#SHAREorg

SHARE
Technology · Connections · Results

Behavioral Analysis & Predictive Analytics meet SLA's - Prevent Problems and Reduce Costs

Jacob P. Ukelson, D.Sc.
CONIC-IT

Wednesday, August 8, 3-4 PM
Session 11558

SHARE in Anaheim 2012

Topics

- What's the problem?
- What is Behavioral Analysis and Predictive Analytics for IT?
- How Does it Apply to Mainframe Performance?
- Is it Useful?
- How Hard is it to Start?

IT Performance Management Challenges

- Too many inaccurate alerts
  - No explicit knowledge of what is “Normal”
  - Static thresholds
  - Symptoms mask root cause

- Users are the first to know

Mainframe Performance Management Challenges

- The complexity of monitoring a large scale, shared, transaction oriented environment is the main cause of mainframe service management problems
  - Shared resources create complex performance interactions in production
  - IT needs real time performance analysis tools not just monitoring tools, to show “where to look”
  - Deep, real time analytics are the only way to get alerts that are meaningful

Behavioral Analysis and Predictive Analytics

- To provide insight into the current state of the system, use smart analysis of previous performance behavior to find anomalies

Look for anomalous behavior to indicate things are amiss

Behavioral Analysis and Predictive Analytics

- Find historical behavior “profile” that describes current behavior – create a behavioral baseline
- Goal: find anomalous performance behavior that may be an early indication of a performance problem
  - Similar to the way an experienced operator would see anomalous patterns and use those as an early warning

