Performance and Cost Evaluation of Public Cloud Cold Storage Services for Astronomy Data Archive and Analysis

Hiroshi Yoshida, Kento Aida

Center for Cloud Research and Development,
National Institute of Informatics

Cloud Cold Storage Services

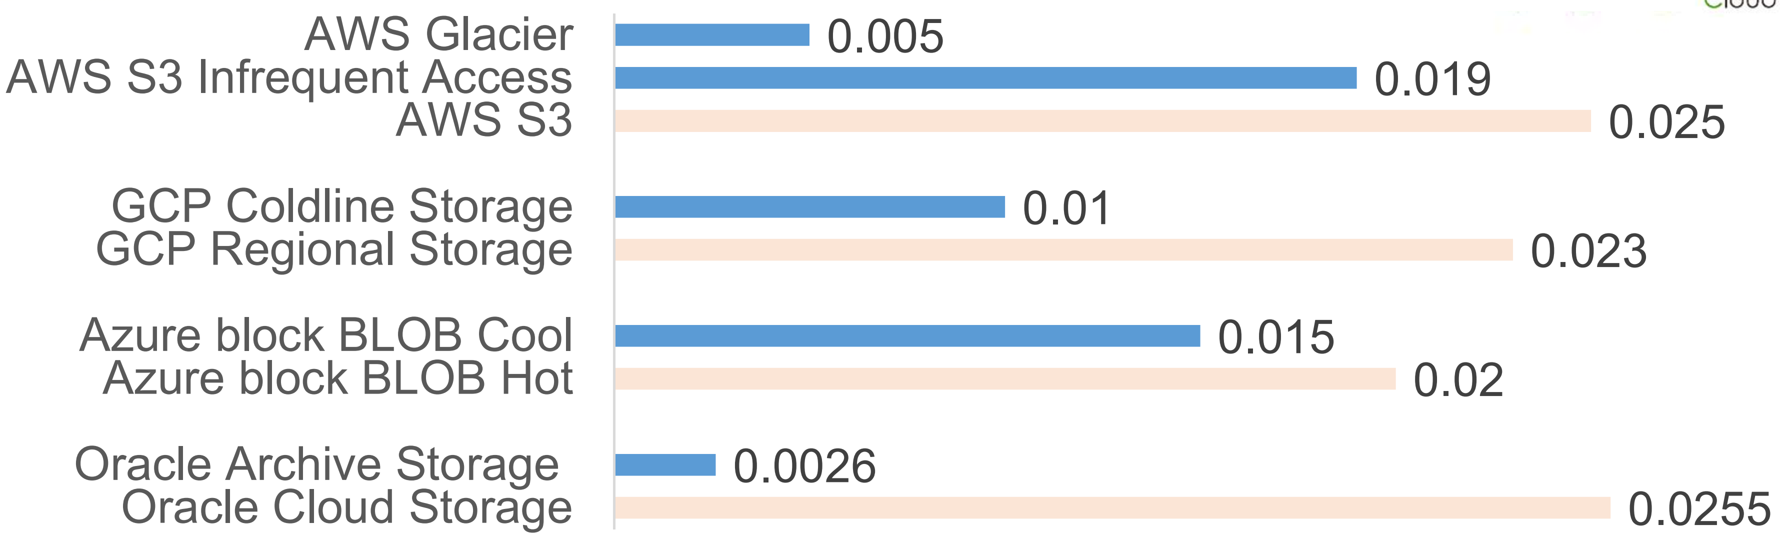

Price unit: USD/(GB*month) [as of July 1, 2018, **Japan region** except Oracle]

- Storage charge per GB*month is relatively inexpensive compared to standard object storage services.
  - 2/3 - 1/10 less expensive
- Drawbacks in exchange for inexpensive storage charge
  - Time consuming restoration process (hours)
  - Extra charge for data retrieval
  - Minimal retention period (30 – 90 days)
  - Limited performance, or extra charge for additional performance
  - Reduced availability

Experiment in Cloud Cold Storage Services

- Issue
  "Is it possible to adopt cloud cold storage to store a large amount of scientific research data for a long time?"
  - Reduction of storage management labor and TCO is expected.
  - However, as very few precedents exist, feasibility in terms of performance, manageability, and cost remains unknown.

- Aim of experiments
  To acquire practical information to determine the suitability of storing research data in cloud cold storage, and to design an overall data storage architecture.
  ➡ Experiments in cold storage services of commercial public clouds

  - Basic benchmark tests including storing up to 1PB data
  - Collaborative case study analyses using **actual research data and applications**
    - High-energy physics (High Energy Accelerator Research Organization)
    - **Astronomy (National Astronomical Observatory of Japan)**

Experiments Using Astronomy Data

GakuNin Cloud

- First step (FY2017 - FY2018)
  - Store observation and analysis data of ALMA radio telescope in cloud cold storage services
  - Evaluate performance, cost, and manageability by porting archive management system "NGAS" to AWS and storing archive data in S3-IA and Glacier

S3-IA: S3 Infrequent Access

- Second step (started in FY2019)
  - Analyze observation data on cloud instances
    - ✓ Expected advantages
      - Reduction of outbound data transfer cost
      - Flexible extension of compute resources in case of on-premise resource shortage
  - Evaluate performance and cost by running analysis software package "CASA" on the instances of public cloud services
  - Investigate optimal selection and usage of instances based on the measurement results

Astronomy Data and Application

GakuNin Cloud

For archive

Data contents	ALMA radio telescope Observation/analysis data provided by NAOJ *1
Quantity	**58.5TiB, 1,380,000 files**
Size	Average 44MiB (falls between smaller than 1MiB and larger than 100GiB)
Application	Archive management: NGAS (Next Generation Archive System)

