

This is the web version of the talk.

Pages that are empty had content we had to remove in order to publish the slides.

We removed some screenshots as well - if you want to see the idb in action, drop us a line.

Landscape photo by Felix Kronlage

Small Kids motives by:

<http://jujus-delivery.com>





RELIABLE  
INFRASTRUCTURES

# Dr. Restless

or

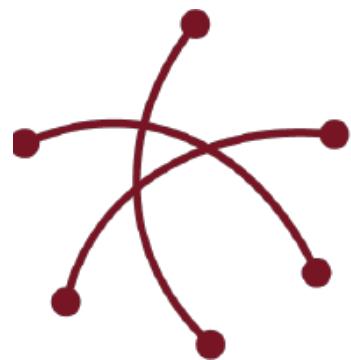
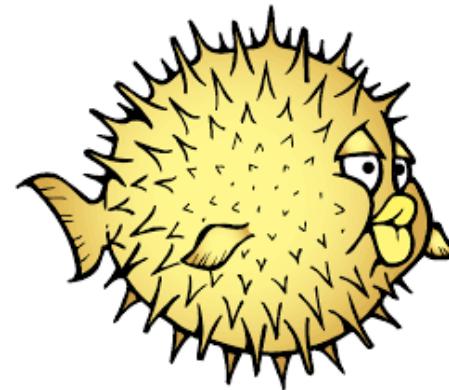
how I stopped worrying about not having proper infrastructure docs at hand!

cfgmgmtcamp 2017 - Gent

felix kronlage <kronlage@bytemine.net>

@felixkronlage

# /whois fkr



**OSB** Open Source  
Business  
**ALLIANCE**

# idb

- infrastructure database
- (CMDB)
- documentation
- hardware management
- universal source of truth



# why

- Documenting is annoying
- Information is often
  - not to be found
  - outdated
  - worst case: simply wrong and misleading
- We want to be techies - *not* bureaucrats
  - yet we still want to play the compliance game
  - and we want to deliver high quality, long-lasting infrastructure work



# how?

- Wiki
- Excel
- odt or ods documents
  - shared across the network
- Inline within other tools
  - Dashboards \* N
- Brain of exactly that guy that is on vacation





**to the rescue**

# how did we start?

- PuppetDB
- Parser for icinga1 configs
- mcollective virsh worker
- Lexware importer



