SDC
STORAGE DEVELOPER CONFERENCE
SNIA ■ SANTA CLARA, 2014

Benchmarking Cloud Storage through a Standard Approach

Wang, Yaguang
Intel Corporation

Legal Notices and Disclaimers

- INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL® PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. INTEL PRODUCTS ARE NOT INTENDED FOR USE IN MEDICAL, LIFE SAVING, OR LIFE SUSTAINING APPLICATIONS.
- Intel may make changes to specifications and product descriptions at any time, without notice.
- All products, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice.
- Intel, processors, chipsets, and desktop boards may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.
- Any code names featured are used internally within Intel to identify products that are in development and not yet publicly announced for release. Customers, licensees and other third parties are not authorized by Intel to use code names in advertising, promotion or marketing of any product or services and any such use of Intel's internal code names is at the sole risk of the user.
- Intel product plans in this presentation do not constitute Intel plan of record product roadmaps. Please contact your Intel representative to obtain Intel's current plan of record product roadmaps.
- Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to http://www.intel.com/performance
- Intel, Intel Inside, the Intel logo, Intel Core, Intel Atom, Pentium and UltraBook are trademarks of Intel Corporation in the United States and other countries.
- **Material in this presentation is intended as product positioning and *not* approved end user messaging.**
- **This document contains information on products in the design phase of development.**
- *Other names and brands may be claimed as the property of others.
-

Agenda

- COSBench Overview & Update
- CDMI Overview
- CDMI in COSBench
- Testing swift through CDMI
- Current Status

COSBench Overview

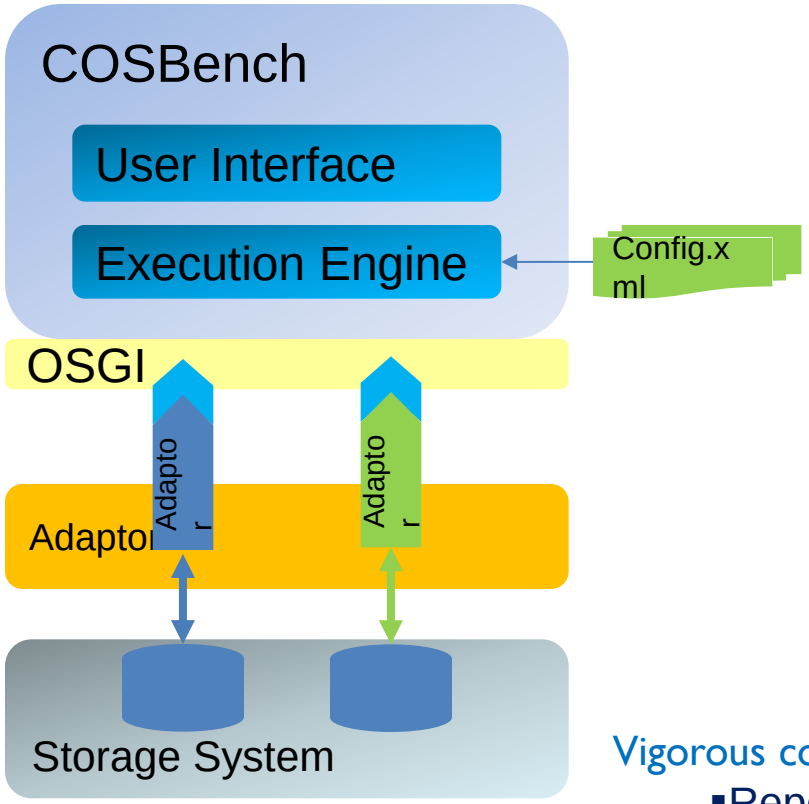

A Cloud Object Storage Benchmarking Tool Announced at the OpenStack design summit 2013 as Open source project.

Supports multiple Object Storage backends

- Swift, Ceph, Amazon S3*, Amplidata*, Scality*, CDMI*

Distributed Model for Traffic Generation

- Drivers: Workload generator
- Controller: coordinates among drivers, collects & aggregates results.

