

1
2
3
4
5
6
7
8
9
10

Scalability Analysis Comparisons of Cloud-based Software Services

Amro Al-Said Ahmad^{1,2} and Peter Andras¹

¹School of Computing and Mathematics, Keele University, Newcastle-under-Lyme, UK

²Faculty of Information Technology, Philadelphia University, Amman, Jordan

{a.m.k.al-said.ahmad, p.andras}@keele.ac.uk

11 **Abstract:**

12

13 Performance and scalability testing and measurements of cloud-based software services
14 are necessary for future optimizations and growth of cloud computing. Scalability, elasticity,
15 and efficiency are interrelated aspects of cloud-based software services' performance
16 requirements. In this work, we use a technical measurement of the scalability of cloud-based
17 software services. Our technical scalability metrics are inspired by metrics of elasticity. We
18 used two cloud-based systems to demonstrate the usefulness of our metrics and compare their
19 scalability performance in two cloud platforms: Amazon EC2 and Microsoft Azure. Our
20 experimental analysis considers three sets of comparisons: first we compare the same cloud-
21 based software service hosted on two different public cloud platforms; second we compare
22 two different cloud-based software services hosted on the same cloud platform; finally, we
23 compare between the same cloud-based software service hosted on the same cloud platform
24 with two different auto-scaling policies. We note that our technical scalability metrics can be
25 integrated into a previously proposed utility oriented metric of scalability. We discuss the
26 implications of our work.

27

28 **Keywords:** Measurement; Performance; Scalability; Software-as-a-Service (SaaS);
29 Metrics;

30

31 **1. INTRODUCTION**

32 Cloud-based applications are increasing rapidly as hosting cost have been reduced and
33 computing resources become more available and efficient. In order to maximize the
34 scalability and performance of any software system, it is essential to incorporate performance
35 and scalability testing and assessment into the development lifecycle. This will provide an
36 important foundation for future optimization and will support the Service Level Agreement
37 (SLA) compliant quality of cloud services [1, 2]. There are three typical requirements that are
38 associated with the performance of cloud-based applications: scalability, elasticity, and
39 efficiency [3, 4].

40 In this study, we adopt technical definitions of these performance features, which were
41 identified by Lehrig et al. [5]. Scalability is the ability of the cloud layer to increase the
42 capacity of the software service delivery by expanding the quantity of the software service
43 that is provided. Elasticity is the level of autonomous adaptation provided by the cloud layer
44 in response to variable demand for the software service. Efficiency is the measure of
45 matching the quantity of software service available for delivery with the quantity of demand
46 for the software service. However, we note that alternative, utility-oriented (i.e. economic
47 cost/benefit focused) approaches are also used in the literature for the conceptualization and
48 measurement of these performance aspects of cloud-based services [6, 7]. Technical
49 scalability measurements and testing is key to assessing and measuring the performance of
50 cloud-based software services [1, 8]. Both elasticity and efficiency aspects depend on
51 scalability performance.

52 Cloud Computing, auto-scaling and load-balancing features provide the support for
53 cloud-based applications to be more scalable, which allows such applications to be able to
54 deal with sudden workload by adding more of instance(s) at runtime. Furthermore, as cloud-
55 based applications are being offered as Software as a Services (SaaS), and the use of multi-

56 tenancy architectures [9]; emphasizes the need for scalability that supports the availability
57 and productivity of the services and on-demand resources.

58 A relevant systematic literature review reports, only a few research works (e.g. project
59 reports, MSc theses) which try to address the assessment of technical scalability of cloud-
60 based software services [5]. However, recently a number of publications addressed the
61 technical measurement of the elasticity of cloud-based provision of software services [5, 10].
62 On the other hand, other recent publications address the scalability of cloud-based software
63 services from utility perspective [5–7, 11].

64 In order to try to improve the scalability of any software system, we need to understand
65 the system's components that effect and contribute to scalability performance of the service.
66 This could help to design suitable test scenarios, and provides a basis for future opportunities
67 aiming to maximize the services scalability performance. Assessing scalability from utility
68 perspective is insufficient for the above purpose, as it works from an abstract perspective
69 which is not necessarily closely related to the technical components and features of the
70 system.

71 In this paper, we use technical scalability measurements and metrics for scalability [12]
72 of cloud-based software services, inspired by earlier technical measures of cloud elasticity
73 [13–15], this work is extended from previous works [12], [16]. We demonstrate the metrics
74 application using two cloud-based software services (OrangeHRM and/or MediaWiki) run
75 through the Amazon EC2 and Microsoft Azure clouds. We perform three comparisons, the
76 first one between the same cloud-based software service hosted on two different public cloud
77 platforms. The second comparison is between two different cloud-based software services
78 hosted on the same cloud platform. The third comparison is between the same cloud-based
79 software service hosted on the same cloud platform with different auto-scaling policies. We
80 show how the metrics can be used to show differences in the system behavior based on

81 different scaling scenarios. We discuss how we can use these metrics for measuring and
82 testing the scalability of cloud-based software services.

83 The rest of the paper is organized as follows: Section 2 presents related works. A
84 description of our approach to measuring the scalability of cloud-based software services and
85 our metrics based on this measurement approach are presented in Section 3. Section 4
86 presents our experiments and analyses using two different usage scenarios, and three sets of
87 comparisons to demonstrate the measurement approach and metrics results. Next, we discuss
88 the implications and importance of the approach and metrics in Section 5. Finally, we present
89 our conclusions and future works in Section 6.


