KT ucloud storage

Two Years of Life with OpenStack Swift

2012. 08. 11 / Jaesuk Ahn, Cloud OS Dev. Team, Korea Telecom

Who I am.

Jaesuk Ahn (安宰奭)

- 2012 ~ : Cloud OS Dev. Team Lead, KT
- 2011 ~ : OpenStack Korea Community Leader
- 2009 ~ 2011: KT Master Project Manager
  - Research on Open Source Cloud Tech.
  - Open Source Community Support

email: js.ahn@kt.com / bluejay.ahn@gmail.com

twitter: @songerie

Who I am.

Who we are.

kt

openstack™

Korea Telecom (KT)

- The First 'Public' Cloud in Korea (2010年 8月)
- A Pioneer in Cloud Computing in Korea
- Common HW and Open SW
  (XenServer, CloudStack,

- The First Company Commercializing OpenStack Swift in Asia
- Cloud Data Center of High Density and Efficiency

kt

openstack™

Software Adoption

S/W	Description	How to Use
XenServer (Free)	Hypervisor	collaboration
CloudStack	Cloud OS (Management Stack)	collaboration
NexentaStor	Storage management	collaboration
MySQL	Database	Building
Splunk	System Log data gathering & analysis	collaboration
Nagios, Zabbix, Collected	Monitoring	Building
Chef	Automation of Cloud Deployment & Configuration	Building
Spring	Application Framework	Building
OpenStack Swift	Object Storage	Building

Internalization (Internal Private Cloud)

- 8,000 vm, 80% Cost Reduction

  from Internal IT Infrastructure.

- Saved 47 Mill. US$ for CAPEX

  after Jun. 2010 from Internalization with Cloud

- From Virtualization to Cloud
- Cultural Changes

kt

openstack™

Public Cloud (https://ucloudbiz.olleh.com)

- 1,400 Customers with ucloud biz.
  - 3,000 VMs for SMB, LE and Gov.

ucloud personal	Real time data sync and backup for mobile/internet subscribers
ucloud office	Real time data sync and backup for B2B customers
ucloud Server	Virtual Server with value-added services (LB, Firewall, etc)
ucloud Server+	Automation and Orchestration (Auto-Scaling)
ucloud CDN	Contents delivery network for storage user's
ucloud VDI	Virtual desktop infrastructure from servers
ucloud Storage	Mass object storage service (OpenStack swift)
ucloud DB	RDB: Relational DB
ucloud Backup	Auto sync and back-up

kt

openstack™

Building Eco System

- Expanding Service to Virtual Private and Hybrid Cloud.
- Unified Monitoring Service (like CloudWatch)
- Additional Feature like ELB, EIP, Availability Zone etc.
- Strengthen Automation and Security
- Providing Cloud Application Architecture Guide
- Building Eco System with Cloud Incubating Center (http://www.cloudincubation.com)
  - 3 Months Free, Providing Office Space and Marketing Support, Use Cases, Technical Documentations, Videos.

kt

openstack™

Who am I ?

Why am I here ?

Let’s Start, Swift!

kt

openstack™

KT ucloud storage:

OpenStack Swift-based Object Storage Service

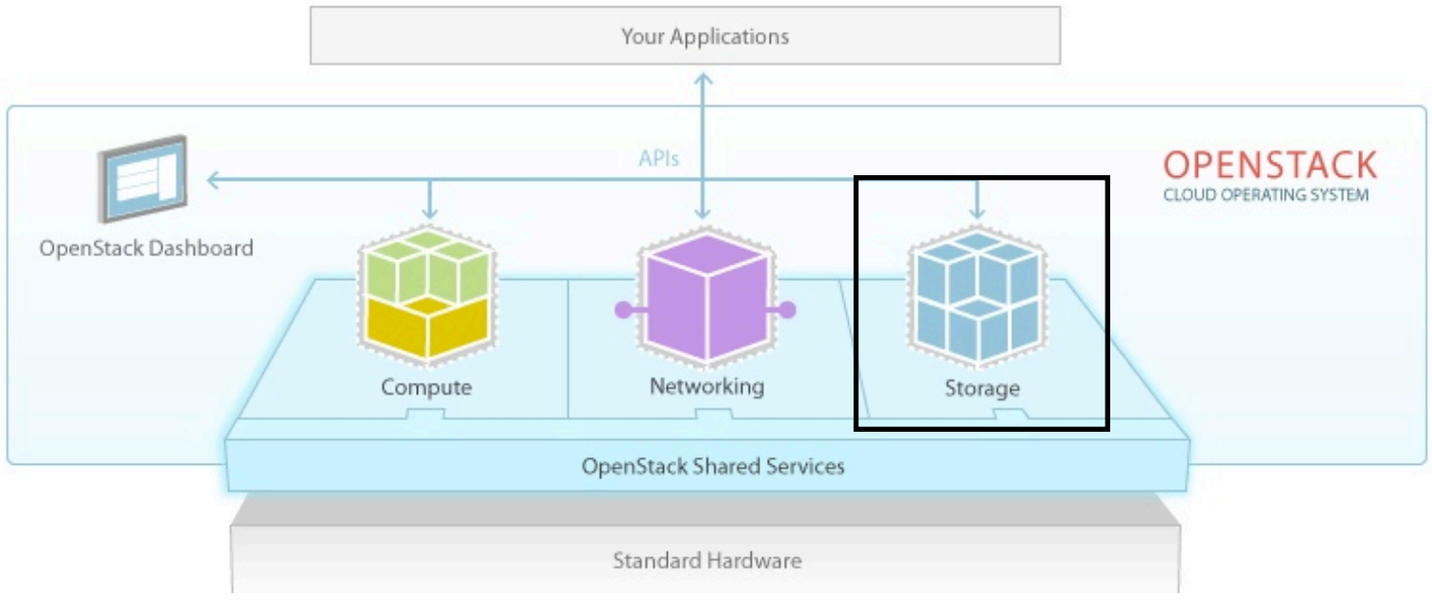

OpenStack Object Storage (a.k.a Swift) - Capabilities

- **Redundant, scalable object storage** using clusters of standardized servers capable of storing petabytes of data
- **Distributed storage system** for static data. Having no central "brain" provides greater scalability, redundancy and durability.
- **Reliability:** Objects and files are written to multiple disk drives spread throughout servers in the data center, with the OpenStack software responsible for ensuring data replication and integrity across the cluster.
- **Scalability:** Storage clusters scale horizontally simply by adding new servers. Because OpenStack uses software logic to ensure data replication and distribution across different devices, inexpensive commodity hard drives and servers can be used in lieu of more expensive equipment.

