

4th annual Zabbix conference
12 - 13 September 2014 | Riga, Latvia

ZABBIX 2014 Conference

Proactive and Reactive Monitoring







Serg Mescheryakov, Doctor of Science, Professor
Dmitry Shchemelinin, Philosophy Doctor

RingCentral Inc., San Mateo, CA, USA



RingCentral IP Telecommunication Company

RC is a fast growing multiservice system:

- ✓ VoIP and MoIP services 
- ✓ Audio and video conferencing 
- ✓ Internet fax 
- ✓ Call recording, forwarding, etc. 
- ✓ Mobile apps, softphone 
- ✓ Other web services 



3.5K production hosts in 2013
(North America and Europe)

5K production hosts in 2014
(40% annual non-stop growth)



RingCentral Distributed Cloud Infrastructure


RC is one of the biggest environments ever monitored by Zabbix



More customers and services – the bigger environment:

- 4 data centers (West and East US Coast, Europe)
- 5K+ hosts (hardware, virtual, aggregated) grouped by 90 components
- Multiplatform, multi-OS (MS Windows, Linux, Mac OS, iOS, Android)
- In-house applications, external providers, 3rd party vendors
- Oracle, MySQL, MongoDB databases
- 300K Zabbix items, 90K triggers per each of 4 locations
- 1.5 minutes average polling interval, 3K values per second traffic

Zabbix Databases



Zabbix DB is like a huge iceberg:

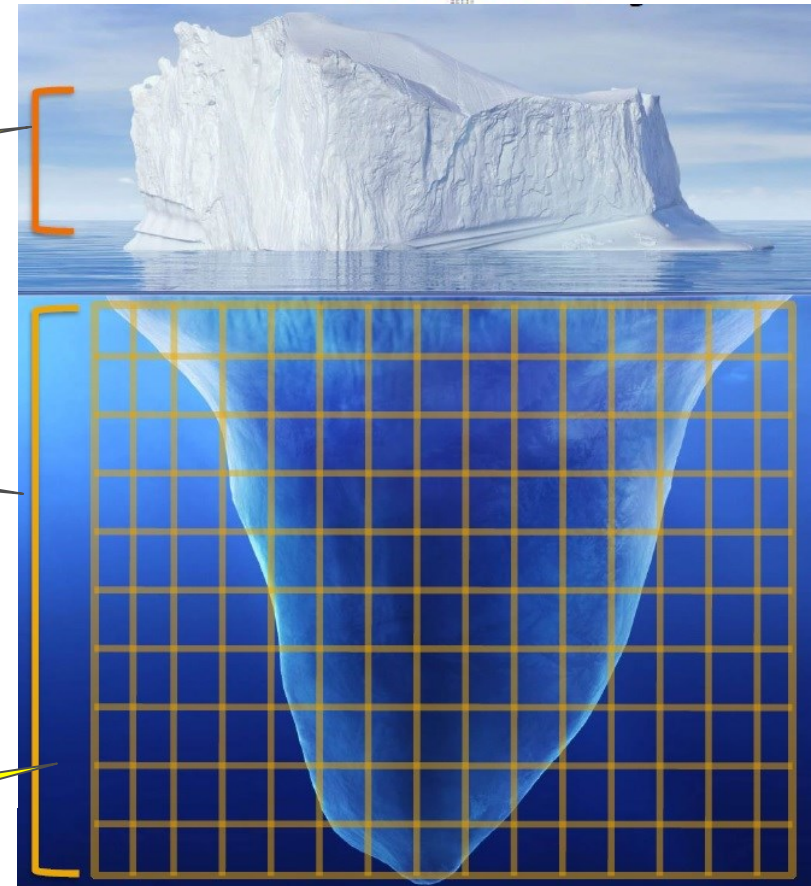
- ❑ MySQL for real time data (1 month)
- ❑ 4 DBs (one for each Zabbix instance)

Real time data may not be enough to analyze historical trends

History data is stored as truncated min, max, and average values

- ❑ NoSQL MongoDB for history (1 year)
- ❑ Distributed architecture for scaling
- ❑ The reads from and the writes into DB are separated for better performance
- ❑ Presented to Zabbix Conference 2013 by Leo Yulnets (RingCentral)

Problem 1:
DB performance is a bottleneck



Standard Approach to Zabbix Monitoring

Custom dashboard is created to monitor Zabbix events from all data locations. Traditional approach is to watch the dashboard, listen to alerts, and escalate.

