

13 JAHRE PUPPET, VON 2 ZU 350 SERVERN

13 YEARS PUPPET, FROM 2 TO 350 SERVERS

KIM KLOTZ

- Principal Engineer Operations
- Since 2011 at makandra
- Contact: kim.klotz@makandra.de

STEFAN XENOPOL

- DevOps Engineer
- Since 2022 at makandra
- Contact: stefan.xenopol@makandra.de

ABOUT MAKANDRA

We plan, develop, operate and design individual web applications. Reliable and of outstanding quality.

- Established in 2009
- Located in southern germany (Augsburg)

MAKANDRA - SERVICES

- Conceptual design
- Web development
- UI- & UX-Design
- DevOps and Hosting

TOPIC OF THIS TALK

Over the past 13 years, we have steadily built up our own hosting.

What initially started as a solution for our internal applications quickly expanded to include our clients as well.

Today, we'd like to share some hopefully interesting insights into how our hosting evolved and how Puppet has helped us make this process more efficient.

Not all timings are correct, not all Puppet updates are included.

2011

2011 - DOCS.PUPPETLABS.COM

The screenshot shows the Puppet Labs Documentation website. At the top left is the Puppet Labs logo. The navigation bar includes links for Documentation, Support, Bug Tracker, Contact Us, Download, and a Google Custom Search box. Below this is a secondary navigation bar with links for Puppet, Services, Resources, Community, and Company. The main header area features the title 'Docs: Puppet Labs Documentation' and links for Docs Home, Quick Nav, and Contribute. The main content area is titled 'Puppet Labs Documentation' and includes a welcome message, a link to a downloadable version, and a note about Marionette Collective documentation. A 'Getting Started' section lists several articles with diamond icons. A 'Contents' sidebar on the right lists 11 items, including 'Getting Started', 'Components', 'Extended Knowledge', 'Puppet Dashboard (Web GUI)', 'Advanced Topics', 'Resource Types', 'Extending Puppet', 'Development Information', 'Auto-generated Docs', 'Other Resources', and 'Help Improve This Document'.

puppet labs

Documentation Support Bug Tracker Contact Us Download Google™ Custom Search Search

Puppet Services Resources Community Company

Docs: Puppet Labs Documentation

Docs Home Quick Nav ▾ Contribute

Puppet Labs Documentation

Welcome to the Puppet Labs documentation site.

An downloadable version of this guide may be found on our [downloads page](#).

NOTE: You can find the Marionette Collective documentation [here](#).

Getting Started


New users should begin here.

- ◆ [An Introduction to Puppet](#)
- ◆ [Supported Platforms](#)
- ◆ [Installing Puppet](#) – from packages, source, or gems
- ◆ [Configuring Puppet](#) – includes server setup & testing


Contents

1. [Getting Started](#)
2. [Components](#)
3. [Extended Knowledge](#)
4. [Puppet Dashboard \(Web GUI\)](#)
5. [Advanced Topics](#)
6. [Resource Types](#)
7. [Extending Puppet](#)
8. [Development Information](#)
9. [Auto-generated Docs](#)
10. [Other Resources](#)
11. [Help Improve This Document](#)

2011 - PUPPET.COM


STUFFEDANIMALS.COM *World's Largest Selection* | PUPPET.COM *100's of Puppets* | BIBLETOYS.COM *Faith-Based Toys* | REDWAGONS.COM *Radio Flyer Experts* | PLUSH.COM *For Retailers* | CUSTOMPLUSH.COM *Custom Toys & Corporate Gifts* | 


ORDER TOLL FREE 1.888.317.9237. Our friendly customer service team is available Monday - Friday 9am to 5pm EST.

 [My Account](#) | [Track My Order](#) | [Contact Us](#) | [Help](#) | [Shopping Cart](#)

Order Toll Free 1.888.317.9237 | Search: [GO](#)

RELIGIOUS PUPPETS | HAND PUPPETS | FINGER PUPPETS | SOCK PUPPETS | HALF & FULL BODY PUPPETS | THEATERS

FREE SHIPPING on \$99 orders!  **On Most Items In Our Store!** [Click Here For Complete Details](#)




Find your New Puppet Today!

Our cast of colorful characters are waiting to come to life! All you need is a little imagination!

