BYRON - AN EVENT-DRIVEN MICROSERVICES FRAMEWORK

thors:		advisors:
oão F. Daniel ¹	Leonardo Lana ¹	Prof. Dr. Alfredo Goldi

Prof. Dr. Eduardo Guerra² BsC. Renato Cordeiro¹ Iman¹

¹ Instituto de Matemática e Estatística - Universidade de São Paulo ² Instituto Nacional de Pesquisas Espaciais

Introduction

In the past twenty year, the Web became ubiquitous to the point where availability and reliability are essencial to online systems. This demand pushed systems to become more complex: there are a lot of aspects that increase the chance of failure and there are a handful of practices that make them highly coupled, when a

What is Byron?

Byron is an Event-Driven Microservices Framework. It is a framework to develop systems following Microservices Architecture that adopts events as its core abstraction in the communication model, in order to provide decoupling between parts. It's a TypeScript framework that generates GraphQL APIs, uses MongoDB as the Cache with Mongoose as Object-Document Mapping and adopts NATS Streaming as Message Broker.

failure in one part causes failures all across the system.

The aim of this project is to create a solution for developing microservices that values decoupling, an architecture that provides decoupling and data consistency. To guide this project, three research questions were defined:

Is it possible to propose an architecture that provides data consistency within a system component, even if it's **Q1** eventual consistency?

Assuming the proposed architecture and considering more than one component, does it sustain the eventual **Q2** consistency across all components?

