

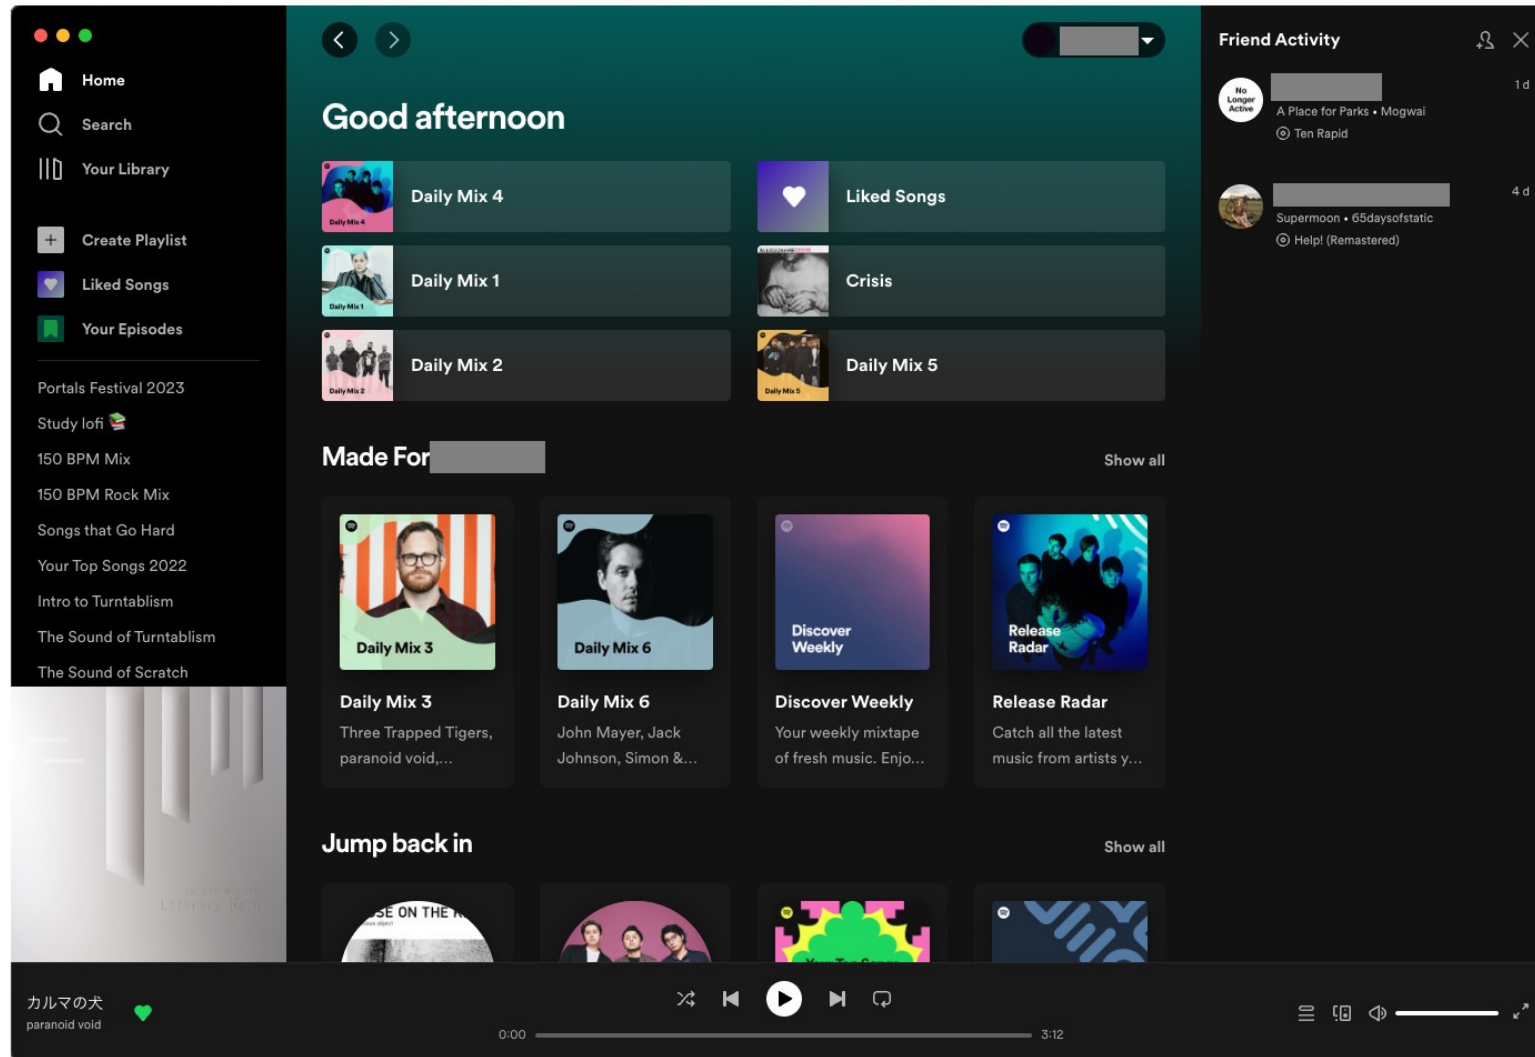
17-423/723: Designing Large-scale Software Systems