	Timestamp	Host	Trigger	Group	POD/CL	Loc.
Information	2013-04-24 21:55:04	p08-cps01	Rebooted	CPS	P08	IAD
Critical	2013-04-24 21:48:43	p08-cps01	ICMP packet loss > 20%	CPS	P08	IAD
Warning	2013-04-24 21:46:29	p08-pws03	Problem with CPS connections	PWS	P08	IAD
Warning	2013-04-24 21:46:26	p08-pws02	Problem with CPS connections	PWS	P08	IAD
Warning	2013-04-24 21:46:24	p08-pws04	Problem with CPS connections	PWS	P08	IAD
Warning	2013-04-24 21:46:23	p08-pws01	Problem with CPS connections	PWS	P08	IAD
Critical	2013-04-24 21:30:05	p08-adb01	FAX Queue "Rendering queued" size more 1000	ADB	P08	IAD
Warning	2013-04-24 21:25:06	p08-adb01	FAX Queue "Rendering required" size more 1000	ADB	P08	IAD
Service Outage	2013-04-24 21:10:06	p08-adb01	SLD: Rendering Queue "Rendering required" size more 1000 during 15 mins	ADB	P08	IAD
Warning	2013-04-24 21:00:06	p08-adb01	FAX Queue "Rendering required" size more 1000	ADB	P08	IAD
Service Outage	2013-04-24 20:55:06	p08-adb01	FAX Queue "Rendering required" size more 1000	ADB	P08	IAD
Warning	2013-04-24 20:50:06	p08-adb01	FAX Queue "Rendering required" size more 1000	ADB	P08	IAD
Warning	2013-04-24 20:45:06	p08-adb01	FAX Queue "Rendering required" size more 1000	ADB	P08	IAD
Critical	2013-04-24 20:10:06	p08-adb01	FAX Queue "Rendering queued" size more 1000	ADB	P08	IAD
Warning	2013-04-24 20:05:06	p08-adb01	FAX Queue "Rendering required" size more 1000	ADB	P08	IAD

Zabbix is the main but not the only data source for escalations

Data Source	Percent
Zabbix alerts	80
Hourly service tests	10
Email from customer support	5
Other monitoring tools	5

Standard Approach to Zabbix Monitoring

Custom dashboard is created to monitor Zabbix events from all data locations. Traditional approach is to watch the dashboard, listen to alerts, and escalate.

	Timestamp	Host	Trigger	Group	POD/CL	Loc.
Information	2013-04-24 21:55:04	p08-cps01	Rebooted	CPS	P08	IAD
Critical	2013-04-24 21:48:43	p08-cps01	ICMP packet loss > 20%	CPS	P08	IAD
Warning	2013-04-24 21:46:29	p08-pws03	Problem with CPS connections	PWS	P08	IAD
Warning	2013-04-24 21:46:26	p08-pws02	Problem with CPS connections	PWS	P08	IAD
Warning	2013-04-24 21:46:24	p08-pws04	Problem with CPS connections	PWS	P08	IAD
Warning	2013-04-24 21:46:23	p08-pws01	Problem with CPS connections	PWS	P08	IAD
Critical	2013-04-24 21:30:05	p08-adb01		ADB	P08	IAD
Warning	2013-04-24 21:25:06	p08-adb01		ADB	P08	IAD
Service Outage	2013-04-24 21:10:06	p08-adb01		ADB	P08	IAD
Warning	2013-04-24 21:00:06	p08-adb01	FAX Queue "Rendering required" size more 1000			
Service Outage	2013-04-24 20:50:06	p08-adb01	FAX Queue "Rendering required" size more 1000			
Warning	2013-04-24 20:40:06	p08-adb01	FAX Queue "Rendering required" size more 1000			
Warning	2013-04-24 20:30:06	p08-adb01	FAX Queue "Rendering required" size more 1000			
Critical	2013-04-24 20:10:06	p08-adb01	FAX Queue "Rendering required" size more 1000			
Warning	2013-04-24 20:00:06	p08-adb01	FAX Queue "Rendering required" size more 1000			