90 **2. RELATED WORK**

91 Related reviews [17, 18] highlight scalability and performance testing and assessment
92 for cloud-based software services, as promising research challenges and directions. Another
93 related mapping study [19] highlights that the majority of the studies in software cloud testing
94 present early results, which indicates growing interests across the field and also the potential
95 for much more research to follow the early results.

96 A relevant systematic literature review [5] covers cloud performance assessments and
97 metrics in terms of scaling, elasticity, and efficiency. Highlights of their key findings are:
98 most of the reviewed papers focus on elasticity, and in the term of scalability, they report that
99 the papers were either early and preliminary result or initial ideas of research students. The
100 review [5] provides the definitions of the key performance aspects (scalability, elasticity, and
101 efficiency) which have been adopted in this study. Other similar recent surveys [20, 21] focus
102 primarily on cloud service elasticity.

103 The majority of the studies focus on measuring the elasticity of cloud services from a
104 technical perspective [4, 10, 15, 22–26]. For example, Herbst et al. [4] sets a number of key

105 concepts that allows measuring cloud service elasticity in technical term (see Fig. 1) such as
106 the quantity and time extents for periods of time when the service provision is either below or
107 above what is required by the service demand. Elasticity measures defined by [4, 22] is: the
108 timeshares and average time lengths in under-provisioned and over-provisioned states; the
109 amounts of the over-provisioned and under-provisioned resources per time unit; the averages
110 of the excess and lacking resources; and the jitter, which is the number of resource
111 adaptations during a specific time of provisioning the service. The up-elasticity and the
112 down-elasticity metrics are defined as the reciprocal value of the product of the average
113 under-provisioned/over-provisioned time length and average lack of resources. Further
114 elaboration [23] that extended the above metrics introduced other factors and ways such as
115 reconfiguration time, functions of resource inaccuracy, and scalability.



116

117 Fig. 1. Key concepts for measuring elasticity.

118 From the utility-oriented perspective of measuring and quantifying scalability, we note
119 the work of Hwang et al. [7, 11]. Their production-driven scalability metric includes the
120 measurement of a quality-of-service (QoS) and the cost of that service, in addition to the
121 performance metric from a technical perspective [7, 11]. This approach is useful from a utility
122 perspective, as it depends on multiple facets of the system (including cost measures), it is
123 improbable to be able to provide useful and specific information in terms of contribution of
124 system components to scalability in a technical perspective.

125 Technical-oriented measurements or metrics for cloud-based software scalability
126 research are limited. Such as [4] provides a technical scalability metric, however, this is a
127 rather elasticity driven metric which measures the sum of over- and under-provisioned
128 resources over the total length of time of service provision. While, Jayasinghe et al. [13, 14]
129 provides a technical scalability measure in terms of throughput and CPU utilization of the
130 virtual machines, but the work does not provide a metric or measure. Jamal et al. [27]
131 describe practical measurements of systems throughput with and without multiple virtual
132 machines (VMs), without clearly formulating specific measurements or metric of scalability.
133 Gao et al. [15] evaluate software as services (SaaS) performance and scalability from the
134 capacity of the system perspective, by using the system load and capacity as measurements
135 for scalability. Another recent work [28] focuses on building a model that helps to measure
136 and compare different deployment configurations in terms of costs, capacity, and elasticity.
137 Brataas et al. [29] offered two scalability metrics, one based on the relationship between the
138 capacity of cloud software services and its use of cloud resources; the second is the cost
139 scalability metric function that replaces cloud resources with cost, in order to demonstrate the
140 metrics, they used CloudStore application hosted in Amazon EC2 with different
141 configurations. In an earlier work, [30] provides a theoretical framework of scalability for
142 mobile multi-agent systems, however, which remains limited to theory and modeling results.

143 In terms of comparisons, we note that [13, 14] compared the performance and scalability
144 of two applications (RUBBoS and/or Cloudstone) on three public clouds (Amazon, Open
145 Cirrus, and Emulab), and three private clouds that have been built using the three mainstream
146 hypervisors (XEN, KVM and CVM). As we mentioned above the comparison were based on
147 CPU utilization and throughput without providing any metric or measure. Similarly, Hwang
148 et al. [7, 11] introduces a set of experiments involving five benchmarks, three clouds, and set
149 of different workload generators. Only three benchmarks were considered for scalability

150 measurements, the comparison was based on the scaling scenarios, and what the effect on
151 performance and scalability. Gao et al. [15] run the same experiments in two different AWS
152 EC2 instance types, one with load-balancing and one without. While Vasar et al. [31]
153 introduces a framework for testing web application scalability on the cloud, run the same
154 experiments settings to measure response time on three different EC2 instance types.

