Apply Google Cloud Takeaways To Research

Group Meeting Presentation

Speaker: Yu Wu
10/13/2017

UCDAVIS
UNIVERSITY OF CALIFORNIA

Google Cloud Platform

- **Google Compute Engine** – IaaS providing virtual machines.
- **Google Container Engine** – Combination of IaaS with Docker container.
- **Google App Engine** – PaaS to host services.
- **Bigtable** – IaaS massively scalable NoSQL database.
- **BigQuery** – SaaS large-scale database analytics.
- **Google Cloud Functions** – FaaS providing serverless functions to be triggered by cloud events.
- **Google Cloud Datastore** - DBaaS providing a document-oriented database.
- **Google Storage** - IaaS providing online file and object storage.

Wikipedia, “Google Cloud Platform".

First Takeaway: It is all about APIs

- **HTTP REST APIs**
  - ✓ Bookstore/shelves/1 GET HTTP 1.0
- **Remote Procedure Call (RPC) APIs**
  - ✓ Bookstore.service/GetShelf(1)

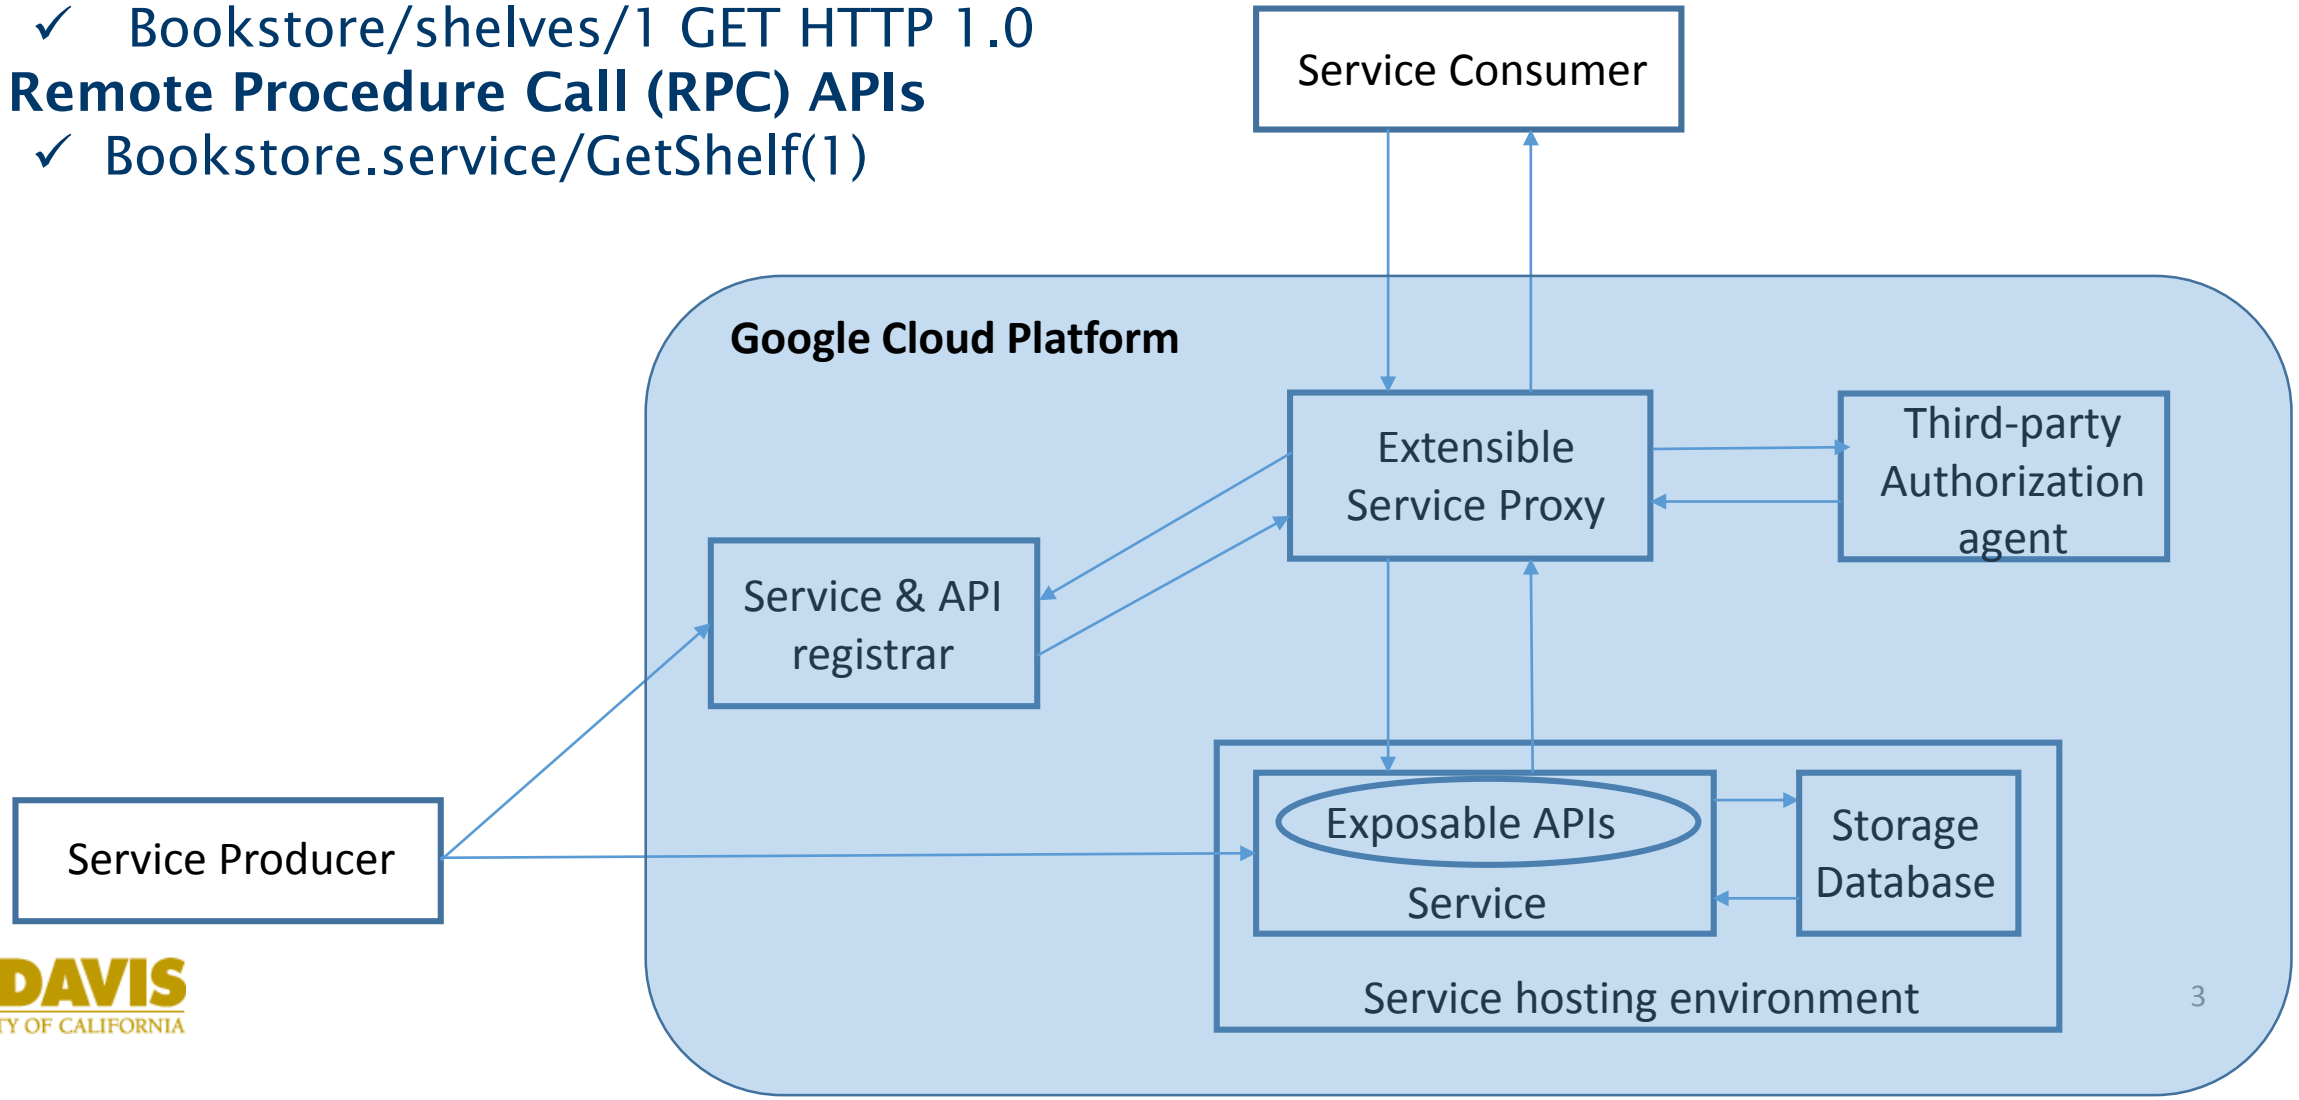

UCDAVIS
UNIVERSITY OF CALIFORNIA

Extensible Service Proxy

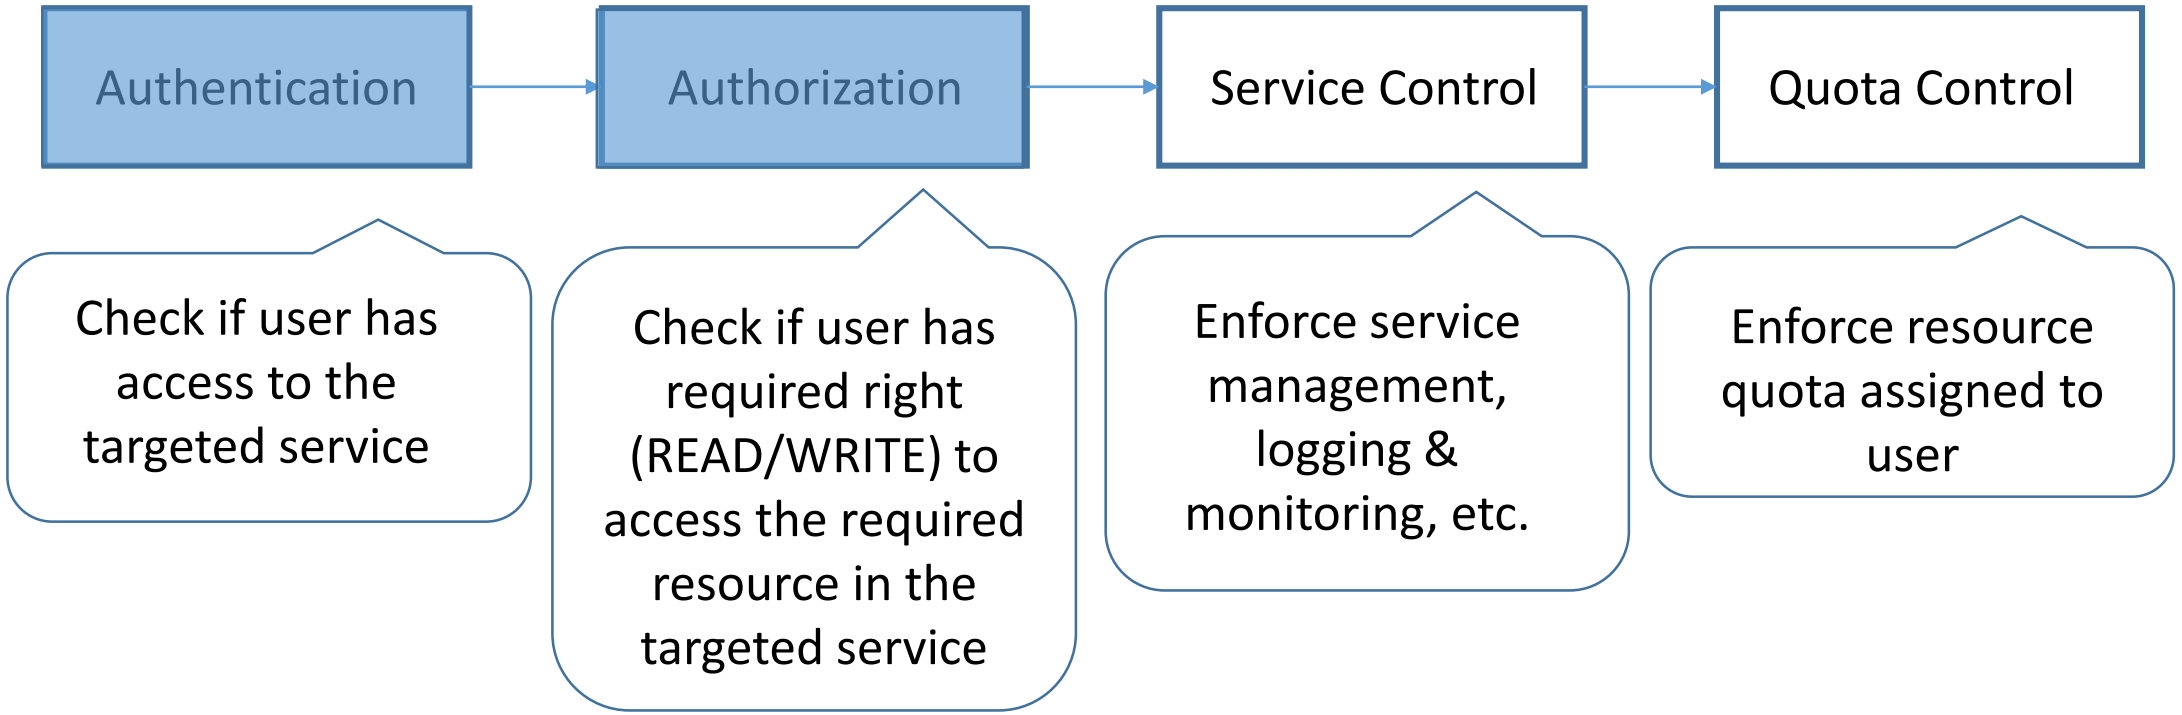

- **What I did in this summer:**
  - ✓ Integrated with Google Identity-Aware Proxy (IAP) to provide authentication.
  - ✓ Improved authorization performance.
  - ✓ Provided RPC support for authorization.

UCDAVIS
UNIVERSITY OF CALIFORNIA

Second Takeaway: Service is pay-per-use

- **Cloud Storage pricing**
  - ✓ Data storage
  - ✓ Network usage
  - ✓ Operational usage
  - ✓ Retrieval and early deletion

UCDAVIS
UNIVERSITY OF CALIFORNIA

Google Cloud Platform, “Google cloud storage pricing".

Data Storage

❑ Multi-regional location pricing

Multi-Regional Storage (per GB per Month)	Nearline Storage (per GB per Month)	Coldline Storage (per GB per Month)
