

HUAWEI ENTERPRISE **A BETTER WAY**

iForesight: Intelligent Cloud Operations and Maintenance

iForesight relies on artificial intelligence, event correlation, and time-based technologies to effectively detect and predict anomalous behaviors during cloud operations to reduce the workload of human operators

Prof. Dr. Jorge Cardoso
Chief Architect for Cloud Operations and Analytics
German Research Center
Munich, Germany

CERN openlab OpenDay
September 21.09.2017
Geneva

enterprise.huawei.com
HUAWEI TECHNOLOGIES CO., LTD.



CERN
openlab



Executive Summary

Goal. Develop the next generation of DevOps tools & services for monitoring and operational automation. Drive **artificial intelligence** into **autonomous cloud operations**

Background. Cloud providers are increasingly interested in exploiting the use of machine learning to build new tools for autonomous cloud operations (Azure Insights, Amazon AWS, Google, HP Operations Bridge, Salesforce NoOps)

Objectives. Provide advanced solutions for autonomous cloud operations for Open Telekom Cloud (OTC), Fusion Cloud Stack (FCS) and FusionSphere

Technologies. **Machine Learning**, **data mining**, **time-series analysis**, and **distributed tracing technologies**

Amazon AWS

- Amazon is interested in local, mini data-centres for its huge cloud business.
- It may be able to disrupt other telecoms/IT businesses, and steer some “edge computing” technology standards and paths
- **AWS Greengrass** powers local edge gateways and appliances. It works as an extension to AWS IoT.



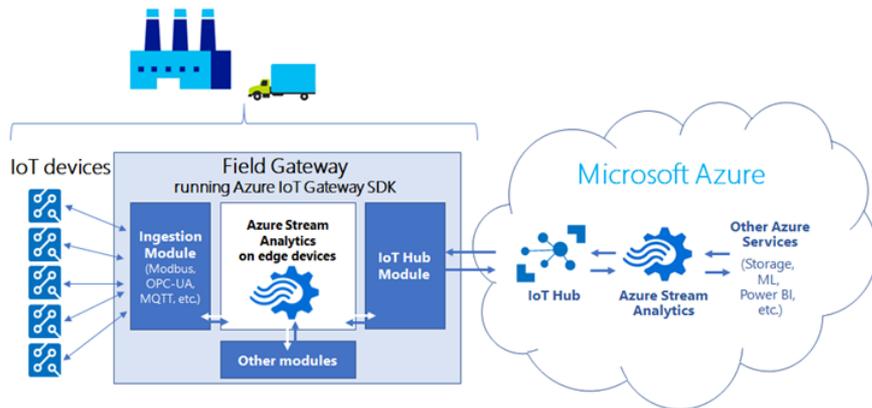
AWS Greengrass has a small footprint that can run on system-on-a-chip devices like Raspberry Pi and BeagleBone powered by ARM processors

Benefits

- Reduced latency between the devices and data processing layer
- Reduced bandwidth costs involved in ingesting large amount of data to the cloud
- Compliance and security by retaining the sensitive data locally

Microsoft IoT Edge

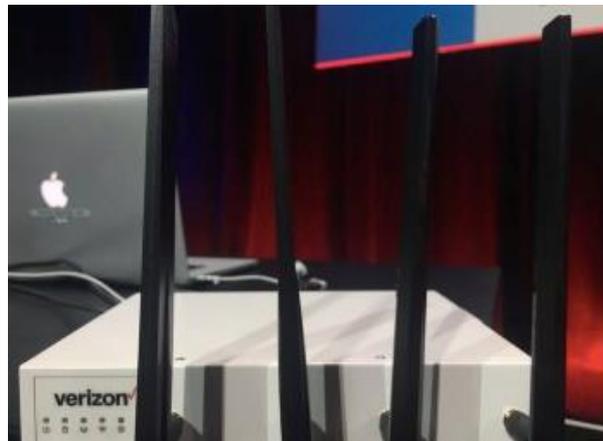
- Run artificial intelligence at the edge
- Perform edge analytics
- Deploy IoT solutions from cloud to edge
- Manage devices centrally from the cloud
- Operate with offline and intermittent connectivity



OpenStack and Edge Computing

The top diagram, titled "OPENSIFT + OPENSTACK & NFV Edge Computing", illustrates a cloud architecture. It features "OPENSIFT 3" and "Kubernetes" at the top, with a "Node" box containing icons for "Pod", "PVC", and "PV". Below these are icons for "Containerized Microservices", "Runtime and Packaging For nst", "Automated Security", "Orchestration", "Networking", "Custom Services", and "Storage". A "RED HAT CLOUDFORMS" logo is in the bottom right. To the right, a flow diagram shows "NFVO" connected to "VNF-M".

