

Simulating and Evaluating a Real-World ElasticSearch System using the RECAP DES Simulator

Malika Bendeche ¹, Sergej Svorobej ², Patricia Takako Endo ³, Adrian Mihai ⁴ and Theo Lynn ⁵

¹ School of Computing, Dublin City University, Dublin 9, Ireland; {malika.bendeche}@dcu.ie

² School of Computer Science and Statistics, Trinity College Dublin, Dublin 2, Ireland; {sergej.svorobej}@tcd.ie

³ Universidade de Pernambuco, Pernambuco, Brazil; {patricia.endo}@upe.br

⁴ Opening.io Company, Dublin 2, Ireland; {amo}@opening.io

⁵ Irish Institute of Digital Business, Dublin City University, Dublin 9, Ireland; {theo.lynn}@dcu.ie

Abstract: Simulation has become an indispensable technique for modelling and evaluating the performance of large scale systems efficiently and at a relatively low cost. ElasticSearch (ES) is one of the most popular open source large-scale distributed data indexing systems worldwide. In this paper, we use the RECAP DES simulator, an extension of CloudSimPlus, to model and evaluate the performance of a real-world cloud-based ES deployment by an Irish small and medium-sized enterprise (SME), Opening.io. Following simulation experiments that explored how much query traffic the existing Opening.io architecture could cater for before performance degradation, a revised architecture was proposed, adding a new virtual machine in order to dissolve the bottleneck. The simulation results suggest that the proposed improved architecture can handle significantly larger query traffic (about 71% more) than the current architecture used by Opening.io. The results also suggest that the RECAP DES Simulator is suitable for simulating ES systems and can help companies to understand their infrastructure bottlenecks under various traffic scenarios, and inform optimisation and scalability decisions.

Keywords: Simulation, Modelling; ElasticSearch; DES; CloudSim; CloudSimPlus; Query; Workload; Search engines.

1. Introduction

Search engines are central to our experience of the Internet and a critical building block in how we navigate an ocean of ever-expanding data. In the UK alone search engines were used by 94% of the adult population [1]. They are central to product and service purchase decisions and as such, are an important method for acquiring and retaining consumers for companies of all sizes [2,3]. While search engine interfaces are designed to be easy to use, there is a complex layer of data storage, classification/indexing, and partitioning technologies working in the background, all optimised for speed and relevant data extraction. Research suggests that the functional performance of a search engine, including aspects such as the relevance and usefulness of search results and speed, exerts a strong influence on user satisfaction [3]. At the same time, search engine providers operate under tight margins, while requiring high quality of service (QoS) levels.

ElasticSearch (ES) [4] is one of the most popular open source large-scale distributed data indexing systems worldwide. Its scalable distributed design makes it capable of near real-time search query processing [5]. Optimally configuring search system resources to achieve fast query processing times, both at large- and hyper-scale, is a significant challenge. Performance optimisation using real production or testing environments can lead to substantial resource provisioning costs and can take significant time to stage and evaluate experiments [6]. Simulation modelling offers a powerful

Citation: Title. *Future Internet* 2021, 1, 0. <https://doi.org/>

Received:
Accepted:
Published:

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Copyright: © 2021 by the author. Submitted to *Future Internet* for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

36 alternative solution that can reduce the effort, cost and risk associated with the optimisa-
37 tion/experimentation phase [7–9]. ES search engine systems typically comprise three
38 main architectural components - web crawling, indexing, and query processing. These
39 components contribute to the efficiency and scalability of online search engines [10]. In
40 order to realistically model and simulate an ES search engine system, the simulation
41 model needs to take into account the following: (a) a realistic workload model in the
42 form of requests sent by users to the search engine, (b) a virtual resource provisioning
43 allocation agnostic of the complexities of the underlying data centre hardware, and (c) a
44 data flow logic that is similar to the system implementation under examination.