Vigorous code repository and community

- Repository: https://github.com/intel-cloud/cosbench
- Mailing-List: http://cosbench.1094679.n5.nabble.com

User Adoption

Quanta Computer, wiwynn, DELL, Seagate, intel, HGST a Western Digital company

TEAMSUN 华胜天成, 中国电信 CHINA TELECOM, at&t, SAKURA Internet, 九州云 99cloud.net, eNovance Cloud & Managed Services Provider

zmanda A Carbonite Company, EVault A Seagate Company, HITACHI Inspire the Next, AMPLIDATA, SCALITY, vmware

Australian National University, redhat, 上海交通大学, BBC, SAMSUNG

(More are anonymous.)

Usage Models

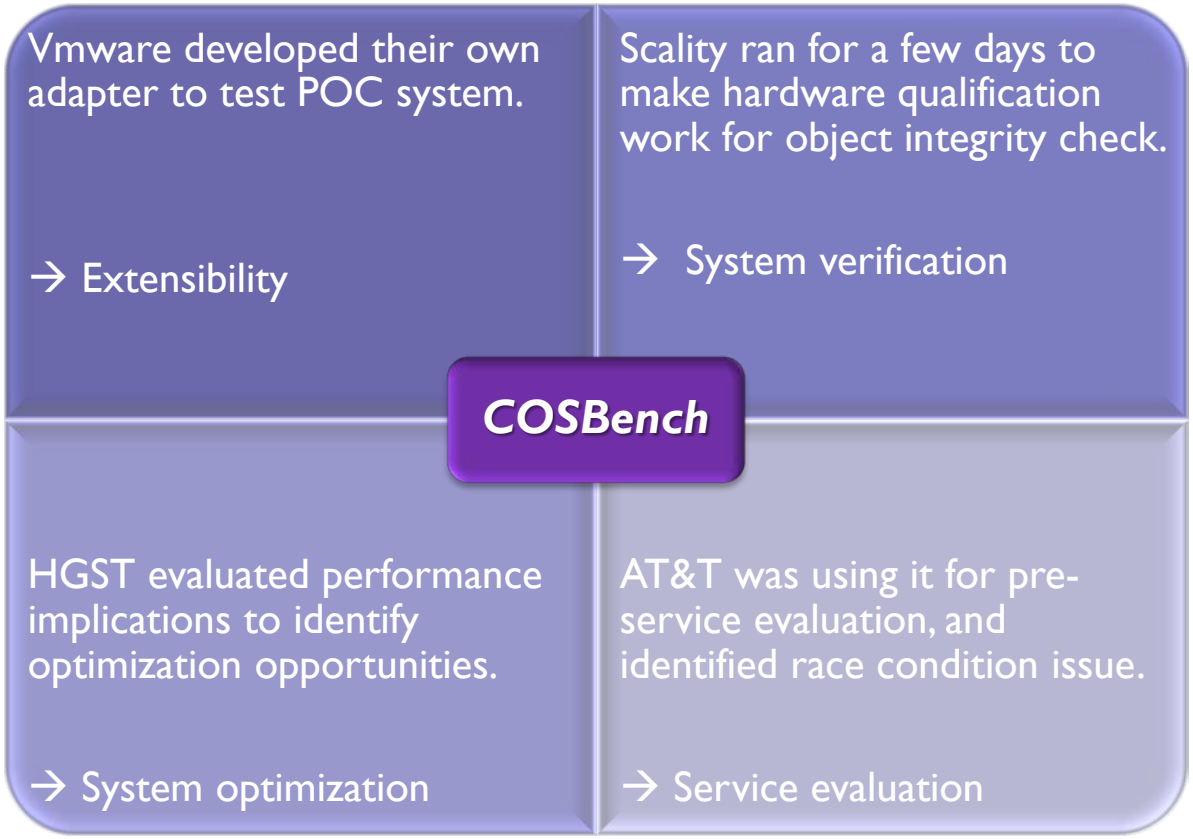

Major Progress

- New Object Store/Auth
  - http basic/digest
  - S3/Ceph (Librados based, or Radosgw based)/Scality sproxyd/CDMI (swift through CDMI middleware, scality)
- Backend
  - Core Functionalities
    - New selectors / new operator
    - Object integrity checking
    - Response time breakdown
  - Storage policy support for swift
  - Enhance on error handling
  - Trigger supporting
- Frontend (UI)
  - Advanced Workload Configuration UI
    - Adds Batch Test Configuration to COSBench
    - Makes COSBench workload configuration more like Iometer
  - Workload management
    - Support to archive or reload workload
    - Re-submit historical or archived workloads
    - Workload reordering