[View All!](#)

New Additions!



[View All](#)

"My kids love the puppets that I ordered from you. When the order arrived it was so

2011 - CHOOSING THE RIGHT TOOL

- Puppet vs. Chef
- Ansible initial release 2012

We don't remember why we chose puppet :)

2011 - STARTING THE FIRST PUPPET ENVIRONMENT

- 2 servers to be configured with Puppet which were manually configured before
- We started managing parts of the servers configurations
- Puppet 2.6/2.7
- Puppet server (using Phusion Passenger) manually set up

2011 - STARTING THE FIRST PUPPET ENVIRONMENT

- All servers manually configured in `nodes.pp`
- Heavy using of ruby DSL
 - Got later deprecated with Puppet 3 because it was "largely ignored"
 - Primarily as we wanted to use loops/iterations
 - Iterations got implemented in Puppet-DSL in Puppet 4
- Puppet changes are done directly on puppetmaster
 - Versioning and backup using `etckeeper`

2012

2012 - PUPPET AUTOAMI

- Using <https://github.com/ccaum/puppet-autoami>
- Creating up-to-date AMIs
 - Start AWS EC2 instance using our latest AMI
 - Run Puppet agent
 - Custom Puppet Report to check if something got changed
 - If there were changes, create new AMI from current server
 - After that the EC2 instance got terminated

2012 - STARTING OVER

- Starting over (still on Puppet 2.7)
- At the same time as an architectural change (more usage of VMs)
- New repository
- Using Puppet environments
- Without any Ruby-DSL code
- Deployment via rake task
 - rake task clones/pulls git repository in branch named environment

2012 - R10K

Why didn't we use r10k instead of a rake task?

- r10k 1.0.0 got released 2013

Why haven't we switched yet to r10k?

- Our workflow is designed with our monorepo in mind
- We have it on our long-term Roadmap to evaluate it again

2012 - HIERA

- Replaces our nodes.pp and inherited role classes
- Start using own facts to set up servers
- Node configuration got a lot better over the next year

2012 - HIERA AND CUSTOM FACTS

- Using a file
- `modules/cloud/lib/facter/cloud_env.rb`

```
Facter.add("cloud_env") do
  setcode do
    if File.exist? '/etc/puppet/cloud_env'
      Facter::Util::Resolution.exec('/bin/cat /etc/puppet/cloud_env')
    else
      nil
    end
  end
end
```

```
hierarchy:
- name: "Cloud Environment"
  path: "hieradata/{facts.cloud_env}.yaml"
```

2012 - PUPPETDB

- PuppetDB 1.0.0 got released
- New Features
 - Fact storage
 - Catalog storage
 - REST Fact retrieval
 - REST Resource querying

2013

2013 - IMPLEMENTING PUPPETDB

- We started using PuppetDB
- Now we have a good overview of our infrastructure in real-time
- Our documentation improved by using PuppetDB queries

2014

2014 - HEARTBLEED

- Bug in OpenSSL
- With the help of PuppetDB we could easily find any affected server
- Let's encrypt started 2015
- All certificates were manually ordered from CAs

2015

2015 - USE EXPORTED RESOURCES WITH PUPPETDB FOR MONITORING

- We implemented Exported Resources to automatically monitor everything
- For Forge modules we've created Wrapper modules
 - `include makandra::redis` instead of `include redis`
 - Monitoring is generated automatically

2015 - USE PUPPETDB TO CHECK FOR FAILED PUPPET AGENT RUNS

We added monitoring to check for Reports in our PuppetDB to see if Puppet agents fail. As it was complicated to check only the latest Puppet agent run, we've later replaced it with a custom Puppet Report.

Custom Report processors are a nice feature to generate additional informations or alerts from Puppet reports and also to add additional automation.

2015 - USE PUPPETDB TO DO UPGRADES ON ALL SERVERS

Until now we had carried out our server updates semi-manually on a weekly basis.

As the number of servers grew we had to automate this. Only 16 terminals fit well on one screen.

We've created a custom software to query PuppetDB for all needed informations (Servers, OS) to automatically upgrade our servers.