45 CloudSim is one of the most widely used platforms in both research and academia
46 for modelling and simulating cloud computing infrastructures and services [11,12].
47 CloudSim has been significantly extended since its first version due to its success.
48 CloudSimPlus, in particular, enhanced many engineering features and aspects such as
49 code maintainability, reusability and extensibility, allowing better precision, simplicity,
50 and flexibility [13]. The load distribution on existing CloudSim models only allow
51 cloudlets to be sent from one sender to one destination at pre-determined times, resulting
52 in a one-to-one mapping between a cloudlet and a processing VM. The ES search engine
53 has a distributed architecture that allows for parallel request processing and aggregation,
54 which is not supported by the CloudSim [14,15] framework. To meet these needs,
55 The RECAP Discrete Event Simulator (DES) was designed and developed [16,17]. The
56 RECAP DES extends and introduces new simulation models to CloudSimPlus that take
57 into account the distributed system behaviour of both the architecture and the workload.
58 The RECAP DES was tested and validated by comparing simulation findings with KPI
59 (Key Performance Indicator) traces collected from a live Elasticsearch cluster deployed
60 in public cloud infrastructure by Linknovate.com [16]. The experimental results showed
61 that the proposed RECAP DES simulator supports a range of aspects and features that
62 may aid in search engine based system deployment and provisioning decisions such as:

- 63 • Modelling and simulating a distributed data flow based on a hierarchical architec-
64 ture;
- 65 • Providing customised policies for distributing workload in a hierarchical architec-
66 ture;
- 67 • implementing synchronisation communication barrier between search engine com-
68 ponents for data aggregation; and
- 69 • Providing easy and flexible modelling that can be easily integrated with other
70 CloudSim extensions.

71 In this paper, our main contribution is the extension of the RECAP DES Simulator
72 in order to model, simulate and analyse data traces for an Irish SME, Opening.io. The
73 goals of this analysis are: (i) validate the RECAP DES in other real-world settings, (ii)
74 provide the case site with feedback and insights on their current architectural design,
75 and (iii) recommend and evaluate an alternative architecture that will enable them to
76 handle larger traffic without service degradation.

77 The remainder of this paper is organised as follows. Section 2 briefly summarises
78 selected related work. Section 3 introduces the case site, Opening.io. Section 4 summar-
79 ises the modelling and simulation approach for the ES search engine used in the case
80 site. Section 5 presents and discusses the simulation results of the existing Opening.io
81 architecture and proposed revised architecture. The paper concludes in Section 6.

82 2. Related Work

83 Web search engines, such as ES, are multi-layered complex systems. Therefore
84 system configuration and performance optimisation can consume significant time and
85 resources. Using a live environment for testing, experimentation, and performance eval-
86 uation are not always feasible due to cost and scale limitations. Simulation frameworks
87 can be used as a better alternative to understand system behaviour and minimise real
88 system testing effort [18].

89 There is a relative dearth of literature on simulating web search engines. There are
90 a few research articles in this field. For instance, the Discrete Event System Specification
91 (DEVS) formalism has been used to build search engine simulation models. Inostroza-
92 Psijas et al. [19] used DEVS to model large scale web search engines. To validate
93 their approach, they compared web search engines through MPI implementation and
94 process-oriented simulation. The DEVS approach was also used for search engine user
95 behaviour modelling by Marin et al. [20]. Combining timed coloured Petri Nets, process-
96 oriented simulation, and circulating tokens that represented search engine sequences
97 of operations, they were able to measure the computational costs of search engine
98 queries [20]. Both Inostroza-Psijas et al. [19] and Marin et al. [20] use DEVS as their
99 approach, and data sets sourced indirectly from hyperscale search engines i.e. Yahoo!
100 and AOL, it is unclear to what extent the search engine reflects an actual real-world
101 implementation and supports real-world feasible decisions. Indeed, the perceived utility
102 of their outputs from the data set source is not confirmed. Nasution et al. [21] pursued a
103 mathematical modelling approach developing an adaptive and selective approach for
104 expressing the characteristics of a search engine based on information space constraints.
105 While interesting academically, the data used to test their model is limited and it lacks
106 real-world validation.

107 As discussed in the previous section, Bendeche et al. [16] validated the RECAP
108 DES using real-world ES trace data with an SME. They found that RECAP DES could
109 predict system performance at different scales in an efficient, precise, and accurate man-
110 ner. Despite the popularity of ES search engine, Bendeche et al. [16] is one of the few
111 works that specifically looks at simulating ES search engines. In this paper, we further
112 validate, extend and adapt the RECAP DES to model and simulate the performance of
113 another real-world SME use case that uses ES search engine on a public cloud. In this
114 respect, the study differs from previous works in that it focuses on data sets directly
115 sourced from an SME to address a live real-world issue for the case site, and whose
116 utility is confirmed by the case site. Furthermore, it uses a simulation framework built
117 upon CloudSim, one of the most popular and accessible cloud computing simulation
118 frameworks. This decision provides researchers and practitioners greater scope for
119 future use and extension.

