

Advanced log monitoring

Giedrius Stasiulionis

Logs

One of the cornerstones in monitoring field

Logs provide visibility into the day-to-day operations of IT systems and applications, helping ensure smooth functioning

Zabbix & logs

log[]

log.count[]

custom items

Zabbix & logs – custom items

When to use it?

Only when you can't achieve desired result with `log[]` or `log.count[]`

Zabbix & logs – custom items

- monitoring rapidly updated files (600k+ lines per minute)
- something that you would collect only to use for calculated items
- multi-line monitoring
- monitoring logs as “Passive item”

Zabbix & logs – custom items

How?

Mimic the agent in simplified way

You must be able to read the log files in similar way as agent does

Zabbix & logs – custom items

Pareto principle (80/20 rule)

Concept: key points

- read log portion-by-portion, just like `log[]` and `log.count[]`
- make it fast – operate in bytes instead of lines
- don't forget log rotation
- choose whether to read or ignore unread data if agent was stopped or item was disabled

Concept: reading

During each run you have to operate with 2 numbers:

- 1) how much data was read during last run
- 2) how much data you have now

Concept: reading

```
2024-07-22 10:13:15 Ilzmmztzjg8C2wzq0DFc7C2Ktdvwlshrc2gbXJFfxutoYmDEcHvE4sK0y8bxgqxH
2024-07-22 10:13:15 PDNqWQ01WbM6SeulqQFSrUmQ0Y4yPz9Jqist4ogojiHNXrYJLEWizNws5x 7tjTA
2024-07-22 10:13:16 V0mgn1 u49swv7D0bCPayZvXyHTmuHlNx5N2lQaqn4drWooFHSjgcpPp 71mFgG5
2024-07-22 10:13:16 ZGgn rt13koIHV0z4mXP3kOPIbX JhgX9oDhhhNCGnY0czZV5AQ5ZErFGI2m0NpF
2024-07-22 10:13:17 61QN8eTV4ZlvGUxFNjzstEyCo 0HSdWjqPzS8tg8sgcJx7shXB7JHMq2D8P4F1M0
2024-07-22 10:13:18 iqwe2My2b4CsmkST0HCqR8FoFeh69mpuxRfyi2xsb24F9gMFYdh9pFJ1WCrmLbz
2024-07-22 10:13:18 YA2saj0chEvcx0Yh3XvkhFzqP0bm0eFJA0ErFTzo1yMtAznpwB0DxEs2alu7XmgQ
2024-07-22 10:13:18 8NUj55Cu00CKlQL6IBsYsNjL47Ppw6M2E3Q58TaHKfDCJDiyeXBj7G0ZHap s0iN
2024-07-22 10:13:18 rjj@dCwW08WzxWg5bKQIQis1Ni0Pb4S4f54YHvA7KTj36bSnK4PSaUqy7qiqQahf
2024-07-22 10:13:18 4CrT8DsD7S2wa5pWTgD6oCaKmpjtYD5Slp9L06HJuHKGu943QfZHa5CZSG53W0H6
2024-07-22 10:13:19 TrJX1sap0UqYbznEtQ8bCv9Jhd3ry5AlhQg6F1Jn6e1HkykJHhofk0G8ZRl0jTsn
2024-07-22 10:13:20 sl3e3QVUaZS2asLdlkYm1VQ1u lQA8BmtZ7bLaG5GY8y9bHUtuLAarcC71r9XS8
2024-07-22 10:13:20 1WCw9f5QKN0Zl0AHDMBooTWhemZLRsvEoCTbgnrNcoUZs sdCwSyZknISuXjbJZt
2024-07-22 10:13:20 bK4uE0o pzx63tbb3wzybL3QblqHhr9SqooMdl0ZSmr7xJzn7TyJ00B16df0jQv6
2024-07-22 10:13:21 y iHR1enmvSKj0YeUBkW9l4RrhZiVdCpzmDQ7JTF17pxDExWnrFnc3sib2a85ffb
2024-07-22 10:13:21 CcS1dIxE6wMRS2LsoZ2Y9P5c qwFLSeKw35AbZtj5B3mkB8UFS4TVGMKQTGSatGP
2024-07-22 10:13:22 A7HccP3YpnFltFIsLLE3t5xpRHd0V2ej7mSj4brE0xWdpknrxPwfCFv57SoZKTwd
2024-07-22 10:13:23 FTd7kacSbykKrL0tTTyiW2uvDrv3dFEeYcQ0Dh9UxsP8T79D976nHbILc2S17aSx
2024-07-22 10:13:24 X9I4pCDYnLQTqrGpVCg0421ZZIYpCEb0uoWfgSCNVENCQjbMA10KSezuAaah0dJ4
2024-07-22 10:13:25 ZWI4miXup9dJG2jymNxZpK8CZByuWuwrxyJPZ695LH ZBv8QXZ01zUB2U1vPS5j6
2024-07-22 10:13:26 MRJ5mTomqmjWMIxjRn13S5N3HGcKHMMW0Nr6K6ojjPGZhwDekuPGD6SX0WShk5U4
2024-07-22 10:13:27 mgsCnUEqeLWgIaXFBZlhkeyyW42hvs5bCD0E14B04eycIaBX6x8W7z HHUeajwI9
```

