

FedCSIS
2015



What has gone wrong in the cloud?

Monitoring the cloud and things that run in it

Wojtek Kozaczynski

Principal Software Development Engineer
Microsoft – Developer Division – Application Insights



Who is Wojtek?

- ⌵ formally educated in Poland
- ⌵ almost entire professional career in the US
 - ⌵ UICC, Andersen Consulting, System S/W Associates, Rational, IBM, Microsoft
- ⌵ Principal S/W Development Engineer
- ⌵ IC (Individual Contributor)
- ⌵ for the last 5+ years working on OS and application monitoring and diagnostics
 - ⌵ WMI, WinRM, PowerShell, ETW, Azure, ...
- ⌵ currently an architect on the ApplicationInsights team, which is developing a public Azure/VSO telemetry service


The shortest version of the talk

1. Real cloud applications get complex quickly
2. The only way to understand what is happening in, or to, a cloud application is to be very deliberate what telemetry data it sends
3. In order to analyze telemetry from a complex cloud application you have to build another cloud application, which may be equally or more complex than the one it monitors

Cloud application



Devices




Services




	Devices	Cloud Services
execution environment	<ul style="list-style-type: none"> • you have no access to the devices • application is self-contained 	<ul style="list-style-type: none"> • you have access to, manage, and pay for the runtime • your application takes dependency on cloud services
application and users	<ul style="list-style-type: none"> • number of instances each owned by a different user • you know very little about the users 	<ul style="list-style-type: none"> • one instance owned by you
telemetry focus	<ul style="list-style-type: none"> • use • usability • failure detection and analysis 	<ul style="list-style-type: none"> • performance • throughput • availability • dependencies on cloud services • resource utilization • failure detection and analysis
telemetry data	<ul style="list-style-type: none"> • interrupted flow • small per instance volume • possibly large cumulative volume 	<ul style="list-style-type: none"> • continues flow • large volumes • about application, execution environment, dependencies

Telemetry service



Telemetry service == cloud application that

- collects
- processes, and
- presents data about other cloud applications



Telemetry Service



About the telemetry data

- ④ What should it look like?
 - ④ what should it represent?
 - ④ how should it represent it?
- ④ How can we make all devices and servers send consistent data so we can
 - ④ collect
 - ④ store, and
 - ④ processit consistently?

Common telemetry representation

- ④ Commonly telemetry items contain two parts
 - ④ data about context in which the item was generated
 - ④ data item itself
- ④ In practice there are only four basic telemetry types
 - ④ measurements
 - ④ single values or aggregations
 - ④ events
 - ④ named collections of properties representing events
 - ④ trace messages
 - ④ crash data
 - ④ exception trees, call stacks, memory dumps

Open source SDKs == a way to consistency

⌵ .Net

⌵ Java

⌵ JavaScript

⌵ node.js

⌵ iOS

⌵ Android

⌵ C++

⌵ PHP

⌵ Ruby

⌵ Python

⌵ OSX

⌵ Server ASP .Net

⌵ Unity

⌵ Xamarin

⌵ Angular JS

[https://github.com/Microsoft/
ApplicationInsights-Home](https://github.com/Microsoft/ApplicationInsights-Home)

README.md

ApplicationInsights-Home

Welcome to the central repository for Application Insights projects. We intend to include any content that applies to all our repositories here, while specific content to each project will be found with that project.

More information about Application Insights can be found [here](#).

Our Repositories

[Android](#)

[Asp.Net v5](#)

[C++ Universal apps](#)

[DotNet Core](#)

[iOS](#)

[JavaScript](#)

[Java](#)

[Node.js](#)

[OSX](#)

[PHP](#)

[Python](#)

[Ruby](#)

WordPress on [GitHub](#) and Release Version on [Wordpress.org](#)

[Xamarin](#)

Python

```
from applicationinsights import TelemetryClient
tc = TelemetryClient('3863d2b7-643a-4c72-a4fd-a34bf90ff441')
tc.track_event('CowSold', { 'breed': 'Lowline' }, { 'price': 42 })
tc.flush()
```

.Net

```
TelemetryClient tc = new TelemetryClient();
tc.Context.InstrumentationKey = "3863d2b7-643a-4c72-a4fd-a34bf90ff441";
```