Behavioral Analysis and Predictive Analytics

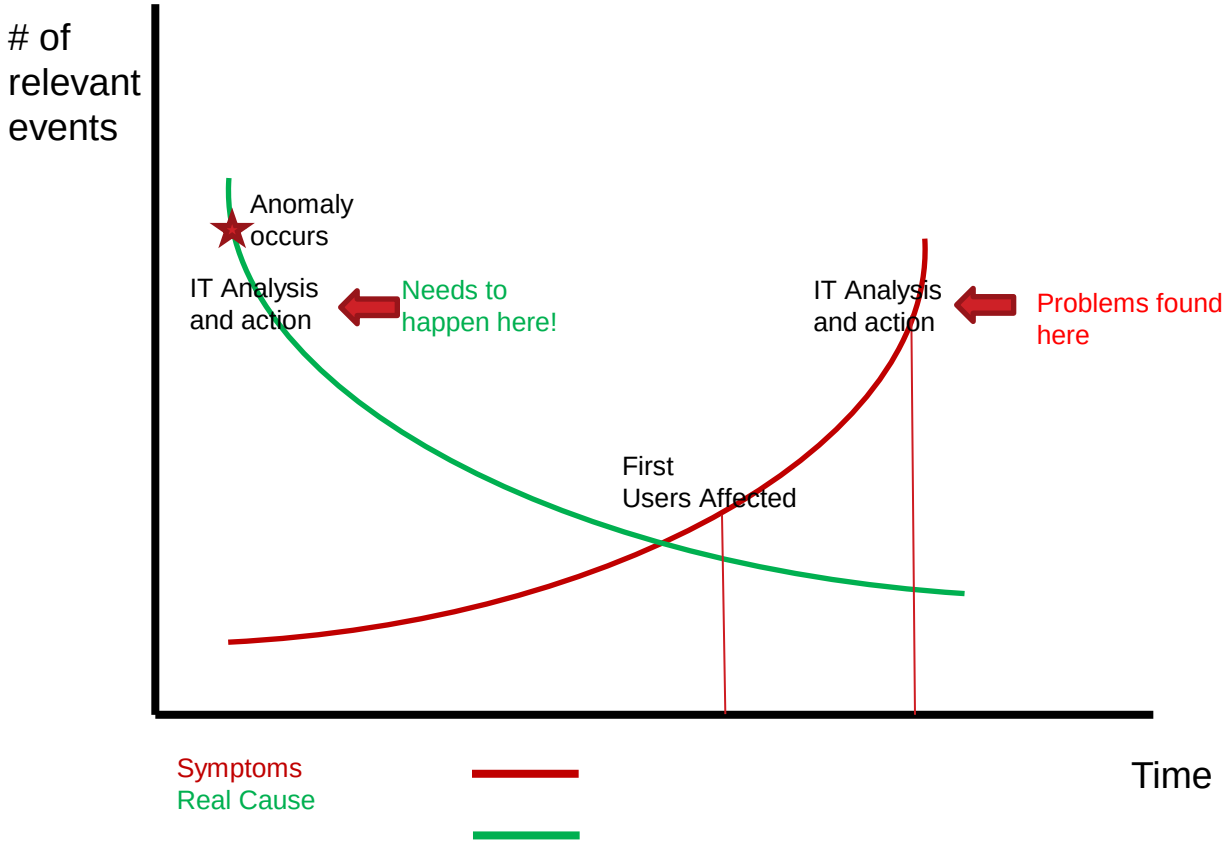

The Early Bird Catches the Worm

What Happens if You are Late

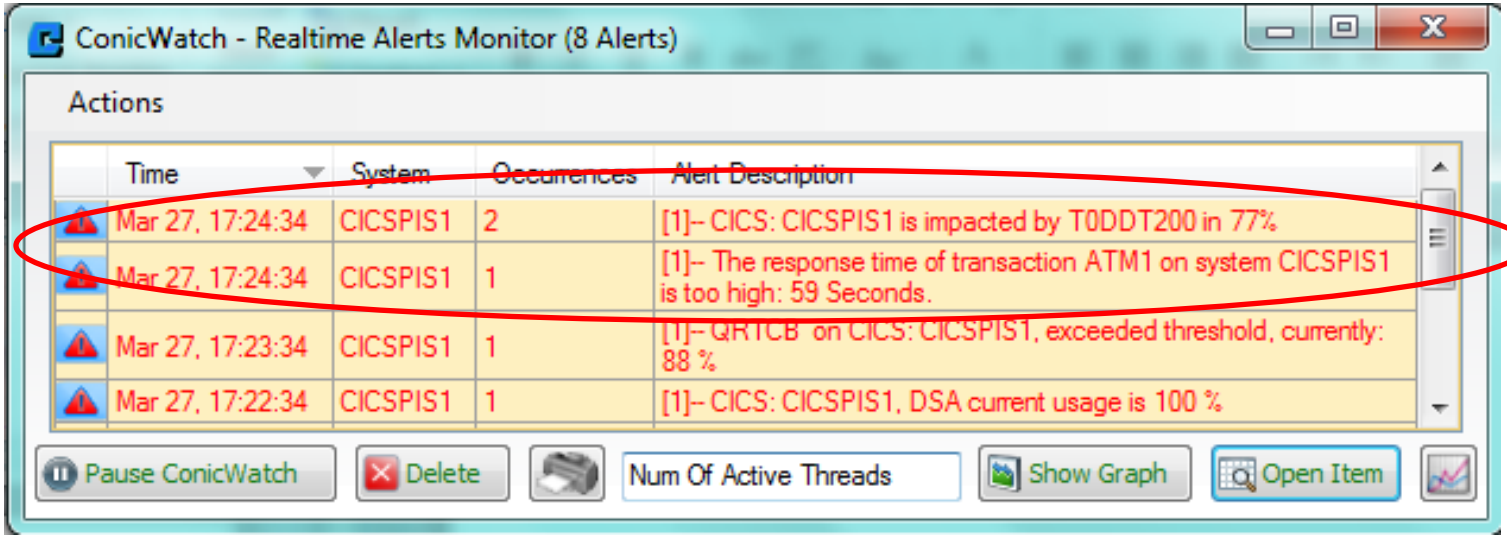

What Happens if You are Late

MVS Overview

System Name	LPAR CPU Usage	LPAR CPU Profile	TCB	SRB	NCL	MVS	IFA	IFC	IIP	IIC
SYSA: Center NY	82	56	60	4	0	10	2	7	15	8
SYSB: Center LA	73	47	57	6	0	10	6	8	16	7
SYSC: Center Frankfurt	51	56	43	5	0	3	12	17	17	14
SYSD: Backup Center	53	69	46	5	0	2	12	15	16	10

Maybe need more capacity?

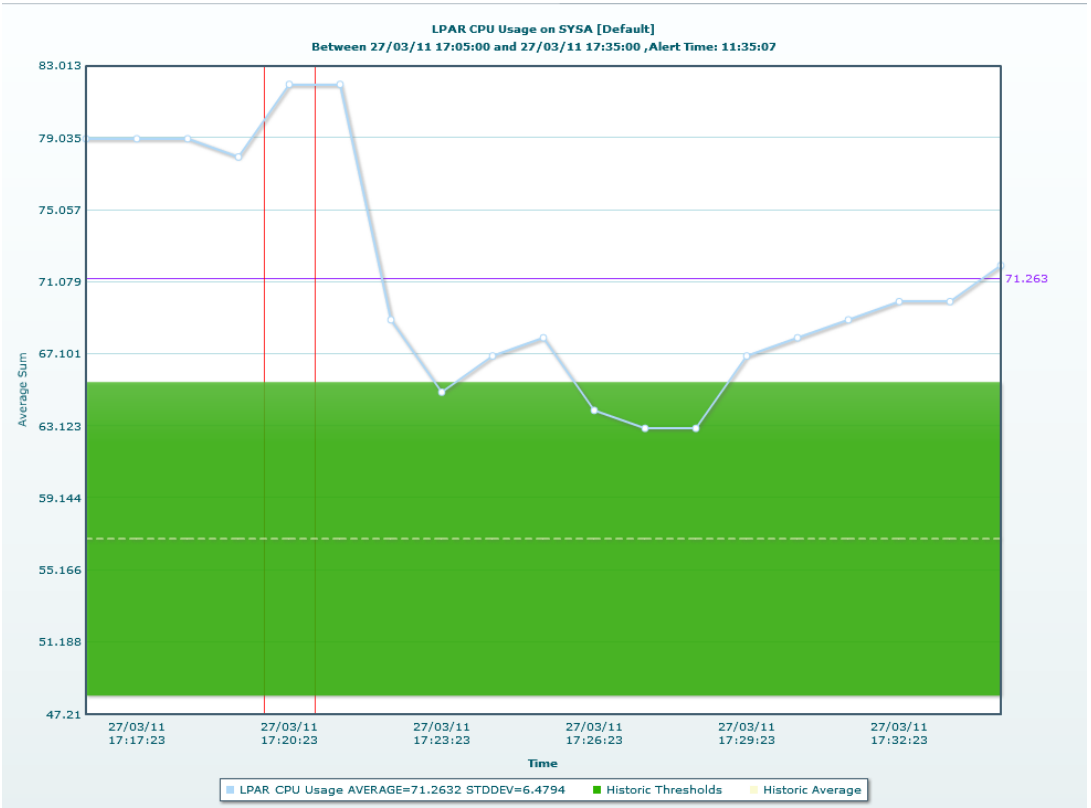

What Happens if You are Late

DB2 Overview

System Name	# Lock Owners	# Lock Waiters	Buffers Status	# Threads	# Active Threads	Log Archiving Rate	Impact
SYSDB2P2	2	6	BP7 (100%)	36	36 (100%)	No change (182 Min)	
SYSDB2P2_B	none	none	No exceptions	22	1 (4%)	No change (22 Min)	

Maybe a problem with the DB?