The bottom part is a video player for "Verizon Product Case Study - OpenStack at the Edge". The video title is "OpenStack at the Edge" and it is dated "8 May 2017". The OpenStack logo is prominent. The video player interface includes a play button, a progress bar at 0:22 / 44:58, and a "SUMMIT BOSTON" banner.



Verizon's cloud-in-a-box offering based on OpenStack delivers computing to the edges
Boston Summit 2017

Companies

- Microsoft's
 - Data centers average failure rate is 5.2 devices/day and 40.8 links/day
 - Median time to repair of ca. 5 minutes and a maximum of 1 week.
- Google
 - New cluster over one year
 - 5 times rack issues 40-80 machines seeing 50 percent packet loss
 - 8 network maintenance events (four of which might cause ~30-minute random connectivity losses).
 - 3 router failures resulting in the need to pull traffic immediately for an hour
- CENIC
 - 500 isolating network partitions with median 2.7 and 32 minutes
 - 95th percentile of 19.9 minutes and 3.7 days, respectively for software and hardware problems

Applications

- MongoDB
 - Separated primary from its 2 secondaries
 - 2 hours later the old primary rejoined and rolled back everything on the new primary
- Redis
 - A network partition isolated the Redis primary from all secondaries.
 - Every API call caused the billing system to recharge customer credit cards automatically, resulting in 1.1 percent of customers being overbilled over a period of 40 minutes.
- MySQL
 - The partition caused inconsistency in the MySQL database.
 - Because foreign key relationships were not consistent, Github showed private repositories to the wrong users' dashboards and incorrectly routed some newly created repositories.
- Elasticsearch
 - For several seconds, Elasticsearch is happy to believe two nodes in the same cluster are both primaries, will accept writes on both of those nodes, and later discard the writes to one side.

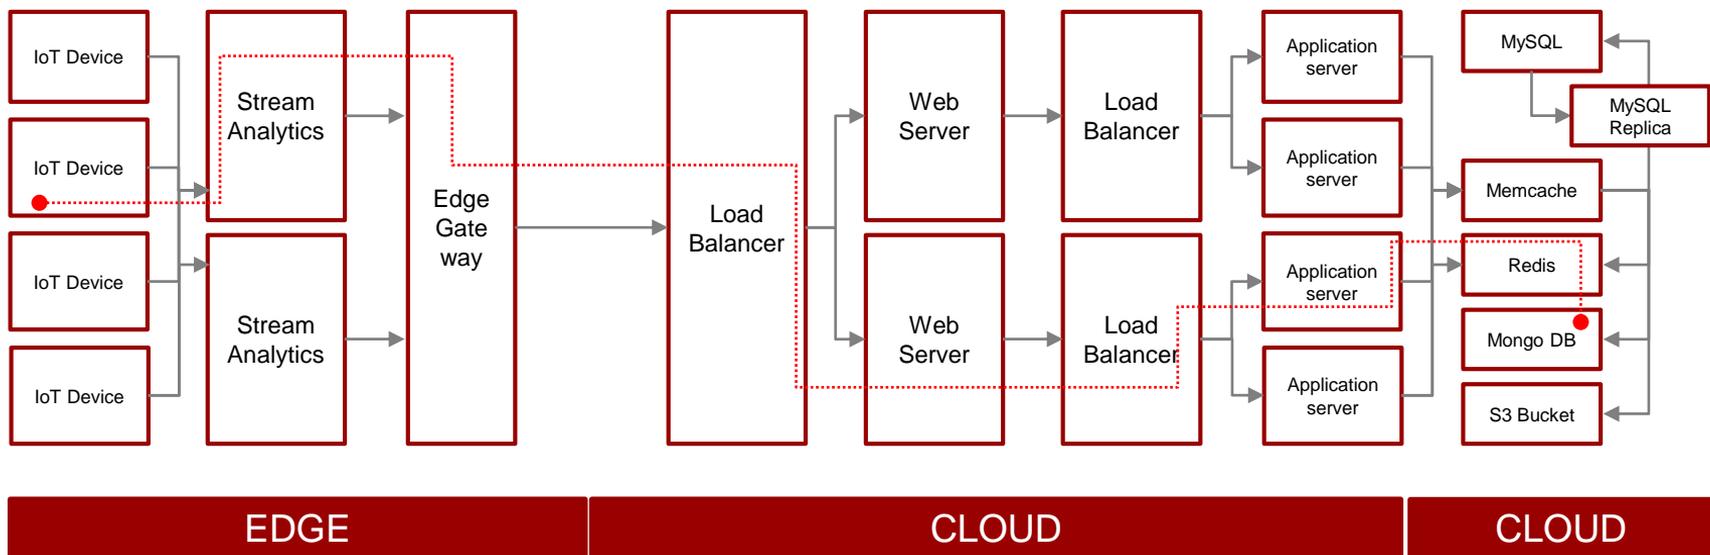
<https://queue.acm.org/detail.cfm?id=2655736>