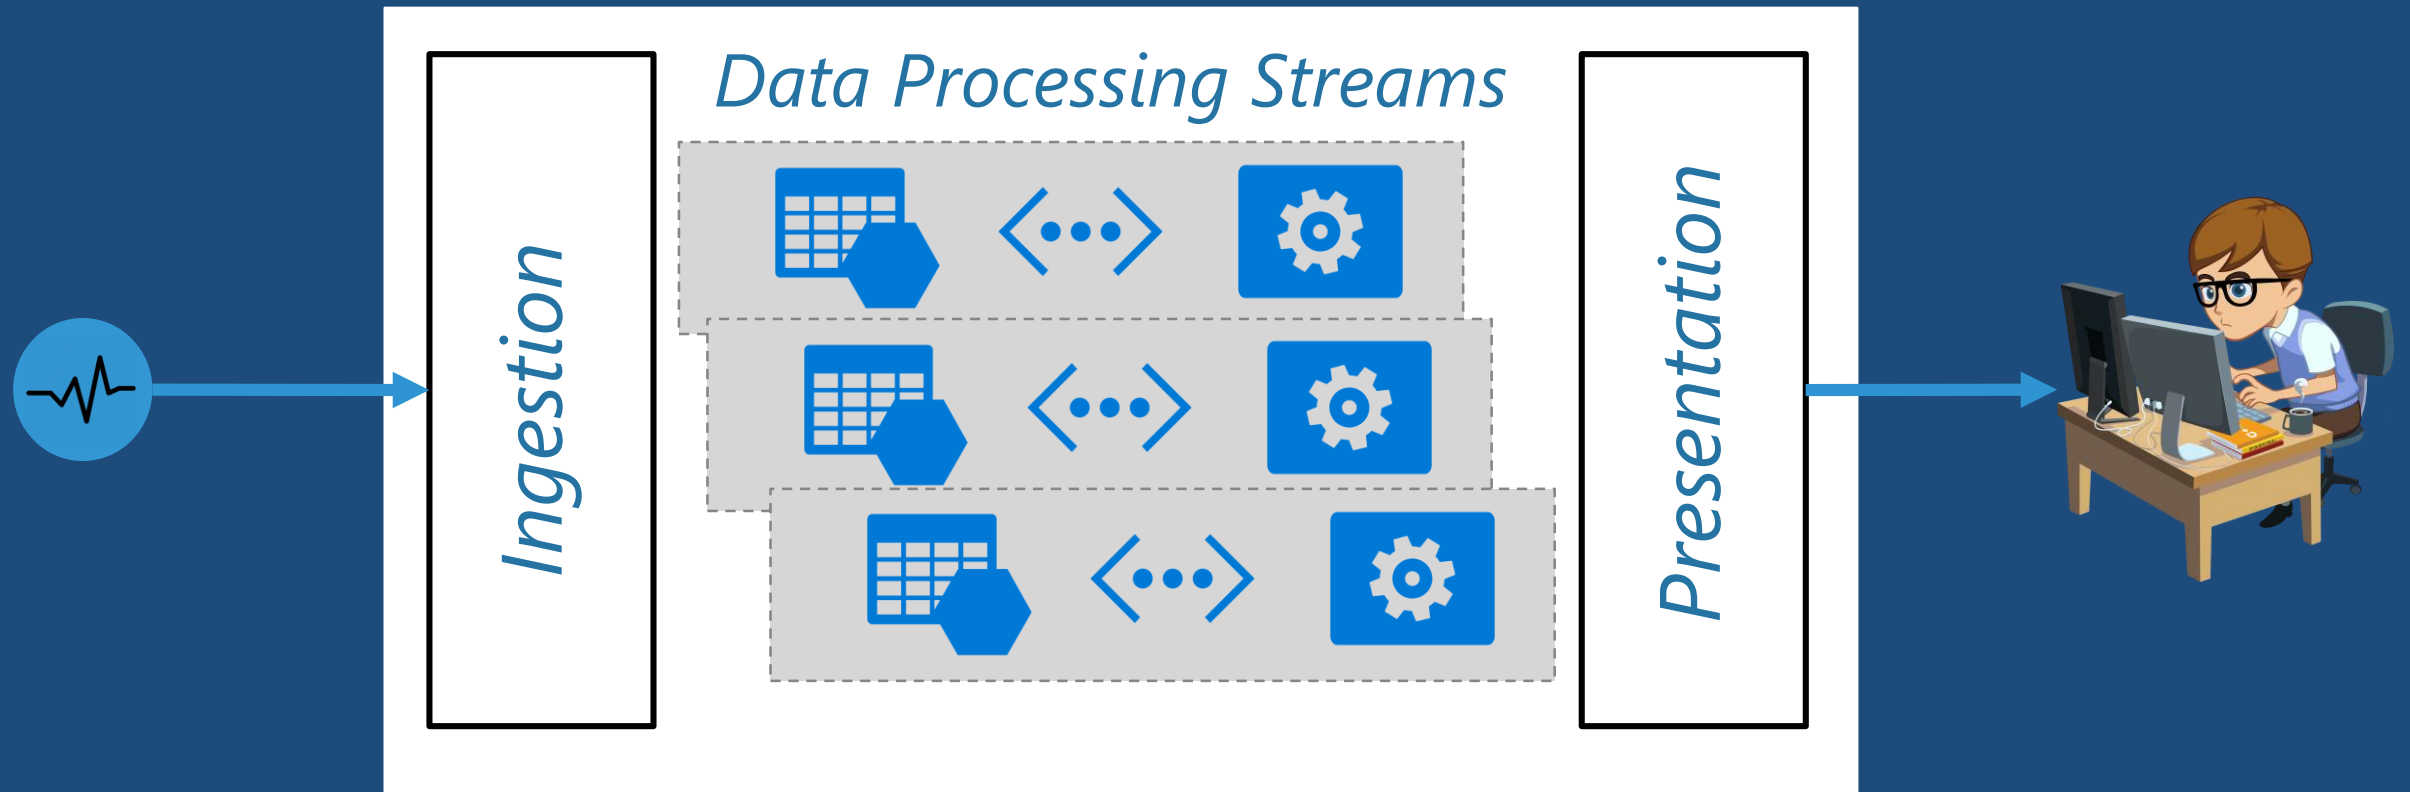
```
EventTelemetry et = new EventTelemetry();
et.Properties.Add("breed", " 'Braford");
et.Metrics.Add("price", 77);
tc.TrackEvent(et);
tc.Flush();
```

