



# ChinaNetCloud

*Running the World's Internet Servers*



## Zabbix

## At Scale



*By Steve Mushero*

*September, 2014*



# Greetings

---



**I'm Steve**

**I'm from Shanghai, China**

**We have a big Internet there**

**We have a big business**

**We have a big monitoring system**


**That's Zabbix**

**Let me tell you more about it . . .**



# First, a Word about Numbers

---



**My Title: 2500 hosts & 1 million items**

**Well, I wrote that in spring**

**We're not quite there yet**

**Today – about 1250 hosts & 300K items**

**With 75K triggers**

**500 templates**

**25+ graphs per host**

**200 screens & maps (60% customer specific)**

**10 proxies around the world**

**About 550 new values per sec (NVPS)**


**Hope you're not too disappointed**

**Still lots to share**



# More Numbers

---



**This is a busy system**

**About 1250 hosts & 300K items**

**Billions of data points in DB**

**Hundreds of users**

**78,000 Alerts/Events per month**

**2,600 per day**

**Mostly false alerts**

**Working hard to reduce these**

**1,000+ Actionable Tickets per month**



# What we do

---



**Internet Managed Service Provider**

**But higher-end than most**

**We build & run large-scale Internet systems**

**In any data center in Asia, USA, and beyond . . .**

**Our customers have about 250M users**

**We have over 100 customers**


**1000 servers for 100 customers <> 1000 for yourself**

**We run about 200 DB servers, and dozens or hundreds of everything else**

**We run every possible service and system (Linux-based)**

# How we do this

---



**About 100 employees, 50% in Operations**

**About 25 integrated systems**

**Zabbix, ticket, customer, procedure, logs, CMDB, etc.**

**7x24 Operations**

**Continual new customers/systems**

**Continual changes to systems/monitoring**

**Customers rarely test things**

**VERY dynamic environment**


**Exact opposite of most Enterprises**

**Think Wild, Wild West**



# How we use Zabbix

---



**Core monitor system for six years**

**Alerts link to our Procedure system**

**Via URL in Trigger, customized**

**We are on 1.8.3 and -> 2.2 this month**

**Other systems feed it via sender (DWM, SEC)**

**Heavily modified dashboard with ticket integration**


**Automation & reporting**

**We have 100 day-time users, about 25 hard-core**

**All our customers use it, varying degrees of use**

**Lots of macros (and reports)**

# Our Ops Area – Zabbix Everywhere



## Image of AR/Ops area





# How we use Zabbix

---



**Very dynamic system**

**Things changing every day, every hour**

**Make changes every day, by dozens of people**  
**Servers coming & going, updating templates**

**Very easy to misalign, misconfigure – have safety report**

**We have 500 templates**  
**About 25 that get used for every customer**  
**By service – Apache, MySQL, etc.**



# How we use Zabbix – Templates

---



## About 500 templates

- Many customer specific or obsolete
- Trying to reduce to < 100

## One core Template for Linux

- 163 Items

- 79 Triggers

- Will soon split to RedHat vs. Ubuntu


## One Template per Service

- MySQL, Apache, Nginx, many, many more

SNMP devices also, messy ports

# How we use Zabbix – Agents

---



## Agents on all non-SNMP hosts

### 100% Passive Agents

#### Server Security issues with Active

No good way to lock down 100 locations

**So every item/new value is an update to DB**

### LOTS of custom scripts – 25+

MySQL – Real-time any variable/config/status

Java – Custom JMX reader

Lots of services – Haproxy, Apache, Node.js, etc.

