Reliability and Performance for OpenStack through SmartOS

Daniele Stroppa, ZHAW
strp@zhaw.ch

MCN
Mobile Cloud Networking

Agenda

1. Mobile Cloud Networking
2. Why a Cloud OS
3. OpenStack
    1. Introduction
    2. Architecture
4. SmartOS
    1. Introduction
    2. Features
5. OpenStack on SmartOS
    1. Current status
    2. Further work

Mobile Cloud Networking

MCN
Mobile Cloud Networking

- Mobile Cloud Networking is a FP7 IP European research project.
- 18 top-tier partners from industry and academia.
- Mobile Cloud Networking running over the period of 36 month (2012 - 2015).
- The total investment is roughly 16M Euro.

Mobile Cloud Networking

MCN
Mobile Cloud Networking

- The MCN project investigates, implements and evaluates the technological foundations for the system of Mobile Network plus Decentralized Computing plus Smart Storage offered as one atomic service: On-Demand, Elastic and Pay-As-You-Go
- Mobile Cloud Networking objectives:
  - Extend concept of Cloud Computing beyond data centres towards mobile End-User
  - Design an 3GPP-compliant Mobile Cloud Networking Architecture that exploits and supports Cloud Computing
  - Design a novel virtualization layer, monitoring system, and general provisioning across the various domains within MCN.
  - Design a Mobile Platform for future MobileCloud-enabled services and applications.
  - Enable a novel business actor, the MCN Provider
  - Deliver and exploit concept of an End-to-End MCN for Novel Applications and Services

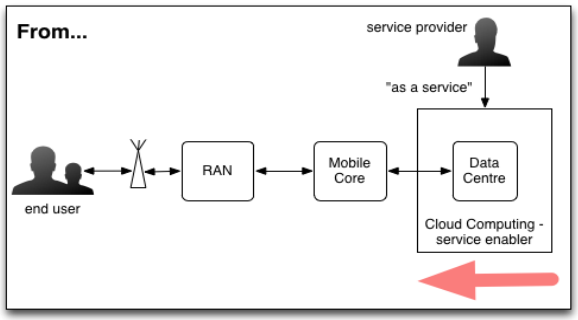

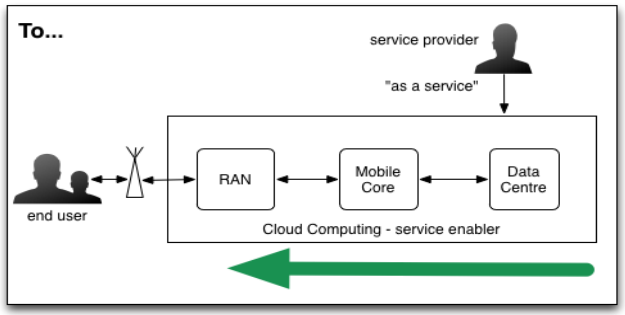

The need for a truly Cloud OS

MCN
Mobile Cloud Networking

- Cloud providers must be able to offer a single server to multiple users without them noticing that that they are not the only users of that machine. This means that the underlying operating system must:
  - be able to provision and de-provision, VMs in a very fast seamless way;
  - allocate physical resources efficiently and fairly amongst the users;
  - support multithreaded and multi-processor hardware;
  - be highly reliable;
  - must provide a way to quickly determine what's causing the system to misbehave.
- A customer of the cloud provider will also expect the server to be:
  - fast, i.e. the observed latency should be minimal;
  - flexible, i.e. get extra power when needed;
  - secure, meaning that neighboring users must not interfere with each other.

OpenStack introduction

MCN
Mobile Cloud Networking

- OpenStack is a Infrastructure-as-a-Service (IaaS) open source project
  - Offers control over large pools of compute, storage, and networking resources.
  - Easily managed through a dashboard that gives administrators control while empowering their users to provision resources through a web interface.
- OpenStack is designed to deliver a massively scalable platform
  - To achieve this, each of the constituent services are designed to work together to provide a complete Infrastructure-as-a-Service (IaaS).
- OpenStack is managed by the OpenStack Foundation
  - Supported by HP, IBM, RedHat, Rackspace, Dell, Intel, NEC, Ericsson, Alcatel-Lucent et al.

