

# SysDB – System DataBase

A system management and inventory collection service

Sebastian 'tokkee' Harl

<tokkee@sysdb.io>

Icinga Camp San Francisco 2014

September 25, 2014





## **WARNING:**

SysDB is still under heavy development  
and not considered stable yet<sup>1</sup>.

Flaming, bashing or other forms of constructive  
feedback are very appreciated ... also, contributions :-)

---


<sup>1</sup> That is, interfaces may still change quickly.





## Background / Motivation













- Database Server Foo
- External Gateway
- Zooperserver





- Database Server Foo
- External Gateway
- Zooperserver


gw1.domain.tld








- Database Server Foo
- External Gateway
- Zooperserver




gw1.domain.tld



gw1.domain.tld

arch=amd64  
lsbdistid=Debian  
loc=BayArea





● Database Server Foo

● External Gateway

● Zooperserver

gw1.domain.tld



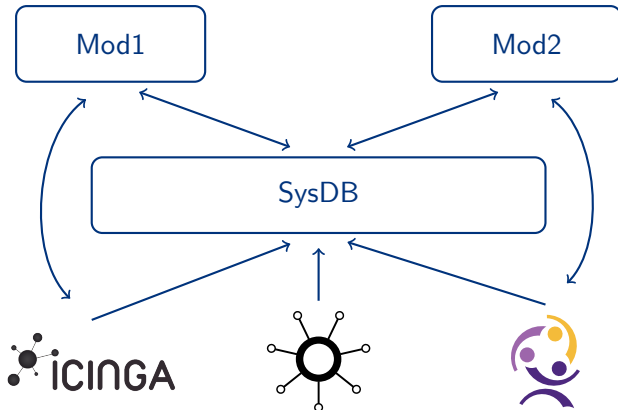
gw1.domain.tld

```
arch=amd64  
lsbdistid=Debian  
loc=BayArea
```



# The System DataBase







- SysDB collects information about arbitrary hardware and software systems for a global overview of the entire infrastructure.
- Focus on meta-data and links to more specific information.
- Simple examples:
  - Hosts and their attributes (“facts”)
  - Services and their attributes
  - Metrics (pointer to performance data)
  - Monitoring information (e.g. current state)
- SysDB collects these information and merges objects from multiple back-ends by hostname.





- <https://sysdb.io/>, <https://github.com/sysdb>
  - CI: <https://travis-ci.org/sysdb/sysdb>
  - ~ 80% code (function) unit-test coverage in the core
- BSD license
- Written in C (UI written in Go WIP)
- Generic store aggregating system information
- Multi-threaded, event-based client/server architecture
- Easy to extend (simple plugin API)
- Simple network protocol and query language
- Most of the code implemented as a library (reusable)





Currently available collectors (back-ends):

- **collectd::unixsock** – query collectd's UNIX socket interface  
→ host metrics
- **mk-livestatus** – query Monitoring systems (Nagios, Naemon, Icinga, Shinken) using Check\_MK Livestatus  
→ hosts and services (metrics planned)
- **puppet::store-configs** – query Puppet  
→ host attributes

Planned: Passive data collection (e.g. using Gearman), Foreman, PuppetDB, \$your\_favorite\_system (send patches!) :-)





Time-series fetchers:

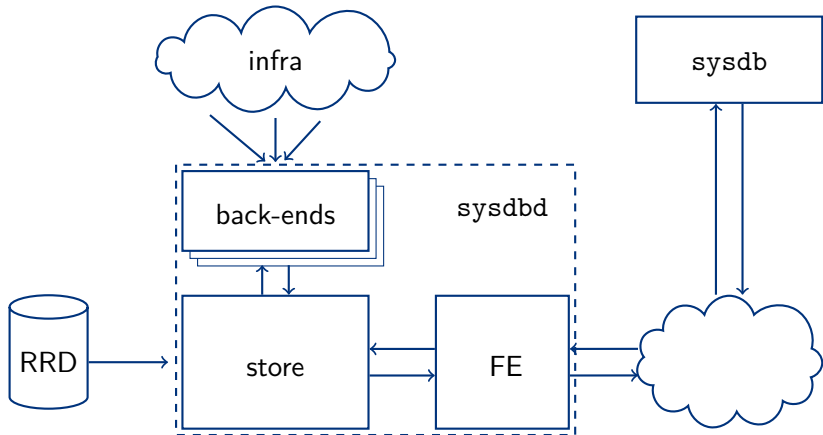
- `timeseries::rrdtool` – fetch time-series from local RRD files

Misc plugins:

