OpenStack & Ceph

An Update of Enterprise Use-Cases at BMW

Andreas Pöschl
Senior Solutions Architect
BMW Group

Dyon van Bavel
OpenStack Solutions Architect
Alphabet / BMW Group

SUSE
We adapt. You succeed.

The environment at BMW

Where we come from

- Huge Linux experience and always driving innovation
  - Using Linux since 2002
  - Migrated from RISC servers to x86 (mostly Linux) until 2010
  - Virtualizing production workloads on SLES®/XEN since 2008
    - Using Fibre Channel SAN storage and host based mirroring (HBM)
    - Unique combination of HBM and live migration ability (block-dmmd)
    - Management via homegrown scripting framework
  - Approx. 4500 Linux instances in place
    - 1300 physical servers, 350 virtualization hosts, 2800 VMs
    - Also virtualized most business critical workloads (production control systems, SAP, databases (PostgreSQL/Oracle)

The environment at BMW

Where we come from (2)

- Huge variety of compute workloads
  - VMs from 1 vCPU / 1GB to 16 vCPU / 64 GB
  - Physical servers up to 60 cores (+HT) and 1 TB memory
- SAN storage on enterprise arrays, connected via 8/16GBit fibre channel SAN (approx. 4,5 PB)
  - Only mirrored setups in separate datacenters
  - Separate LUNs for application storage and VM instances
- NAS storage (NFS) for parallel access (especially in the web server environment (approx. 5,7 PB)

The environment at BMW

Storage issues …

- Currently used (enterprise) storage is safe, fast and worth it’s price, but lacks agility on migrations and changing environment
- Currently we are trying to gain more flexibility by using NPIV with FC based SAN
- Migrations and necessary reconfigurations are still an issue in an FC based SAN environment, even when using automation
- NFS is widely used and might resolve some of the issues, but is not a cloud storage system (from our perspective)

The environment at BMW

… how to resolve them

We have been looking for:

- An IP based storage architecture,
- Which keeps us independent from storage vendors (and underlying technology at all),
- Is disaster safe,
- Has no single point of failure
- And is “native” to us (integration into Linux distribution)

Ceph/Rados Comes Into Play

Our expectations

How we plan to benefit from Ceph/Rados

- Increased flexibility by using Ethernet instead of Fibre Channel (easy routing in case of changes)
- No necessity for resource intensive migration tasks in the future as Ceph handles the (re-)placement
- Possibilities like thin provisioning, cloning and snapshotting without vendor specific storage extensions
- Future proof concept (developed for the cloud) and ready also for future use cases

Our goals – 3 years ago

How we planned to use Ceph/Rados

- We wanted to limit the scope of Ceph/Rados for the first implementation (just use the Rados block layer (rbd) to store images for virtual servers)
- No CephFS or NFS replacement intended (yet)
- Leverage the possibilities only, where you can live with the drawbacks (in terms of maturity and admin experience)

➔ Try to replace at least a part of the SAN storage used for VMs with low criticality and load

Our goals – now

How we plan to use Ceph/Rados

- Ceph is still a hot option – but is it sufficient to just to introduce a new storage?
- We have to rethink not only the storage type, but also the storage usage!
- If a workload is suitable for Ceph/Rados, it might as well be suitable for further evolution (-> Cloud operations)
- Use Ceph/Rados for non-critical workloads (performance and business impact) in the legacy environment (to learn) and completely bet on it within OpenStack

OpenStack & Ceph - The Complete Picture

OpenStack & Ceph at BMW

The reason to use it

- Classic IT is highly optimized, stable and cost efficient, but it is not necessarily suitable for all new workloads (especially in terms of agility).
- New use cases do not only require new technology, but especially new operation and usage models.
- For these new “cloud” use cases we wanted to have
  - an open IaaS framework
  - with future proof concepts,
  - with industry standard APIs and functionality,
  - supported by a trusted partner.

OpenStack & Ceph at BMW

Our concept

- Classic IT and Private Cloud - based on OpenStack - are considered as separate worlds.
- Customers/projects have to fulfill prerequisites to get access to OpenStack:
  - Customers need the skill and mindset to operate and maintain their workloads properly – on their own.
  - Workloads have to be “cloud aligned”:
    - Real need for agility and dynamics, not just “forgot to order”
    - No critical dependency on single instance uptime
    - Automatic deployment of workloads, redeployment as base concept
    - Scale-out architecture instead of scale-up
- Customers gain more freedom, flexibility and agility.