155 **3. SCALABILITY PERFORMANCE MEASUREMENT**

156 Scalability is the ability of the cloud-based system to increase the capacity of the software
157 service delivery by expanding the quantity of the software service that is provided when such
158 increase is required by increased demand for the service over a period of time during which
159 the service is exposed to a certain variation in demand for the service (i.e. a demand scenario)
160 [5]. Our focus is whether the system can expand in terms of quantity (scalability) when
161 required by demand over a sustained period of service provision, according to a certain
162 demand scenario. We are not concerned with short-term flexible provision of the resources
163 (elasticity of the service provision) [22]. The purpose of elasticity is to match the service
164 provision with actual amount of the needed resources at any point in time [22]. Scalability is
165 the ability of handling the changing needs of an application within the confines of the
166 infrastructure by adding resources to meet application demands as required, in a given time
167 interval [5, 32]. Therefore, the elasticity is scaling up or down at a specific time, and
168 scalability is scaling up by adding resources in the context of a given time frame. The
169 scalability is an integral measurement of the behavior of the service over a period of time,
170 while elasticity is the measurement of the instantaneous behavior of the service in response to
171 changes in service demand. Furthermore, we are not concerned with the efficiency of the
172 cloud-based software services delivery, which is usually measured by the consumption of
173 resources (i.e. cost and power consumption) required to complete the desired workload [5].

174 The increase of cloud capacity usually happens by expanding the volume of service

175 demands served by one instance of the software or by providing a lower volume of service
176 through multiple instances of the same software, or a combination of these two approaches.
177 Generally, we expect that if a service scales up the increase in demand for service should be
178 matched by the proportional increase in the service's provision without degradation in terms
179 of quality. In this work, the quality of the service may be seen for example in terms of
180 response time.

181 The ideal scaling behavior of the service system should be substantial over a sufficiently
182 long timescale, in contrast with cloud elasticity that looks at short-term mismatches between
183 provision and demand. If the system does not show ideal scaling behavior, it will increase the
184 volume of the service without changing the quality of that service. Ordinarily, real systems
185 are expected to behave below the level of the ideal scaling and the aim of scalability testing
186 and measurements is to quantify the extent to which the real system behavior differs from the
187 ideal behavior.

188 To match the ideal scaling behavior, we expect that the system will increase the quantity
189 of the software instances proportionately with the rise in demand for the software services, i.e.
190 if the demand is doubled, we would ideally expect the base number of software instances to
191 also double. We also expect that the system maintains quality of service in terms of
192 maintaining the same average response time irrespective of the volume of service requests, i.e.
193 if demand was increased by 25%, we would ideally expect no increase in average response
194 time. Formally, let us assume that D and D' are two service demand volumes, $D' > D$. Let I
195 and I' be the corresponding number of software instances that are deployed to deliver the
196 service, and let t_r and t'_r be the corresponding average response times. If the system scales
197 ideally we expect that for any levels of service demand D and D' we have that

$$198 \quad D' / D = I' / I \quad (1)$$

199


$$t_r = t'_r \quad (2)$$

200 Equation (1) means that the volume of software instances providing the service scale up
201 linearly with the service demand. Equation (2) means that the quality of service, in terms of
202 average response time, remains the same for any level of service demand.

203 In order to measure the values of I and t_r the system must perform the delivery of the
204 service over a period of time, such that short-term variations corresponding to system
205 elasticity do not influence the measurements. This means that the measurements should be
206 based on an average number of software instances and average response time measured
207 regularly (e.g. every second) during the execution of a demand scenario following a particular
208 pattern of demand variation. The same demand pattern should be executed multiple times to
209 get reliable averages.

210 Demand scenarios may follow certain patterns expected to test the scalability of the system
211 in specific ways. Two kinds of demand patterns that appear as natural and typical choices are
212 the steady increase followed by a steady decrease of the demand with a set level of the peak,
213 and the stepped increase and decrease, again with a set peak level of demand. The second
214 scenario is a stepped increase and decrease, again with a set peak level of demand; with this
215 scenario, we schedule to start with 10% of the demand size, then stepped increase 10%
216 through time, while stepped down 10% through time. These two demand scenarios are shown
217 in Fig. 2. The purpose of having two scenarios is to see how the auto-scaling service (services
218 that automatically help to ensure that an application has the proper number of instances
219 dynamically, can handle the workload during runtime [33, 34]) handles cloud-based software
220 services with different patterns of growth of workloads and to verify that the cloud resources
221 covers the target system's needs without experiencing a drop in performance. A demand
222 scenario is characterized by a summary measure of the demand level, which may be the peak

223 level or the average or total demand level. This characteristic of a demand scenario is denoted
224 as D.



225

226 Fig. 2. Demand scenarios: A) Steady rise and fall of demand; B) Stepped rise and fall of demand.

227 In general, real-world cloud-based systems are unlikely to deliver the ideal scaling
228 behavior. Given the difference between the ideal and the actual system scaling behavior, it
229 makes sense to measure technical scalability metrics for cloud-based software services using
230 as reference the ideal scalability behavior defined in equations (1) and (2).

