

agoda.com



OpenStack in a Large Deployment

Sharkrit Impat – IT Server Infrastructure

22 September, 2016

Overview

- **Background information:**
 - **About Agoda**
 - What is Agoda? How big is it?
 - **Where we come from.**
 - Technology Transition : Conventional Infrastructure → Virtualization → Private Cloud
 - Large Scale IT Projects , Demanding of server requests
- **OpenStack deployment at Agoda:**
 - What is OpenStack, Why we use it, and What is it not.
 - OpenStack in house implementation
 - Cloud Infrastructure Integration Tools
 - Physical bare metal deployment
 - VMs Deployment Techniques
 - Upgrade process
 - Current Statistics / Visibility / Reports
 - Your Opportunity

About Me:

Education:

- Bachelor's Degree Computer Science, Rajamangala University of Technology
- Electronics Technician, Military Technical Training School

Working Experiences:

- More than 13 years experiences on Infrastructure implementations and administrations.
- Passionate with virtualization and cloud technology.
- Specialize in HP hardware products and Microsoft technology.
- Joined Agoda since 2009, an 8 year journey with Agoda.
- In the last 4 years intensive experiences of Linux and Open Source Technology.
- Positive Thinking, Willing to learn new thing, Enjoy every journey, Self-Learning, Self-Starter, Team Player.

Recognitions and Certifications:

- Red Hat OpenStack Administration Training I & III (CL110, CL310)
- Public Speaker at Microsoft MVP IT Camp “Multi-Server Management Using Windows 2012 R2”
- Public Speaker at Windows Storage Day 2014 “Windows Scripting for Manage Storages”
- Microsoft Certified Technology Specialist, Business Desktop Deployment Planning Services

About Agoda

เว็บไซต์ของอโกดามีความรวดเร็ว ง่ายต่อการใช้งาน

Agoda.com คืออะไร | ข่าวประชาสัมพันธ์ | ติดต่อเรา | ข้อมูลพันธมิตรโรงแรมและลี้กอิน | พันธมิตรบริษัทค้า

เกี่ยวกับอโกด้า 



และใช้เทคโนโลยีระดับโลกจนทำให้ได้รับรางวัลมาแล้ว

อโกด้า (www.agoda.com) เป็นหนึ่งบริษัทผู้ให้บริการเว็บไซต์จองห้องพักในโรงแรมที่เติบโตเร็วที่สุดในโลก โดยมีโรงแรมในเครือข่ายกว่า 100,000 แห่งและมีเว็บไซต์ที่ได้รับการแปลเป็นภาษาต่างๆ ถึง 38 ภาษา Agoda Company Pte. Ltd. บริษัทซึ่งอยู่เบื้องหลังการดำเนินงานของ Agoda.com ก่อตั้งขึ้นในปี พ.ศ. 2548 โดยสองผู้คร่ำหวอดในวงการธุรกิจท่องเที่ยวออนไลน์ ในปี 2007 อโกด้าได้กลายเป็นส่วนหนึ่งของเครือไพรซ์ไลน์กรุป (Priceline Group) บริษัทผู้ให้บริการจองห้องพักทางออนไลน์ที่ใหญ่ที่สุด ทั้งนี้ หุ้นของไพรซ์ไลน์กรุปมีการซื้อขายในตลาดหลักทรัพย์ NASDAQ (Nasdaq: PCLN) และเป็นส่วนหนึ่งของดัชนี S&P 500

อโกด้ามียุทธศาสตร์ที่ขยายฐานการท่องเที่ยวมากกว่า 2,000 คนจากประเทศต่างๆ มากกว่า 20 ประเทศ โดยมีสำนักงานตั้งอยู่ในสิงคโปร์ กรุงเทพฯ กัวลาลัมเปอร์ โตเกียว ฮิดันย์ ฮองกง และบูดาเปสต์ รวมถึงเมืองสำคัญต่างๆ ทั่วเอเชีย แอฟริกา ตะวันออกกลาง ยุโรป และอเมริกา ผู้จัดการบริหารกลุ่มลูกค้าที่ทุ่มเทของเรารักษาความสัมพันธ์อันดีกับโรงแรมพันธมิตรของอโกด้าทั่วโลก พร้อมทั้งสร้างสรรค์โปรแกรมพิเศษและแผนการตลาดเพื่อช่วยให้ลูกค้าสามารถมอบข้อเสนอที่ดีที่สุดทางอินเทอร์เน็ตให้แก่ลูกค้า นอกจากนี้ ยังมีนโยบายการันตีราคาที่ดีที่สุดซึ่งช่วยคอยสนับสนุนความได้เปรียบที่กล่าวมาข้างต้น

เว็บไซต์ของอโกด้ามุ่งเน้นความรวดเร็ว ง่ายต่อการใช้งาน และใช้เทคโนโลยีระดับโลกจนทำให้ได้รับรางวัลมาแล้ว เราให้บริการยืนยันการจองห้องพักทันทีสำหรับโรงแรมกว่าแสนแห่งทั่วโลก นอกเหนือจากประเภทสถานที่พักและห้องพักที่หลากหลายแล้ว อโกด้ายังมีรีวิวโรงแรมที่มาจากลูกค้าตัวจริง 100 เปอร์เซ็นต์ ปัจจุบัน เรามีรีวิวซึ่งลูกค้าส่งมาให้เราหลังจากเข้าพักที่โรงแรมแล้วกว่าหลายล้านรีวิว สุดท้ายนี้ ฝ่ายลูกค้าสัมพันธ์ของเรายินดีให้ความช่วยเหลือลูกค้าตลอดทุกวัน 24 ชั่วโมง โดยให้บริการในภาษาต่างๆ มากมายหลายภาษา ท่านจึงมั่นใจได้ถึงการสนับสนุนที่รวดเร็ว

อโกด้ามุ่งมั่นที่จะเป็นบริษัทที่น่าเชื่อถือ เข้าถึงได้และมีเครือข่ายพันธมิตรมากมาย เราแนะนำโรงแรมที่หลากหลายและมุ่งมั่นที่จะมอบราคาที่ดีที่สุดให้ลูกค้าอยู่เสมอ โดยมุ่งหวังที่จะช่วยให้ทุกคนเดินทางท่องเที่ยวได้ง่ายขึ้นในราคาที่ไม่แพง พร้อมทั้งสนุกสนานกับการเดินทางได้มากยิ่งขึ้น

AGODA® เป็นเครื่องหมายการค้าจดทะเบียนของ AGIP LLC ภายใต้การอนุมัติจาก Agoda Company Pte. Ltd.



