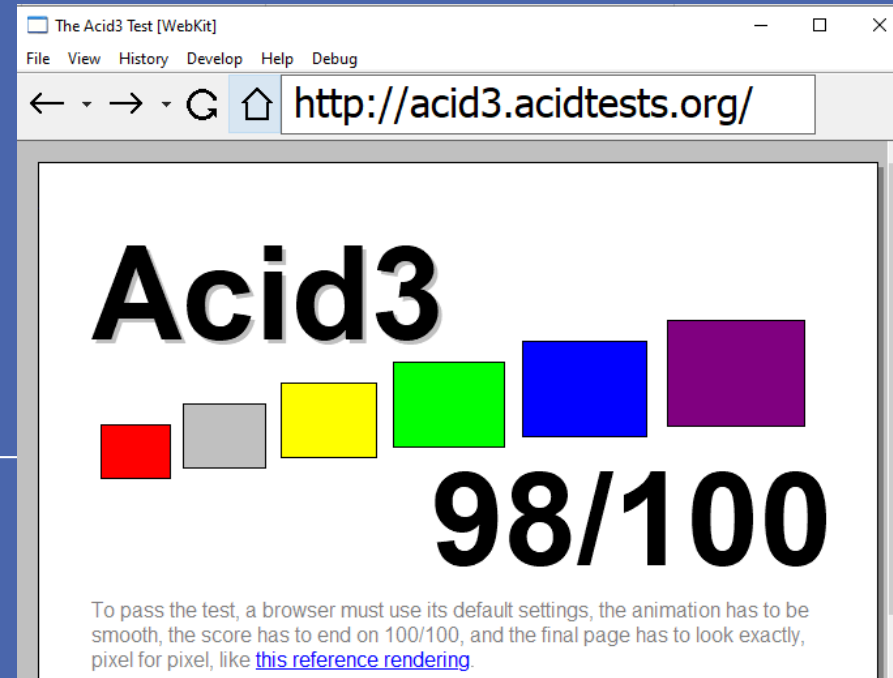


AMAZON ELASTICSEARCH FOR RUBY ON RAILS

By Spencer Peloquin

Curtis Senior High School




Engineering Goal

Engineering Goal. My engineering goal is to create a framework on Ruby for the Amazon ElasticSearch Relational Database. I tested the framework on Windows 10 and Android. My port operates online in a controlled environment using a package manager such as RubyGems, or offline through sideloading. This offers greater security and portability as a piece of software.

In the end, my goal of creating an efficient, cross-platform, database was successful, as it can run on ARM and WebAssembly, as well as traditional x86-64 based processors without problems.

Portage of API Reference Data



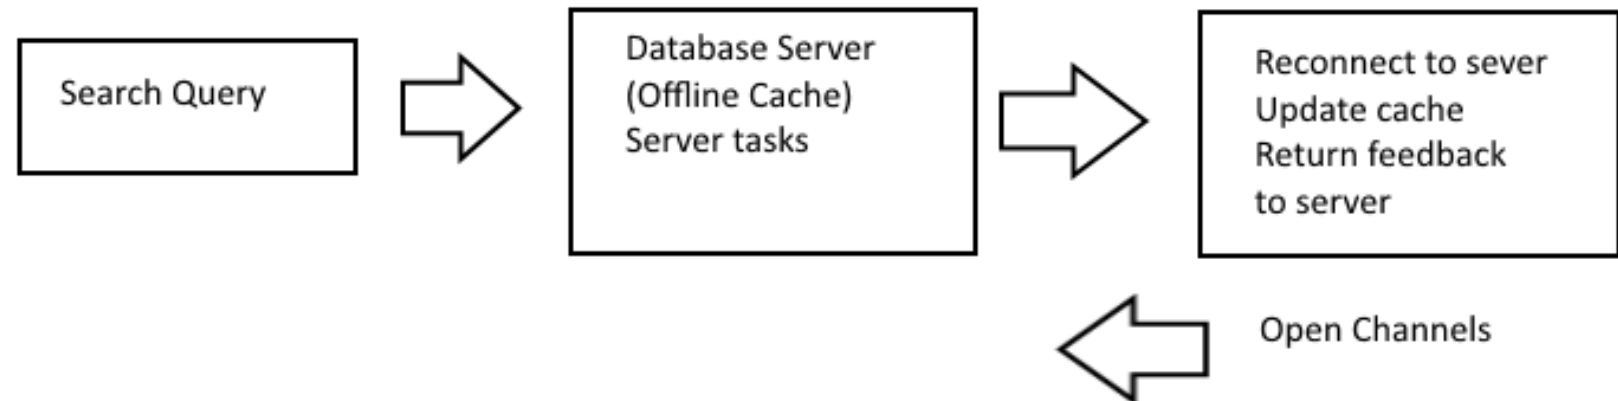
Problems:

- Proprietary Software
- High CPU Usage
- Uses assembly code (Hinders portability)

Benefits:

- Safer (fewer vulnerabilities)
- Efficient use of CPU
- Supports more Architectures (ARM, webAssembly)

REST API Usage



Features of ElasticSearch.net

- One-to-one mapping with REST API
- Generalized, pluggable architecture
- Configurable, automatic discovery of cluster nodes.
- Persistent, Keep-Alive connections.
- Load balancing across all available nodes.
- Child client support.
- TypeScript support out of the box.

elasticsearch-rails-origin 7.2.4

Ruby on Rails integrations for Elasticsearch.

VERSIONS:

7.2.4 - May 31, 2022 (7 KB)
7.2.3 - March 26, 2022 (7 KB)
7.2.2 - March 23, 2022 (7 KB)

[Show all versions \(4 total\) →](#)

DEVELOPMENT DEPENDENCIES (16):

```
bundler >= 0  
cane >= 0  
lograge >= 0  
minitest >= 0  
mocha >= 0  
oj >= 0  
pry >= 0  
rails > 3.1  
rake ~> 12  
require-prof >= 0  
ruby-prof >= 0  
shoulda-context >= 0  
simplecov >= 0  
test-unit >= 0
```

 Star 2,953

TOTAL DOWNLOADS

1,420

FOR THIS VERSION

390

GEMFILE:

```
gem 'elasticsearch-'
```

INSTALL:

```
gem install elastic-
```

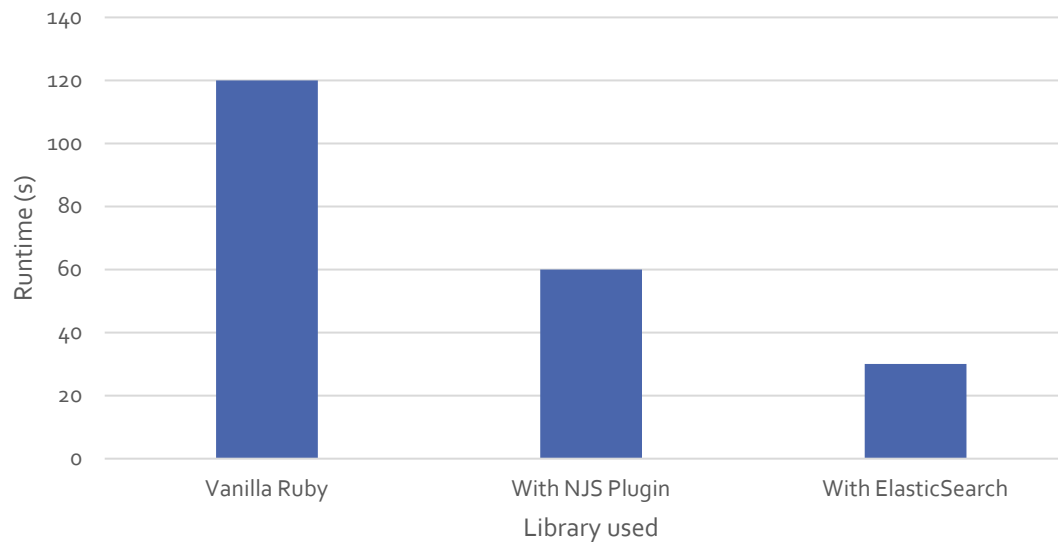
LICENSE:

APACHE 2

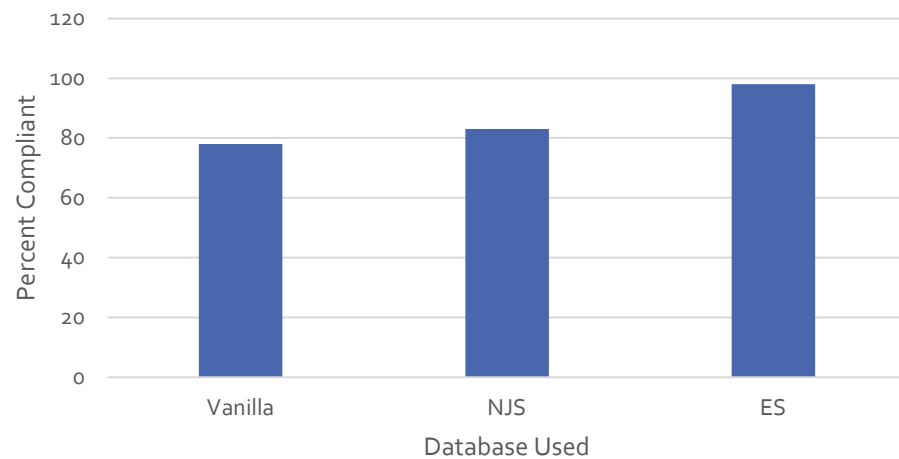
Usage Statistics

Performance

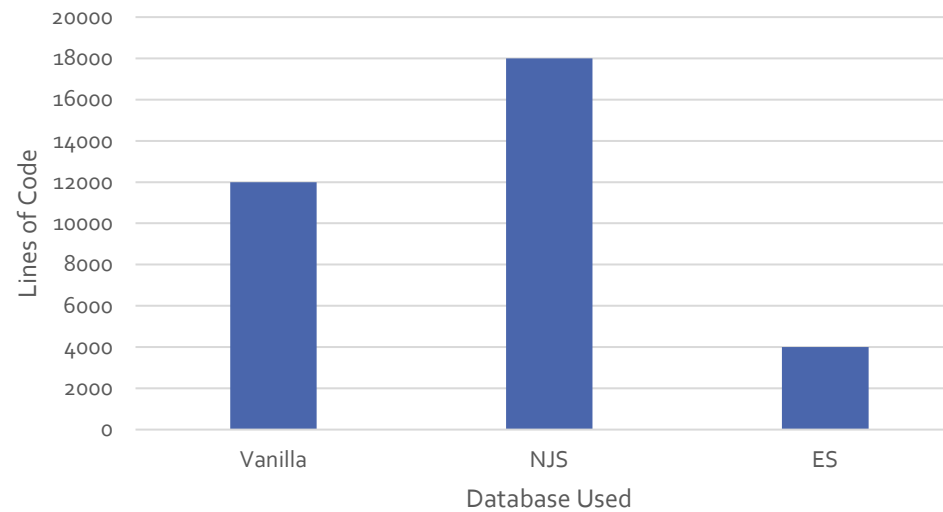
Project Runtime Comparison




ECMA Compliance



Library Bloat



SEARCH DATA (300 SEARCHES)



MODULE REFERENCE

```
elasticsearch-rails > lib > elasticsearch > rails > instrumentation.rb
21 module Elasticsearch
22   module Rails
23
24     # This module adds support for displaying statistics about search duration in the Rails application log
25     # by integrating with the `ActiveSupport::Notifications` framework and `ActionController` logger.
26     #
27     # == Usage
28     #
29     # Require the component in your `application.rb` file:
30     #
31     #   require 'elasticsearch/rails/instrumentation'
32     #
33     # You should see a output like this in your application log in development environment:
34     #
35     #   [2015-01-22 15:00:00] INFO -- Elastic Search (321.3ms) {"index": "article", "type": "article", "body": {"query": {}} }
36     #
37     # Also, the total duration of the request to Elasticsearch is displayed in the Rails request breakdown:
38     #
39     #   Completed 200 OK in 615ms (Views: 230.9ms | ActiveRecord: 0.0ms | Elasticsearch: 321.3ms)
40     #
41     # @note The displayed duration includes the HTTP transfer -- the time it took Elasticsearch
42     #        to process your request is available in the `response.took` property.
43     #
44     # @see Elasticsearch::Rails::Instrumentation::Publishers
45     # @see Elasticsearch::Rails::Instrumentation::Railtie
46     #
47     # @see http://api.rubyonrails.org/classes/ActiveSupport/Notifications.html
48     #
49     #
50     module Instrumentation
51     end
52   end
53 end
```