120 3. Use Case: Opening.io

121 Established in 2015 in Dublin, Ireland, Opening.io was named as one of Europe's
122 20 super AI, SaaS and enterprise start-ups in 2018. Opening.io leverages machine in-
123 telligence on top of existing large-scale recruitment processes to accelerate and inform
124 recruitment decision making and job-role matching. Most resumes are structured follow-
125 ing a logical pattern, individual coherent sections covering particular distinct snapshots
126 (time-wise and/or activity delimited) and are typically sorted by importance. Most
127 highlight various skills, abilities and major experience gained over time. The core of
128 Opening.io's processes relies on an analysis of these relationships - both observable as
129 well as latent - in jobs, candidates and interactions data, either unstructured information
130 (raw job descriptions and resume text) or pre-existing relations (employment history,
131 internal categorisation). Opening.io's underlying system performs a number of functions
132 including parsing and analysing resumes, information extraction, matching and rank-
133 ing of candidates in relation to jobs, candidate skills inference and recommendations,
134 CV summaries, salary recommendations, and intelligence around talent pools and the
135 human capital ecosystem at large. In the process of distilling knowledge, vast amounts
136 of information are being indexed. The ES cluster is a central component within the
137 wider Opening.io system and can suffer relatively heavy traffic due to batch ingestion
138 processes (jobs/ candidates imports) and high numbers of search queries performed in
139 parallel. When such scenarios emerge Opening.io observes traffic spikes in the order of
140 300 req/s (writes, ingestion) and 100 req/s (reads, queries).

141 In this work, we model and simulate the ES cluster within the Opening.io system.
 142 In particular, we simulate high traffic scenarios resulting from search queries (read,
 143 queries).

144 4. Modelling and Simulation of Opening.io

145 The Opening.io virtual infrastructure consists of an ES cluster composed of one ES
 146 Node and three Data Nodes. The ES node is also called the coordinator (master node)
 147 that is responsible for: (i) Distributing the queries among the Data Nodes; (ii) Managing
 148 and coordinating the query search results of different Data Nodes; and (iii) returning the
 149 query results to the user.




Figure 1. A high level overview of the Opening.io architecture.

150 Three virtual Data Nodes are used for distributed data storage. The Data Nodes are
 151 responsible for storing and processing old and fresh data. The four nodes (ES and Data
 152 Nodes) are hosted in three virtual machines; VM0, VM1, and VM2, respectively. The
 153 three VMs are hosted on-premise in the same physical machine. Note that both the ES
 154 Node and one of the Data Nodes (Data Node 1) are hosted by the same VM (VM0), as
 155 shown in Figure 1.

156 As previously mentioned, Bendechache et al. [16] demonstrated the RECAP DES
 157 using real-world ES trace data from Linknovate.com, one of the RECAP partners. How-
 158 ever, despite having the same goal of offering an ES service, Opening.io has its own
 159 system specificity, and therefore this study validates the applicability of the RECAP DES
 160 on a discrete system, independent of the RECAP project, and the extensibility of the
 161 RECAP DES in order to accommodate new idiosyncratic requirements. Indeed they
 162 differ in terms of where data resides, the number of nodes that can be hosted in a VM
 163 at any one time, and the nature of their query traffic. In the Linknovate model, the ES
 164 Node collects partial query results from all Data Nodes in parallel then aggregate the
 165 results. In contrast, with Opening.io's model, the entire query result resides in only one
 166 Data Node, therefore the ES Node forwards the query and get the results from a unique
 167 Data Node. Furthermore, in the Linknovate model, VMs can only host one node - an
 168 ES Node or Data Node at one time. In the Opening.io model, the VMs can host one or
 169 two types of nodes - Data Nodes only, or a Data Node and an ES Node. Furthermore,
 170 Linknovate is characterised by only one type of traffic query, (read, queries), whereas

171 the Opening.io model has two types of query traffic, (read, queries) and (write, queries).
 172 Finally, the Opening.io model is characterised by three types of read queries that differ
 173 from each other with respect to their service times; the Linknovate model only has one
 174 type of read query.