 agoda

 priceline.com

 KAYAK

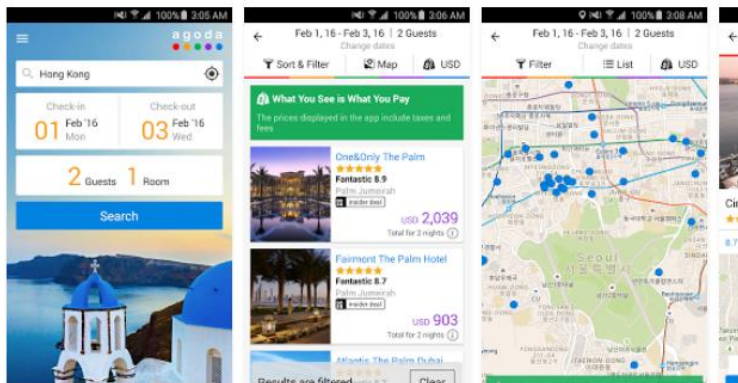

 Rentalcars.com

 Booking.com

 OpenTable



Agoda Apps on Mobile – Easy to use, Fast and Secure



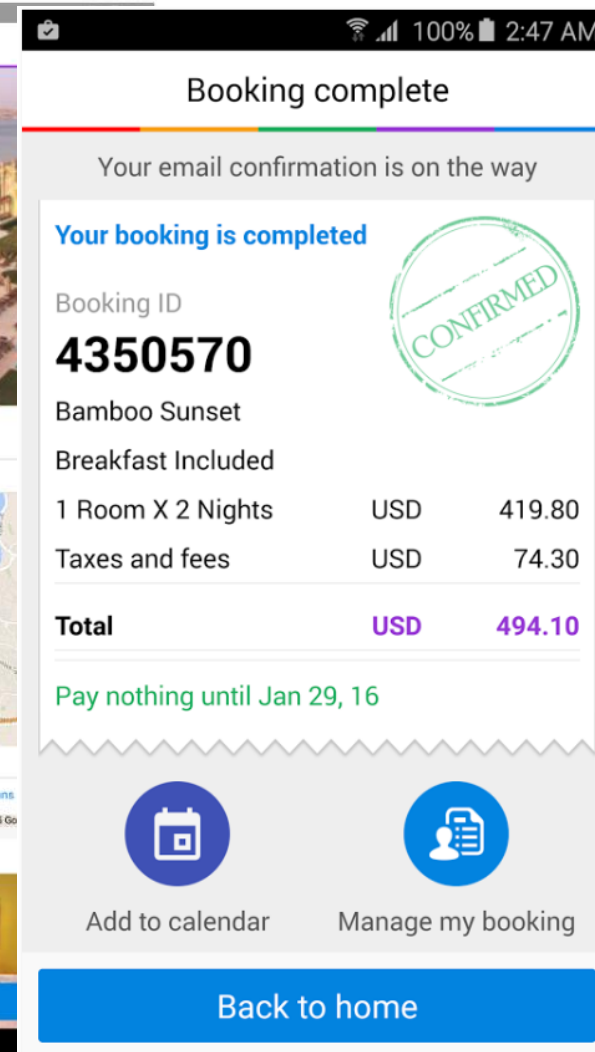
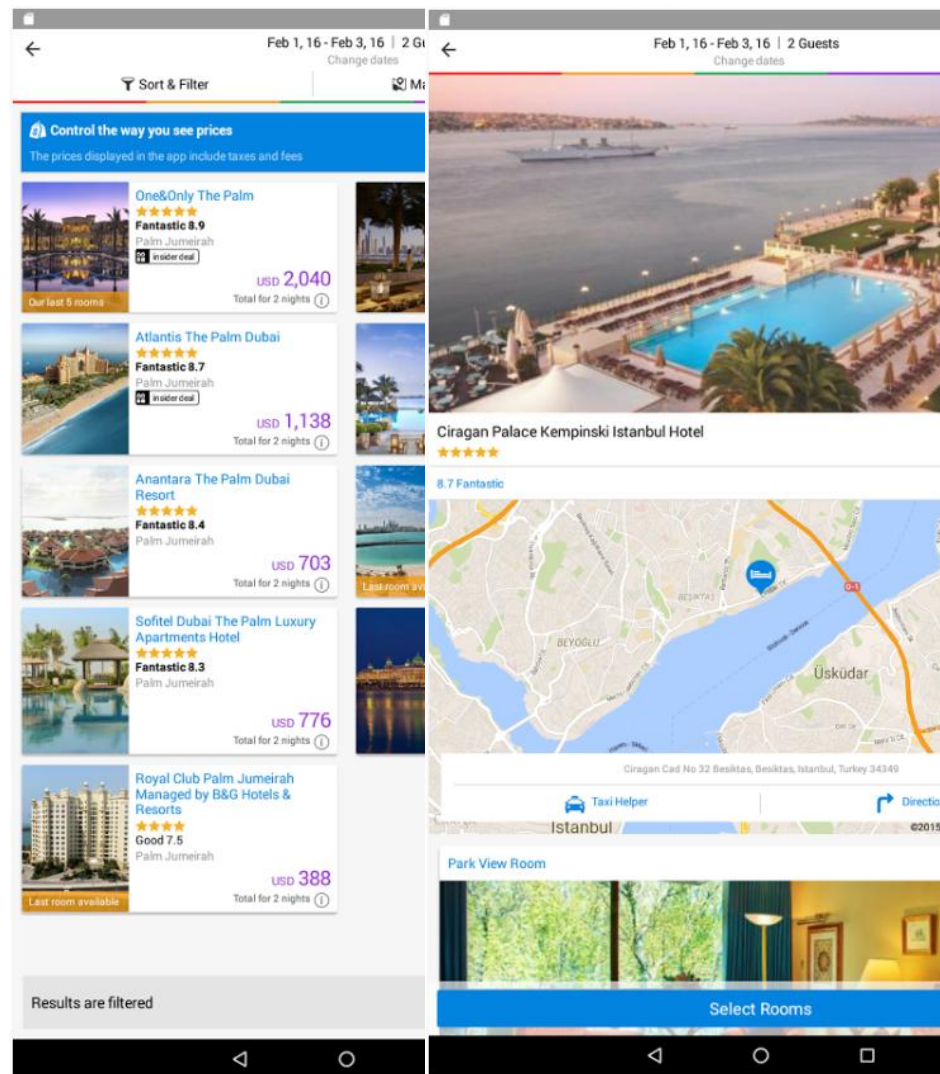
Whether you need a last minute room for tonight or are planning your next holiday, finding and booking the best deals on hotels and other accommodations is fast and easy with the new Agoda app.

Plan and book your trip anytime, anywhere:

- SEARCH. Find hotels & accommodations near you right now or search by city, landmark or property name. Use the map to compare prices based on location.
 - DECIDE. Use our filters, high quality photos and 9+ million verified traveler reviews to select the perfect place to stay for your budget and style.
 - SAVE. Get additional savings and promotions exclusive to the app.
 - BOOK. Make your reservation quickly & securely in only a few taps with your saved information.
- Access all of your booking information for printer-less planning and seamless check-in.

All backed with the Agoda advantage:

- EVERYWHERE. Choose what kind of accommodation fits your needs from our selection of more



Who are Agoda?

- Top Hotels booking website in Asia
- More than 2,000 staff, HQ in Bangkok, presence in over 20 countries.
- More than 400,000 hotels on website, translated into 38 languages.
- More than 30 Millions customers visit our website per month.
- More than 30,000 Agoda apps download per month.
- Event messaging framework generates > 20K events per second
- Customer Service 24x7 and NOC 24x7
- 5 Datacenters in 3 Continents, contain 154 racks. More than 2,500+ physical servers, 4,000+ VMs
- 5 OpenStack Production Clusters: Total of vCPU: 36,000+ cores, RAM: 61TB, Storage: 700+TB
- 2 OpenStack Q/A and CI Clusters: Total of vCPU: 2300+ cores, RAM: 6+TB, Storage: 35+TB

Where we come from

Conventional Infrastructure...