LINKED LIBRARIES (VS CODE 2022)

```
1 # Licensed to Spencer Peloquin under one or more contributor
2 # license agreements. See the NOTICE file distributed with
3 # this work for additional information regarding copyright
4 # ownership. Spencer Peloquin licenses this file to you under
5 # the MIT License (the "License"); you may
6 # not use this file except in compliance with the License.
7 # You may obtain a copy of the License at
8 #
9 # https://mit-license.org/
10 #
11 # Unless required by applicable law or agreed to in writing,
12 # software distributed under the License is distributed on an
13 # "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY
14 # KIND, either express or implied. See the License for the
15 # specific language governing permissions and limitations
16 # under the License.
17
18 require 'pathname'
19 require 'elasticsearch'
20
21 subprojects = ['elasticsearch-rails', 'elasticsearch-persistence']
22 subprojects << 'elasticsearch-model' unless defined?(JRUBY_VERSION)
23
24 __current__ = Pathname(File.expand_path(__dir__))
25
26 def admin_client
27   $admin_client ||= begin
28     transport_options = {}
29     test_suite = ENV['TEST_SUITE'].freeze
30
31     if hosts = ENV['TEST_ES_SERVER'] || ENV['ELASTICSEARCH_HOSTS']
32       split_hosts = hosts.split(',').map do |host|
33         /(http:\/\/\w\/)?(\S+)/.match(host)[2]
34       end
35
36       host, port = split_hosts.first.split(':')
```


CHANGELOG (2021)

```
1  ## 7.2.1
2
3  * The default git branch `master` has been renamed to `main`
4  * Adds compatibility with Ruby 3 [Pull Request](https://github.com/elastic/elasticsearch-rails/pull/992)
5
6  ## 7.2.0
7
8  * Updates specs and dependency to use with `elasticsearch` v7.14.0.
9  * Update README, remove Virtus (unmaintained)
10 * Updates `Bundler.with_clean_dev` (deprecated) to `with_unbundled_env` [commit](https://github.com/elastic/elasticsearch-rails/commit/e4545e4fe2a1ce80009206c831d5740360bad6c2)
11 * Deal with `nil` document types in Multimodel [commit](https://github.com/elastic/elasticsearch-rails/commit/cd9c309b78de443d2e37760998418616ba34276d)
12 * Update dependency to explicitly support version 7 [commit](https://github.com/elastic/elasticsearch-rails/commit/65942e3da9cabad2f6965e69c8ef6a0994da9408)
13 * Stop emitting FATC log when checking existence of indices [commit](https://github.com/elastic/elasticsearch-rails/commit/20ca39865d77f7109360ca7f0e31212)
14 * Remove `InetSocketAddress` exception test on index health [commit](https://github.com/elastic/elasticsearch-rails/commit/ce57cc17e304b0a4af123c1599f37fb892a5d93a)
15 * Removes dependency on extensions [commit](https://github.com/elastic/elasticsearch-rails/commit/ed070b8329ca48b4cb12b513ac81ed78c88acc61)
16 * Fixes basic template elasticsearch dependency [commit](https://github.com/elastic/elasticsearch-rails/commit/a4ec07b2d097545ca41c13686c9cbfc9eab9e639)
17
18 ### ActiveRecord
19
20 * Fixes indexing to use right logger in client
21 * Updates ES client spec for client 7.14.0
22 * Updates transport references
23
24 ## 7.1.1
25
26 * Fix: Ruby 2.7 deprecation warning on `find_in_batches`
27 * Updates README for generating app with template. Related: #938
28
29 ### ActiveRecord
```