OpenStack Swift - Basic Architecture

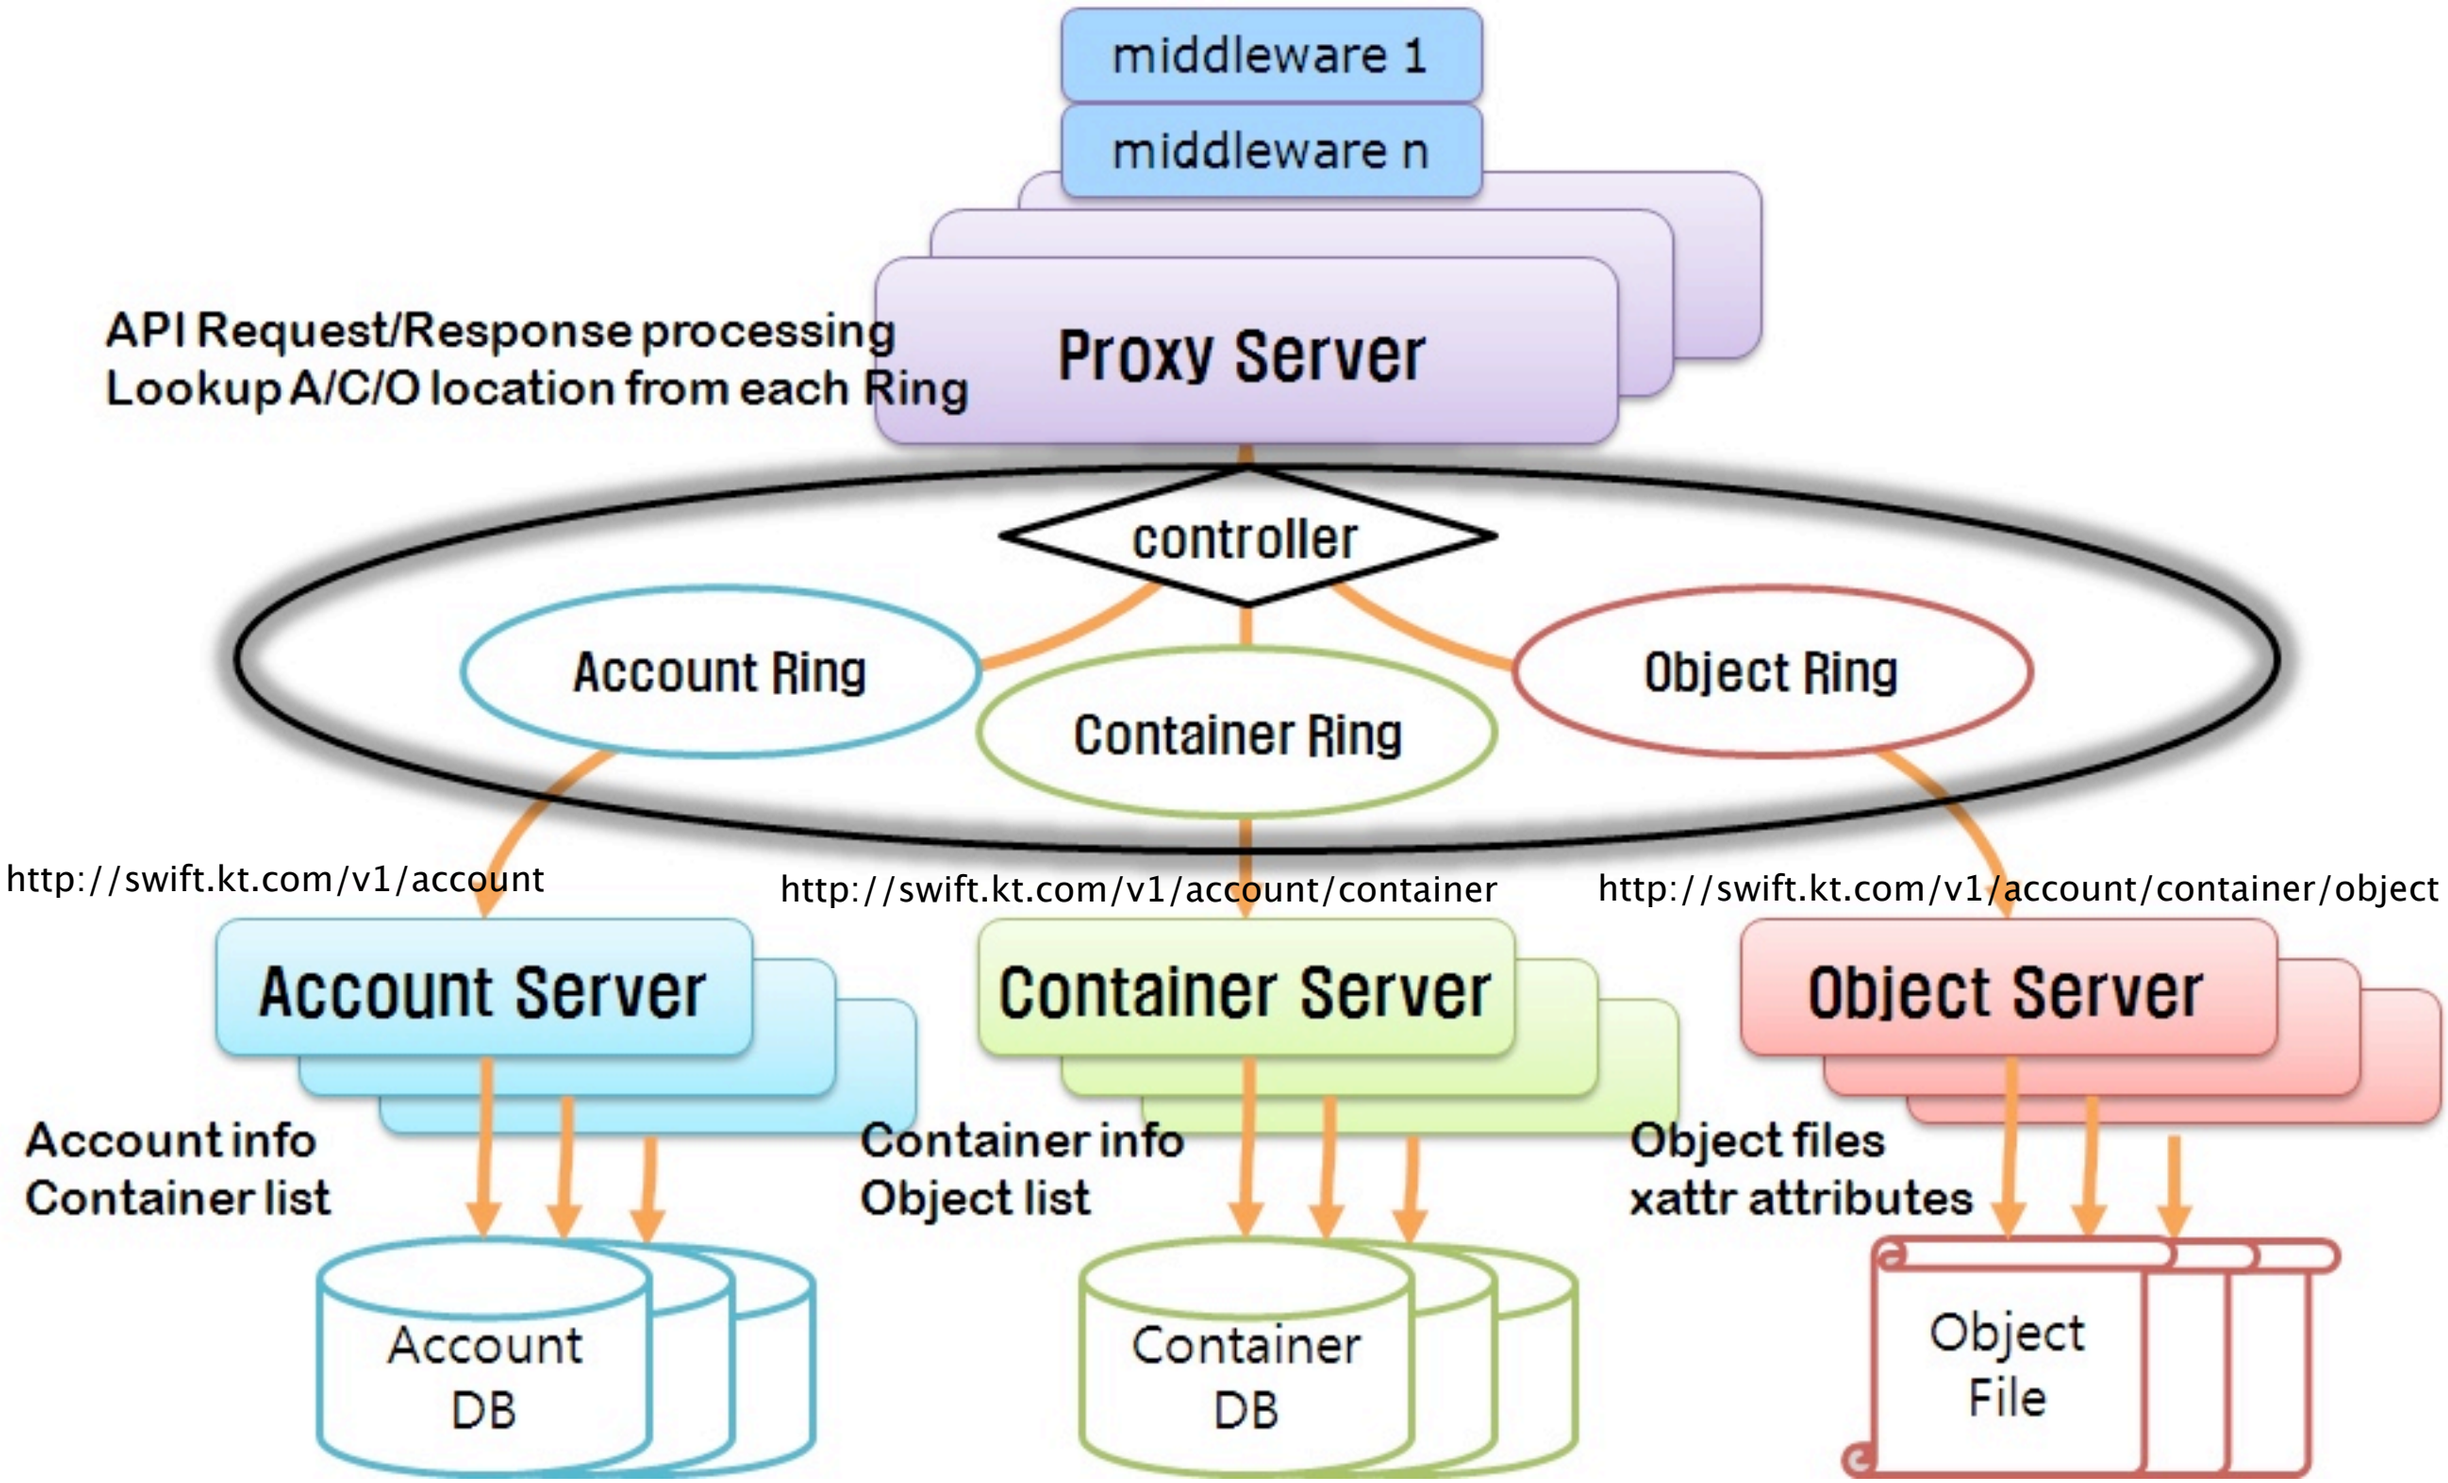

OpenStack Swift - RESTful API

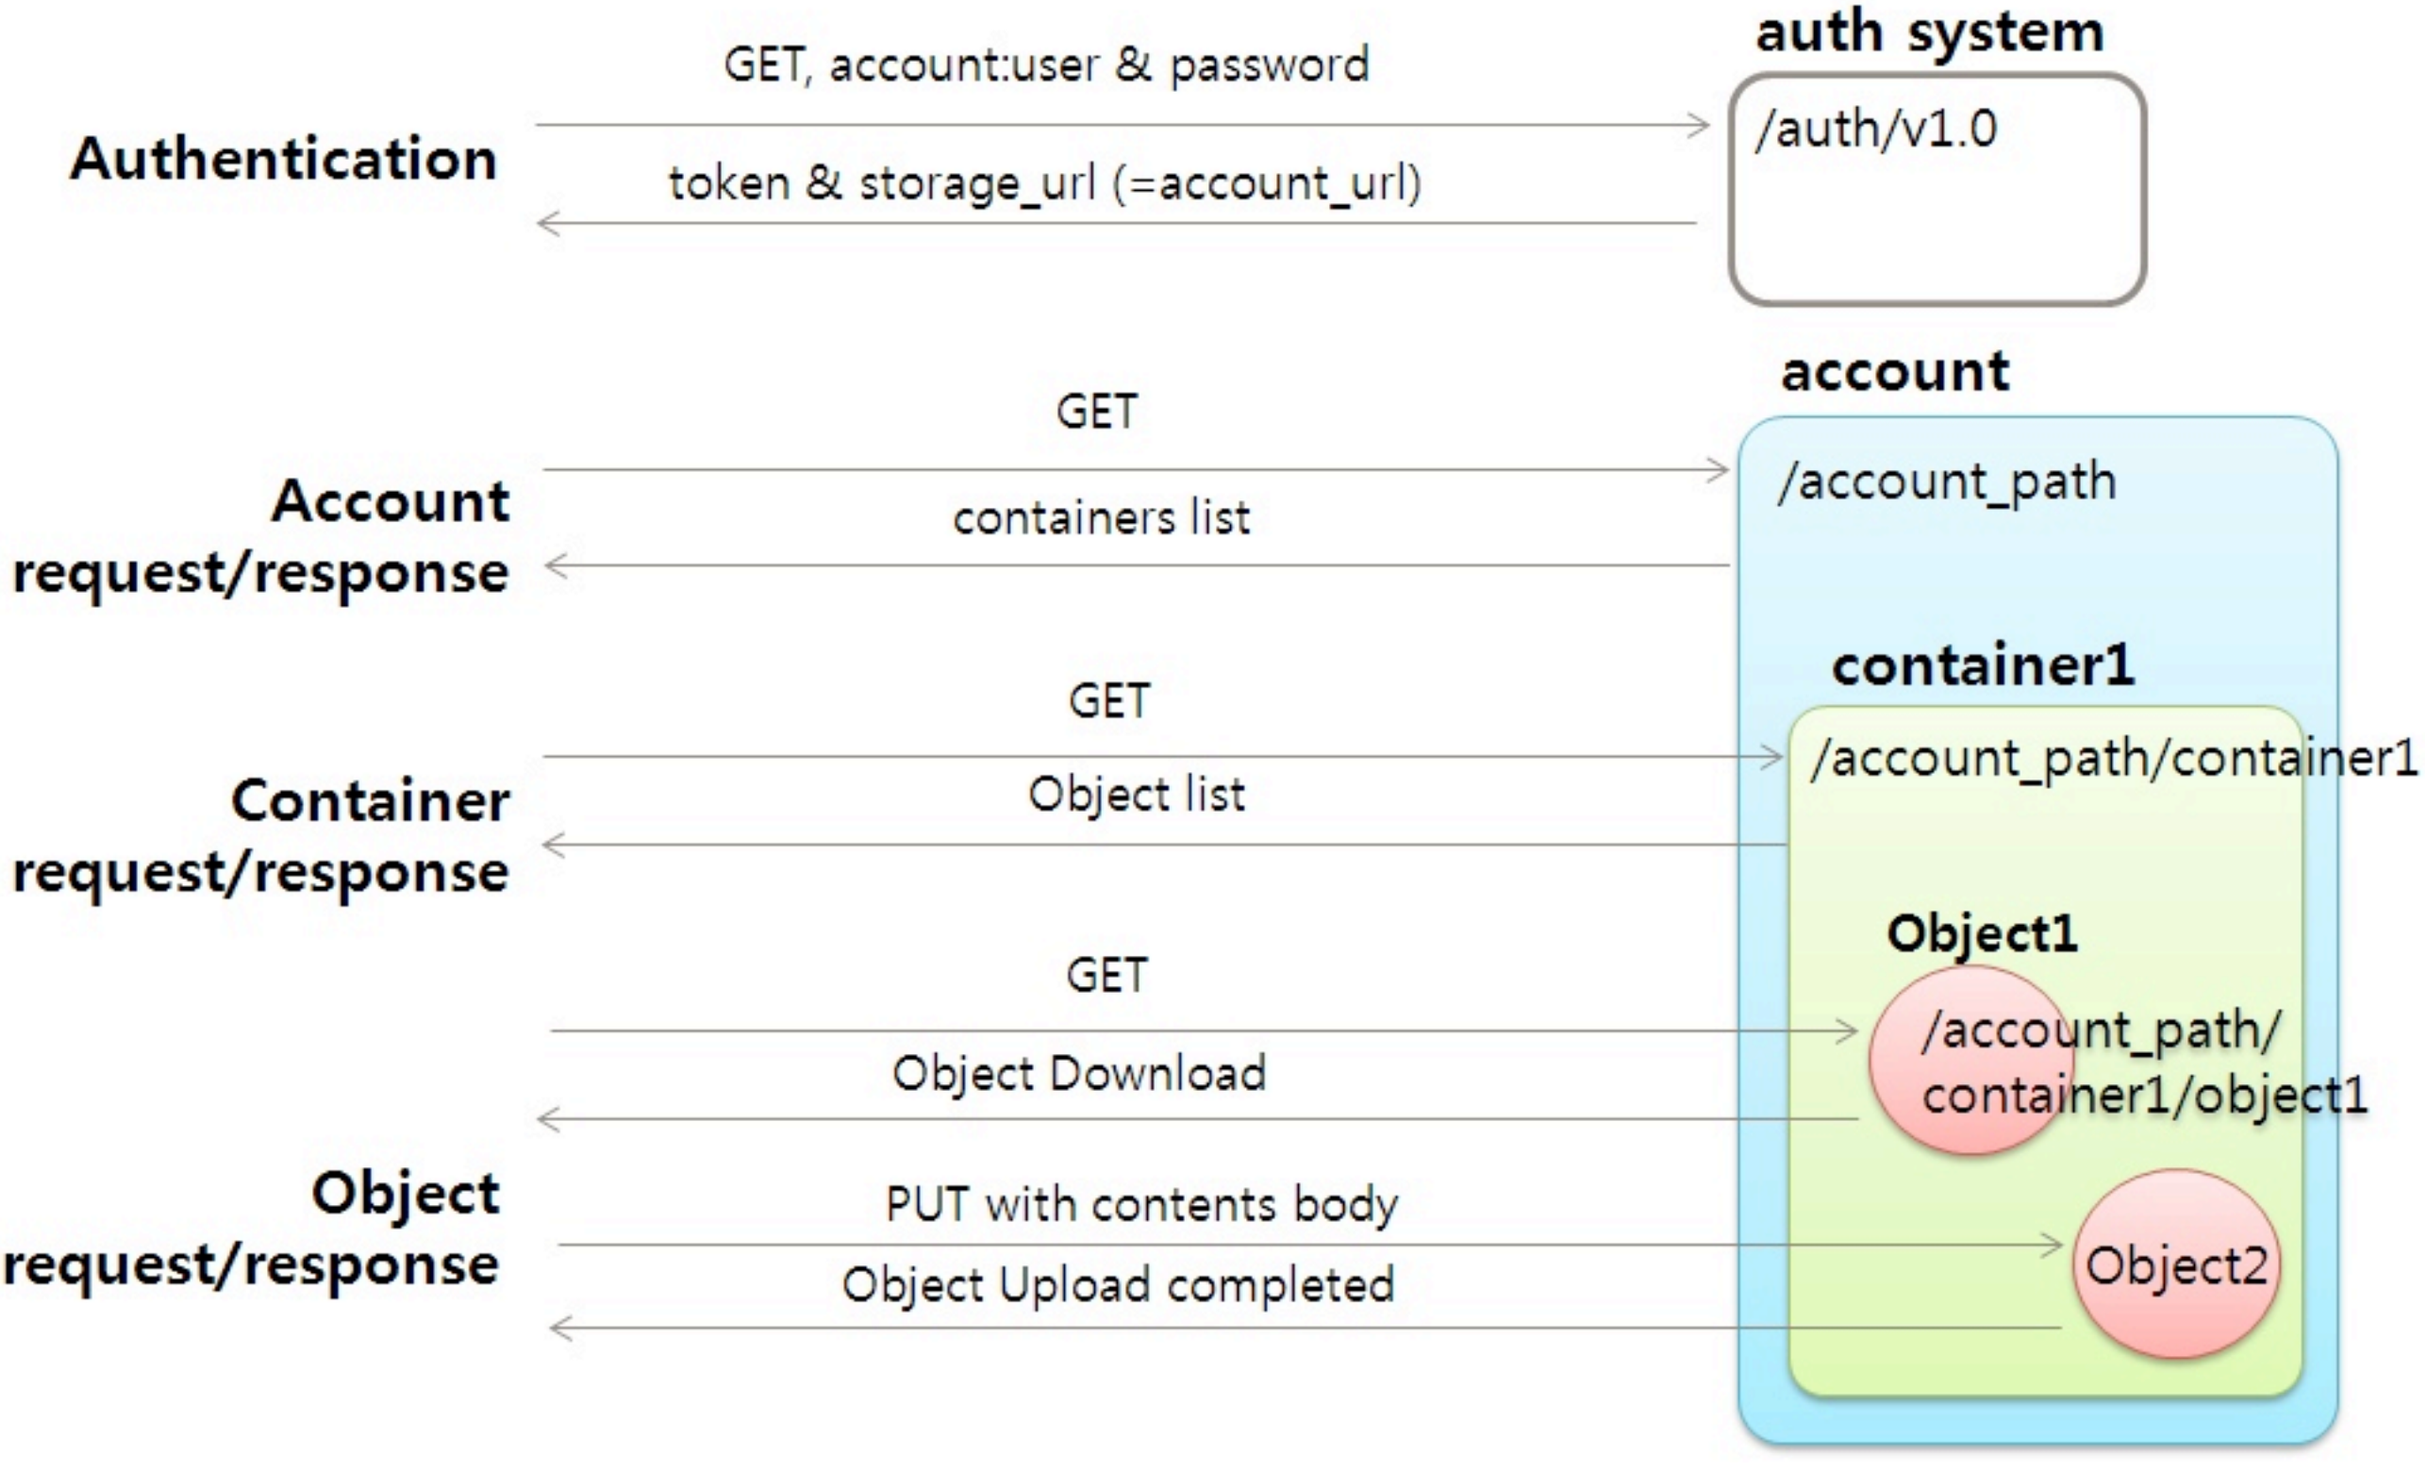

OpenStack Swift - What KT has been added.

