

hammer-io

DESIGN DOCUMENT

Team: sdmay18-19

Client/Adviser: Lotfi Ben-Othmane

Members:

Erica Clark – Data Analytics Lead & Website/Content Management

Nathan De Graaf – Asana Expert & Weekly Status Reports

Nathan Karasch – Project Manager & Technical Writing

Jack Meyer – Communications & Software Architecture

Nischay Venkatram – UI Lead & Node.js SME

Email: sdmay18-19@iastate.edu

Website: sdmay18-19.sd.ece.iastate.edu

Revised: 15 Oct 2017 (version 1)

Contents

i. List of Tables and Figures	3
ii. Acronyms and Definitions	3
1 Introduction	4
1.1 Problem Statement	4
1.2 Solution Approach	4
1.2.1 Purpose	4
1.2.2 Goals	4
1.3 Operational Environment	5
1.4 Intended Users and uses	5
1.5 Assumptions and Limitations	6
1.5.1 Assumptions	6
1.5.2 Limitations	6
1.6 Expected End Product and Deliverables	6
1.6.1 Automated DevOps Process	7
1.6.2 Deployment Monitoring	7
1.6.3 NodeJS Microservice Development Framework	7
1.7 Deliverables Timeline	7
2 Specifications and Analysis	8
2.1 Functional Requirements	8
2.1.1 Automated DevOps process for Node.Js applications	8
2.1.2 Framework to develop Node.Js microservice applications	8
2.1.3 Monitoring Interface	8
2.2 Non-functional Requirements	8
2.3 Standards	9
2.4 Proposed Design	9

2.5 Design Analysis	10
3 Testing and Implementation	12
3.1 Interface Specifications	12
3.2 Hardware and software	12
3.3 Process	12
3.4 Results	13
4 Closing Material	14
4.1 Conclusion	14
4.2 References	14

i. List of Tables and Figures

Figure 1 First Semester Timeline

7

ii. Acronyms and Definitions

CI	Continuous Integration
CD	Continuous Deployment
CLI	Command Line Interface
DevOps	DevOps is a software engineering practice that aims at unifying software development (Dev) and software operation (Ops). ¹
Docker	A container platform used for deploying distributed software applications.
GUI	Graphical User Interface
MS	Microservice
Node.js	A framework for running standalone JavaScript applications.
npm	Node Package Manager
PR	Pull Request
UI	User Interface

¹ <https://en.wikipedia.org/wiki/DevOps>

1 Introduction

1.1 PROBLEM STATEMENT

There is a tendency to develop software as a collection of managed microservices. The microservice architecture involves a fair amount of complexity that may intimidate small teams with limited resources, limited time, or limited domain knowledge. Often the microservices need to be deployed to the cloud. As a result, another constraint on a small team is maintaining the system as its various components are updated independently of each other. Traditionally this is managed by a separate “DevOps” team; however, again, a small team with limited resources may not be able to dedicate people to the task. These two primary concerns—the overhead involved in setting up a microservices architecture and the resources involved in maintaining one—outline the main problem this project is addressing: that a microservices architecture may not be feasible for small teams, such as students or startups.

1.2 SOLUTION APPROACH

We will create a framework for developing Javascript-based microservices that exhibit specific quality attributes, as well as provide an automated DevOps process for managing the development and deployment of these JavaScript microservices. The project will be built in Node.js, an open-source server framework for running JavaScript code.

1.2.1 Purpose

By creating a framework for Node.js microservices, teams with limited knowledge, resources, and time can quickly get started with a simple infrastructure. We also provide the tools and structure to monitor and maintain the health of their applications in development and in production. This lets more developers start making cool things faster.

1.2.2 Goals

The main goal of our project is to create an environment to develop JavaScript microservices as well as a DevOps process to manage and monitor its deployment.

Specifically, we hope to create an app that requires minimal configuration and setup. This means the user would download one or two different tools and then have scripts create most of the configuration files needed with account info (e.g. GitHub, Docker Hub, etc.) supplied by the developers. Easy-to-use CLI/GUI tools would be available to push, deploy, and create new microservices.

Another goal that fits into this process is creating a seamless deployment workflow. This looks like a user pushing code to a master branch and watching his or her changes reflected in the live product a short while later. Once the application is deployed, the user then tracks the status of the application through a Data Monitoring app, which can notify interested parties when a service goes offline, display metrics related to server performance, etc. In addition, the Data Monitoring tool should be able to perform load balancing between microservices.