90+ issues resolved

CDMI Overview

- CDMI: Cloud Data Management Interface
- A Specification defined by SNIA, and accepted by ISO as standard.

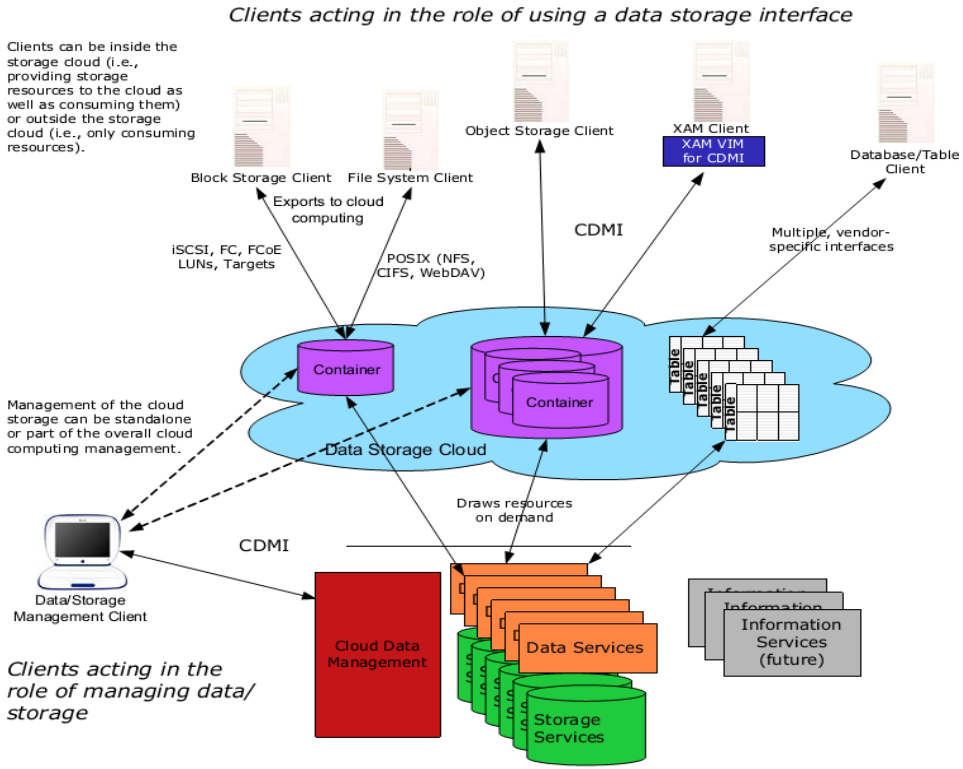

Benefits from standard approach

Nirvanix Inc. Key Developments

Nirvanix Announces Shut Down of its Service

Sep 30 13

Nirvanix has officially shut down its service, the company said that it is actively winding down its business, and is scrambling to transfer data on the cloud storage service to other providers.

- Easier manage different data source, reduce vendor lock-in.
- Consistent and controlled protocol
  - 100+ APIs from programmable web (http://www.programmableweb.com/apitag/storage)
- Reuse development investment

Access Method

- Access Method
  - Path-based style
    - Container/object
  - ID-based style
    - UUID
    - ID can be applied to container or object depending on capabilities.
  - E.g.
    - http://cloud.example.com/root/MyDataObject.txt
    - http://cloud.example.com/root/cdmi_objectid/00006FFD001001CCE3B2B4F602032653

Content Format

- Non-CDMI content type
  - Raw content in body
  - Similar to S3
- CDMI-content type
  - Content is wrapped into a json structure in body
  - Could be with different encoding like UTF-8/Base64…
  - See next page.
- Multi-part
  - Json structure is wrapped into multiple parts:
    - header in the first part
    - content in following part (similar to non-cdmi form)
    - ending bracket “}” in part 3
    - A zero-sized part as the closing part.