2016

2016 - IMPROVE CUSTOM FACTS

- We started using the fqdn for our custom fact
- Example c42.\$FACTVALUE.makandra.de

```
fqdn = Facter.value(:fqdn)
if fqdn =~ /^c\d\d\.(.*)\.makandra\.de$/
  fqdn.split('.')[1]
end
```

2016 - OCTOCATALOG-DIFF

GitHub published [octocatalog-diff](#).

It compiles a Puppet catalog from 2 branches and compares them.

You can see the catalog changes, which adds a lot of value to a `puppet agent --noop` run.

2016 - OCTOCATALOG-DIFF

TODO ZU NOTZIEZEN ANPASSEN

We've created a script to get a set of different servers and run `octocatalog-diff`.

- At least one server for each OS, OS-Release, virtual/hardware.
- We can also trigger `puppet agent --noop` runs using this list
- The list is autogenerated out of PuppetDB, no manual updates needed!

Later we even improved it to get at least one server for each specific configuration of our infrastructure (Rails Appserver, PHP Appserver, PostgreSQL, MariaDB, Proxy, Monitoring, ...).

2017

2017 - FINALLY UPGRADING TO PUPPET 4

- Puppet "AIO" install
 - All dependencies are included
- New Version of Puppet Language
 - Iteration
 - Type-checking
 - Probably the Puppet upgrade where we found the most bugs in our code
- Puppetserver instead of Passenger Puppetmaster

2017 - IMPROVE OUR FACTS TO UTILIZE HIERA

- Still depends on fqdn
- Can get overwritten by file on server
- `${type}${number}-${stage}.${customer}.example.com`
- `db42-staging.acmeinc.example.com`
 - Type: db
 - Number: 42
 - Stage: staging
 - Customer: acmeinc

2017 - IMPROVE OUR FACTS TO UTILIZE HIERA

```
app23-prod.acmeinc.example.com  
db42-prod.acmeinc.example.com  
app00-stage.acmeinc.example.com
```

hierarchy:

- name: "common"
path: 'common.yaml'
- name: "Per customer data"
path: "customer/\${facts.makandra.customer}.yaml"
- name: "Per type data"
path: "customer/\${facts.makandra.customer}/\${facts.makandra.type}.yaml"
- name: "Per stage data"
path: "customer/\${facts.makandra.customer}/\${facts.makandra.stage}.yaml"
- name: "Per type-stage data"
path: "customer/\${facts.makandra.customer}/\${facts.makandra.type}/\${facts.makandra.stage}.yaml"

2017 - USING PUPPETDB QUERIES IN PUPPET

- We've grown to a complex infrastructure supporting a variety of applications/configurations
- Exported Resources got complicated for some tasks
 - E.g. creating databases using exported resources
 - Exported Resources for database creation will get exported from different servers with different configuration
 - Duplicate declaration on any database change for up to 30 minutes

But fortunately we can query PuppetDB right from our Puppet code.