Concept: reading




```
log_read=$(dirname "${0}")/.$(basename "${my_log}").${skip}.read
current_size=$(wc -c < "${my_log}")

bytes_read=$(cat "${log_read}")
echo "${current_size}" > "${log_read}"

tail -c +${(bytes_read+1)} "${my_log}" | head -c ${current_size-bytes_read})
```

Concept: reading in bytes

Reading in lines is more understandable for human eye



```
[root@linux ~]$ wc -l /var/log/messages
6969 /var/log/messages
[root@linux ~]$ tail -1 /var/log/messages
Jul 22 09:51:02 linux systemd: Started Session 102386 of user root.
[root@linux ~]$
```

Concept: reading in bytes

But reading in lines is very slow when you have big files!

Reading in bytes is crucial for performance

Concept: reading in bytes



```
[root@linux ~]$ time wc -c /stuff/monitor/FACILITY/local0-log  
11738026507 /stuff/monitor/FACILITY/local0-log
```

```
real    0m0.002s  
user    0m0.000s  
sys     0m0.002s
```

```
[root@linux ~]$ time wc -l /stuff/monitor/FACILITY/local0-log  
25865054 /stuff/monitor/FACILITY/local0-log
```

```
real    0m37.671s  
user    0m1.027s  
sys     0m6.634s
```

```
[root@linux ~]$
```

Concept: log rotation

Don't forget log rotation – when log file becomes smaller than it was during last run, it means it's rotated



```
# if rotated, let's read from the beginning
```

```
[[ ${bytes_read} -gt ${current_size} ]] && bytes_read=0
```

Concept: read vs. skip

If you stop agent for a while, you have to choose what to do when it's started again: read everything or skip?

Similar to “maxdelay” setting for `log[]` / `log.count[]`

Concept: read vs. skip



Implementation

So, after putting those four ideas into one script, we have a working frame already

```
1 #!/bin/bash
2
3 my_log="${1}"
4 skip="${2}"
5 skip_time="${3}"
6
7 [[ "${my_log}" == "" || ! -f "${my_log}" ]] && exit 1
8
9 log_read=$(dirname "${0}"/.$(basename "${my_log}").${skip}.read
10
11 # get current file size in bytes
12
13 current_size=$(wc -c < "${my_log}")
14
15 # remember how many bytes you have now for next read
16 # when run for first time, you don't know the previous
17
18 [[ ! -f "${log_read}" ]] && echo "${current_size}" > "${log_read}"
19
20 # if not ran for a while, consider you want to skip to fresh data
21
22 if [[ "${skip}" == "skip" && $(( $(date +%s) - $(stat -c %Y "${log_read}") )) -gt "${skip_time}" ]]; then
23     echo "${current_size}" > "${log_read}"
24 fi
25
26 bytes_read=$(cat "${log_read}")
27 echo "${current_size}" > "${log_read}"
28
29 # if rotated, let's read from the beginning
30
31 [[ ${bytes_read} -gt ${current_size} ]] && bytes_read=0
32
33 # get the portion
34
35 tail -c +${(bytes_read+1)} "${my_log}" | head -c ${current_size-bytes_read}
36
37 exit 0
```

Implementation

It would just read data and output fresh lines without any further processing

What you do next with this ability to collect data depends on your needs and imagination


Time to see it in action!