Linux /proc reader – TCP, conntrack


iostat/vmstat

RAID/Hardware

CatchAll low disk filesystem – Very important

# How we use Zabbix – Pushing Data

---



## Several monitors that push data

### We patched Server to allow direct reporting

**Normally must push data to proxy**

Very painful globally

### Syslog via SEC

**Pushes fixed set of items like OOM**

**Clear manually in Dashboard GUI**

### Distributed Web Monitors (DWM)

**Push status, failures to server**

# How we use Zabbix – Web Checks

---



**We don't use any more**

**Not distributed (until 2.2)**

**Can't get error reasons**

**Makes troubleshooting harder**

**Easy to break triggers**

**Built our own Distributed Web Monitor**

**Runs on nodes / proxies**

**Hosts are in zones/nodes**

**Nodes pull config**

**Parallel checks on nodes**

**Detailed errors/triggers**


**Single node per host**

**Soon to be multi-node with voting**

**More sensitive, less false-alerts**

# How we use Zabbix – Other Parts

---



## Screens

- Used heavily
- Hard to build, need automation
- XML Import/Export
- Not used much
- But powerful

## API

- Use for adding servers
- Use to pull graphs

## Maps

- Use a few
- Automate in future

**IT Services & Reports – Don't use**

**Discovery – Thinking how to use**



---

# Experience

# Our General Experience

---



**Overall it's very good**

**False alerts kill us**  
**Especially unreachables**

**Proxies kill us**

**Lots of patches/changes**

**Long list of improvements**

**We will be world's largest user**





# Issues & Improvements

---



## **Time checks – Fuzzy proxies**

**Totally useless globally**

**No solution yet, just SQL reports**

## **What is unreachable, when, how ?**

## **Many agent improvements needed**

**Linux memory RSS**

**iostat devices**

**We have long list**

## **Latching – Need a way to latch**

**To keep short-lived events up**

Can use hysteresis, but not great



# Scale issues

---



**Housekeeper**

**Large templates – core Linux template**

**Database size**

**Disk I/O challenges**

**Operations time out**

**Or lockup system**

**Proxies challenging**

**Poor Permission System**





---

# Zabbix Details

# Our Zabbix Architecture

---



**Central Server in Shanghai**

**Proxies in 10 locations in 5 countries**

**Most support many customers**

**Some dedicated in customer data centers**

**System consists of 3 large VMs:**

**Web GUI**

**Zabbix Server**

**MySQL Database**

**Looking at geographic HA options**

**Need to move all hosts to proxies first**

# Hardware History

---



**2008 - Dell R200 Virtualized – up to 350 hosts**

**2010 - Dell 1950 Virtualized – up to 750 hosts**  
**Integrated single larger VM – GUI, Server, DB**



**2013 - Dell R420 Virtualized with SAS – up to 1250 hosts**  
**Split VMs – GUI, Server, DB**  
**96GB or RAM**


**2014 - Dell R420 Virtualized with SSD – should go 5-10X**  
**Split VMs – GUI, Server, DB & slave DB for reporting/backup**  
**128GB or RAM, looking at 256GB**