SSL performance tuning

KT API for management

KT middleware
- integrating with KT's Single Sign On System

middleware 1

middleware n

Proxy Server

API Request/Response processing
Lookup A/C/O location from each Ring

Metering data for specific billing requests

Fast Ring Access with SSD

controller

Real Time Usage Reporting instead of eventually-consistent usage reporting

Account Ring

Container Ring

Two separate deployments
2 replications + 3 replications

Account Server

Container Server

Object Server

Account info
Container list

Container info
Object list

Object files
xattr attributes

Account DB

Container DB

Object File

kt

openstack™

What KT had to do to commercialize OpenStack Swift

- Product service requirements
- HW selection
- Network design
- HW architecture design
- Facility Planning
- HW standup
- SW provisioning
- System Configuration
- Load balancing
- Authentication integration
- Utilization & billing integration
- Additional “Value Added” services & integration
- Monitoring development and integration
- Operational Tooling
- Operator Training and Documentation

kt

openstack™

What KT had to do to commercialize OpenStack Swift

- **Product service requirements** Public Object Storage Service
- **HW selection** Commodity Servers, Disks (JBOD), Switches (Arista)
- **Network design** 10G for public service, 1G for management network
- **HW architecture design** Similar to Rackspace Reference Architecture (5 zone 3 replication)
- **Facility Planning** Zone Separation, H/W Scale-Out , 3 system (dev, staging, production)
- **HW standup**
- **SW provisioning** Automated Test Environment (Chef-Vagrant/VirtualBox), Production (Git+Chef)
- **System Configuration** Fully Automated Deployment/Configuration based on Chef
- **Load balancing**
- **Authentication integration**
- **Utilization & billing integration** KT customized logging for realtime reporting & billing
- **Additional "Value Added" services & integration** CDN integration, advanced CLI tool, various language bindings, Management API, Cloud Portal
- **Monitoring development and integration** KT monitoring system based on Zabbix & Collectd
- **Operational Tooling**
- **Operator Training and Documentation**

kt

openstack™

Commercial Deployment

- 2010.7 ~ : Research & Study (KT R&D Department)
- 2011.1Q ~ : Commercial Development/Deployment (KT Cloud Business Unit)
- 2011.2Q : Commercialization (First in Asia)

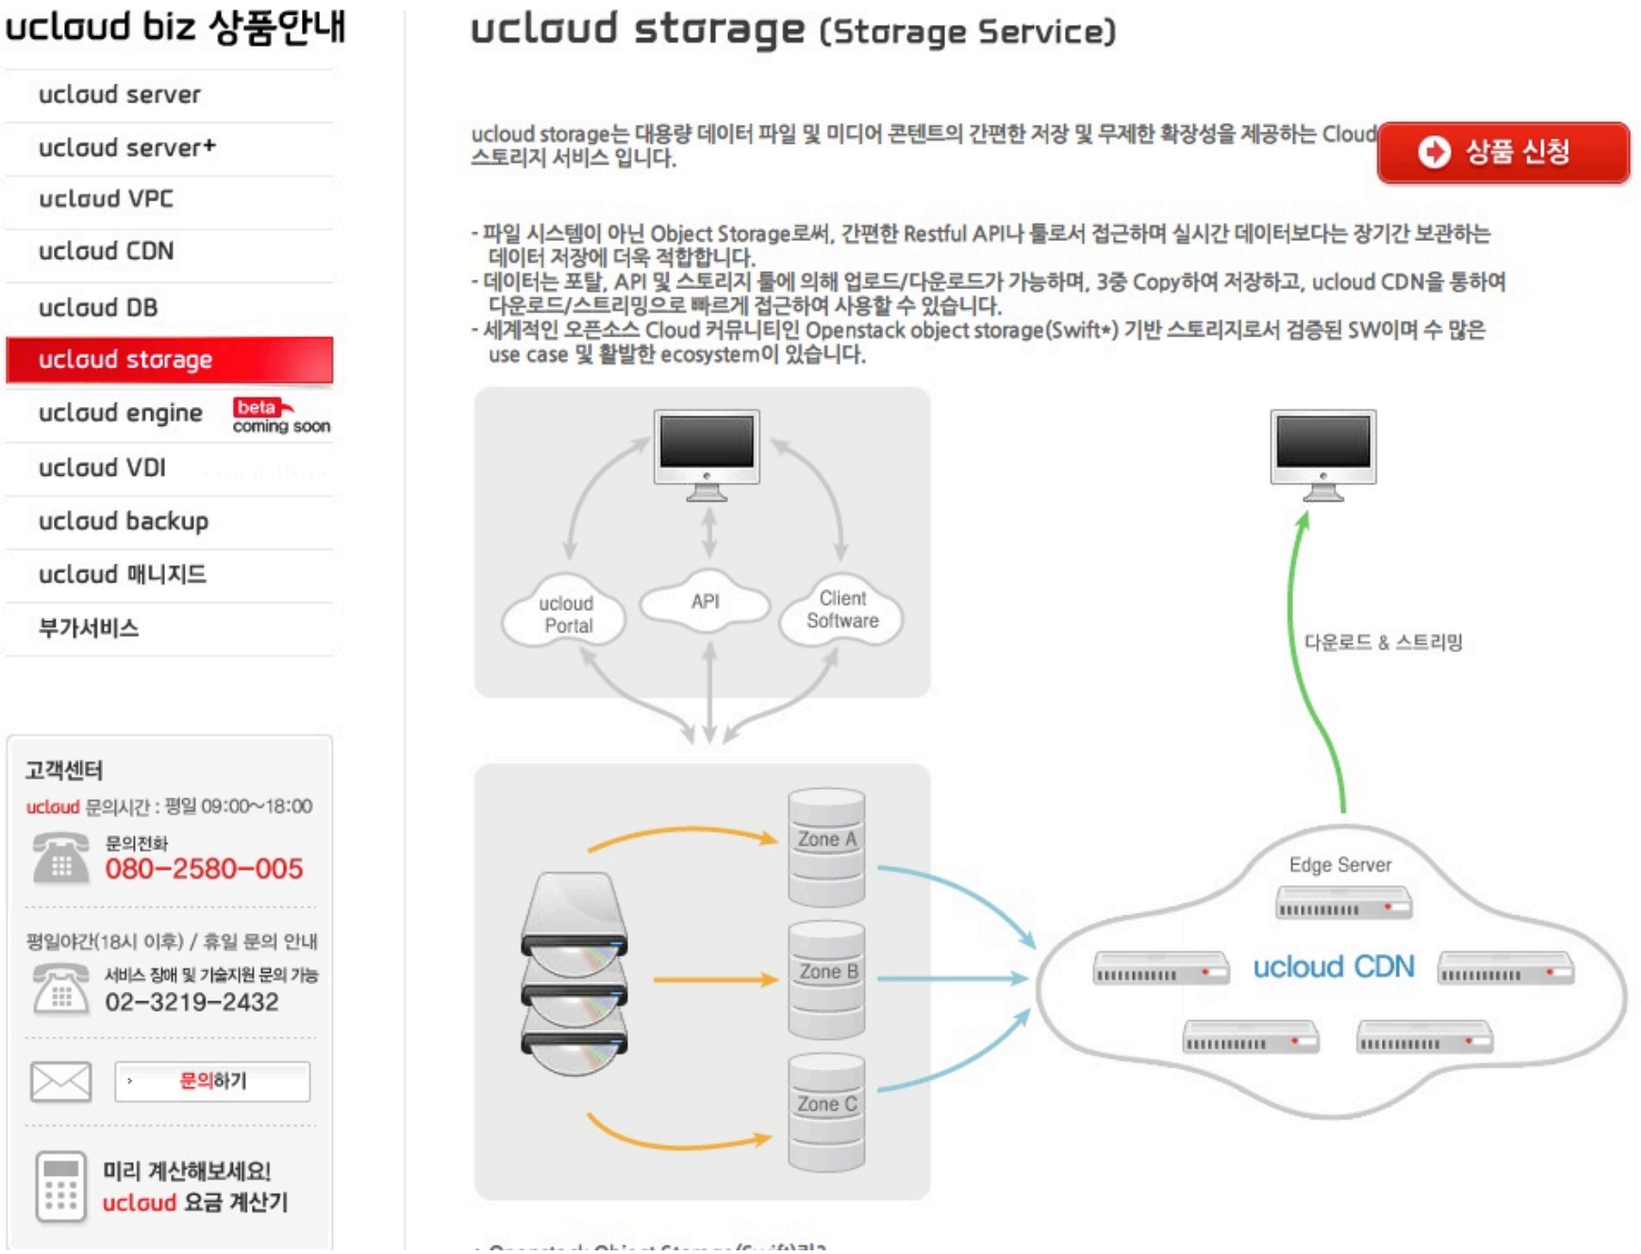

kt

openstack™

KT ucloud storage - architecture

- 5 zone - 3 replications
- Commodity Hardware based on JBOD
- currently Diablo version -> preparing upgrade to Folsom version

