Karlsruher Institut für Technologie

Configuration and Deployment Management at KIT(*)

Andreas Petzold

STEINBUCH CENTRE FOR COMPUTING - SCC

Infrastructures at KIT

- Basic IT services
  - email, webservers, ESX, SAP, campus software, campus network
  - OpenStack-based cloud
- High Performance Computing
  - 3 running clusters
  - 2 large new systems in the pipeline
- Data Intensive Computing
  - GridKa Tier-1
  - Large Scale Data Facility
  - Smart Data Innovation Lab

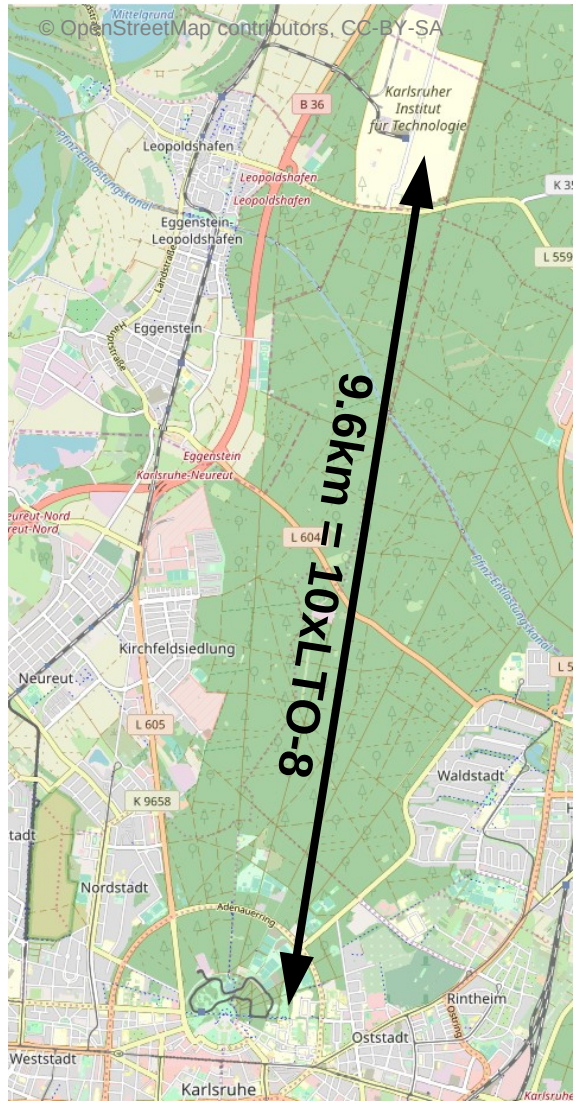

Configuration & Deployment Management

- basic IT services
  - many different approaches including **Ansible**
- High Performance Computing
  - homegrown system system: **KITE**
- Data Intensive Computing
  - **Katello/Satellite/Foreman + Puppet**

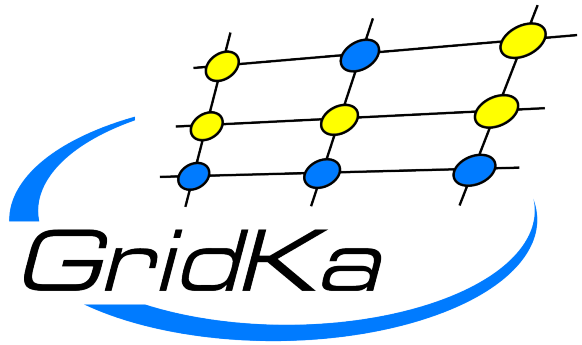

Smart Data Innovation Lab

LSDF

Configuration & Deployment Management

- basic IT services
  - many different approaches including **Ansible**
- High Performance Computing
  - homegrown system system: **KITE**
- Data Intensive Computing
  - **Katello/Satellite/Foreman + Puppet**

