



From Monolithic Services to a Microservice Architecture

Designing a Reference Architecture

Colin Milhaupt, Purdue University

Project Mentors: AKM Hassan and Michael Mitchell, Org. 9354

Problem:

Sandia National Labs develops web applications supporting multiple internal and external customers. As common in many IT organizations, applications are often developed in siloed teams. This independency can lead to redundancy between service components, adding unnecessary time to each application's development lifecycle. Some web applications utilize monolithic services, which may not scale well or may not be resilient to network and infrastructure disruptions.

Objective:

An ideal structure for these web apps would be to have them utilize a centralized core of services, which handled the most common tasks between all web apps. This would allow developers to focus on developing the software specific to their project. This centralized core of services would be accessible across the NSE, and they would be easily accessible, scalable, and agile, and serve as an example of best practices for logging.

Approach:

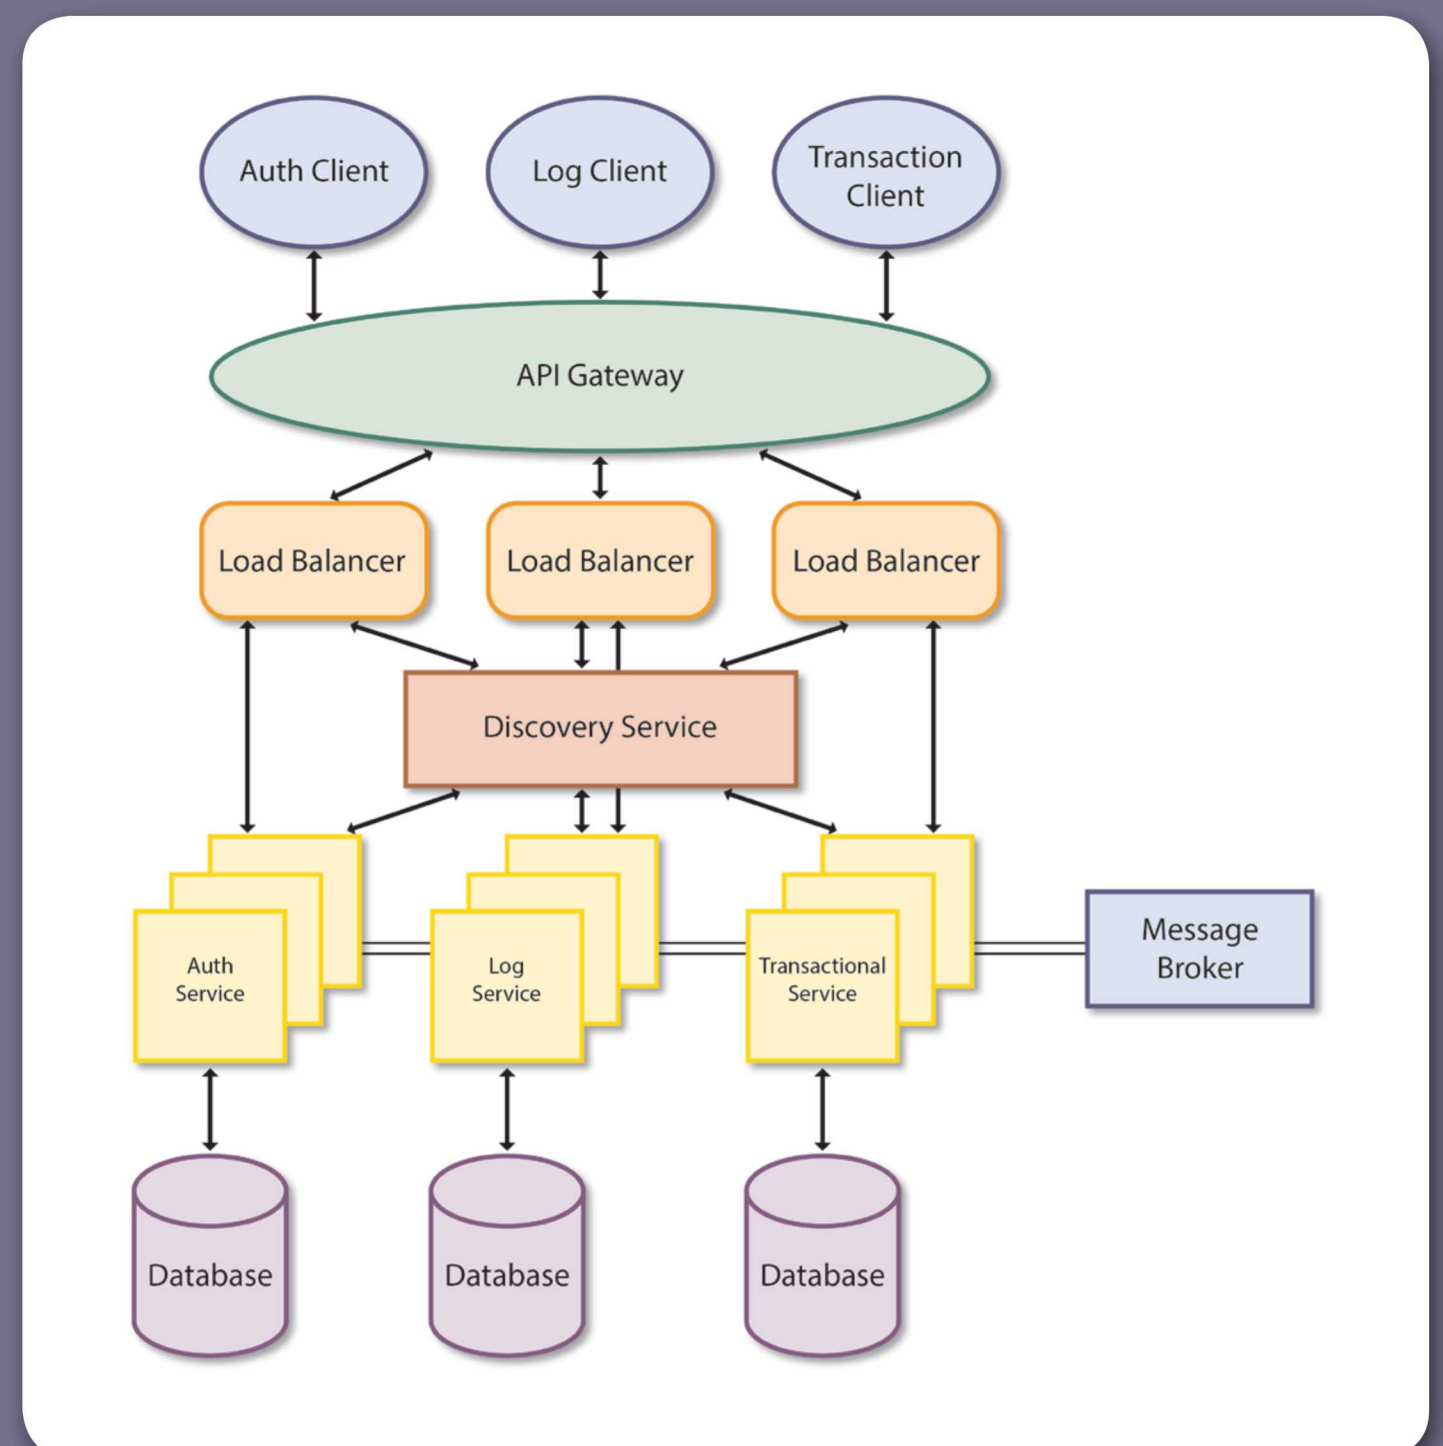
To demonstrate how useful and easy it would be to integrate a microservice into an existing web app, we built a proof-of-concept application and reference architecture that would show off the ease of containerization, load balancing, and centralized logging. The application also contains a load simulator demonstrating the scalability of microservices under disparate load conditions.