175 Table 1 summarises the capacities of the different VMs comprising the architecture
 176 of Opening.io.

Table 1: Opening.io VM features

| VM-ID | Data Node | CPU (#Cores) | RAM (GB) | STORAGE (TB) |
|-------|-----------------------|--------------|----------|--------------|
| VM0 | ES Node & Data Node 1 | 4 | 64 | 1 |
| VM1 | DataNode 2 | 4 | 64 | 1 |
| VM2 | DataNode 3 | 6 | 64 | 1 |

177 4.1. Infrastructure Model

178 The infrastructure simulation model of Opening.io was designed and implemented
 179 to capture the actual analytics engine architecture design provided by the company.
 180 Opening.io deploys its infrastructure in VMs residing in one single powerful physical
 181 machine with VMs communicating with each other through the physical machine's
 182 virtual network. As such, the infrastructure simulation model does not focus on the
 183 network architecture *per se*; it only focuses on the application level.

184 4.2. Application and Workload Propagation Models

185 Application behaviour for the simulation experiment is realised by implementing a
 186 modelling concept that captures data flow through multiple interconnected, distributed,
 187 components (nodes) in the Opening.io ES search engine.

188 Opening.io's architecture contains a set of VMs which can host one or two nodes
 189 (Data Nodes only or Data Node with ES Node) as shown in Figure 1. Each Data Node
 190 contains data files that are indexed by the ES Node. The ES Node takes advantage of
 191 the indexing to speed up the redirection of each query to the Data Nodes that contain
 192 the targeted data files. Requests considered by the Opening.io ES search engine are
 193 characterised by the fact that their entire result resides in only one Data Node. Therefore,
 194 the ES Node forwards the query and gets the result from a unique Data Node.

195 The ES search engine is characterised by a distributed architecture with parallel
 196 request processing and coordination behaviour. The ES Node forwards the query to the
 197 targeted Data Node, collects the result, and returns it as the result of the query. The ES is
 198 also co-hosted in the same VM as one of the Data Nodes (Data Node 1) to which it sends
 199 a query/gets results. While this has the advantage of avoiding the latency of inter-VM
 200 network connections, it might lead to an overload of the CPU and RAM usage of the ES
 201 and the Data Node (See Figure 1).

202 Events and task in the RECAP DES simulator are defined by cloudlets. A cloudlet
 203 represents a submitted job. As such, to simulate the Opening.io workload, each query
 204 is modelled as a series of cloudlets moving through the system's nodes (as shown in
 205 Figure 2).

206 Provided that a node is capable of running a cloudlet, the cloudlet's execution time
 207 is determined by: (i) the cloudlet's overall computational cost, (ii) the total amount of
 208 available CPU per VM, (iii) the number of cores the cloudlet can use in parallel, and (iv)
 209 the amount of CPU and RAM instructions the cloudlet can use at any given time.

210 In ES, the workload is distributed among the nodes depending on a variety of
 211 criteria e.g. the data can be distributed based on the frequency of access to a particular
 212 type of data that resides in a particular node in the system. In the case of Opening.io, the
 213 workload (requests) is distributed over the Data Nodes depending on the type of query
 214 they address (three types based on their service time). Figure 2 shows the application
 215 and workload propagation models. As shown, a collection of cloudlets is created and




Figure 2. Modelling the application and workload propagation of Opening.io using the RECAP DES.

216 executed sequentially when a query is launched. The first cloudlet is executed at the ES
 217 Node. Then, another cloudlet is sent to the targeted Data Node with the appropriate
 218 file (query results). Finally, another cloudlet is sent from the Data Node to the ES as the
 219 result of the query. Therefore, our simulation produces a total number of cloudlets equal
 220 to $Cloudlet_no = 2 + n$ to model a query load, where n is the number of Data Nodes
 221 queried by ES. In the case of Opening.io architecture, $n = 3$.

