

Application Performance in the Cloud


Understanding and ensuring application performance in highly elastic environments

Albert Mavashev, CTO
Nastel Technologies, Inc.
amavashev@nastel.com

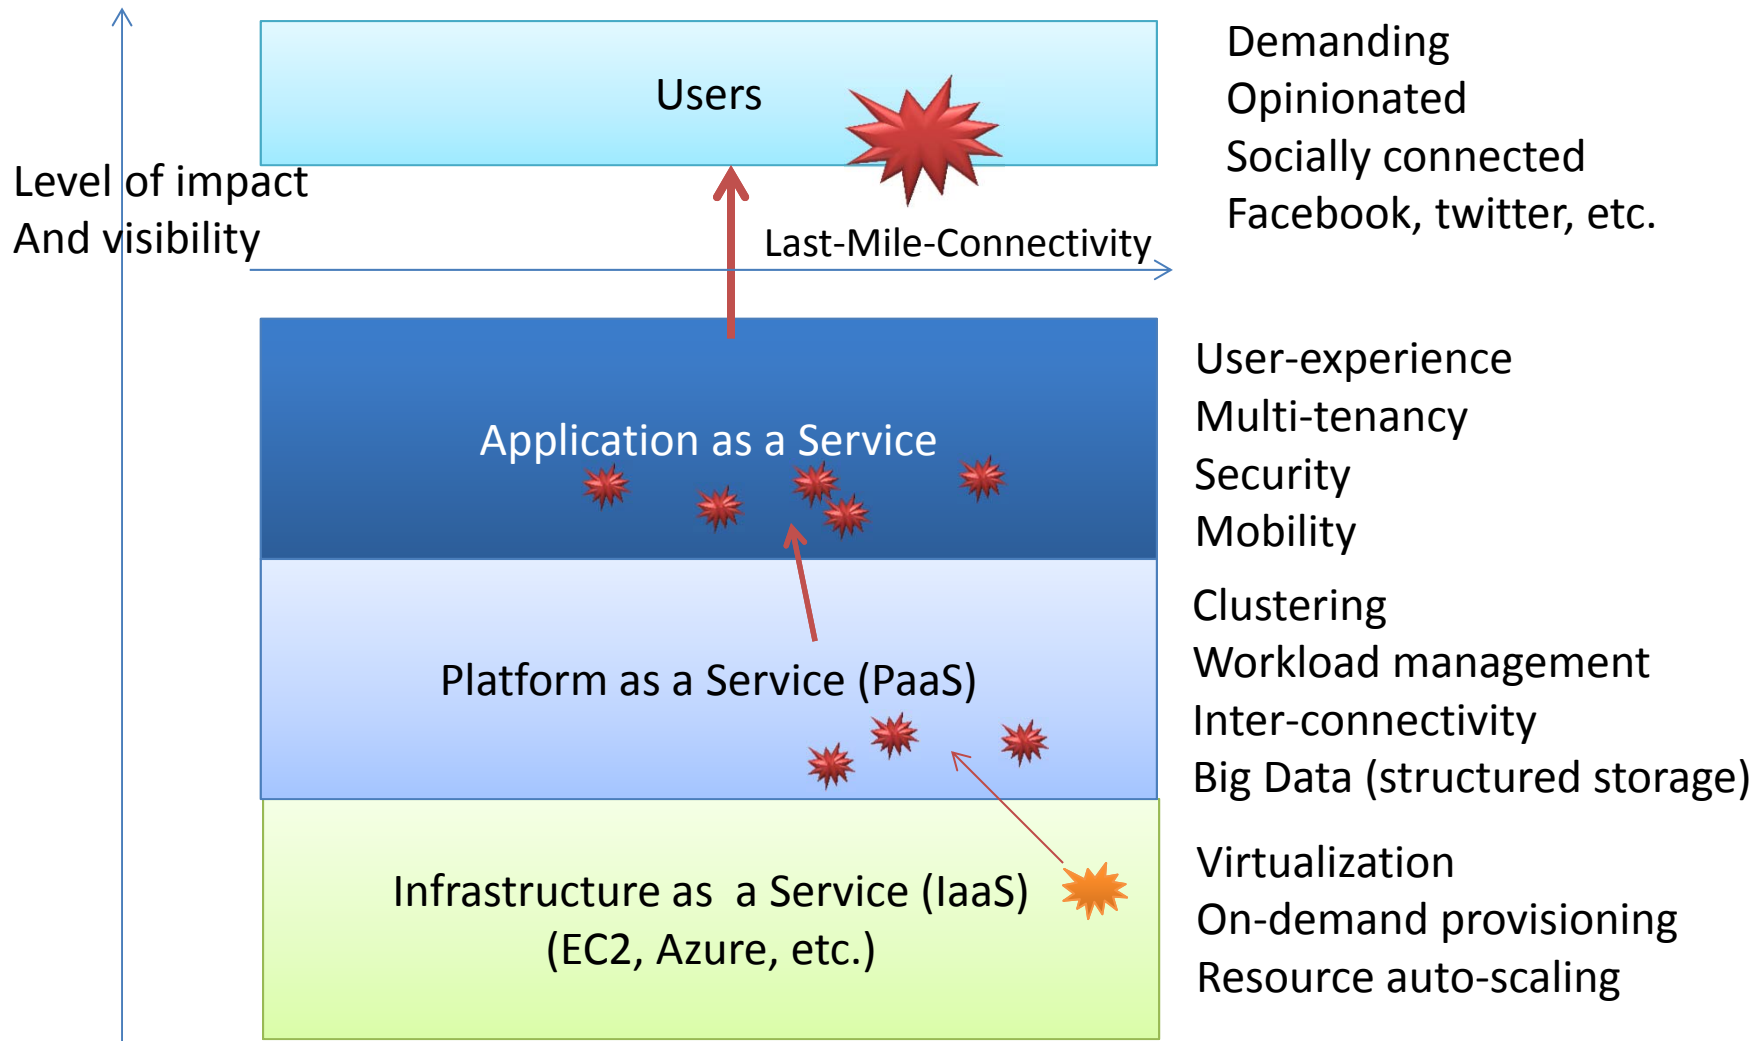
What is Cloud?

- Typically defined by its properties
 - Elastic – on-demand scalability
 - Agile – quick to deploy provision
 - Reliable – business continuity, ability to survive failures
 - Maintenance – single deployment, no need to deploy on each user
 - Independence -- device, hardware, location independence
 - Cost – typically pay-per-use
 - **Performance – could be very well defined for IaaS, less so for PaaS, and even less so for SaaS**
- Clouds can be public, private, hybrid, federated
 - Hybrid and federated are the most complex
 - Public and Private -- off premise vs. on-premise (similar but differ in delivery mechanism)

What does it mean Application Performance in the Cloud?