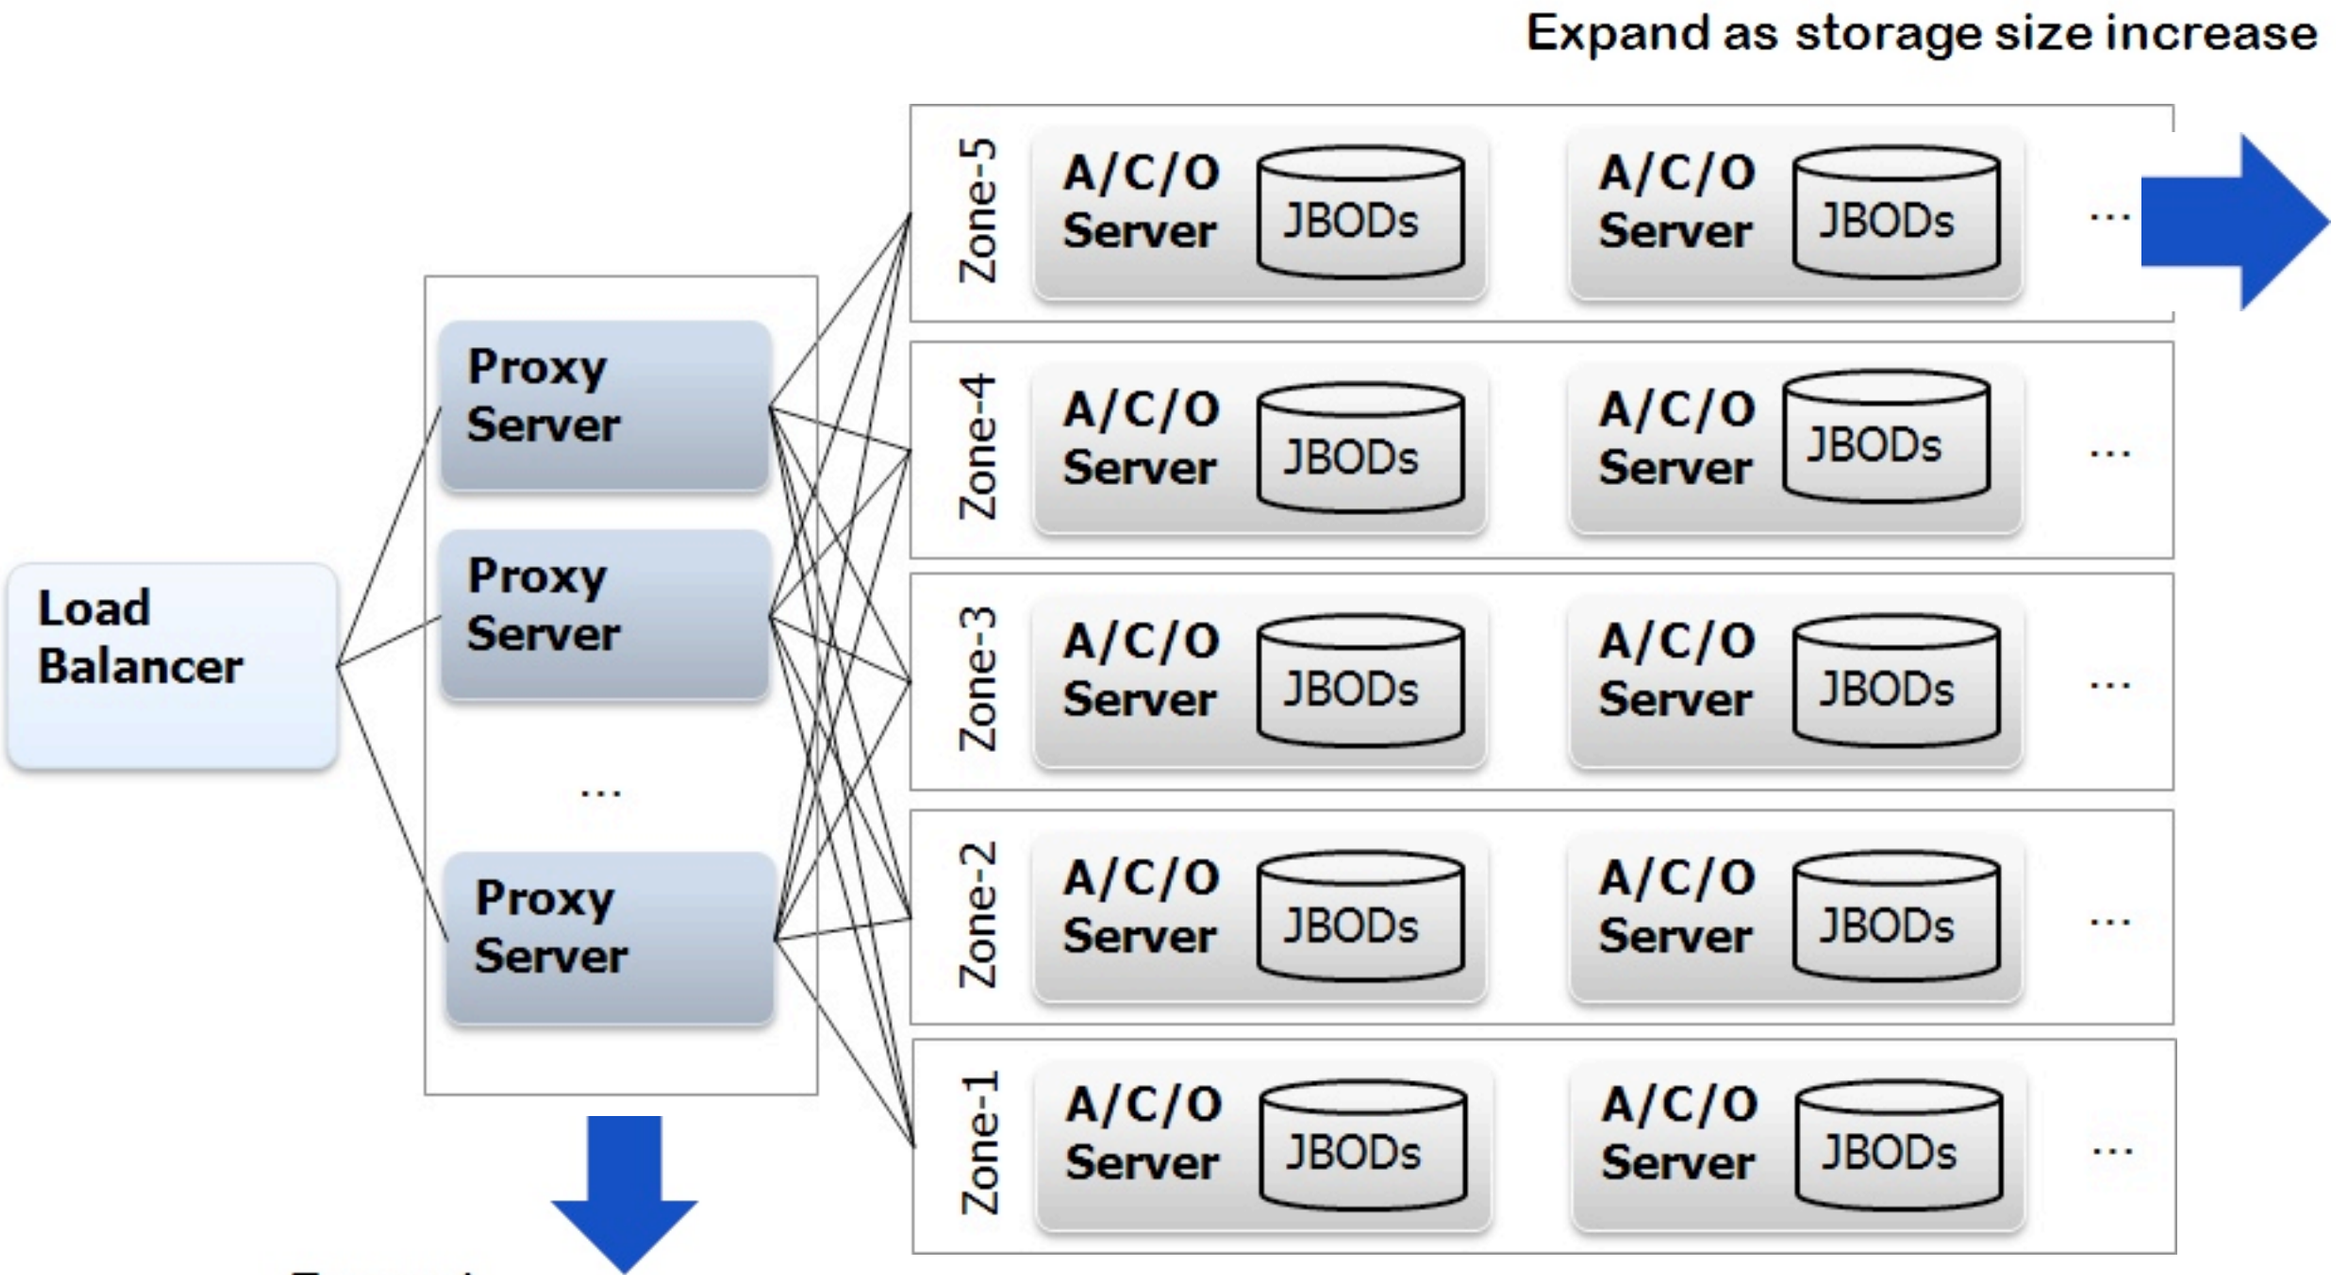

KT ucloud storage - Service Process

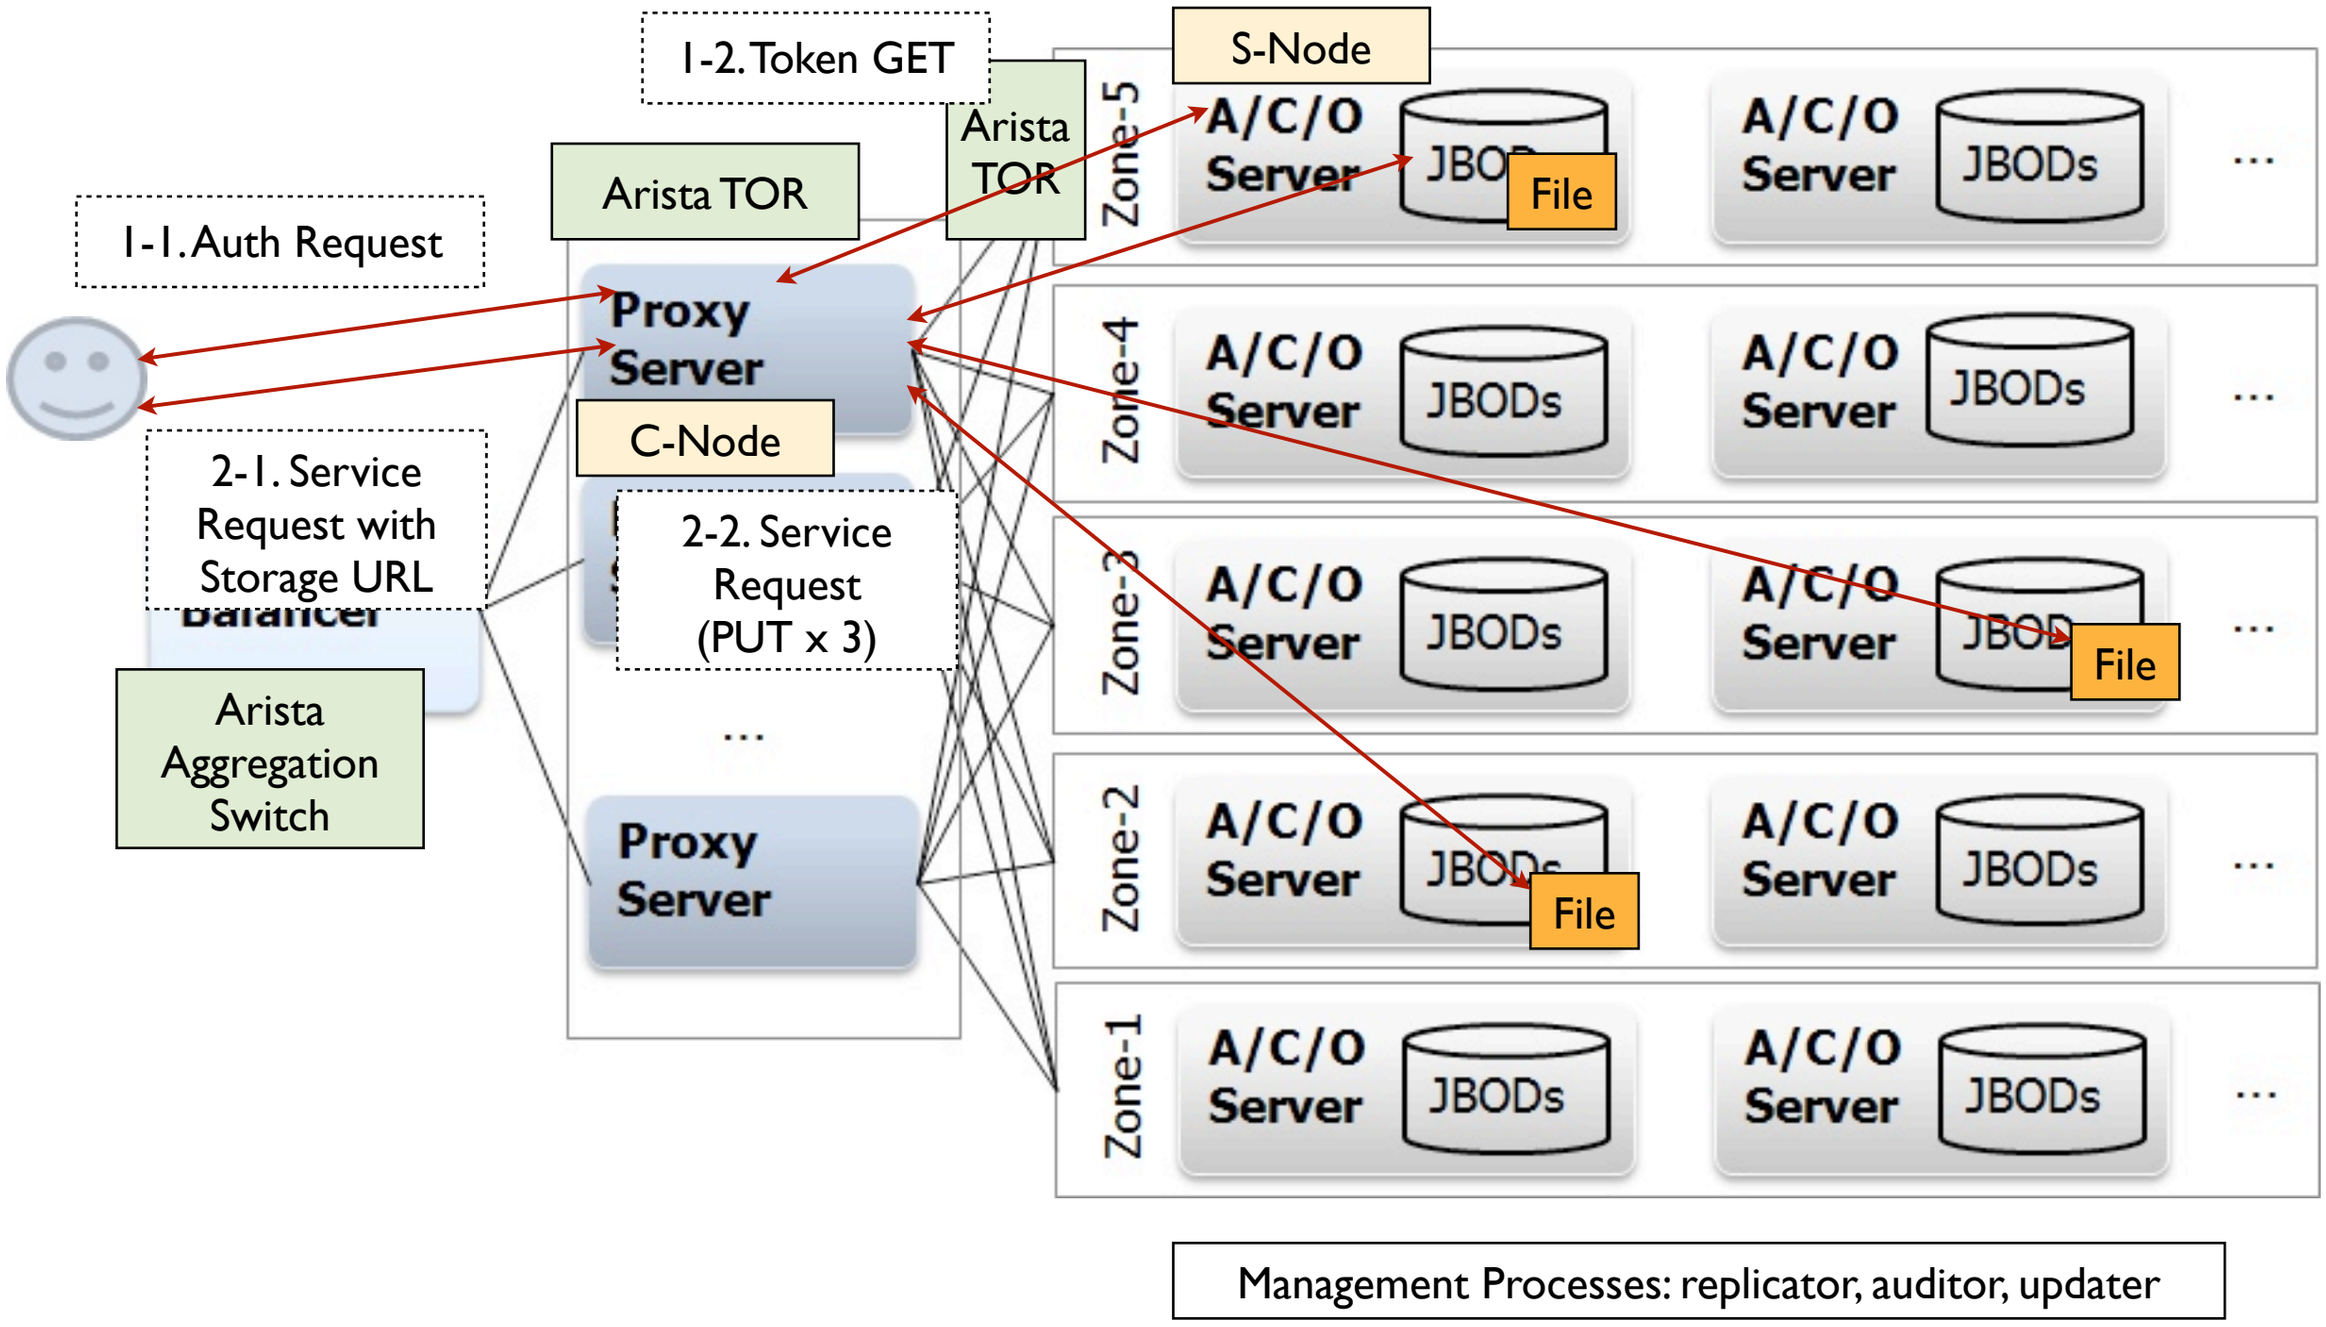

Automatic Deployment (General)