Peter Bailis, UC Berkeley

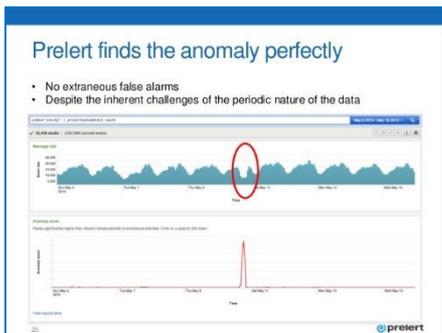
Kyle Kingsbury, Jepsen Networks

Troubleshooting Multi-clouds

A service request failed
Which path was followed?

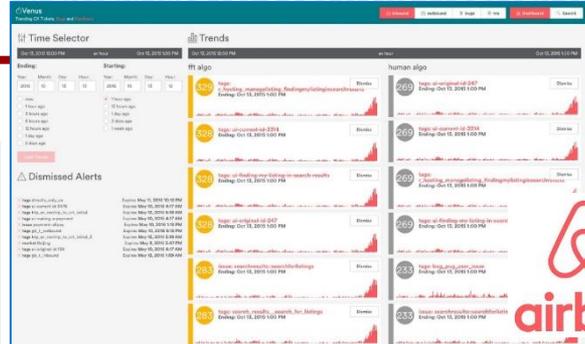


Anomaly Detection Services

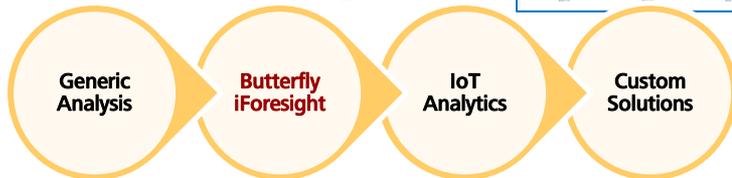


http://www.prealert.com/docs/splunk_app/latest/

<https://www.sumologic.com/resource/datasheet/anomaly-detection/>

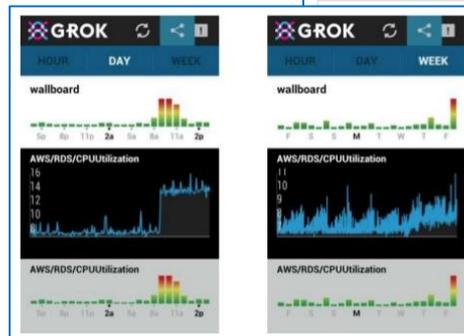


<https://medium.com/airbnb-engineering/how-airbnb-manages-to-monitor-customer-issues-at-scale-b883301ca461>



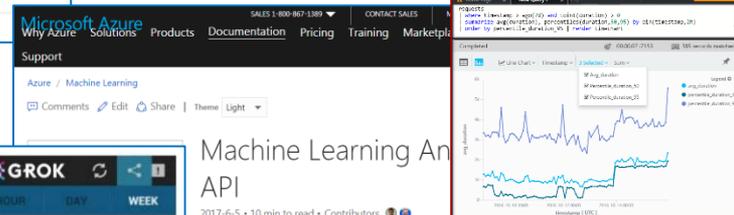
- Generic Analysis**
 - Decision trees
 - SVN
 - For experts
 - Too difficult
- Butterfly iForesight**
 - Flexible
 - Explore and Fix
 - AIOps as Code
- IoT Analytics**
 - Specific Analysis
 - Uses Forms
 - Predefined Domains
- Custom Solutions**
 - Too expensive
 - Too specific