222 5. Simulation Results

223 The main goals of modelling and simulation of the Opening.io system are:

- 224 1. to validate the RECAP DES in a real-world context;
- 225 2. to evaluate the response time while varying the query traffic in the system to see
 226 how much query traffic can be handled by the Opening.io system;
- 227 3. to propose a new and improved architecture for Opening.io that can cope with a
 228 large increase in the number of requests per second.
- 229 4. to evaluate the response time while varying the query traffic in the improved
 230 Opening.io architecture to show the advantages of the proposed architecture over
 231 the existing design.

232 5.1. Data Set

233 To run our experiments, we used a real query data set provided by Opening.io. The
 234 queries were submitted to the Opening.io search engine between 13:28:00 and 13:47:00
 235 on October 15, 2019, and there are three types of queries. The queries differ from each
 236 other in terms of service time (provided in the data set). Each query targets one Data
 237 Node to return the query result. In this paper, we refer to the three queries in three
 238 different colours, red-query, blue-query and green-query. Table 2 summarises the mean
 239 service time of each of these three queries types.

240 5.2. Existing Opening.io Architecture: Response Time vs Query Traffic

241 In this section, we evaluate the performance of the Opening.io architecture by
 242 looking at the query response time while varying the query traffic (number of queries

Table 2: Characteristics of the queries used

| Query Name | Targeted Data Node | Mean Service Time (ms) |
|-------------|--------------------|------------------------|
| Red-query | Data Node 1 | 21 |
| Blue-query | Data Node 2 | 17 |
| Green-query | Data Node 3 | 19 |

243 received by the system at the beginning of the simulation). The goal of this experiment
 244 is to determine the number of queries per second the system can handle.

245 We monitor the query response time while varying the number of queries/requests
 246 per second received by the system. Note that we use the terms queries and requests
 247 interchangeably. Figure 3 represents a box plot (min, max, lower quartile, upper quartile)
 248 that shows the query response time based on the number of queries the system receives.
 249 Note that the prominent black dots in the figures show outliers of simulation results.






Figure 3. Query response time achieved with different query traffic volumes. The red dotted line represents the threshold (barrier) of 1 second beyond which queries affect the response time of those received in the following second.

250 The query response time increases gradually with the query traffic. The query
 251 response times for the blue and green queries take less than a second to run up to 140
 252 Requests (Reqs), i.e., 46 Reqs for each type. The response time for the red queries starts
 253 to exceed the one second threshold (the barrier represented by the red dotted line in
 254 Figure 3) for running the 140 Reqs traffic. Below this threshold, the system is capable of
 255 handling 140 Reqs with no requests impacting on those arriving in the following seconds.
 256 Beyond that threshold, the execution of requests arriving at a time t would impact the
 257 response time of those arriving at time $t + 1$, and thus creating a snowball effect on the
 258 response time.

259 Figure 4a shows the performance degradation slope within 10 seconds time interval
 260 for the query traffic of 140 Reqs/s. The goal of this experiment is to show that the
 261 increase in query response time for the traffic of 140 Reqs/s is slow. For the first second,
 262 only one type of query (i.e., red) exceeds the one second barrier. As we progress in time,
 263 the other types of queries are also affected and take more time to finish. For example, we
 264 can see that in the second 2, the green queries take longer than one second. The response
 265 time of the blue queries increases slowly with time and does not exceed one second until



(a) Traffic of 140 requests during ten seconds. (b) Traffic of 150 requests during ten seconds.



(c) Traffic of 180 requests during ten seconds.

Figure 4. Query response time achieved for a traffic of (a) 140, (b) 150 and (c) 180 requests during ten seconds

266 the 9th second. This means that the traffic of 140 Reqs/s is the load at which the system
 267 starts to slow down but with a moderate slope.

268 Figure 4b shows that as we increase the query traffic to 150 Reqs/s, the query
 269 response time increases faster. This is due to the fact that the system is more stressed
 270 and slowing down faster.

271 Figure 4c shows that the query response time increases even faster as we increase
 272 the query traffic to 180 Reqs/s. This is explained by the fact that we are approaching the
 273 saturation point. In fact, all three queries exceed one second execution time from the
 274 first second. All the queries in the system are delayed due to the waiting time.