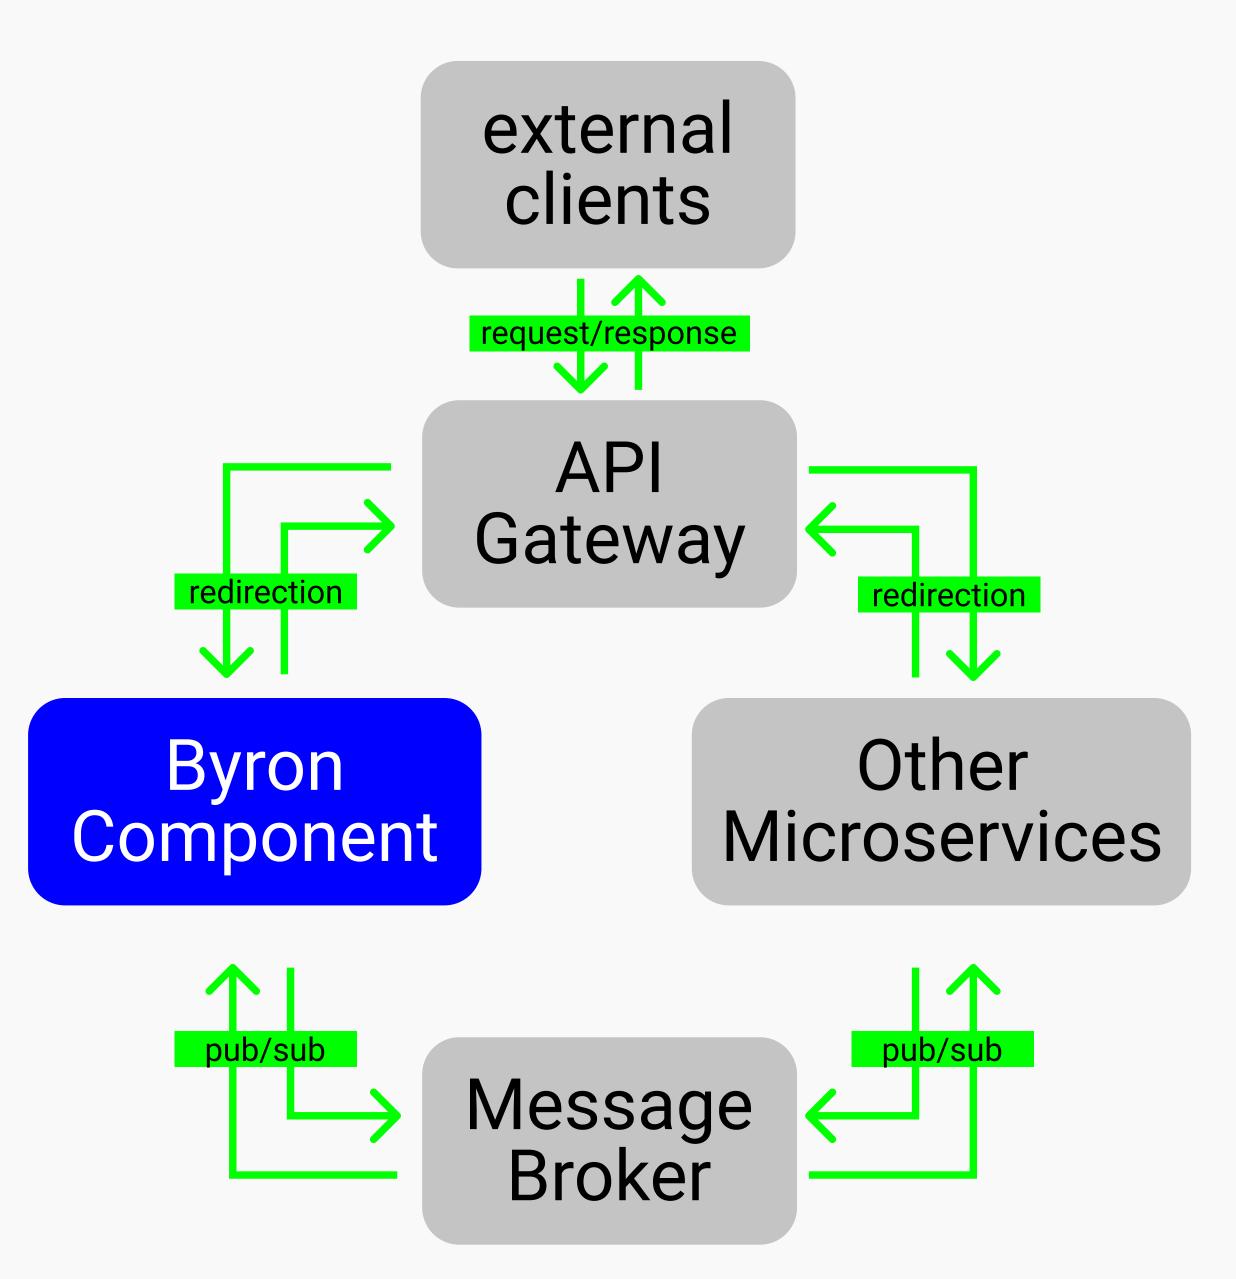
Does the proposed architecture provide decoupling between components, so that in case of one failing, the **Q3** others doesn't fail too?