2017 - USING PUPPETDB QUERIES IN PUPPET

```
class makandra::dbcollector (
  String $dbbackend,
  String $type,
) {

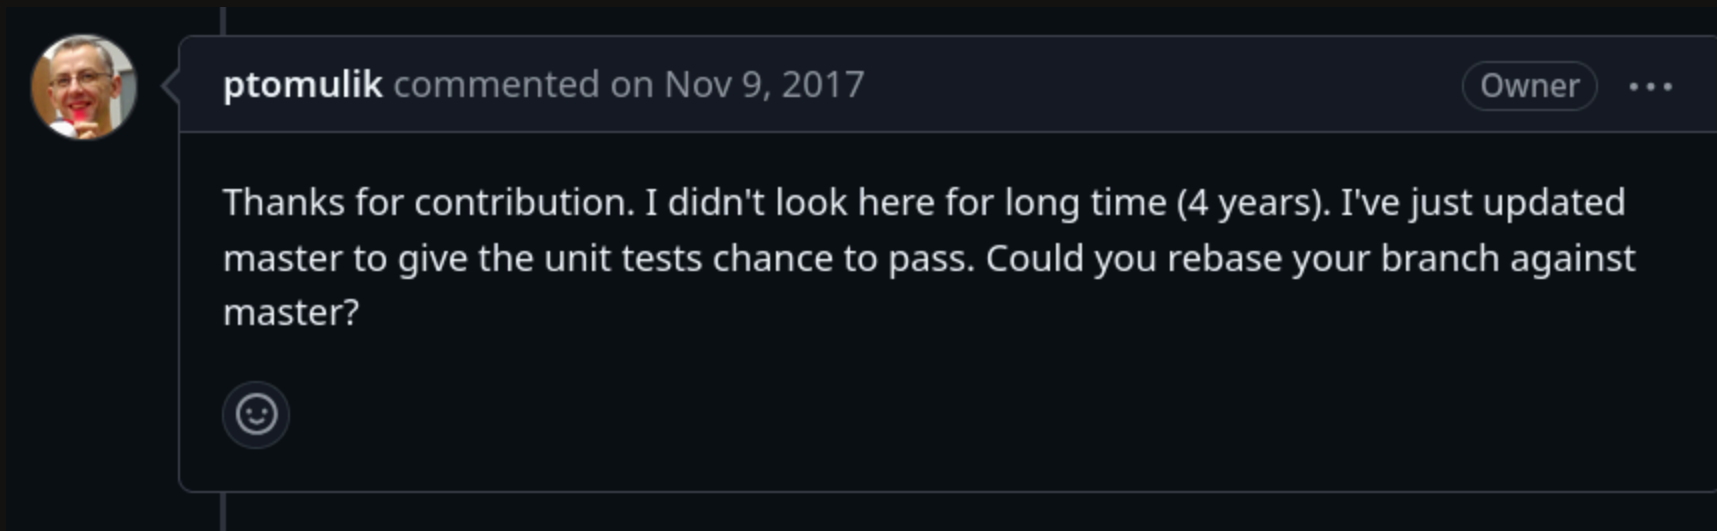
  $vhost_types = [ '\Environments::Elixir::Vhost\'',
                  '\Environments::Rails::Vhost\'',
                  '\Environments::Php::Vhost\'',
                  ]

  $vhosts = puppetdb_query("resources {type in ${vhost_types} and parameters.dbbackend = '${dbbackend}' and parameters.createdb = true}")

  $vhosts.each |$vhost| {
    ensure_resource('environments::railscomplete::db', $vhost['parameters']['dbuser'], {
      'password' => $vhost['parameters']['dbpassword'],
      'type'      => $type,
      'extensions' => $vhost['parameters']['dbextensions'],
    })
  }
}
```

2017 - PUPPET MODULE CONTRIBUTIONS

- When possible we try to get fixes/improvements we do to the upstream module
- Getting fixes in other peoples Puppet modules is often a nice experience



2018

2018 - DEPLOYMENT VIA CI JOB

- Atomic creation of environments
 - We had problems when running puppet agents manually "too soon" while the deploy process was still running
- Better cache clearing of the Puppetmaster cache
- Generate puppet types for better environment isolation
- Combined with pre-receive git hooks to check for:
 - linting (.pp, .rb, .erb, .yaml, ...)
 - correct commit messages
 - custom checks for valid configuration in hiera
 - and more

2019

2019 - IMPROVE OUR VM SETUP

- We've create a new tool to start VMs in our infrastructure
- The script asks PuppetDB which VM Hosts are available and were VMs can get started
- We check if Puppet catalog can compile before the VM get startet with octocatalog-diff
- IP adresses are assigned from a database which got updated automatically from the PuppetDB content
- An allowlist entry is generated for the puppetserver so the new VM certificate will get signed automatically

2019 - PUPPET HELPS IF YOU SCREW UP

- Puppet agent overwrites file
- You notice there was content in this file which should not get overwritten
- You don't have a backup as you modified it just one hour ago
- Puppet to the rescue!
 - Your file is backed up in `/opt/puppetlabs/puppet/cache/clientbucket`
 - It only happened once for us but instructions are documented here:
<https://makandracards.com/operations/29357-get-backup-files-deleted-changed-puppet>

2020

2020 - FACTS AGAIN, HELPING WITH "MULTI-CLOUD"

- We manage servers for a lot of customers in different datacenters

```
[...]
return 'aws' if Facter.value(:ec2_metadata) && (Facter.value(:ec2_metadata)['services']['partition'] == 'aws')
return 'vagrant' if Facter.value(:puppet_server) == 'vagrant-puppetmaster'
return 'localtesting' if Facter.value(:puppet_server) == 'puppetmaster.lxd'

# try interfaces
Facter.value(:networking)['interfaces'].each_key do |interface|
  [...]
```

2021

2021 - USE PUPPET BOLT AND PUPPETDB TO DO BETTER AND FASTER UPGRADES

We've created Puppet Bolt Tasks to do the upgrades on our servers.

