

# Performance Analysis and Capacity Planing

collectd – the system statistics collection daemon

Sebastian 'tokkee' Harl  
<sh@teamix.net>

teamix GmbH / collectd core team

Libre Software Meeting 2012  
July 10, 2012



## Solid IT-Infrastructure

Location: Nuremberg, Munich, Frankfurt

Open-Source

Monitoring

Network

N-IX

NetApp

Juniper

Riverbed

VMWare

Trainings



collectd Overview

Overview


Main Features

Feature Overview


Extending collectd

Extras


- **collectd** collects performance data of systems
- some (simple) examples:
  - CPU utilization
  - memory utilization
  - network traffic
- **collectd** collects and stores the performance data
- stored data is usually used to generate graphs
- → performance analysis, capacity planing
- not to be confused with *monitoring!*
- Homepage: <http://collectd.org/>




- daemon
- free software (mostly GPL)
- portable (Linux, \*BSD, Solaris, ...)
- scalable (OpenWrt, ..., Cluster / Cloud)
- sophisticated network support
- efficient (default resolution: 10 seconds)
- flexible architecture
- modular (more than 100 plugins in Version 5.1)




- daemon
- free software (mostly GPL)
- portable (Linux, \*BSD, Solaris, ...)
- scalable (OpenWrt, ..., Cluster / Cloud)
- sophisticated network support
- **efficient** (default resolution: 10 seconds)
- flexible architecture
- modular (more than 100 plugins in Version 5.1)



# 10 seconds resolution



- daemon
- free software (mostly GPL)
- portable (Linux, \*BSD, Solaris, ...)
- scalable (OpenWrt, ..., Cluster / Cloud)
- sophisticated network support
- efficient (default resolution: 10 seconds)
- flexible architecture
- **modular** (more than 100 plugins in version 5.1)





## available plugins (version 5.1)

amqp	apache	apcups	apple_sensors	ascent
battery	bind	conntrack	contextswitch	cpu
cpufreq	csv	curl	curl_json	curl_xml
dbi	df	disk	dns	email
entropy	ethstat	exec	filecount	fscache
GenericJMX.java	gmond	hddtemp	interface	ipmi
iptables	ipvs	irq	java	libvirt
load	logfile	lpar	madwifi	match_empty_counter
match_hashed	match_regex	match_timediff	match_value	mbmon
md	memcached	memcached	memory	modbus
Monitorus.pm	multimeter	mysql	netapp	netlink
network	nfs	nginx	notify_desktop	notify_email
ntpd	numa	nut	olsrd	onewire
openvpn	OpenVZ.pm	oracle	perl	pinba
ping	postgresql	powerdns	processes	protocols
python	redis	routers	rrdcached	rrdtool
sensors	serial	snmp	swap	syslog
table	tail	tape	target_notification	target_replace
target_scale	target_set	target_v5upgrade	tcpconns	teamspeak2
ted	thermal	threshold	tokyotyrant	unixsock
uptime	users	uuid	varnish	vmem
vserver	wireless	write_graphite	write_http	write_mongodb
write_redis	xmms	zfs_arc		

- daemon collects data locally  $\Rightarrow$  runs on every client system (exceptions: SNMP, databases, etc.)
- one or more central servers
- clients push their data to the central servers
- first steps: install; select plugins; start daemon; enjoy ;-)

## C4


collection 4

- All instances
- All graphs
- Host "[redacted]"

## Graph "CPU utilization"

Instance "[redacted]"/0"

Instance: "[redacted] / cpu - 0 / cpu - all"

 Search

Hour ▾

JSON (gRaphaël)

RRDtool

Go

[collection 4.0.0](#)