*1: National Astronomical Observatory of Japan

For analysis

Dataset #	Dataset id	Number of objects	Data size (GiB)
Small (analysis time≃ 1 hour)	**1**	**147**	**0.5**
	2	**267**	**0.4**
	3	**99**	**0.6**
Medium (analysis time ≃5 hours)	**4**	**2,076**	**3.9**
	5	**4,000**	**3.5**
	6	**240**	**2.2**
Large (analysis time≃ 1 day)	**7**	**3,384**	**26.1**
	8	**3,879**	**16.7**
	9	**2,421**	**9.0**
Extra large (analysis time»1 day)	**10**	**456**	**87.3**

- Application: CASA (Common Astronomy Software Applications)

Experiment Environment

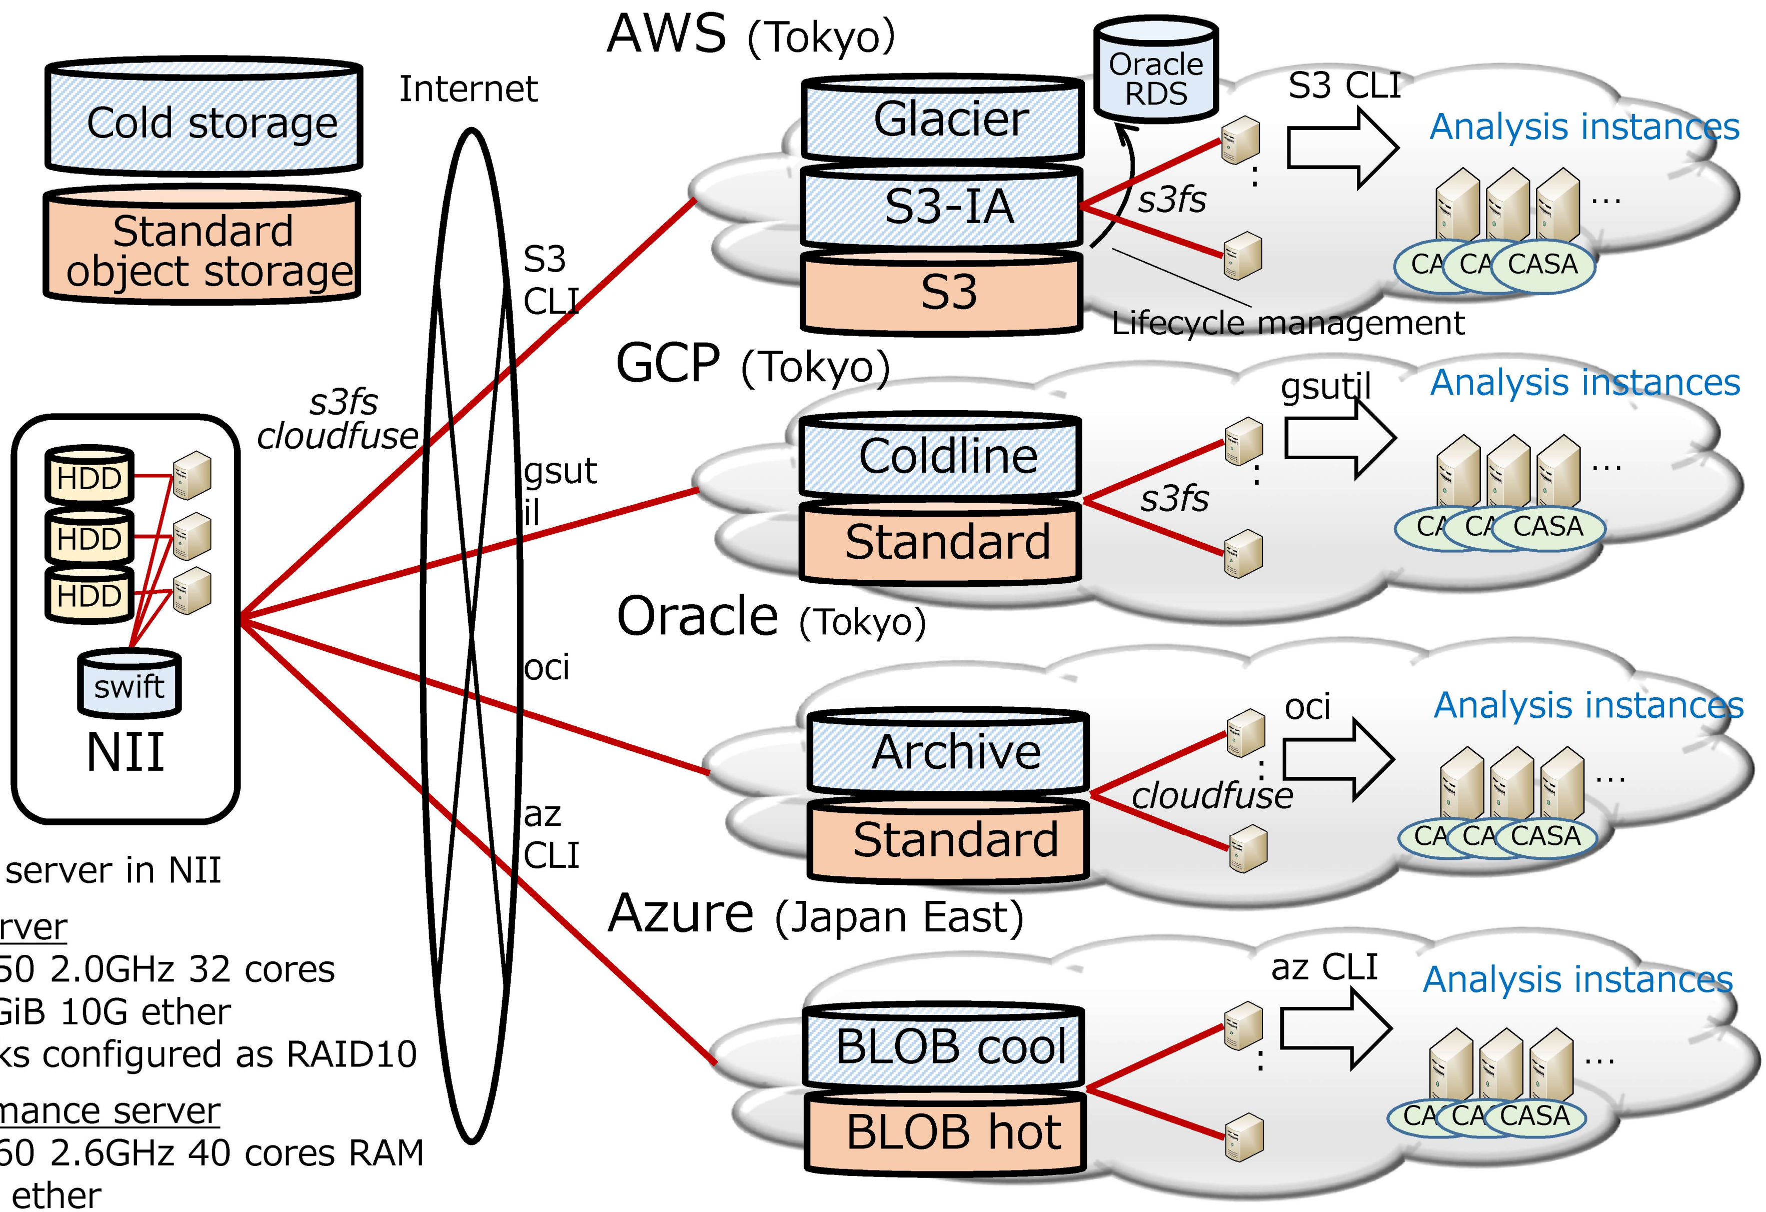

On-premise server in NII

<u>Standard server</u>
Xeon E5-2650 2.0GHz 32 cores
Memory 64GiB 10G ether
External disks configured as RAID10

<u>High performance server</u>
Xeon E5-2660 2.6GHz 40 cores RAM 128GiB 10G ether
External RAID connected via iSCSI

1st Step Experiment: NGAS on AWS

- Store observation and analysis data of ALMA radio telescope in cloud cold storage services
- Evaluate performance, cost, and manageability by porting archive management system "NGAS" to AWS and storing archive data in S3-IA and Glacier