- `cname::dns` – canonicalize host-names using DNS
- `syslog` – syslog logging









- Stores generic objects (in memory atm)
- Canonicalization of hosts (by their name as provided by back-ends)
- Each object stores the time-stamp of the last update and the (automatically calculated) update interval
- Interface to query data
- Generic access to time-series data
- JSON is the external data representation

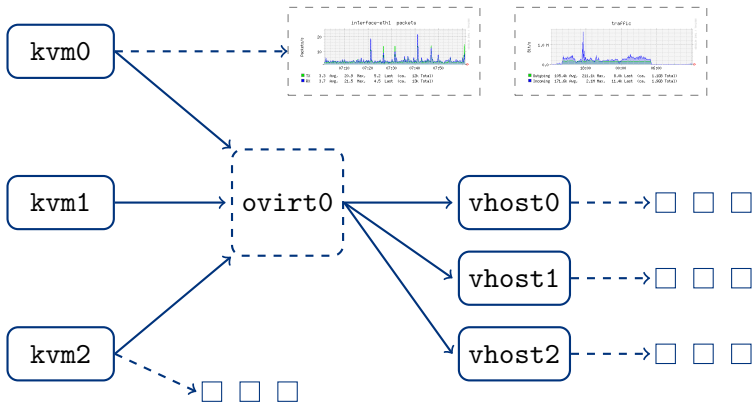




- Host – any kind of physical resource
- Service – any kind of service
- Metric – identifier for metrics data
  - time-series data is **not** stored in SysDB
- Attribute – attributes of hosts and services
  - values support multiple types
- Common core attributes:
  - unique name
  - time-stamp of last update
  - update interval
  - list of back-ends providing the object



# The SysDB Store – Example



(Full graph-support not implemented yet.)





- Interactive client program for SysDB
- Connects to a SysDB daemon
- Interactive command shell
- Receives and displays asynchronous log messages





- Remotely similar to SQL
- But it's not SQL<sup>1</sup>
- List and query hosts
- Fetch time-series information (from some back-end data-store)

---

<sup>1</sup> Yes, a DSL might limit options but it was the easiest for now – believe it or not ;-)





- LIST hosts
- FETCH host <hostname>
- LOOKUP hosts MATCHING <condition>
- ... FILTER <condition>
- TIMESERIES <host>.<metric>  
START 2014-01-01 END 2014-12-31

→ <https://sysdb.io/manpages/head/sysdbql.7.html>





```
sysdb=> LIST hosts FILTER :age < 2 * :interval;
[ {
  "name": "monitor.lxc.tokkee.net",
  "last_update": "2014-04-03 10:26:41 +0200",
  "update_interval": "5m4s",
}, {
  "name": "puppet.lxc.tokkee.net",
  "last_update": "2014-04-05 11:04:08 +0200",
  "update_interval": "5m2s"
}]
```





## The SysDB Query Language – Example (2)



```
sysdb=> LOOKUP hosts WHERE attribute.architecture = 'amd64'  
                                AND service.name =~ 'postgres';  
{ "name": "db.lxc.tokkee.net",  
  "last_update": "...", "update_interval": "10s",  
  "attributes": [{  
    "name": "architecture", "value": "amd64",  
    "last_update": "...", "update_interval": "5m3s"  
  },{ ... }],  
  "services": [{  
    "name": "PostgreSQL Master",  
    "last_update": "...", "update_interval": "5m"  
  },{ ... }],  
  "metrics": [{...}]  
}
```





Search and filter conditions:

```
MATCHING attribute.architecture = 'amd64'  
        AND service.name =~ 'postgres'
```

```
MATCHING host =~ '\.tokkee\.net$'
```

```
FILTER :age >= 3 * :interval
```

```
FILTER :backend != 'mk-livestatus'
```



## Use Cases





- More flexible web front-ends combining multiple back-ends
  - correlate and annotate information
  - central dashboard
  - link to specialized front-ends
- Compare the back-ends (monitoring)
  - Which hosts / services are missing in a back-end?
  - Base for business-processes: What is the global status of all Windows systems in some data-center?
- Extend your CMDB with dynamic annotations or cross-checks



## Future Directions





- Better and more integration with other systems
- Persistent store: RDBMS, graph database (?)
- Interface to query other live data (e.g., monitoring status) from back-ends
- Distributed architecture (HA and load-balancing)
- Extensible Web-Interface
- Extend the type system and filters
- ...





Thank you for your attention!  
Questions, comments, rants?



**<https://sysdb.io/>, <https://github.com/sysdb>**  
🌐 <https://sysdb.io/> + 📘 <https://sysdb.io/facebook>  
🐦 @SystemDatabase

