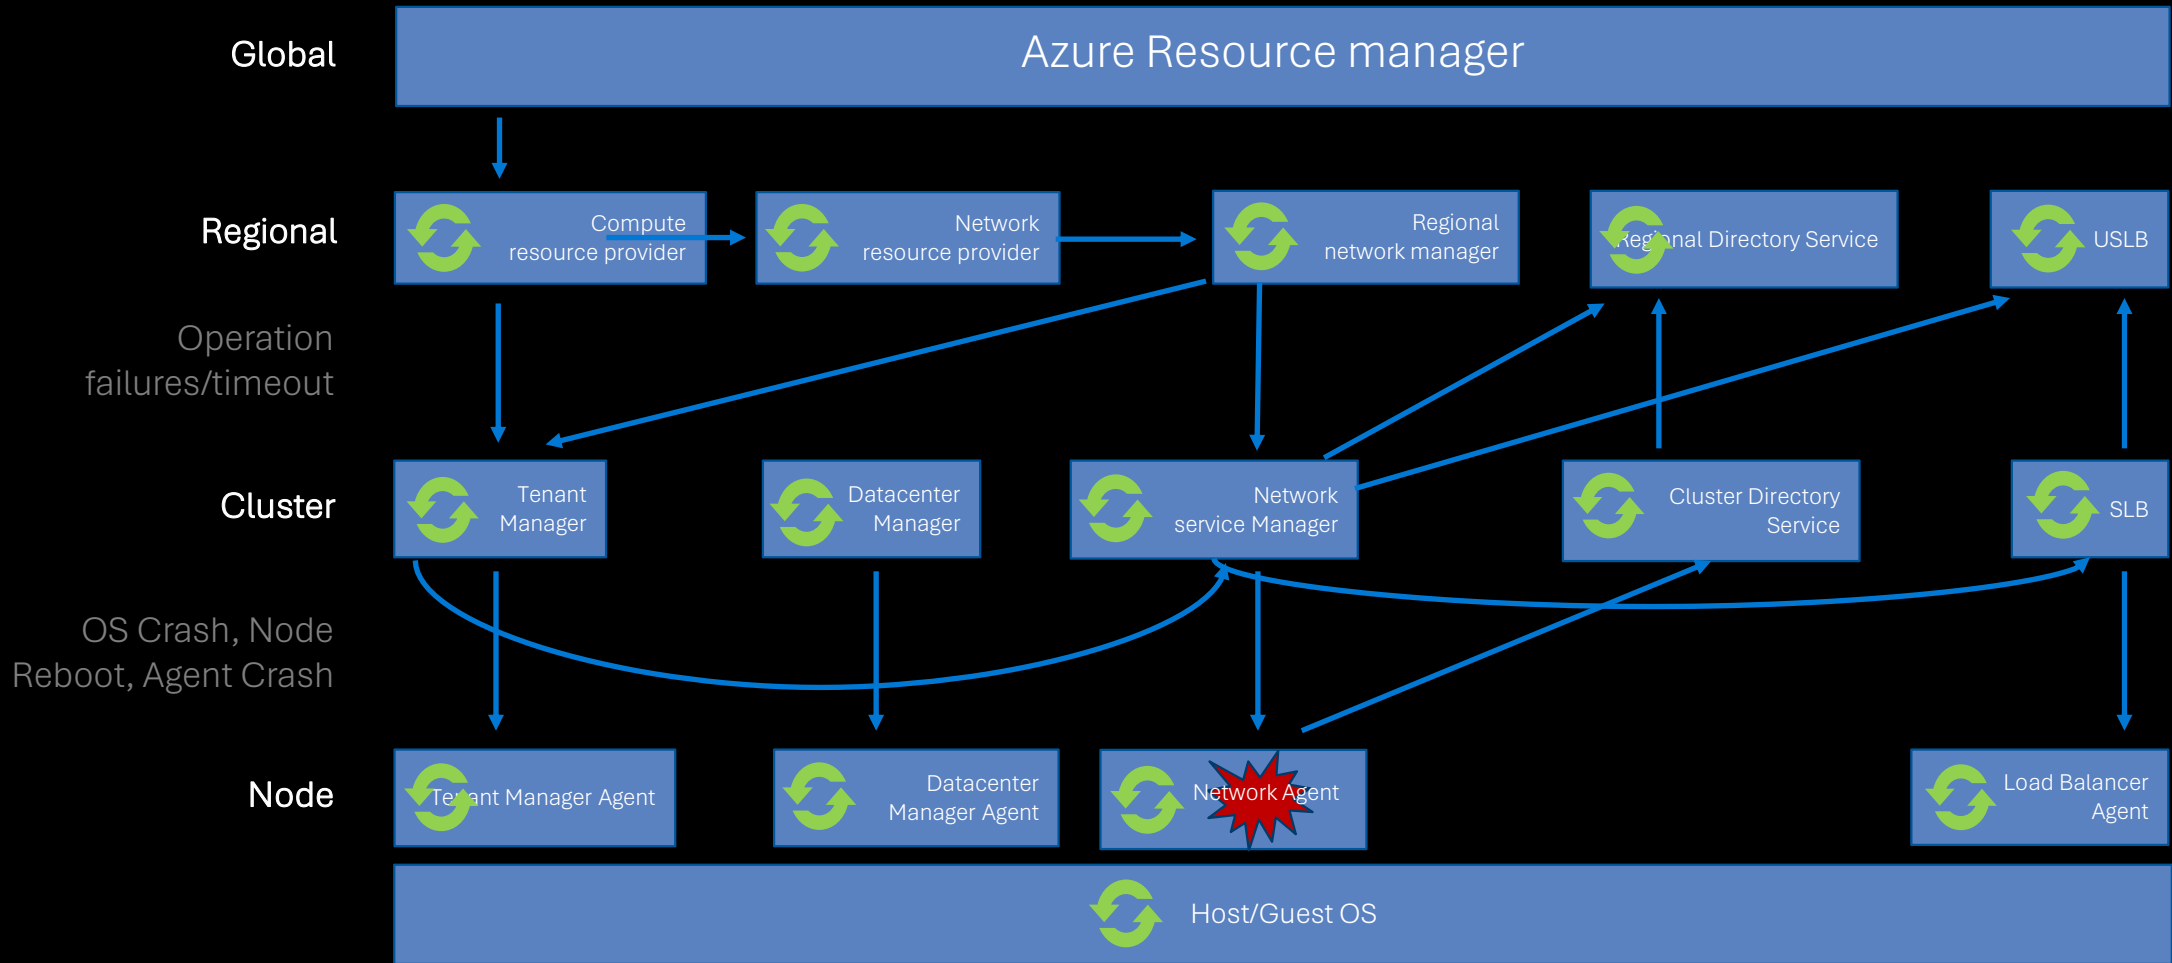
Effective and efficient action-taking: integrating intelligence into Azure DevOps