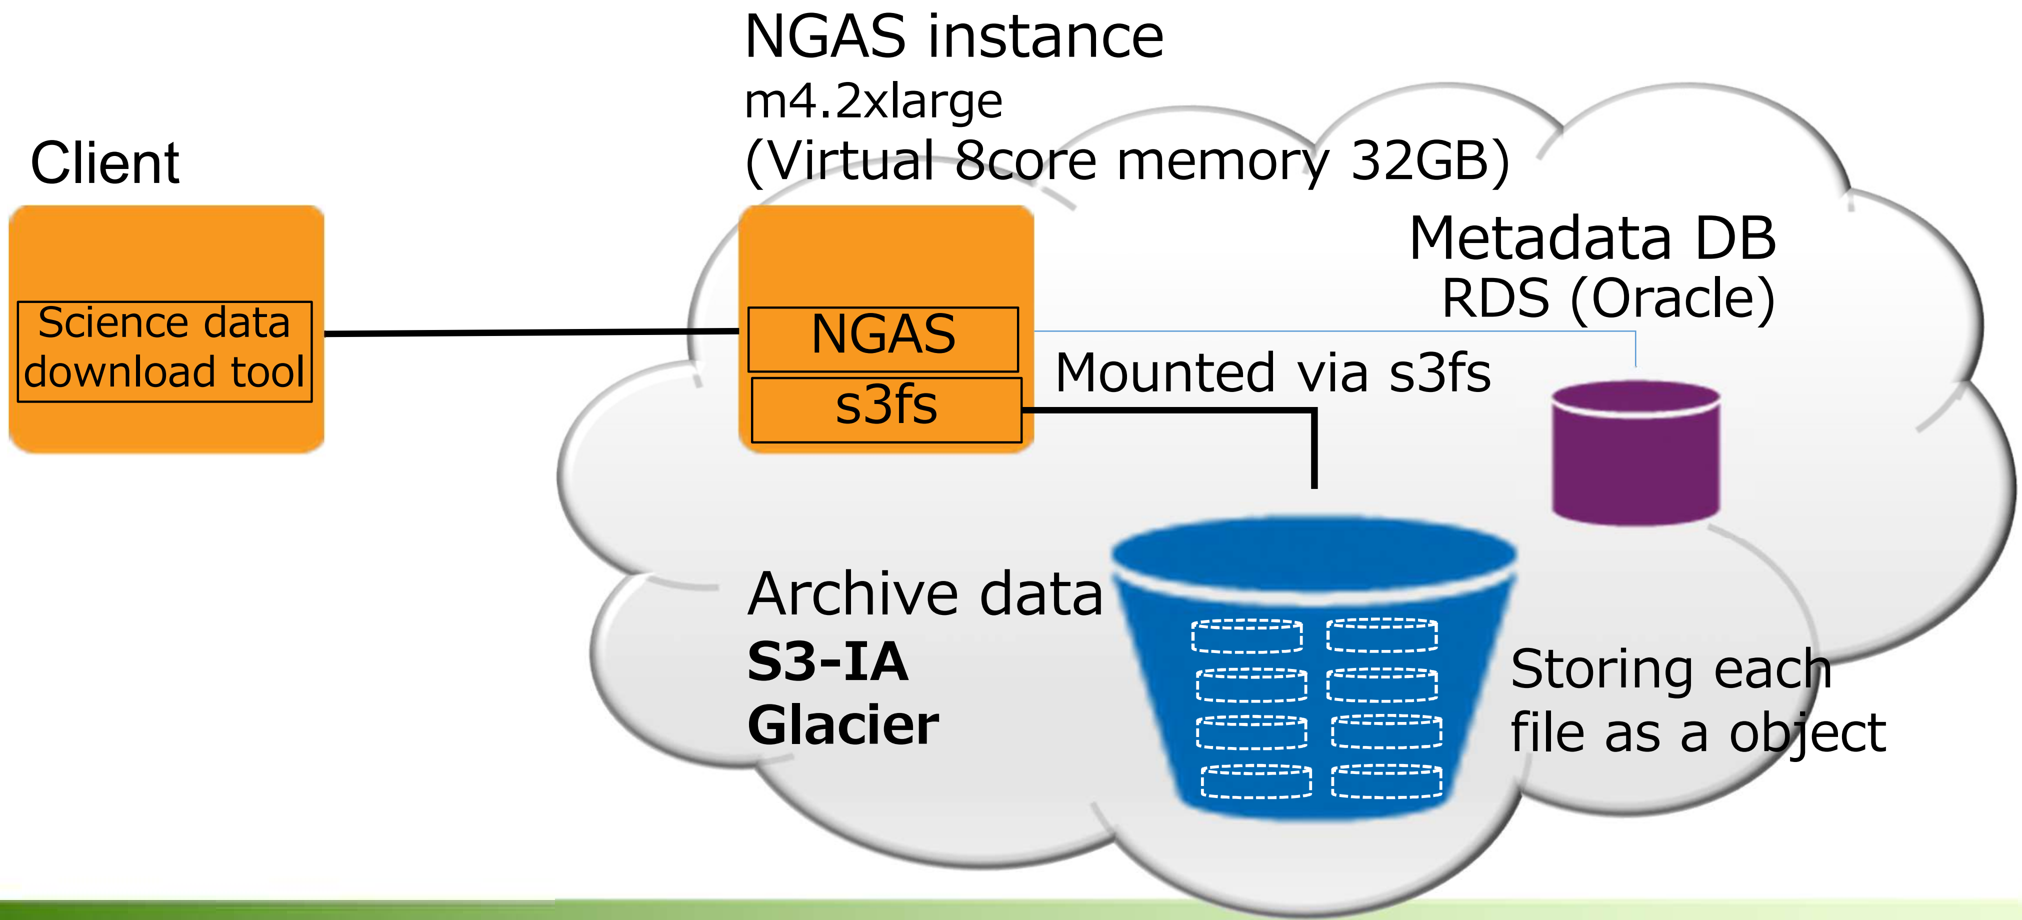

Performance of Uploading Archive Data

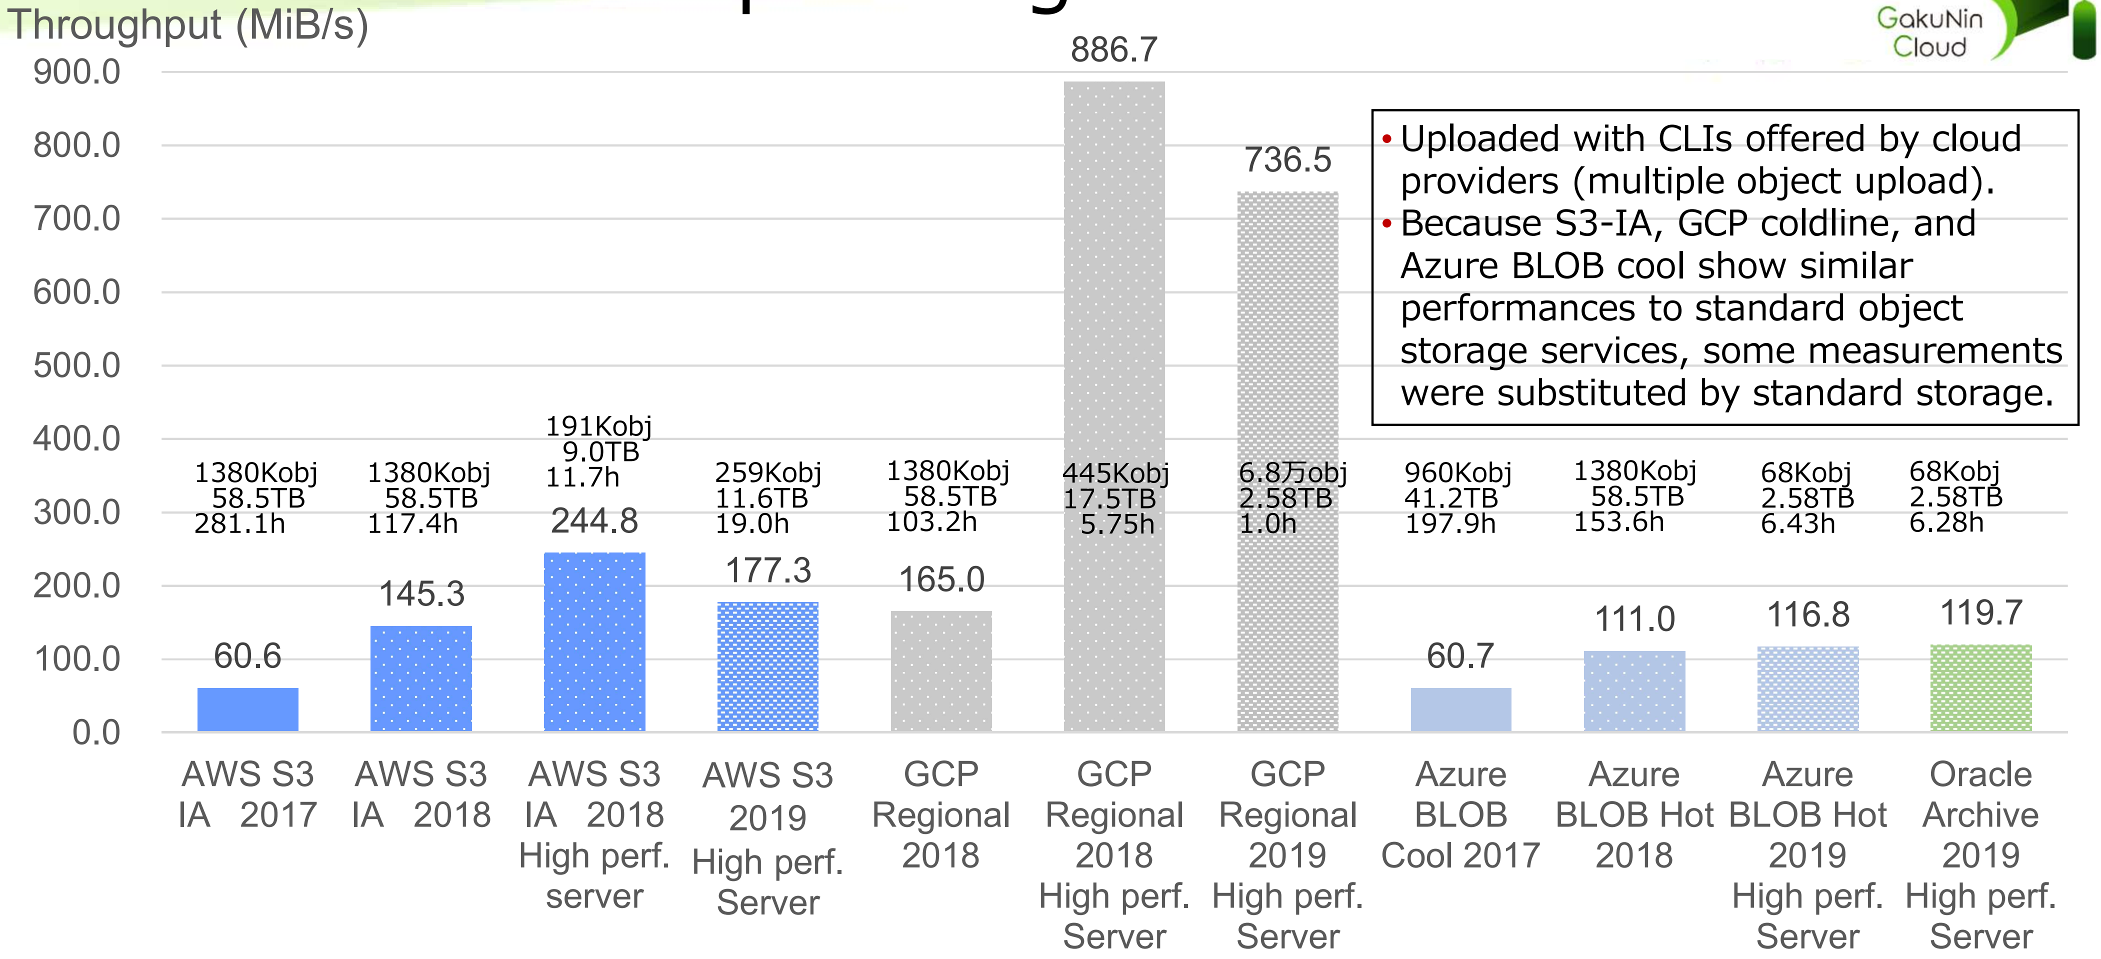

- Significant Performance improvement between 2017 and 2018
  - More performance improvement in the higher performance server cases
- Performance seems to be degraded between 2018 and 2019.