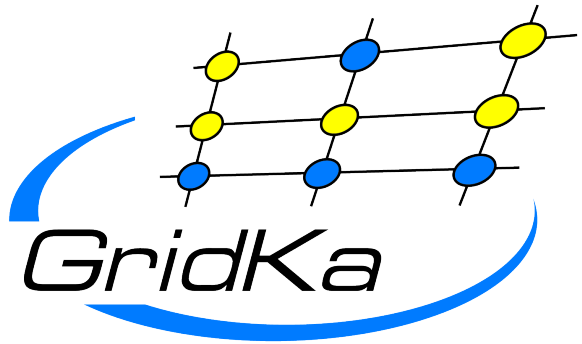

Smart Data Innovation Lab

LSDF

Our Goals

- Centralized control of configuration
  - everything should be in GitLab
- Self-service machine deployment for bare metal and VMs
  - unified interface to all resources
- Centrally managed RPM repositories
  - local mirrors of external repos & private repos
- Share knowledge & workflows
  - benefits for everyone

History

- CFEngine for “Config Management”
- ROCKS for GridKa Cluster deployment + homegrown solution
- many scripts & lists of commands in wiki etc.

- 2013: started looking into Puppet & Foreman
- First used for managing CEs
- GitLabCI integration
- Separate Puppet&Foreman instances inside&outside of GridKa
- RH Satellite also for ppc64(le)
- Foreman integration with Power HMC
- 2017: Spectrum Scale servers & dCache (+everything new/redeployed)
- Foreman integration with KIT DNS
- early 2019: ROCKS still manages SL6 WNs and GridKa DNS + DHCP

GitLab + CI

- Puppet Modules
  - homegrown
  - forks of public modules
  - direct import of public modules (puppet librarian)
- Hiera Data
- Foreman templates
- Commit triggers pipelines for syntax check and deployment
- Feature branches → Puppet environments
- push to deployment ~1m