OpenStack & Ceph at BMW

Topics to consider

- Additional supporting infrastructure
  - Install server (based on our internal “classic” deployment mechanism (for integration in regular processes))
  - Admin server (for OpenStack deployment – crowbar master)
  - Quorum monitor server (for storage in just two locations)
- DNS (for clients) – how to consolidate namespaces of cloud and classic world
- Naming conventions for physical hosts and VMs (leverage existing names)
- Process integration (prerequisites in terms of planning, accounting, blueprints, approvals, …)

OpenStack & Ceph at BMW

Our base environment

- 2 Control nodes
  - 16 physical cores (HT) / 256 GB memory
- 10 Compute nodes (to be extended soon)
  - 6x 16 physical cores (HT) / 256 GB memory
  - 4x 30 physical cores (HT) / 512 GB memory
- 8 Storage nodes (to be extended soon)
  - 16x physical cores (HT) / 256 GB memory
  - 10x 1.2 TB HDD + 2x 100 GB SSD (for journals)
- SUSE OpenStack Cloud 5
- SUSE Storage 1.0

OpenStack & Ceph at BMW

We are growing fast

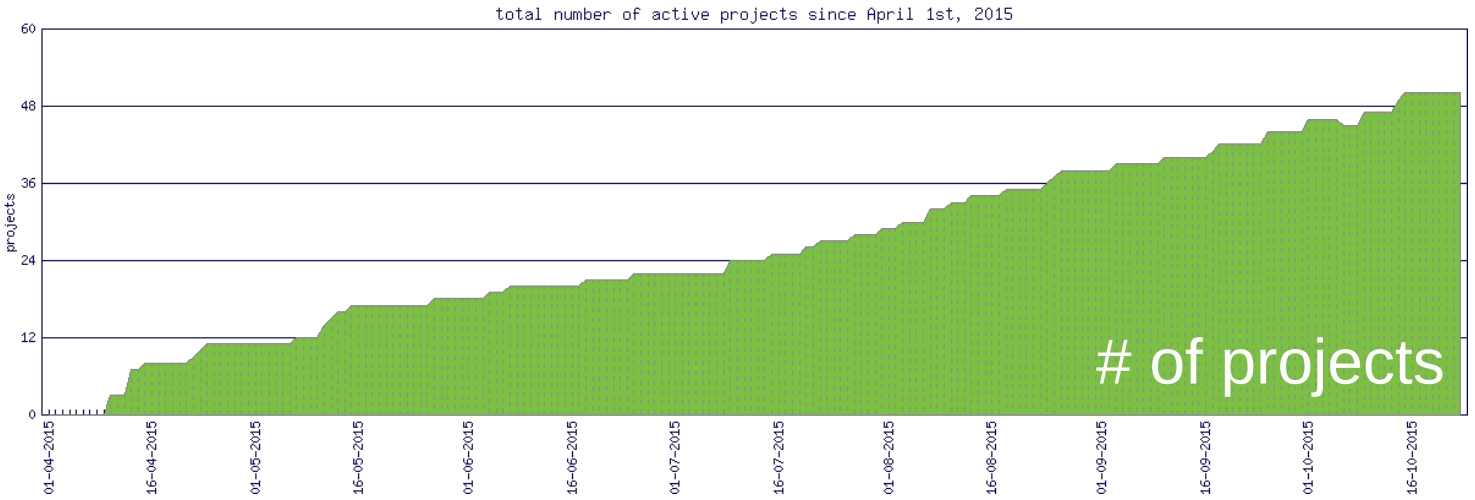

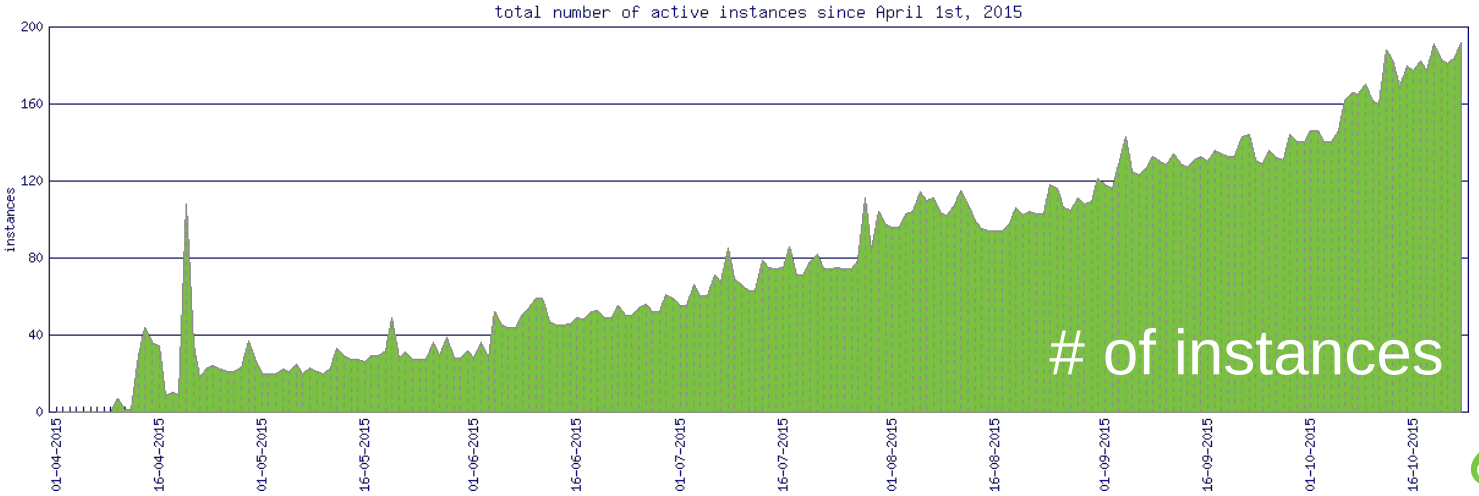

OpenStack & Ceph at BMW

Our workloads are evolving

Our assumptions in the beginning:

- Mainly used for Continuous Integration
- Several PoC's and evaluation systems

→ Simple flat networking, regular storage requirements

What we see now:

- Increasing number of Big Data and analytics
- More complex setups

→ More dynamic networking, extensive storage needs

OpenStack & Ceph at BMW

Our experiences so far