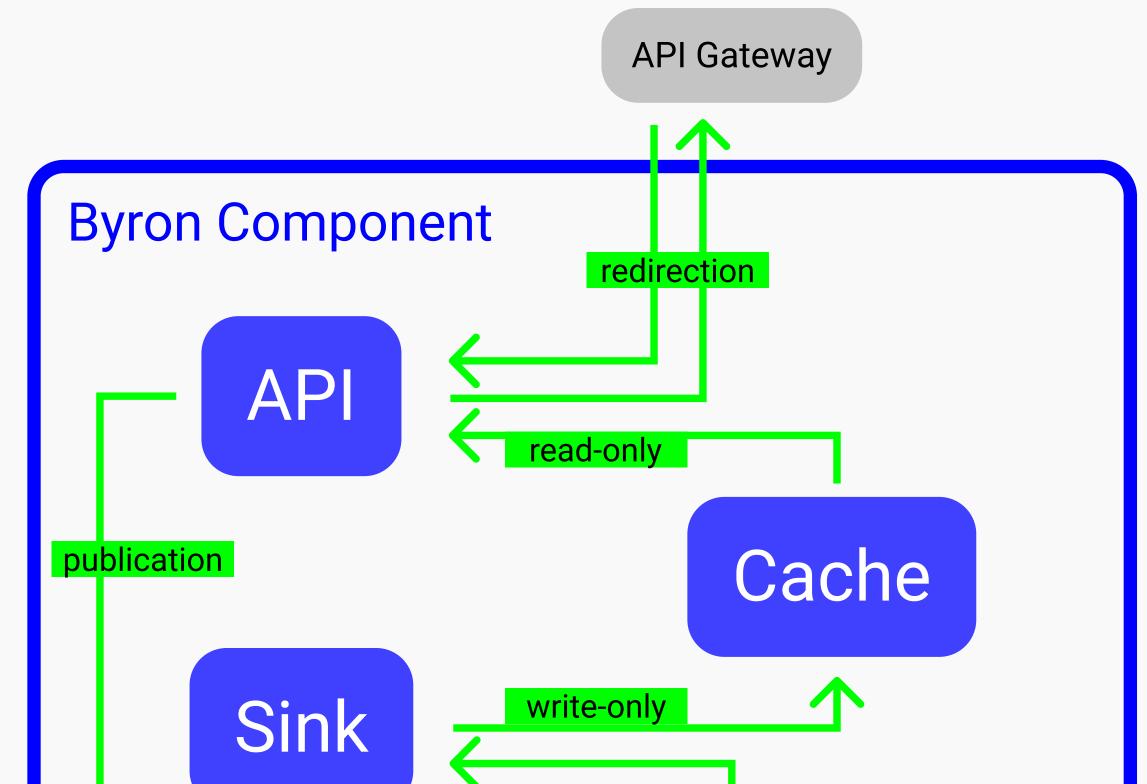
To reach this objective, it was built Byron – an event-driven microservices framework. It is based on three major principles: reactivity addresses the demand issue, stressing qualities about response times and internal communication standards; Cloud-Native focus on guiding the development of an application that is conceived to run with an automated, fully managed infrastructure; and Fast-Data models data handling as soon as it arrives.

Architecture

Byron's architecture is complex and it's better explained following the C4 Model (https://c4model.com) for visualizing software architecture. This model adopts an "abstraction-first" approach, similarly to what Google Maps does: it's possible to zoom in or zoom out, depending on the interest – more detail or more context, respectively.

A Component is the highest abstraction defined, it's supposed to implement a subset of the system's business logic. It's deployed within a back-end of a web server, being target of a request from an external client. An API Gateway can be used to abstract externally the inner division. Whenever there's an event, the Component interacts with the broker, via a pub/sub protocol.





A Component a set of three microservices: API, Sink and local Cache. When a request comes in, the API is responsible for extracting information from it and run a command. If it needs stored data, it reads the Cache. If the request causes any change in the application state, the API is also

When to use

To discuss the situations when the adoption of Byron is beneficial and when it's not, this section presents three different cases.

Social Media - beneficial

Suppose there's a social media with a lot of users across the world, with a great amount of development teams and a huge demand for content all over the days. This social media supports some features, such as creating a profile and an institutional page, posting to a news feed, linking people and interests but does not have a chat system. This chat is to be developed and for that they intend to adopt Byron.

- the development staff is big, with many teams, which
- benefits from autonomy provided by microservices architecture;
- the system is used all across the world with a great demand, also benefiting from system scalability;
- the feature of chat tolerates eventual data consistency and, thus, asynchronous messages.