openstack™
CLOUD SOFTWARE

OpenStack Conceptual Architecture

MCN
Mobile Cloud Networking

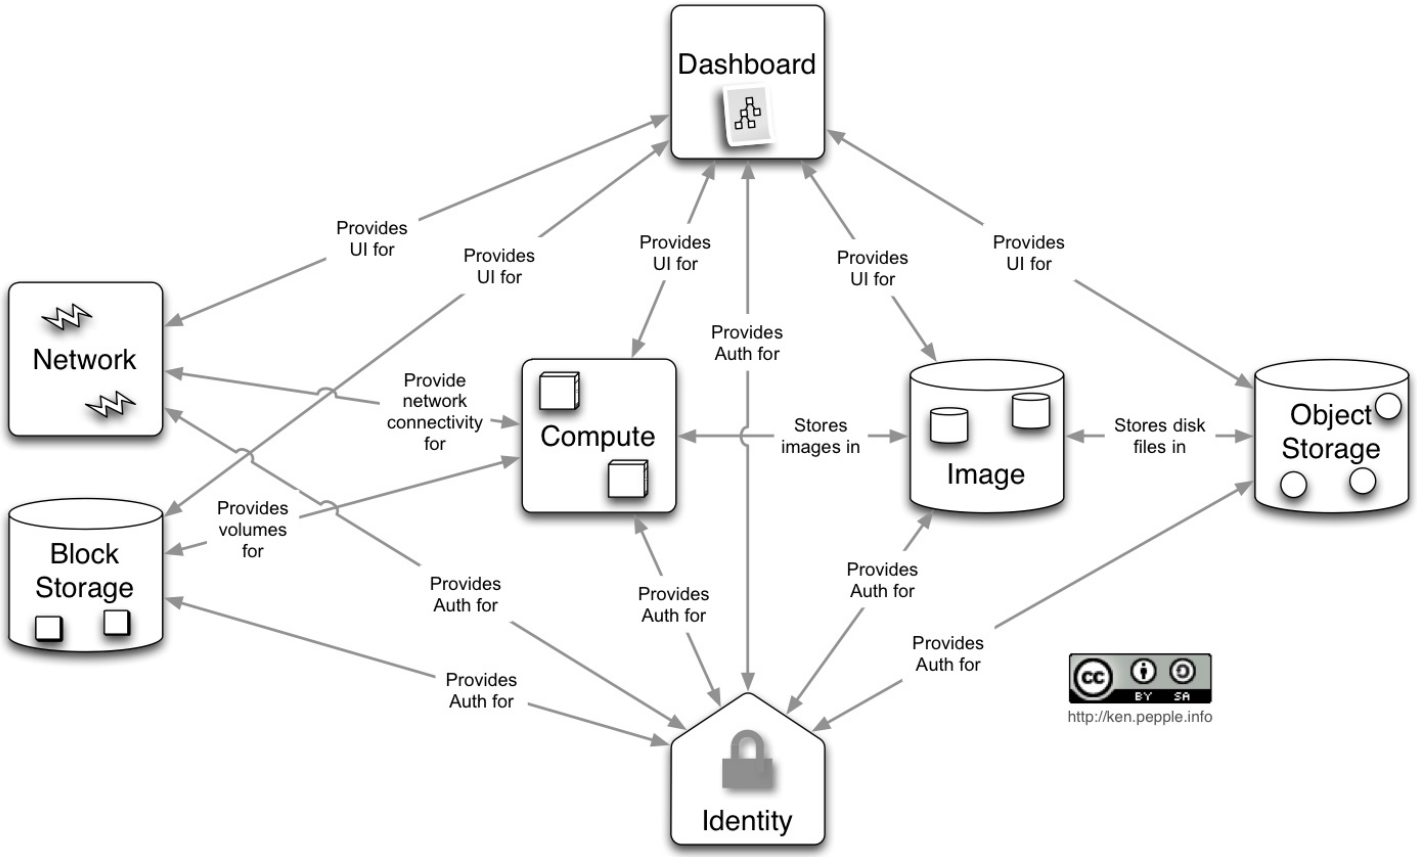

openstack™
CLOUD SOFTWARE

SmartOS introduction

MCN
Mobile Cloud Networking

- SmartOS is an open source type 1 hypervisor platform
  - Based on Illumos, a descendant of OpenSolaris, and developed by Joyent
- SmartOS is a live operating system,
  - It can be booted via PXE, USB or an ISO image
  - Runs entirely from memory, leaving the full space on the local disk to be used for virtual machines.
  - This type of architecture makes it very secure, easy to upgrade and recover