**lift off!**

# Modules



RELIABLE  
INFRASTRUCTURES

# Modules

# Modules

# Adapter



RELIABLE  
INFRASTRUCTURES

# Importer



RELIABLE  
INFRASTRUCTURES

# a matter of perspective

- physical view
- logical view
- informations are mapped together



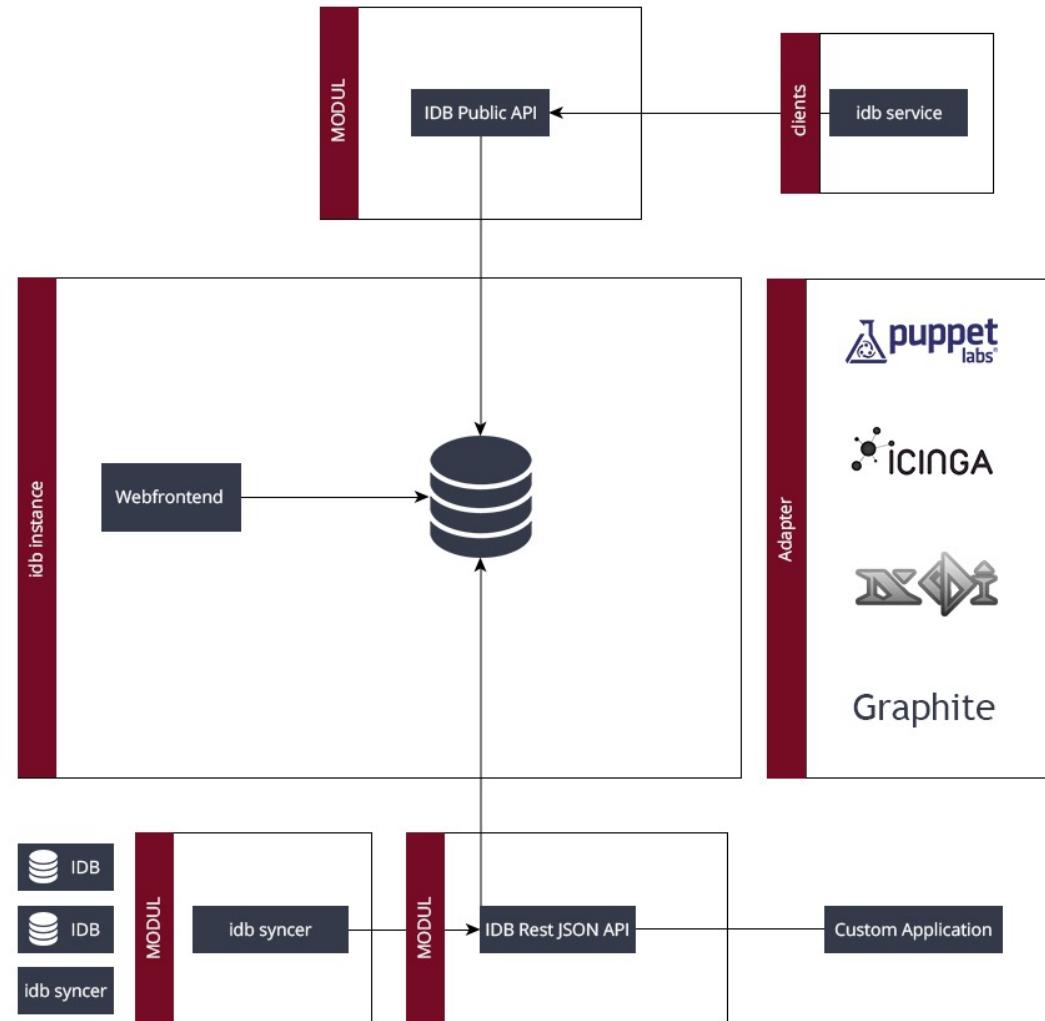
RELIABLE  
INFRASTRUCTURES

# idb-core

- Ruby on Rails web application
  - Rails 5
- SQL Backend
- directory service for users
  - ActiveDirectory
  - LDAP
- *privacyIdea* integration
- REST JSON API