Problem 2: manual supervision, nice to have been automated

Zabbix is the main but not the only data source for escalations

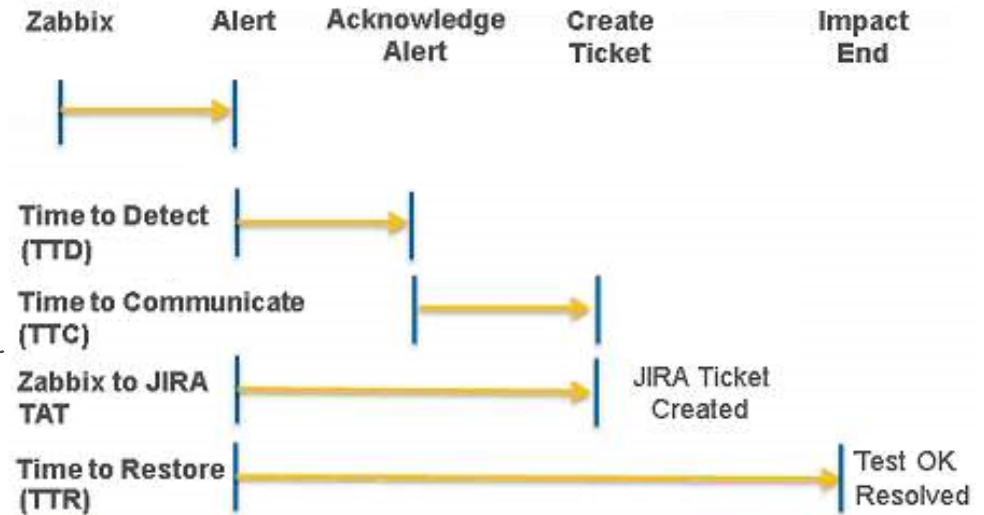
Data Source	Percent
Zabbix alerts	80
Hourly service tests	10
Email from customer support	5
Other monitoring tools	5



Escalation Process Workflow



Escalation process starts together with Zabbix alert



Problem 3:
data and other delays

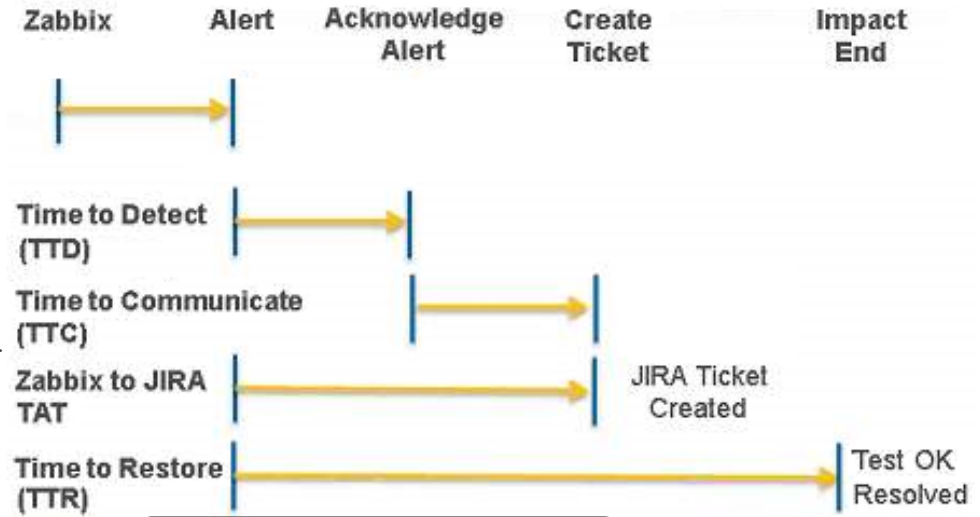
Overall impact is too big
and should be reduced