231 In terms of provision of software instances for the delivery of the services, the scaling is
232 deficient if the number of actual instances is lower than the ideally expected number of scaling
233 instances. To quantify the level of deficiency we pick a demand scenario and start with a low
234 level of characteristic demand D_0 and measure the corresponding volume of software instances
235 I_0 . Then we measure the number of software instances I_k corresponding to a number (n) of
236 increasing demand levels D_k following the same demand scenario, we can then calculate how
237 close are the I_k values to the ideal I_k^* values (in general we expect $I_k < I_k^*$). Following the ideal
238 scalability assumption of equation (1) we get for the ideal I_k^* values:

239

$$I_k^* = (D_k / D_0) \cdot I_0 \quad (3)$$

240 Considering the ratio between the area defined by the (D_k, I_k) values, $k = 0, \dots, n$, and the
241 area defined by the (D_k, I_k^*) values we get the metric of service volume scalability of the
242 system η_I :

243
$$A^* = \sum_{k=1, \dots, n} (D_k - D_{k-1}) \cdot (I_k^* + I_{k-1}^*) / 2 \quad (4)$$

244
$$A = \sum_{k=1, \dots, n} (D_k - D_{k-1}) \cdot (I_k + I_{k-1}) / 2 \quad (5)$$

245
$$\eta_I = A / A^* \quad (6)$$

246 where A and A^* are the areas under the curves evaluated piecewise as shown in Fig. 3A
 247 calculated for actual and ideal I values and η_I is the volume scalability performance metric of
 248 the system. The system is close to the ideal volume scalability if η_I is close to 1. If the
 249 opposite is the case and η_I is close to 0, then the volume scalability of the system is much less
 250 than ideal.

251 We define the system quality scalability in a similar manner by measuring the service
 252 average response times t_k corresponding to the demand levels D_k . Here, the system average
 253 response time measures as the average time that the system takes to process a request once it
 254 was received. We approximate the ideal average response time as t_0 , following the ideal
 255 assumption of equation (2). The system quality scalability is less than ideal if the average
 256 response times for increasing demand levels increase, i.e. $t_k > t_0$. By considering the ratio
 257 between the areas defined by the (D_k, t_k) values, $k = 0, \dots, n$, and the area defined by the (D_k, t_0)
 258 values we get a ratio that defines a metric of service quality scalability for the system η_t :


259
$$B^* = \sum_{k=1, \dots, n} (D_k - D_{k-1}) \cdot t_0 = (D_n - D_0) \cdot t_0 \quad (7)$$

260
$$B = \sum_{k=1, \dots, n} (D_k - D_{k-1}) \cdot (t_k + t_{k-1}) / 2 \quad (8)$$


261
$$\eta_t = B^* / B \quad (9)$$

262 where B and B^* are the areas under the curves evaluated piecewise as shown in Fig. 3B
 263 calculated for actual and ideal t values and η_t is the quality scalability performance metric of
 264 the system. If η_t is close to 1 the system is close to ideal quality scalability. On the other hand,
 265 if η_t is close to 0 the quality scalability of the system is far from the ideal.

266



267



268 Fig. 3. The calculation of the scalability performance metrics: A) the volume scalability metric is η_v , which is
269 the ratio between the areas A and A^* – see equation (6); B) the quality scalability metric is η_q , which is the ratio
270 between the areas B^* and B – see equation (9). The red lines indicate the ideal scaling behavior and the blue
271 curves show the actual scaling behavior.

272 Figure 3 illustrates the calculation of the two scalability performance metrics. In Fig. 3A,
273 A^* is the area under the red line showing the ideal expectation about the scaling behavior (see
274 equation (1)) and A is the shaded area under the blue curve, which corresponds to the actual

275 volume scaling behavior of the system. The blue curve is expected in general to be under the
276 ideal red line, indicating that the volume scaling is less efficient than the ideal scaling. In Fig.
277 3B, B^* is the shaded area under the red line indicating the expected ideal behavior (see
278 equation (2)) and B is the area under the blue curve, showing the actual quality scaling
279 behavior of the system. Again, in general, we expect that the blue curve is above the ideal red
280 line, indicating that the quality scaling is below the ideal. We chose nonlinear curves for the
281 examples of actual scaling behavior (blue curves in Fig. 3) to indicate that the practical scaling
282 of the system is likely to respond in a nonlinear manner to changing demand.

283 The above-defined scalability metrics allow the effective measurement of technical
284 scalability of cloud-based software services. These metrics do not depend on other utility
285 factors such as cost and non-technical quality aspects. This allows us to utilize these metrics in
286 technically focused scalability tests that aim to spot components of the system that have a vital
287 impact on the technical measurability, and additionally the testing of the impact of any change
288 in the system on the technical system scalability. The scalability performance refers to the
289 service volume and service quality scalability of the software service; these two technical
290 measurements reflect to the performance of the scalability of the cloud-based software
291 services.

292 Applying these metrics to different demand scenarios allows the testing and tuning of the
293 system for particular usage scenarios and the understanding of how system performance can
294 be expected to change as the pattern of demand varies. Such application of these metrics may
295 highlight trade-offs between volume scaling and quality scaling of the system that characterize
296 certain kinds of demand pattern variation (e.g. the impact of the transition from low-frequency
297 peak demands to high-frequency peak demands or to seasonal change of the demand).
298 Understanding such trade-offs can help in tailoring the system to its expected or actual usage.