- provide collected data through JSON
- different frontends possible
- efficiently handles large amounts of data
- flexible configuration of graphs

graphite

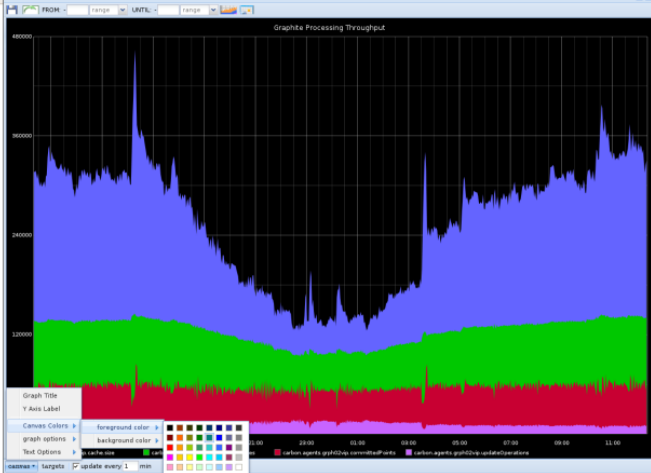
[Command Line Interface](#)  
Logged in as odars, [logout](#) ([edit profile](#))  
[Documentation](#)

[production](#)  
[pre-production](#)

Tree Search Auto Complete

- Graphite
  - PROD
  - Perf
  - afs
  - aws
  - carbon
    - agents
      - graph01ng
        - \*
          - avgUpdateTime
          - cache
          - committedPoints
          - creates
          - errors
          - pointsPerUpdate
          - updateOperations
        - graph02ng
          - \*
            - avgUpdateTime
            - cache
            - \*
              - @queries
              - @users
              - size
              - committedPoints
              - creates
              - errors
              - pointsPerUpdate
              - updateOperations
  - yyep@Utilisation
  - yyep
  - database
  - bestimg
  - ife
  - ods
  - My Graphs
    - ag/Test
    - test123
  - User Graphs

Graphite Processing Throughput



## collectd Overview

### Feature Overview

CPU, memory, network I/O

Networking Support

RRDtool Support

Generic Plugins (Overview)

### Extending collectd

### Extras

- specialized read plugins
  - CPU, memory, network interfaces, ...
- IO plugins
  - network plugin
  - RRDtool, RRDCheckedD
  - Graphite
  - MongoDB, Redis
  - AMQP
- generic plugins
  - SNMP
  - tail
  - PostgreSQL
- filter-chains

## configuration synopsis


```
LoadPlugin "cpu"  
LoadPlugin "memory"  
LoadPlugin "interface"
```


## configuration synopsis


```
LoadPlugin "cpu"  
LoadPlugin "memory"  
LoadPlugin "interface"
```

```
<Plugin interface>  
  Interface lo  
  Interface sit0  
  IgnoreSelected true  
</Plugin>
```









## modes of operation

- send data (*“client”*)
- receive data (*“server”*)
- forward data (*“proxy”*)
- Unicast (*“point-to-point”*)
- Multicast (*“point-to-group”*)
- IPv4 and IPv6

## rule them all

Modes may be mixed arbitrarily.

## synopsis: client

```
LoadPlugin "network"
```

```
<Plugin "network">
```

```
  Server "collectd0.example.com"
```

```
  Server "collectd1.example.com"
```

```
  Server "ff18::efc0:4a42"
```

```
</Plugin>
```

### synopsis: server

```
LoadPlugin "network"
```

```
<Plugin "network">
```

```
  Listen "collectd0.example.com"
```

```
  Listen "ff18::efc0:4a42"
```

```
</Plugin>
```

### synopsis: proxy

```
LoadPlugin "network"
```

```
<Plugin "network">
```

```
  Listen "collectgw.extern.example.com"
```

```
  Server "collectd1.intern.example.com"
```

```
  Forward true
```

```
</Plugin>
```

- writes data to RRD files **efficiently** → caching
- functionality now also available in RRDtool as stand-alone RRD Caching Daemon (RRDCacheD)

### synopsis

```
LoadPlugin "rrdtool"
```

```
<Plugin "rrdtool">  
  DataDir "/var/lib/collectd/rrd"  
</Plugin>
```




### configuration synopsis


```
<Plugin "rrdtool">  
  DataDir "/var/lib/collectd/rrd"  
  
  CacheTimeout 3600 # 1 hour  
  CacheFlush 86400 # 1 day  
  
  WritesPerSecond 30  
</Plugin>
```

- FLUSH command allows for graphing of current values

- idea: generic approaches rather than specialized solutions
- → user configuration determines behavior
- ⇒ new equipment does not require a new version of **collectd**
- examples: SNMP, tail, curl, DBI, PostgreSQL