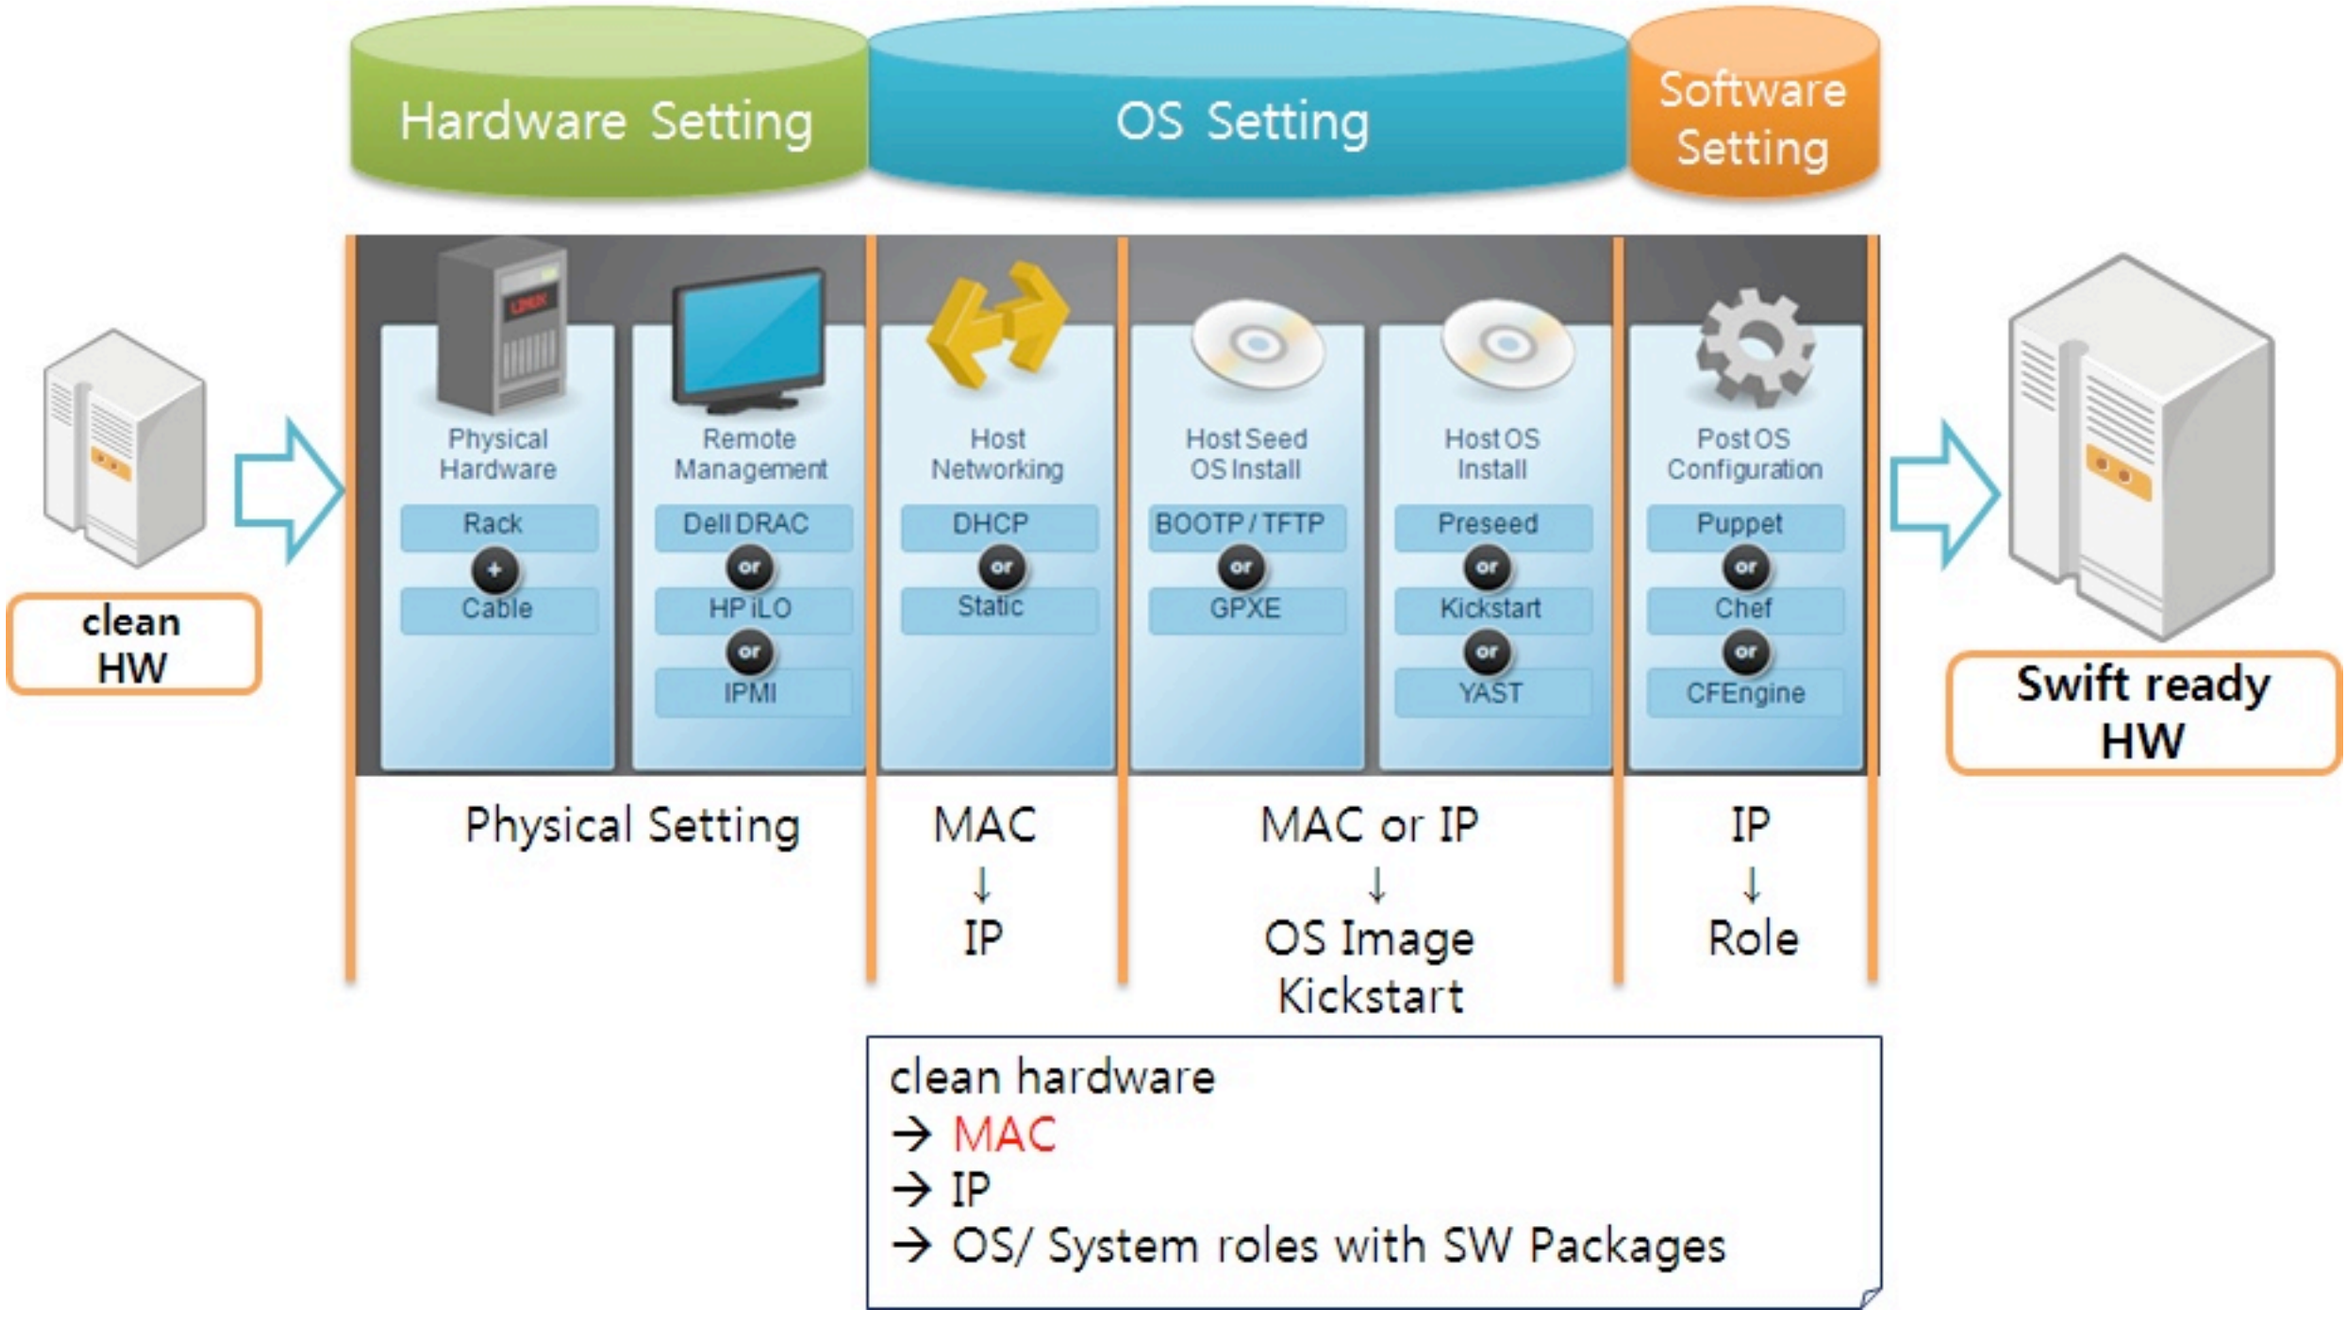

KT ucloud storage - Automatic Deployment

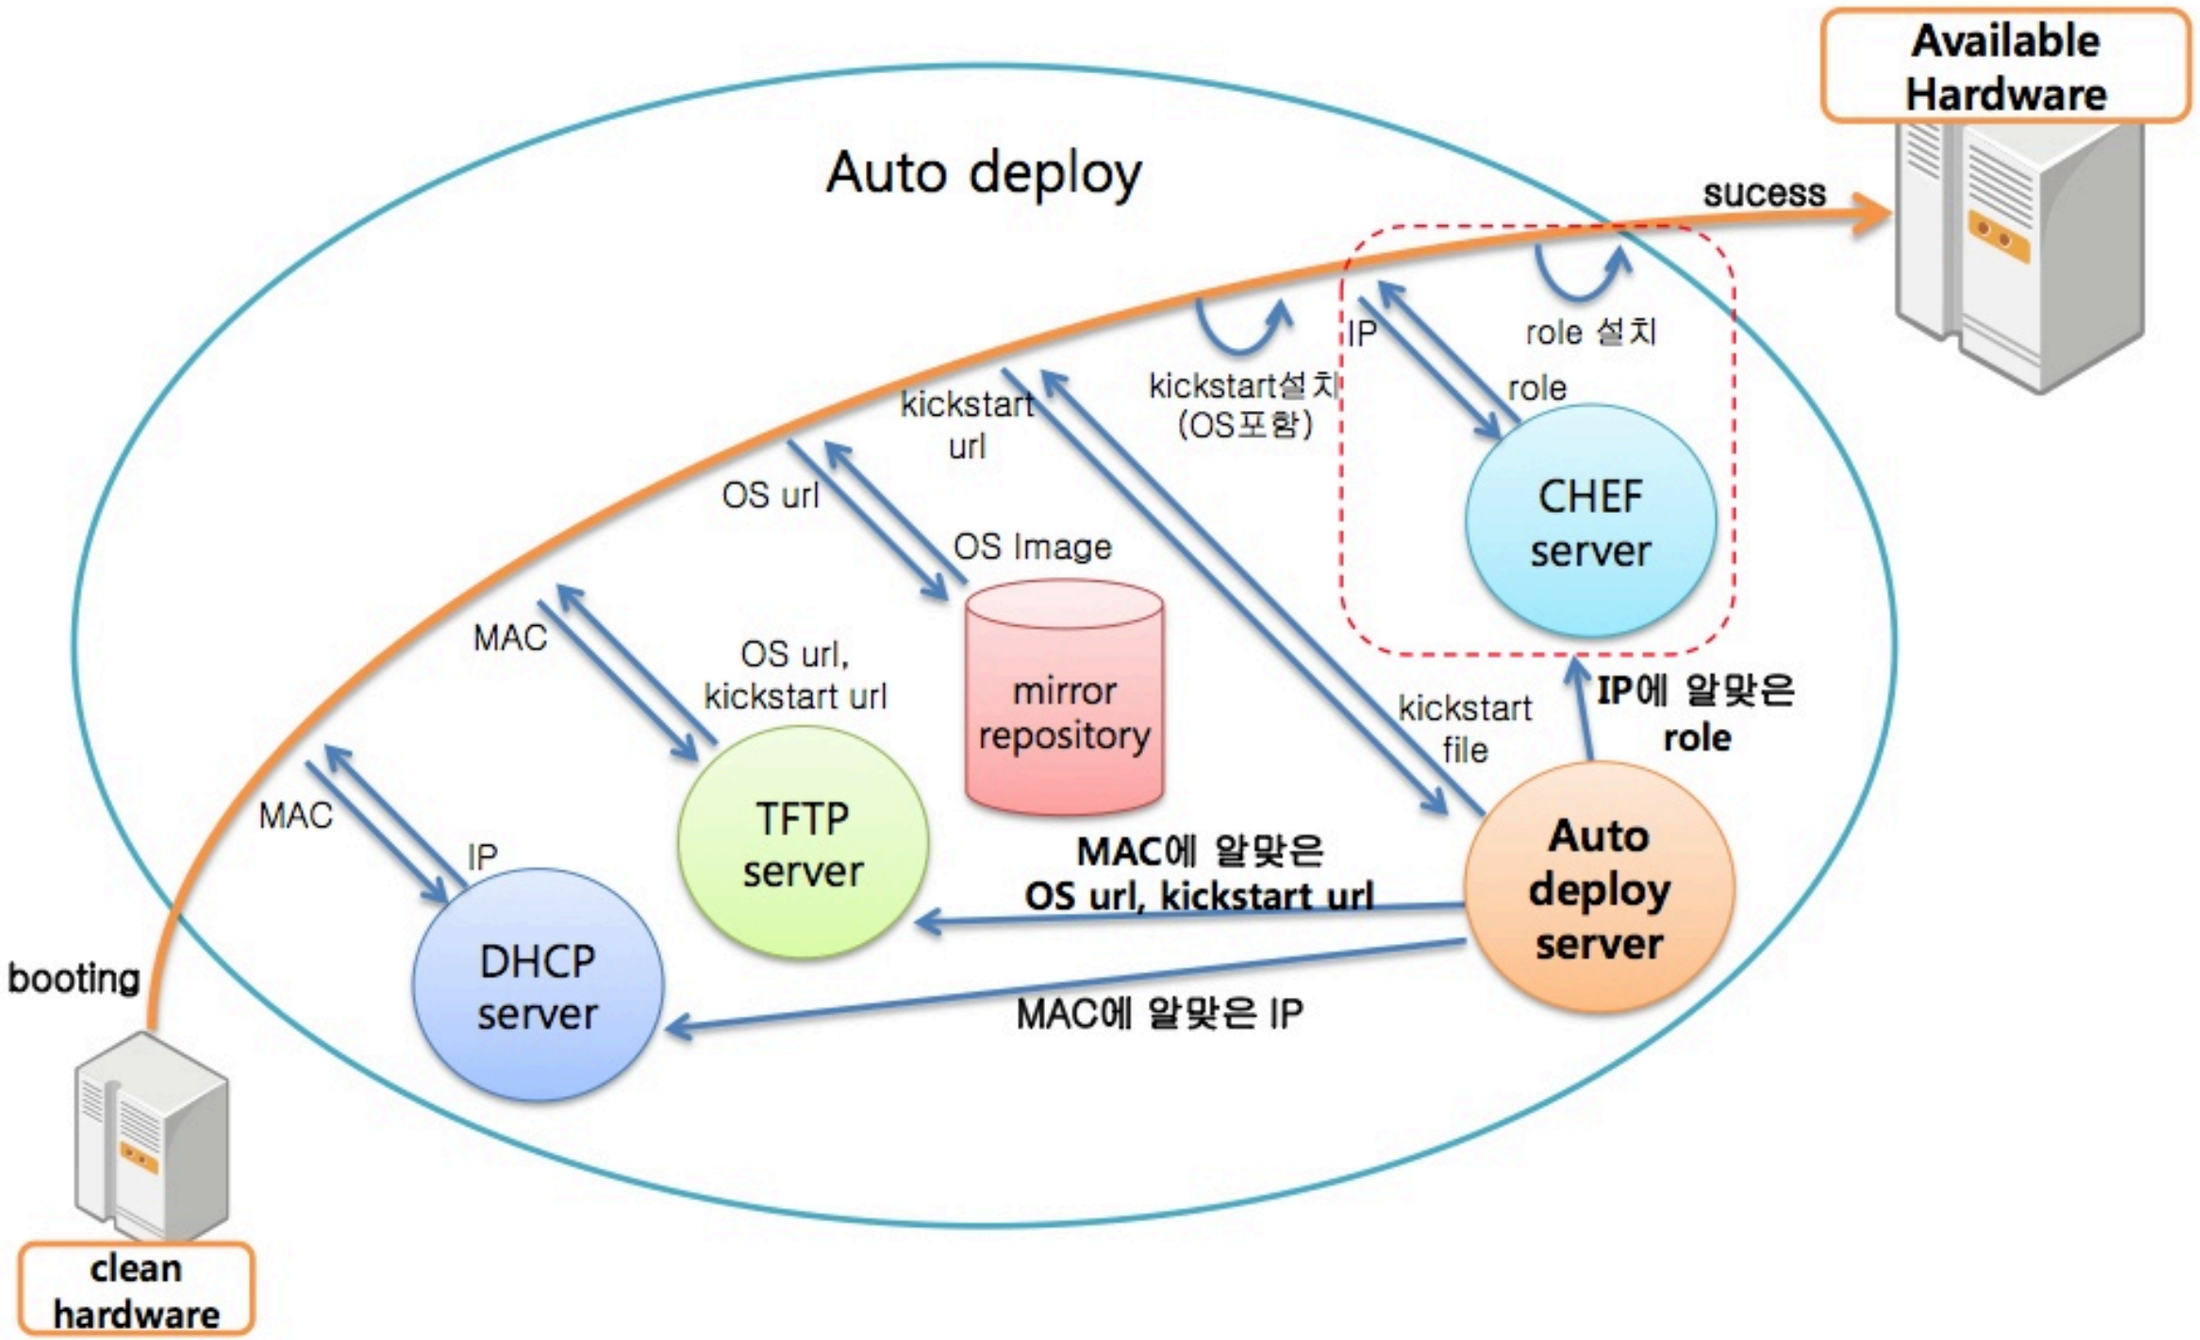

KT ucloud storage - Automatic Deployment

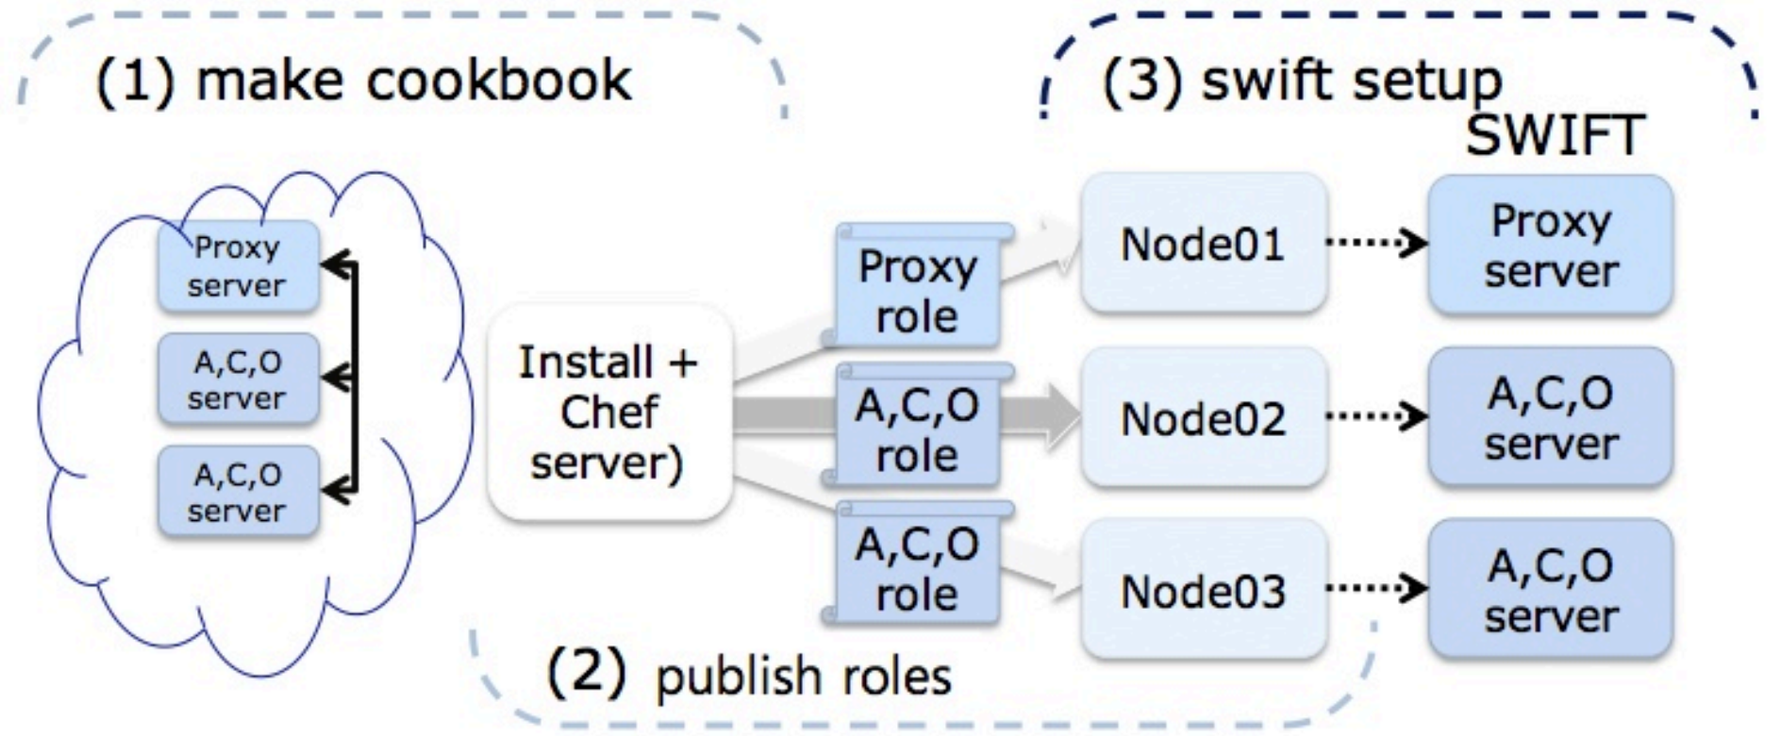

1. Default installation (PXE boot)
   - ✓ LAN boot, IP allocation
   - ✓ Kick start file (OS, chef client)
2. Role install (swift-multi)
3. Ring build

KT ucloud storage - Client Tools

- Benefits using OpenStack -
  There are many.. many.. client tools to provide.