299 **4. EXPERIMENTAL SETUP AND RESULTS**

300 To validate the volume and quality metrics, we performed experiments on Amazon AWS
301 and Microsoft Azure cloud platforms, we used OrangeHRM and Mediawiki as cloud-based
302 software services. Mediawiki is an open-source wiki software system available from
303 <https://www.mediawiki.org>, OrangeHRM is an open source human resource management
304 software system available from <https://www.orangehrm.com>. The reason for using these two
305 cloud-based software services (OrangeHRM and MediaWiki) is based on the REST-based
306 nature of the applications, which is highly adopted by cloud and application providers. As the
307 architecture of these applications support REST caching to improve performance and
308 scalability; by caching the data and the code, which will reduce the amount of time required to
309 execute each HTTP request and therefor improving response times by serving data more
310 quickly [35, 36].

311 The purpose is to check the scalability performance of cloud-based applications using
312 different cloud environments, configuration settings, and demand scenarios. We applied the
313 similar experimental settings for the same cloud-based system (OrangeHRM) in two different
314 cloud environments (EC2 and Azure). We have changed the parameters for Mediawiki, which
315 runs a different type of instance on AWS EC2 environment. Table 1 illustrates the hardware
316 configurations for both cloud platforms.

317 TABLE 1: HARDWARE CONFIGURATIONS FOR CLOUD PLATFORMS

Platform	Type	CPU Credits/hr	V-CPU(s)	RAM	Price (\$/ Hr)
Amazon EC2 (London)	t2.micro (Linux)	6	1	1	0.0132
	t2.medium (Linux)	24	2	4	0.052
MS Azure (UK South)	Standard A1 (Linux)	6	1	1.75	0.06

318 To provide the scaling of the services we relied on the Auto-Scaling and Load-Balancer
 319 services provided by both Amazon AWS and Microsoft Azure. Furthermore, Amazon
 320 CloudWatch and Azure Monitor services have been configured in order to monitor the
 321 parameters. The Auto-scaling polices (the default policies that are offers by the cloud
 322 providers when setting up an auto-scaling group) that have been used for the first two set of
 323 experiments are given in Table 2.

324 TABLE 2: AUTO-SCALING POLICES

Auto-Scaling Policies	
Add Instance	When $80\% \geq \text{CPUUtilization} < +\text{infinity}$
Remove Instance	When $30\% \leq \text{CPUUtilization} > -\text{infinity}$

325 In this study, we perform three kinds of comparisons, one between the same cloud-based
 326 software hosted on two different cloud platforms (EC2 and Azure). The second comparison is
 327 between two different cloud-based software services hosted on the same cloud platform
 328 (EC2). The third is between the same cloud-based software service hosted on the same cloud
 329 platform (EC2) with different Auto-scaling polices. The parameters of these experiments are
 330 listed in Table 3.

331 TABLE 3: CLOUD-BASED SERVICES, WORKLOAD, AND CLOUD PLATFORM

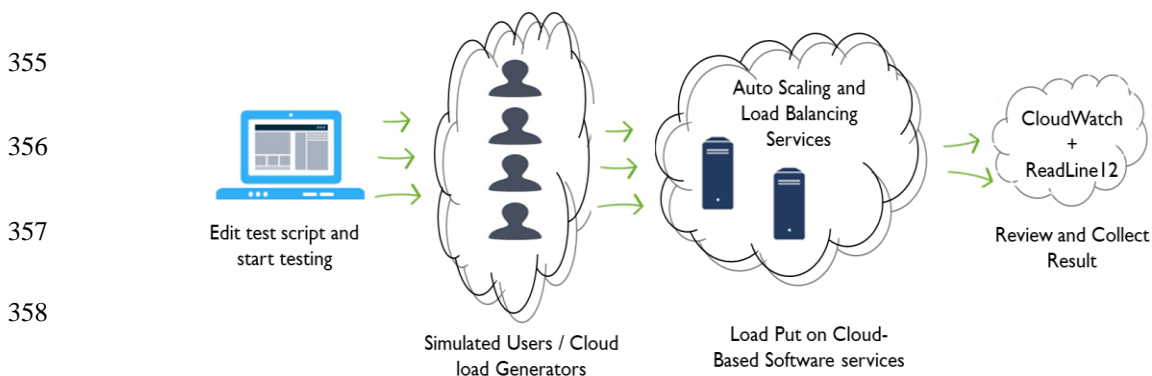
System service	Cloud provider / Instance type	Workload generator
OrangeHRM	Amazon EC2 / t2.micro	JMeter script run by Redline13 services.
OrangeHRM	Microsoft Azure / Standard A1	JMeter script run by Redline13 services.
Mediawiki	Amazon EC2 / t2.medium	Redline13

332 For OrangeHRM experiments (hosted on EC2 and Azure), we simulate the workload using
 333 an Apache JMeter script (<http://jmeter.apache.org/>) and run through Redline13 services after
 334 connecting our cloud accounts to the service (<https://www.redline13.com>).

335 We used Redline13 services by uploading the test script into our account; which allows us
336 to easily deploy JMeter test plans inside our cloud domain and repeat the tests without the
337 need to reset the test parameters again. This allows efficient extraction of the data. The
338 experimental data has been collected through both Redline13 management portal and the
339 monitoring services from EC2 and Azure. The service requests consisted of an HTTP request
340 to all pages and links of OrangeHRM by gaining login access using the following steps via the
341 Apache JMeter:

- 342 • Path = /.
- 343 • Method = GET.
- 344 • Parameters = username, password and login button.