1.3 OPERATIONAL ENVIRONMENT

Our project is a software project that will operate in 3 different environments:

1. A Command Line Interface (CLI), which will be installed on a user's computer
2. A user's linux server
3. A cloud-based hosting service

The CLI gives the user the ability to generate a Node.js microservices project. Since the CLI will be used to generate a new user project, it is necessary that it be installed on the user's computer. Having software which can be installed on a user's computer gives a semi-unpredictable operating environment. To make it more predictable and limit the scope of the project, our software supports running on the three most common operating systems (Windows 10, Mac OSX, and Debian-based distributions of Linux).

The second operating environment would be on a user's linux server. Our monitoring service will have the ability to be run on any hosting service a user should want to use, including their own servers. We must ensure that our monitoring service has the ability to be deployed on a user's server.

Finally, we provide a cloud-hosted solution for the online project generator and for the user project monitoring service. That means, in addition to the user's monitoring service being independently deployed to his own linux server, we will have a central monitoring service run on our own host.

Our software uses the Node.js programming language. Since we choose Node.js, this means our software will run in a JavaScript runtime environment built on Chrome's V8 JavaScript engine. Node.js requires the user of our software to have Node.js v5 or newer and npm installed on their computer/linux server.

Our project strives to be fully open-source, which makes distribution of the software slightly different than traditional software. We primarily use Node Package Manager (npm) to distribute our software, but all of the code is available from GitHub and can be built from source. The framework is also available via npm and can be installed from source. Finally, the monitoring service can also be distributed and installed from source.

1.4 INTENDED USERS AND USES

For many companies, there is a lot of overhead in creating a new product. Our tool seeks to alleviate some of this overhead by providing a development pipeline out of the box. Our application and framework will allow developers to begin programming immediately without waiting to set up their deployment pipeline and have to worry about the intricacies of all the configurations. While many developers may benefit from Hammer, it is intended for students and small companies who need to deliver a working product with limited resources and time.

The framework will be highly customizable, with popular options like Express.js, GitHub, and TravisCI to fit a diverse set of needs. These options may be chosen initially or may be integrated after the project has already begun development.

1.5 ASSUMPTIONS AND LIMITATIONS

1.5.1 Assumptions

- We assumed the major audience that would be interested in using this product will be students, small teams, and startups that are looking for solutions that are open source, easy to use, and scalable. Most of the decisions that drive the project, are based primarily on this assumption. We want users to be able to scaffold projects quickly and scale without hitting any resource constraints. So all the technologies that we use in this project are open source and backed by a large community of developers.
- We assume the user is looking to build a microservice application in NodeJS, along with a continuous integration pipeline using TravisCI, containerization using Docker, backend using popular NodeJS web frameworks like ExpressJS, and deployment using hosting services like AWS, Heroku, and Digital Ocean. Hence, currently we only support generation of projects that use some or all of these technologies.
- If the user chooses to use any of these technologies, we assume they already have the necessary software installed to be able to generate and develop an application in these technologies. This includes having software like Git, NodeJS, etc installed and necessary accounts on Github, TravisCI, AWS, etc created beforehand.

1.5.2 Limitations

- The major limitation our project is that there is a limit to the processes we can automate. This can be seen directly from the assumptions. We cannot automate processes such as creating Github, TravisCI, or AWS accounts and also cannot simply install a variety of software on their machine. We merely assume they already have done that or fail gracefully.
- Since project revolves around bringing together open source software and many 3rd party services and API's, it is extremely difficult to test all the different pieces combined. This would make the application more prone to bugs.
- Since we are focusing on a certain set of technologies to provide to the user, we are possibly eliminating larger companies from the user group because they may want services or software options that we do not provide.

1.6 EXPECTED END PRODUCT AND DELIVERABLES

By the end of development for the project, we will have created three products:

1. An automated DevOps process for NodeJS microservice applications
2. Application to monitor the health and status of the deployed NodeJS applications
3. A framework to develop NodeJS microservice applications

The first product we are developing is the automated DevOps process for NodeJS microservice applications. In this product, the user should have the ability to do the following:

- Manage their development workflow (Create Github repo, push code, etc)
- Have their code automatically pushed to a CI environment to have it tested and built

- Have their code automatically containerized and deployed to hosting services
- Manually deploy and revert deployments of their application
- Perform security analysis

1.6.1 Automated DevOps Process

The DevOps application can be used in two different ways. The first is as a command line interface. Here, the user will be able to have the initial code base for their project automatically generated (a.k.a. scaffolded). This should include the configuration files needed for the many services our application will be connecting to based on the user's choices. We will also hook up any 3rd party applications as needed.

