

Resilient puppet in the cloud

OSAD 2022



Dr. Jan Bundesmann

October 4th, 2022

- 1 Cloud & Puppet: Is this a thing?
- 2 Stacked puppetserver infrastructure
- 3 Target focussing: Agents as stacks & units
- 4 Conclusions and Outlook

- 1 Cloud & Puppet: Is this a thing?
- 2 Stacked puppetserver infrastructure
- 3 Target focussing: Agents as stacks & units
- 4 Conclusions and Outlook

Why wouldn't it be a thing?



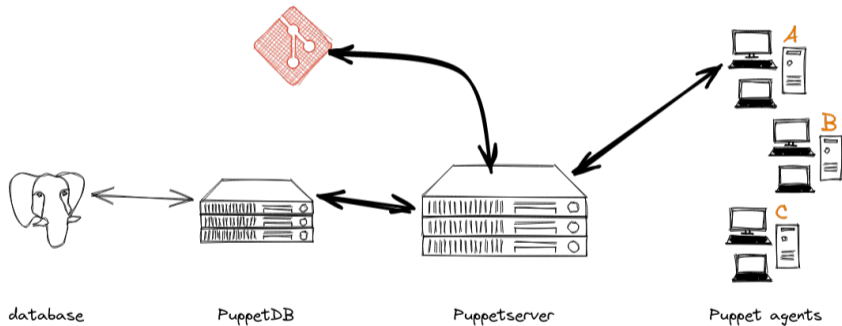
- ▶ Old-fashioned: no cloud when Puppet was invented
- ▶ Slow: 30min heartbeat vs. seconds in GitOps

- ▶ No immediate spin-up
 - ▶ No immutable infrastructure
 - ▶ auto-deployment
 - ▶ arbitrarily complex infrastructure
-
- ▶ No pets, no cattle
 - ▶ Manageable complexity

- 1 Cloud & Puppet: Is this a thing?
- 2 Stacked puppetserver infrastructure**
- 3 Target focussing: Agents as stacks & units
- 4 Conclusions and Outlook

What? Stacks?

Stacks: components of (distributed) applications, consisting of one or more units



Why? Stacks?



- ▶ De-individualize instances
 - ▶ Care for working units
 - ▶ Redeploy non-working units
 - ▶ Stack broken?
- ▶ For puppet: use cloud mechanisms, map cloud features
 - ▶ DBaaS, GITaaS
 - ▶ Scalability
 - ▶ Identity and Access Management, Secret Management

How? Stacks?



- ▶ Single source of truth
- ▶ One CA, multiple puppet servers
 - ▶ Access rights
 - ▶ Backup
 - ▶ Unique identifiers
- ▶ Race conditions
 - ▶ Tolerance for DB inaccessibility
 - ▶ Logging / Monitoring: expected failures (regarding PuppetDB in particular)
- ▶ Puppet manages everything

How? Stacks?



- ▶ Single source of truth → control repository, r10k
- ▶ One CA, multiple puppet servers
 - ▶ Access rights
 - ▶ Backup
 - ▶ Unique identifiers
- ▶ Race conditions
 - ▶ Tolerance for DB inaccessibility
 - ▶ Logging / Monitoring: expected failures (regarding PuppetDB in particular)
- ▶ Puppet manages everything


- ▶ Single source of truth → control repository, r10k
- ▶ One CA, multiple puppetmasters → shared folder, each instance considers itself singular
 - ▶ Access rights
 - ▶ Backup
 - ▶ Unique identifiers
- ▶ Race conditions
 - ▶ Tolerance for DB inaccessibility
 - ▶ Logging / Monitoring: expected failures (regarding PuppetDB in particular)
- ▶ Puppet manages everything

- ▶ Single source of truth → control repository, r10k
- ▶ One CA, multiple puppetmasters → shared folder, each instance considers itself singular
 - ▶ Access rights
 - ▶ Backup
 - ▶ Unique identifiers
- ▶ Race conditions → Eventual consistency through foreman modules
 - ▶ Tolerance for DB inaccessibility
 - ▶ Logging / Monitoring: expected failures (regarding PuppetDB in particular)
- ▶ Puppet manages everything

- ▶ Single source of truth → control repository, r10k
- ▶ One CA, multiple puppetmasters → shared folder, each instance considers itself singular
 - ▶ Access rights
 - ▶ Backup
 - ▶ Unique identifiers
- ▶ Race conditions → Eventual consistency through foreman modules
 - ▶ Tolerance for DB inaccessibility → `soft_write_failure`
 - ▶ Logging / Monitoring: expected failures (regarding PuppetDB in particular)
- ▶ Puppet manages everything