PreAlert: 1st class citizen: Log lines (text patterns) + AI
 SumoLogic: 1st class citizen: Log lines (text patterns) + AI
 Grok: 1st class citizen: Sys. Metrics (text patterns) + AI
 Twitter: 1st class citizen: Log lines (text patterns) + AI
 AirBnB: 1st class citizen: JSON (isolated time-series) + AI
 Azure: 1st class citizen: JSON (isolated time-series) + AI
iForesight: 1st class citizen: Distributed Traces + AI



<https://grokstream.com/>

Twitter S-H-ESD Algorithm
 • <https://github.com/nlittlepole/thermometr>



Machine Learning Anomaly API

2017-6-5 • 10 min to read • Contributors
 In this article
 Overview
 API Deployment



Microsoft Insights

<https://docs.microsoft.com/en-us/azure/machine-learning/machine-learning-apps-anomaly-detection-api>

Introducing practical and robust anomaly detection in a time series

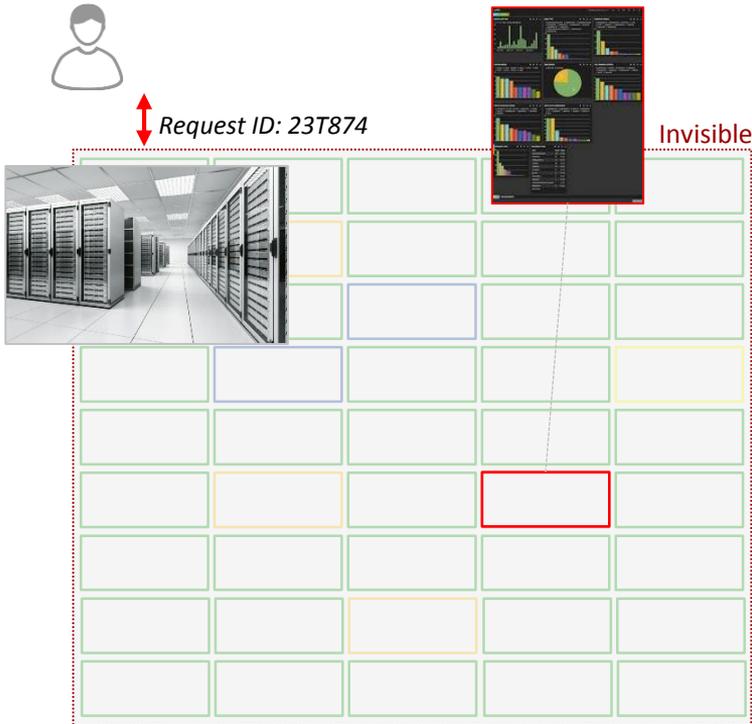
By @arun_kejariwal
 Tuesday, 6 January 2015



Log/Metrics vs. Distributed Traces Monitoring

- **SLA/SLO:** OTC replies within 0.5 seconds of receiving a request
- **Scenario:** The request for the execution of a service was 50% slower when compared to 12h ago.
- **Root cause:** NOT KNOWN