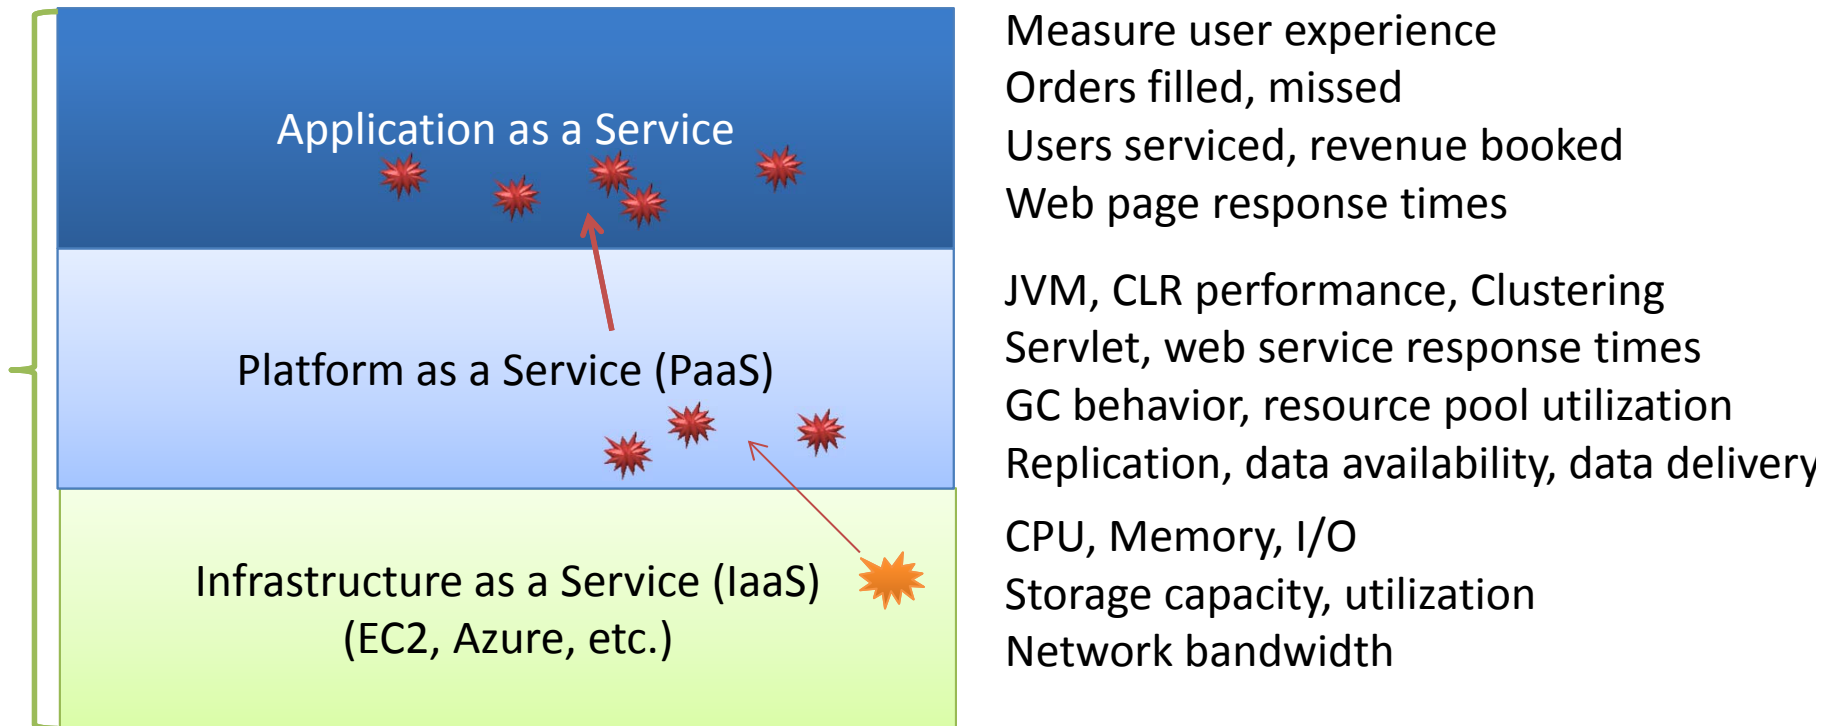


Application in the Cloud



Application Performance: Take 1

Monitoring Performance of each stack




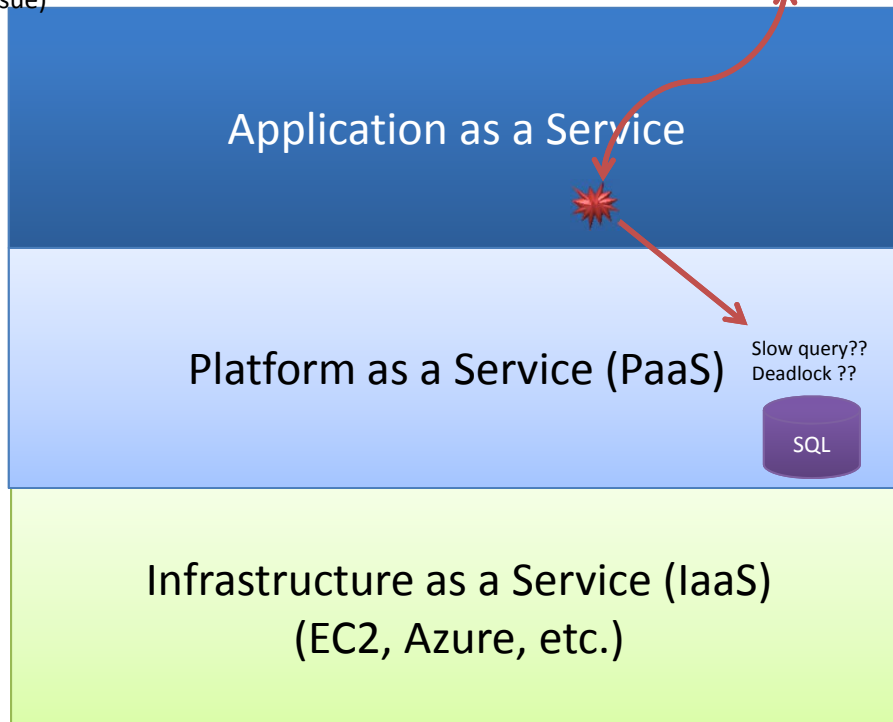
Application Performance: Take 2

What happens when problems are between stacks?

What if performance impact is data dependent?

What about Last-Mile-Connectivity?