- A lot less complex than our custom software.
 - Easy to parallelize/scale
1. Bolt connects via SSH to all servers to create a list of updates
 2. We check the list manually for possible problems
 3. We start the update process on all servers
 4. Updates are executed in stages depending on puppet facts, so that e.g. staging servers get updated first

2021 - EVEN MORE BOLT

After we started using Bolt, we now use it regularly for a lot of tasks:

- Check servers for BIOS Updates
- Check Software for available updates
- Roll out changes "at the same time" to different servers
 - Database password changes
 - Force puppet runs
 - Restart services
 - Search for processes
 - Remount network filesystems

2022

2022 - TESTS WITH LOCAL CONTAINERS

- Using LXC
- Fully automated setup including puppetmaster container
- Starting local containers using production configuration
- Simple overrides for local differences using hiera

2023

2023 - MODULE UPGRADE PAIN

We needed a new version of puppetlabs-stdlib but also more time to replace old functions.
(Could help with your Puppet 8 upgrade, but don't tell anyone you doing this and we didn't test it).

2023 - MODULE UPGRADE PAIN

[#16204] puppetlabs/stdlib 9: Add compatibility module

```
---
modules/stdlib_compat/lib/puppet/functions/ensure_packages.rb | 61 ++++++
modules/stdlib_compat/lib/puppet/functions/merge.rb          | 112 ++++++
modules/stdlib_compat/lib/puppet/parser/functions/fqdn_rand_string.rb | 41 ++++++
modules/stdlib_compat/lib/puppet/parser/functions/has_key.rb  | 41 ++++++
modules/stdlib_compat/lib/puppet/parser/functions/hash.rb     | 48 ++++++
modules/stdlib_compat/lib/puppet/parser/functions/is_function_available.rb | 32 ++++++
modules/stdlib_compat/lib/puppet/parser/functions/is_hash.rb  | 28 ++++++
modules/stdlib_compat/lib/puppet/parser/functions/is_integer.rb | 52 ++++++
modules/stdlib_compat/lib/puppet/parser/functions/prefix.rb   | 57 ++++++
modules/stdlib_compat/lib/puppet/parser/functions/validate_ip_address.rb | 62 ++++++
10 files changed, 534 insertions(+)
```

2023 - PERFORMANCE PROBLEMS

- To create monitoring configuration we've used puppetlabs/nagios_core
- We have >15000 services, which meant >15000 resources
- Puppet runs took more than 15 minutes
- We've replaced this with a custom module which uses PuppetDB Queries and generate a single file
- Puppet runs now take less than 30 seconds

2024

2024 - OS-RELEASE UPGRADE

- We use Puppet and a ruby-gem to automate dist-upgrades on our Ubuntu servers
- The ruby-gem handles dist-upgrade itself
 - Utilized PuppetDB to create tasks
 - Check hiera for specific informations about the affected server
- Our Puppet module does the configurations before and after the dist-upgrade using facts
 - Execute upgrade-specific changes before and after the dist-upgrade
 - Installs our dist-upgrade-gem and its dependencies