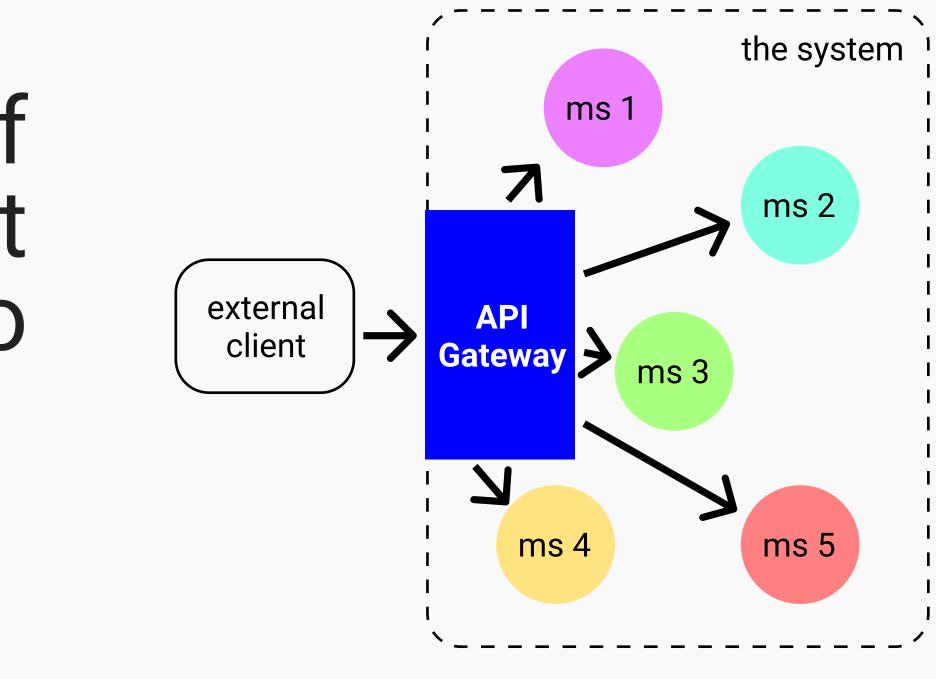
Bank - non-beneficial

Suppose there's a bank that is re-building its monolith system now adopting microservices architecture. It's a big bank supporting uncountable online transactions per day, with a large development staff. The core of the system deals with money, transactions and security. This core is to be rebuilt and for that they intend to use Byron.