345 We used the Redline13 Pro services to test Mediawiki, which allows us to test the targeted
346 application by covering HTTP requests for all pages and links, including getting
347 authentication (log in) to the application's admin page. In this paper, we report the behavior of
348 the service software in response to the most basic service request, i.e. a generic HTTP request.
349 The JMeter script allows us to send an HTTP/HTTPS request to the targeted application, and
350 parses HTML files for images and other embedded resources (i.e. applets, stylesheets (CSS),
351 external scripts, frames, iframes, background images...etc.), and sends HTTP retrieval requests
352 [37]. For our purposes it was sufficient to issue the simplest HTTP Request, i.e. logging in to
353 the software service and getting in response an acceptance of the login request. Figure 4
354 illustrates our way to test the scalability of cloud-based software services.



359 Fig. 4. Scalability testing process.

360 **4.1 Experimental Process**

361 The cloud resources must be adequately configured to measure up to the workload in order
362 to achieve efficient performance and scalability. We considered two demand scenarios as
363 shown in Fig. 2. The first scenario follows the steady rise and fall of demand pattern (see Fig.
364 2A). The second scenario consists of a series of stepwise increases and falls in demand as
365 shown in Fig. 2B. Examples of the two kinds of experimental demand patterns (users running)
366 are shown in Fig. 5. Figure 5. A is an example of experiments on Mediawiki in AWS EC2 and
367 Fig. 4.B is an example of experiments on OrangeHRM in Microsoft Azure. We varied the
368 volume of demand and experimented with four demand scenarios: 100, 200, 400 and 800
369 service requests in total.



370
371 Fig. 5. Typical experimental demand patterns: A) Mediawiki/EC2 - Steady rise and fall of demand;
372 B) OrangeHRM/Microsoft Azure - Series of step-wise increases and decreases of demand.

373 All experimental settings were repeated 20 times, in total 640 experimental were
374 conducted. The average number of simultaneously active software instances and the average
375 response time for all service requests for each experimental run has been calculated. In this
376 study, the system average response time was measured as the average time that the targeted

377 system takes to process an HTTP request once it was received. The averages and standard
378 deviations of simultaneously active software instances and average response times over the 20
379 experimental runs have been calculated. The standard deviations are included alongside the
380 averages in the results graphs.

381 **4.2 Measured Cloud-based software Services Result**

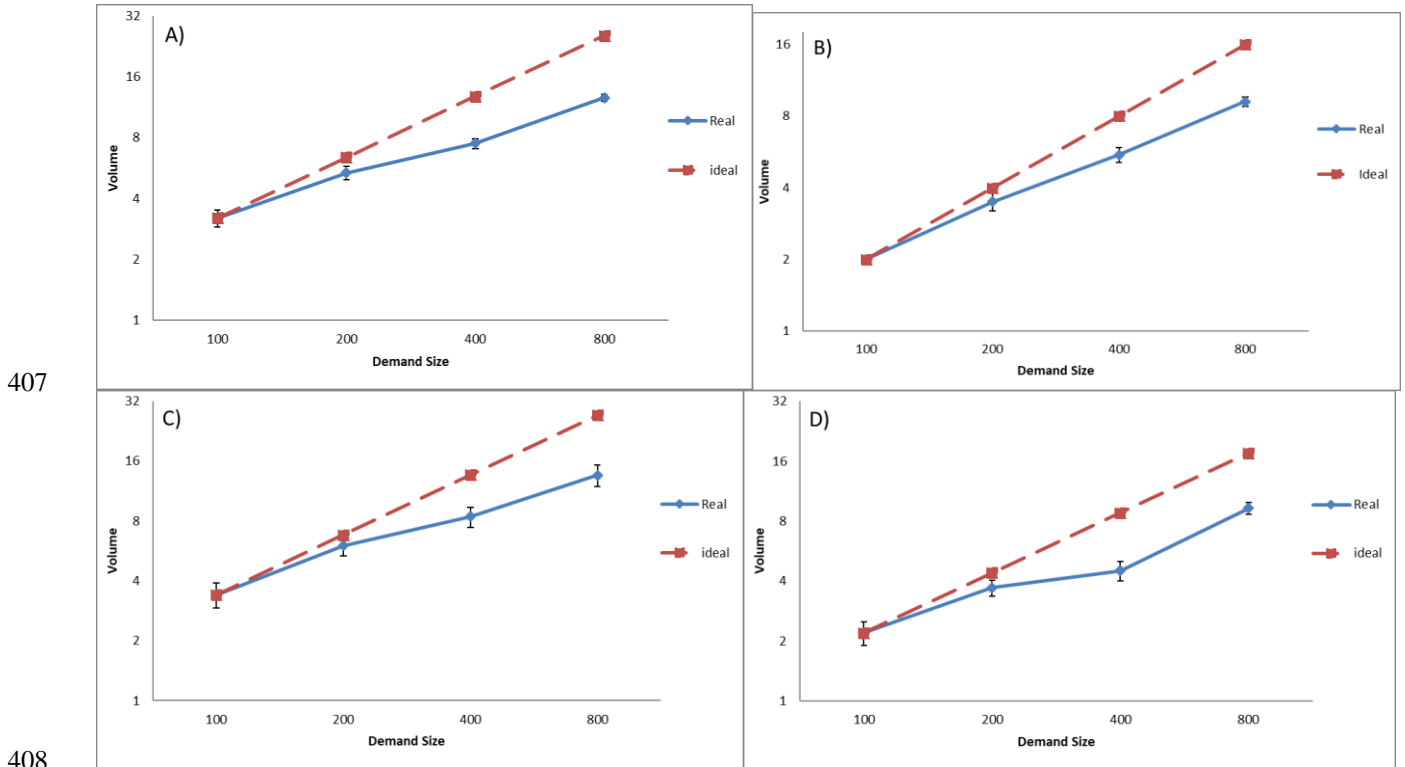
382 *4.2.1 Results for the same cloud-based software system on EC2 and Azure*