Key Performance Indicator	Time
1. Zabbix data delay	2 minutes
2. Time to detect	2 minutes
3. Time to communicate	2 minutes
4. Time to acknowledge	4 minutes
5. Time to restore and test	6+ minutes
Total impact	10+ minutes

Escalation Process Workflow



Escalation process starts together with Zabbix alert



Problem 3: data and other delays

Key Performance Indicator	Time
1. Zabbix data delay	2 minutes
2. Time to detect	2 minutes
3. Time to communicate	2 minutes
4. Time to acknowledge	4 minutes
5. Time to restore and test	6+ minutes
Total impact	10+ minutes

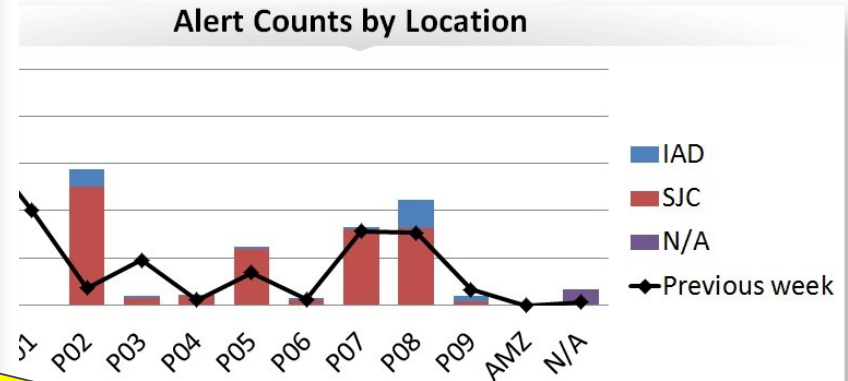

Overall impact is too big and should be reduced

Solution: To resolve the alerts faster, monitoring should be proactive


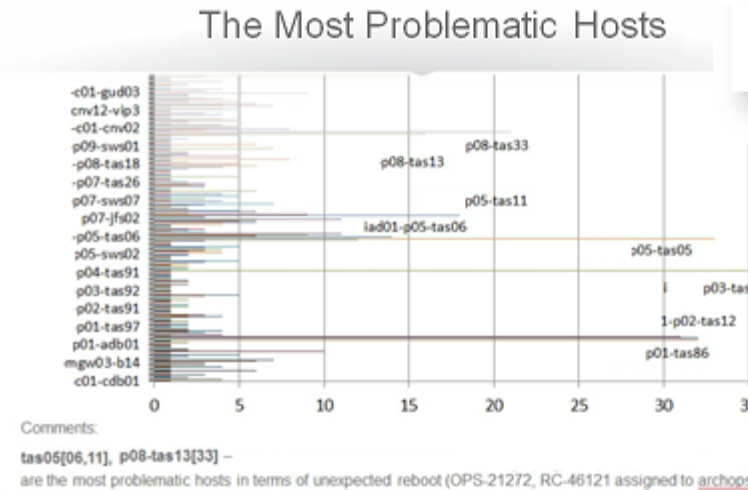


Regular Analysis of Zabbix Statistics


The goal of analysis is to find out repeatable alerts, problematic hosts, bottlenecks.