- must be solid, stable, reliable... “unbreakable”
- is not normally agile
- needs a large shared storage, IOPS congestion, space limitation, and expensive.

Challenge:

- Implement a flexible infrastructure that supports a constantly changing application portfolio
- Continue to provide traffic routing flexibility (capable of absorbing traffic from other datacenters transparently)

Where we come from

Operating System

- Shift from Windows to Linux
 - Before Windows 90% : Linux 10%
 - Now Windows 30% : Linux 70%

Server deployments

- Installed manually
- Administered manually

Application deployments

- Installed manually
- Configured manually
- Administered manually

Server Management

- Centralized
- Inventory in Spreadsheets

Manual installation/configuration is not scalable

- It's ok for 20 servers
- It's not ok for 1,000's of servers across multiple continents

Server requests ~ 20x machines (in 5x Datacenters)

Ad hoc requests, Tiny requests, generic orders

Where we come from

Large Scale IT Projects , Demanding of server requests

- Frontend services (www, mobile APIs)
- Partner services (XML, RESTful-APIs)
- Caching Farm (memcache, session cache, content cache)
- Hotels prices & property search
- Big Data, Data mining, Machine Learning
- Log management, Graphing
- Automation Patch processing, Automate Deployment
- Centralized Server Management, Monitoring


Where we come from

Technology Transition

- **1st Generation - Conventional Infrastructure – build to sustain the load & HA**
 - Physical server dedicated for single role
 - SAN connections
 - Static Network
- **2nd Generation - Virtualization Infrastructure – better resources utilization**
 - One Physical server service for many roles.
 - Better resource utilizations
 - SAN connections still in use
 - Hardware independency
 - Implement TOR Switches
- **3rd Generation – Cloud Infrastructure – Infrastructure as a Service**
 - Scalable, Flexible, and Elasticity
 - Quick delivery & Self service
 - Programmability, API support for Automation, Infrastructure as a code
 - Replace expensive SAN with Software-Defined Storage

Where we come from

“Blast from the past” I don’t see these stuff for long time.




As well as these rack layout

[illegible]

Rack Layout Design - Scalable and Simple

- **Old Rack Layout Design**

[illegible]


- **New Rack Layout Design**

OpenStack Rack		OpenStack Rack		Big Data & Database Rack		Big Data & Database Rack	
Pos	Device	Pos	Device	Pos	Device	Pos	Device
46		46		46		46	
45		45		45		45	
44	OOB Switch 16	44	OOB Switch 16	44	OOB Switch 16	44	OOB Switch 16
43		43		43		43	
42	TOR Switch Arista 100	42	TOR Switch Arista 100	42	TOR Switch Arista 100	42	TOR Switch Arista 100
41		41		41		41	
40	TOR Switch Arista 100	40	TOR Switch Arista 100	40	TOR Switch Arista 100	40	TOR Switch Arista 100
39		39		39		39	
38		38		38		38	
37	OPENSTACK COMPUTE	37	OPENSTACK COMPUTE	37		37	
36	OPENSTACK COMPUTE	36	OPENSTACK COMPUTE	36	JBOJ	36	JBOJ
35	OPENSTACK COMPUTE	35	OPENSTACK COMPUTE	35		35	
34	OPENSTACK COMPUTE	34	OPENSTACK COMPUTE	34	DATABASE SERVER	34	DATABASE SERVER
33	OPENSTACK COMPUTE	33	OPENSTACK COMPUTE	33		33	
32	OPENSTACK COMPUTE	32	OPENSTACK COMPUTE	32	JBOJ	32	JBOJ
31	OPENSTACK COMPUTE	31	OPENSTACK COMPUTE	31		31	
30	OPENSTACK COMPUTE	30	OPENSTACK COMPUTE	30	DATABASE SERVER	30	DATABASE SERVER
29	OPENSTACK COMPUTE	29	OPENSTACK COMPUTE	29		29	
28	OPENSTACK COMPUTE	28	OPENSTACK COMPUTE	28	HADOOP_DATA_NODE	28	HADOOP_DATA_NODE
27	OPENSTACK COMPUTE	27	OPENSTACK COMPUTE	27		27	
26	OPENSTACK COMPUTE	26	OPENSTACK COMPUTE	26	HADOOP DATA NODE	26	HADOOP DATA NODE
25	OPENSTACK COMPUTE	25	OPENSTACK COMPUTE	25		25	
24	OPENSTACK COMPUTE	24	OPENSTACK COMPUTE	24	HADOOP DATA NODE	24	HADOOP DATA NODE
23	OPENSTACK COMPUTE	23	OPENSTACK COMPUTE	23		23	
22	OPENSTACK COMPUTE	22	OPENSTACK COMPUTE	22	HADOOP DATA NODE	22	HADOOP DATA NODE
21	OPENSTACK COMPUTE	21	OPENSTACK COMPUTE	21		21	
20	OPENSTACK COMPUTE	20	OPENSTACK COMPUTE	20	HADOOP DATA NODE	20	HADOOP DATA NODE
19	OPENSTACK COMPUTE	19	OPENSTACK COMPUTE	19		19	
18	OPENSTACK COMPUTE	18	OPENSTACK COMPUTE	18	HADOOP DATA NODE	18	HADOOP DATA NODE
17	OPENSTACK COMPUTE	17	OPENSTACK COMPUTE	17		17	
16	OPENSTACK COMPUTE	16	OPENSTACK COMPUTE	16	HADOOP DATA NODE	16	HADOOP DATA NODE
15	OPENSTACK COMPUTE	15	OPENSTACK COMPUTE	15		15	
14	OPENSTACK COMPUTE	14	OPENSTACK COMPUTE	14	HADOOP DATA NODE	14	HADOOP DATA NODE
13	OPENSTACK COMPUTE	13	OPENSTACK COMPUTE	13		13	
12	OPENSTACK COMPUTE	12	OPENSTACK COMPUTE	12	HADOOP DATA NODE	12	HADOOP DATA NODE
11	OPENSTACK COMPUTE	11	OPENSTACK COMPUTE	11		11	
10	OPENSTACK COMPUTE	10	OPENSTACK COMPUTE	10	HADOOP DATA NODE	10	HADOOP DATA NODE
9	OPENSTACK COMPUTE	9	OPENSTACK COMPUTE	9		9	
8	OPENSTACK COMPUTE	8	OPENSTACK COMPUTE	8	HADOOP DATA NODE	8	HADOOP DATA NODE
7	OPENSTACK COMPUTE	7	OPENSTACK COMPUTE	7		7	
6	OPENSTACK COMPUTE	6	OPENSTACK COMPUTE	6	HADOOP_DATA_NODE	6	HADOOP_DATA_NODE
5	OPENSTACK COMPUTE	5	OPENSTACK COMPUTE	5		5	
4	OPENSTACK COMPUTE	4	OPENSTACK COMPUTE	4	HADOOP DATA NODE	4	HADOOP DATA NODE
3	OPENSTACK COMPUTE	3	OPENSTACK COMPUTE	3		3	
2	OPENSTACK COMPUTE	2	OPENSTACK COMPUTE	2	HADOOP_DATA_NODE	2	HADOOP_DATA_NODE
1	OPENSTACK COMPUTE	1	OPENSTACK COMPUTE	1		1	