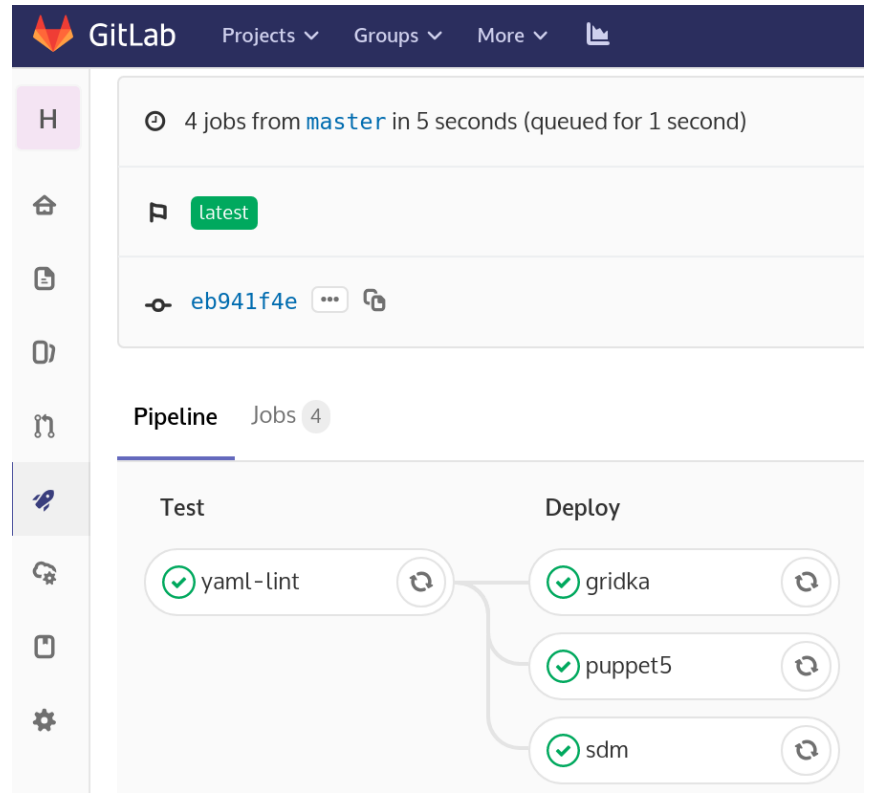

Infrastructure

- Foreman/Katello manage all aspects of host life cycle
  - bare metal, VMs (RHV/oVirt/OpenStack/ESX/LPARs)
  - integration with local inventory DB (IdoIT)
- SmartProxies
  - multiple instances help with scaling
  - used to communicate with separate/private networks
  - DNS (bind + KIT DNSVS)
  - DHCP
  - *Content* (Repositories, Errata), RedHat Subscriptions
  - BMC
  - Puppet masters
- Plan to migrate all hosts into a unified Katello instance with SmartProxies (Capsules)

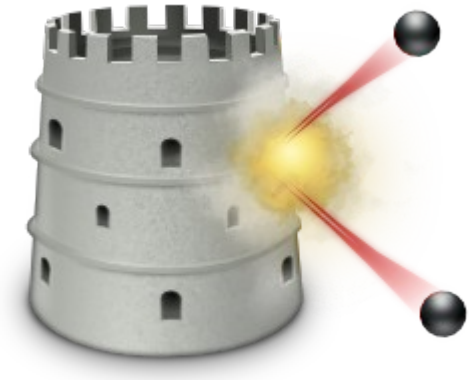

Why Katello?

- Pulp: Repository and Errata Management
- Repositories
  - External repos can be centrally synced
  - Local repos can be centrally managed (API)
- Content Views
  - customizable views of repos: frozen in time/filtered + versioning
  - can be assigned to hosts/groups
- Errata management
  - katello agent on node reports back
  - overview of patch status of all hosts

Name	Subscription Status	Installable Updates	OS
compute-23.sdil.kit.edu	✓	0 ▲ 0 🐞 0 ✚ 0 📦	RHEL 7.6
foreman-test-rhel7-ppc64le.sdil.kit.edu	✓	11 ▲ 17 🐞 2 ✚ 113 📦	RHEL 7.6
foreman-test-rhel7.sdil.kit.edu	✓	8 ▲ 13 🐞 1 ✚ 84 📦	RHEL 7.6