- You have to learn together with your customers
  - Start with a minimal viable product
  - Let them experience the borders and push you into the right direction
  - Let them participate in this development process (better identification with cloud environment)
- You always have to do a pre-check if your customer is cloud aware and cloud ready.
- Basic OpenStack services are mainly stable
  - Keystone, Neutron, Cinder, Glance, Nova
  - The installation mechanism had some “integration issues”

OpenStack & Ceph at BMW

Where our software partner helped us

- Great support on integrating OpenStack into our enterprise environment
- Functionalities of “Kilo” backported by SUSE:
  - The ability of being able to “disable” a certain image version
  - Multiple authentication backend support
  - Fixed Cinder/Ceph integration
- Functional improvements
  - Creation of an “Project Admin” role
  - Adjusted security concepts to meet enterprise needs (SSL for VNC due to flat provider network necessary)
- Good discussion partner for general concepts

OpenStack & Ceph at BMW

Our next steps

- Move on to SUSE OpenStack Cloud 6
- Increase the OpenStack infrastructure resources
  - CPU / Memory / Storage
- Improve the flexibility of network
  - Move from enterprise flat structure to something more dynamic (currently we are using Linux bridges)
- Introduce new functionality
  - Shared storage service
  - Backup offering
  - Extended reporting for customers

Questions?

Annotations?

Thank you.

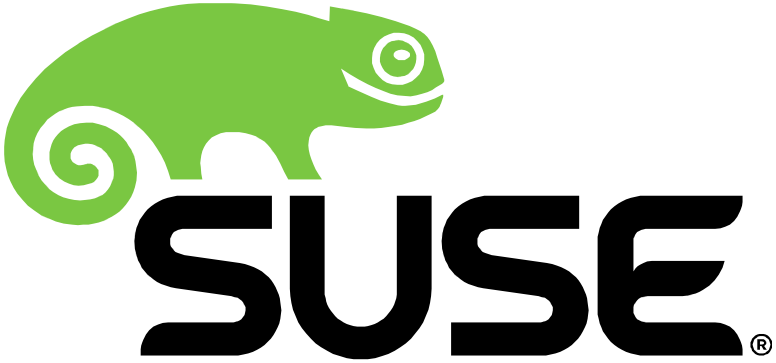
SUSE®
We adapt. You succeed.