Name	Description	URL	Cloud storage
Cyberduck	Open source FTP, SFTP, WebDAV, Cloud Files, Google Docs & Amazon S3 Browser for Mac & Windows.	Cyberduck.ch	Amazon S3 FTP, WebDav Google Docs Cloudfiles Window Azure **Openstack** Etc…
Gladinet	Seamless access, aggregate, and backup to cloud storage	Gladinet.com	Amazon S3 FTP, WebDav Google Docs Cloudfiles Internap **Openstack** Etc…
Cloudfuse	Cloudfuse is a FUSE application which provides access to Rackspace's Cloud Files (or any installation of Swift).	ohloh.net/p/cloudfuse	Cloudfiles **Openstack**(Swift)

KT ucloud storage - Use Case

- Origin Store for CDN

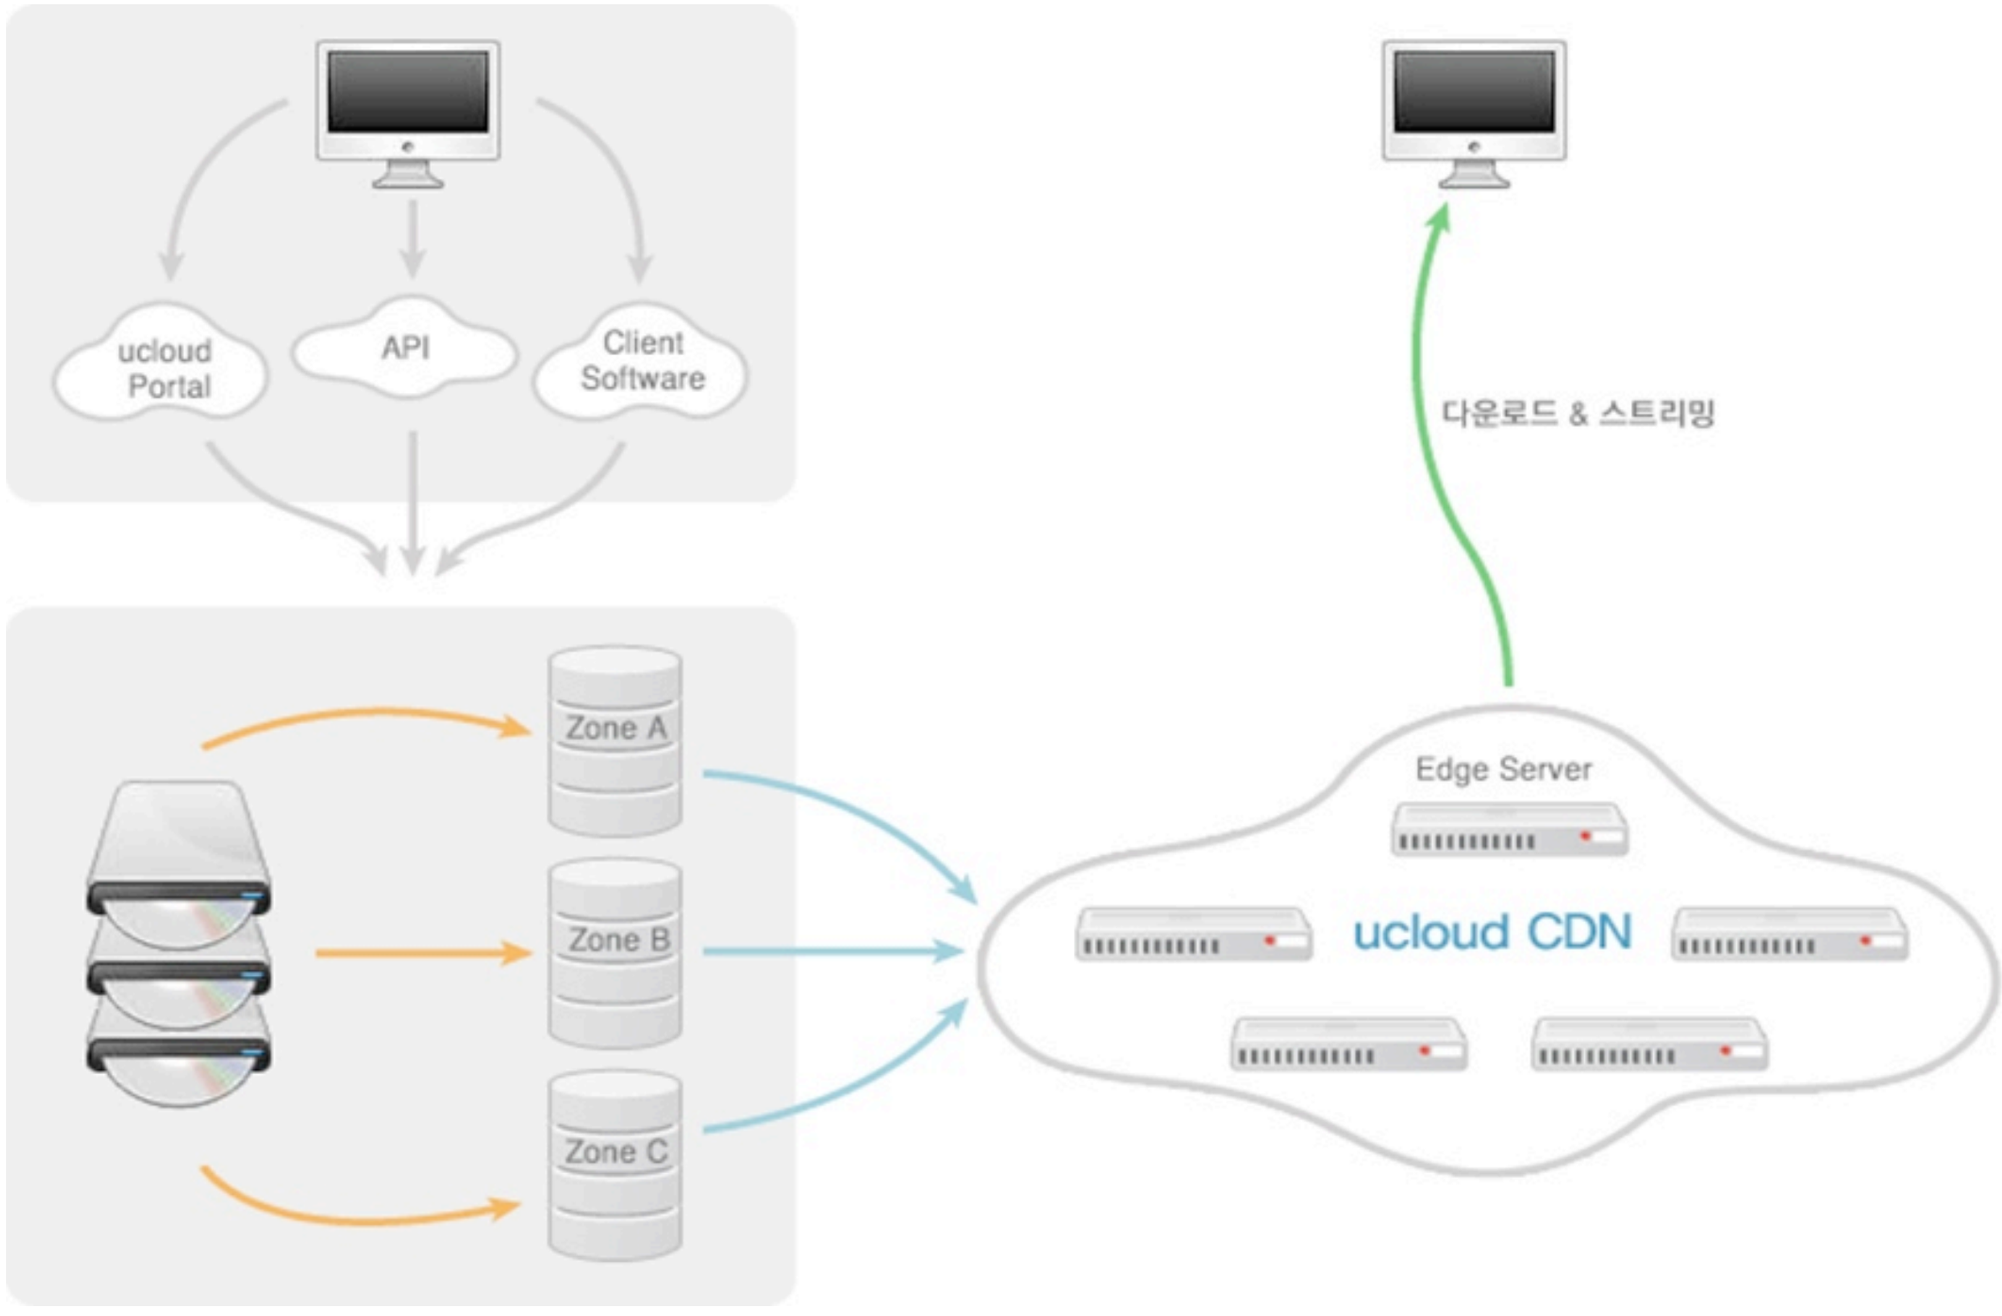

KT ucloud storage - Use Case

- Backend-Storage for SNS service
- Media Archiving
- Data Backup
- Snapshot/VM image backup store
- Backend Storage for Cloud File Service
- Addition feature in Personal NAS storage

Some Test Data

- Just for Fun, Result can be different with different setting and circumstance.

- Upload/Download from/to its own VM service.
- What kind of server is used.
  - ucloudbiz VM CentOS 5.8
  - ucloudbiz VM Ubuntu 11.04
  - Amazon EC2 CentOS 5.8
  - Rackspace Cloud Servers CentOS 5.8
  - Rackspace Cloud Servers Ubuntu 11.04

Some Test Data

Avg. Upload Speed (Mb/s)	500MB	100MB	10MB	1MB	100KB	50KB	10KB
ucloudbiz Ubuntu	208.87	168.16	151.22	49.20	7.10	3.88	1.21
ucloudbiz CentOS	182.16	146.20	139.69	41.92	6.47	3.45	1.02
Amazon CentOS Tokyo	146.84	144.14	118.23	38.57	4.96	2.68	0.54
Amazon CentOS N.California	136.86	140.64	118.47	38.79	5.26	2.95	0.59
Rackspace CentOS US	96.22	94.85	76.37	32.97	5.86	2.71	0.61
Rackspace Ubuntu US	95.00	93.80	79.00	28.08	4.28	2.27	0.43
HP CentOS Arizona	29.19	22.84	18.70	6.32	0.72	0.25	0.08
HP Ubuntu Arizona	57.22	56.66	35.53	8.83	0.96	0.49	0.10

Avg. Download Speed (Mb/s)	500MB	100MB	10MB	1MB	100KB	50KB	10KB
ucloudbiz Ubuntu	522.28	463.88	314.99	80.48	16.46	8.92	1.70
ucloudbiz CentOS	523.91	460.15	293.16	70.56	12.89	6.78	1.31
Amazon CentOS Tokyo	351.51	332.19	206.13	50.40	6.83	3.50	0.78
Amazon CentOS N.California	340.30	314.45	224.09	53.70	7.39	3.79	0.80
Rackspace CentOS US	384.07	295.44	182.10	68.93	12.19	6.69	1.03
Rackspace Ubuntu US	269.81	249.79	178.05	49.22	7.37	3.85	0.65
HP CentOS Arizona	37.21	22.16	16.46	5.40	0.74	0.31	0.07
HP Ubuntu Arizona	74.40	58.03	26.54	7.98	1.01	0.51	0.09

