

Locating Inter-communication Performance Anti-patterns in Microservices

Masoumeh Nourollahi

Polytechnique Montréal DORSAL Laboratory

Introduction

- Performance anti-patterns (SPAs)
 - Bad practices which result in performance degradation
 - May not cause system failure, but impact performance



Locating Inter-communication Performance Anti-patterns in Microservices

2/14 – dorsal.polymtl.ca

Agenda

- Problem statement
- Proposed artichecture and detection method
- Experiments
- Results
- Discussion and future work

Problem <u>statement</u>

- Excessive Messaging
- Blob
 - High communication overhead through remote communication between central blob component and other components
- Empty Semi-truck
 - Large number of small messages transmitted between two components as part of a single user request







Proposed Architecture



5/14 – dorsal.polymtl.ca

SPA Detection method

- 1. Decision making on the required tracing setup (load test and tracing scope)
- 2. Trace collection
- 3. Data Aggregation: The data is aggregated over fixed time windows
- 4. Using heuristics to detect SPA

LOAD PROFILES

- 1.single user-single action
- 2.single user- mixed action
- 3.increasing load-single action
- 4.increasing load- mixed action
- 5.max load- single action
- 6.max load- mixed action

TRACING SCOPE

- 1.Kernel space
- Kernel events
 - net_'*',sock_'*',napi_poll,skb_'*'
 - Sched_'*'
- 2.User space
- Function call entry/exit

Intercommunication SPAs- Detection Method

- Excessive Messaging
 - Load profile: Increasing load- single action
 - Tracing scope: Kernel events: net_if_receive_skb, net_if_receive_skb
- Blob
 - Load profile: stress load- single action
 - Tracing scope: Kernel events: net_if_receive_skb, net_if_receive_skb
- Empty Semi-truck
 - Load profile: Single user-single action
 - Tracing scope: Kernel events: net_if_receive_skb, net_if_receive_skb, Sched_*, 2 UST Tracepoints- action related method entry/exit

Tracepoint	Definition
Net_dev_queue	A network packet is sent
Net_if_receive_Skb	A network packet is received
Sched_switch	Transition from one task to another on a CPU

Experiments

- TestBed:
 - DeathStarBench- Social Network application



SPA Injection- Excessive messaging

 A loop that performs 1000 consecutive queries to MongoDB. This excessive messaging significantly increases the load on the MongoDB server during read operations.

Original

```
int mongo_start = start + post_ids.size();
std::multimap<std::string, std::string> redis update map;
if (mongo start < stop) {</pre>
 // Instead find post ids from mongodb
 mongoc client t *mongodb client = mongoc client pool pop(
      mongodb client pool);
 if (!mongodb client) {
    ServiceException se;
   se.errorCode = ErrorCode::SE_MONGODB_ERROR;
    se.message = "Failed to pop a client from MongoDB pool";
    throw se:
 }
 auto collection = mongoc_client_get_collection(
     mongodb_client, "user-timeline", "user-timeline");
 if (!collection)
    ServiceException se;
    se.errorCode = ErrorCode::SE_MONGODB_ERROR;
    se.message = "Failed to create collection user-timeline from MongoDB";
    throw se:
```

SPA Injected

```
int mongo start = start + post ids.size();
std::multimap<std::string, std::string> redis_update_map;
if (mongo_start < stop) {</pre>
  // Excessive messaging: Multiple queries to MongoDB
 for (int i = 0; i < 1000; ++i) {</pre>
    mongoc client t *mongodb client = mongoc client pool pop(
        mongodb_client_pool);
    if (!mongodb client) {
      ServiceException se:
      se.errorCode = ErrorCode::SE_MONGODB_ERROR;
      se.message = "Failed to pop a client from MongoDB pool";
      throw se;
    auto collection = mongoc_client_get_collection(
        mongodb_client, "user-timeline", "user-timeline");
    if (!collection) {
      ServiceException se;
      se.errorCode = ErrorCode::SE_MONGODB_ERROR;
      se.message = "Failed to create collection user-timeline from MongoDB";
      throw se;
    3
```

Results

- Identification of intercommunication performance anti-patterns by kernel traces
- Reduced pre-tracing effort
- Low-cost tracing (2 to 6 kernel events required to identify the SPAs)

Metric	Trubiani (2023) [1]	Avritzer (2002) [6]	Wert (2014) [10]	Our method
Pre-tracing effort	UST instrumentation	UST instrumentation	UST instrumentation	1 UST TP
Tracing cost/overhead	Not reported	Not reported	High for EST pattern	Between 2 to 6 kernel events
Detection method	Load testing and profiling	Queuing performance models	Algorithmic approach	Statistical analysis and Algorithmic approach
Ассигасу	F-measure > 85%	F-measure > 44%	True positive ≈ 1 Notable False positive	True positive ≈ 1 Notable False positive

Discussion

- 1. Considering time-window analysis instead of per-request
- 2. Previous works all used UST traces to detect intercommunication anti-patterns, we are mostly using kernel traces
 - Less effort pre tracing (does not need any or very little instrumentation)
 - Avoiding unknown impact on system performance
 - Taking 2 to 6 kernel events which is very low-overhead tracing
- 3. We implemented a number of methods to compare
 - Wert statistical heuristics to detect performance anti-patterns [10]
 - Avritzer queuing model for anti-pattern detection [6]

Discussion

• Future directions:

- Metrics based comparison of application performance to detect SPAs
- Use of time-series analysis methods by using the available metrics
- Some suggested metrics to measure in each time window
 - msg_service_time
 - request_throughput
 - message_throughput
 - request_arrival_rate
 - msg_arrival_rate
 - msg_response_time_variance
 - msg_waiting_time
 - msg_residence_time

• ...