OpenStack deployment at Agoda:

What is OpenStack

- Infrastructure as a service (IaaS)
- Infrastructure by API

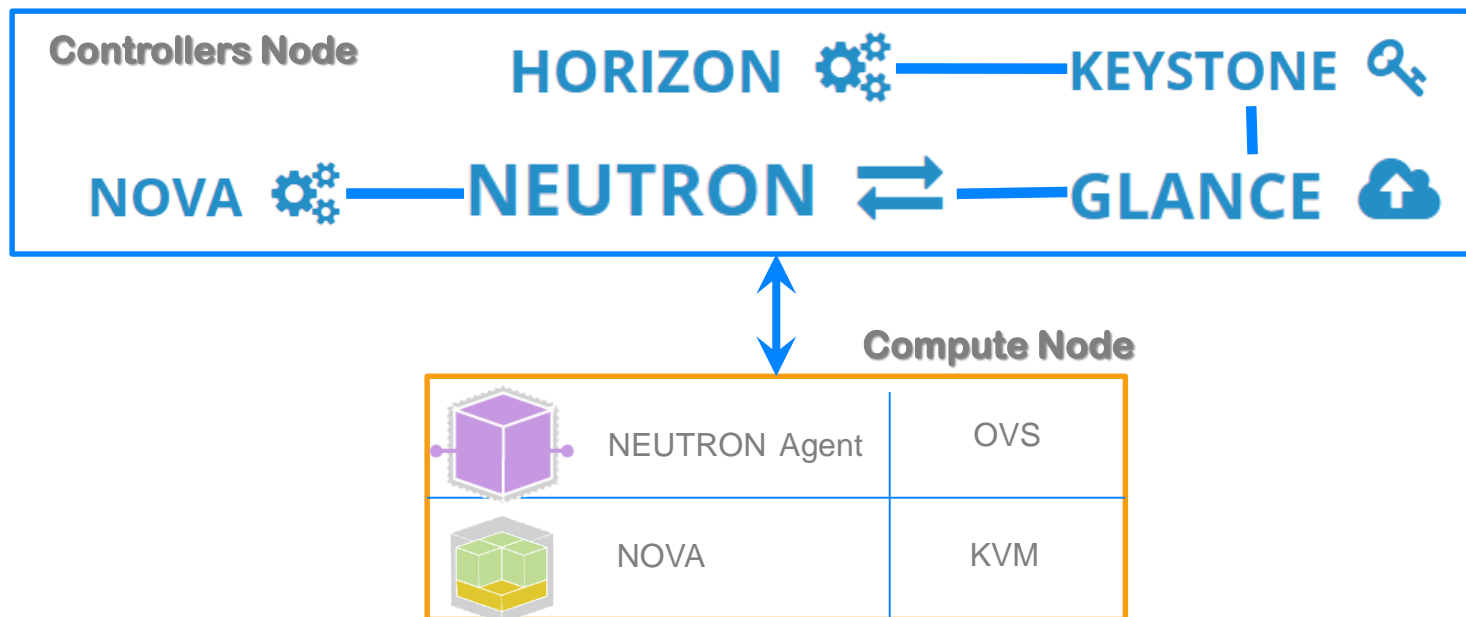


Why OpenStack

- Open Source and no licensing cost
- Flexibility (Vendor)
- Customizable
- Active Community
- Mailing list
- Bug tracing
- Industry Standard
- OpenStack Foundation
- HP, Intel, Cisco
- Development driven by the user
- Proven
- Used in some of the worldwide largest clouds
- The next mainstream cloud deployment

OpenStack deployment at Agoda:

- Simplicity is the key : Simple design, easy to understand and easy to maintain
- First spawned production Instance: 2014-09-23 14:16:16
- Using multi FlatNetwork
- Started with Havana with CentOS 6.x
- Currently deploying on Juno release (working on the upgrade to Mitaka)
- Hypervisor: KVM - Open source and very stable
- Open vSwitch: Opensource, stable and many vendor plugins support
- 5x separate production clouds
- 3 Continents
- 1x Continuous Integration cloud / QA
- In house Development
 - Hooks for DNS entry
 - Developers use APIs and Vagrant
 - Our own dashboard



OpenStack deployment at Agoda:

OpenStack Networking - Multi-FlatNetwork

- **Linux bridge qbr...**

The Linux bridge is needed to apply firewall rules to the instances. To block or allow traffic OpenStack uses iptables. Currently openvswitch doesn't support iptables, therefore a linux bridge is a workaround for this.

- **Openvswitch br-int**

The openvswitch br-int provides the switch for internal communication. If there are multiple instances in one compute node the traffic between this instances is managed by br-int.

- **Openvswitch br-bond0**

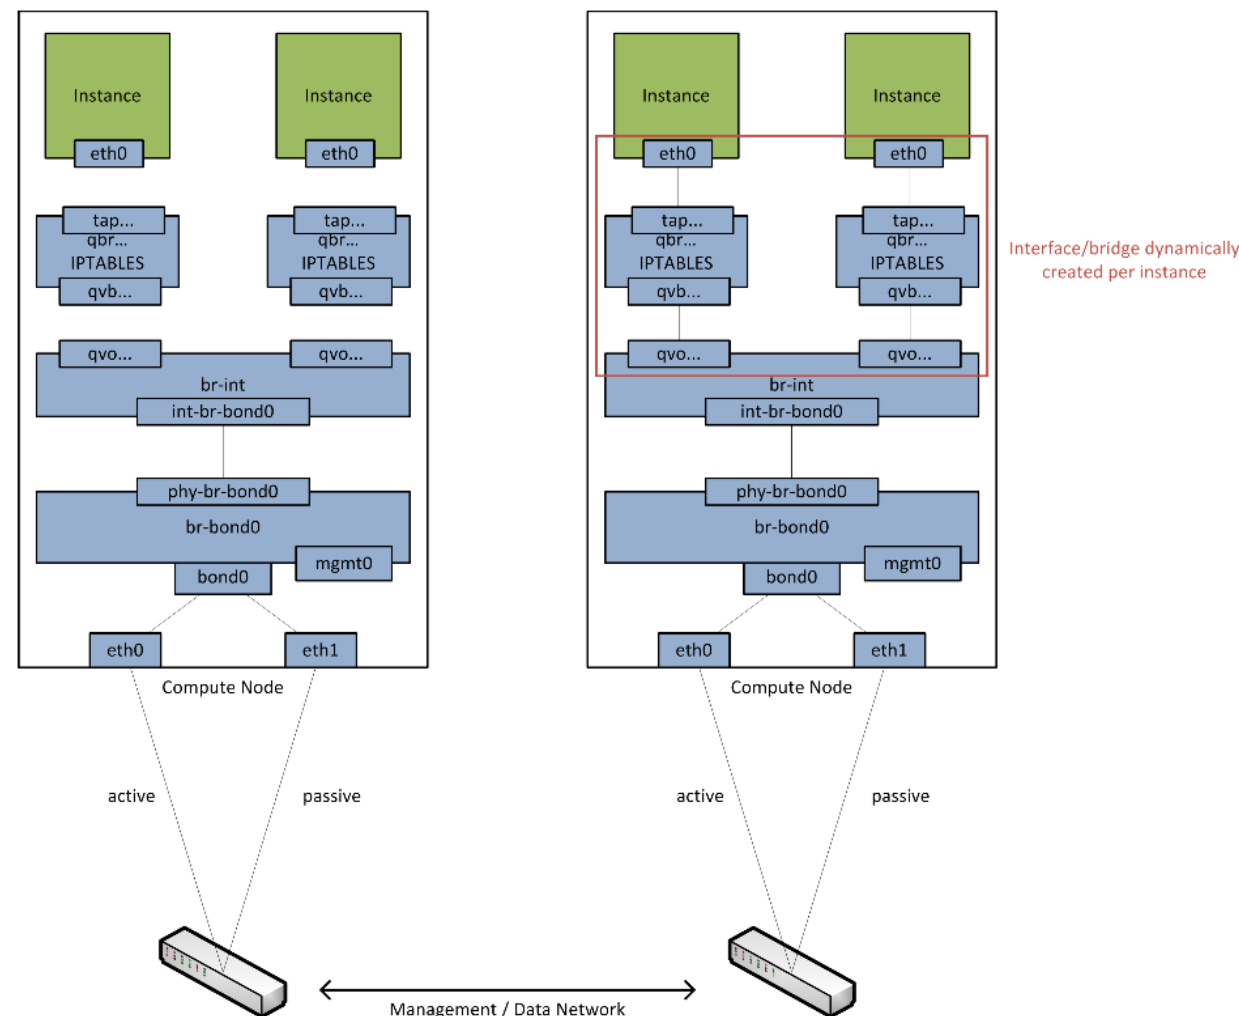
The openvswitch br-bond0 provides connectivity to the bonding interface bond0. It connects to br-int by the internal ovs ports int-br-bond0 and phy-br-bond0.