## tail plugin: SSH brute-force attack



collectd Overview

Feature Overview

Extending collectd

- UNIXSOCK interface

- Custom Extensions

Extras

- **collectd** API: C, Perl, Python, Java
- external programs may use unixsock or exec plugins

### UNIXSOCK plugin

- opens a UNIX socket
- text and line based protocol  
(e. g. PUTVAL, FLUSH, LISTVAL, GETVAL)
- query and submit values
- collectd-nagios
- collectdctl (since version 5.0)
- cussd.pl: *“collectd UNIX socket shell”*

```
-> | GETVAL "FQDN/load/load"  
<- | 3 Values found  
<- | shortterm=4.000000e-02  
<- | midterm=6.000000e-02  
<- | longterm=7.000000e-02  
  
-> | PUTVAL FQDN/users/users 1341851406:42  
<- | 0 Success: 1 value has been dispatched.
```



## overview

- integrates a Perl interpreter (similar to Apache's `mod_perl`)
- compiles Perl code only once
- exports internal API (→ full flexibility available)

## also available

- Java and Python


## overview

- executes arbitrary programs / scripts
- parses STDOUT of the program
- respawns process upon termination  
(cf. `init`)

```
#!/bin/bash
INTVL=${COLLECTD_INTERVAL:-10}
CHOST="${COLLECTD_HOSTNAME:-localhost}"
IDENT="$CHOST/magic/magic_level"
while sleep $INTVL
do
    V='magic --level'
    echo "PUTVAL \"$IDENT\" interval=$INTVL 'date +%s':$V"
done
```


- aggregation of arbitrary values
- per-plugin interval settings
- Nagios support (pushing information to Nagios)
- store data in RMDBSs
- ...

<http://collectd.org/wiki/index.php/Roadmap>



Thank you for your attention!

Questions?



Contact:

Sebastian “tokkee” Harl  
teamix GmbH, Nuremberg  
<sh@teamix.net>

<collectd@verplant.org> — irc.freenode.net/#collectd — <http://identi.ca/collectd>  
<http://tokkee.org/events.html>

collectd Overview

Feature Overview

Extending collectd

Extras

- SNMP plugin

- RRDCacheD Plugin

- collectd Internals

- Perl Plugin

## overview

- queries network equipment via SNMP
- *generic*: not specialized for any particular hardware
- queries multiple systems in parallel

## configuration

- “Data” Block
- “Host” block



## synopsis: data block


```
<Plugin "snmp">
  <Data "ifmib_if_octets64">
    Type "if_octets"
    Table true
    Instance "IF-MIB::ifName"
    Values "IF-MIB::ifHCInOctets" \
          "IF-MIB::ifHCOutOctets"
  </Data>
</Plugin>
```

## synopsis: host block


```
<Plugin "snmp">  
  <Host "switch0.intern.musterfirma.de">  
    Address "10.0.42.2"  
    Version 1  
    Community "public"  
    Collect "ifmib_if_octets64"  
    Interval 60  
  </Host>  
</Plugin>
```

## overview


- generalizes caching idea of the rrdtool plugin
- standalone daemon
- integrated into RRDtool (available since version 1.4)
- extended feature set, e. g. journal
- benefit: restart **collectd** without losing cache




graphic © Florian "octo" Forster



graphic © Florian "octo" Forster



graphic © Florian "octo" Forster



graphic © Florian "octo" Forster

```
core: read_thread()  
  
p = get_next_plugin();  
sleep_until_due(p);
```






```
core: read_thread()
```

```
p->callback();
```



```
plugin1.c:read()
```





```
plugin1.c:read()
```

```
...
```

```
dispatch_values(v);
```

```
core: dispatch_values()
```

```
...  
plugin_write(v);
```



```
plugin1.c:read()
```

```
...  
dispatch_values(v);
```



```
core: dispatch_values()
```

```
...
```

```
plugin_write(v);
```



```
plugin1_thread()
```

```
plugin2.c:write()
```

```
...
```

```
package Collectd::Plugin::Magic;
use Collectd qw( :all );
sub magic_read
{
    my $v1 = { plugin => 'magic',
               values => [Magic->getCurrentLevel ()] };
    plugin_dispatch_values ('magic_level', $v1);
}
plugin_register (TYPE_READ, 'magic', 'magic_read');
```