$0.026	$0.01	$0.007

Storage plans for storing infrequently accessed data.

❑ Regional location pricing

Tokyo

Regional Storage (per GB per Month)	Nearline Storage (per GB per Month)	Coldline Storage (per GB per Month)
$0.023	$0.016	$0.01

UCDAVIS
UNIVERSITY OF CALIFORNIA

Network Usage

Monthly Usage	Network (Egress) Worldwide Destinations (excluding China & Australia, but including Hong Kong) (per GB)	Network (Egress) China Destinations (excluding Hong Kong) (per GB)	Network (Egress) Australia Destinations (per GB)	Network (Ingress)
0-1 TB	$0.14	$0.23	$0.19	Free
1-10 TB	$0.14	$0.22	$0.18	Free
10+ TB	$0.12	$0.20	$0.15	Free

Operational Usage

Storage Class	Class A operations (per 10,000 operations)	Class B operations (per 10,000 operations)	Free operations
Multi-Regional and **Regional**	$0.05	$0.004	Free
Nearline and **Durable Reduced Availability**	$0.10	$0.01	Free
Coldline	$0.10	$0.05	Free

- Class A operations include object adds and listings.
- Class B operations include object gets and metadata retrieving.
- Free operations include object deletes.

Retrieval And Early Deletion

	Nearline Storage	Coldline Storage
Data retrieval	$0.01 per GB	$0.05 per GB
Early deletion	30 days	90 days

- A retrieval cost applies when you read data or metadata that is stored as Nearline Storage or Coldline Storage. This cost is in addition to any network charges associated with reading the data.
- A minimum storage duration applies to data stored as Nearline Storage or Coldline Storage. You can delete the file before it has been stored for this duration, but you are charged as if the file was stored for the minimum duration.