(May be not so much in US, but other countries LMC is a big issue)



User experiences a timeout



All SaaS indicators are normal



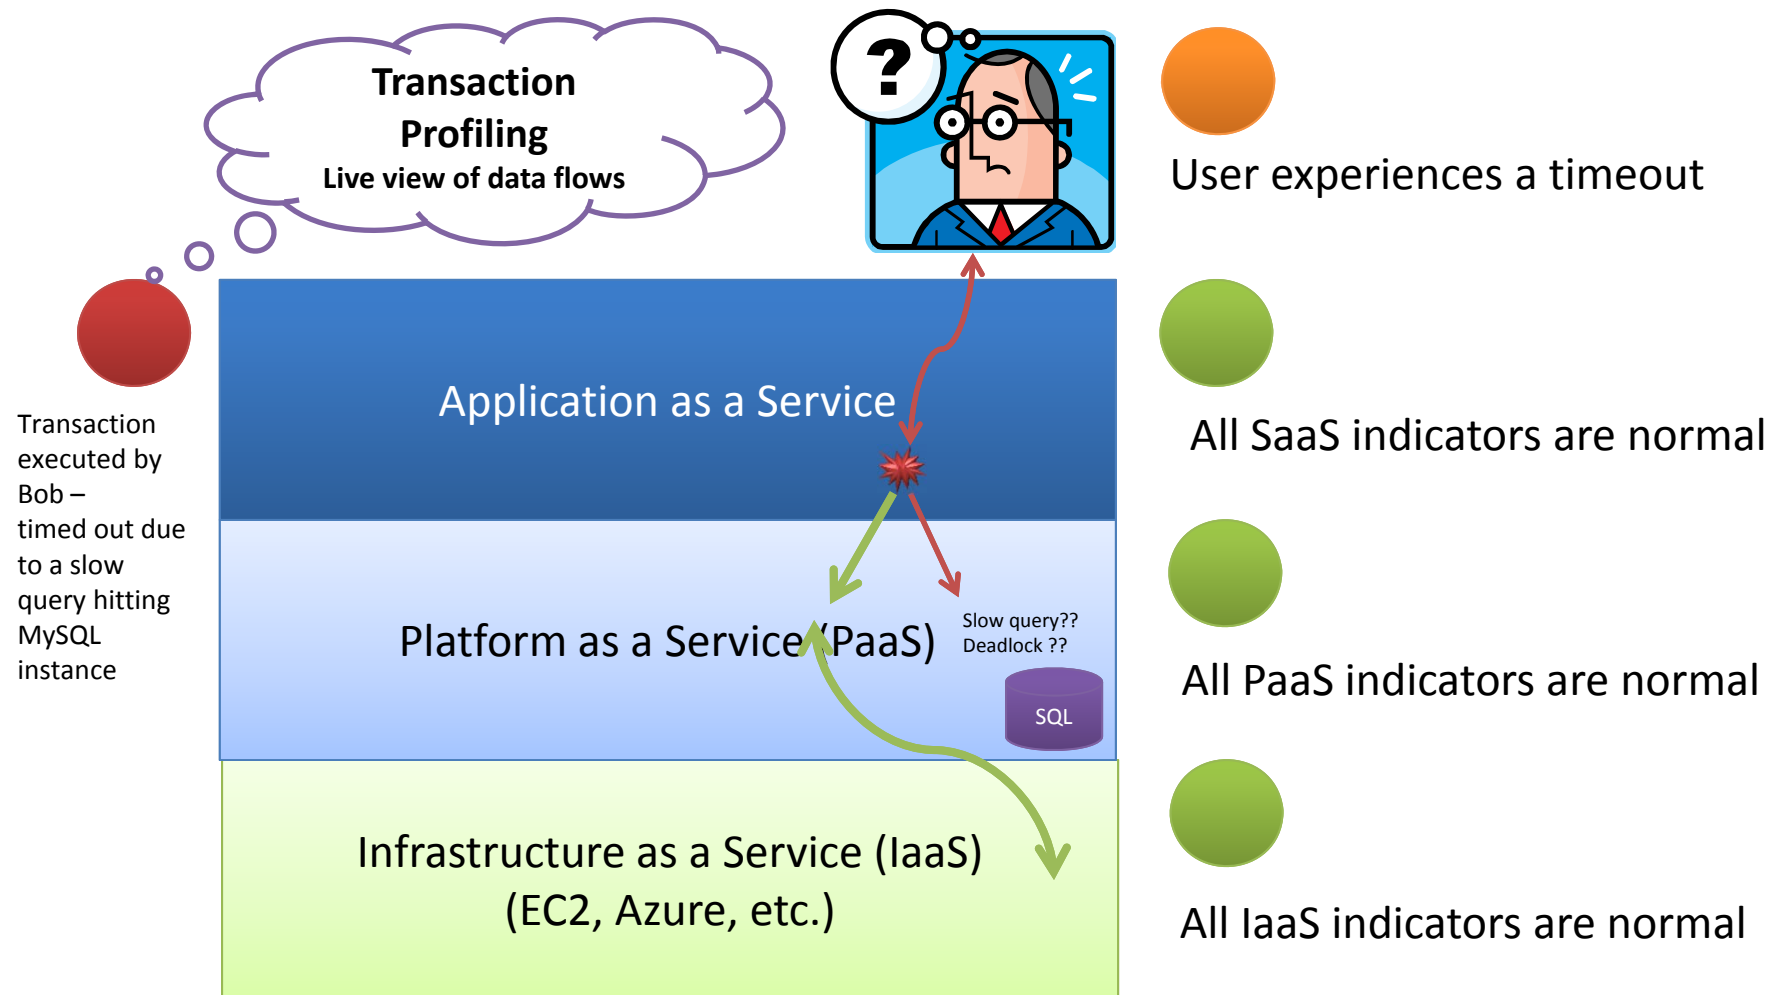
All PaaS indicators are normal