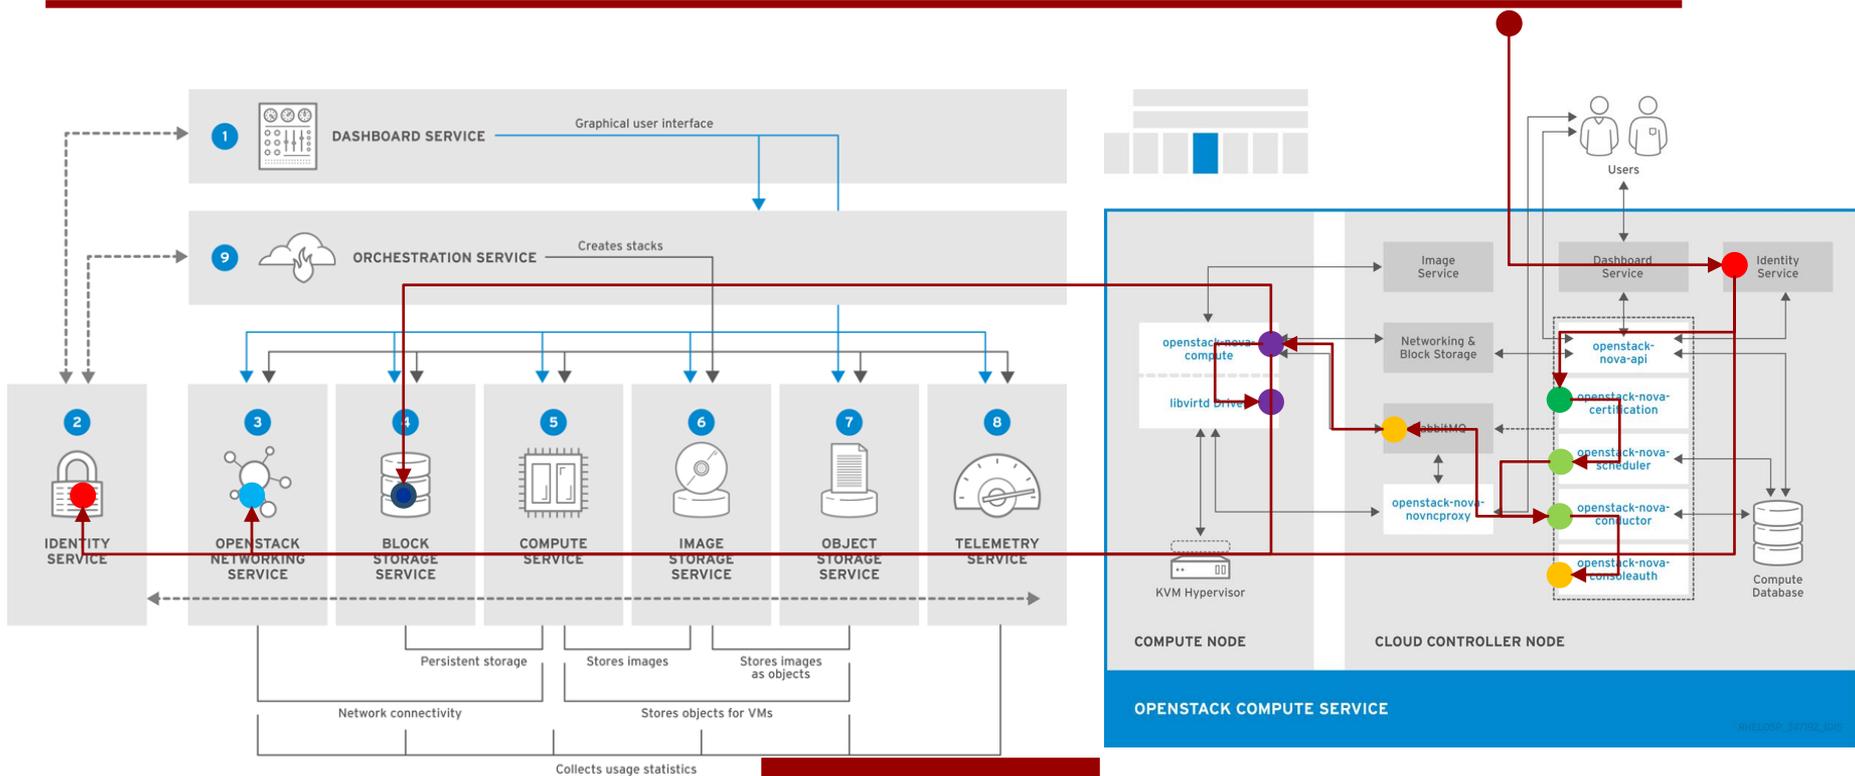
- **SLA/SLO:** OTC replies within 0.5 seconds of receiving a request
- **Scenario:** The request for the execution of a service was 50% slower when compared to 12h ago.
- **Root cause:** HEALTH OF NODES N1 AND N4 IS LIMITED



Log/Metrics Monitoring

Distributed Tracing Monitoring

Distributed Tracing in OpenStack



- A**
 (18, 17, 17)
- B**
 (23, 18, 17)
- B**
 (23, 18, 17)
- A**
 (18, 17, 17)
- C**
 (14, 17, 17)
- D**
 (26, 14, 17)
- E**
 (34, 26, 17)
- E**
 (34, 26, 17)
- D**
 (26, 14, 17)
- C**
 (18, 17, 17)



Advanced Trace Analysis

1. Trend Analysis

- Script: *detect_trend_transitions.py*
- How does the performance of FusionSphere subsystems change over time?