Zabbix is a powerful tool for analysis



Example 1: Monitoring Java Resources of JEDI Servers



Problem: memory degradation may lead to application crash

Certain time is needed to restore the service

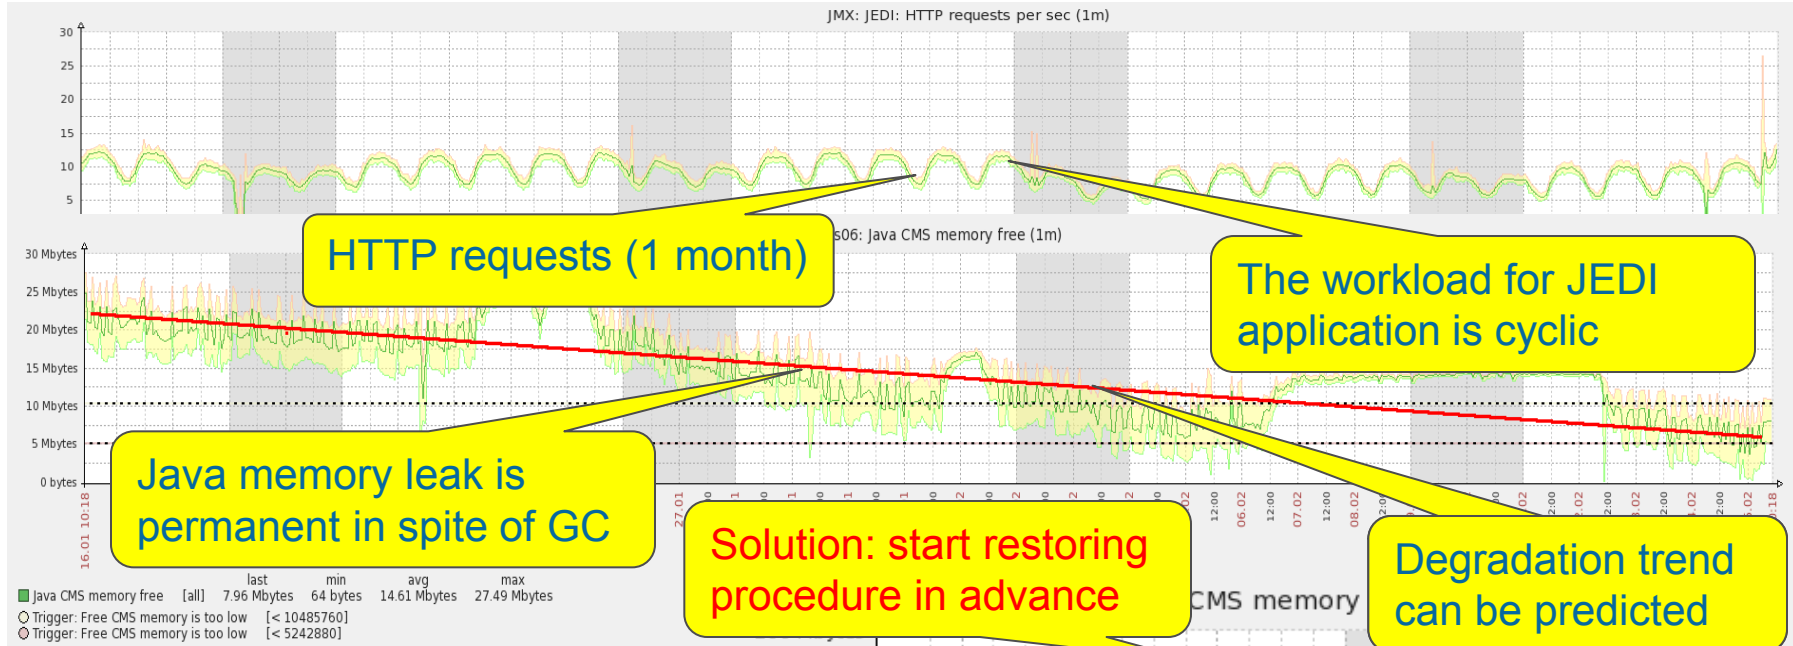
Monitor

CPU

Memory



Example 1: Detecting Critical Degradation of Java Memory




$$y = -30.52x + 407.36; R^2 = 0.79$$

Ratio of variance:

$$R^2 = 1 - \frac{V(y|x)}{V(y)} = 1 - \frac{\sigma^2}{\sigma_y^2}$$

Standard deviation:

$$\sigma^2 = \sum_{t=1}^n (y_t - \hat{y}_t)^2$$



$$y = -24.67x^2 + 2 \cdot 10^6x - 4 \cdot 10^{10}; R^2 = 0.98$$

$$y = 8.57x^2 - 141.91x + 667.27; R^2 = 0.96$$

$$y = -10.14x^3 + 10^6x^2 - 5 \cdot 10^{10}x + 7 \cdot 10^{14}; R^2 = 0.89$$

Example 1: Using MS Excel to Evaluate the Trend



Example 1: JMX: Introducing New Zabbix Items


Name	Last check	Last value	Change
JMX_JEDI (6 Items)			
JMX: JEDI: HTTP requests processing time	14 Mar 2013 13:55:53	406	-15
JMX: JEDI: HTTP requests per sec	14 Mar 2013 13:55:53	0	-
JMX: JEDI: Active sessions	14 Mar 2013 13:55:53	47	-
JMX: JEDI: Session alive time (average)	14 Mar 2013 13:55:53	0	-
JMX: JEDI: Session alive time (maximum)	14 Mar 2013 13:55:53	0	-
JMX: JEDI: Web processing active thread count	14 Mar 2013 13:55:53	0	-
XPDB (4 Items)			
JMX: JEDI: Active XPDB connections count	14 Mar 2013 13:55:53	0	-
JMX: JEDI: Total XPDB connections count	14 Mar 2013 13:54:24	3	-
JMX: JEDI: XPDB objects count	14 Mar 2013 13:55:53	57026	+13
JMX: JEDI: Active XPDB sessions count	14 Mar 2013 13:55:53	88	+1

Why JMX for JEDI?


The benefits of JMX metrics:

- JMX – Java Management eXtensions (refer to Zabbix documentation for more details).
- JMX tools provide access to Java API internal objects, classes, services.
- JMX metrics allow measuring real workload and actual utilization of Java resources.

Example 1: JMX: Measuring Actual JEDI Utilization



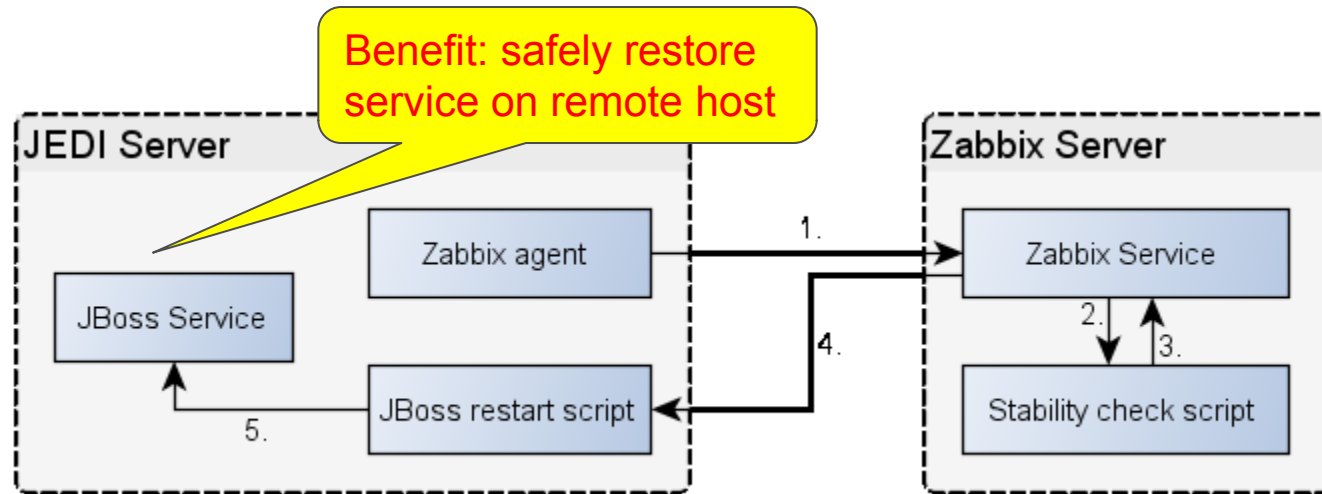
Example 1: Detecting Critical Degradation of a Host