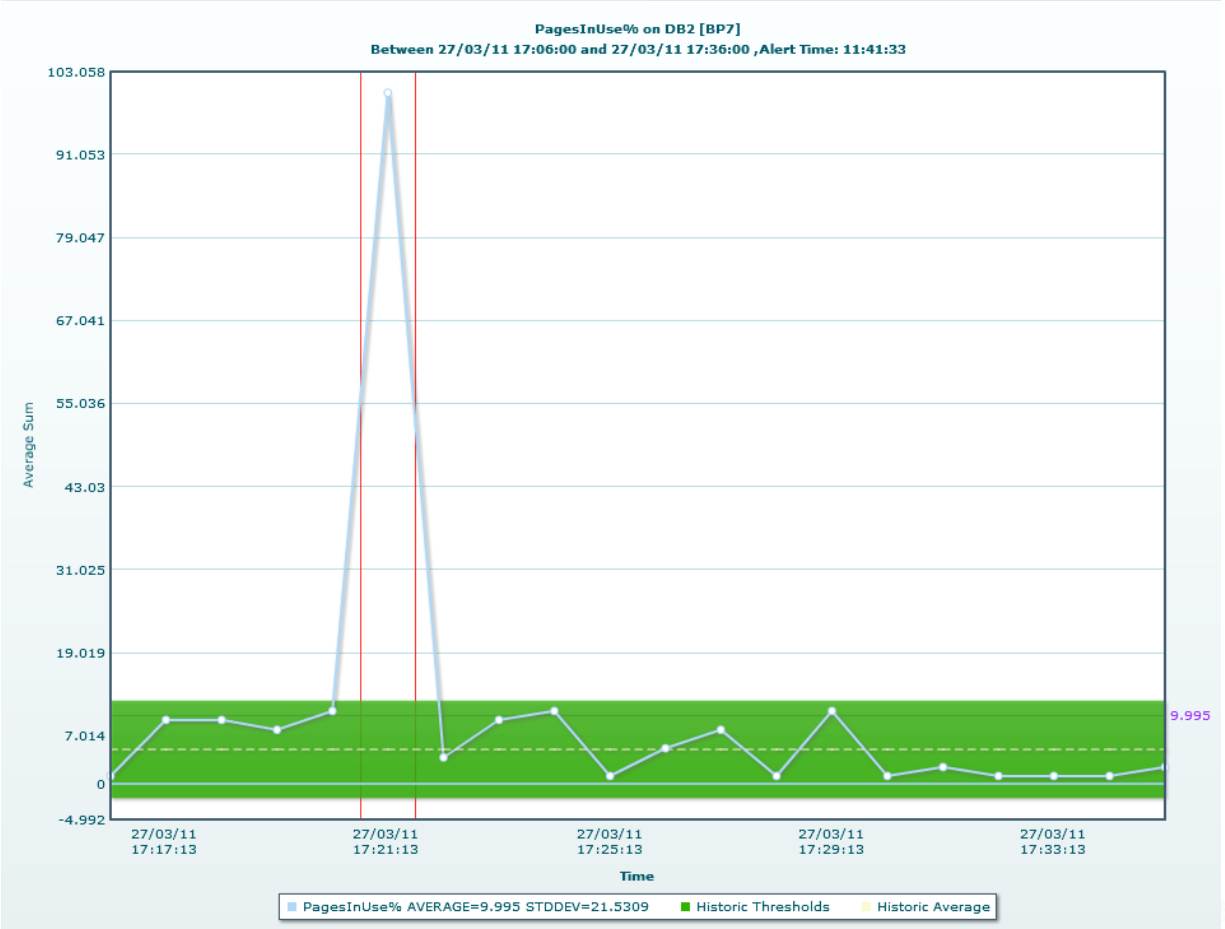

But If You Could Turn Back Time…

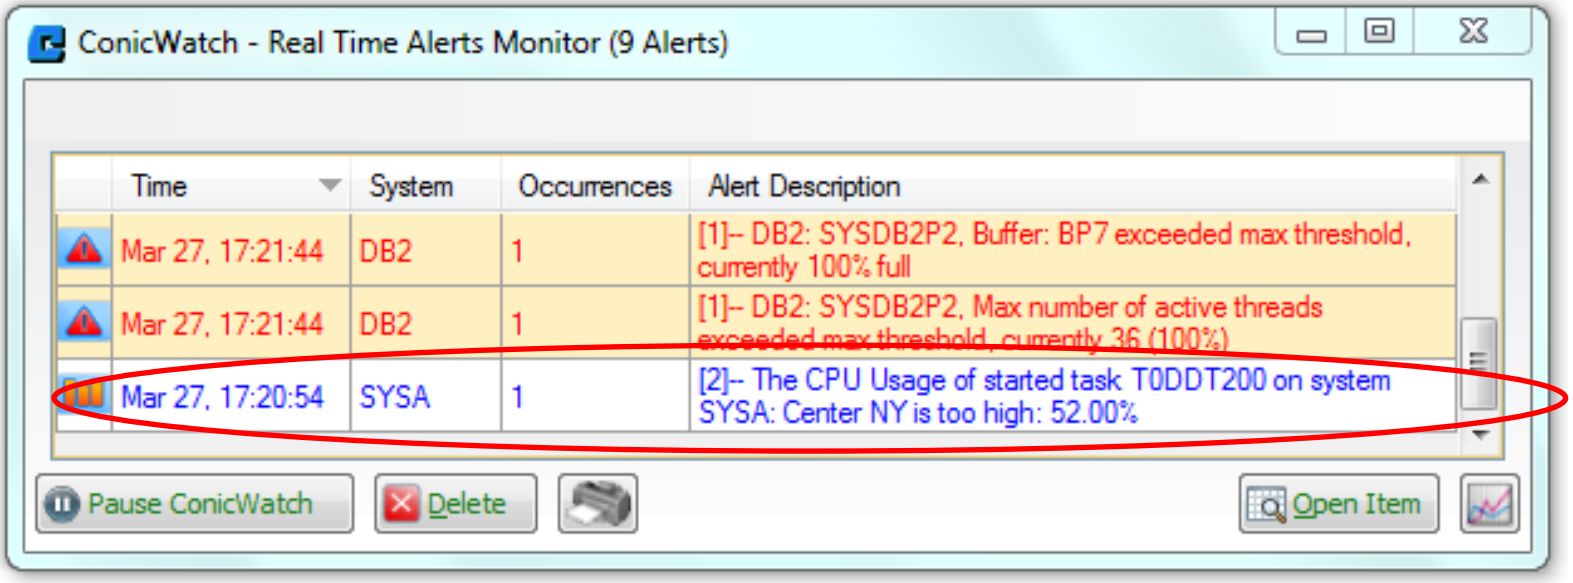

Become a Timelord

The Problem is Much Simpler

Watching the right place capturing the right information at the right time (and with a dash of understanding)