Test and Redesign

- The prototype used a lot of memory and relied heavily on 3rd party software and had a less permissive license. The need for a redesign was evident from the beginning. However, feedback from on AWS allowed me to improve the efficiency of the framework.
- The first thing I did was to create a template on visual studio that would allow me to effectively rewrite the framework. Using original code and help from AWS, I managed to reduce the library size to under a Megabyte. I then compiled the repository for approval from rubygems.
- The result was extremely positive, with a 50% faster compiling time, and a 20% faster interpretation period.

Future Research

- ▶ I have published the gem to the RubyGems Repository under the name “ElasticSearch-Rails-Origin.”
- ▶ Next, I may try to implement a RESTful API to make networking more efficient.
- ▶ Finally, I would spend more time debugging and reducing vulnerabilities.

Bibliography

“Advantages of Web Development with Ruby on Rails.” Onix Systems, <https://onix-systems.com/blog/advantages-of-web-development-with-ruby-on-rails>.

Elastic. “Elastic/Elasticsearch-JS: Official Elasticsearch Client Library for Node.js.” GitHub, <https://github.com/elastic/elasticsearch-js>.

Peloquin, Spencer M. “Programming-Ruby.pdf.” Github.io, 2021.

“Rails Crash Course.” Rails Crash Course, <https://railscrashcourse.github.io/>.

Sam Ruby on March 22, et al. Ruby2JS, <https://www.ruby2js.com/>.

terrasky064. “TERRASKY064/Elasticrails-Origin: Updated Repository to [Https://Github.com/terrasky064/Elasticrails](https://Github.com/terrasky064/Elasticrails).” GitHub, <https://github.com/terrasky064/elasticrails-origin>.

License

- The MIT License (MIT)
- Copyright © 2022 Spencer Peloquin
- Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:
- The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.
- THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

JSON Output (Rails)

{ } modulexrubykendra.json 1

{ } module2nodejs.json

{ } module.json ×

{ } module.json

```
1  module Search
2    class TagSerializer
3      include FastJsonapi::ObjectSerializer
4
5      attributes :id, :name, :hotness_score, :supported, short_summary, :rules_html, :time_51
6    end
7  end
8
9  tag = Tag.find(id)
10 serialized_data = Search::TagSerializer.new(tag).serializable_hash.dig(:data, :attributes)
11 SearchClient.index(id: tag.id, index: "tags_development", body: serialized_data)
12
13 SearchClient.get(id: tag.id, index: "tags_development")
14
```

JSON Output (NJS)

} *modulexrubykendra.json* 1

{ } module2nodejs.json ×

{ } module2nodejs.json

```
1  module Search
2    class TagSerializer
3      include FastJsonapi::ObjectSerializer
4
5      attributes :id, :name, :hotness_score, :supported, short_summary, :rules_html, :time_90
6    end
7  end
8
9  tag = Tag.find(id)
10 serialized_data = Search::TagSerializer.new(tag).serializable_hash.dig(:data, :attributes)
11 SearchClient.index(id: tag.id, index: "tags_development", body: serialized_data)
12
13 SearchClient.get(id: tag.id, index: "tags_development")
14
```

Vanilla Ruby (Kendra)

{ } modulexrubykendra.json 1 ×

{ } module2nodejs.json

{ } module.json

{ } modulexrubykendra.json

```
1  module Search
2    class TagSerializer
3      include FastJsonapi::ObjectSerializer
4
5      attributes :id, :name, :hotness_score, :supported, short_summary, :rules_html, :time_120
6    end
7  end
8
9  tag = Tag.find(id)
10 serialized_data = Search::TagSerializer.new(tag).serializable_hash.dig(:data, :attributes)
11 SearchClient.index(id: tag.id, index: "tags_development", body: serialized_data)
12
13 SearchClient.get(id: tag.id, index: "tags_development")
14
```