The same data on-the-wire

The screenshot shows a network analysis tool interface with the following tabs: Statistics, Inspectors, AutoResponder, Composer, Log, Filters, and Timeline. Below these are view options: Headers, TextView, WebForms, HexView, Auth, Cookies, Raw, JSON (selected), and XML. The main area displays a JSON tree structure:

```
JSON
├── {}
│   └── data
│       ├── baseData
│       │   ├── measurements
│       │   │   └── price=42
│       │   ├── name=CowSold
│       │   ├── properties
│       │   │   └── breed=Lowline
│       │   └── ver=2
│       └── baseType=EventData
├── iKey=3863d2b7-643a-4c72-a4fd-a34bf90ff441
├── name=Microsoft.ApplicationInsights.Event
├── sampleRate=100
├── tags
│   ├── ai.device.id=Wojtek-Carbon
│   ├── ai.device.locale=en_US
│   ├── ai.device.osVersion=6.2.9200
│   ├── ai.device.type=Other
│   └── ai.internal.sdkVersion=py3:0.10.0
├── time=2015-08-23T23:44:41.555967Z
└── ver=1
```

Anatomy of Telemetry Service



About ingestion volumes

⌚ How much data an application sends?

App Size	Documents/sec	Documents/month
Small	< 200	< 518,400,000
Medium	200 - 500	~ 0.5 – 1.3 Bil
Large	500 – 1,000	~ 1.3 – 2.3 Bil
Very large	1,000+	2.3+ Bil
Super large	n*1K	n * 2.5 Bil

Ingesting at super large scale

- ⌵ A very popular internet phone application 😊
 - ⌵ n*100 Mil users/devices and >10 K servers
 - ⌵ 1 Mil [docs/sec] at peaks
 - ⌵ 5-10 [TB/hour] => 150 [TB/day]
 - ⌵ 5 geo-distributed ingestion regions
 - ⌵ n*1 K machines (Azure services) ingesting data
 - ⌵ > 500 Azure storage accounts
 - ⌵ ~50 – 100 K storage TPS



Official
Technology Partner



Microsoft and Real Madrid are joining together to transform the way the best fútbol club in the world connects with its **450 million international fans**. Through Microsoft Cloud technology and devices, we will change the way Real Madrid fútbol is played, coached, watched and ultimately experienced.