Message Format

- Request Line
- Header
- Metadata
- Content

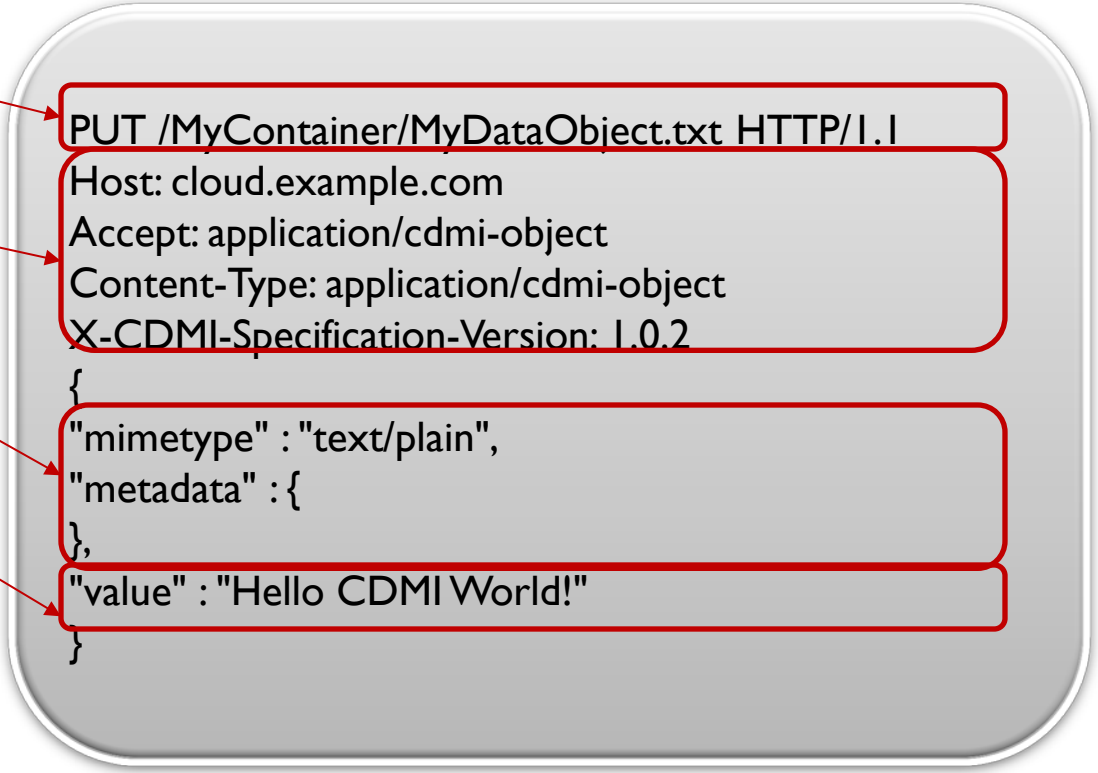

Authentication API

- Http Basic & Digest
- Token-Based
  - E.g., OpenStack swift authentication.