2. Controller Efficiency Comparison

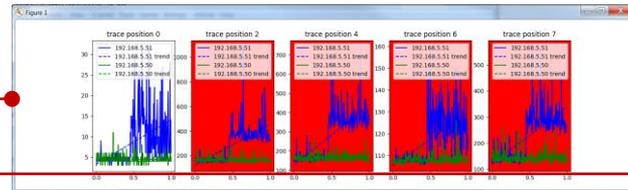
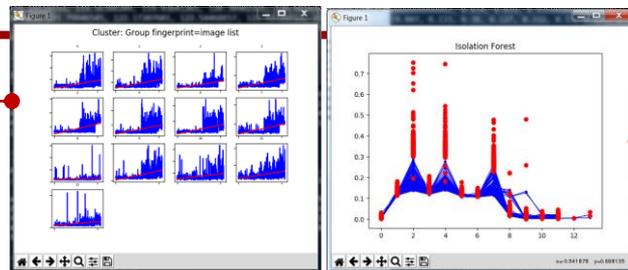
- Script: *controller_anomalies.py*
- Identify FS controller nodes with poor performance

3. Outlier Detection for SLO

- Script: *detect_outliers.py*
- Identify traces with abnormal performance characteristics

4. Root Cause Analysis

- Script: *trace_troubleshooting.py*
- Which FS components are responsible for poor performance?



```

iforesight service
Start trace and transition analysis using linear regression? [Y/n] y
Select the time window (in days) to analyze: 2
Anomalous traces received with a trend degree>15 over the last 2 days
TID: 7674b8b6340ade05a83c0d856f832149, Tenant-ID: 97130491-e102-4676-b8a9-f47b149ba577
TID: abc f891e548cb6d8e849f4f55af9167, Tenant-ID: 9c436f88-8eef-449d-99cf-dcad1b58145e
Getting transitions from anomalous traces
> 28 transitions received
Select the threshold for labeling anomalies (0-90): 0
Anomalous transition: image_3d7079e
Anomalous transition: image_b337fa9
Anomalous transition: image_2711e78
Anomalous transition: image_551dadc
Anomalous transition: image_5445add
Anomalous transition: image_3344890
Anomalous transition: image_f05ab0e
Anomalous transition: image_5f147fa
Anomalous transition: image_81774ed
Anomalous transition: image_03bf6c2
Anomalous transition: image_35d88b4
Getting resources associated with transitions
2 resources to consider: glance, keystone
Root cause of anomalies.
Resource: keystone (support: 17)
Resource: glance (support: 5)
    
```