All IaaS indicators are normal

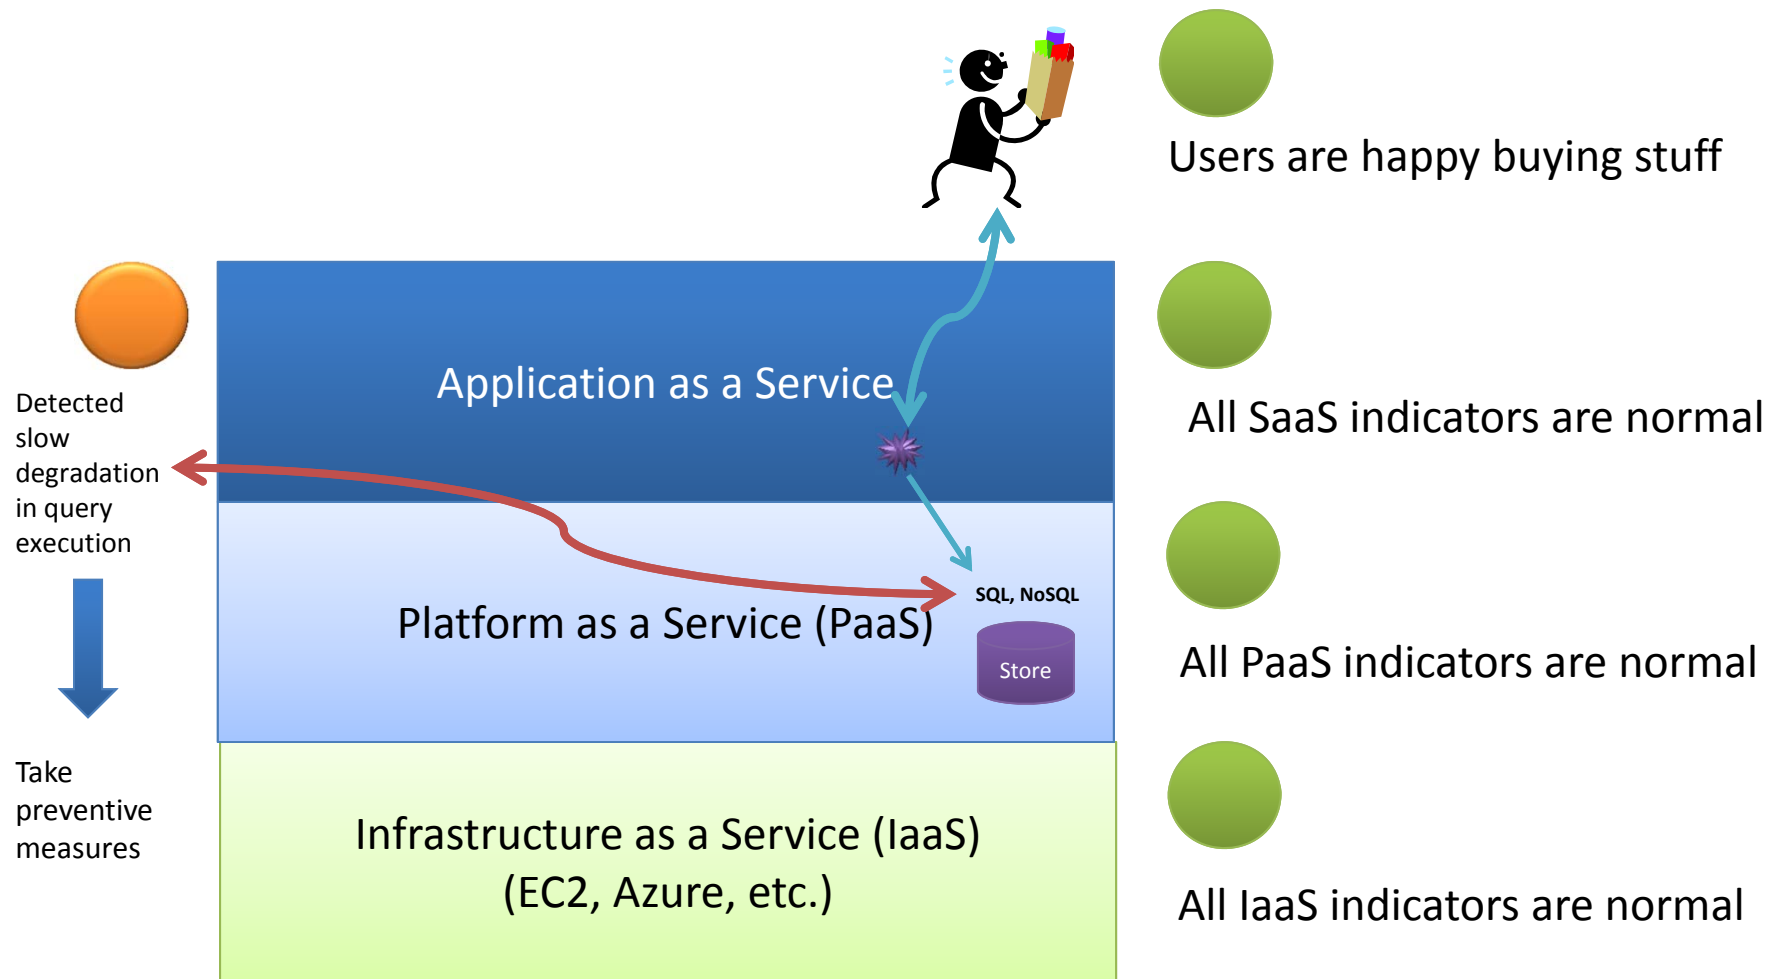
Application Performance: Take 3

Monitoring each stack as well as transactions that flow between stacks?