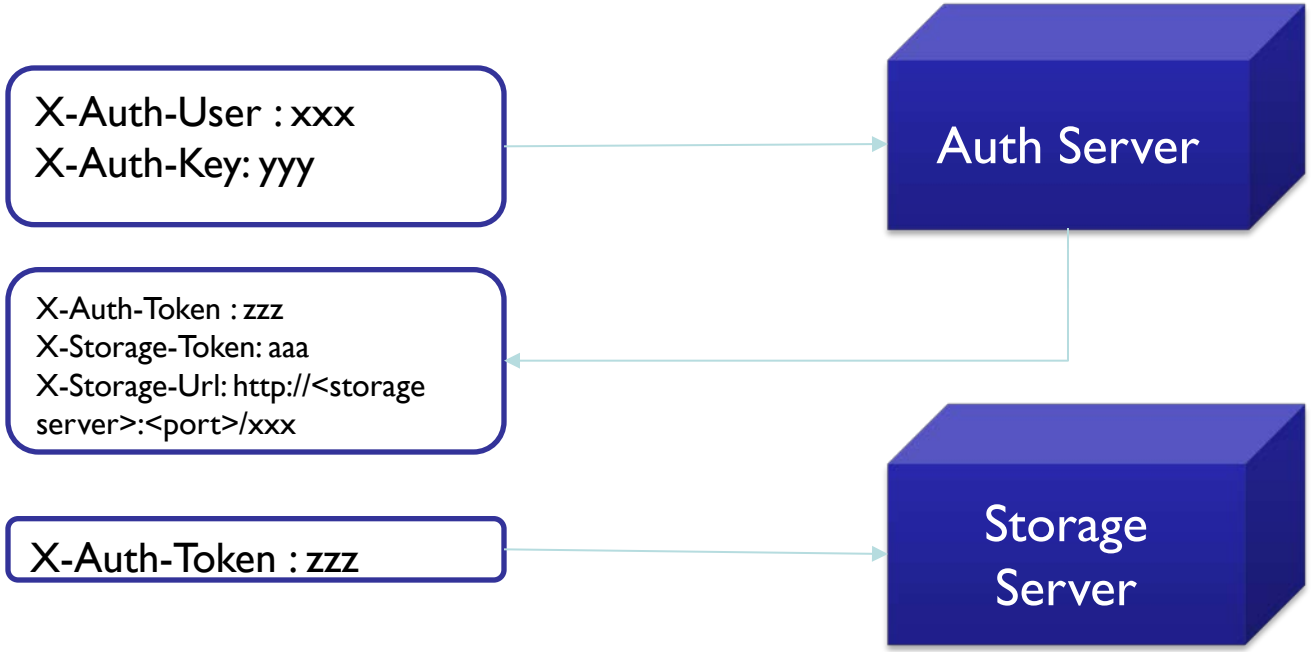

Storage API

- Storage API
  - Import parameter list
    - Two basic paramters
      - Token
      - endpoint
    - From the return of auth system
      - Openstack (auth_token, storage_url)
    - Provided by user
      - Direct (auth and storage system sits together).
      - Security depends on Transport like TLS.
    - Combined
      - CDMI
        - Token: http basic/digest/hmac…
        - Endpoint: user

CDMI in COSBench

- CDMI Util (utilitiy)
- CDMI Base (general)
- CDMI Flavor (vendor specific)