Example no. 1

Limit of 600k lines per minute

Both `log[]` and `log.count[]` share same limitation of 600k lines to be analyzed per minute

Example no. 1



Example no. 1



```
# get the portion
```

```
tail -c +${(bytes_read+1)} "${my_log}" | head -c ${current_size-bytes_read} | wc -l
```

Example no. 1: configuration

```
UserParameter=log.count.custom[*],/etc/zabbix/zabbix_agentd.d/zbx_scripts/log_count.sh "$1" "$2" "$3"
```

* Name

Type

* Key

Type of information

Units

* Update interval

Custom intervals

Type	Interval	Period	Action
<input type="text" value="Flexible"/> <input type="text" value="Scheduling"/>	<input type="text" value="s58"/>		Remove

[Add](#)

Example no. 1: performance

```
[root@linux ~]$ for i in $(seq 1 6); do time ./log_count.sh /stuff/monitor/FACILITY/local0-log && sleep $((i*20)); done
0
real    0m0.015s
user    0m0.008s
sys     0m0.006s
184320

real    0m0.268s
user    0m0.040s
sys     0m0.183s
393819

real    0m0.455s
user    0m0.102s
sys     0m0.380s
602513

real    0m0.702s
user    0m0.161s
sys     0m0.657s
863529

real    0m0.875s
user    0m0.243s
sys     0m0.870s
1022463

real    0m0.900s
user    0m0.277s
sys     0m0.988s
[root@linux ~]$
```

Example no. 1: result



Example no. 2

Collecting something just for the sake of feeding it to calculated items later on

AND


having lots of that “something”

Example no. 2

Each line of log has duration component – say duration of request processing

You want to check, what's average duration per minute

Example no. 2



Example no. 2

But what to do, when you have 10k+ of such lines per minute?

Storing all of it just for the sake of calculating that one average value is of course inefficient