SmartOS

SmartOS features - Virtualization

MCN
Mobile Cloud Networking

- SmartOS offers two types of OS Virtualization:
  - Container-based virtualization
    - A container is the combination of resource controls and Solaris zones.
    - Provides a lightweight solution and a complete and secure user space environment on a single global kernel.
    - Can easily scale vertically, something that is more difficult with KVM virtualization.
  - KVM virtual machines
    - KVM virtual machines provide a solution for running a variety of guest OS's, including Linux and Windows, in a full, hardware-assisted virtualization.
    - KVM images on SmartOS run as a process inside of a zone.

SmartOS

SmartOS features - Storage

MCN
Mobile Cloud Networking

- ZFS
  - Combines file system and logical volume manager.
  - Fast file system creation - allowing to add new tenants almost instantly - and data integrity guarantee.
  - ZFS also includes storage pools - that simplify the task of expanding storage capacity - copy-on-write snapshot creation and snapshot cloning - allowing fast and reliable backup and cloning of virtual machines.
- Disk I/O throttling
  - Introduced to overcome a drawback in Solaris.
  - All zones/applications are ensured to get a reliable turn at reading/writing to disk.
  - When a system is under heavy load from multiple tenants disk I/O throttling comes into effect, while during quiet time, tenants can enjoy faster I/O without affecting other users.

SmartOS

SmartOS features - Networking

MCN
Mobile Cloud Networking

- Network Virtualization
  - SmartOS makes use of Crossbow to provide a network virtualization layer.
  - Fully integrated with the virtual machine administration tool of SmartOS, i.e. vmadm.
  - Allows each virtual machine can get up to 32 virtual network interfaces (VNICs).
- SmartOS inherits multipath IP (IPMP) from Solaris.
- SmartOS can leverage data link multi-pathing (DLMP), i.e. trunk aggregation.

SmartOS

SmartOS features - Observability

MCN
Mobile Cloud Networking

- DTrace is a performance analysis tool included by default in different operating system, amongst them Illumos and Solaris and therefore SmartOS.
- DTrace can instrument code by modifying a program after it has been loaded into memory.
- Not limited to use with user-space application, but can be used to inspect the OS kernel and device drivers.
- In SmartOS, DTrace can be used to analyze and troubleshoot issues across all zones in a server or within an entire datacenter.

SmartOS

SmartOS features - Reliability

MCN
Mobile Cloud Networking

- Fault Management Architecture (FMA)
  - Helps detect, report and diagnose any fault or defect that can occur on a SmartOS system.
- Service Management Facility (SMF)
  - Dependencies between services
  - Parallel starting
  - Automatic restart upon
  - Delegation of services to non-root users.
- Virtual Router Redundancy Protocol (VRRP)
  - Highly available load balancing.
  - Provides a means to implement hot-failover via virtual IP sharing.

SmartOS

OpenStack on SmartOS – Current Status

MCN
Mobile Cloud Networking

- Blueprinted: https://blueprints.launchpad.net/nova/+spec/smartos-support
- Working with latest OpenStack Grizzly release:
  - Nova-compute
  - Nova-network
- Code publicly available on Github: https://github.com/dstroppa/openstack-smartos-nova-grizzly

openstack
CLOUD SOFTWARE

SmartOS

OpenStack on SmartOS – Further work

MCN
Mobile Cloud Networking

- Integrate with OpenStack Quantum
  - Possibly integrate with OpenFlow controllers
- Make DTrace metrics accessible from OpenStack Ceilometer
  - Monitoring and billing
- Integrate ZFS features
  - VM cloning (snapshot)
  - Block-type storage

openstack
CLOUD SOFTWARE

SmartOS

THANK YOU!

MCN
Mobile Cloud Networking

© 2012-2015 MCN. All Rights Reserved