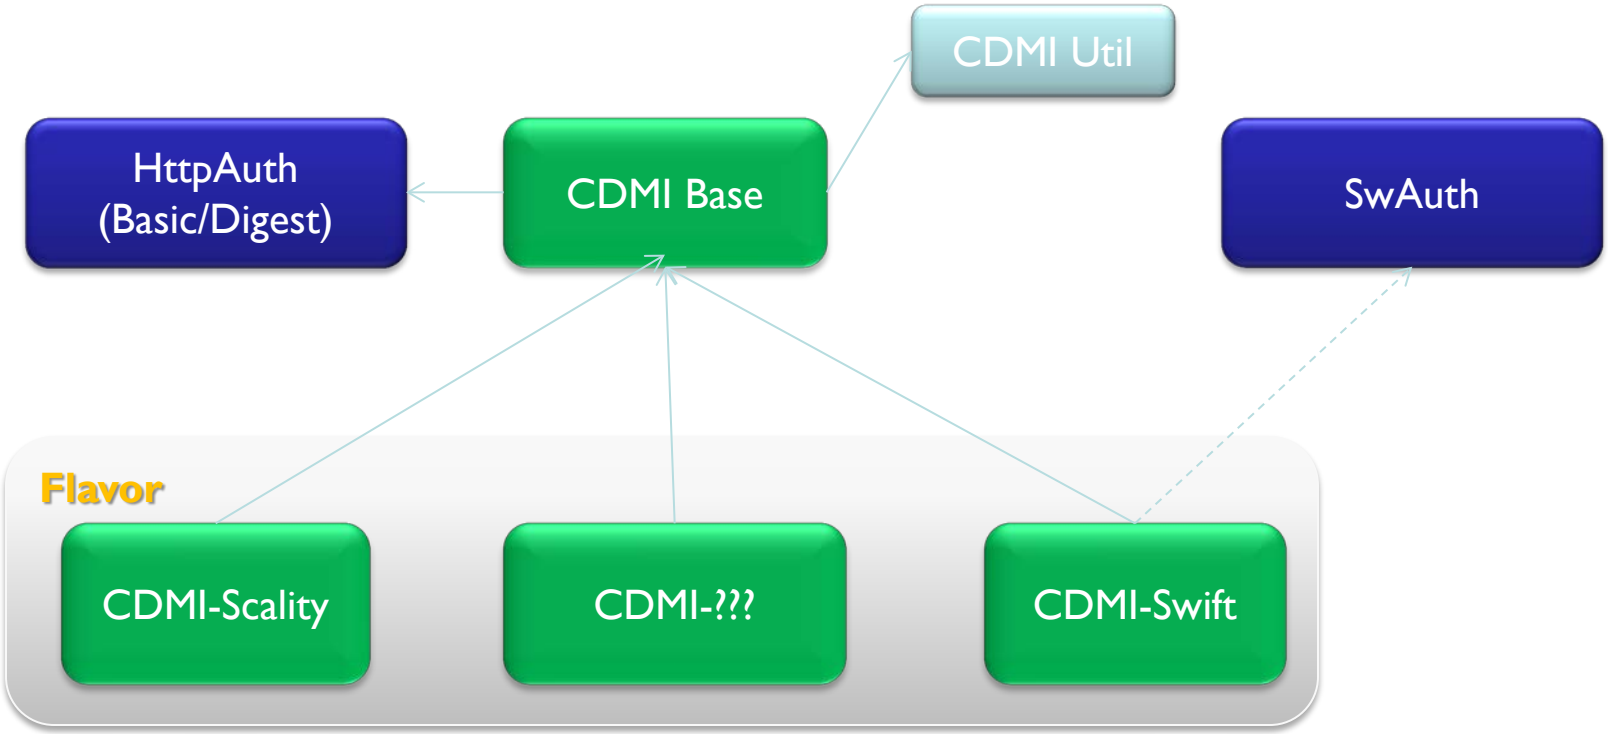

Using CDMI adapters

- CDMI-Base

<storage type="cdmi" config="type=<cdmi|non-cdmi; custom_headers=<header:value_reference>" />

Parameter	Type	Default	Comment
type	String	"cdmi"	Options: "cdmi" or "non-cdmi", it indicates the content type to be used, "cdmi" means the storage access will follow cdmi content type, "non-cdmi" means the storage access will follow non-cdmi content type.
Customer_headers	String		This is an experimental parameter to see if possible to support cdmi derivatives, which may require additional headers. The parameter may be removed without notification.

- CDMI-Swift

<storage type="cdmi_swift" />

Testing swift through CDMI

- Swift + CDMI middleware
  - https://github.com/osaddon/cdmi
- Authentication:
  - Swift specific swauth/keystone
- Two approaches:
  - Through CDMI-Swift adapter
  - Through CDMI-Base adapter, but with "custom_headers" parameter

Custom headers

custom_headers=<header:value_reference> →

Config="custom_headers=X-Auth-Token:token"