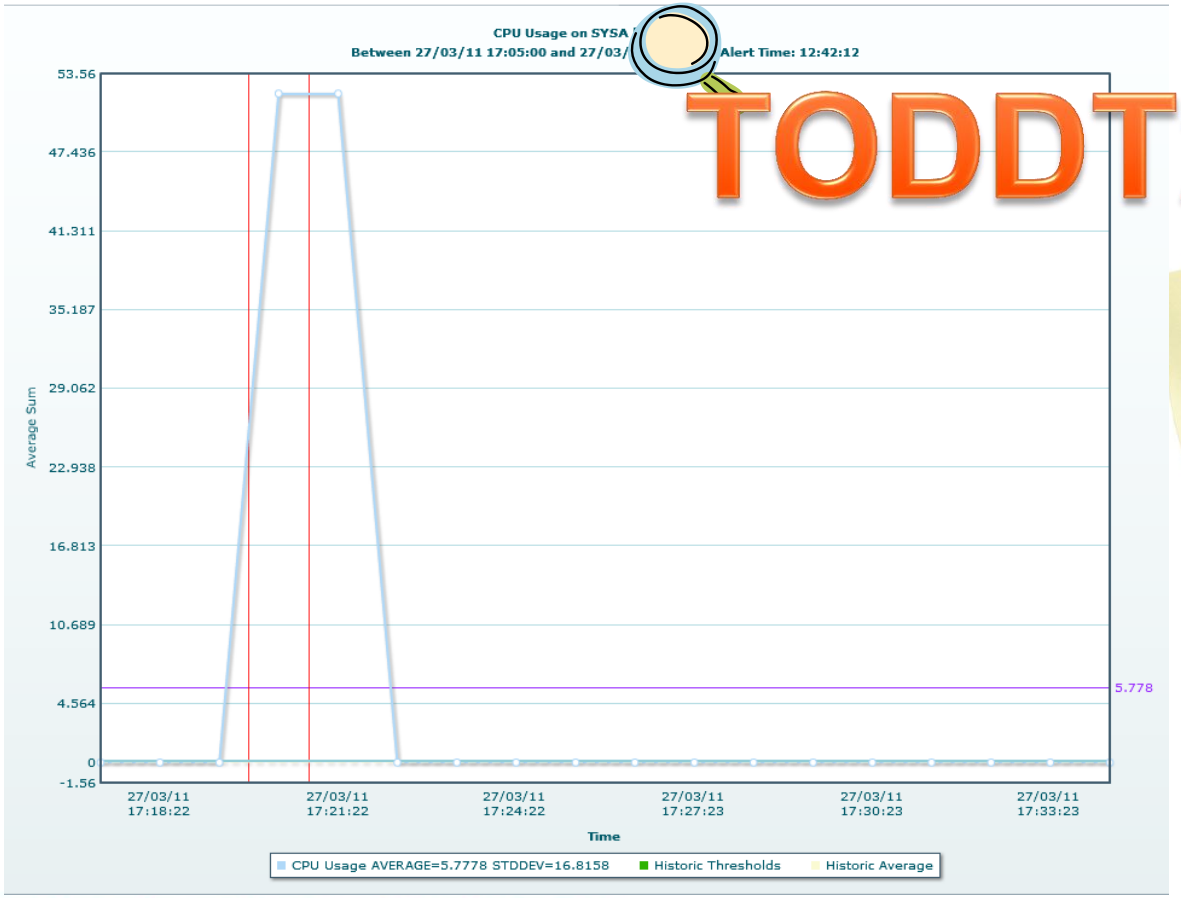

TODDT200

Fix task

No Capacity or DB problem

So How Can It be Done?

- Purely statistical approach ?
  - Uses lots of “low quality” models and attempts a best fit to the current state. Prediction based on selected model
  - Also called descriptive statistics
  - Good News: Generic and can be used with any system
  - Bad News: Doesn’t really work

So How Can It be Done?

- Mathematical modeling + domain expertise works
  - Hierarchy of domain expertise
    - Computers – e.g. CPU must be positive
    - Performance – e.g. Locked resources
    - Mainframe – e.g. Queues
    - Not generic, but can be architected to separate concerns

Domain Expertise

- Which variables and combination of variables are of critical importance
  - There are hundreds, or even thousands of parameters that could be considered
- Definition of anomalies
  - e.g. time related parameters
  - e.g. work related parameters
- The power of dynamic thresholds
  - Do nothing, Collect and Aggregate, Acquire more data (increase confidence), Alert

Performance Informed Prediction Model

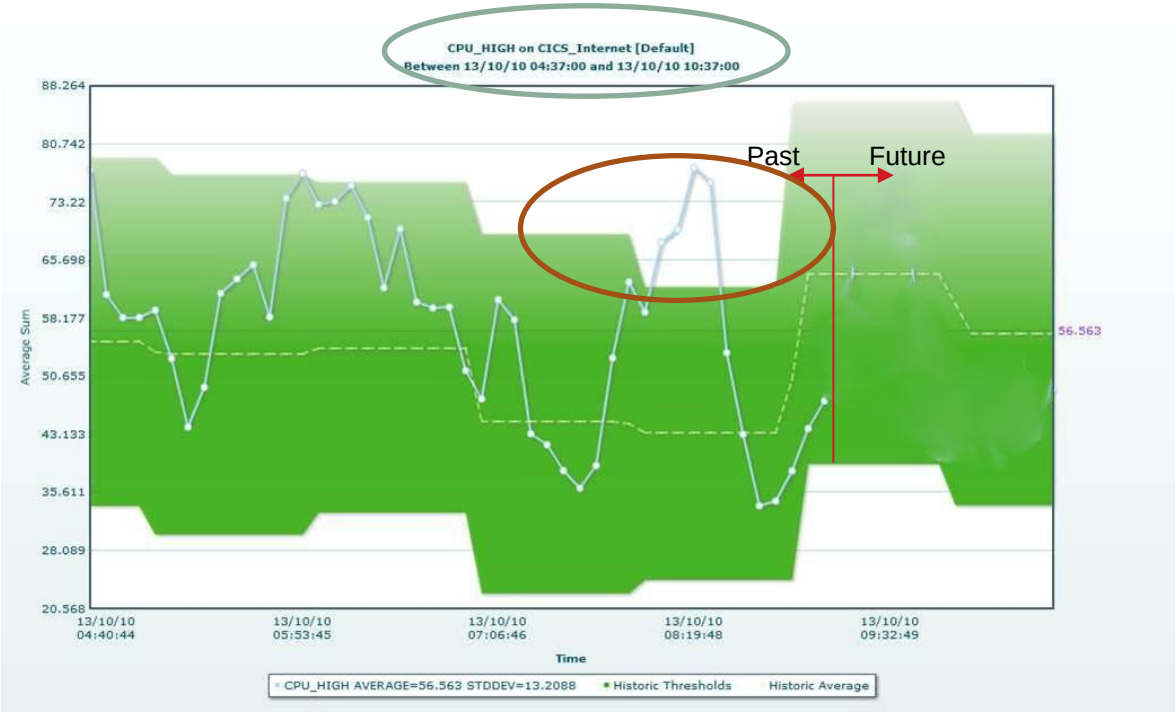

+ Domain Awareness

Parameter Selection

- System Specific Parameters
- Synthetic Parameters

Alerting Rules
(Alert Worthy?)

+ Customer Specific Configuration

Examples of Information Collected for Analysis

- Information from **MVS** monitor
  - LPARs CPU, 4HRA CPU - 4 hour rolling average, Specialty processor usage, Common Storage Activity
  - Disk activity
- Information from **DB2** monitor
  - Buffer pool usage, Thread activity, Archiving rate, Locking (owners and waiters), Deadlock, (Lock), Timeout
- Information from **CICS** monitor
  - Transaction rate (per CICS), CPU Consumption (per CICS), Average response time (per CICS), “QRTCB” values (per CICS), “ENQ locks owner/waiters” values (per CICS)

Rule Examples

- Amount of started task CPU usage over\under expectations (highway)
- Low capture ratio
- Most utilized disk related to CPU overage
- Cumulative resource utilization for transactions across LPARS\CICS