Hala Madrid!

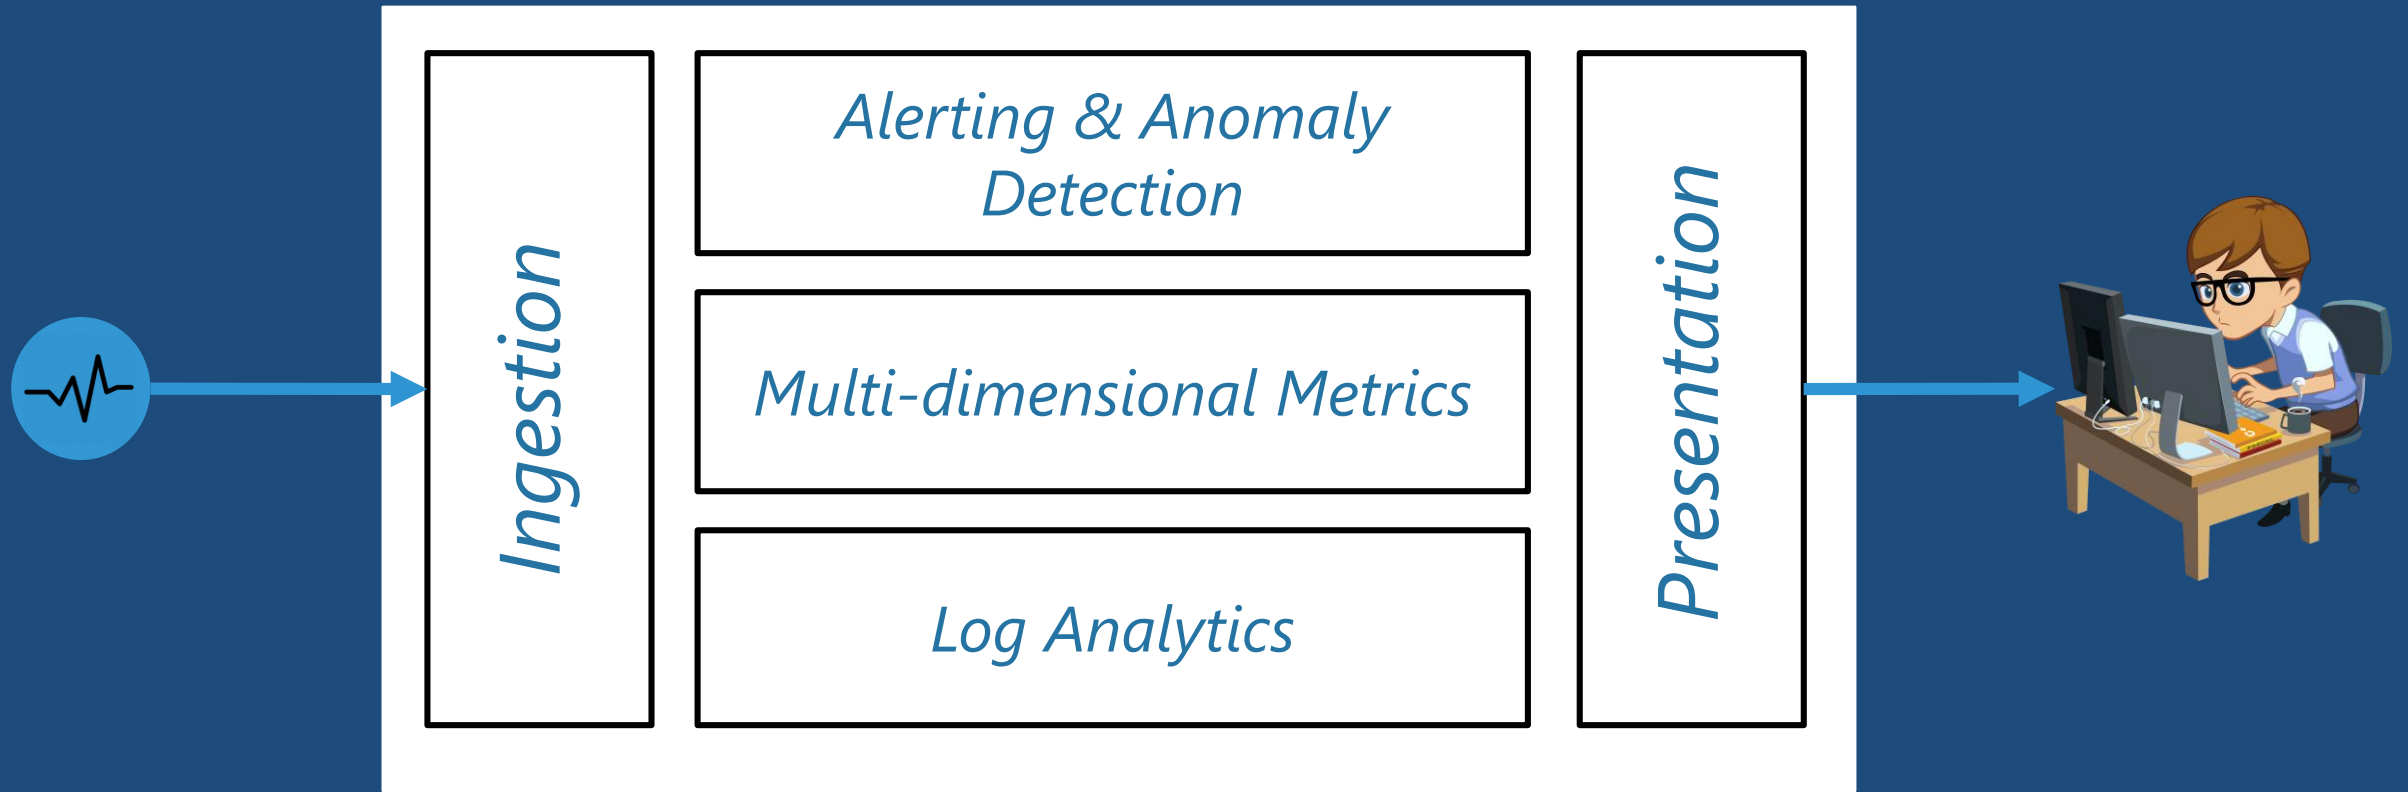
Ingesting at large scale

- ⌵ Real Madrid CF
 - ⌵ 8 geo-distributed ingestion regions
 - ⌵ Ingested to 6 data clusters, each with ~100 nodes/machines
 - ⌵ 1.5 – 2 K [docs/sec] sustained ingestion rate from devices
 - ⌵ > 2 K [docs/sec] sustained ingestion rate from servers

Common use of telemetry data

- ④ APM (Application Performance Management)
 - ④ What is the volume of data transactions?
 - ④ What is the average response time? What is the 90th-percentile response time?
 - ④ What is the load on my servers (CPU + Memory)?
 - ④ What is the size of request queue? . . .
- ④ Issue detection
 - ④ If you detect average CPU usage across cluster higher than 80%, let me know ASAP.
 - ④ If you detect abnormal memory usage in any cluster node, let me know ASAP.
- ④ Diagnostics
 - ④ What failed and when?
 - ④ What happened before the failure?
 - ④ What was the root cause? . . .
- ④ Usage
 - ④ Who uses my application, where and when?
 - ④ What are the most used feature?
 - ④ What are the least used features? . . .