The other way the DevOps capabilities are consumed is through a web application that allows users to view statistics about the DevOps process, view development artifacts (build statuses, code coverage, test results, test run history, etc.), manage the build and deployment of the various services, and perform the same functions as the CLI. The DevOps application can be deployed on the user's own server or be consumed through our cloud service.

1.6.2 Deployment Monitoring

The second part is the monitoring application. This overlaps with the Devops web application. We plan to have it as part of the same web application. The user should be able to view the health of all the deployed applications, have access to crash logs/reports, memory usage, CPU usage, application load, and uptime of each application.

1.6.3 NodeJS Microservice Development Framework

The third part is the framework that will allow users to write these microservice applications in NodeJS. The framework should have capabilities like service discovery, inter-service messaging, request load balancing, and service presence and health. This is still being researched. We are looking to leverage existing frameworks and modify it for our use case. More information regarding this will come at a later date.

1.7 DELIVERABLES TIMELINE

Here is a timeline of the deliverables for the first semester

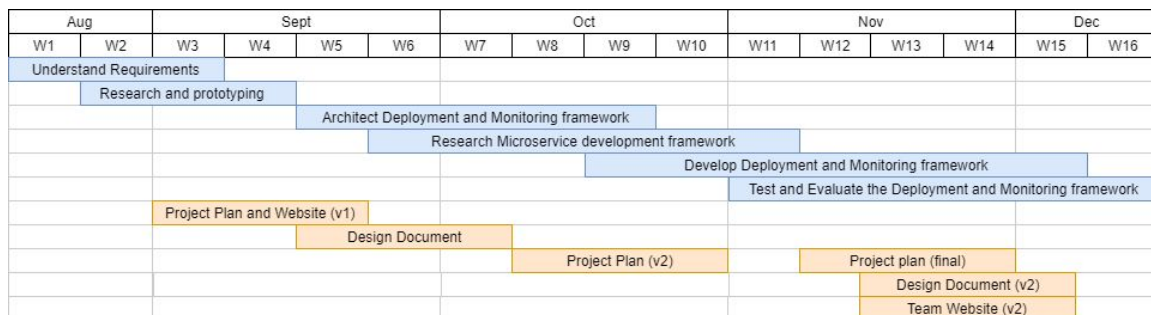


Figure 1. First Semester Timeline

2 Specifications and Analysis

2.1 FUNCTIONAL REQUIREMENTS

2.1.1 Automated DevOps process for Node.js applications

- Provides a CLI, which contains commands to:
 - Setup a new project (generates default configuration files) based on the services that a user wants to include
 - Serve the web application version of the CLI, which must:
 - Allow the user to select which services he/she wants to include in a new project
 - Download the default files for a new project based on the services that a user wants to include
 - View reports, results, and statistics about the DevOps process
 - Manage the build and deployment of services
 - Automate delivery of code to the CI environment, which must:
 - Build the application
 - Package the application in a Dockerfile
 - Run test suites
 - Report status and results
- Ability to configure various services with a configuration file
- Provides documentation for configuration

2.1.2 Framework to develop Node.js microservice applications

- Provides a CLI, which contains commands to:
 - Generate Node.js templates for new microservices
 - Integrate new microservices into the existing microservice architecture
 - Configure the microservices (host, port, etc)

2.1.3 Monitoring Interface

- A web application, which must:
 - Provide useful statistics, reports, and analytics about the deployed microservices
 - Details TBD
 - Allow reports and logs to be downloaded

2.2 NON-FUNCTIONAL REQUIREMENTS

- Usability
 - The system will only support the English language
 - The CLI must have a clean, consistent look and feel
 - The web application must have a clean, consistent look and feel
 - The application will be usable by those who have a limited understanding of DevOps
- Supportability

- The system will support Unix-based systems (e.g. Mac or Linux)
- The system will support Node.js version 8.x
- Reliability
 - The deployed web applications will run 24 hours a day, 7 days a week
 - We expect to have an update of greater than 99%
- Security
 - The system will not store any plain-text passwords in configuration files or elsewhere
- More non-functional requirements TBD