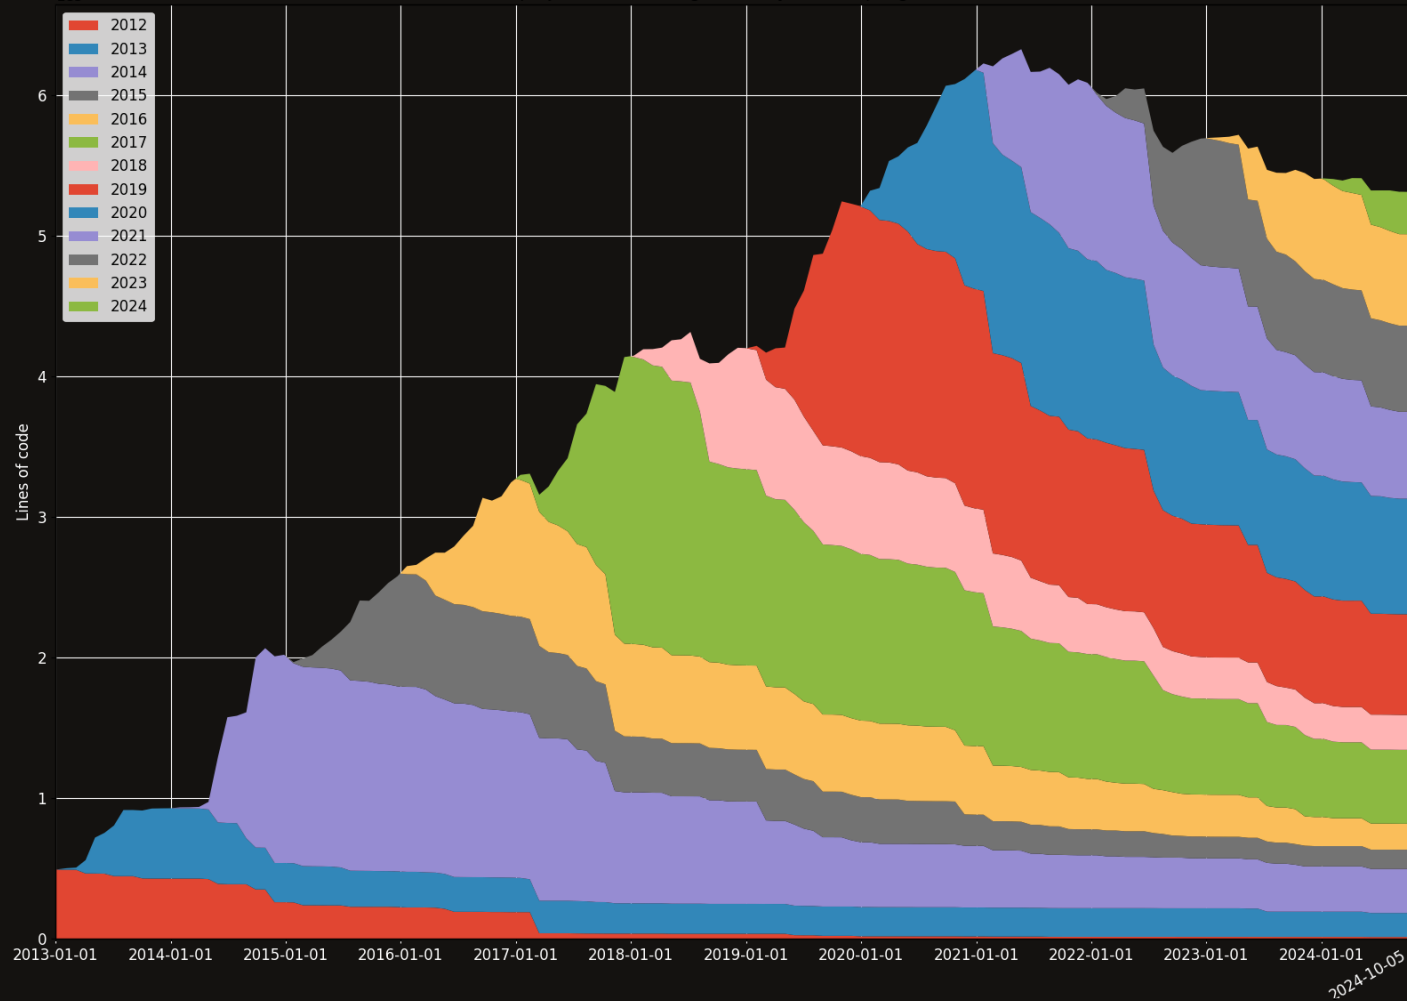
2024 - CURRENT CUSTOM FACTS

Over the years we've added more custom facts, for example:

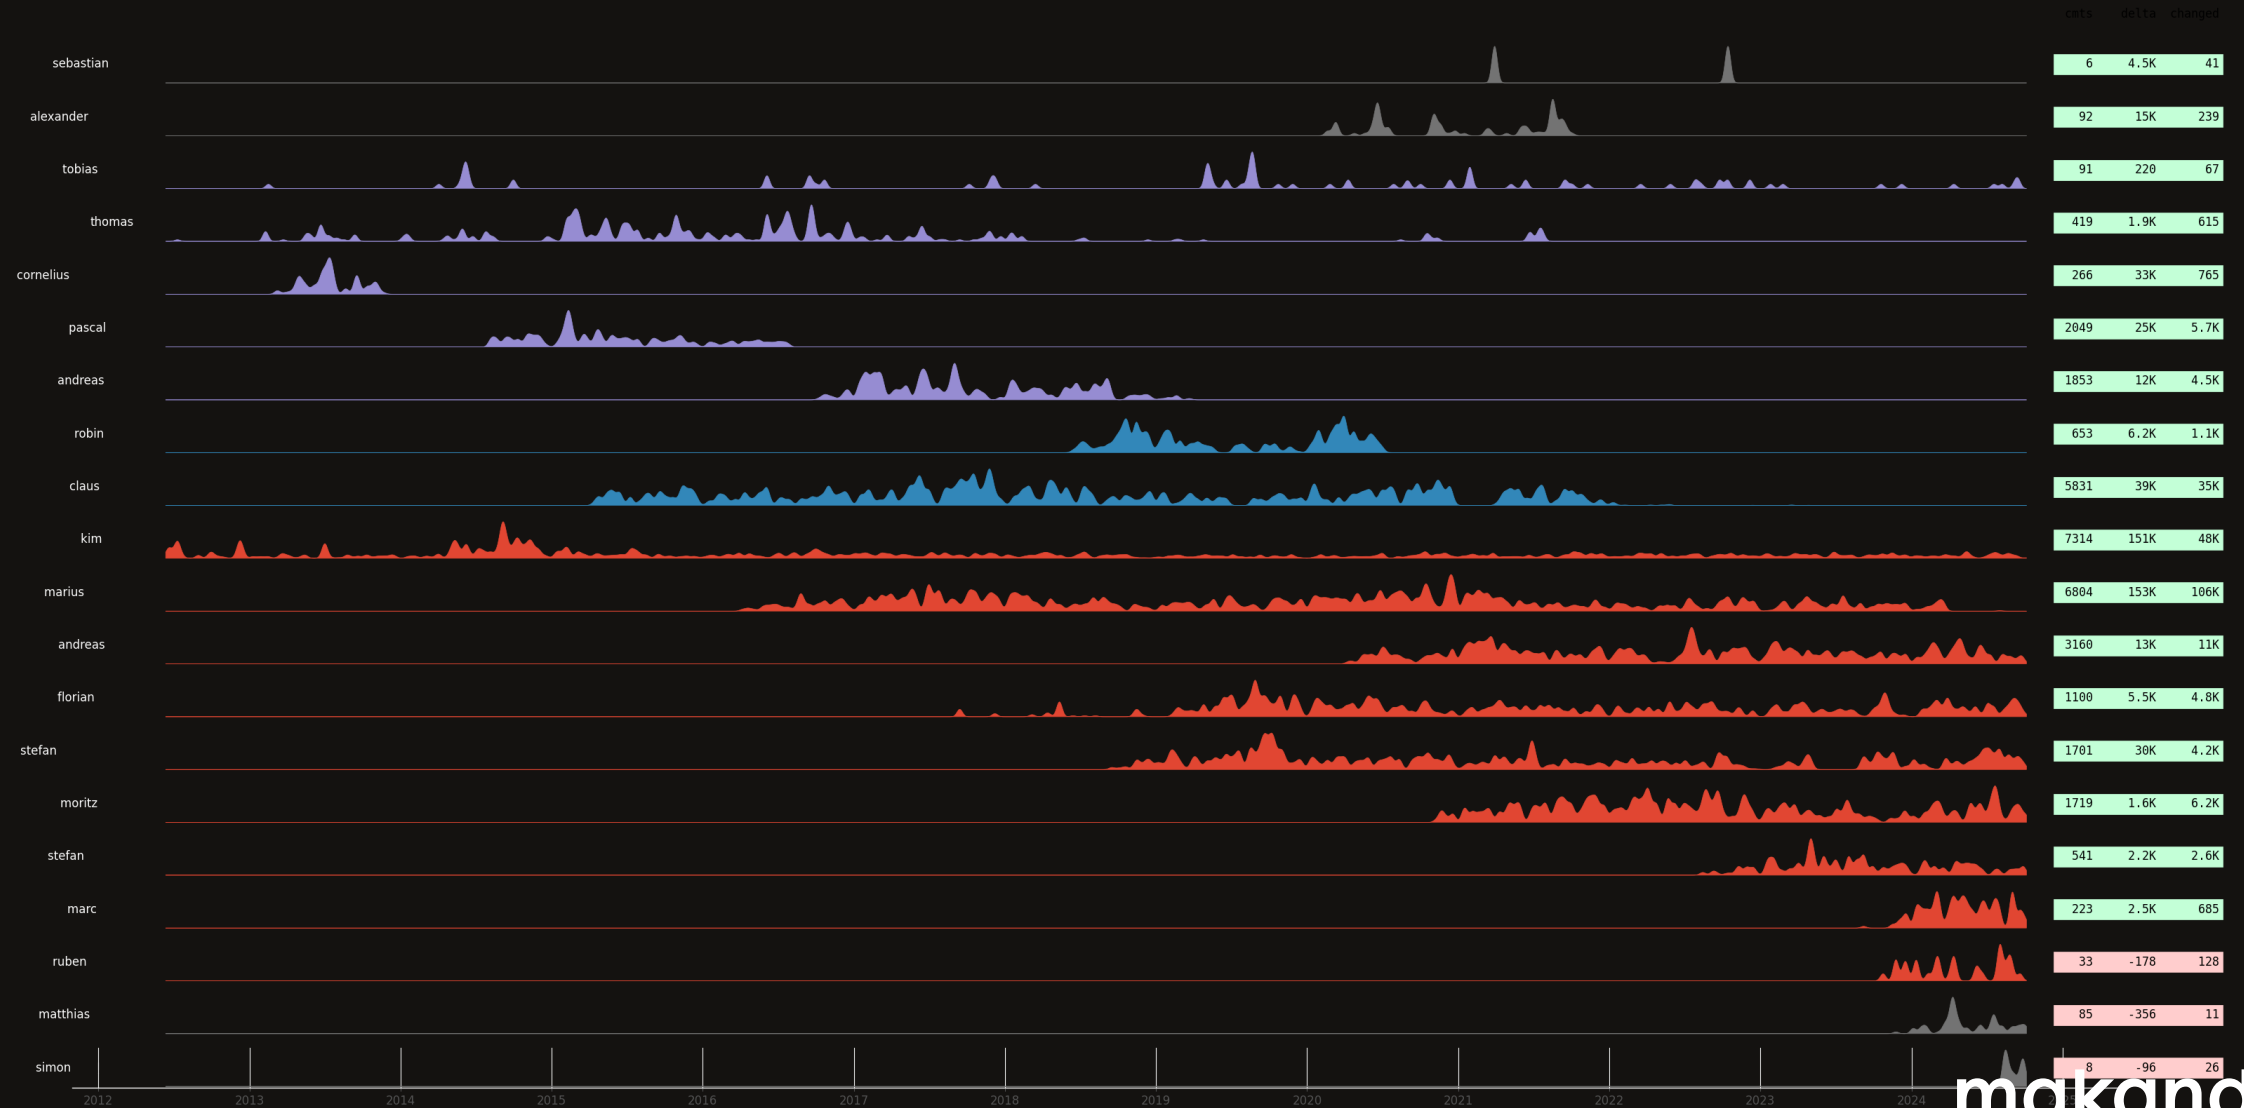
- Public Keys to configure secure Backups
- IPMI (BIOS) Version (for Updates)
- Host Server on which a VM is running
- and more

SUMMARY

SUMMARY - GIT LINES OF CODE



SUMMARY - DEVOPS TEAM



SUMMARY

- More than 20 people worked on our codebase
- Same codebase from Puppet 2.7 to 7
- Puppet is easy upgradable
 - This code works on any release

```
package { 'vim':  
  ensure => present,  
}
```

THE FUTURE

- More tests
 - Tests in Puppet modules
 - octocatalog-diff in our pipeline
- renovate bot for puppet modules
- Additional hiera backends
- Puppet 8
 - Automated Certificate Renewal
 - Unchanged Resources Excluded from Reporting
 - We have servers with over 6000 Resources
 - Our PuppetDB Database is >72GB

END

QUESTIONS?