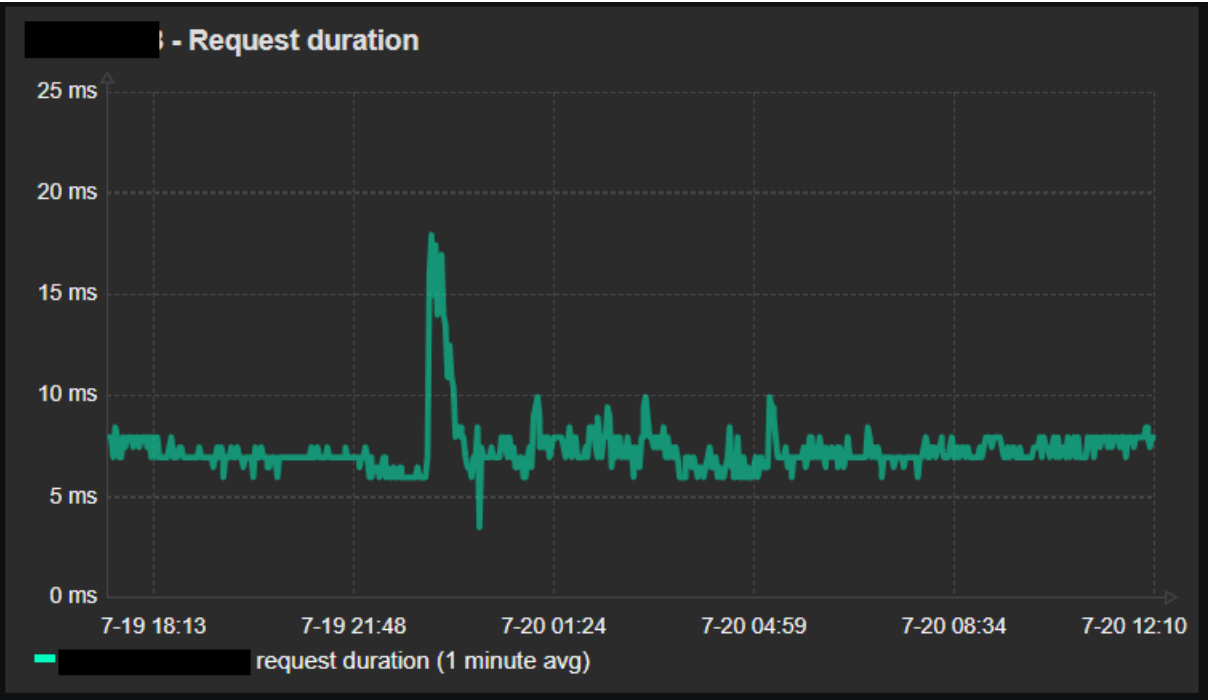
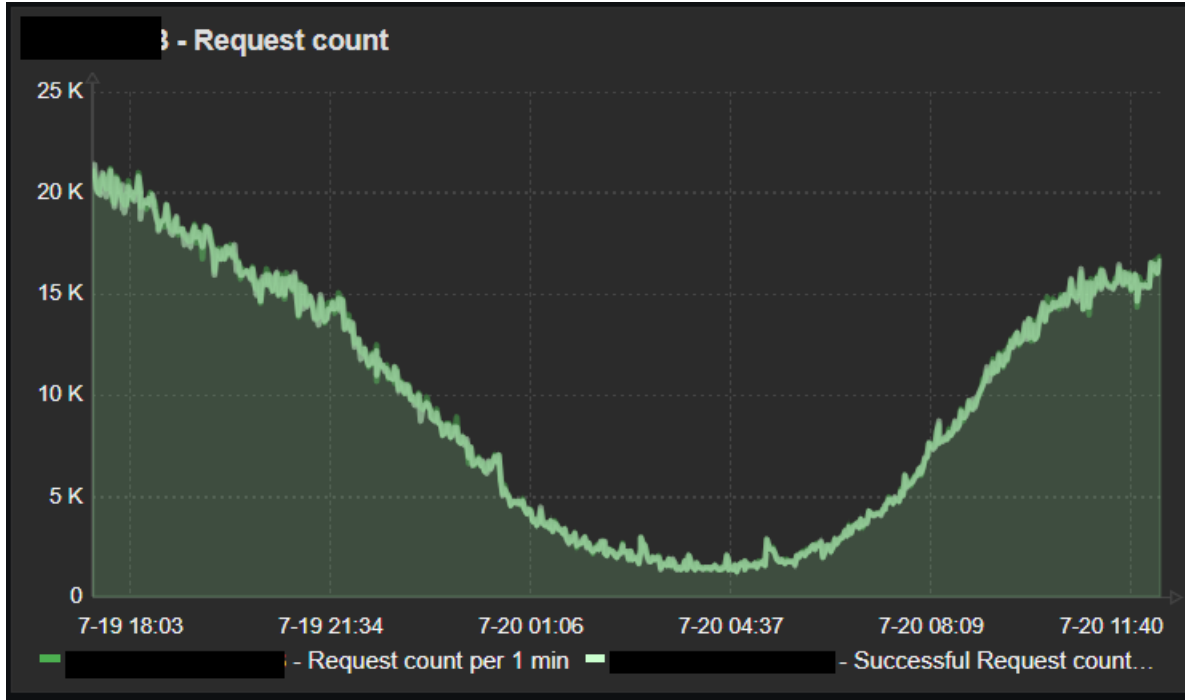
Displaying 10k+ dots is also not a good idea...