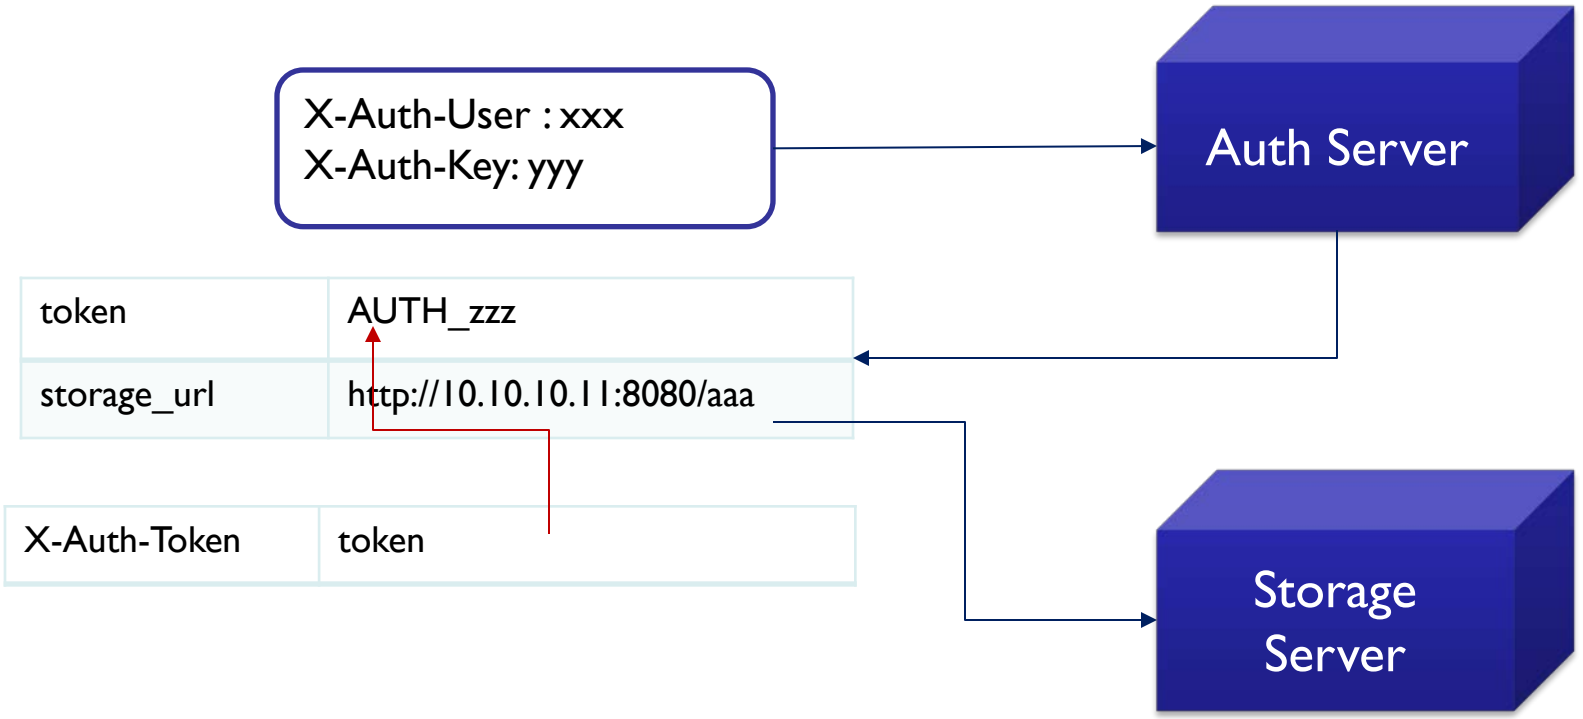

Current Status

- Path-Style & ID-Style
  - Only Path-Style so far
  - **ID-Style supporting has relevant impact to current design, as it requires to track all IDs.**
- CDMI vs non-CDMI content type vs multi-part
  - CDMI and non-CDMI content type are tested with Openstack Swift + CDMI middleware
  - **No multi-part code yet.**
- http basic & digest or token-based
  - Basic and Digest are tested with Tomcat web server separately.
  - Token-based is tested with Openstack swift.
- More flavor adapters

Call for Action

- Welcome more CDMI vendors to contribute flavor adapters into COSBench, or try COSBench to help identify issues.
- Contact me for any COSBench related issues or suggestions:

  yaguang.wang@intel.com