275 5.3. Proposed Opening.io Architecture

276 Following our initial simulation and the insights derived from the experiment
 277 results, a revised architecture for the Opening.io ES architectures was proposed with the
 278 goal of improving its performance.

279 The basic architecture of Opening.io contains three Data Nodes, one of them (Data
 280 Node 1) resides in the same VM as the ES Node (VM0). Therefore, these two nodes are
 281 sharing/competing for the same resources (CPU and RAM). Having fewer resources
 282 results in the response time of queries executed in Data Node 1 (i.e. red) being more
 283 impacted than the response time of other types of queries. This can be seen in Figure 4c.
 284 Furthermore, sharing resources in VM0 also negatively impacts the performance of the
 285 ES Node. While there is no direct monitoring of this performance, indications of this
 286 impact are reflected in the response time of queries run on the other Data Nodes (blue
 287 and green).

288 In this section, we propose and evaluate a new architecture that separates the ES
 289 Node from Data Node 1 by migrating Data Node 1 to another separate VM (VM3) as
 290 shown in Figure 5. The proposed Opening.io architecture contains four VMs instead
 291 of three. In this architecture, VM0, which hosts the ES Node, acts only as an ES Node,

292 indexing and coordinating the Data Nodes. Data Node 1 is relocated into its own VM
 293 (VM3) with 4 CPU cores, 64 GB of RAM and 1 TB of storage. The main goal is to alleviate
 294 and distribute the load on the ES Node and Data Node 1, and verify how the system
 295 will handle larger query traffic.






Figure 5. High level overview of the proposed revised Opening.io architecture

296 We repeat the same experiments to evaluate the query response time while varying
 297 the query traffic using the same three types of queries (red, blue, and green) as done in
 298 the sub-section 5.2. Figures 6a and 6b compare the query response time achieved by the
 299 existing Opening.io architecture and the proposed architecture.



(a) Existing Opening.io architecture.



(b) Proposed architecture

Figure 6. Comparison of query response time achieved with different query traffic volumes. The red dotted line represents the threshold (barrier) of 1s, beyond which queries affect the response time of those received in the following second.

300 Figure 6b shows that the system reaches the threshold (red discontinuous line) at a
 301 later point than with the existing system. In fact, the system executes up to 240 Reqs/s
 302 of all three types of queries in less than one second. This means that the system can
 303 handle comfortably up to 240 Reqs/s. This is due to the ES Node and Data Node 1 each
 304 running in separate VMs giving them more computational resources (CPU and RAM) to
 305 process their tasks.

306 5.4. Discussion

307 The current Opening.io architecture assumes one ES Node and three Data Nodes.
 308 These nodes are hosted in VMs, and the ES Node shares a physical node with one
 309 Data Node (Data Node 1). As the ES Node is the coordinator, and therefore a central
 310 component of the Opening.io architecture, we propose a new arrangement to improve
 311 the system performance as a whole. By introducing an additional VM to separate the

312 ES Node from Data Node 1, simulation results suggest that the proposed architecture
313 could handle greater query traffic than the current architecture used by Opening.io.
314 Isolating the ES Node (separating the ES Node and Data Node 1 into two different
315 VMs) dissolved the bottleneck and, as a result, the ES Node was able to process more
316 requests since it had more computational resources available for executing its tasks.
317 In the proposed architecture, it could respond to 100 additional read Reqs/s. As we
318 increase the query traffic beyond 240 Reqs/s, there is a divergence in query response
319 times in line with the results from the simulation of the existing Opening.io architecture.
320 By adopting a simple strategy of adding a new VM, we could improve the Opening.io
321 architecture performance by 71% more requests. Since the addition of a VM may result
322 in more cost for the owner, a simulation tool acts as a decision support system regarding
323 infrastructure changes. In this case, no changes were proposed at the software level, for
324 examples, new resource allocation algorithms, and therefore additional strategies and
325 mechanisms could be proposed to further improve performance.

326 6. Conclusion

327 In this work, we have described and modelled Opening.io architecture. First, using
328 and adapting the RECAP DES, we simulated the ES cluster implementation based on
329 historical data provided by Opening.io. Next, after the simulation experiment results
330 analysis, a revised architecture for Opening.io was proposed and modelled. Finally,
331 the new application model then was used to evaluate the performance of the proposed
332 architecture using the updated RECAP DES.