2.3 STANDARDS

The team will use the following standards during development of the project:

- Version Control System
 - Git will be used as the primary means of version control for all project code.
 - The Google software suite (Docs, Sheets, etc) will be used for all formatted documentation, such as planning and design documents. It has version control features that can be accessed from the menu (File → Version history).
- Code Review
 - Development will be done in feature branches.
 - When a feature branch is ready to be merged to the master branch, the author will assign one or more of the other project members as a reviewer.
 - Reviewers must check to ensure the code:
 - correctly implements the desired functionality
 - contains sufficient tests to ensure correctness
 - passes tests
 - is free of errors
 - integrates successfully with the existing software
 - Reviewers will communicate code issues to the author, who is then responsible for addressing all issues.
 - Once the code has passed inspection, it can be merged to master.
- Testing Standards
 - *To Be Determined*

2.4 PROPOSED DESIGN

The team is currently working on the automated DevOps CLI (the first major Functional Requirement). The following has been accomplished thus far:

- When the user starts the CLI, it asks a series of questions for scaffolding a new project. It asks for the following:
 - Project Name
 - Description
 - Version
 - Author
 - License
 - Tooling Options

- Choice of Continuous Integration Tool
 - Currently we support TravisCI
 - Choice of Containerization Tool
 - Currently we support Docker
 - Choice of Hosting Service
 - Currently we support Heroku
 - Choice of Web Application Framework
 - Currently we support ExpressJS
 - GitHub username/password
 - For each tooling option selected above, the CLI asks follow up questions based on the user's choice (e.g. if Heroku is chosen for hosting, it will prompt for the user's Heroku email, username, and password)
- Once the prompts are complete, the CLI generates the initial file structure:
 - It creates a new directory for the user's project
 - It generates the "package.json" and "src/index.js" files for the NodeJS project
 - It generates a ".gitignore" file
 - If ExpressJS is chosen as the web application framework, the CLI will add Express to the module dependencies and generate the initial "src/routes.js" file
 - If TravisCI was chosen, the CLI will generate a ".travis.yml" file for the project
 - If Docker was chosen, the CLI will generate a "Dockerfile" and ".dockerignore"
- Once the files have been initialized, the CLI integrates the project with the selected third-party applications
 - It initializes a new git repo in the project folder, commits the files, and pushes them to a new repository in the user's account
 - It enables the project in the user's TravisCI account
 - Etc.

2.5 DESIGN ANALYSIS

[This section will continue to grow in the following versions.]

Our team has created the minimum viable framework for the automation of the creation and deployment of a simple JavaScript project. The work done so far mainly consists of linking various services together through a simple CLI for the user. Targeted research and careful deliberation was used to pick the different tools our framework would use.

The simple CLI we have created has functionality for creating a sample project along with all the necessary configuration files for the selected services. So far, we provide TravisCI as an option for continuous integration and Docker as an option for a containerization platform. GitHub will act as the software development platform for the user's project. After the user provides their account details, our CLI sets up both the user's local git structure along with a connected remote repository on their account. This GitHub project and repository is then linked with TravisCI and in turn with Docker. This overall allows the user to go from local development to their remote project all the way to a container within Docker. The CLI has been able to successfully perform each step of this functionality.

Our team makes the conscious choice towards opinionated development. This means, for example, that rather than work to support all forms of continuous integration services, we choose to support

TravisCI in the best way possible and leave the option for other tools to be added in the future. This practice of choosing the simplest or most popular tool used in the industry best fits the needs of our client who wants the minimum hassle possible in setting up and using the services of their project. This means that at each part in this process we have the option to continue to add options and services for our users.

As we move forward, we will continue to build out the services and functionality for DevOps and deployment of the JavaScript services. We will continue to build out the monitoring and DevOps functionality through a more extensive front end website to supplement the simple CLI we have built so far.

3 Testing and Implementation

3.1 INTERFACE SPECIFICATIONS

As we develop our project, we add tests to verify consistent functionality after each change. We use Mocha for our test suite and run it manually and with TravisCI when new pull requests are made.

However, Mocha is unable to test a lot of our integrations like if a repo was successfully created on GitHub. So we have to manually check whether changes to GitHub still result in a remote account being created or if the repository was enabled on TravisCI.