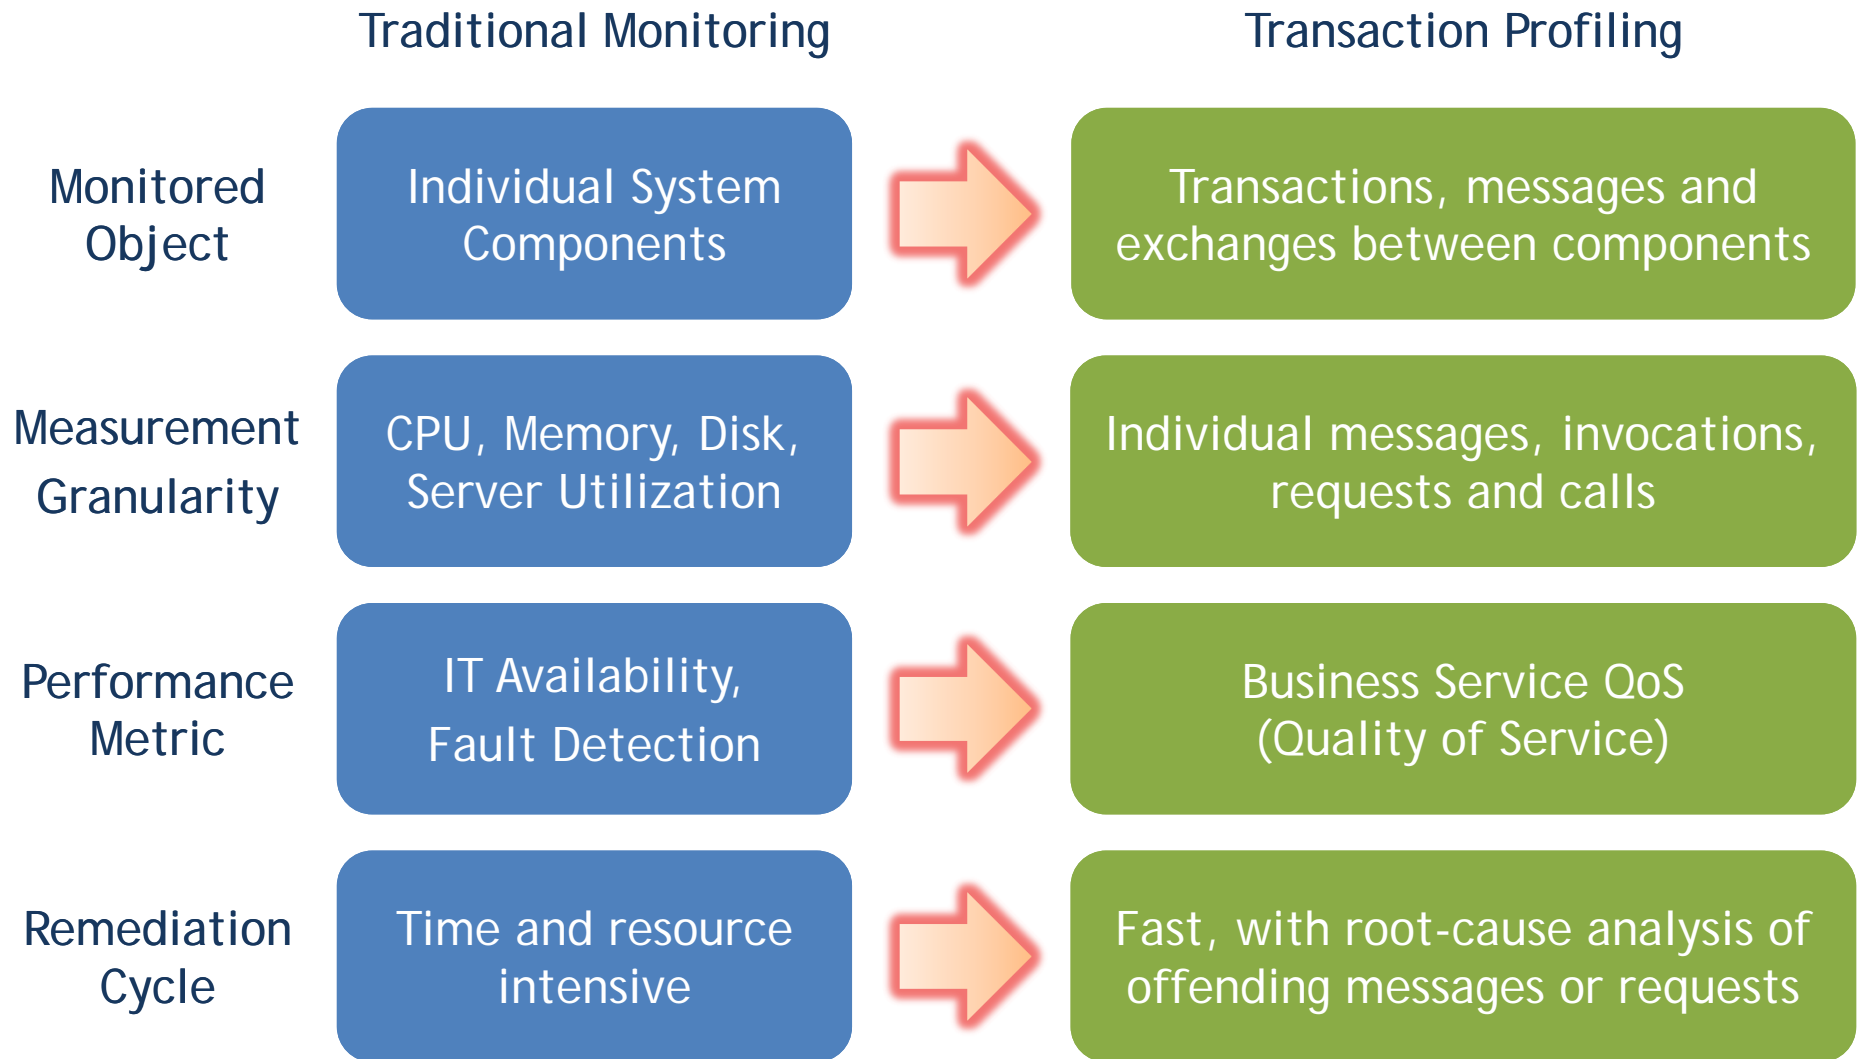


Application Performance: Take 4

What about proactively monitoring trends before service degrades?



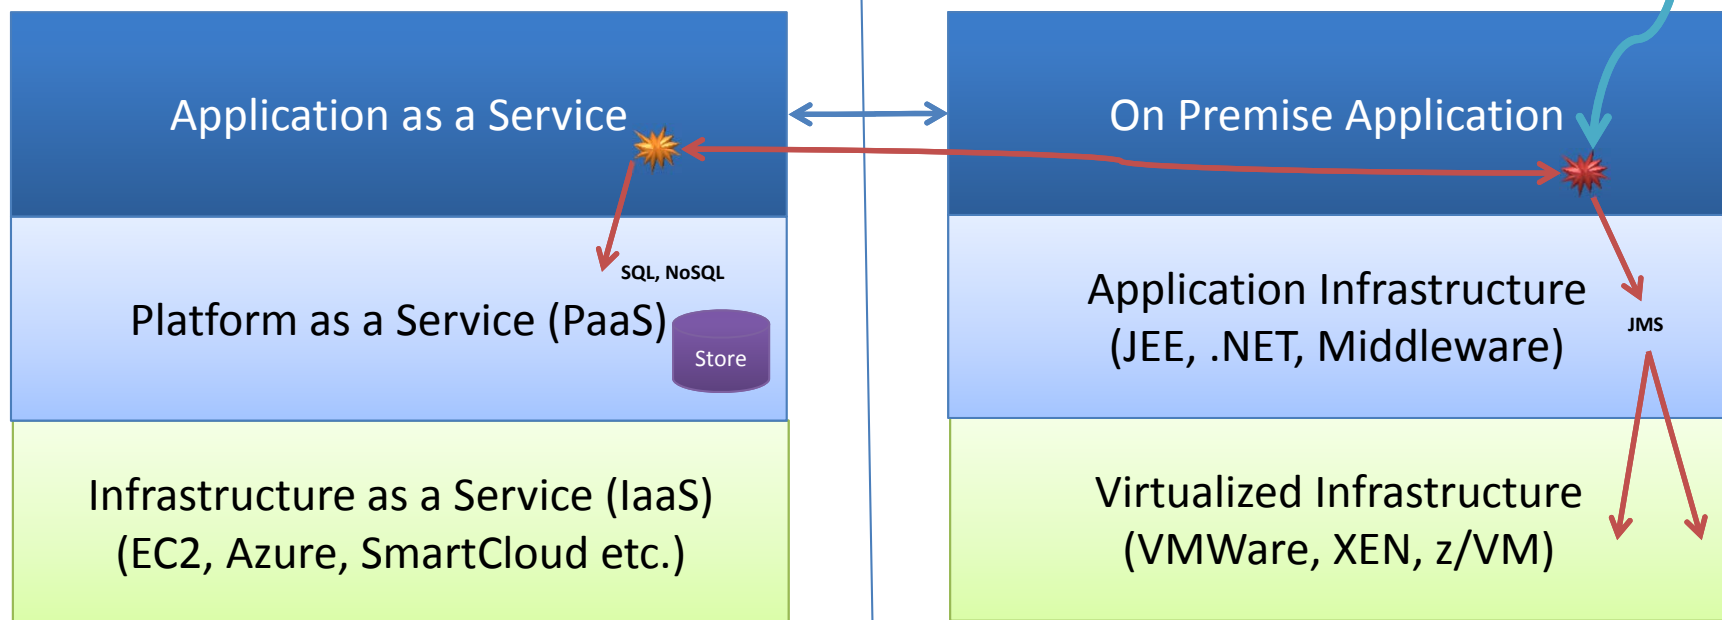
Traditional Monitoring vs. Transaction Profiling



What about Hybrid deployments?

(More complex)

Last-Mile-Connectivity



- On-premise and off-premise (Cloud Service) resources need to be monitored
- Monitoring itself could be delivered as on-premise or off-premise
- Monitoring instrumentation must be cloud/firewall ready/friendly (HTTP/HTTPS)
- Inter-cloud interactions need to be monitored as well

USE CASE: TREND ANALYSIS

Ways to detect performance trends

- Measure relevant application performance indicators
 - Orders filled, failed, missed
 - JMV GC activity, memory, I/O
- Create a base line for each relevant indicator
 - 1-60 sampling for near real-time baseline
 - 1, 10, 15 min → daily, weekly, monthly for short, long term baseline
 - Samples can range anywhere from 1-60 seconds depending on level of required resolution
- Apply analytics to determine trends and behavior
 - Can vary from simple to complex
 - **Prefer KISS approach (Keep It Simple and Stupid)**

3 Simple methods to detect trends (No complex math required)

- **Bollinger Bands**

- Determine high and low bands based on available baseline
- Defines a normal channel which is typically within 2 standard deviations from the mean
- Compute STDDEV, Mean, Current sample

- **% Change**

- Sample to sample, day-to-day, week-to-week, etc.