References

- 1. Trubiani, Catia, et al. "Automated detection of software performance antipatterns in Java-based applications." *IEEE Transactions on Software Engineering* (2023).
- 2. VanDonge, Riley, and Naser Ezzati-Jivan. "N-Lane Bridge Performance Antipattern Analysis Using System-Level Execution Tracing." 2022 IEEE 22nd International Working Conference on Source Code Analysis and Manipulation (SCAM). IEEE, 2022.
- 3. Chalawadi, Ram Kishan. "Automating the characterization and detection of software performance antipatterns using a data-driven approach." (2021).
- 4. Stefanakos, Ioannis. Software analysis and refactoring using probabilistic modelling and performance antipatterns. Diss. University of York, 2021.
- 5. Smith, Connie U. "Software performance antipatterns in cyber-physical systems." *Proceedings of the ACM/SPEC International Conference on Performance Engineering*. 2020.
- 6. Avritzer, Alberto, et al. "Scalability assessment of microservice architecture deployment configurations: A domain-based approach leveraging operational profiles and load tests." *Journal of Systems and Software* 165 (2020): 110564.
- 7. Trubiani, Catia, et al. "Exploiting load testing and profiling for performance antipattern detection." *Information and Software Technology* 95 (2018): 329-345.
- 8. Heger, Christoph, et al. "Expert-guided automatic diagnosis of performance problems in enterprise applications." 2016 12th European Dependable Computing Conference (EDCC). IEEE, 2016.
- 9. Keck, Philipp, et al. "Antipattern-based problem injection for assessing performance and reliability evaluation techniques." 2016 IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW). IEEE, 2016.
- 10. Wert, Alexander, et al. "Automatic detection of performance anti-patterns in inter-component communications." *Proceedings of the 10th international ACM Sigsoft conference on Quality of software architectures.* 2014.

Thanks for your attention

```
Masoumeh.nourollahi@polymtl.ca
https://github.com/mnourollahi
```

Blob Injection

 Instead of fetching only the necessary data (such as post IDs) from MongoDB, it fetches entire documents for the user's timeline. This approach unnecessarily transfers large amounts of data between the database and the application, leading to inefficient resource utilization and slower performance.

bool found = mongoc_cursor_next(cursor, &doc); if (found) {
bson iter t iter 0;
bson iter t iter 1;
bson iter t post id child;
bson_iter_t timestamp_child;
int idx = 0;
bson_iter_init(&iter_0, doc);
<pre>bson_iter_init(&iter_1, doc);</pre>
while (bson_iter_find_descendant(&iter_0,
<pre>("posts." + std::to_string(idx) + ".post_id").c_str(),</pre>
<pre>&post_id_child)</pre>
&& BSON_ITER_HOLDS_INT64 (&post_id_child)
<pre>&& bson_iter_find_descendant(&iter_1,</pre>
<pre>("posts." + std::to_string(idx) + ".timestamp").c_str(),</pre>
<pre>&timestamp_child)</pre>
&& BSON ITER HOLDS INTO 4 (&LIMESTAMD CHILd) } {
auto curr post 1d = bson iter into4(Apost 1d child);
if (idu = mestamp = bson_iter_int64(×tamp_cniid);
nost ide ambigo state ;
post_rus.emprace_back(curr_post_ru),
}
redis update map.insert(std::make pair(std::to string(curr timestamp),
std::to string(curr post id)));
bson iter init(&iter 0, doc);
<pre>bson iter init(&iter 1, doc);</pre>
idx++;
}
}

<pre>bool found = mongoc_cursor_next(cursor, &doc); if (found) {</pre>
bson iter t iter;
<pre>if (bson_iter_init(&iter, doc)) {</pre>
bson_iter_t array;
if (bson_iter_find_descendant(&iter, "posts", &array)) {
been iten t neet.
while (bson iter next(sarray)) {
bson iter recurse (&array, &post);
bson iter t post id iter;
<pre>if (bson_iter_find_descendant(&post, "post_id", &post_id_iter)) {</pre>
income is the income is income.
Into4_t post_id = bson_iter_as_into4(&post_id_iter);
post_ids.push_back(post_id);
}
}
}
}
}

Empty-semi-truck Injection

- After retrieving data from Redis, a large vector containing 100 MB of unnecessary data is created
- By introducing this anti-pattern, we illustrate how transferring large amounts of unnecessary data can lead to increased resource consumption, slower performance, and wasted resources.

```
// Simulate empty semi-truck anti-pattern: transferring large amounts of data
unnecessarily
std::vector<char> unnecessary_data(1024 * 1024 * 100, '0'); // 100 MB of
unnecessary data
LOG(info) << "Transferring unnecessary data of size: " << unnecessary_data.size()
<< " bytes";</pre>
```

```
// Transfer the unnecessary data over the network
std::ofstream outfile("unnecessary_data.txt", std::ios::binary);
outfile.write(unnecessary_data.data(), unnecessary_data.size());
outfile.close();
```

//empty semi-truck finished

• After creating the data, it is then written to a file named "unnecessary_data.txt" in binary mode. This file write operation simulates the transfer of the unnecessary data over the network, which impacts the network resources.