Infrastructure

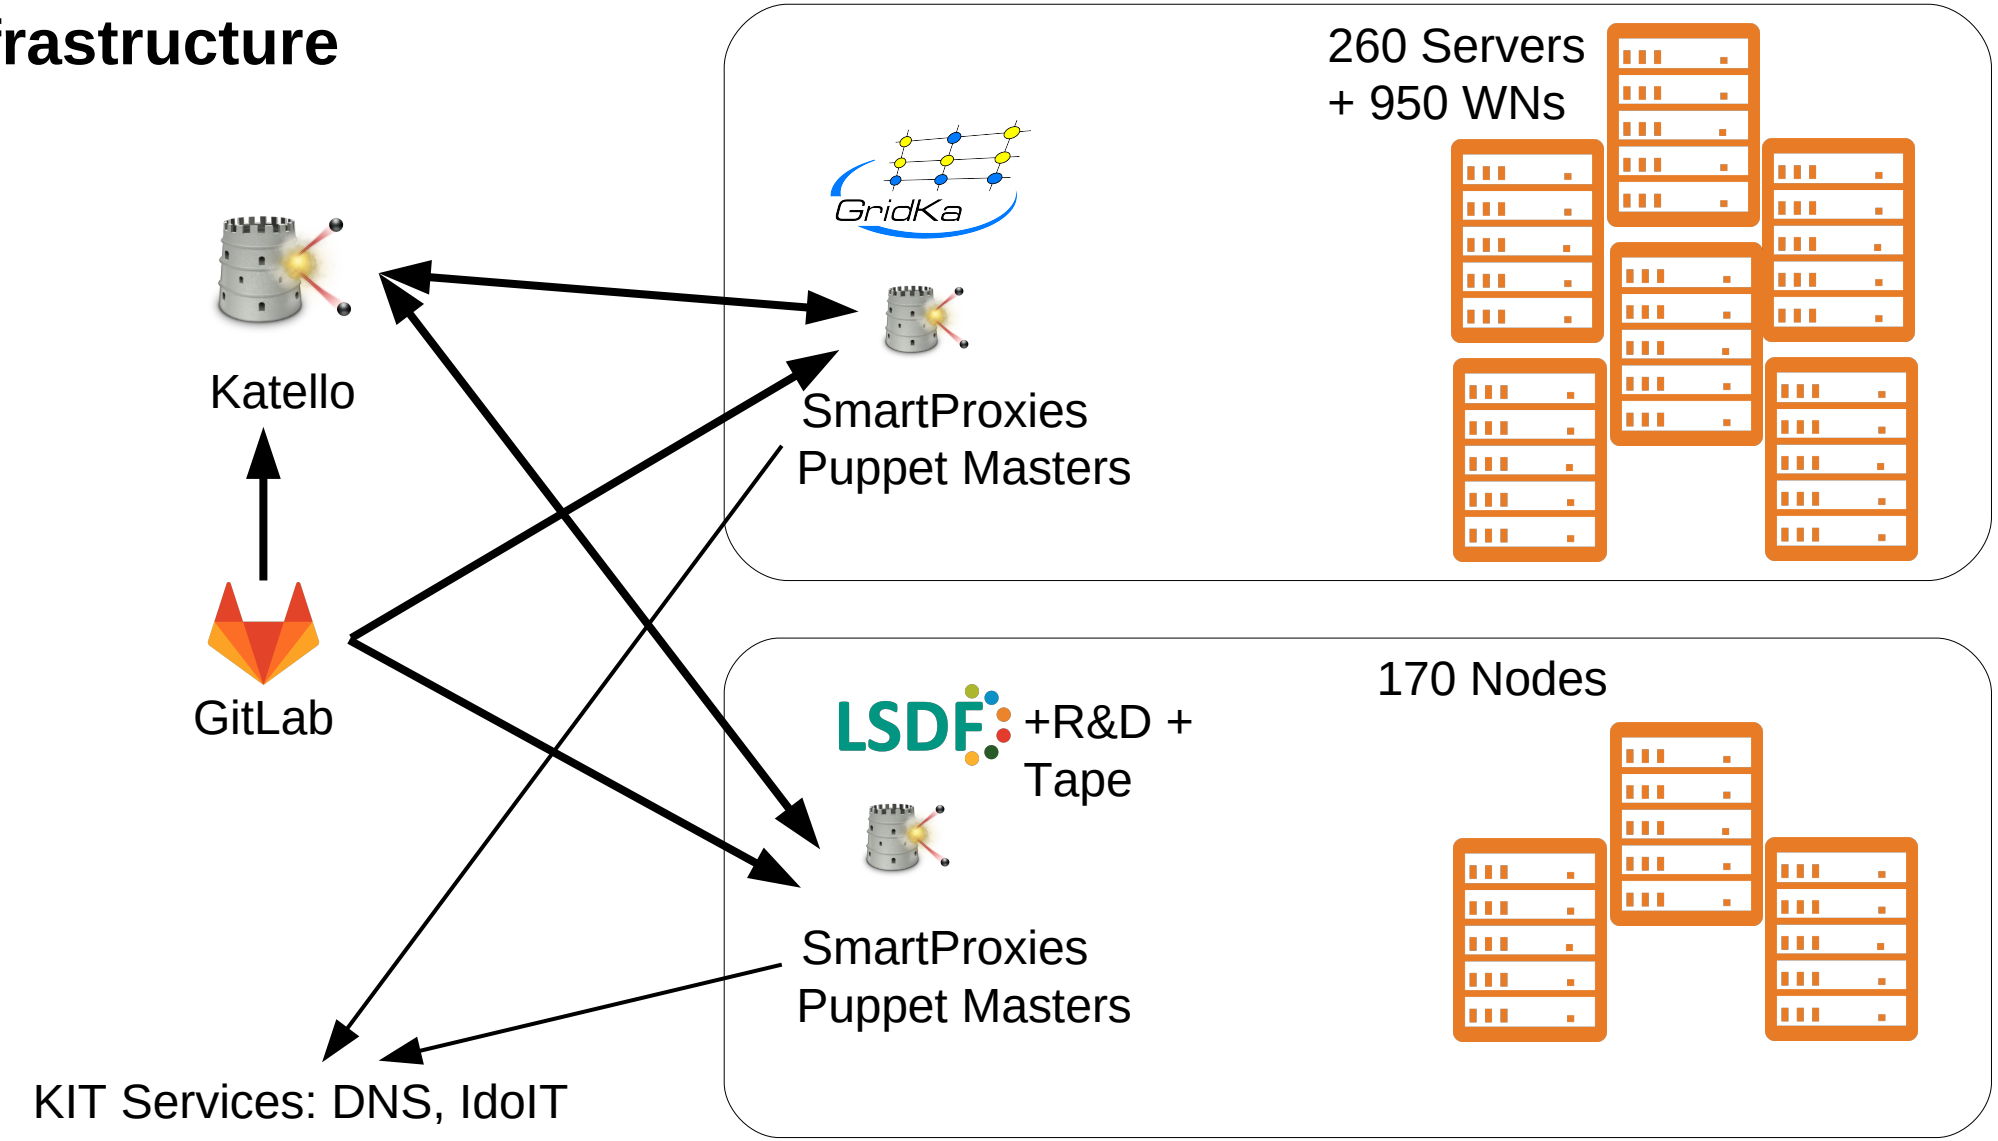

Deployment Workflow (Bare Metal)

- New host boots Foreman Discovery Image
  - queries inventory DB for location, IP, name
  - reports hardware details to Foreman
  - DHCP reconfigured
- Discovered host can be provisioned

Discovered hosts » Discovered host: mac90b11c08f529

Back | Select Action | Delete

Highlights

architecture	x86_64
discovery_bootif	90:b1:1c:08:f5:29
discovery_bootip	10.97.106.112
discovery_subnet	GridKa internal net a01-106 (10.97.106.0)
ipaddress	10.97.106.112
macaddress	90:b1:1c:08:f5:29
manufacturer	Dell Inc.
memorysize	47.12 GB
memorysize_mb	48254.66
physicalprocessorcount	2
processorcount	24
productname	PowerEdge R420

Interfaces

Type	Identifier	MAC address	IP address
	em1	90:b1:1c:08:f5:29	10.97.106.112
	em2	90:b1:1c:08:f5:2a	N/A
	ipmi	90:b1:1c:08:f5:2b	192.168.1.112

Deployment Workflow (Bare Metal)

- New host boots Foreman Discovery Image
  - queries inventory DB for location, IP, name
  - reports hardware details to Foreman
  - DHCP reconfigured
- Discovered host can be provisioned
  - will test autoprovisioning rules for WNs
  - DNS entries provisioned
- Kickstart installation
- 1st Puppet run
- Reboot

Human Side

- Puppet Modules
  - 2013: everyone should write Puppet modules
  - 2019: most people only ever interact with Hiera. Few people develop puppet modules
- Pets vs. Cattle
  - The bad guys taking away people’s pets and give them cattle instead
  - Takes time to get the new philosophy across
- New users want the feature you wanted to implement next week
  - Need to learn that iterations are OK
- Requires a project manager with good soft skills to get the users on board
  - weekly “puppet clinic” to quickly address specific questions

Next Steps

- WN deployment
  - +150% number of hosts
- Migration to Puppet 6
  - reduction of technological debt
- Separate Hiera repositories
  - priviledge separation
- Migration to central Katello
  - errata management

Summary

- We have centralized control
- We have common workflows for all infrastructures
- Soft skills are important
- We still have a lot of work ahead