- **Velocity**

- Number of measured units per unit of time (example: response time drops from 10 to 20 seconds over 5 sec interval – means $(20-10/5)=2$ units/sec.

Typical Usage

- **High Band**
 - Given a set of metrics, alert when one or more are above High band for at least 2+ samples
 - Indication of abnormal activity over a period of time
 - **Caution: abnormal can become the new normal**
- **% Change**
 - Useful indicator for near real-time monitoring of resources (such as heap, memory, CPU, storage)
 - Useful indicator for long term trends (daily, weekly)
- **Velocity**
 - Very useful for monitoring metrics that measure usage of resource that have a finite upper bound (memory, storage, table space etc.)
 - Measuring velocity can help measure when upper limits can be reached

Required instrumentation

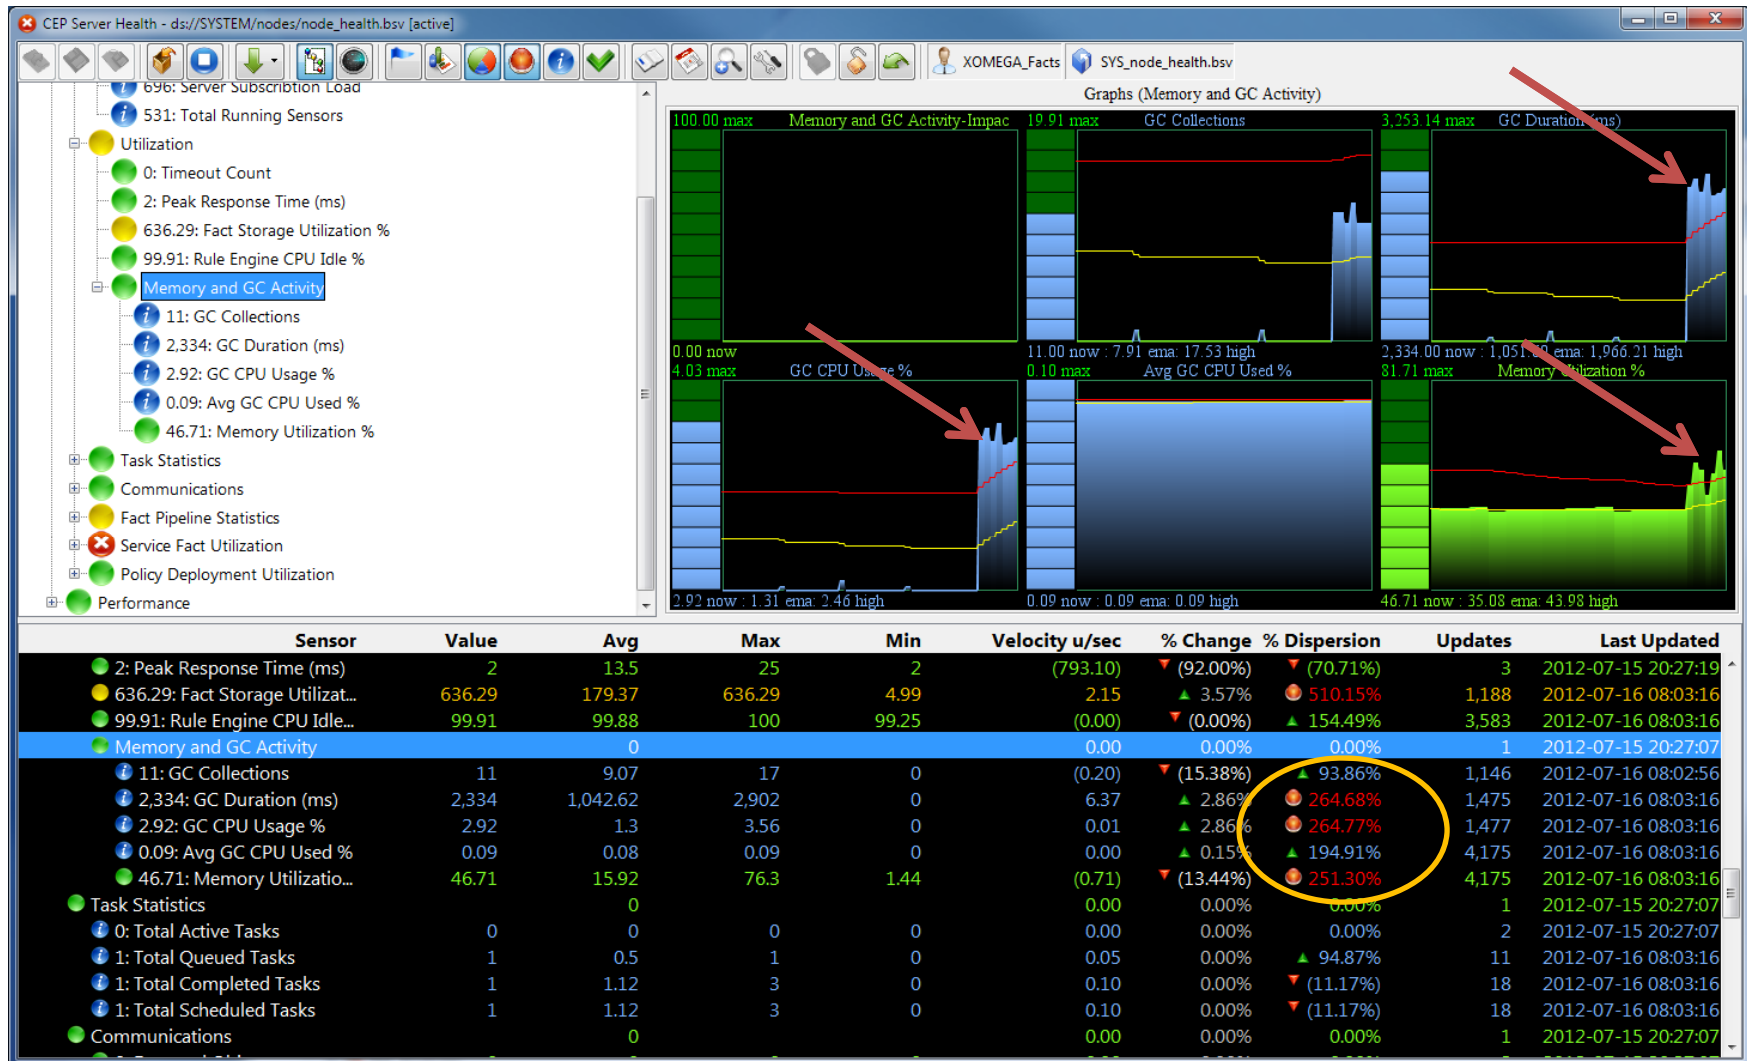
- **Data collectors**
 - Attempt to collect all relevant indicators within the same time tick
 - Response time, GC activity, memory usage, CPU usage
- **Build a history for each collected metric**
 - Either in memory for near real-time analysis
 - Storage for short, long term (min, hours, days)
- **Pattern matching, analytics**
 - Need to scan and pattern match application metrics (such as find all applications whose GC is above High Bollinger Band for 2+ samples)
 - Run as a continuous query, which is executed as metrics are collected and updated
- **Actionable Outcome**
 - Alerts, notifications, actions
 - Visualization, dashboards