- Effective monitoring and diagnosis

- Thousands of high-quality tickets filed every year

- Increased deployment velocity by ~4 times


Preventing Code Regressions: Challenges



Increasing Host Machine Resilience




Goal – minimize VM reboots due to host failures by triggering Live Migration (moving VMs to healthy node with only a few seconds of blackout time) and other protection methods



Increasing Host Machine Resilience (Cont'd)

Project Narya

Predictive and adaptive failure prevention



Further read

- Azure blog: <https://azure.microsoft.com/en-us/blog/advancing-failure-prediction-and-mitigation-introducing-narya/>
- Sebastien Levy, et. al., Predictive and Adaptive Failure Mitigation to Avert Production Cloud VM Interruptions, OSDI 2020

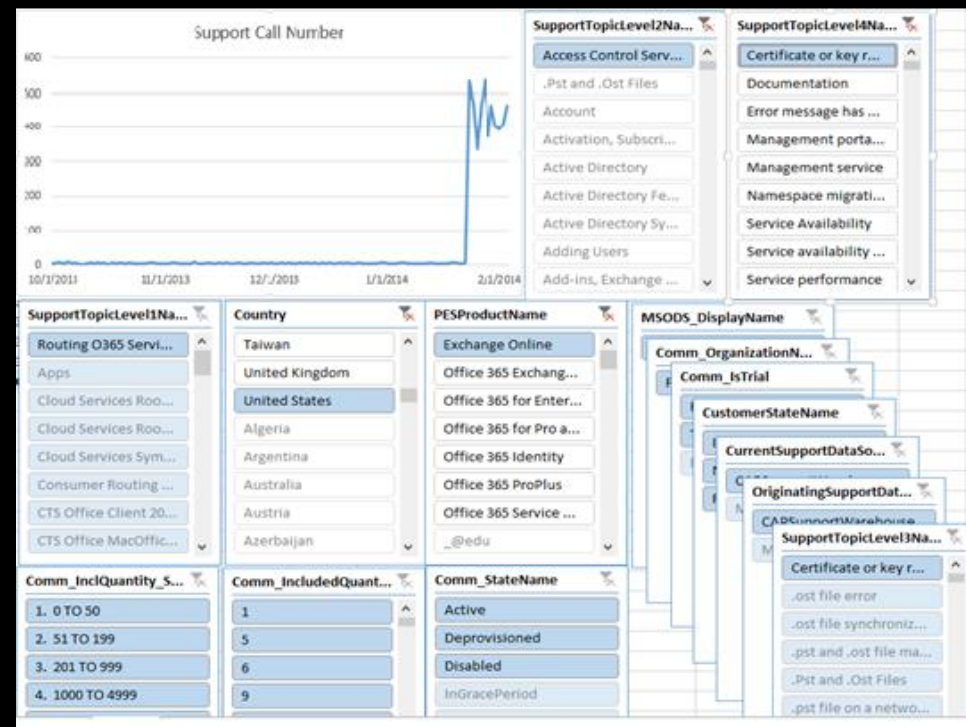
Multi-dimensional Anomaly Detection

Common Practice

- Manually identify monitor combinations with pivot table
- Set up pipelines to monitor hundreds of thousands of time series
- Total Time Series: 100,000+*

Our solution


- Formulated as a "combinatorial optimization problem"
- Solved by a specific-tailored "meta-heuristic search" method
- Details see the paper* from our Microsoft Research partners



*paper: "Efficient Incident Identification from Multi-dimensional Issue Reports via Meta-heuristic Search", FSE 2020


Multi-dimensional Anomaly Detection - Motivation

(Example case of Azure VM Interruptions)



Time	Interruption
2019-04-01	100
2019-04-02	99
2019-04-03	103
2019-04-04	97
2019-04-05	103
2019-04-06	99
2019-04-07	98
.....

Overall KPI not having obvious spike



Time	Country	Datacenter	Disk Type	Interruption
2019-04-10	USA	DM1	SSD	1
2019-04-10	Australia	MEL21	SSD	1
2019-04-10	USA	DC1	HDD	4
2019-04-10	India	BL1	SSD	10
2019-04-10	UK	SN6	Hybrid	3
2019-04-10	USA	DM1	HDD	0
.....

Spike observed in a particular pivot

AI Ops and LLM (Large Language Model)

- **AskBrain Copilot:** Infuse generative AI into how we design, build, and operate cloud services for delightful customer experience and engineering efficiency



AIOps Benefits & Results

- First version of Brain deployed in Production in early 2019
- Major Time To <x> (TTx) improvement
- Incident/alert auto-correlation -> Less noise

72%

TTM Reduction

58%

TTN Reduction

100%

Auto-Comm
percentage
increase

25%

Incident noise
reduction

98.26%

Detection Recall

98.83%

Detection Precision


Our journey has just started.

Brain's mission is to leverage AI & automation to significantly reduce impact due to outages & eliminate noise.



Q&A

Invent with purpose.



Microsoft