Critical degradation of Java memory on a single host

User activity on JEDI

Example 1: Auto-remediation Workflow for JEDI



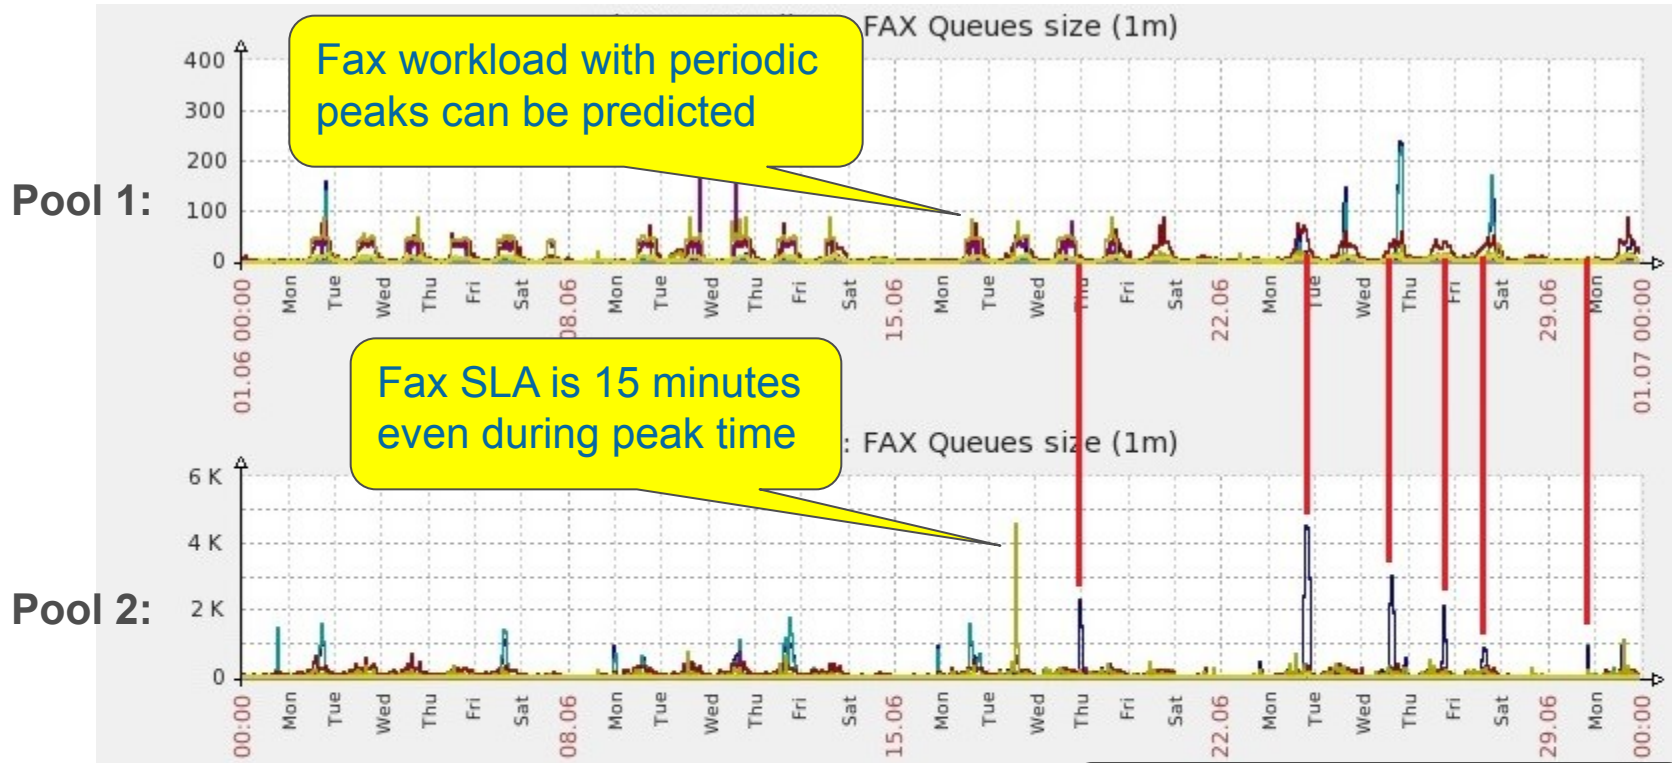
Benefit: safely restore service on remote host