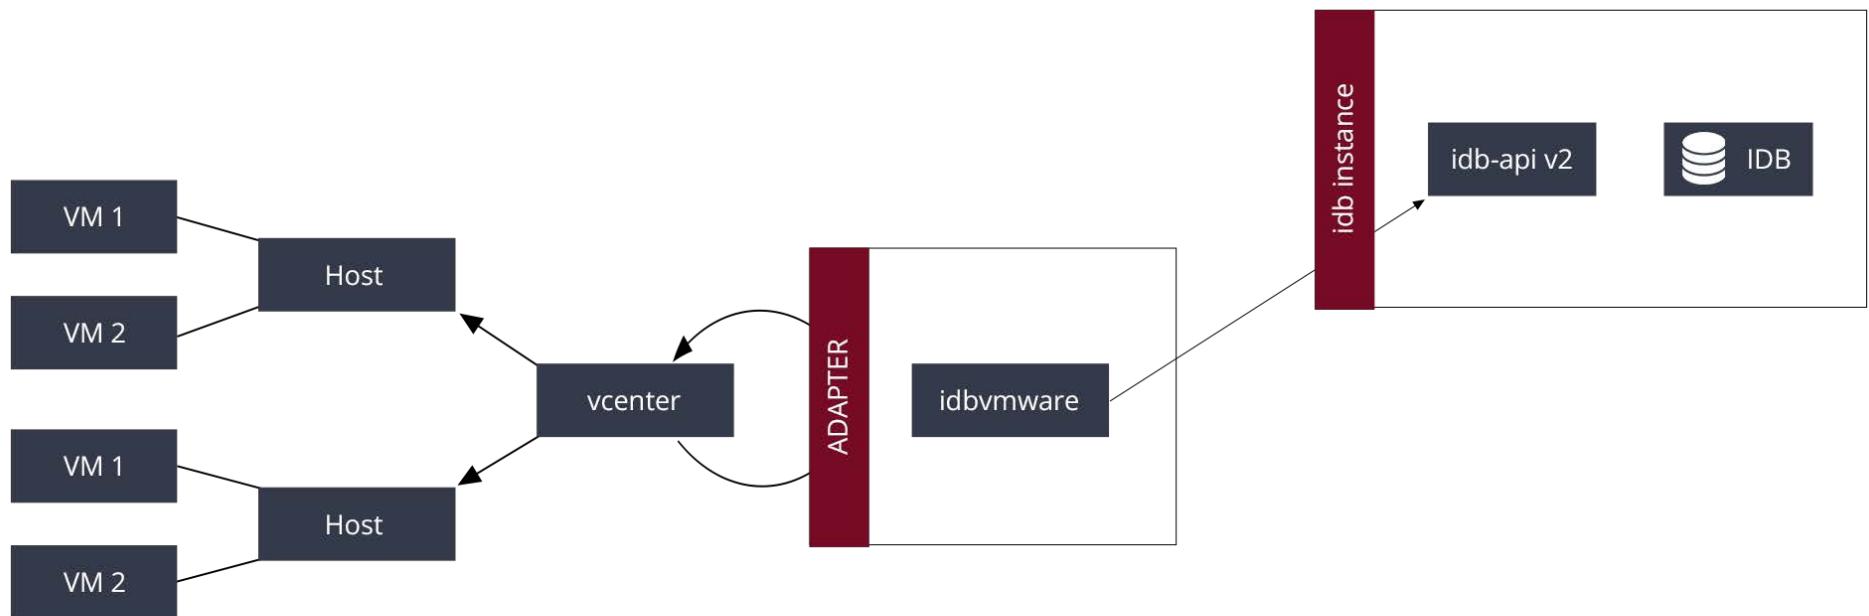


# Architecture



**all your APIs are belong to us**

# Architektur



# Procedure

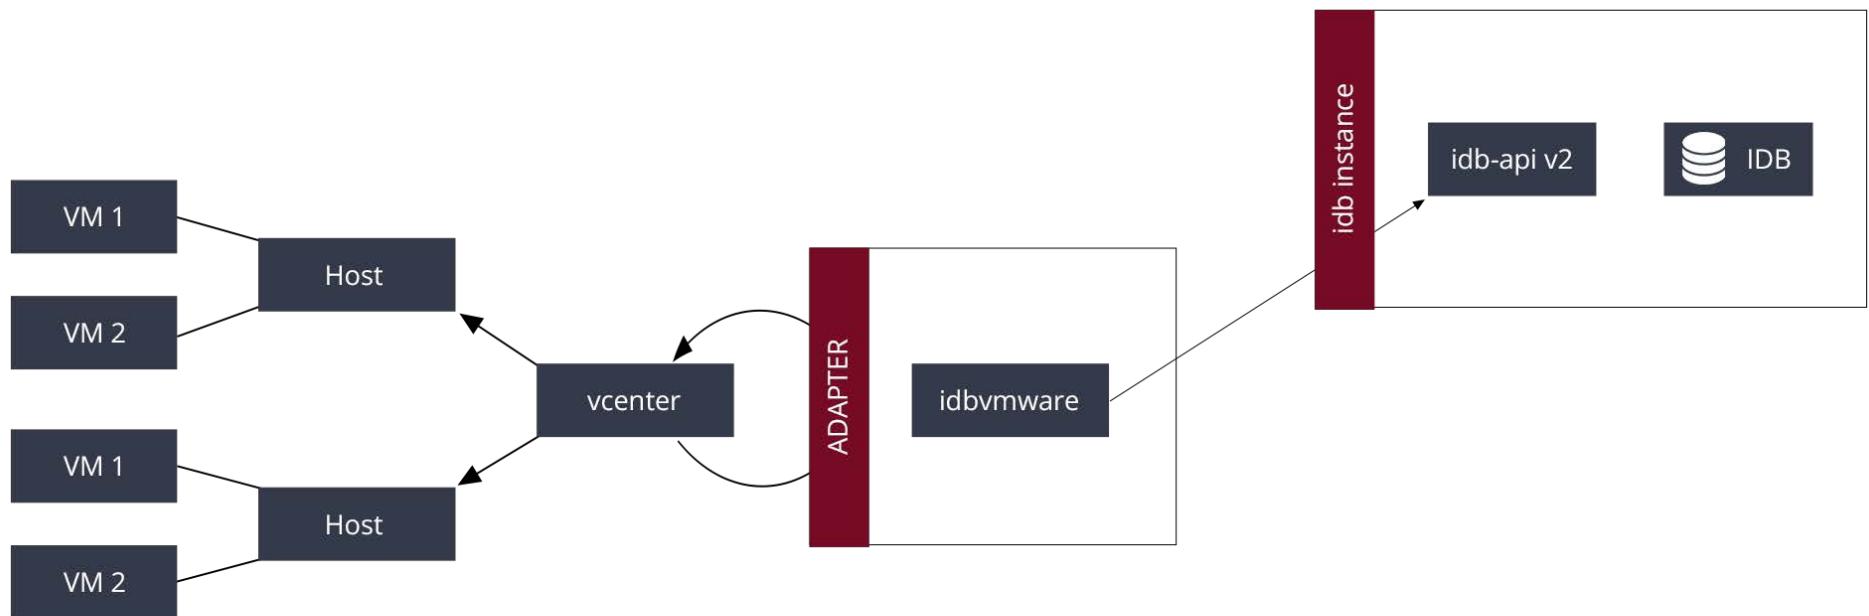
- Adapter crawls source of information
- Adapter submits json filled with information to idb rest api
  - Adapter requests certain actions
    - eg. create machine if non-existent
- idb decides what to do
- idb has the final say on what is done with the information



```
{  
  "fqdn": "oliver.hazardous.org",  
  "cores": 6,  
  "nics": [  
    {  
      "ip_address": {  
        "addr": "199.10.13.15",  
        "netmask": "255.255.255.254"  
      },  
      "name": "eth0"  
    }  
  ],  
  "aliases": [  
    {  
      "name": "girotti.hazardous.org",  
      "remove": true  
    }  
  ],  
  "create_machine": "true"  
}
```

```
curl -X "PUT" "https://idb-dev.office.bytetime.net/
api/v2/machines" \
-H "X-IDB-API-Token: WeuY7geish6uota" \
-H "Content-Type: application/json;
charset=utf-8" \
-d "{\"fqdn\":\"oliver.hazardous.org\",\"cores\":6,\"nics\":[{\"ip_address\":{\"addr\":\"199.10.13.15\",
\"netmask\":\"255.255.255.254\"},\"name\": \"eth0\"}],\"aliases\":[{\"name\": \"girotti.hazardous.org\"},\"remove\":true}],
\"create_machine\":true\"}"
```

# Architektur



# Cloud Adapter

- Python
- based on libcloud
- fast prototyping possible
- low hanging fruits
  - AWS
  - Azure
  - Digital Ocean
  - libvirt
  - Profitbricks