Some Test Data

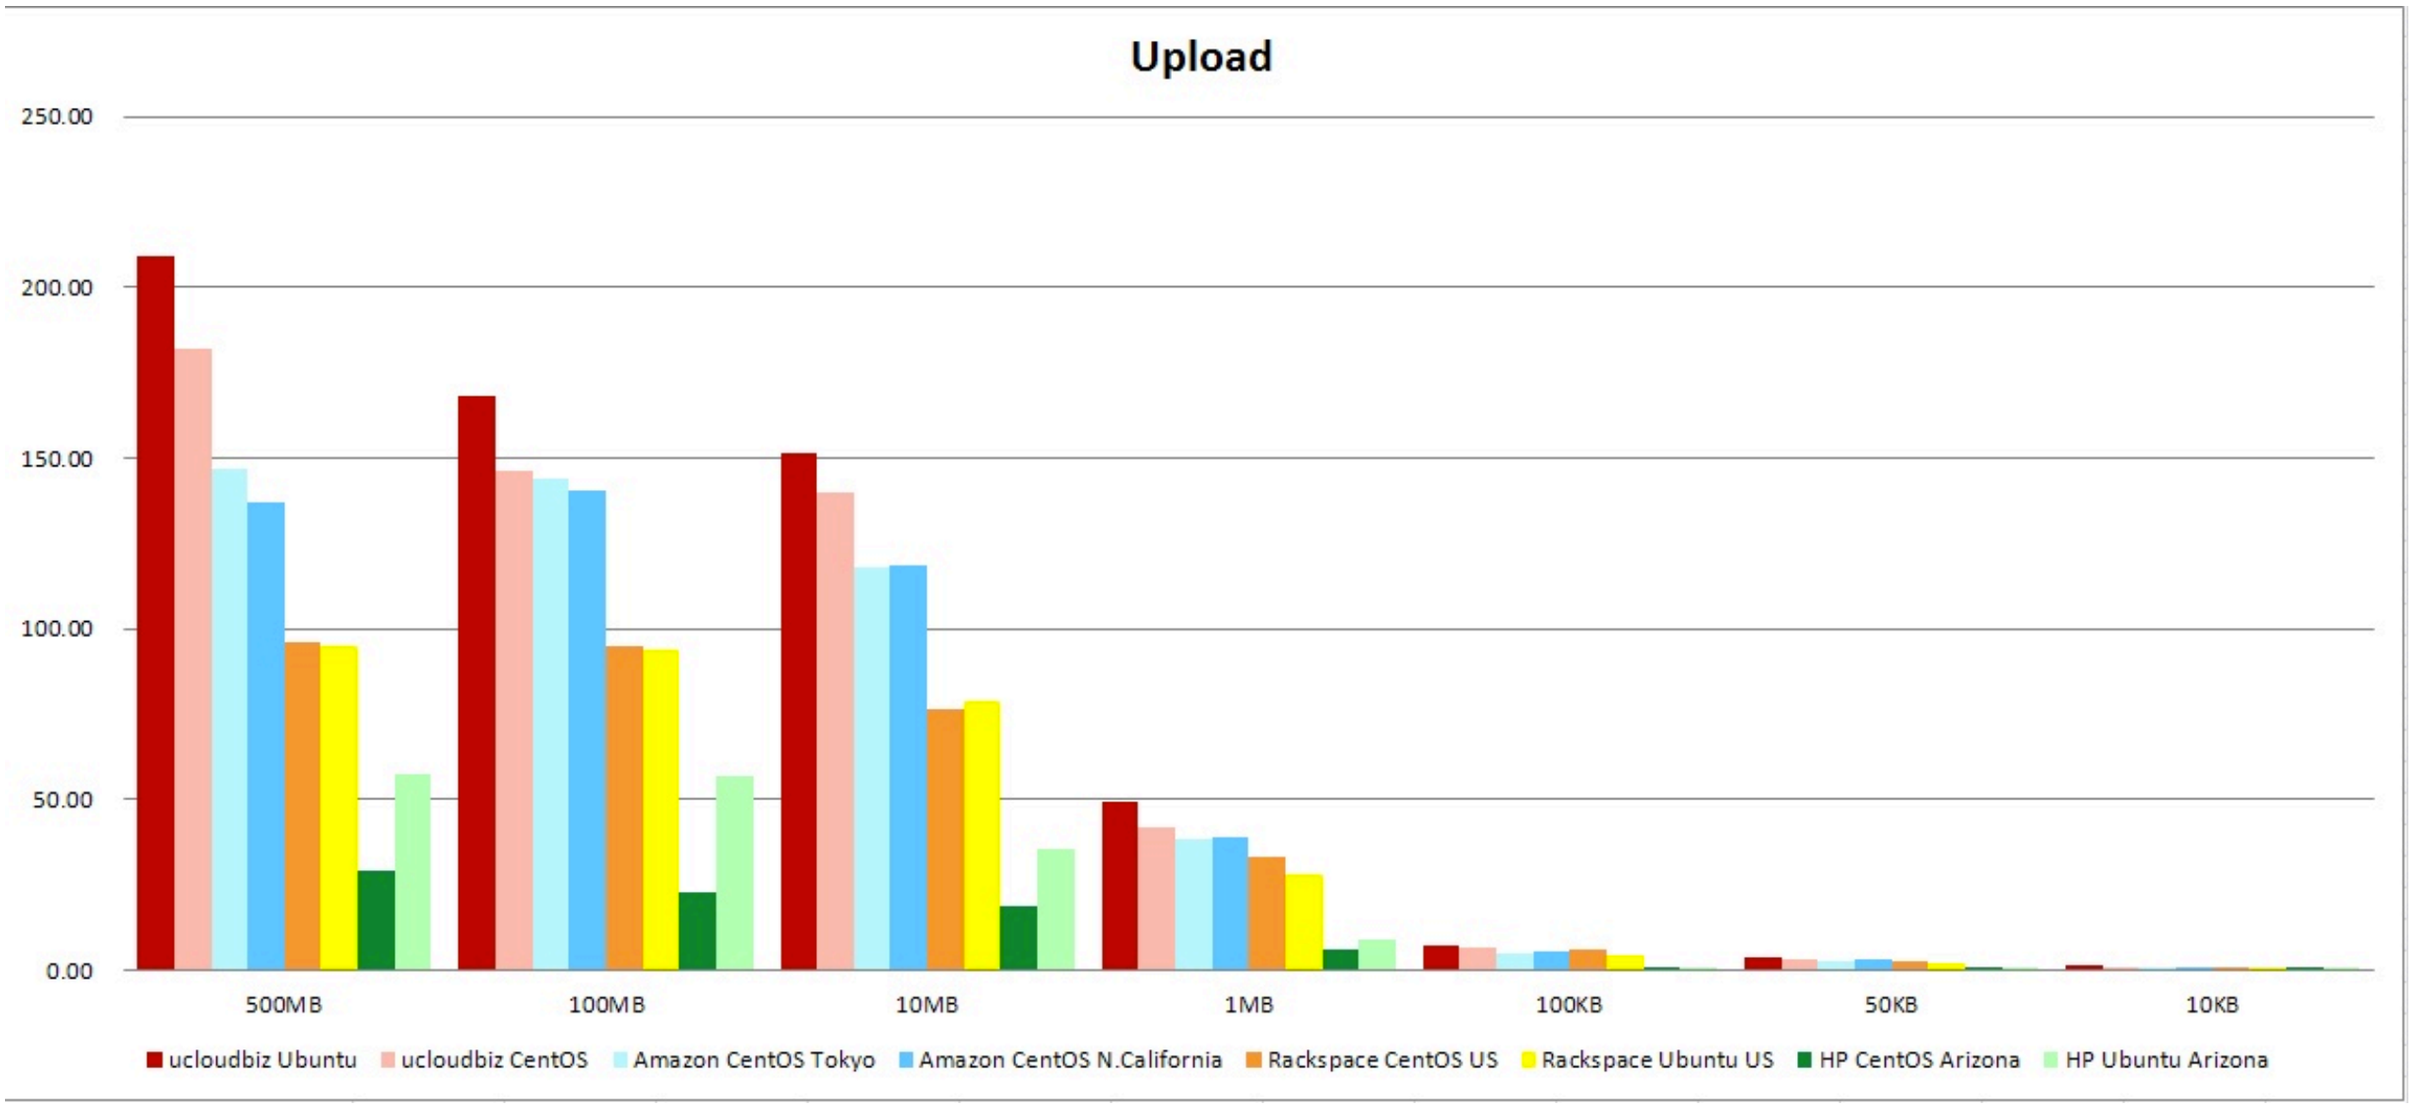

kt

openstack™

Some Test Data

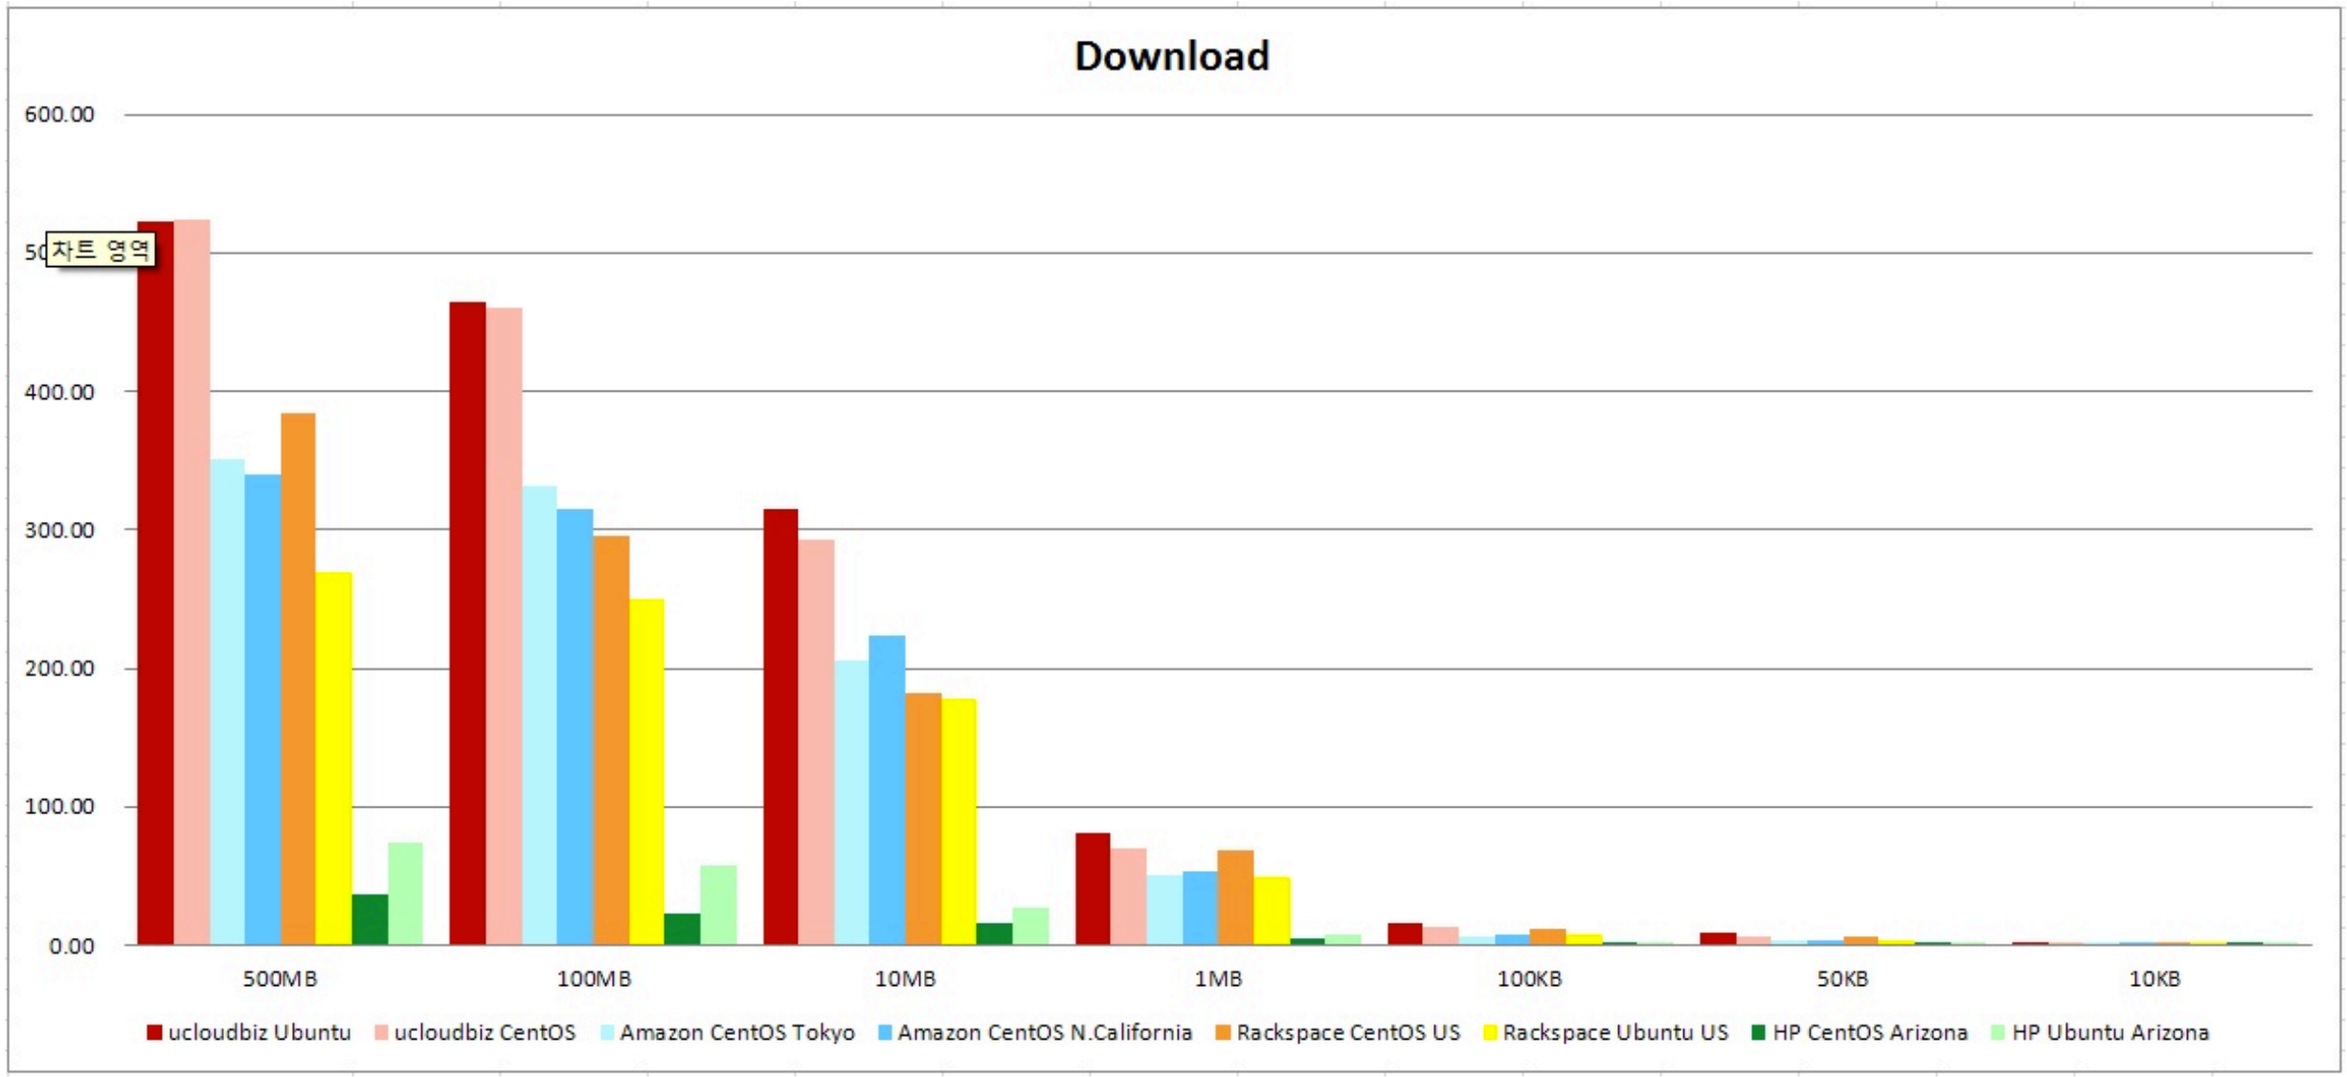

kt

openstack™

What I am trying to say is...

- Upload, Download, Big Size, Small Size, Many files in a container, Many containers, etc.
- There can be lots of use cases.
- Gather your usage data and Analyze patterns, then focus on what you need to improve the most.

Lesson Learned

- OpenStack Swift is a very stable and mature software stack.
- However, prepare for the various problems, when it deployed on Hardware

- In terms of business,
- People Does Not Know “Object Storage”

kt

openstack™

- It has been
  - NOT “selling a service”
  - BUT “explaining what it is and how to use”
- and frankly speaking,
  - KT has been experiencing difficulties to increase revenue from object storage service.

- Educate the customer and Find Business Model with them.
  - This is a new concept for the customer. You should put “object storage” concept at their initial service design
- Grow your potential customer pool with active support
  - AWS S3 took 3 years. Be patient.
  - Various promotion for younger developers & students

- Finally, OpenStack Community is Fantastic.

Questions?

if you have any question on openstack swift deployment, just send me email. :)

kt

openstack™