# Operations @ Scale

Dean Lorenz, Eran Raichstein,
Hillel Kolodner, Kathy Barabash,
Liran Schour, Shelly Garion

## The Challenge: Operating a Cloud-Scale Service

**Monitoring Today**

**Volume & Velocity**: scale of data, rate of data
- 1,000,000s of requests per second, 10,000s of servers & network elements,
- 100s of services, multiple data centers, multiple layers, PBs of data

**Variety & Veracity**: multiple sources & forms of data, uncertainty & noise
- Logs, app metrics, HW telemetry, network (sflow/netflow), complaints
- Multiple systems, owners & admins (app, server, storage, net)
- Multiple paths, alternative systems – built for failure HA approach

**Operations Today**

**Can show what is happening, what has happened**
- Collect all data, search, visualize, alert based on static thresholds

**But… too much Information, too much hands on, too much expertise**
- Lots of dashboards, reports, and alerts
- Hard to understand, opaque, requires interdisciplinary DevOps skills
- Important problems take too long to solve or are overlooked
- Trends and imminent problems are not detected before service is affected

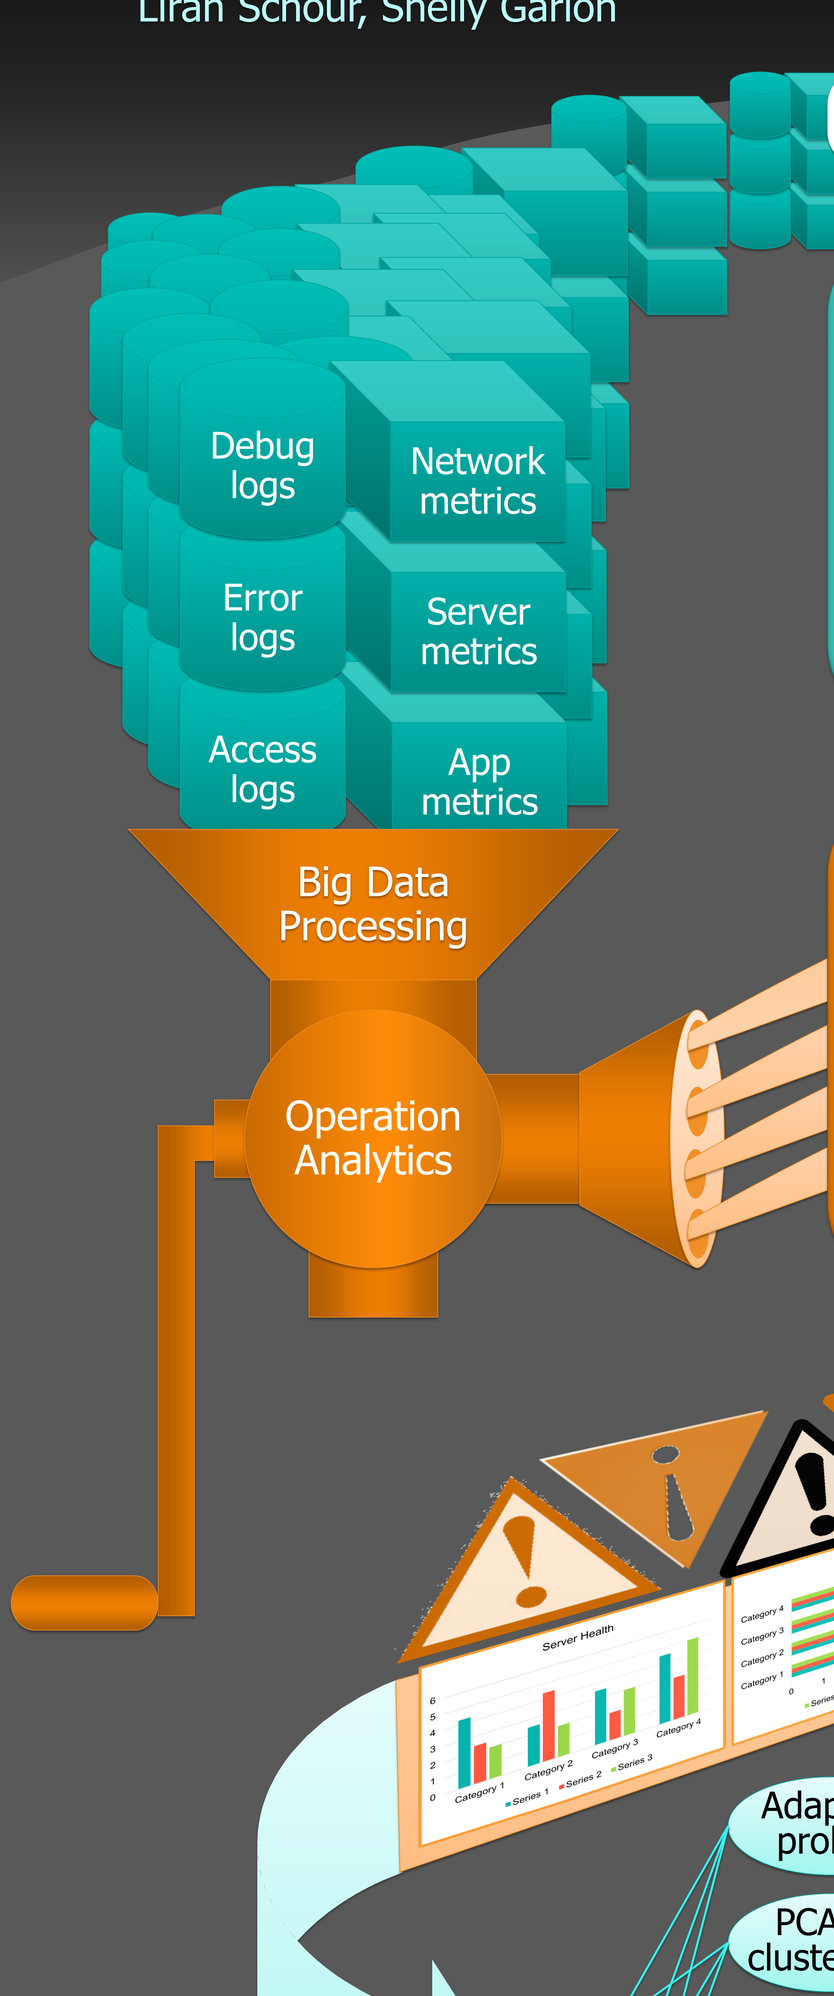


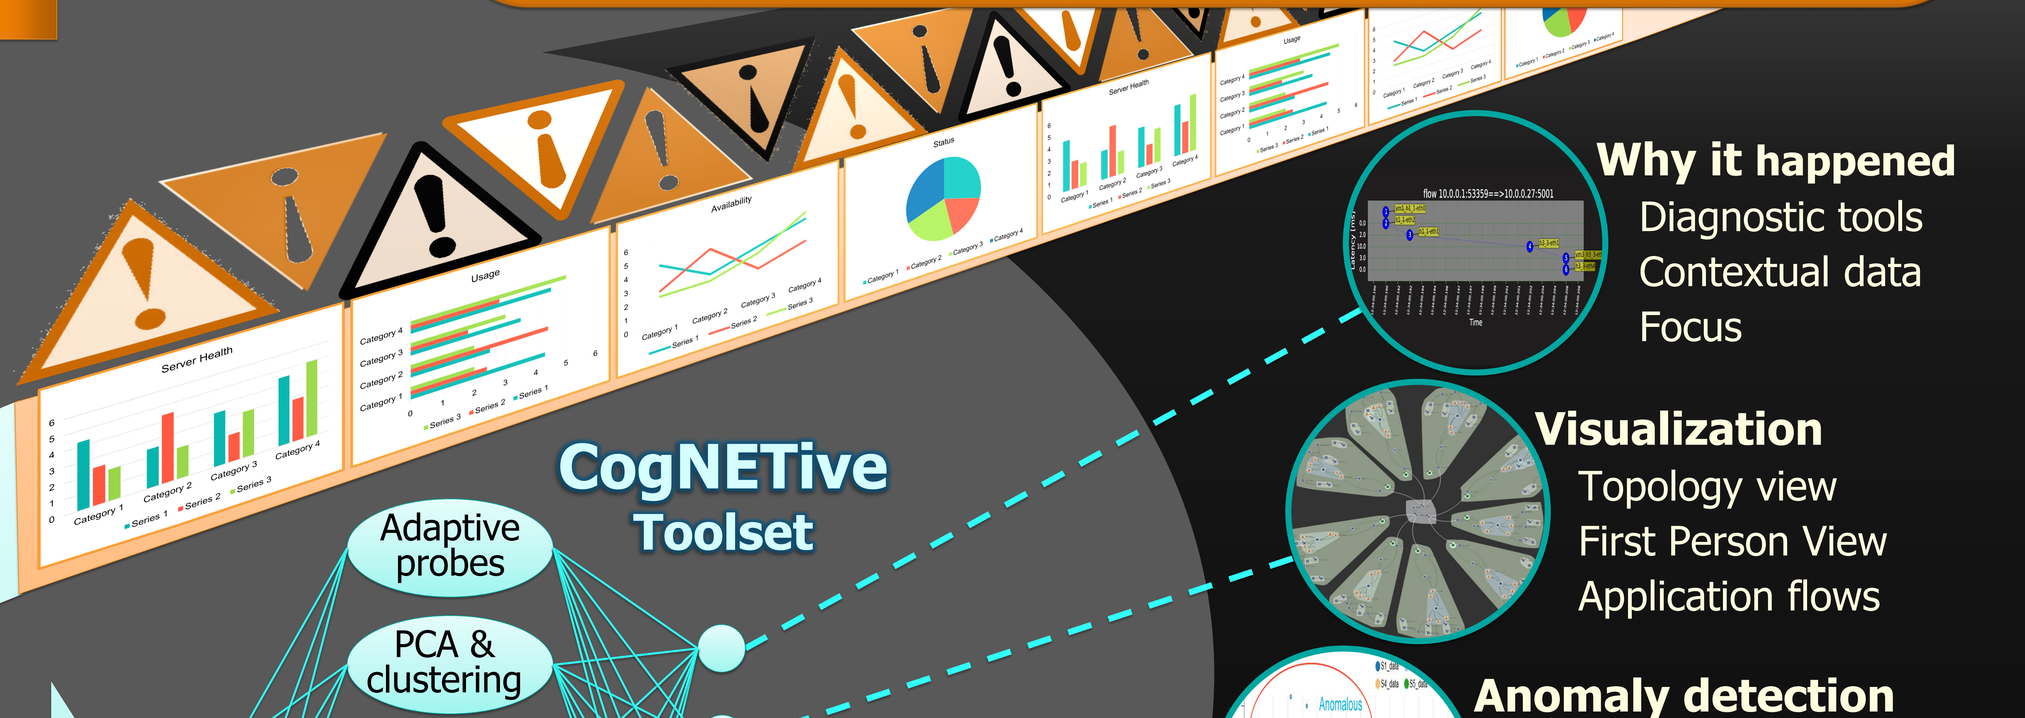


## CogNETive Toolset

- Adaptive probes
- PCA & clustering
- Robust adaptive filtering
- Skydive
- Time series
- Network analytics
- Machine learning

**IBM Cloud**
Object Storage

**BEACON**
Federated Cloud Networking




**Why it happened**
- Diagnostic tools
- Contextual data
- Focus



**Visualization**
- Topology view
- First Person View
- Application flows

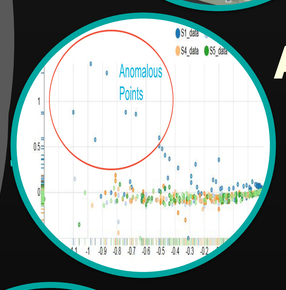


**Anomaly detection**
- Automatic monitoring
- Smart, real-time alerts
- Dynamic thresholds

**Cognitive**
- Business value
- Insights to system behavior
- Relationships between metrics

**What may happen**
- Predictions, forecasts
- Models, risks, priorities
- Implications

**What can I do**
- Prescriptive, informed decisions
- Proactively minimize problems
- Remedies, avoidance, adaptations

**IBM Research – Haifa**
Computing as a Service
Cloud Architectures and Networking