383 To achieve fair comparisons between two public clouds, we used similar software
384 configurations, hardware settings, and a workload generator in the experiments. To measure
385 the scalability for the proposed demand scenarios for the first cloud-based software service
386 (OrangeHRM) hosted in EC2 and Azure. The average number of OrangeHRM instances for
387 both scenarios and for the four demand workloads are shown in Fig. 6. The average response
388 times for both scenarios and four demand workloads are shown in Fig. 7. In both figures, the
389 ‘Ideal’ lines show the expected value of average response time, assuming that the scaling of
390 the software service works perfectly. The ‘Real’ curves show the actual measured average
391 response times.

392 We note that there are variations in average response times for the same cloud-based
393 application hosted on two different cloud platforms (EC2 and Azure). So we checked all
394 configurations for instances, Auto-Scaling, and Load-Balancer services for both cloud
395 accounts, to make sure that all settings match. We re-ran a number of tests to make sure that
396 the variations in results are not caused by configuration differences.

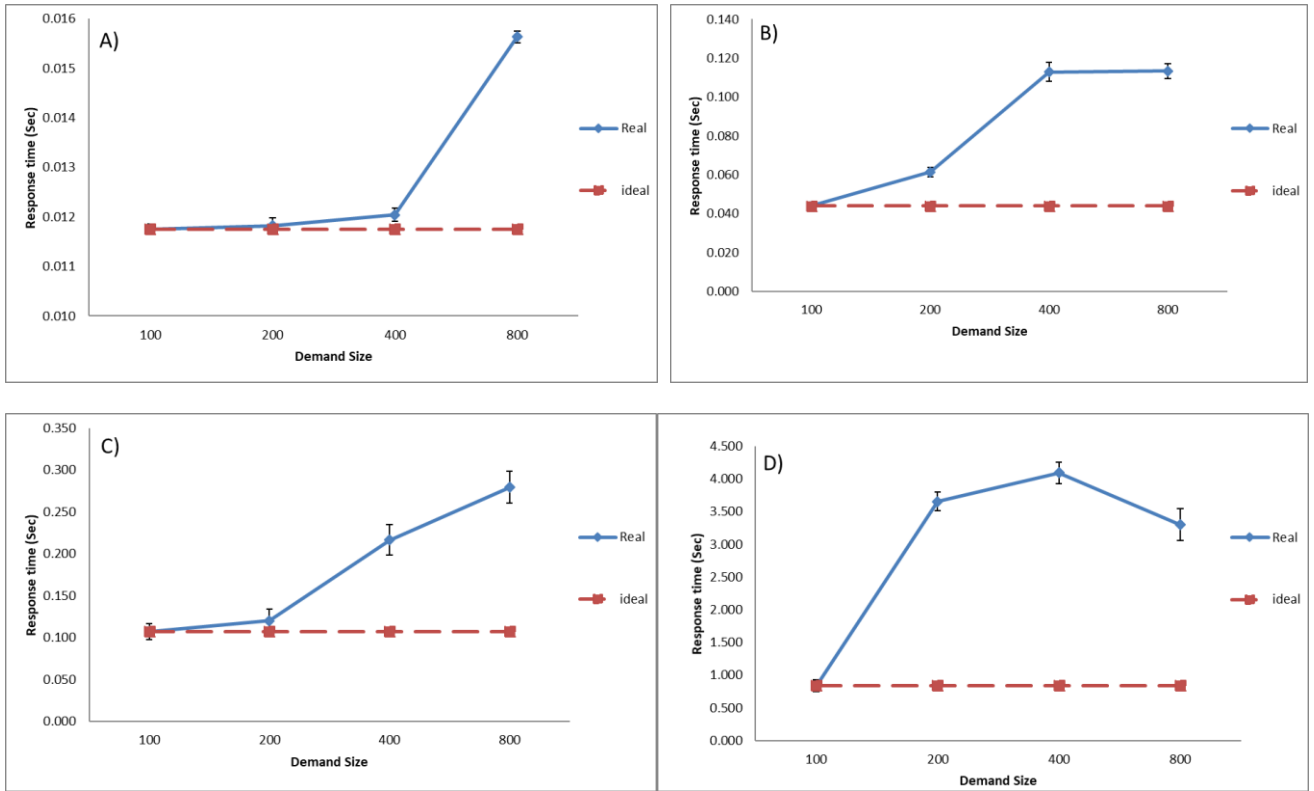
397 We note that there have been other investigations about variations in average response
398 times for cloud-based applications by [38, 39]. There are a number of factors that could cause
399 variations such as: bursty workload, software component management strategies, bursts in
400 system consumption of hardware resources, and network latency. However, all software
401 configurations, hardware settings, and workload generator are similar in our experiments.