A Detailed Pricing Example

Pricing Category	Type of Usage	Amount
Data storage	Multi-Regional Storage	60 TB
Data storage	Nearline Storage (stored in a multi-regional location)	100 TB
Network	Egress to the Americas and EMEA	25 TB
Network	Egress to the Asia-Pacific	25 TB
Network	Ingress	30 TB
Operations	Class A operations (object adds, bucket and object listings)	100,000 operations
Operations	Class B operations (object gets, retrieving bucket and object metadata)	10,000,000 operations
Nearline fees	Data retrieval (the Nearline Storage portion of your overall data egress)	10 TB

A Detailed Pricing Example Cont'd

Pricing Category	Calculation	Cost
Data Storage	60 TB (61440 GB) Multi-Regional Storage * $0.026 per GB	$1,597.44
Data Storage	100 TB (102400 GB) Nearline Storage * $0.01 per GB	$1,024.00
Network	25 TB total egress to the Americas and EMEA: (0-1 TB tier): 1TB (1024GB) egress * $0.12 per GB (1-10 TB tier): 9TB (9216GB) egress * $0.11 per GB (10+ TB tier): 15TB (15360GB) egress * $0.08 per GB	 $122.88 $1,013.76 $1,228.80
Network	25 TB total egress to the Asia-Pacific: (0-1 TB tier): 1TB (1024GB) egress * $0.12 per GB (1-10 TB tier): 9TB (9216GB) egress * $0.11 per GB (10+ TB tier): 15TB (15360GB) egress * $0.08 per GB	 $122.88 $1,013.76 $1,228.80
Network	Ingress is free	$0.00
Operations	100,000 Class A operations * $0.05 per 10,000 operations	$0.50
Operations	10,000,000 Class B operations * $0.004 per 10,000 operations	$4.00
Nearline fees	10 TB (10240 GB) data retrieval * 0.01 per GB	$102.40
Total		**$7459.22**

UC DAVIS
UNIVERSITY OF CALIFORNIA

How To Apply What I learned To Research

- **To be specific, how to apply what I learned to ICC paper journal extension?**

- **ICC paper recap:**
  - ✓ Given inputs such as traffic distribution, green energy availability distribution, disaster damage probability distribution, etc., the goal is to place the content in a way that minimizes brown energy consumption constrained by risk and latency.
  - ✓ We compared the performance of Content Replication (CR) scheme and Content Fragmentation (CF) scheme in terms of different metrics such as content popularity, risk threshold, content redundancy level.

How To Apply What I learned To Research

- **Existing research works fall into three categories:**
  - ✓ Looked into the content placement and rearrangement problem from content provider's perspective.
  - ✓ Tackled one goal at a time without shedding light on a more valuable plan in which multiple goals are considered together.
  - ✓ Treated the content placement problem and the rearrangement problem separately.

Journal Plan

- **What I want to do:**
  - ✓ Solve the content placement and rearrangement problem from cloud infrastructure provider's perspective.
    - ❖ Content provider tends to migrate to cloud.
    - ❖ Content provider has less knowledge than cloud infrastructure provider both data-wise and resource-wise to do global optimization.
    - ❖ As long as the required QoS is satisfied and the total cost remains low, content provider would like to focus more on service itself and let cloud infrastructure provider help manage the content service.
  - ✓ Tackle multiple goals together according to their importance order.
    - ❖ Revenue
    - ❖ QoS (resiliency and latency)
    - ❖ Green energy (responsibility taken by cloud infrastructure providers)
  - ✓ Design fast meta-heuristic algorithm for content placement, and adapt it to content rearrangement problem.

Goals

- **Revenue**
  - ✓ The profit collected from content provider is relevantly stable.
  - ✓ Maximizing revenue becomes **minimizing** Operational Expenses (OPEX), i.e., **the recurring electricity cost of Data Centers (DCs) to operate content service.**
    - ❖ DC cost: (1) Capital Expenses (CAPEX), i.e., the investments for infrastructure and equipment that were made upfront, and (2) Operational Expenses (OPEX), i.e., the recurring costs of operating the content-hosting DCs. **CAPEX is considered to be fixed**.
    - ❖ OPEX can be further categorized into operational cost (such as repairs, maintenance, salaries, etc.) and electricity cost. Considering the fact that large cloud infrastructure providers, such as Google, are usually considered as Tier 1 ISPs and not usually charged when using Internet connections for transit purposes according to peering agreements [1], it is also reasonable to assume the **average operational cost to be fixed**.
- **QoS**
  - ✓ **Resiliency:** content replication vs Content Fragmentation.
  - ✓ **Latency:** user to DC delay.
- **Green energy:** reduce the brown energy consumption.

UCDAVIS
UNIVERSITY OF CALIFORNIA

[1] Wikipedia, “Internet service provider".

Thank you!