Example: Monitoring Java Application by examining GC Activity

- Java Application running in a standalone JVM container
- Monitoring JVM GC (Garbage Collection) as a byproduct of application activity
 - Sample GC every 10 seconds
 - # GC Samples
 - GC Duration (ms.)
 - GC CPU Usage %
 - Avg. GC CPU Usage (since JVM startup)
 - JVM Heap Utilization %

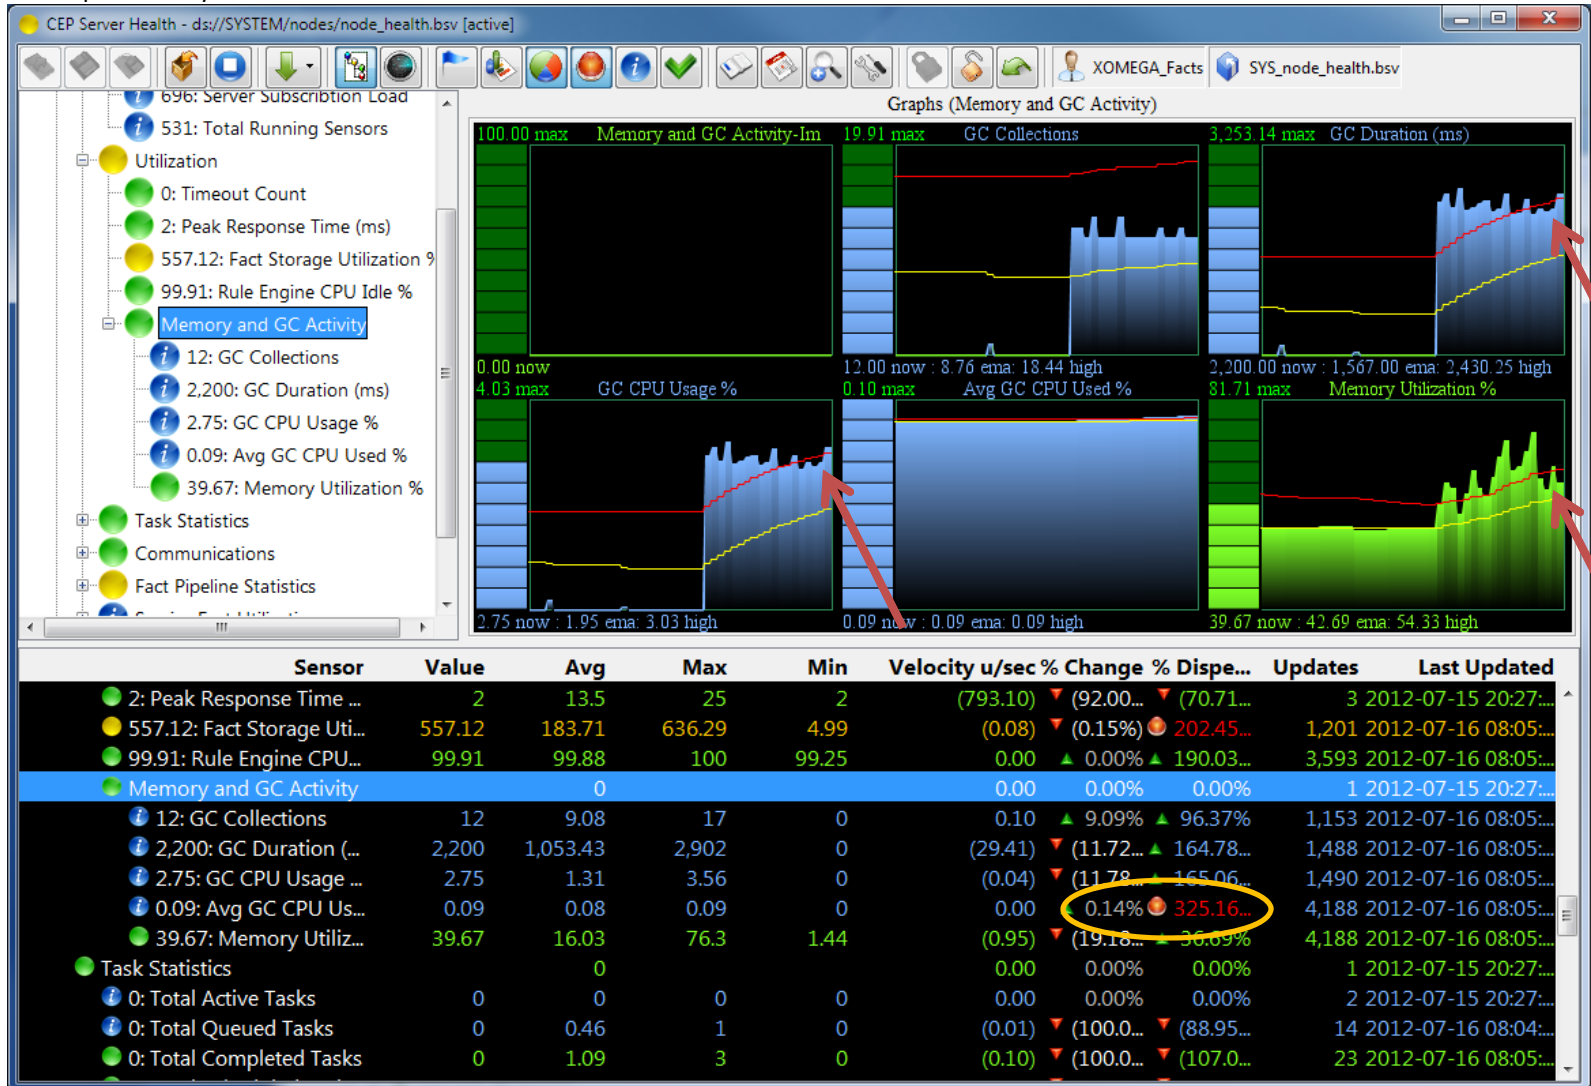
Example 1: Java Application, Sudden Spike in Activity

Data provided by Nastel AutoPilot™



Example 2: Java Application, Adjustment to new workload – The New Normal

Data provided by Nastel AutoPilot™





Resource Leak Detection

Detecting Leaks using Trend Analysis
(Java Example)