So all you need is just 1 average value per minute

Example no. 2

```
● ● ●  
  
# get the portion  
  
duration_sum=0  
count=0  
  
while read duration; do  
  
    [[ ${duration} == "" ]] && continue  
  
    duration_sum=$((duration_sum+duration))  
    count=$((count+1))  
  
done <<< "$(tail -c +${(bytes_read+1)} "${my_log}" | head -c ${current_size-bytes_read} | grep -P "${request_pattern}"  
| grep -Po "(?<=DURATION\\=)(\\d+)(?=\\")")"  
  
[[ ${count} -eq 0 ]] && duration_avg=0 || duration_avg=$((duration_sum/count))
```

Example no. 2: result



Example no. 3

Multi line analysis

Your interest to find something meaningful might sometimes lay in multiple lines

Example no. 3

Say you have a request with some unique identifier PAYMENTID=...
logged as one line and later some answer / response should be logged
for the same PAYMENTID

What if response is never logged and you want to know it?

Example no. 3: configuration

* Name

Type

* Key

Type of information

* Update interval

Custom intervals

Type	Interval	Period	Action
<input type="checkbox"/> Flexible <input checked="" type="checkbox"/> Scheduling	<input type="text" value="50s"/>	<input type="text" value="1-7,00:00-24:00"/>	Remove

[Add](#)

Example no. 3



```
#!/bin/bash
```

```
my_log="${1}"
```

```
skip="${2}"
```

```
skip_time="${3}"
```

```
request_pattern="${4}"
```

```
response_pattern="${5}"
```

```
identifier="${6}"
```

```
timeout="${7}"
```

```
...
```

Example no. 3



```
[[ ! -d $(dirname "${0}")/$(basename "${my_log}").tmp ]] && mkdir $(dirname "${0}")/$(basename "${my_log}").tmp  
tmp_dir=$(dirname "${0}")/$(basename "${my_log}").tmp
```

Example no. 3



```
while read line; do

    [[ "${line}" == "" ]] && continue

    uniqueid=$(grep -Po "${identifier}=\K(\d+)" <<< "${line}")
    echo "${uniqueid}" > "${tmp_dir}/${uniqueid}"

done <<< "$(tail -c +${(bytes_read+1)} "${my_log}" | head -c ${((current_size-bytes_read))} | grep -P "${request_pattern}")"
```

Example no. 3

```
while read line; do

    [[ "${line}" == "" ]] && continue

    uniqueid=$(grep -Po "${identifier}=\K(\d+)" <<< "${line}")
    rm -f "${tmp_dir}/${uniqueid}" 2>/dev/null

done <<< "$(tail -c +${((bytes_read+1))} "${my_log}" | head -c ${((current_size-bytes_read))} | grep -P "${response_pattern}")"
```

Example no. 3



```
while read line; do

    [[ "${line}" == "" ]] && continue

    if [[ $(( $(date +%s) - $(stat -c %Y "${line}")) ) -gt "${timeout}" ]]; then
        echo "No response for ${identifier}=$(cat ${line})"
        rm -f "${line}" 2>/dev/null
    fi

done <<< "$(ls -1 ${tmp_dir}/* 2>/dev/null)"
```