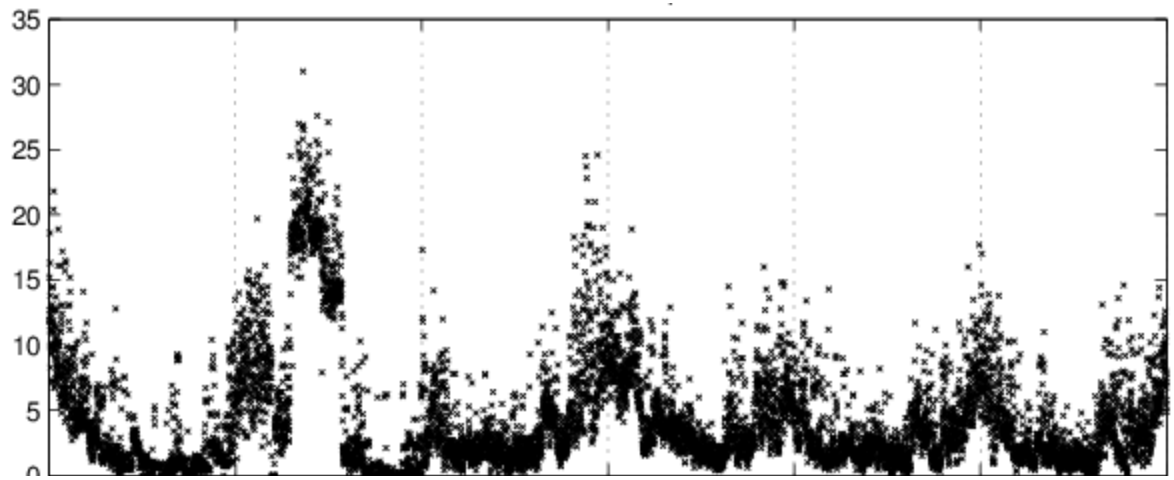
35
30
25
20
15
10
5
0

Fits Time Related Data Profile

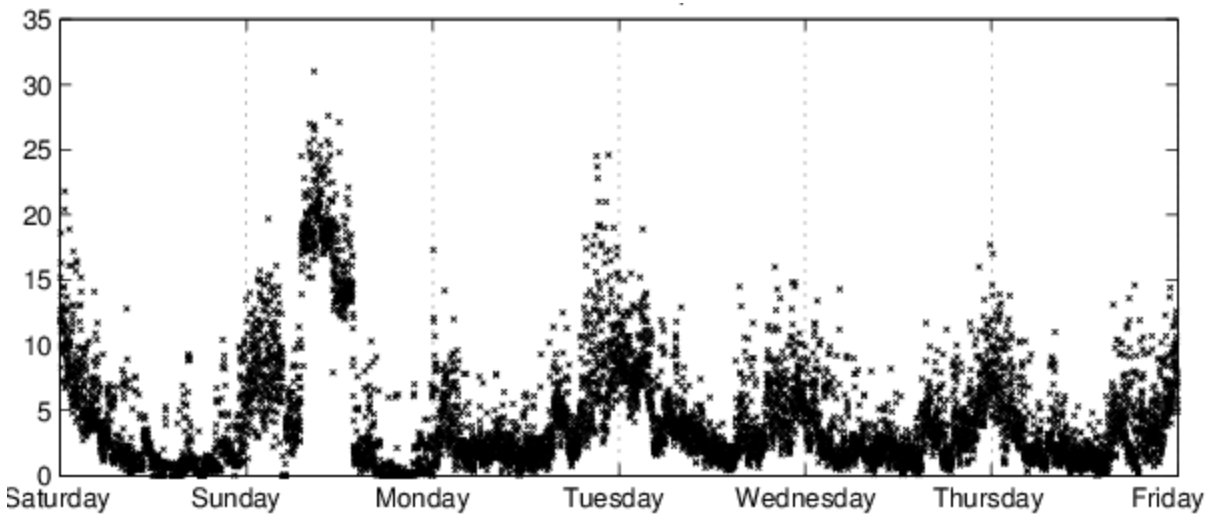

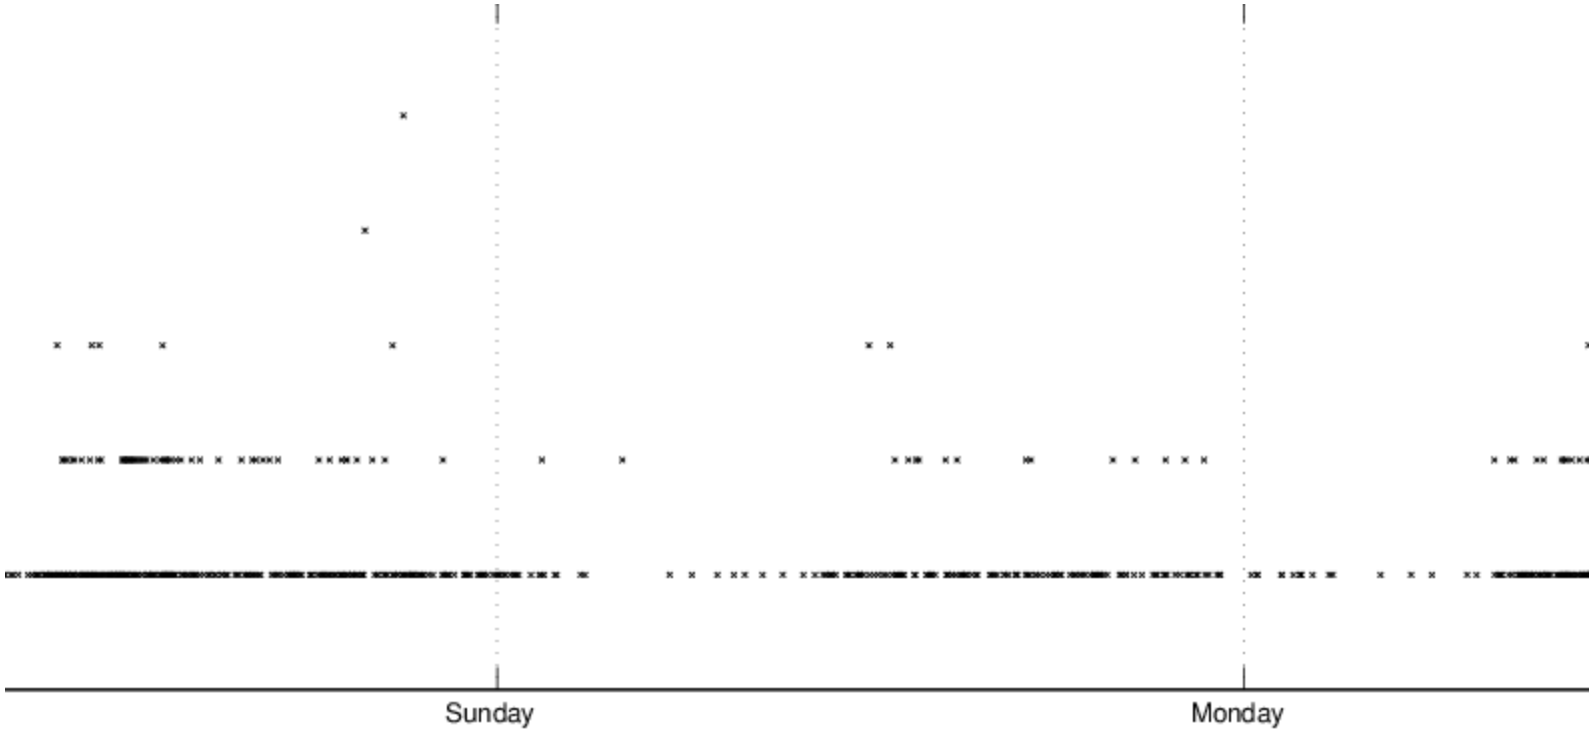
Sunday
Monday

Fits Work Related Data Profile

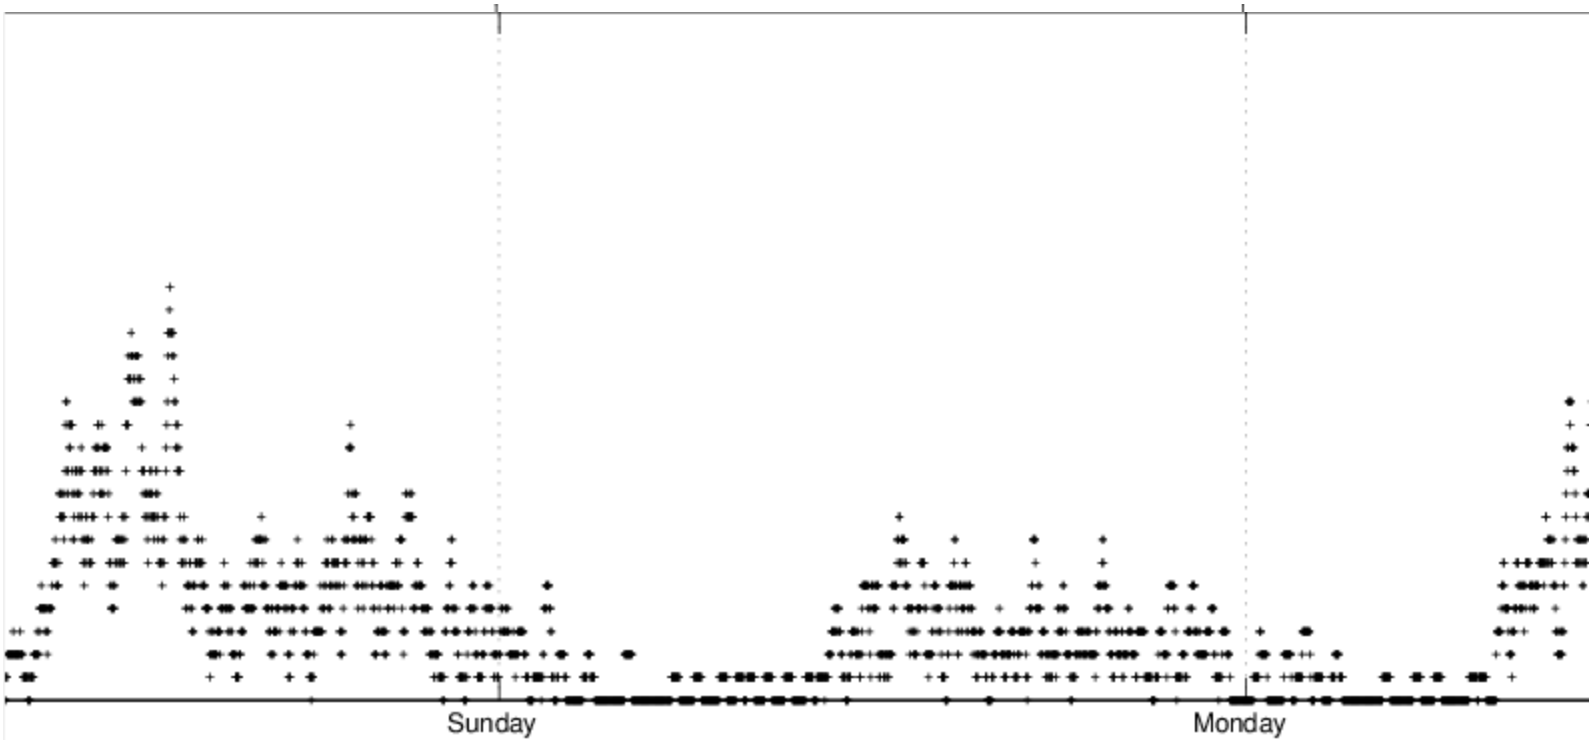