  - Internet connection of providers' data centers might be more congested because other measurements of internal transfer (copy objects between S3 buckets) in AWS show the same performance (≒ 240MB/s).

Cost Estimation of NGAS on AWS

- Store 900TiB archive data in S3- IA and retrieve 550TiB per year
- Store 900TiB archive data in Glacier and retrieve 550TiB per year

Archive data in S3-IA $268.516

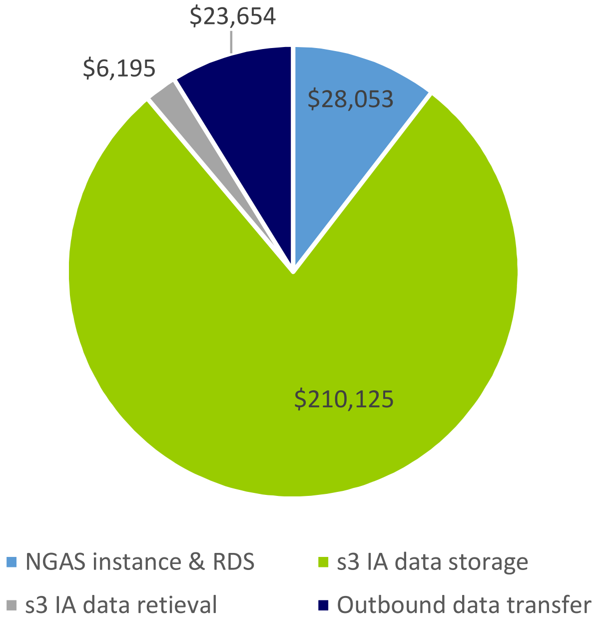

Archive data in Glacier $117,128

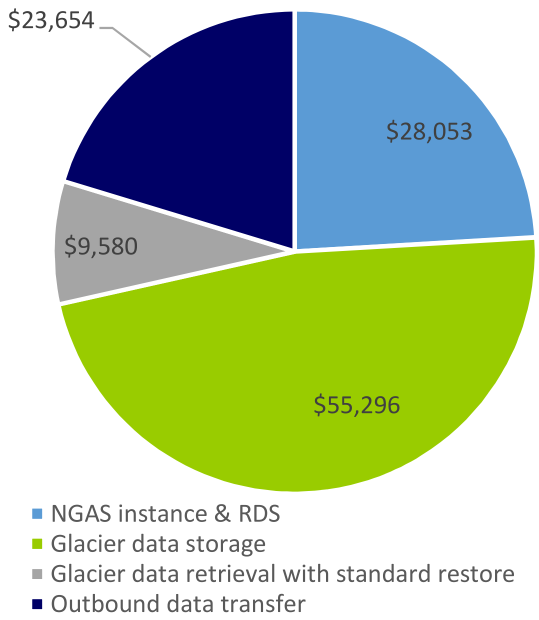

- Yearly operation cost: 249,033USD
- Storage cost is 84% of the total cost.