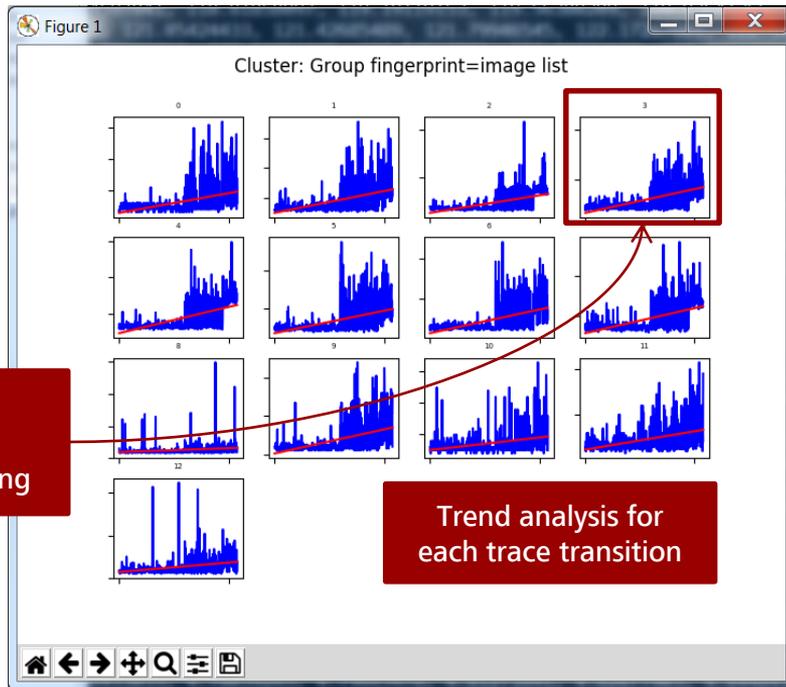
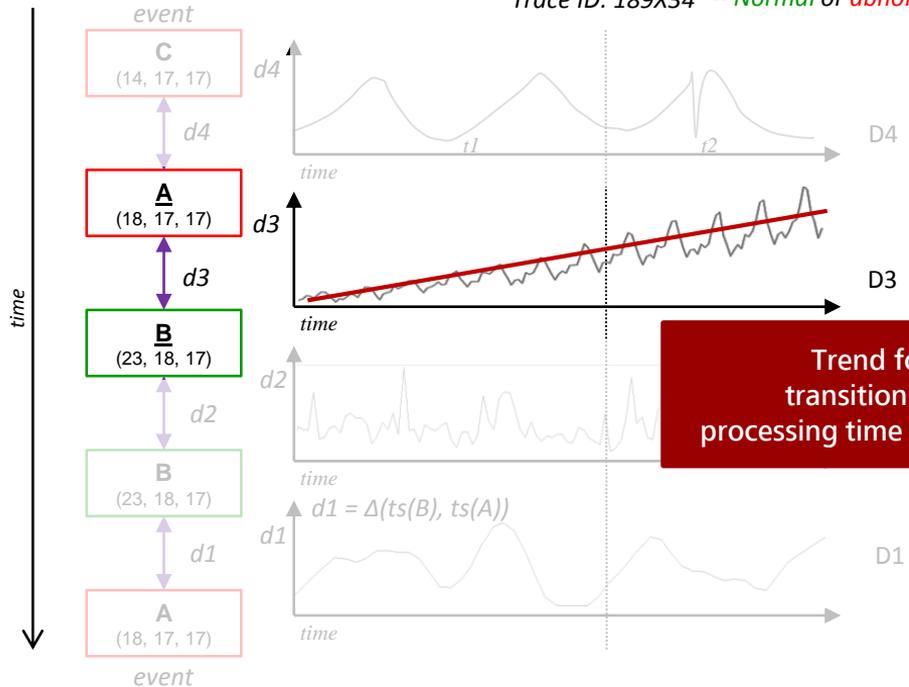
AI Algorithms	Status
Trend Analysis	OK (-)
Outlier Detection	OK (-)
Root Cause Analysis	OK (-)
Spikes & Dips	TBD

AI Algorithms	Status
Mean Shift	TBD
Ramp up	TBD
Distribution change	TBD
Deep Learning	TBD

Advanced Trace Analysis

Distributed Trace

Trace ID: 189X34 -- Normal or abnormal?



Advanced Trace Analysis

Identify FusionSphere controller nodes with poor performance

Workload analysis

Time Window (windows)

- time interval of traces to consider

Cluster (transition, start-end)

- place similar traces together into clusters

Aggregate (project, service, path, query, sql, rpc)

- aggregates trace events to analyze

Group by (host, ip)

- Group aggregates based on some property
- group_by(left.ip)

Execute Machine Learning algorithm

- ML_TSA_method(...)

Retrieve Results

- get_results(...)