Example

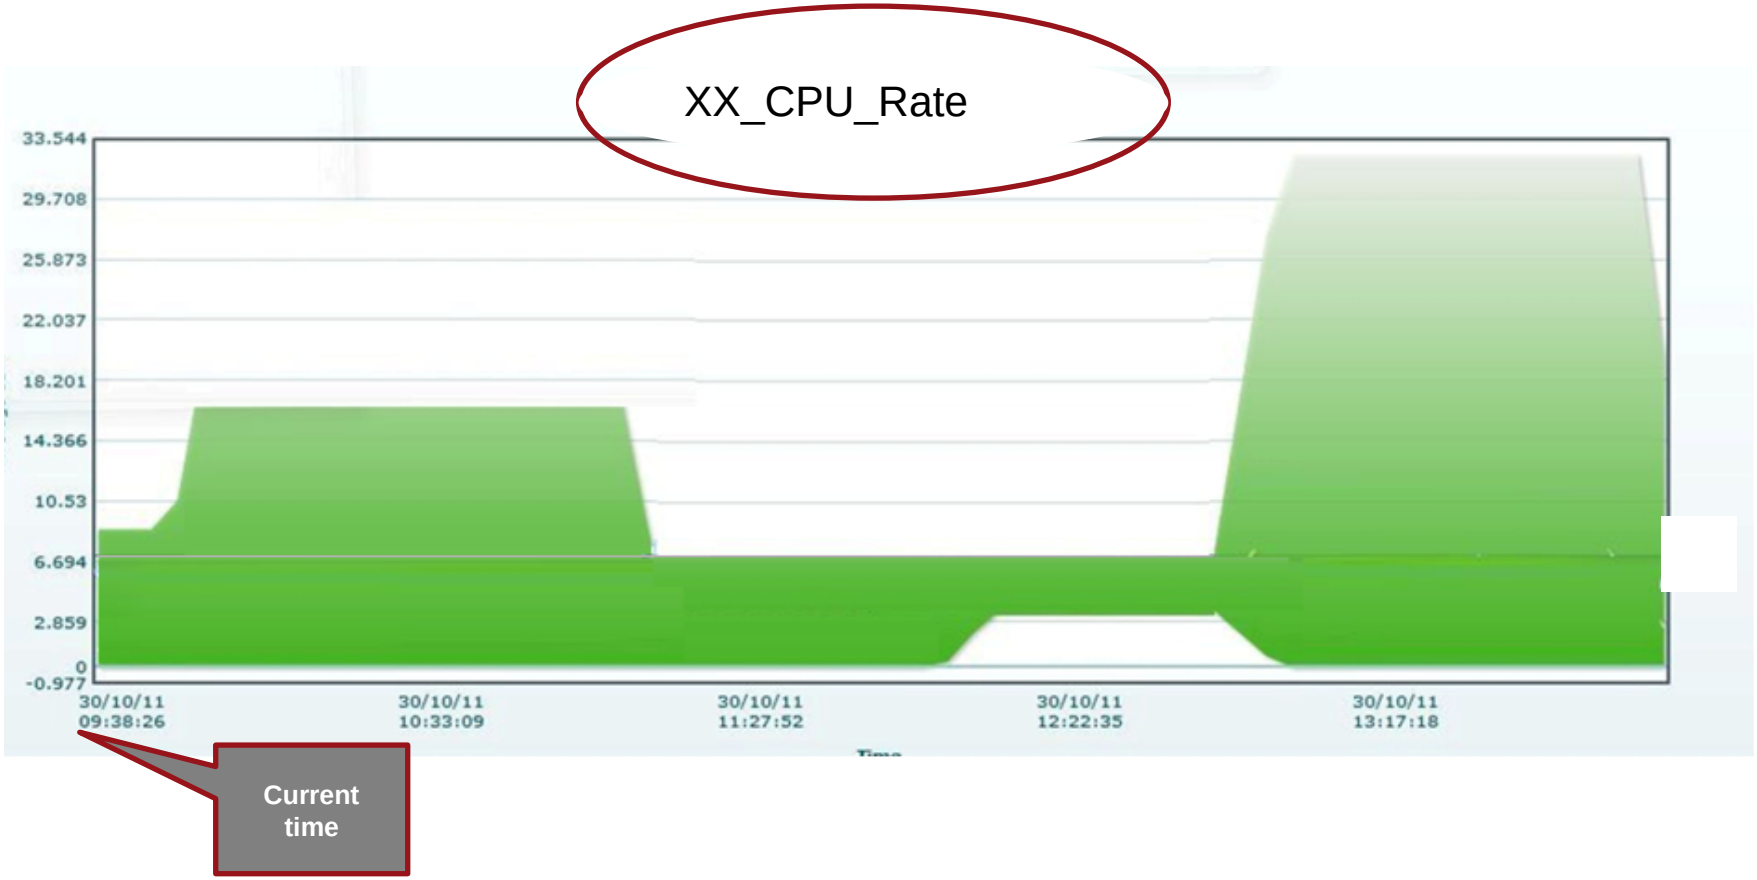
XX_CPU_Rate
33.544
29.708
25.873
22.037
18.201
14.366
10.53
6.694
2.859
0
-0.977
30/10/11 09:38:26
30/10/11 10:33:09
30/10/11 11:27:52
30/10/11 12:22:35
30/10/11 13:17:18
Current time