Common data streams



Alerting & anomaly detection

- ⌚ Low latency (aka near-real-time) data stream with no data retentions
- ⌚ Alerts/notifications based on two kind of rules
 - ⌚ Aggregated metric below or above a threshold
 - ⌚ Anomaly detection

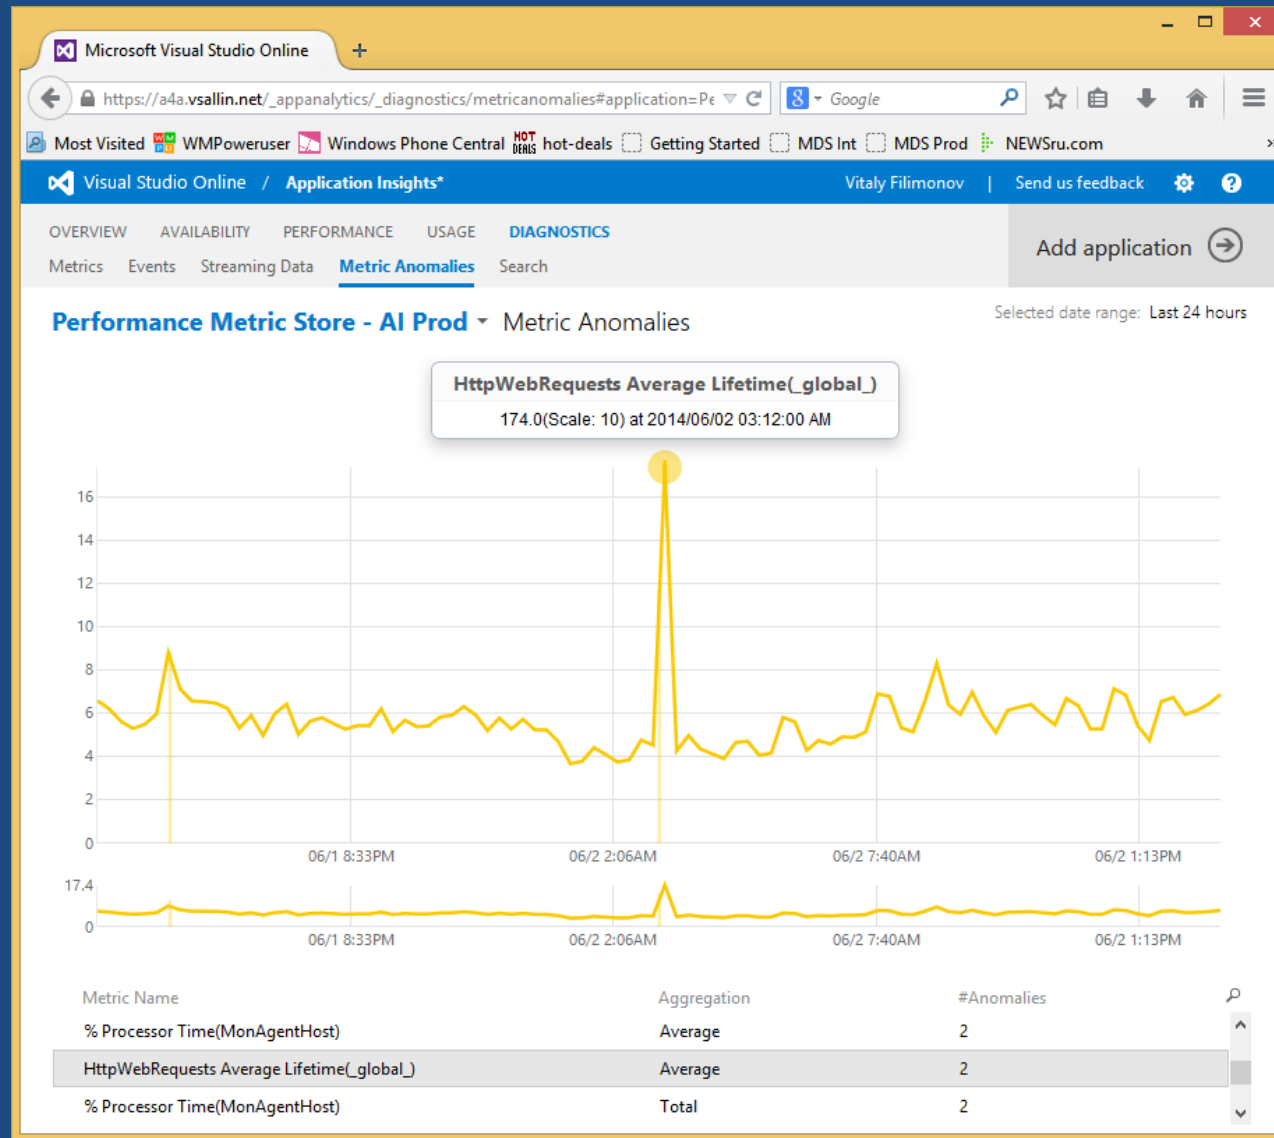
...