the development staff is big, with many teams, so it benefit

Background

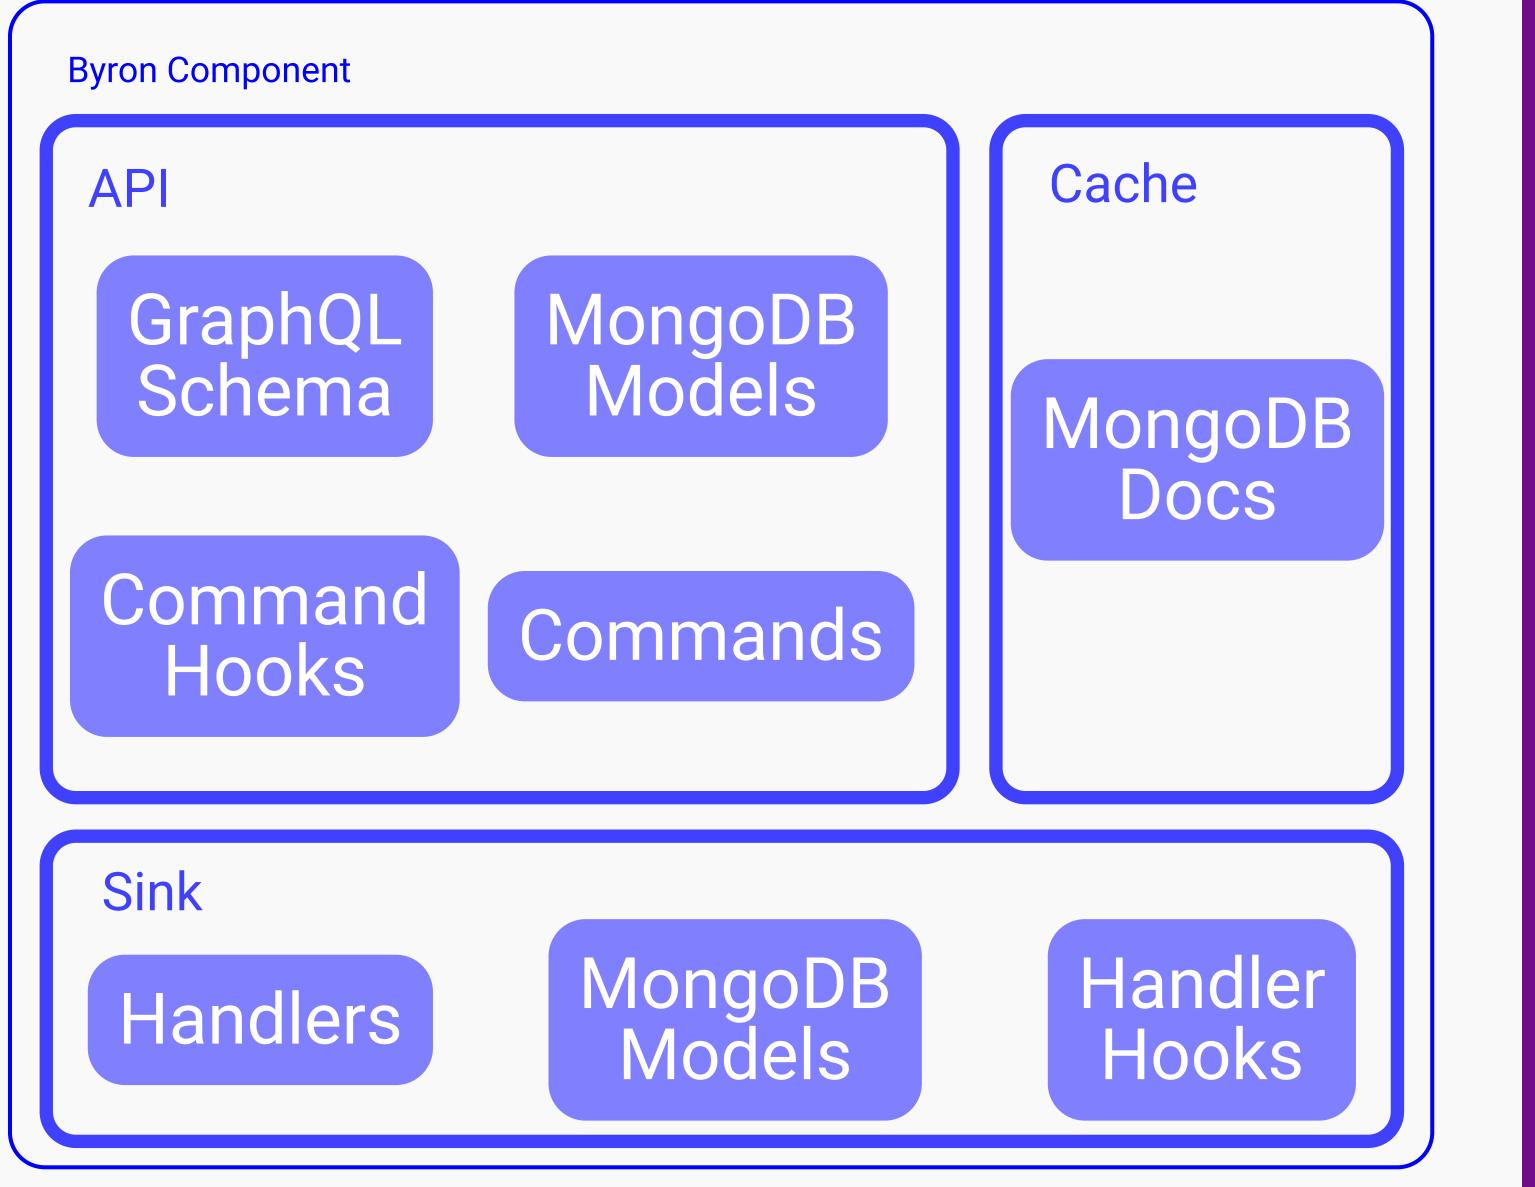
adopting API Gateway, the system itself routes the requests to the correct microservices, unifying the address to the external client