Example

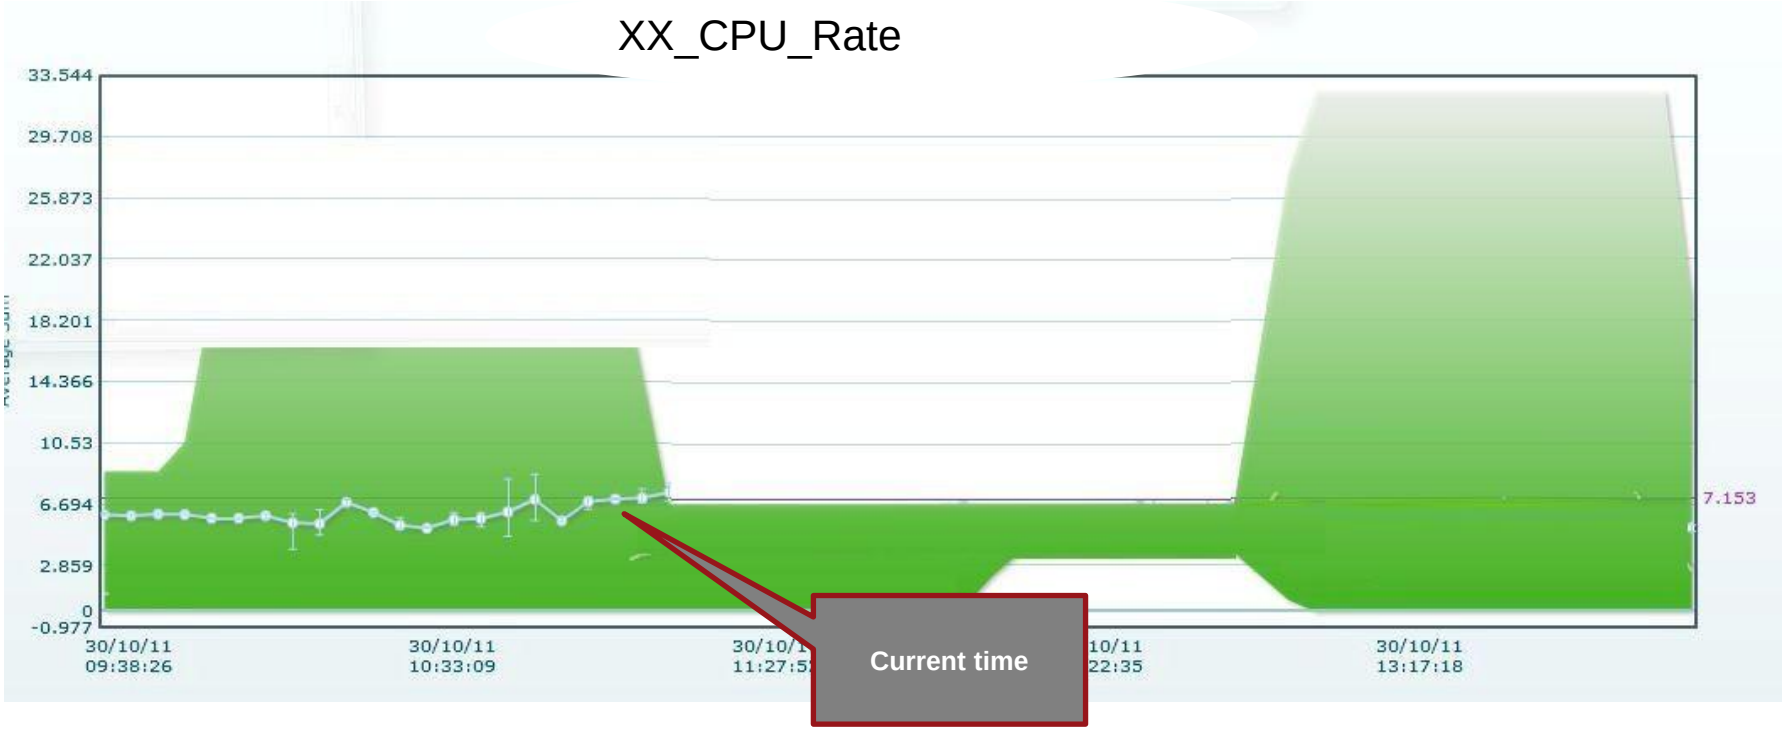
XX_CPU_Rate
33.544
29.708
25.873
22.037
18.201
14.366
10.53
6.694
2.859
0
-0.977
7.153
30/10/11
09:38:26
30/10/11
10:33:09
30/10/11
13:17:18
Current time