- ▶ Single source of truth → control repository, r10k
- ▶ One CA, multiple puppetmasters → shared folder, each instance considers itself singular
 - ▶ Access rights
 - ▶ Backup
 - ▶ Unique identifiers
- ▶ Race conditions → Eventual consistency through foreman modules
 - ▶ Tolerance for DB inaccessibility → `soft_write_failure`
 - ▶ Logging / Monitoring: expected failures (regarding PuppetDB in particular)
- ▶ Puppet manages everything → next section

- ▶ Infrastructure is
 - ▶ Ephemeral
 - ▶ Idempotent
 - ▶ Not Immutable
- ▶ Cloud manages scaling of stacks
 - ▶ Autoscaling, load balancing
 - ▶ Unique identifiers
- ▶ Redeployed within 20 minutes
 - ▶ Raw OS images
 - ▶ Accelerate with your specialized images




Agenda



- 1 Cloud & Puppet: Is this a thing?
- 2 Stacked puppetserver infrastructure
- 3 Target focussing: Agents as stacks & units**
- 4 Conclusions and Outlook

How does puppet work



▶ Server chooses (very old)

```
1 node 'www1.example.com' {  
2   include common  
3   include apache  
4   include squid  
5 }
```

▶ Server chooses (nowadays, using hiera lookup)

```
1 hiera_include('classes')
```

Determine classes from hiera



hiera.yaml:

```
1 ---
2 version: 5
3 defaults:
4 hierarchy:
5   - name: "Per-node data"
6     path: "nodes/{::trusted.certname}.yaml"
7   - name: "Other YAML hierarchy levels"
8     path: "common.yaml"
```

Hiera structure:

```
1 .
2 +-- common.yaml
3 +-- nodes
4   +-- www1.example.com.yaml
5   +-- www2.example.com.yaml
```

common.yaml:

```
1 classes:
2   - common
```

nodes/www1.example.com.yaml:

```
1 classes:
2   - apache
3   - squid
```

nodes/www2.example.com.yaml:

```
1 classes:
2   - nginx
3   - certbot
```

- ▶ Write facts
 - ▶ custom facts (written in ruby)
 - ▶ external facts (written in any language)
 - ▶ hostname based
 - ▶ ip-range based
 - ▶ query API of inventory system
 - ▶ simple plain JSON or YAML files
- ▶ Which facts determines the stack?
- ▶ Use hiera to aggregate classes as in roles, profiles or **stacks!**

Obtain list of profiles using mapped_paths



```
1 ---
2 version: 5
3 defaults:
4 hierarchy:
5   - name: "roles"
6     mapped_paths:
7       - stacks
8       - stack
9       - "stack/{role}.yaml"
10  - name: "Generic defaults"
11    path: "common.yaml"
```

```
1 .
2 +-- common.yaml
3 +-- stack
4   +-- jitsi.yaml
5   +-- pizza_delivery.yaml
6   +-- puppetdb.yaml
7   +-- puppetserver.yaml
8   +-- wordpress.yaml
```

Example: puppetserver



```
/etc/facter/facts.d/  
roles.json:
```

```
1 {  
2   "stacks": [  
3     - "puppetserver"  
4   ]  
5 }
```

```
puppetserver.yaml
```

```
1 ---  
2 classes:  
3   - puppet  
4   - puppet::server::puppetdb  
5  
6 puppet::server: true  
7 puppet::server_ca: true  
8  
9 # important: the module by default tries to contact a  
10  # foreman instance  
11 # the following three parameters prevent this  
12 puppet::server_foreman: false  
13 puppet::server_foreman_facts: false  
14 puppet::server_external_nodes: ''  
15  
16 # settings for storing catalogs and facts in puppetdb  
17 puppet::server_storeconfigs: true  
18 puppet::server_reports: store,puppetdb  
19 puppet::server::puppetdb::server: puppet-db.server.in.the  
20  .cloud  
21 puppet::server::puppetdb::soft_write_failure: true
```

```
/etc/facter/facts.d/  
roles.json:
```

```
1 {  
2   "stacks": "wordpress"  
3 }
```

```
wordpress.yaml
```

```
1 ---  
2 classes:  
3   - wordpress  
4   - reverse_proxy  
5  
6 reverse_proxy::virtual_server:  
7   - name: %{:facts.fqdn}  
8     proxy_backend: http://localhost:8080  
9  
10 wordpress::installed_plugins: recipes
```

Agenda



- 1 Cloud & Puppet: Is this a thing?
- 2 Stacked puppetserver infrastructure
- 3 Target focussing: Agents as stacks & units
- 4** Conclusions and Outlook

- ▶ De-cluster using stacks!
- ▶ Use stacks to strengthen and fasten your infrastructure!
- ▶ Empower your agents to choose stacks!