- Yearly operation cost: 95,435USD
- Restore time (standard: 3.3 hours) is required for each retrieval.

2nd Step Experiment: Analyses on AWS

- Evaluate data analysis performance and cost by running analysis software package "CASA" on the VM instances of public cloud services.
- Analysis datasets are actually stored in S3-IA and Glacier and retrieval operations from NGAS were simulated.

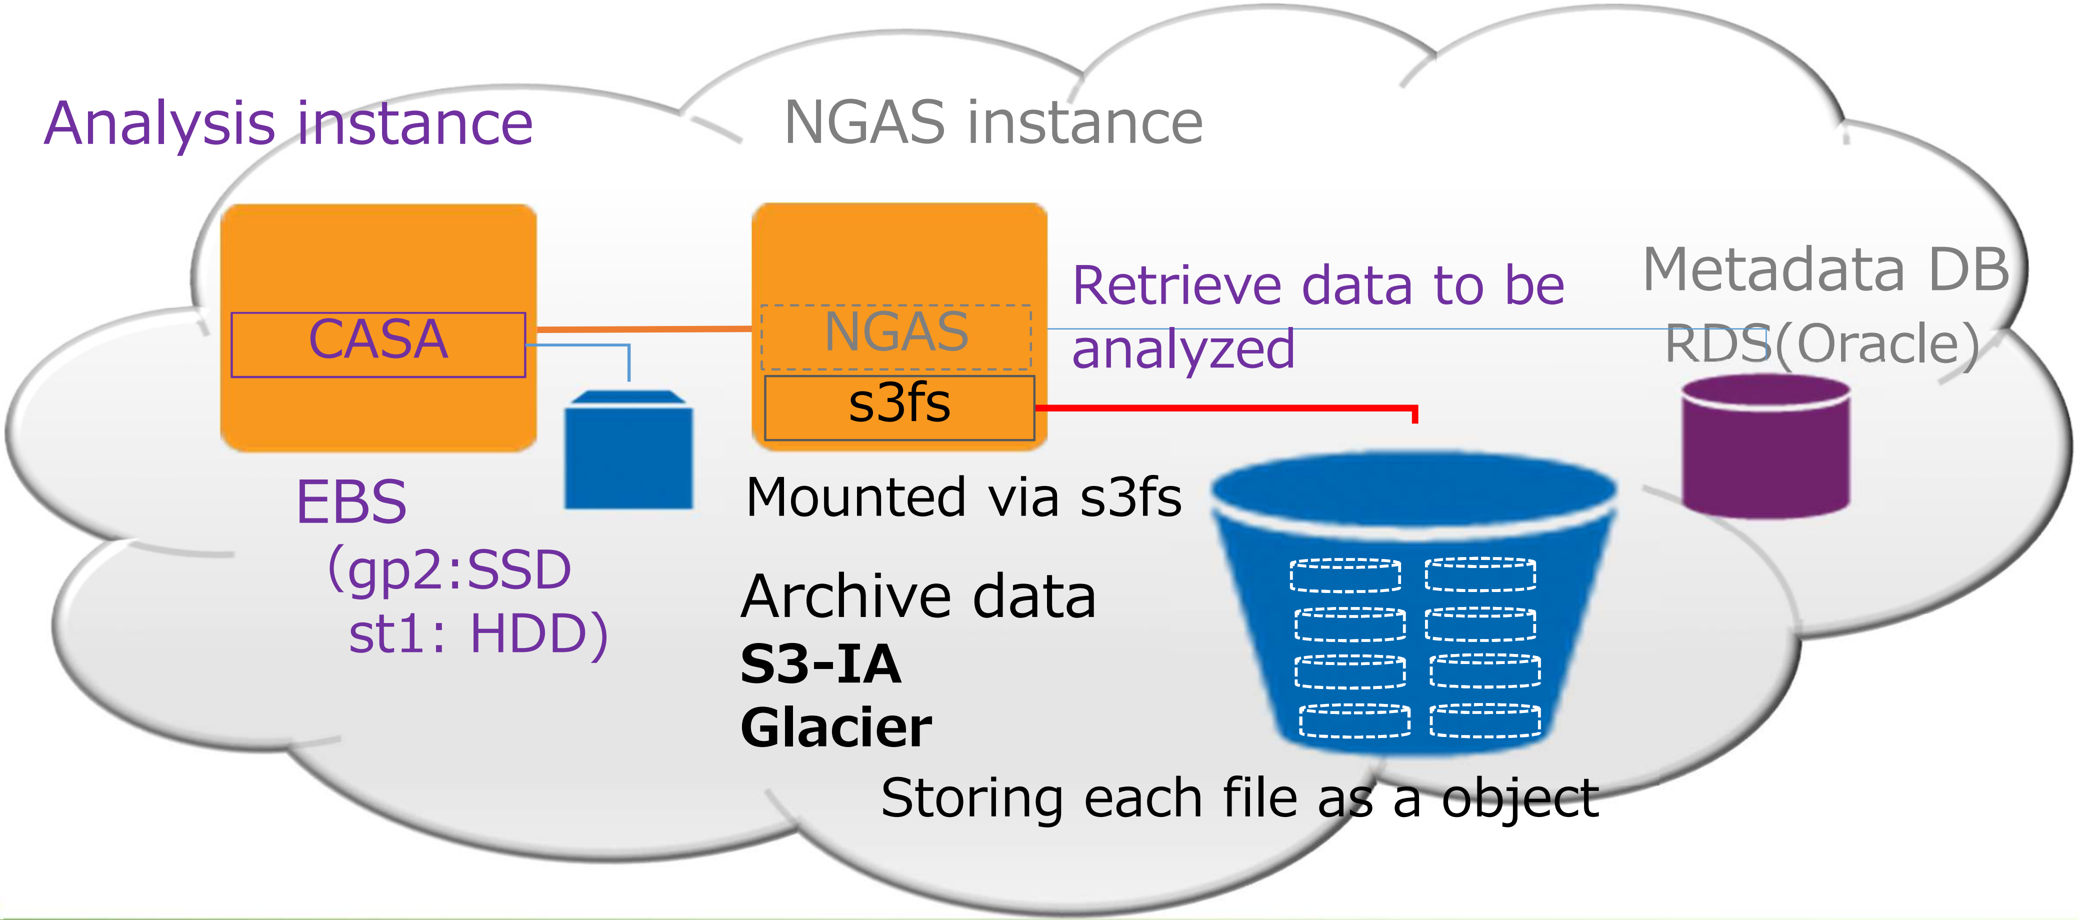

Performance and Cost of Data Analyses in AWS

- Elapsed times and costs of analyses using AWS instances with various memory capacities (32 – 244GB)
  - Achieved performance is suitable for practical use.
