

Grids and Clouds in the Czech NGI

Jan Kandrát

jan.kundrat@cesnet.cz

Martin Adam, Dagmar Adamová, Jiří Chudoba, Tomáš Kouba,
Miloš Lokajíček, Alexandr Mikula, Václav Říkal, Jan Švec,
Rudolf Vohnout

Institute of Physics of the CAS, and CESNET

September 30, 2015

Computing in the Czech Republic

- ▶ HPC for physicists: Institute of Physics at the CAS
- ▶ General-purpose HPC: MetaCentrum
- ▶ IaaS Cloud Computing: CESNET's FIWARE Lab
- ▶ Many other options

Computing in the Czech Republic

- ▶ **HPC for physicists: Institute of Physics at the CAS**
- ▶ General-purpose HPC: MetaCentrum
- ▶ **IaaS Cloud Computing: CESNET's FIWARE Lab**
- ▶ Many other options

Prague Tier-2 LHC Site

- ▶ Main site at Institute of Physics
 - ▶ 258 WNs, 4756 CPU cores, 4.25 PB disk storage, a tape library
 - ▶ 62 m² floor space
 - ▶ 1 × 200 kVA + 2 × 100 kVA UPS, 350 kVA diesel
 - ▶ 3 × 10 GbE external IP connection (LHCONE, PASNET, CESNET)
- ▶ Some HW hosted at CESNET's premises
 - ▶ 31 WNs, 920 CPU cores, 2416 GB RAM, 180 TB disk storage
 - ▶ close to a CESNET POP

War Stories: Hardware

- ▶ Power distribution and blown fuses
- ▶ Switch HW troubles, dying ports
- ▶ Physical network link bottlenecks due to the network topology

War Stories: Software and Operations

- ▶ Xen → KVM
- ▶ Manual PXE → Spacewalk
- ▶ shell scripts → cfengine v2 & SVN → Puppet & git
- ▶ Monitoring: Munin, Ganglia, Nagios w/multisite, MRTG, Netflow

War Stories: IPv6

- ▶ One of the first sites whose file-transfer services are 100% IPv6-enabled
- ▶ Fun stories
 - ▶ IPv6 problems at an unnamed remote site
 - ▶ Just two machines out of a big server pool were hitting errors
 - ▶ Turned out to be a misconfigured router on the remote site

IaaS: OpenStack cluster at CESNET

- ▶ FI-PPP, FIWARE Lab
 - ▶ Federated architecture
 - ▶ Central services provided by international partners
- ▶ OpenStack on an RHEL clone
 - ▶ Puppet, StackForge
 - ▶ OpenStack Grizzly on SL6 → OpenStack Kilo on CentOS 7

Hardware

- ▶ 192 AMD 6344 cores, 48 E5-1650v2 cores, 1 TB RAM
- ▶ 8.25 TB 10k rpm HDDs, 3.84 TB SATA SSDs, 1 TB PCIe SSDs
- ▶ SuperMicro FatTwin, Apple Mac Pro
 - ▶ Mac OS X' licensing issues
- ▶ Dual-10Gbps internal network
 - ▶ Cisco Nexus 3172PQ
 - ▶ Juniper EX4300-24T

Services

- ▶ Infrastructure as a Service: OpenStack
- ▶ Continuous Integration for SW developers
 - ▶ Gerrit, Zuul – inspired by the OpenStack Infrastructure project
 - ▶ Proof-of-concept demo for the KDE project
 - ▶ Access to the OS X platform for developers
- ▶ Stories from the trenches
 - ▶ EOL-ed version of OpenStack
 - ▶ VM migration
 - ▶ IPv6 support (again)
- ▶ Migrating to CEPH

Monitoring Setup

- ▶ Built on top of `collectd`
 - ▶ 10,000+ metrics
 - ▶ By default, gathered every 10 s
 - ▶ Excellent performance thanks to `rrdcached`
- ▶ Alerting by Nagios

Questions

Thank you for your attention!