Recitation

Exercises: Design for Scalability

Mar 22, 2024

Exercise #1: Design Spotify



Exercise #1: Design Spotify

- Sketch a design of a music streaming application (e.g., Spotify), focusing on two operations: **Searching for songs** and **playing a playlist of songs**
- **Questions to discuss:**
 - What data do we need store for the operations?
 - What data storage model (relational vs. document) for which data?
 - What are major components in the system?
 - What type of scaling (vertical, horizontal, or both) do we apply?
 - How do we distribute data (replication, partitioning, or both)?
 - If replication is used, synchronous vs. asynchronous replication?
 - Which data do we cache? Where is the cache stored?

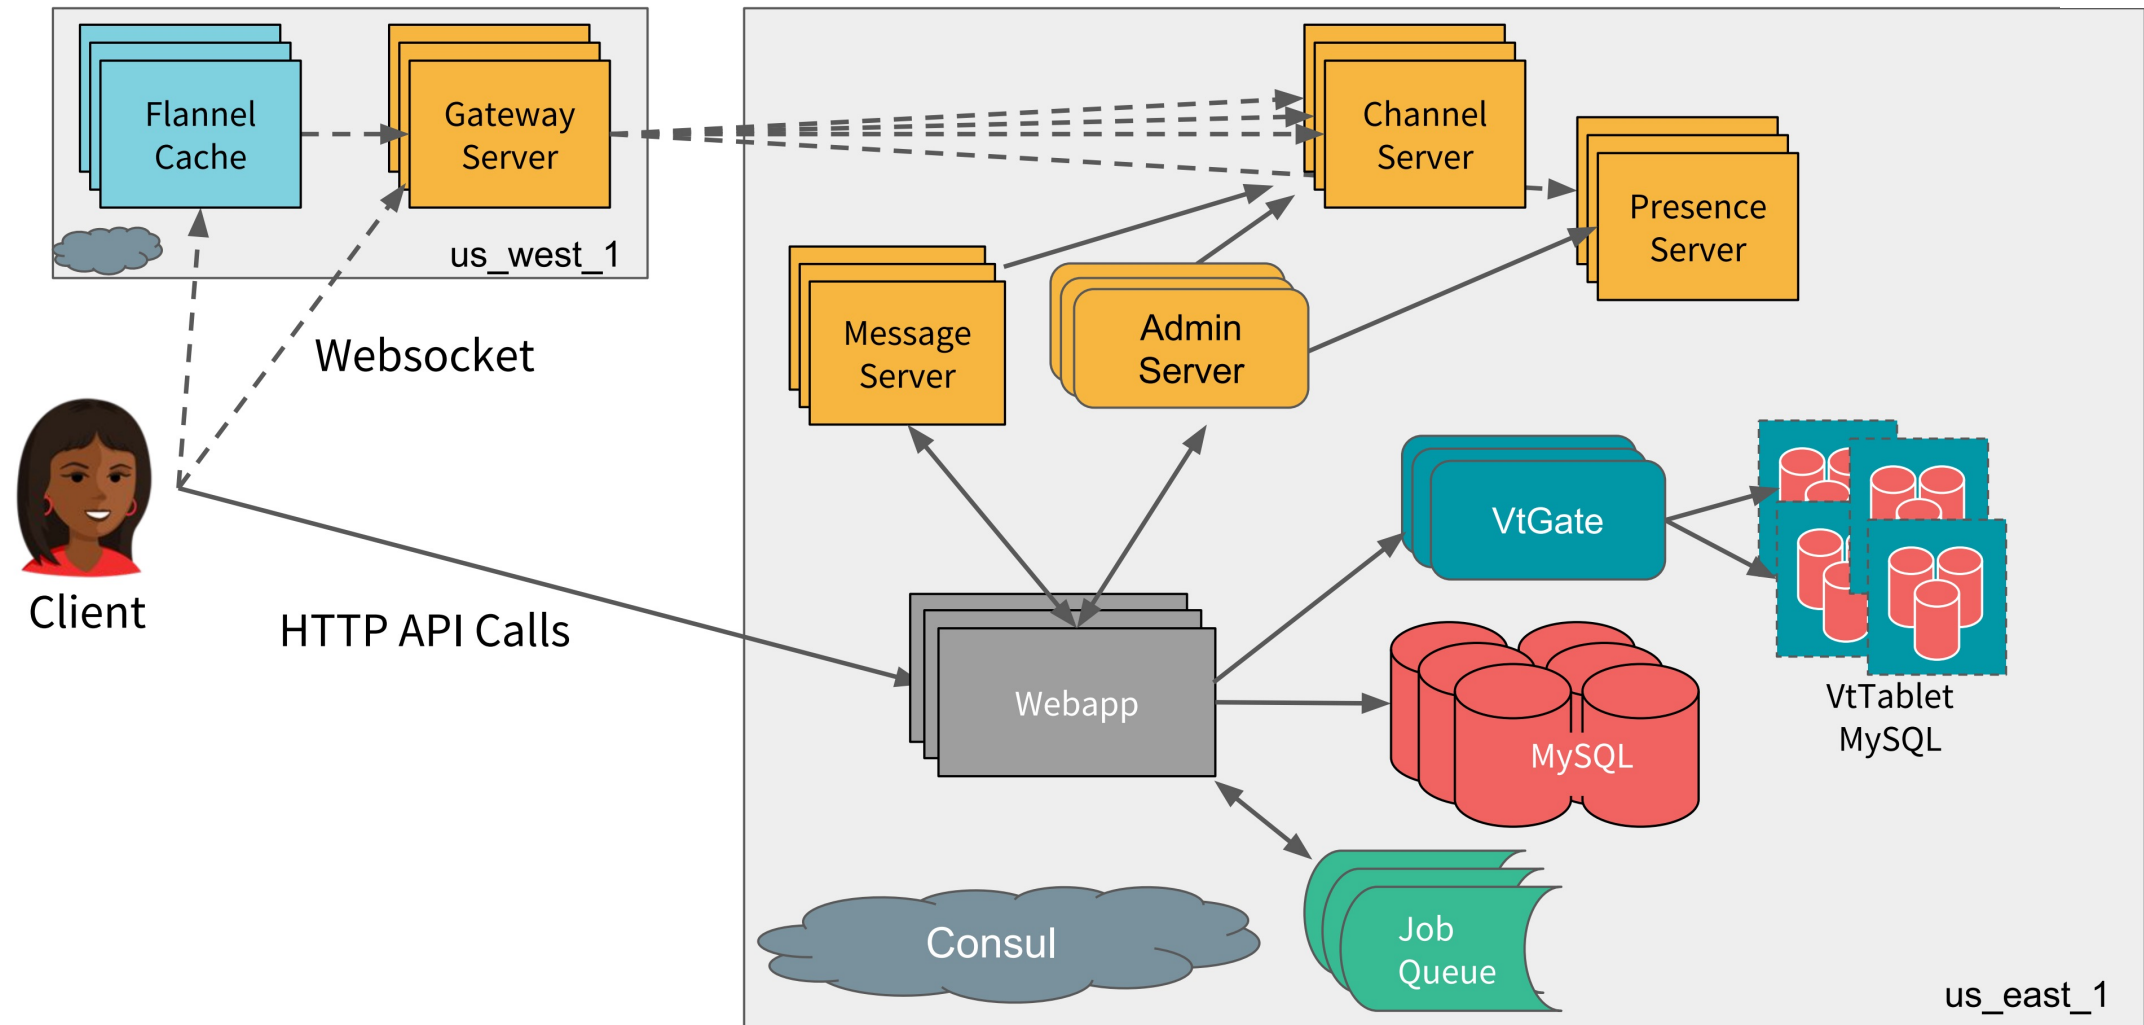
Exercise #2: Design Slack

The screenshot shows a Slack interface for the #client-proposal channel. On the left is a dark sidebar with navigation options: 'Acme Sites' (Victoria Thomas), 'All Threads', 'STARRED' (design-work, summer-campaign, Cory, Tina, Carl), 'CHANNELS' (accounting-costs, brainstorming (1), business-ops, client-proposal (highlighted), design-chat, marketing, media-and-pr, sonic-fanfic, triage-issues), and 'DIRECT MESSAGES' (slackbot and three others). The main channel view shows a search bar, channel name '#client-proposal', and a list of messages. Victoria Thomas posts: 'Hey team, hoping to have that proposal ready for the Alaska clients by 3pm today, how are we doing? I can chip in wherever needed!'