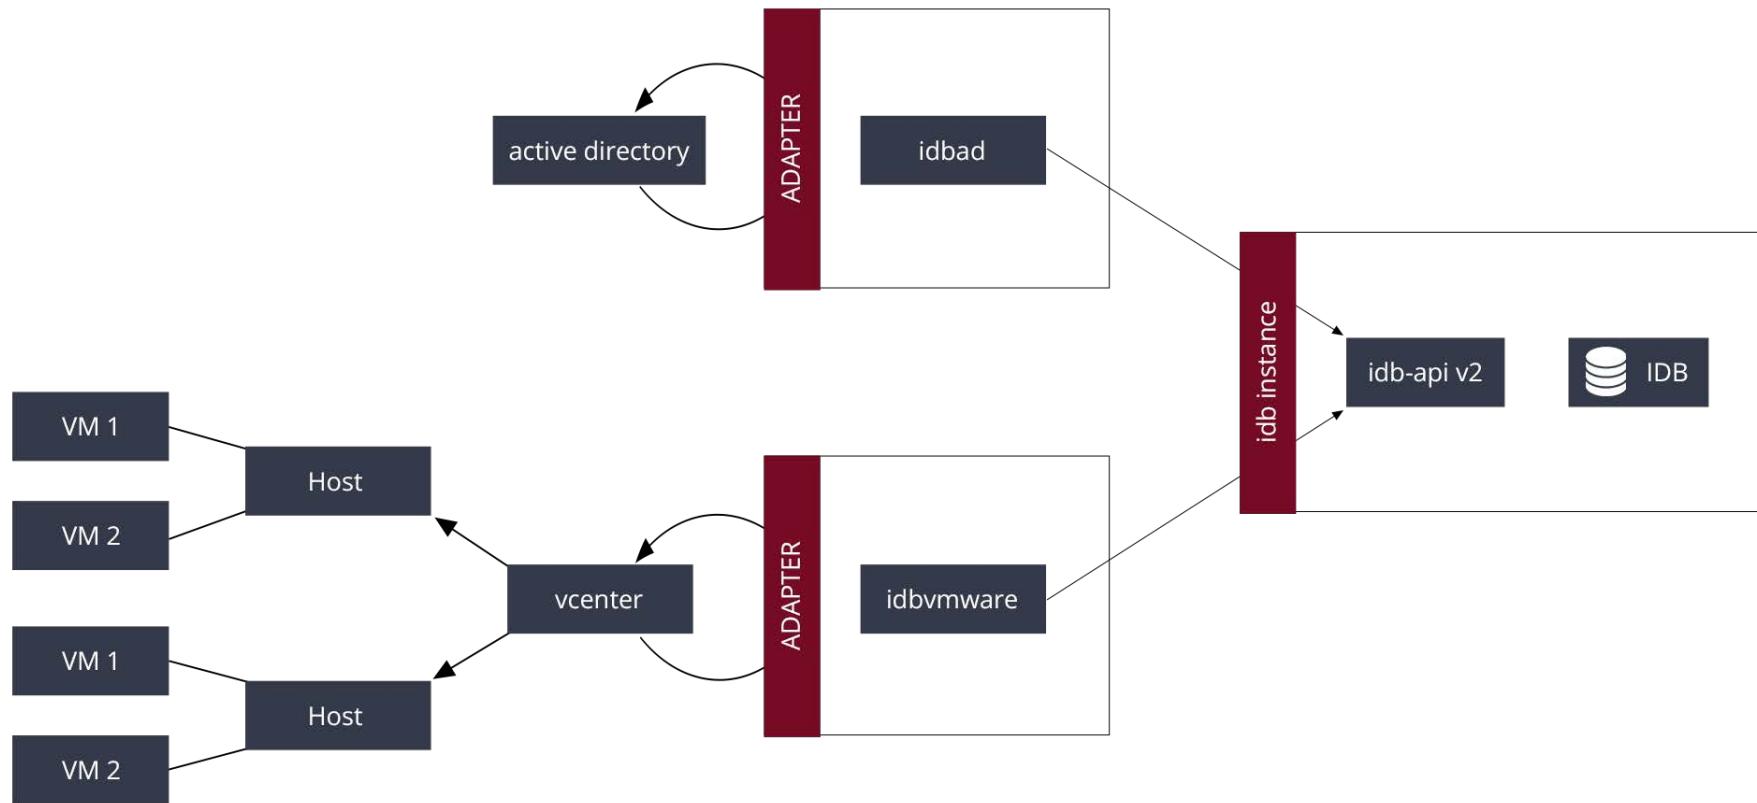


# idbvmware

- go (golang)
- crawls the rest api of vcenter/vsphere
- crawls the environment in order to
  - find new resources
  - find changed resources



# Architektur



# idbad

- PowerShell
- crawls the ActiveDirectory
  - based on a filter looks for ComputerObjects
  - submits these to the idb
- Assumption: AD is clean and up to date ;)