  - Performance differences are mainly caused by the differences of instance generation rather than memory capacity, because the chosen datasets don't require large memory capacity for analyses.

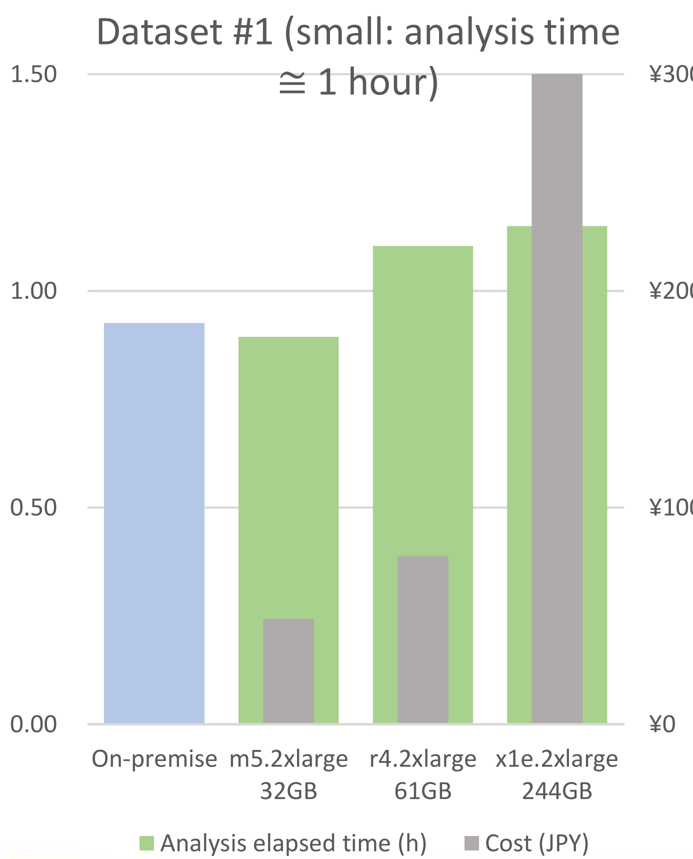

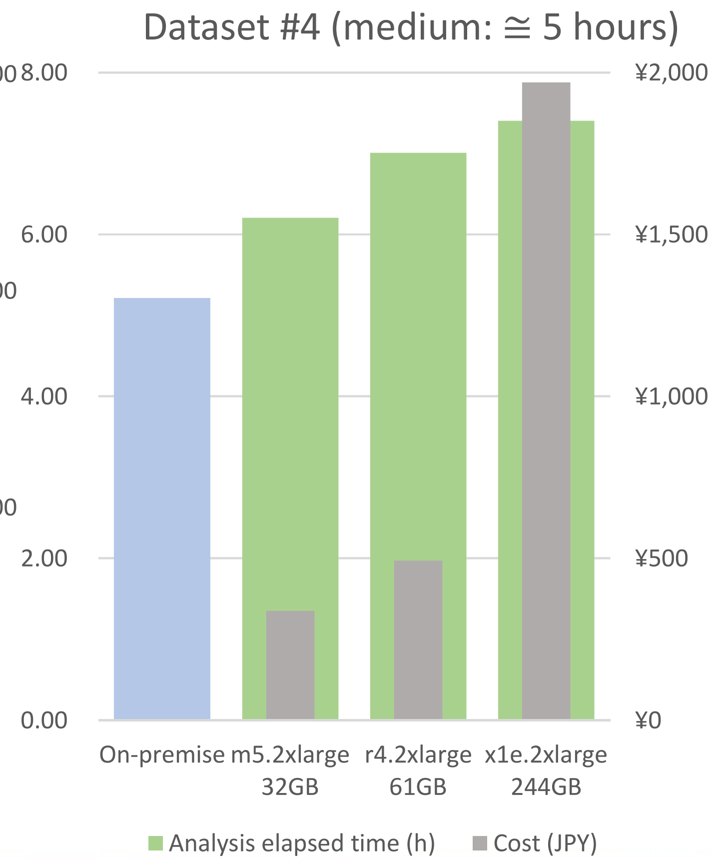

Hybrid Configuration

GakuNin Cloud

- To optimize performance and cost, it is worth considering a hybrid cloud configuration including on-premise environment and cloud environment with tiered storage.
- Data storing locations and analysis locations should be determined based on data retrieval/analysis performance, data transfer time/cost, cost of cloud/on-premise, and expandability/flexibility.