333 The goal of modelling and simulating the Opening.io system was to validate the
334 RECAP DES using real-world data from a system independent of the original RECAP
335 project. Specifically, we wished to explore the utility of the RECAP DES in understanding
336 system bottlenecks under various traffic scenarios in order to correctly provision and
337 scale up Opening.io ES-based services. The analysis aimed to understand the software
338 performance by analysing cluster metrics and employ the findings to fine-tune ES and
339 supporting indexing infrastructure.

340 To this end, the goals of the study were achieved. The study demonstrated that
341 the RECAP DES could be used in different ES implementations to provide insights on
342 deployed ES systems and to optimise architecture design. First, our simulation provided
343 important insight into Opening.io capacity planning identifying threshold points at
344 which QoS starts degrading and additional resources must be provisioned i.e. at 140
345 Reqs/s. This assisted the company to understand its existing system architecture and
346 informed the revised architecture. Second, the RECAP DES enabled us to evaluate the
347 performance of the revised architecture and justify re-architecting the existing system.
348 The proposed architecture resulted in a significant performance improvement i.e. 71%
349 more requests before service degradation when compared to the existing architecture.

350 Moving forward, we plan to extend the analysis to consider other performance
351 metrics, such as delay and processing time, and examine the impact of additional
352 VMs and different on these metrics. We also intend to use our proposed simulator
353 to model and simulate larger real-world use cases and extend the complexity of the
354 models to explore a wider set of performance metrics. As indicated earlier, we focus
355 explicitly on changes at the virtual level, future work can explore additional software-
356 level interventions.

357 **Acknowledgments:** This work was funded by European Union's Horizon 2020 research and innov-
358 ation programme under grant agreement No. 732667 (RECAP). The author Malika Bendeche
359 is supported, in part, by Science Foundation Ireland (SFI) under the grants No. 13/RC/2094_P2
360 (Lero) and 13/RC/2106_P2 (ADAPT). The author Sergej Svorobej is partially supported by SFI
361 grant 16/SP/3804 (ENABLE).

362 **Conflicts of Interest:** Declare conflicts of interest or state "The authors declare no conflict of
363 interest."