**Bottleneck is always DB RAM & I/O**  
**CPU always okay**

**Eventually 256-512GB of RAM and 500GB SSD - Happiness**

# Database Details

---



**MySQL Percona 5.5 w/ very optimized config**

**Will go to 5.6 when 5.6 is more stable**

**No query cache**

**Main DB about 200GB in size**

**Expect 1TB in 2015**

**History is about 2 billion rows of data, 157GB size**


**Trends is 500 million, 37GB size**

**Backups take 8 hours (including compress/encrypt)**

**Has 48GB of Innodb Buffer (will expand to 64+ soon)**


**Will go to 96, 128, 256 over time**

**Generally runs well, except I/O bottleneck**



# Zabbix Server Config

---



**Nothing that special**

**Pollers – Use a lot**

**Max num ~500 in most cases**

**Can be 90-100% busy**

**DB Sync is slow, screws up time**

**Data time-stamped on db sync, not received**

**Housekeeper On, Scheduled**

**100% passive**

**Agents locked down to server/proxies**

**Don't use Actions**

# Housekeeper

---



**Total rows in history tables: 2 BILLION**

**Hourly housekeeping delete: 7 MILLION**

**Takes 3 hours to run**

**Killing us if 8 hour backups are running**

**Did not partition yet, problems in 2.0/MySQL**

**Changed code to make it configurable**  
**Still very slow, as item by item**

**Now only run during day, avoid backups**

**Hourly (shortest option)**

**Will go to partitions in 2.2**







**We hate proxies, sometimes**

**Great in theory, but**

**Always getting stuck – very poor network handling in 1.8**

**Fuzzytime() useless for time check – no solution**

**Server time is not data gathered time**

**We run them in many data centers, countries**

**With often poor connections to main servers**

**We build SQL tools, with PHP GUI coming soon**

**Poller Queue, Sending Queue, Oldest unsent value**

**Coping**

**Use cron to check queue and auto restart**

**Use auto re-route to improve traffic via TCP relays**

**Hoping for better TCP/IP in 2.x**




---

# Customizations

# Customizations – GUI Dashboard

---



**Mostly change columns & colors**

**We only use Top Issues section**

**Added custom columns**

**Mostly ticket system integration**

**Ticket numbers & status/history**

**ACKs include person's initials**

**Status if someone logs into host**

**Changed popup menus**

**Links to various systems**

**Create ticket popups**

**Triggers Dynamic link to procedures**

**Sends hostname and profile fields**

**New dynamic priority system**

**Rule & time-based priority changes**

**Auto escalation, tracking**



# Customizations – GUI Dashboard



**Popup Menu Customized**

**Ping – Via ssh/get/snmp**

**Server Summary**

**Custom history, owner, contact**

**Create Ticket**

**Wikis for each level**

The screenshot shows a table titled "Last 20 Issues" with two columns: "Host" and "Issue". The "Host" column contains entries like "srv-...", and the "Issue" column contains entries like "High IO...", "MySQL...", "Distribu...", "erver s...", "erver s...", "erver s...", "erver s...", "MySQL...", "oad av...", "MySQL...", "MySQL...", "erver s...". A custom popup menu is overlaid on the "Host" column, listing the following items: **Tools**, Ping Host, Traceroute, **Links**, Latest data, Profile, Server Summary, Create Ticket, View Triggers, **Wiki**, W-Customer, W-System, and W-Server. The "Tools", "Links", and "Wiki" items are highlighted in blue.

Host	Issue
srv-...	High IO...
srv-...	MySQL...
srv-...	Distribu...
srv-...	erver s...
srv-...	erver s...
srv-...	erver s...
srv-...	erver s...
srv-...	MySQL...
srv-...	oad av...
srv-...	MySQL...
srv-...	MySQL...
srv-...	erver s...

# Customizations – GUI Dashboard



## Several Dashboards

**Main Dashboard - Internal**

**Secondary Dashboard**

We push things we can't fix

**Customer Dashboard**

More limited view

## Graphs

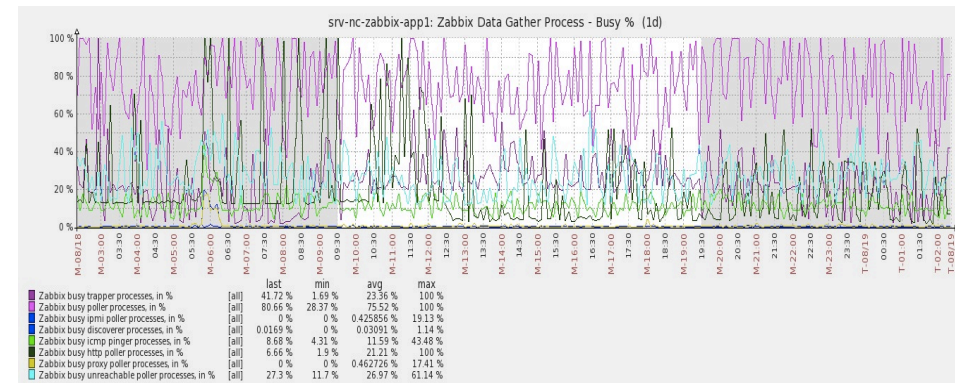
**American Dates**

**Fonts & Sizes**

**Many more coming**

Annotations, Max limits


**LOTS more for future improvements**



# Customizations – Graphs



srv-nc-zabbix-app1: Zabbix Data Gather Process - Busy % (1d)