Smart Script: controller_anomalies.py
iForesight SDK

```
# Windowing: set a temporal window to analyze traces
i4sight.window(when.past(weeks=8), when.past(weeks=1),
               wtype='train')

# Clusters traces with the same structure
i4sight.cluster(itype='structure', similarity=1.0)

# Aggregate the events of traces into smaller structures
i4sight.aggregate(itype='transition')

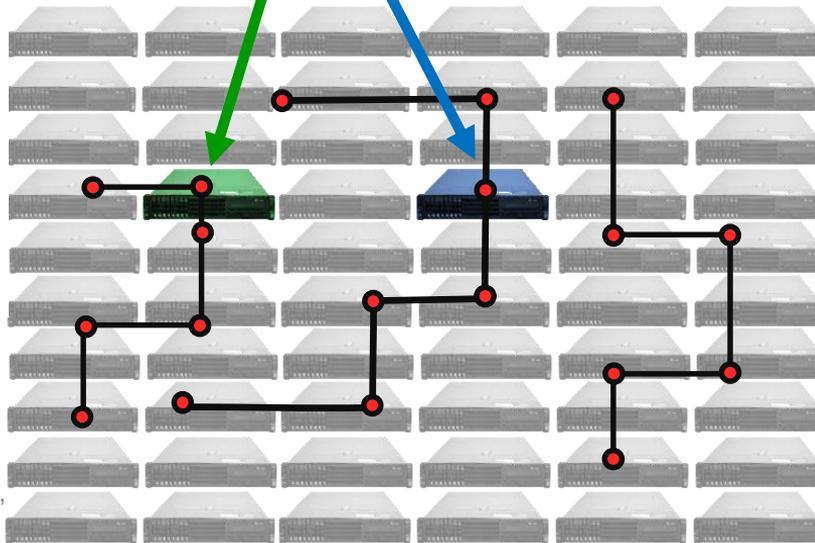
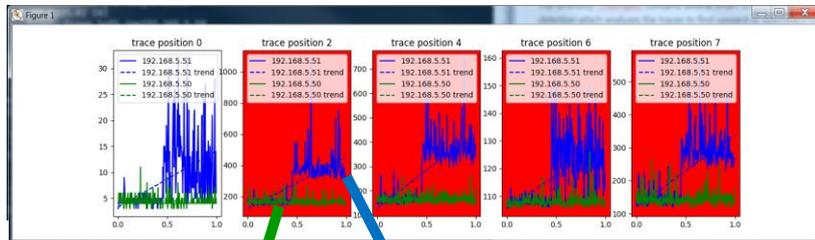
# Group transitions by IP Address
i4sight.group_by(['left.ip'])

# Execute linear regression to each group
i4sight.regression()

result = i4sight.get_result()
```

FS controller latency
for each transition

detected
increased
processing
latency of
blue
controller



FusionSphere,
FCS, OTC

Controller Nodes



Isolation Forest

- **Unsupervised** machine learning algorithm, does not require labeled data for training.

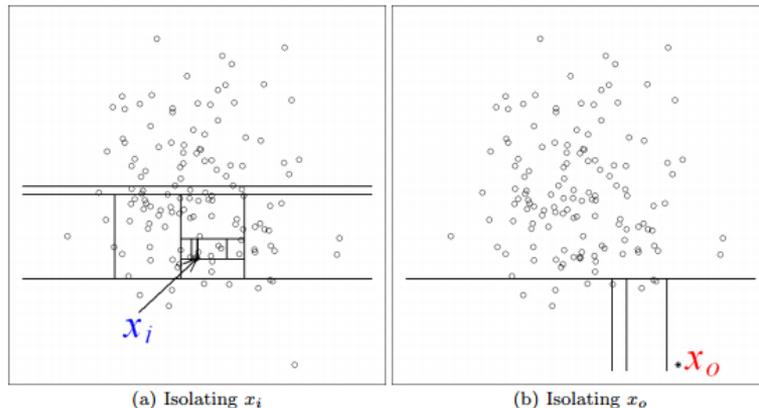
Algorithm

The algorithm 'isolates' traces by randomly selecting a transition and then randomly selecting a split value for its elapsed time between the maximum and minimum values.

1. Recursive (and random) partition of data represented as a tree
2. The end of the tree is reached once the recursive partition of elements is finished. The distance taken to reach the outlier is far less than that for the normal elements.
3. The distance of the tree path is averaged and normalized to calculate an anomaly score.
 - A score=1 is labeled as an outlier.
 - Values close to 0 are considered normal.

Parameters

- *Feature vectors*. Elapsed time of each trace transition
- *Contamination*. Estimated fraction of outliers in the training data set.





HUAWEI ENTERPRISE ICT SOLUTIONS **A BETTER WAY**

Copyright©2015 Huawei Technologies Co., Ltd. All Rights Reserved.

The information in this document may contain predictive statements including, without limitation, statements regarding the future financial and operating results, future product portfolio, new technology, etc. There are a number of factors that could cause actual results and developments to differ materially from those expressed or implied in the predictive statements. Therefore, such information is provided for reference purpose only and constitutes neither an offer nor an acceptance. Huawei may change the information at any time without notice.