- **Instance and the interfaces**

For each new instance on a compute node create a new pair of eth0 - tap... and qvb... - qvo... interfaces.

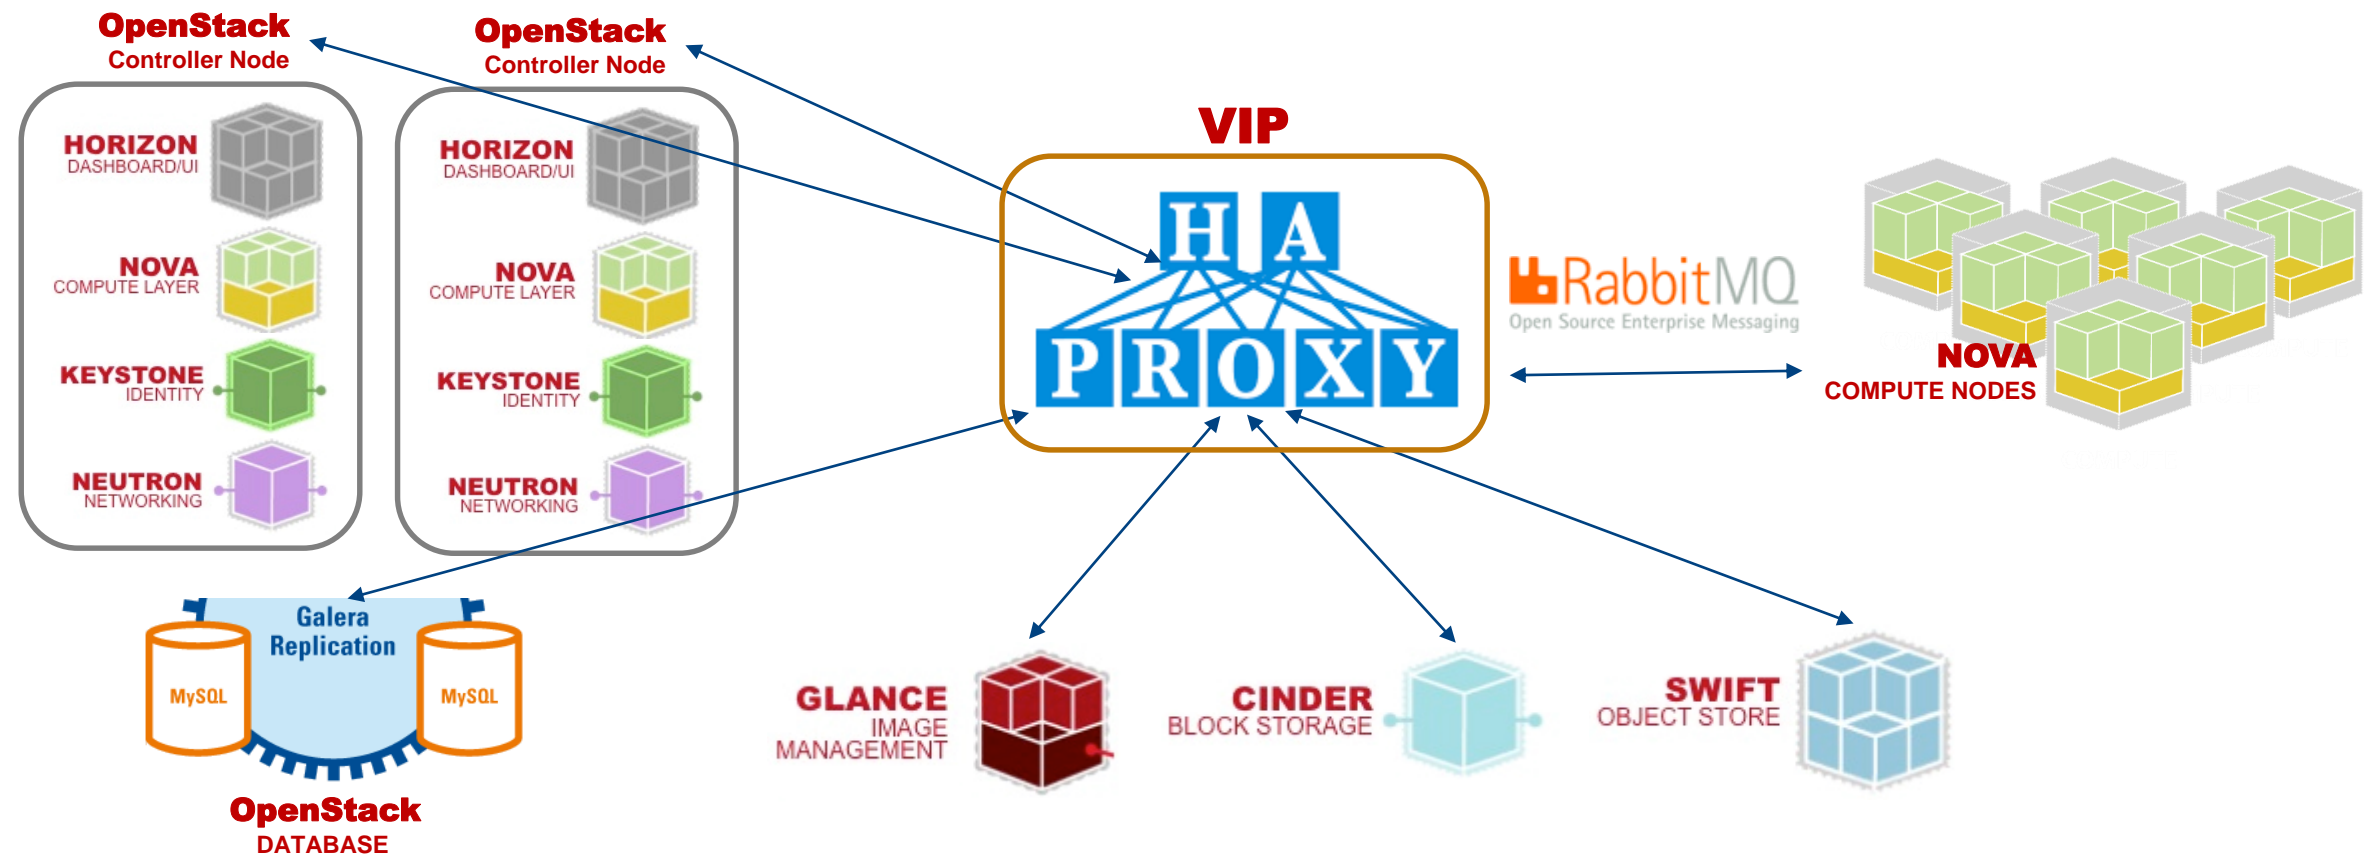
- **Compute Node management interface mgmt0**