# idbvamp

- go
- Adapter for bacula community
- crawls the database
  - only bacula enterprise has a proper api
  - backup of which type was run last
- Answers:
  - how much space does a machine take up
  - how much space do all machines of a client take up



# idbvamp

IDB Machines Maintenance Switches Networks Owners Apps Inventory Location Felix Kronlage

audrey.hazardous.org Host Owner Network Backup Advanced Unattended Pending updates Maintenance History Graphs Edit Delete

**Backup type: Bacula**

Last full run:	Last incremental run:	Last differential run:
2016-09-04 02:58:32	2016-09-06 01:35:27	
Last full size:	Last incremental size:	Last differential size:
103 GB	620 MB	104 GB
Total backup size:		

# idbvamp

The screenshot shows a web-based application interface for managing inventory. At the top, there is a navigation bar with links: IDB, Machines, Maintenance, Switches, Networks, Owners, Apps, Inventory, and Location. On the far right of the top bar, it says "Felix Kronlage". Below the navigation bar, there is a search bar with the placeholder "Owners" and buttons for "+ New Owner", "CSV Import", and "Deleted". The main content area is a table titled "Owners". The table has columns: ID, Name, Nickname, Customer ID, Wiki, Machines, and Total backup size. There are two rows of data. The first row is a header row with empty cells. The second row contains data: ID 12, Name "Kronlage, Felix", Nickname "fkr", Customer ID "21127", Wiki (empty), Machines "1", and Total backup size "104 GB". The "Name" column is currently selected, as indicated by a blue border around the cell containing "Kronlage, Felix".

ID	Name	Nickname	Customer ID	Wiki	Machines	Total backup size
12	Kronlage, Felix	fkr	21127		1	104 GB

# idb-puppetintegration

- custom facts
- */usr/lib/update-notifier/apt-check* etc.
- */var/run/reboot-required*
- based on these, created facts for puppet
  - *idb\_pending\_updates*
  - *idb\_pending\_security\_updates*
  - *idb\_pending\_updates\_package\_names*
  - ....



# **idb-puppetintegration**

# **idb-puppetintegration**

# idb-puppetintegration

# Time to demo



# Eco System

- Creating an eco system
  - idbclient
    - ‘Framework’ in go
    - <https://github.com/idb-project/idbclient>
  - idb-puppetintegration
  - Make it nice and appealing to create more
  - is it worth it?



RELIABLE  
INFRASTRUCTURES

# Implementing it on your own

- what are your information sources?
  - do you already have some sort of central machine registry?
  - if not, use this as a starting point
    - configuration management!
- Include your co-workers
- First results are the best arguments
- ...and there is never (always) the perfect moment



# Sources

- <https://github.com/idb-project/>
- Ubuntu 14.04 or 16.04
- Ruby 2.2 (via rvm...)
- Apache + mod-passenger
- MySQL
- OpenLDAP
- Redis



# Packages

- Ubuntu 14.04 and 16.04
- CentOS / RHEL 6 und 7
- idb-core and all current adapters available through us
- Packages from us come with support



# Roadmap

- 1.6.3 - released end of 2016
  - AGPLv3
- Throughout 1.6.x
  - Software <-> Maschine-Mapping
  - Cloud-Provider finalisation
  - Visibility through which adapter which information came



# Roadmap

- 1.7.0
  - Integration with UCS
  - Nedi
  - Oxidized
  - Suse SMT
  - icinga2 api adapter
- 1.8.0
  - i18n
  - Multi-Tenacy



# Roadmap

- Configuration and administration via BackOffice Module
- Adapter via plugin architecture
- And a whole new perspective on things



RELIABLE  
INFRASTRUCTURES



# The infamous (almost) last page

- Thanks to the organizers!
- ...and thanks to all the helpers!



RELIABLE  
INFRASTRUCTURES

# and thank YOU for listening

<https://blog.bytemine.net/>

<https://github.com/idb-project/>

kronlage@bytemine.net

@felixkronlage



RELIABLE  
INFRASTRUCTURES

**oh...and STICKERS!**



*RELIABLE  
INFRASTRUCTURES*