Example no. 3

```
[root@linux ~]$ cat /tmp/multiline.log
2024-07-23 09:34:42 request PAYMENTID=100001 yJUQr2m6JNFshgACHwo429TkogJSUffjPMLnaTdCJqIJSn8RIB5hlnAidDgdrUGT
2024-07-23 09:34:47 request PAYMENTID=100002 JzJvP1AQI4rDTLG6kPXg0iMA2b4xAdqjiLQj7dfwbIb8yL5GT9V1b9iqsx1dZr8Z
2024-07-23 09:34:51 request PAYMENTID=100003 rq5kTuycIgrszUrGrrzjUFkcDxMEN0ZHAh2lqlg0z0b6fzPLtRwccw7jwb9 q0FK
2024-07-23 09:34:56 request PAYMENTID=100004 7nhM8KRT60FbsLt3mzsWi1JuHNcywC7T1CHR6wiMsntFFzhwfy2Vkw7UAILT2YZ1
2024-07-23 09:35:00 request PAYMENTID=100005 gGalrcnNmvesCTB99piS0wwlsUBsDBeuXTVDwmp0fMrXlCBo vjleAmKM ur9cHc
2024-07-23 09:35:07 response PAYMENTID=100001 XzE05JKG7ELFtjfrsiw9pF8fhu9QcgvKoHHwLWYoaRtJ1r071vnwVpMAS08tODDw
2024-07-23 09:35:11 response PAYMENTID=100002 ic2DMYQhy5rpxvCElfz76ESPjKFxIluI5nNWhoY S4osQJ8nlwWUZ56pZKWaurA
2024-07-23 09:35:14 response PAYMENTID=100003 nD HyKJEMuqhKbZQs 3czVvmqBIZ8nrhSmz2KrGAvpdPUFyMxKYKMVRSdOXAmJwU
2024-07-23 09:35:18 response PAYMENTID=100004 SDj3psaElQ7DBejae3n2RNDsr Hdsik4NxAfYikM62yALP44flkrfDDBceIFgKdn
2024-07-23 09:35:23 request PAYMENTID=100006 FdWZsHrldQV8GKJMFizGjNDUtHrbESoudardJ gIMDcnnkj6DbKtubKKvzeU3Bzk
2024-07-23 09:35:26 request PAYMENTID=100007 kU4EpmJaAJYiLZLPJyn10SdhKD9QdifIYCl6bQH9Hg8 yJMBZ00ghlx9K4R4qknH
2024-07-23 09:35:30 request PAYMENTID=100008 If uAsc7fY5pdouKUV3juvCIV8Nwaw0oHwtl4i11dJJzPgEMMc7IfFMudQVCBU5
2024-07-23 09:35:40 response PAYMENTID=100006 UfXSCc6kbscX2RTAGwwfPX3vU0il016tFDhWThakindneSwrCQFiTEq4py09yfyj
2024-07-23 09:35:45 response PAYMENTID=100008 ITL0U9jff3hI6Ujyu3bvoHhXblbLe1b99mlf7jAXZMbGyYKYP8NKHc9dgxUhDCFM
[root@linux ~]$
```

Example no. 3

```
[root@linux ~]$ ls -ltr /etc/zabbix/zabbix_agentd.d/zbx_scripts/multiline.log.tmp/  
total 12  
-rw-rw-r--. 1 zabbix zabbix 7 Jul 23 09:35 100005  
-rw-rw-r--. 1 zabbix zabbix 7 Jul 23 09:35 100007  
[root@linux ~]$
```

Example no. 3: result

Timestamp	Value
2024-07-23 09:36:43	No response for PAYMENTID=100007
2024-07-23 09:36:03	No response for PAYMENTID=100005

Example no. 3: trigger

*** Name**

Event name

Operational data

Severity Not classified Info Warning Average High Disaster

*** Expression**

[Expression constructor](#)

OK event generation Expression Recovery expression None

PROBLEM event generation mode Single Multiple

Allow manual close

Example no. 3: result

Time ▼	<input type="checkbox"/>	Severity	Info	Host	Problem
09:36:43	<input type="checkbox"/>	Warning		████████████████████	<u>No response for payment 100007</u>
09:36:03	<input type="checkbox"/>	Warning		████████████████████	<u>No response for payment 100005</u>

Run as passive item

With this approach you can analyze logs with passive item type

Why is it important?

Run as passive item



Run as passive item

* Name

Type

* Key

Type of information

Units

* Update interval

Custom intervals

Type	Interval	Period	Action
<input type="button" value="Flexible"/> <input checked="" type="button" value="Scheduling"/>	<input type="text" value="s58"/>	<input type="text"/>	Remove

[Add](#)

Run as passive item

Another use case – active-passive clusters

Add virtual hostname to point to current active node and run such item as passive

This is how you won't have unsupported / useless item (trigger) on passive and no manual work needed to control this

<https://github.com/b1nary1/zabbix>



Thanks!