Each compute node has a Open-vSwitch port on br-bond0 which is used as the management interface for the compute node. The interface has the local IP address which the DNS from the compute note is pointing at.



OpenStack deployment at Agoda:

OpenStack HA Architecture



OpenStack deployment at Agoda: HAProxy Load balancer Stats

HAProxy version 1.5.14, released 2015/07/02

Statistics Report for pid 25414

➤ General project information

```
pid = 25414 (process #), utime = 0
uptime = 0:02.00.00.00
system limit: core size = unlimited, ulimit-c = 4096
memstack = 4096K, maxproc = 20000, maxpipe = 0
current core = 32K, current pipe = 32K, core rate = 1 time
Running task: 1/004; idn = 100 %
```

active UP
 active UP, going down
 active DOWN, going up
 active or backup DOWN
 active or backup DOWN for maintenance (FANET)
 active or backup SOFT STOPPED for maintenance
 Note: INCLUSTDOWN = UP without load-balancing disabled
 backup UP
 backup UP, going down
 backup DOWN, going up
 not checked

- Scope:
- Hide DOMAIN servers
- Refresh now
- CSV report

- [Primary site](#)
- [Updates v1.2](#)
- [Online manual](#)

[illegible]

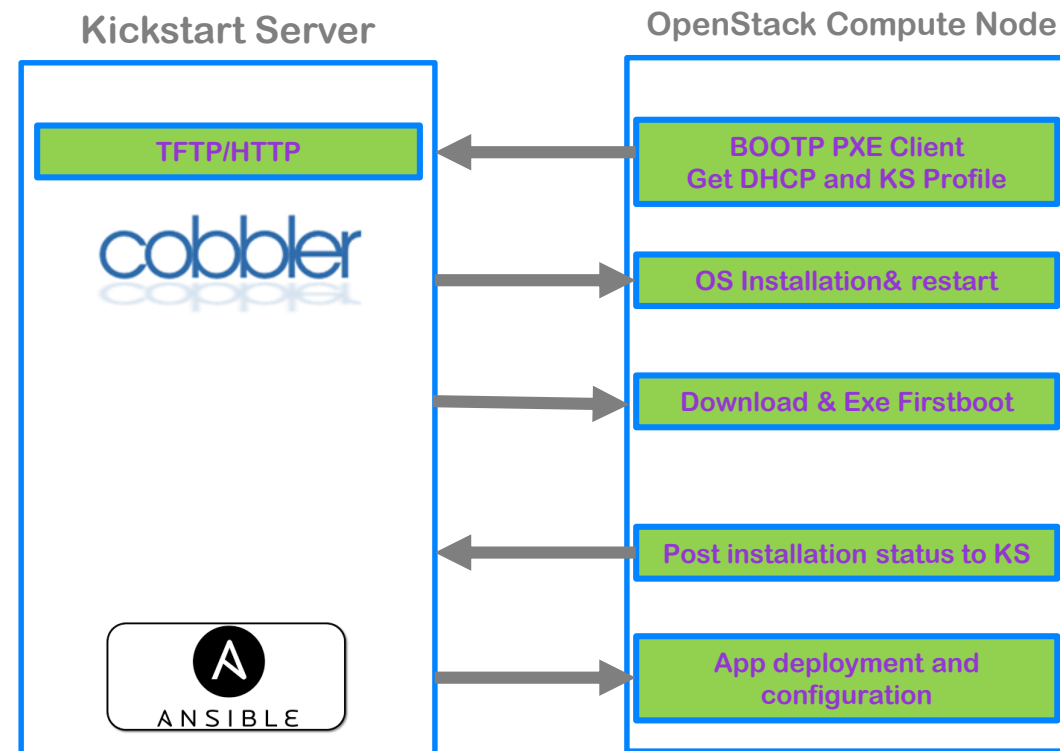
OpenStack deployment at Agoda:

- **OS/Compute Nodes Deployment**
 - **Most frequent task**
 - **Same process is used for our controller plane deployment and other projects.**
 - **Own repository**
 - **Packages management/versioning**
 - **Fast when deploy OS and packages, everything is local.**
 - **Consistency: Replication to all DCs around the world**
 - **Automation is the way to go**
 - **To cover rapid deployment.**
 - **Good for repetitive tasks.**
 - **Less error with less human interaction**

OpenStack deployment at Agoda:

Baremetal Server Deployment

- **Rack & Go:**
 - Out-of-Band gets IP from DHCP
 - Hardware Discovery & Inventory
- **Kickstart:**
 - Create DNS record
 - Configure machine's profile and disk layout
 - Pxe boot and OS installation
- **Firstboot: shell/bash script**
 - Install all the packages and tools
 - Join our domains
 - Configure kernel parameter
 - Install and configure apps specifically to hardware vendor
 - Network test
 - Firmware upgrade
- **Ansible: deploy configure apps, not limit to OpenStack packages**
 - Deploy and configure OpenStack packages: Nova, KVM, Neutron and Open vSwitch