To verify code style, we use Airbnb's Linter.

3.2 HARDWARE AND SOFTWARE

Mocha: A test suite for Node.js that results in organized, and easy-to-read tests. The tests are run each time there is a pull request to master to ensure that the master branch always has functional code.

Linter: Enables consistent style between multiple developers and keeps track if there are any remaining unused imports, etc to keep the code base clean.

3.3 PROCESS

The testing we've implemented so far has been Unit Testing and Style Analysis Testing (Linting). Both of these are run on the developer's machine as he or she is developing code. They are also run by TravisCI when a pull request (PR) is made to merge code into the master branch of the code repository. If the tests fail in TravisCI, the PR cannot be merged until the developer fixes the code that caused the tests to fail. Once the corrected code has been committed and the tests pass in TravisCI, the PR can be merged after it has been reviewed by another developer.

One of the testing challenges we face is that we rely on tight integration with third-party distributed applications for our software to function. For example, the CLI creates a new GitHub repository through GitHub's REST API and enables TravisCI on that repository using TravisCI's REST API. It doesn't stop there. We will need to integrate with hosting services (e.g. Heroku), containerization services (e.g. DockerHub) and other services that are part of the DevOps pipeline. These integrations are vital to meeting our functional requirements, but they are very difficult to automate. We are currently in the process of creating a test strategy to support development in this area. At the moment we are limited to manual and ad hoc testing.

Immediate testing shortfalls that need correction:

- Integration testing automation (if possible)
- Manual testing documentation strategy
- Regression testing strategy

[More test types will be included in future iterations of this document, as well as the mapping of tests to requirements and a flow diagram of the testing process.]

3.4 RESULTS

We have established a rudimentary test cycle centered around unit testing, style analysis testing, code reviews, and manual testing. The team is adhering well to this initial process, and we are currently researching additional test methods for the various third-party integrations.

The build and test logs from TravisCI can be viewed at <https://travis-ci.org/hammer-io/tyr> .

[This part will be refined in the 492 semester where the majority of the implementation and testing work will take place]

4 Closing Material

4.1 CONCLUSION

The popularity, simplicity, and usefulness of JavaScript microservices means there is a need from users for an easy to use and an all encompassing framework and devops systems to build and deploy these JavaScript services. The goal of this project is to meet these needs. Hammer seeks to provide three main parts: an automated deployment process, a Node.js framework to accommodate and enable microservices, and a data monitoring platform to monitor the health and metrics of the entire system.

In order to achieve this functionality, our team plans to implement the following three parts of the overall product. First we build a DevOps application to manage the code base of the service. This can be managed through either a CLI or a web application. Building these applications will be the first step of the project. Once finished with this portion, the user should be able to deploy their own service to be consumed by their users through our application. Second, we will create a monitoring application that will build upon the DevOps web application. This shows the users all relevant information about their information including logs, load, uptime, etc. This monitoring application is part of the features that makes our application better for our clients than managing each of these services individually. Finally, we create a framework to allow users to write these microservices in NodeJS. More details and plan for implementation to come.

Hammer, our solution to all JavaScript microservice development needs, seeks to meet the specific needs of our users in a way that other similar services cannot match. Our application's opinionated approach of setting up a suggested approach without the user needing any specialized knowledge makes Hammer the perfect choice for students or small teams that cannot afford to waste time or resources setting up and managing the logistics of their application. By focusing on simplicity, ease of use, and specializing in JavaScript microservices, our application can provide a better service than any other general DevOps or development framework service.

4.2 REFERENCES

- Babel Documentation: <https://babeljs.io/docs/usage/babelrc/>
- Docker Documentation: <https://docs.docker.com/develop/sdk/>
- DockerHub API: https://docs.docker.com/v1.4/reference/api/docker-io_api/
- ESLint Documentation: <https://eslint.org/docs/user-guide/>
- ExpressJS Documentation: <https://expressjs.com/en/4x/api.html>
- GitHub API: <https://developer.github.com/v3/>
- Heroku API: <https://devcenter.heroku.com/categories/platform-api>
- Mocha Documentation: <https://mochajs.org/>
- NodeJS Documentation: <https://nodejs.org/en/docs/>
- NPM Documentation: <https://docs.npmjs.com/>
- TravisCI API: <https://docs.travis-ci.com/api>
- TravisCI Documentation: <https://docs.travis-ci.com/>