Our application has a front-end UI with two main views: a service dashboard and interactive log viewer. The dashboard contains a card with the service name, a live counter for the number of instances running, a graph of the load over time, and then controls beneath the card which simulate nominal load and peak load conditions, in terms of messages per second. The log viewer searches the most recent logs and contains predefined indexes for the services, errors, messages, heartbeats, etc.

Results:

For the dashboard, each card can be clicked which links to a page with the actual dashboard generated by Kibana. This dashboard contains metrics about the service. These metrics are collected by various Beats provided by Elastic such as network and metric beats. This provides real time monitoring of microservices which allow the user to verify scaling and load in a simulated environment.

For the log viewer, there's a search bar at the top of the page and preset searches below it. When a phrase is entered into the bar, it queries Elasticsearch and searches all indices for matching strings in the message field of the index and lists them in descending order. When preset searches are performed, it queries the field that preset is tied to. Logs are ingested into Elasticsearch through Logstash and can be loaded via a file beat, or over TCP/HTTP. Logs are separated logically and physically to maintain isolation between services, and sessions are generated when a user logs in so they can only see their logs. This serves as one of the many examples we demonstrate in our application for best practices.



Impact and Benefits:

The impact of this web app is in the foundation it lays for future microservice development. On top of the increased speed of development, and easily accessible, scalable, and agile microservices, the reference app serves as a testing ground for future technologies to ensure these properties are maintained.