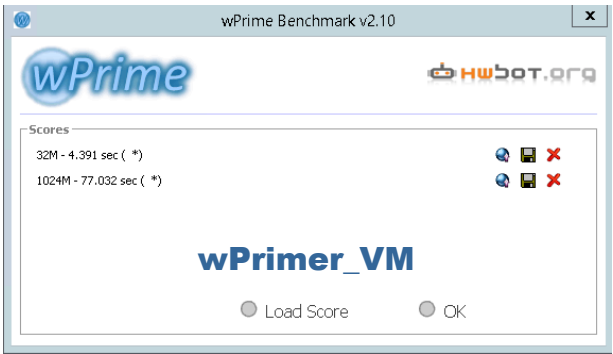
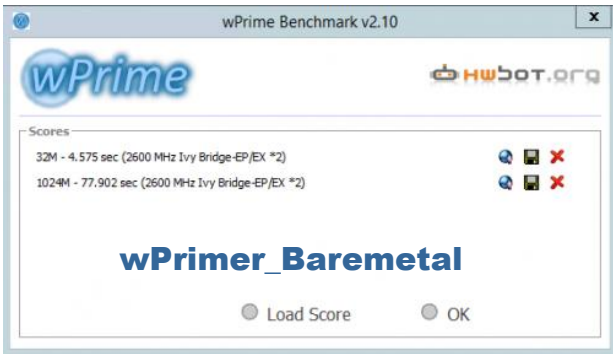
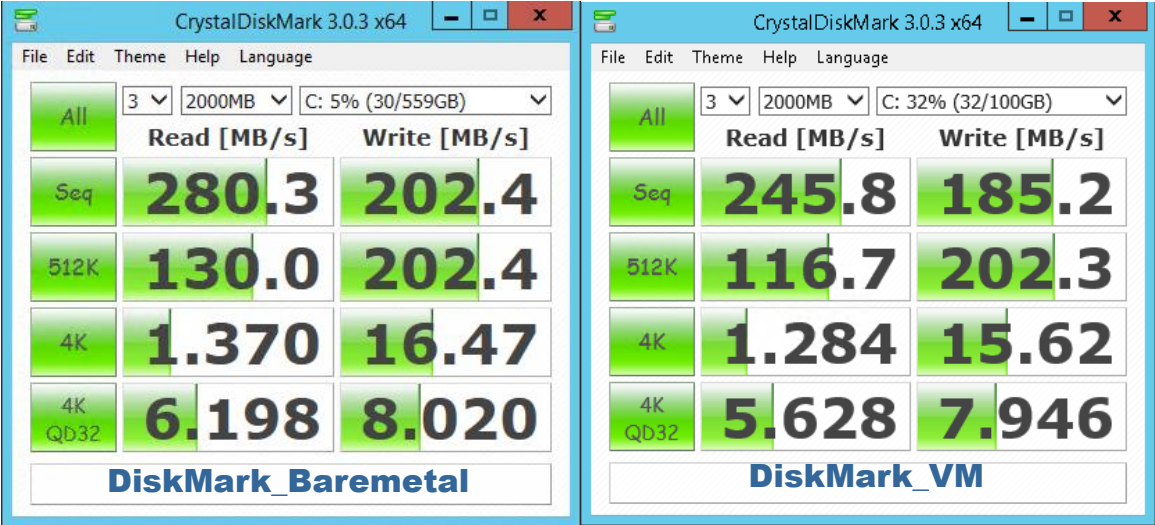
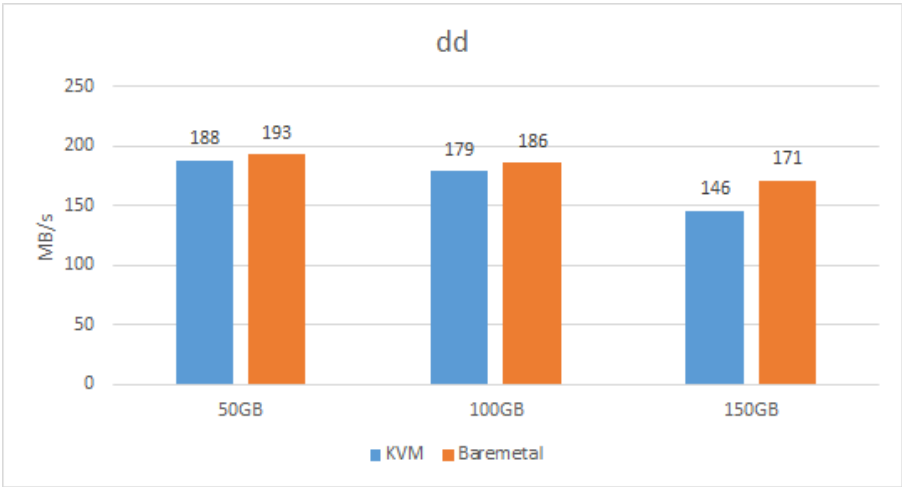

OpenStack deployment at Agoda:

VM Deployment

- **Linux as a Guest:**
 - Cloud-init
 - Get IP address
 - Configure hostname
- **Windows as a Guest**
 - VirtIO driver requires during setup Windows Images
 - Cloudbase-init (<https://cloudbase.it/cloudbase-init>)
 - Support Sysprep with plugins injection
- **Firstboot:** same as above with condition check for VMs

OpenStack deployment at Agoda:


OpenStack BenchMark



OpenStack deployment at Agoda:

OpenStack upgrade process


- Icehouse to Juno, at the same time upgrade compute nodes from CentOS6 to 7
 - Require minimal service interruption
 - Live migrate VMs from old compute nodes to a new one
 - Many issues with live-migration
 - Wrote wrap around script to handle all the bugs and error
 - Issues: Few pings drop while switching over, hardware incompatibility, memory change too often and KVM incompatibility.
- Reinstall the compute node with new OS and packages
- All are fully automated
- 100+ nodes per week



OpenStack deployment at Agoda:

Current Statistics / Visibility / Reports

- 5 OpenStack Production Clusters: Total of vCPU: 36,000+ cores, RAM: 61TB, Storage: 700+TB
- 2 OpenStack Q/A and CI Clusters: Total of vCPU: 2300+ cores, RAM: 6+TB, Storage: 35+TB



OpenStack deployment at Agoda:

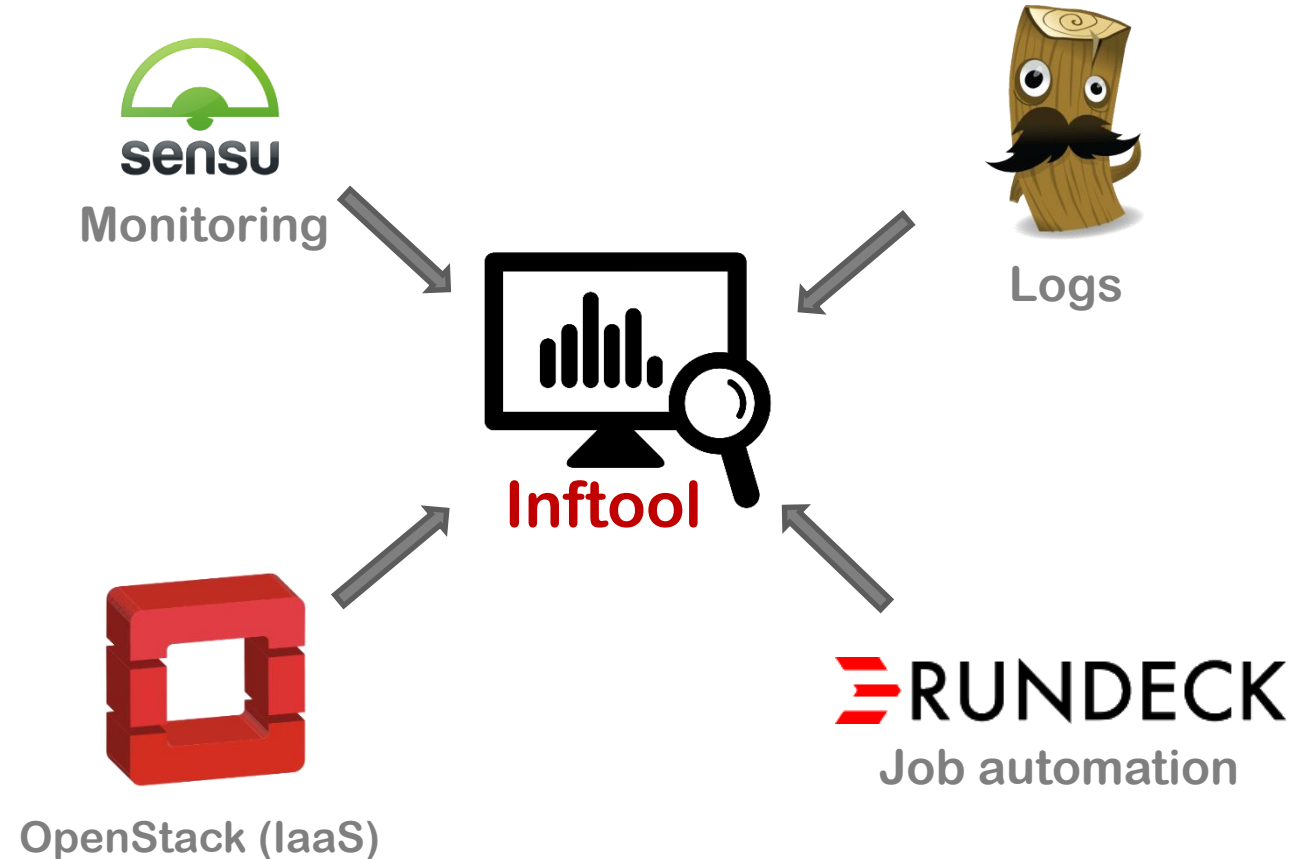
Cloud infrastructure integration

- **Integration in default infrastructure services needed**
 - DNS
 - Log management
 - Monitoring
 - Patching & maintenance
- **Independent API endpoint**
- **No default integration with existing services**
- **Tooling & process management needed**
- **Extensive system inventory**

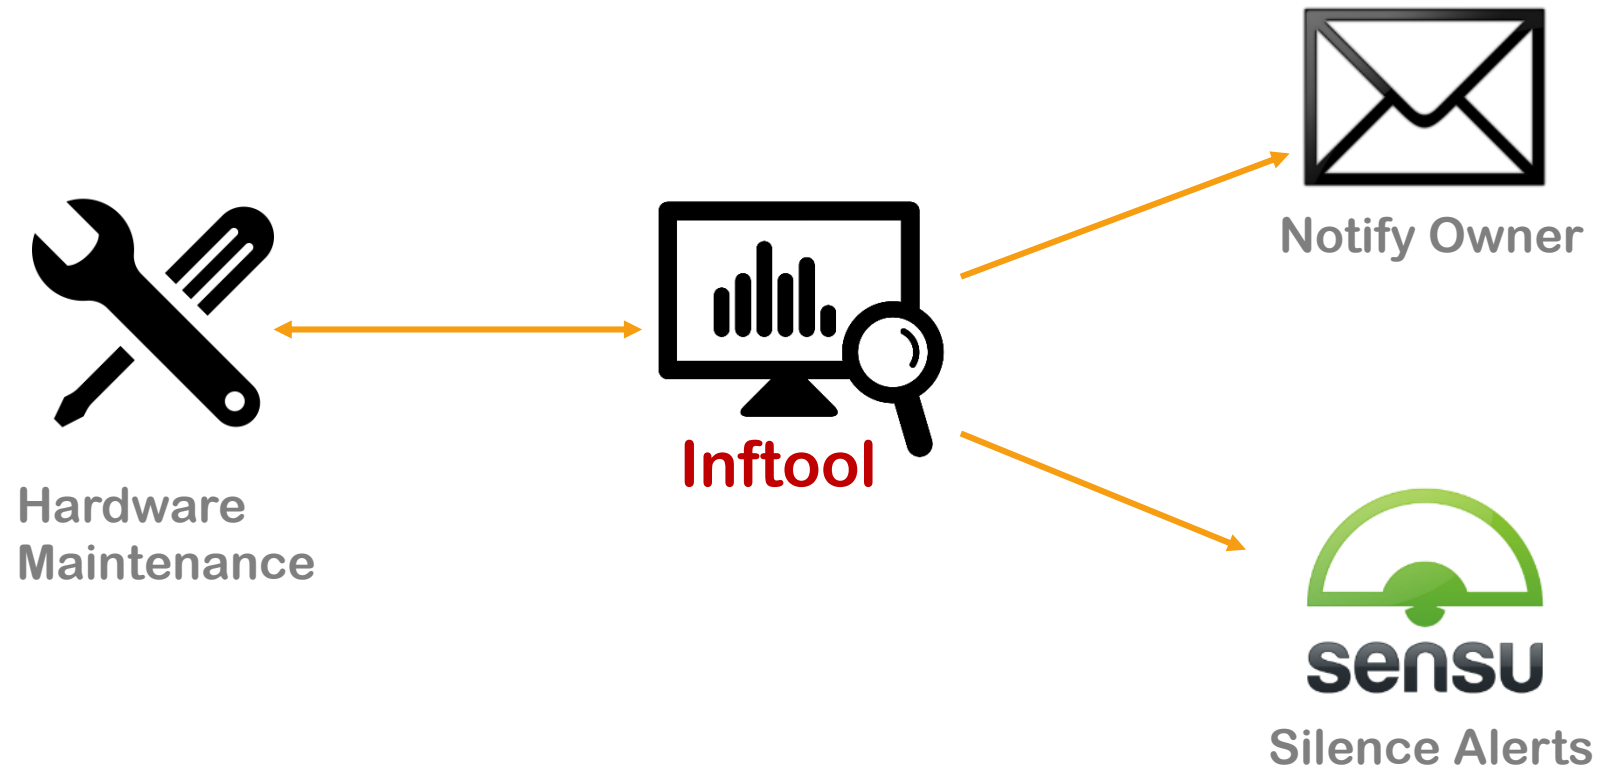
OpenStack deployment at Agoda:

In house development “Inftool” for Infrastructure services

- Centralized server management tool
- Stores metadata information
 - Network Interface detail
 - Physical location
 - Product & Application
 - Contact information
 - Insight hardware information
- Controls infrastructure services
 - Patching
 - Maintenance
 - etc.



Example: Hardware maintenance.



OpenStack deployment at Agoda:

Does it scale well? Real examples:

- Racks of servers can be installed in minutes
- Joined the top 5% largest OpenStack clusters within 1 year
- A developer requested 100 test servers to test some software
- A team requested 160 servers in 4 locations to scale their app into production
- A new business partner required double capacity for our hotel search APIs

Next Step

- More automation needed
- OpenStack infrastructure expansion and upgrades
- Many new technologies out there to test for business benefits

Your opportunity

We are hiring !!!

DevOps Engineer

- Experienced with
 - Python, Ruby, Bash, PowerShell
 - System administration experience

Full Stack Developer

- Experienced with
 - Python, Flask, Message Queue , jQuery, FE design

Environment

- International team: Thai, USA, New Zealand, Cambodian, Indian, German
- Cutting edge technologies, rooms to play

Interested? E-Mail CV to sharkrit.impat@agoda.com

**WE'RE
HIRING!**