adopting CQRS, on which is possible to better fit the requirements of read and writes that might be different, by scaling independently ms 1 or ms 2

responsible for emitting an event notifying the rest of the change. Any Sink interested in that event Message Broker then reads it and updates the

Going deeper into details, each component's the microservices holds its own set of files. The API holds all commands and all command hooks. The Sink holds all handlers and all handler hooks. As the API and the Sink access the Cache, both hold a copy of Mongoose Models.



Cache with the payload data.

from the autonomy for development teams provided by microservices architecture;

the core business does not tolerate eventual data consistency – in that case, a better approach would be to adopt a synchronous communication model between the core microservices.

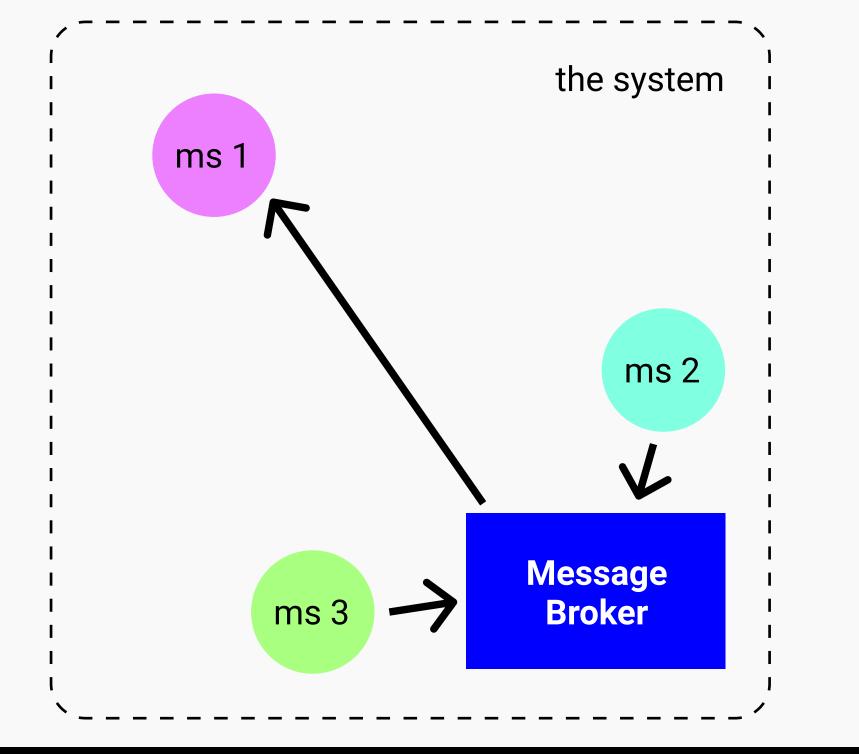
Startup on validation stage - non-beneficial

Suppose there's a small startup building its validation system. The development staff only has a five or six developers and the idea is not yet validated in the market. This system is to be first built and for that they intend to use Byron.

the development staff is small, so it doesn't demand a lot

- of autonomy and it's better not to use microservices, due to the complexity added to system;
- the system doesn't have a big demand yet, so system scalability is not an issue either.

adopting Event Sourcing, on which ms2 and ms3 don't have to wait for ms1 to process the request.



Conclusion and Future Work

In order to verify that the framework provided decoupling and

Reterences