. Carl Benting replies: 'I'm just about finished putting together the estimate portion of it, I could use some feedback. Here's the google doc I'm working on... docs.google.com/bin/proposal'. A Google Drive document 'Q3 OOH – Cost Estimate' is shared. A green checkmark with '1' indicates a reaction. Victoria Thomas replies: 'The numbers look pretty good, I tweaked a few things, but we're good to go!'. Reena Baines replies: 'I'm just wrapping up the sketches, I'll post them here once I'm done!'. On the right, an 'About #client-proposal' panel shows channel details, pinned items, 12/19 members, shared files, and notification preferences. At the bottom is a text input field with a plus icon and a smiley face emoji icon.

Exercise #2: Design Slack

- Sketch a design of a group communication platform (e.g., Slack), focusing on two operations: **Viewing a channel** and **sending a message to a channel**
- **Questions to discuss:**
 - What data do we need store for the operations?
 - What data storage model (relational vs. document) for which data?
 - What are major components in the system?
 - What type of scaling (vertical, horizontal, or both) do we apply?
 - How do we distribute data (replication, partitioning, or both)?
 - If replication is used, synchronous vs. asynchronous replication?
 - Which data do we cache? Where is the cache stored?

Slack Design (2018)



Scaling Slack - The Good, the Unexpected, and the Road Ahead (Mike Demmer, 2018)