④ Metric value

- ④ selected metrics
- ④ for a selected combinations of dimension values
- ④ in a tumbling time window
- ④ aggregation value (count, min, max, average, stddev) is above/below a threshold

④ Anomaly detection

- ④ watches a stream of aggregated metrics
- ④ learns what "normal" looks like and self adjusts
- ④ detects abnormal sequences
 - ④ sharp rises or drops
 - ④ unexpected values
 - ④ spikes



MDM (multi-dimensional metrics)

- ⌚ Medium-to-high-latency data stream with fixed data retention (days or weeks)
- ⌚ Collects all metrics
 - ⌚ commonly pre-aggregates them in small time windows
 - ⌚ often receives pre-aggregated metrics from the application
- ⌚ Constructs time series of metric values (count, min, max, average, stddiv) for
 - ⌚ a given metric
 - ⌚ a set of dimensions and their values

...

④ Example request may look like this

④ "Show me distribution of

④ Max values of *CPU* metric

④ for the last two hours divided into 10-minut intervals

④ for (`cpu.clusterName == "C0T0"`) and (`cpu.serverName == "A45"`)"

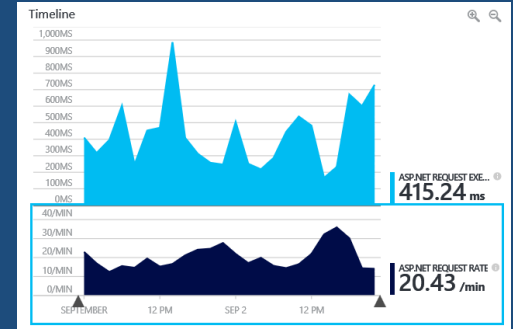
④ Comes down to indexing potentially **very** large set of metric documents

`(metricName, timeInterval, dimName, dimValue) → {metric documents}`

④ Dimension capping problem

$$\text{indexSize} \leq (\sum_{d_1}^d (\text{valueSet}^k).\text{Size}()) * \#\text{metrics} * \#\text{intervals}$$

④ capping is limiting number of dimensions and their values



Log analytics

- ④ Data stream that
 - ④ stores all telemetry data for (practically) unlimited time
 - ④ supports user queries against that data
- ④ What makes is possible?
 - ④ The data is never updated
 - it can be indexed and compressed during ingestion
 - ④ Most of the data is never looked at and when it is, it is looked at in a time interval
 - the data is organized in time-bound blocks
 - only relevant blocks are touched to compute a query
 - ④ Data is not updated => query results can be cached
 - ④ Older data is aggregated into progressively larger time windows (days, weeks, months)

	Alerting & Anomaly Detection	Multi-Dimensional Metrics	Log Analytics
APM		What is the average response time? What is the load on my servers? What is the size of the request queue?	What is the volume of date transactions?
Issue Detection	If you detect abnormal memory usage in any cluster node, let me know ASAP. If you detect average CPU usage across cluster higher than 80%, let me know APSP.	Show me average CPU usage across cluster for the last 2 hours.	
Diagnostics	What failed and when?	What filed and when? What happened before the failure?	What was the root cause?
Usage			Who uses my application, where and when? What are the most used features? What are the least used features?

Making \$\$ analyzing other peoples' telemetry

loggly

New Relic®

HOCKEYAPP

+ **sumologic**
Big Data for Real-time IT™

FLURRY


dynatrace


AppDynamic


Twitter Analytics

splunk®

Before you do it, think about ...

⌵ Multi-tenancy

⌵ SLAs

⌵ Reliability

⌵ Availability

⌵ Security

⌵ ...

⌵ Compliance

⌵ Scalability

⌵ COGS

⌵ Billing

⌵ ...

Back to my initial points

1. Real cloud applications can get very complex
2. You have to make your application send really good telemetry to have any chance of knowing what is happening in and to it
3. Building a good telemetry service may be a very sizable investment

wojtek@microsoft.com