		last	min	avg	max
Zabbix busy trapper processes, in %	[all]	41.72 %	1.69 %	23.36 %	100 %
Zabbix busy poller processes, in %	[all]	80.66 %	28.37 %	75.52 %	100 %
Zabbix busy ipmi poller processes, in %	[all]	0 %	0 %	0.425856 %	19.13 %
Zabbix busy discoverer processes, in %	[all]	0.0169 %	0 %	0.03091 %	1.14 %
Zabbix busy icmp pinger processes, in %	[all]	8.68 %	4.31 %	11.59 %	43.48 %
Zabbix busy http poller processes, in %	[all]	6.66 %	1.9 %	21.21 %	100 %
Zabbix busy proxy poller processes, in %	[all]	0 %	0 %	0.462726 %	17.41 %
Zabbix busy unreachable poller processes, in %	[all]	27.3 %	11.7 %	26.97 %	61.14 %

# Customizations - Other

---



## Housekeeper

### Configurable schedule

Avoids conflict with backups

## Direct data sender

### Bypass proxies

Major improvement

## Ping Host – Can't use behind proxy

Changed to use `zabbix_get`, `ssh-key`

Use SNMP tool for network hosts

## A LOT more changes planned for 2.2

Dozens, including agent, proxy UI





---

# Automation & Reporting

# Automation

---



**Lots of direct DB interaction**

**Dozens of Reports**

Safety checks

Monthly reports

Metrics reports

**Auto Ticket**

Creates tickets based on rules

**Auto Closer**

Closes tickets if alert disappears

**Auto Priority/Escalate**

Rules-based event driver

Dynamic priority engine


**Push new hosts in via API**

Some still direct DB inserts



# Reporting

---



**Two key types:**

**Customer Reports – Monthly / on demand**

**Custom selection by server, customer, graph**

**Use API to pull real graphs**

**Pull data from DB for metrics**

**Auditing – SQL-based**


**Ad Hoc – Lots of SQL-based**

**Safety Reports – Look for problems**



# Safety Reports

---



**Remember – We are in the Wild West**

**Look for mistakes**

**Direct from DB**

**See my DB talk for details**

**So many changes by so many people**

**Not enough automation yet**

**Dozens of people making changes**

Not all as careful as we'd like

**Easy to break things**

**Hard to know or detect**



# Safety Reports – Examples

---



**Items that differ from template**

**Missing templates**

**Disable items/hosts, forget to enable**

**Alerts with no URL/Wiki**

**Hosts missing profile data**

**Items disabled conflict with trigger**


**Web alerts with no trigger**

**Web alerts with long/short timeouts**

**Hosts in wrong, duplicate, conflicting groups**

**Servers in Zabbix, not core system**





**Pretty good**

**Very good for customers**

- One group per customer**
- Read-only rights**

**BUT, not granular enough**

- Either read or write**
- Can't have different level employees**
- Can't limit tasks / functions**
- Useless for ops company**

**Audit system great**

- But GUI & reports useless**
- We have custom reports**


**We'll modify this in 2.2**

- Not sure how yet**



# Future

---



**Lots more basic improvements to make**

**Changing event processing, rule-based**

**Thinking correlation engines**

**Adding new Agent features in 2015**

**New GUI & Security & Audit & Reports**

**Thinking about 10X scale for 2015-16**

**10,000 servers, millions of items**

**Architecture, Federation, ???**

**Wondering about 100X scale**



# Summary

---



**Zabbix is a great system**

**It's critical to our business**

**Still the best system out there**

**We've invested a lot, more to go**

**Still lots of improvements to make**


**Glad to see vibrant user & developers**

**Happy to be part of the Zabbix ecosystem**

**Happy to share all we know & have learned**



# Thanks from ChinaNetCloud



**Pioneers in OaaS – Operations as a Service**



**Shanghai Headquarters:**

X2 Space 1-601, 1238 Xietu Lu

Shanghai, 200032 China

T: +86-21-6422-1946 F: +86-21-6422-4911



**Beijing Office:**

Lee World Business Building #305

57 Happiness Village Road, Chaoyang District

Beijing, 100027 China



**Silicon Valley Office:**

California Avenue

Palo Alto, 94123 USA