Example

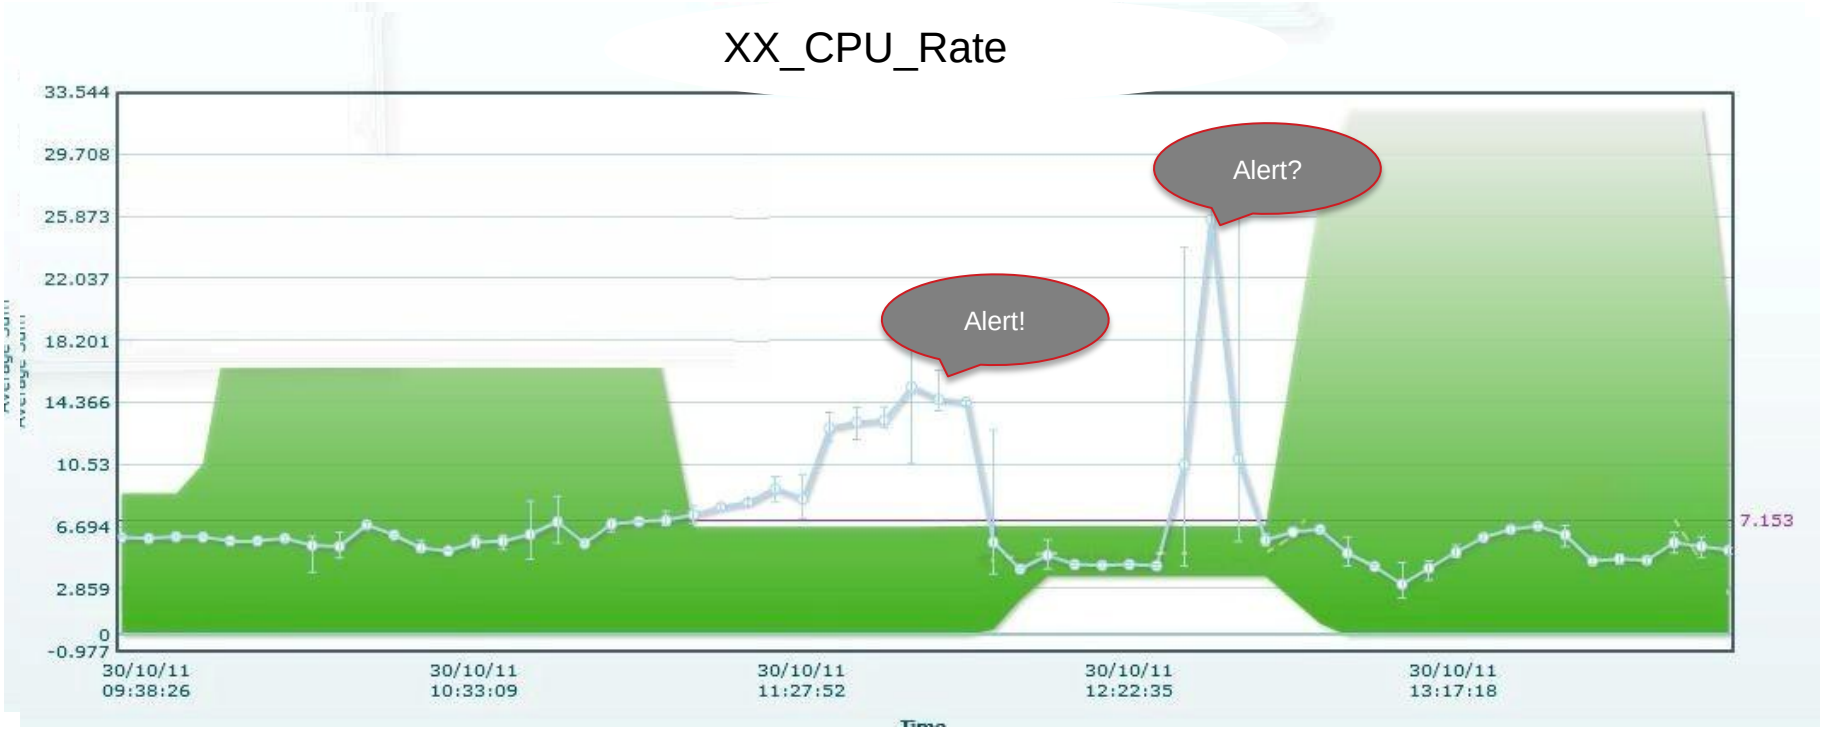
XX_CPU_Rate
33.544
29.708
25.873
22.037
18.201
14.366
10.53
6.694
2.859
0
-0.977
7.153
Alert?
Alert!
30/10/11
09:38:26
30/10/11
10:33:09
30/10/11
11:27:52
30/10/11
12:22:35
30/10/11
13:17:18

Lowering Costs using Behavioral Analysis

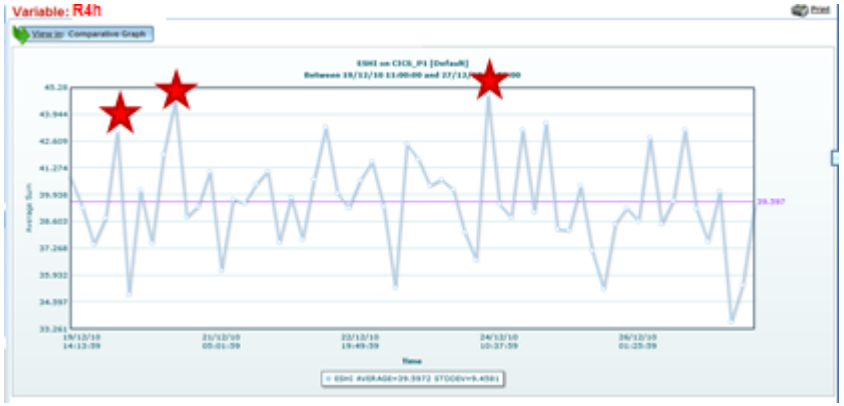

Results

- Improved Service
  - Find anomalies before they degrade service for users.
  - Users call about service degradation **after** technical support already knows about the problem and **is working on resolving it.**

- Lower Cost
  - Improved service quality with less cost
    - Save expert's time; engage only the specific, relevant experts to fix problems
  - Reduced lost opportunity costs
    - Forestall outages and faster resolution
  - Save MIPS
    - Discover application runtime defects that use more resources than needed

#SHAREorg

SHARE
Technology · Connections · Results

Thank you!

In the US contact our distributor:

Software Diversified Services (SDS)
sales@sdsusa.com
Phone: 763-571-9000

Outside the US :

info@conicit.biz
http://www.conicit.biz

SHARE in Anaheim 2012