364 **References**

- 365 1. Adults: Media use and attitudes report 2019. <https://www.ofcom.org.uk>. Accessed:
366 2019-06-07.
- 367 2. Vuylsteke, A.; Wen, Z.; Baesens, B.; Poelmans, J. Consumers' search for information on the
368 internet: how and why China differs from Western Europe. *Journal of Interactive Marketing*
369 **2010**, *24*, 309–331.
- 370 3. Sirdeshmukh, D.; Ahmad, N.B.; Khan, M.S.; Ashill, N.J. Drivers of user loyalty intention
371 and commitment to a search engine: An exploratory study. *Journal of Retailing and Consumer*
372 *Services* **2018**, *44*, 71–81.
- 373 4. Elasticsearch B.V. Open Source Search Analytics - Elasticsearch, 2019.
- 374 5. Kononenko, O.; Baysal, O.; Holmes, R.; Godfrey, M.W. Mining Modern Repositories with
375 Elasticsearch. Proceedings of the 11th Working Conference on Mining Software Repositories;
376 ACM: New York, NY, USA, 2014; MSR 2014, pp. 328–331. doi:10.1145/2597073.2597091.
- 377 6. Buyya, R.; Ranjan, R.; Calheiros, R.N. Modeling and simulation of scalable Cloud computing
378 environments and the CloudSim toolkit: Challenges and opportunities. 2009 international
379 conference on high performance computing & simulation. IEEE, 2009, pp. 1–11.
- 380 7. Svorobej, S.; Takako Endo, P.; Bendeche, M.; Filelis-Papadopoulos, C.; Giannoutakis,
381 K.M.; Gravvanis, G.A.; Tzouvaras, D.; Byrne, J.; Lynn, T. Simulating Fog and Edge Computing
382 Scenarios: An Overview and Research Challenges. *Future Internet* **2019**, *11*, 55.
- 383 8. Bendeche, M.; Svorobej, S.; Takako Endo, P.; Lynn, T. Simulating Resource Management
384 across the Cloud-to-Thing Continuum: A Survey and Future Directions. *Future Internet* **2020**,
385 *12*, 95.
- 386 9. Ashouri, M.; Lorig, F.; Davidsson, P.; Spalazzese, R.; Svorobej, S. Analyzing Distributed
387 Deep Neural Network Deployment on Edge and Cloud Nodes in IoT Systems. 2020
388 IEEE International Conference on Edge Computing (EDGE), 2020, pp. 59–66. doi:
389 10.1109/EDGE50951.2020.00017.
- 390 10. Cambazoglu, B.B.; Baeza-Yates, R. Scalability and Efficiency Challenges in Large-Scale Web
391 Search Engines. Proceedings of the 39th International ACM SIGIR Conference on Research
392 and Development in Information Retrieval; ACM: New York, NY, USA, 2016; SIGIR '16, pp.
393 1223–1226.
- 394 11. Calheiros, R.N.; Ranjan, R.; Beloglazov, A.; De Rose, C.A.; Buyya, R. CloudSim: a toolkit
395 for modeling and simulation of cloud computing environments and evaluation of resource
396 provisioning algorithms. *Software: Practice and experience* **2011**, *41*, 23–50.
- 397 12. Byrne, J.; Svorobej, S.; Giannoutakis, K.M.; Tzouvaras, D.; Byrne, P.J.; Östberg, P.O.;
398 Gourinovitch, A.; Lynn, T. A review of cloud computing simulation platforms and related
399 environments. International Conference on Cloud Computing and Services Science.
400 SCITEPRESS, 2017, Vol. 2, pp. 679–691.
- 401 13. Silva Filho, M.C.; Oliveira, R.L.; Monteiro, C.C.; Inácio, P.R.; Freire, M.M. CloudSim Plus: A
402 cloud computing simulation framework pursuing software engineering principles for im-
403 proved modularity, extensibility and correctness. 2017 IFIP/IEEE Symposium on Integrated
404 Network and Service Management (IM). IEEE, 2017, pp. 400–406.
- 405 14. Mehmi, S.; Verma, H.K.; Sangal, A. Simulation modeling of cloud computing for smart grid
406 using CloudSim. *Journal of Electrical Systems and Information Technology* **2017**, *4*, 159–172.
- 407 15. Hicham, G.T.; Chaker, E.A. Cloud Computing CPU Allocation and Scheduling Algorithms
408 Using CloudSim Simulator. *International Journal of Electrical & Computer Engineering (2088-*
409 *8708)* **2016**, *6*.
- 410 16. Bendeche, M.; Svorobej, S.; Endo, P.T.; Mario, M.N.; Ares, M.E.; Byrne, J.; Lynn, T. Model-
411 ling and simulation of Elasticsearch using CloudSim. 2019 IEEE/ACM 23rd International
412 Symposium on Distributed Simulation and Real Time Applications (DS-RT). IEEE, 2019, pp.
413 1–8.
- 414 17. Spanopoulos-Karalexidis, M.; Papadopoulos, C.K.F.; Giannoutakis, K.M.; Gravvanis, G.A.;
415 Tzouvaras, D.; Bendeche, M.; Svorobej, S.; Endo, P.T.; Lynn, T. Simulating Across the
416 Cloud-to-Edge Continuum. In *Managing Distributed Cloud Applications and Infrastructure*;
417 Springer, 2020; pp. 93–115.
- 418 18. Moysiadis, V.; Sarigiannidis, P.; Moscholios, I. Towards Distributed Data Management in
419 Fog Computing. *Wireless Communications and Mobile Computing* **2018**, 2018.
- 420 19. Inostrosa-Psijas, A.; Wainer, G.; Gil-Costa, V.; Marin, M. DEVs modeling of large scale web
421 search engines. Proceedings of the Winter Simulation Conference 2014. IEEE, 2014, pp.
422 3060–3071.

-
- 423 20. Marin, M.; Gil-Costa, V.; Bonacic, C.; Inostrosa, A. Simulating search engines. *Computing in*
424 *Science & Engineering* **2017**, *19*, 62.
- 425 21. Nasution, M.K. Modelling and simulation of search engine. *Journal of Physics: Conference*
426 *Series*. IOP Publishing, 2017, Vol. 801, p. 012078.