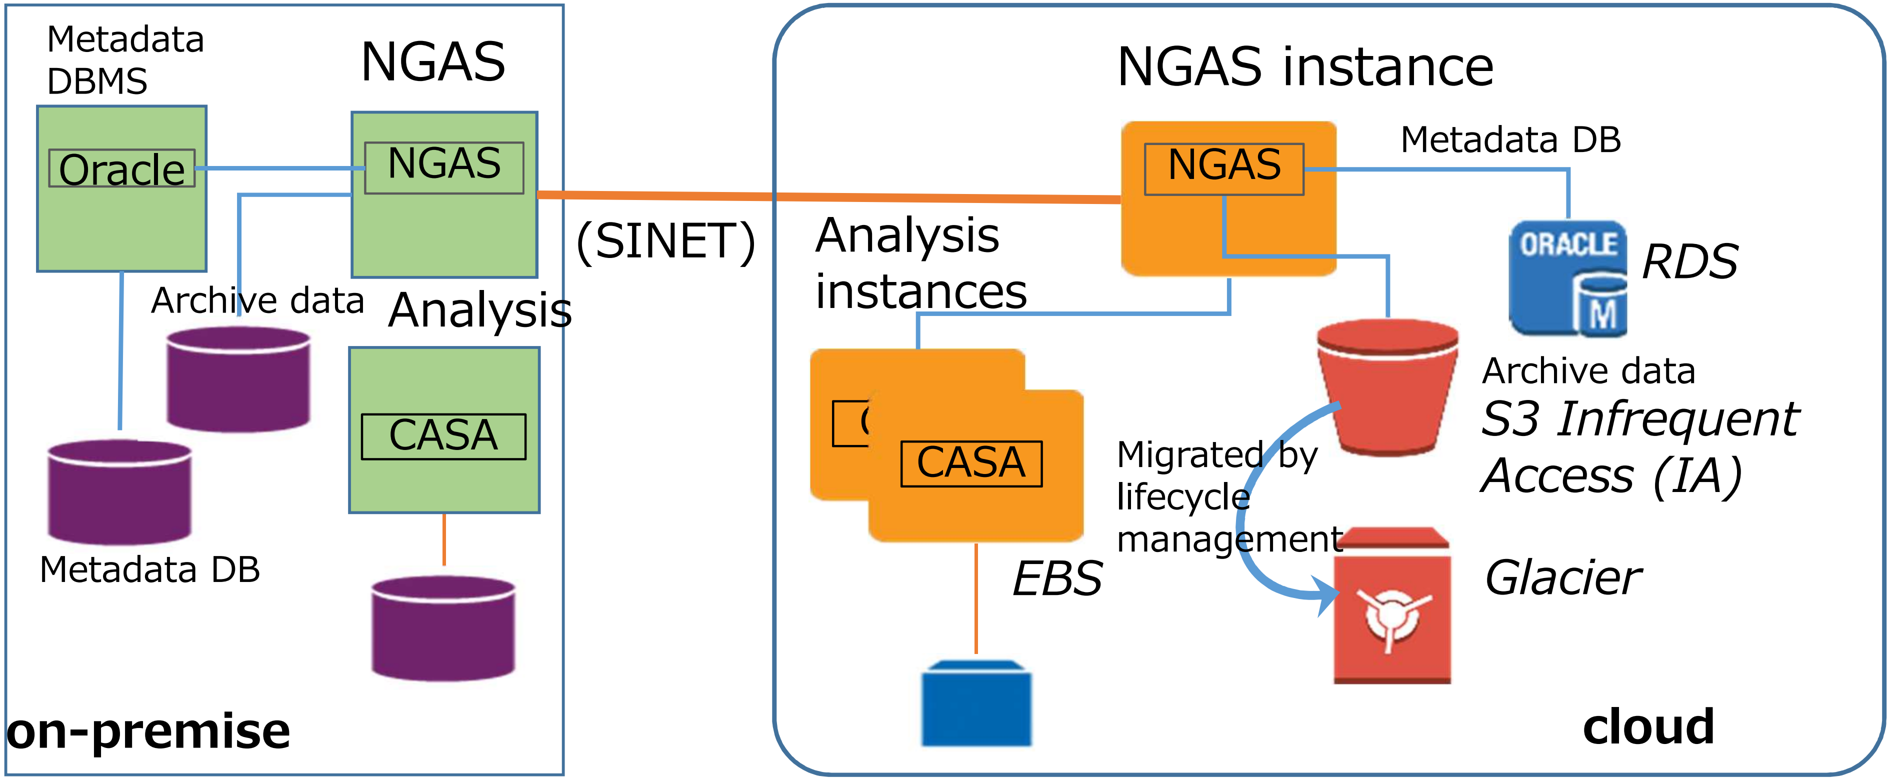

Next Steps

- Establish methodology to estimate required number of cores and memory capacity based on the dataset characteristics to choose optimal instance
- Investigate optimizations of applications and usages of services considering the characteristics of cloud cold storage
  - Optimize mapping between files and objects
    to improve handling of multiple objects
    and to reduce time and cost of restore processing
    e.g., 1 file to 1 object ➡ multiple files accessed at a time to 1 object
  - Adopt cloud-native object storage API
    to improve performance, stability, and predictability of charge
  - Instance Swapping
- **Share the practical information on performance and cost and the best practices of cloud usage with researchers of other scientific field**

GakuNin Cloud

NII

Inter-University Research Institute Corporation / Research Organization of Information and Systems

National Institute of Informatics