The exact threshold numbers depend on the time required to safely restore the service

Auto-remediation procedure is initiated in case of critical degradation of Java memory:

- Java free memory tends to degrade below critical threshold of 5 MB in 15 minutes.
- Java virtual allocated memory is higher than allowed (2 GB for 32-bit VM, 6 GB for 64-bit).
- The JBoss service status on remote host is not responding (for at least 1 minute).
- The JBoss service on the other JEDI hosts in pool is up and running (stability check).

Example 2: Monitoring Fax Capacity in Pool



Pool 1:

Pool 2:



Fax workload with periodic peaks can be predicted

Fax SLA is 15 minutes even during peak time

Problem: big redundancy of fax servers (10% utilization)


		last	min	avg	max
FAX Rendering InProgress	[max]	8	0	2.37	20
FAX Rendering Queued	[max]	42	0	47.63	4.52 K
FAX Rendering Required	[max]	43	28	50.33	1.76 K
FAX Sent InProgress	[max]	60	11	84.6	664
FAX Sent Queued	[max]	1	0	1.45	766
FAX Sent Required	[max]	27	0	26.35	4.6 K
FAX Send Not Rendered	[avg]	0	0	0	0
FAX Senders Count	[max]	12	0	8.1	87

Example 2: Managing Fax Capacity in Pool



	last	min	avg	max
FAX Rendering InProgress	max] 9	0	3.24	11
FAX Rendering Queued	max] 157	0	33.22	402
FAX Rendering Required	max] 54	0	23.26	342
FAX Sent InProgress	max] 237	14	128.46	556
FAX Sent Queued	max] 3	0	11.13	738
FAX Send Required	max] 80	0	71.09	3.55 K
FAX Send Not Rendered	max] 0	0	0	0
FAX Senders Count	max] 53	0	21.33	129

KPI metrics are introduced to detect the peak load and reallocate fax standby resources when needed



Benefit: extra resources can be removed without a risk of outage



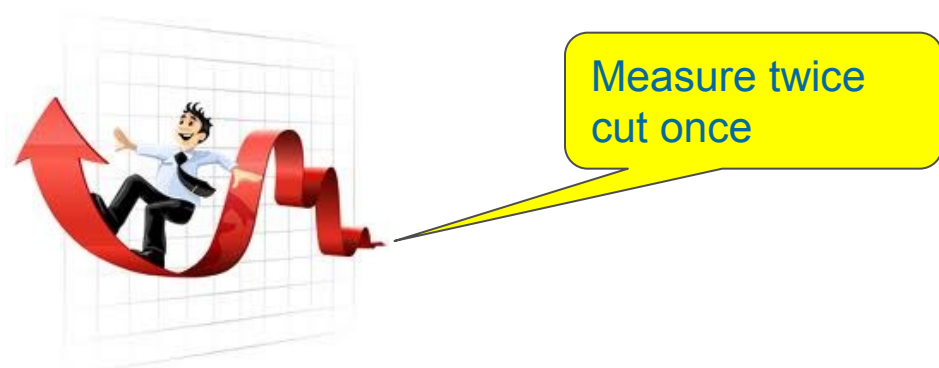
Conclusions

- Zabbix is not only monitoring, but is also a reactive centralized control system, especially in a big distributed system like RingCentral telecommunication company.
- Reaction to the alert is often not fast enough. Monitoring should be more proactive.
- New metrics to detect the critical degradation trend in advance are proposed.
- Corresponding action items and triggers are implemented in Zabbix.
- The 2 real-world examples of proactive monitoring are presented:

1. Preventive auto-restoring the Java service on JEDI host when critical degradation of Java memory is detected.



2. Allocating additional computing resources “on the fly” if the estimated fax capacity of entire pool is not enough.



Thanks!