Typical causes of Java leaks

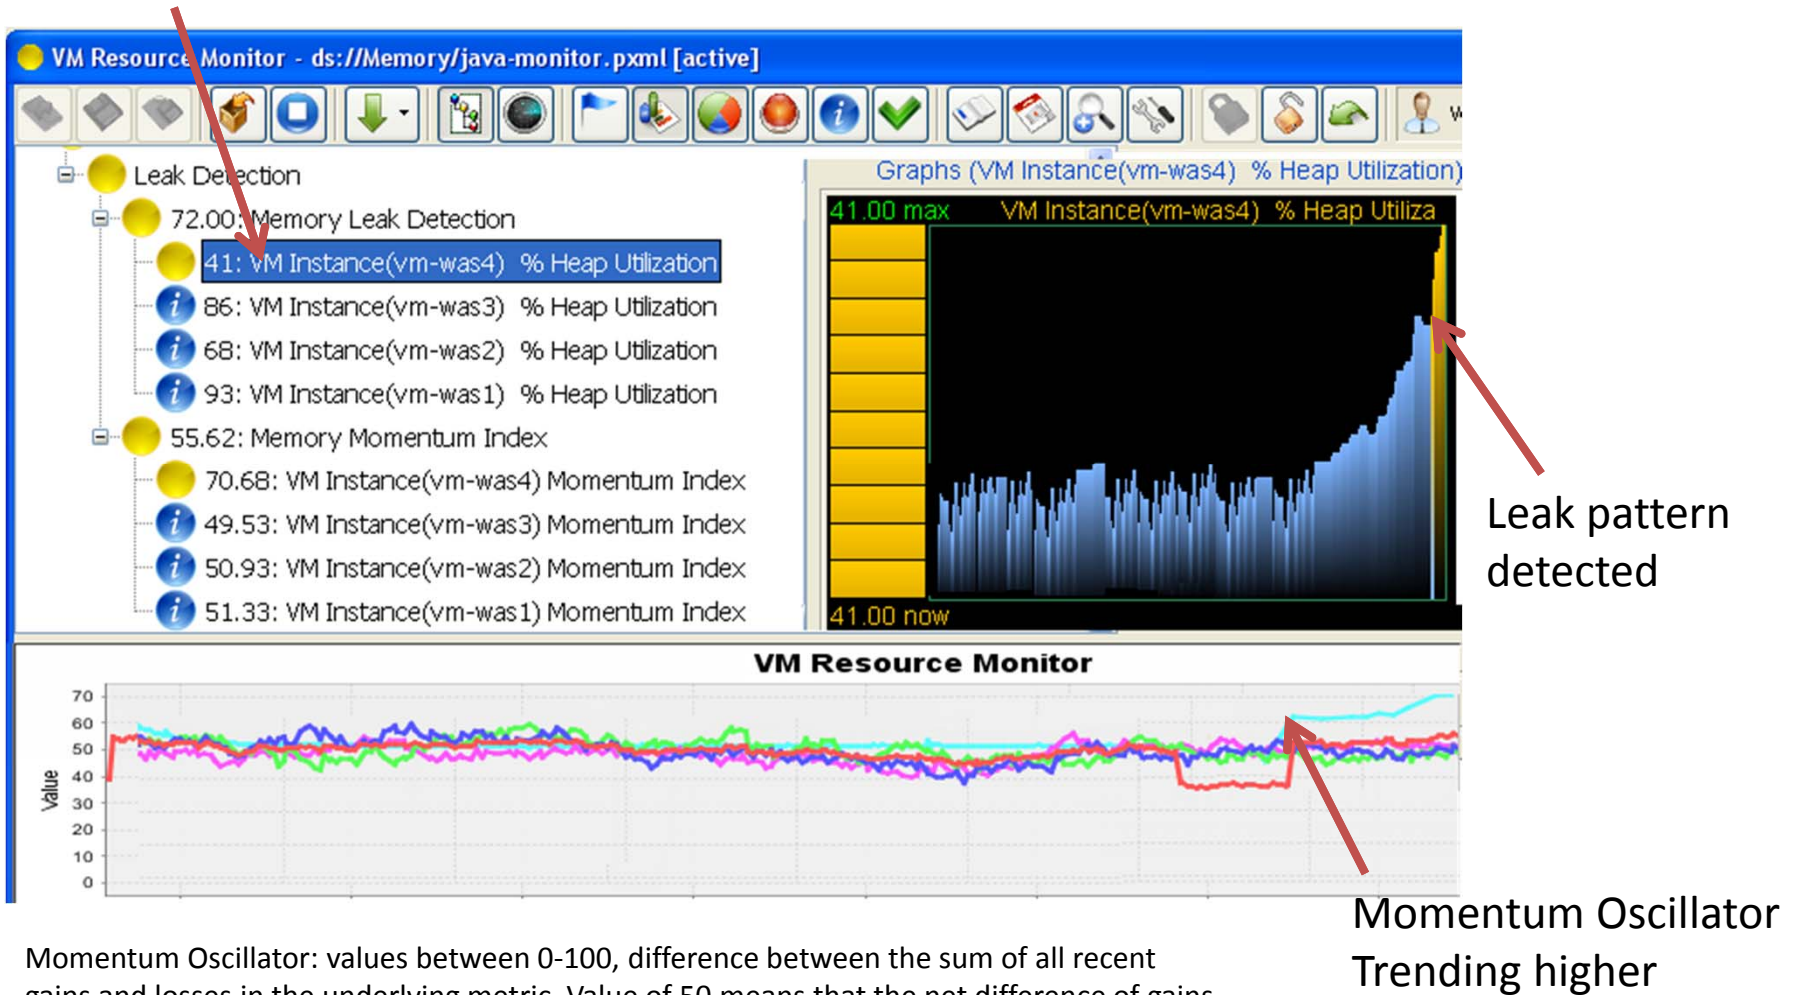
- Programming errors, bugs
 - Unchecked array, list, hash map growth
 - Not closing JDBC Prepared Statements
 - Not closing Sockets, File handles
 - Thread leaks, handle leaks
 - Class loader leaks
 - Resources allocated outside JVM

Leaking Chart Pattern – Detecting Resource Accumulation



Detecting Resource Leaks using Momentum Oscillator

Heap not yet exhausted



Momentum Oscillator: values between 0-100, difference between the sum of all recent gains and losses in the underlying metric. Value of 50 means that the net difference of gains and losses is zero – 0 net gain and loss.

Conclusion: Monitoring Elastic Environments

- **Elastic Applications can't be monitored using static models**
 - Static thresholds
 - Static data/transaction flow models
- **Complex systems layered on top of complex systems**
 - Too many constantly changing variables
 - Makes root cause analysis very difficult
 - Requires extensive cross technology expertise
- **Preferred approach – Holistic Application Monitoring**
 - Granular data collection:
 - Application and infrastructure metrics
 - Analytics, automated base lines
 - Real-time and historical
 - Resource monitoring coupled with Transaction Profiling
 - Visualization that connects different teams:
 - Application support, DevOps, IT Support

For more information

- Visit us at:
 - www.nastel.com
- Questions:
 - info@nastel.com
- Twitter:
 - twitter.com/nastel
- FaceBook:
 - facebook.com/NastelTechnologies
- LinkedIn:
 - linkedin.com/companies/nastel-technologies
- Phone:
 - +1.800.580.2344