402 The observed average response time values for Azure for the stepped rise and fall of
 403 demand scenario are shown in Fig. 7D. Starting from the demand size of 200 the response
 404 time increases significantly. Once the demand size reached 800 the average response time
 405 began to decline significantly. In contrast, response time values for EC2 for the same scenario
 406 which shown in Fig. 7C, have increased gradually with less variation.



409 Fig. 6. The average number of software instances. A) OrangeHRM/EC2 – Steady rise and fall of demand
 410 scenario. B) OrangeHRM/Azure - Steady rise and fall of demand scenario. C) OrangeHRM/EC2– Series of step-
 411 wise increases and decreases of demand scenario. D) OrangeHRM/Azure– Series of step-wise increases and
 412 decreases of demand scenario.

413



414

415

416 Fig. 7. The average response times. A) OrangeHRM/EC2 – Steady rise and fall of demand scenario. B)
 417 OrangeHRM/Azure - Steady rise and fall of demand scenario. C) OrangeHRM/EC2– Series of step-wise
 418 increases and decreases of demand scenario. D) OrangeHRM/Azure– Series of step-wise increases and
 419 decreases of demand scenario.

420 We calculated the scalability metrics η_I and η_f for the two demand scenarios for the cloud-
 421 based application for both cloud platforms. The values of the scalability metrics are shown in
 422 Table 4. The calculated metrics for EC2 show that in terms of volume scalability the two
 423 scenarios are similar, the scaling being slightly better in the context of the step-wise increase
 424 and decrease of demand scenario. In contrast, Azure shows better volume scaling in the first
 425 scenario (Steady rise and fall) with around 0.65, while in the second scenario the volume
 426 scaling performance for the Azure is slightly less than the corresponding performance for the
 427 EC2.

428 In terms of quality scalability, the EC2 hosted system scales much better in the context of
 429 the first scenario, steady rise and fall of demand, than in the case of the second scenario with

430 step-wise increase and decrease of demand. In contrast, Azure shows lower quality scalability
 431 than EC2 in this respect, with the metric being 0.45 in the first scenario, and 0.23 for the
 432 second scenario.

433 TABLE 4: SCALABILITY METRICS VALUES

Cloud Provider	Scenario	Metric	
		η_I	η_t
Amazon EC2	Steady rise and fall	0.5687	0.9041
	Step-wise increase and decrease	0.5882	0.5201
Microsoft Azure	Steady rise and fall	0.6532	0.4526
	Step-wise increase and decrease	0.5592	0.2372

434 We note from the values of both metrics η_I and η_t for both clouds that software system
 435 performed better with respect to both volume and quality in the first scenario, steady rise and
 436 fall of demand, which is more realistic and simpler demand scenario for many cloud-based
 437 software services. In general, we conclude that OrangeHRM performed better in Amazon
 438 EC2, in the terms of quality scalability, while performed slightly better in Azure in the terms
 439 of volume scalability for the steady rise and fall demand scenario. In the case of the variable
 440 rise and fall of demand, the OrangeHRM performs considerably better on the EC2 than on the
 441 Azure.

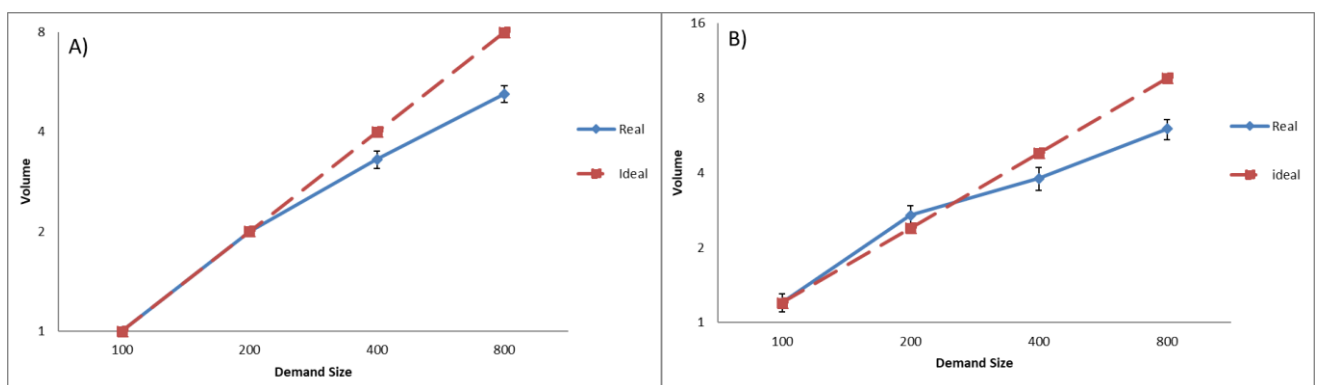
442 The big difference in the average response times for the software system running on the
 443 two cloud platforms indicates that either the software system is tailored better to the provisions
 444 of the EC2 system or that the Azure might have issues with the speed of service delivery for
 445 the kind of service software systems like the OrangeHRM (or for some particular kind of
 446 technical aspect of this software system). Both options raise interesting questions and

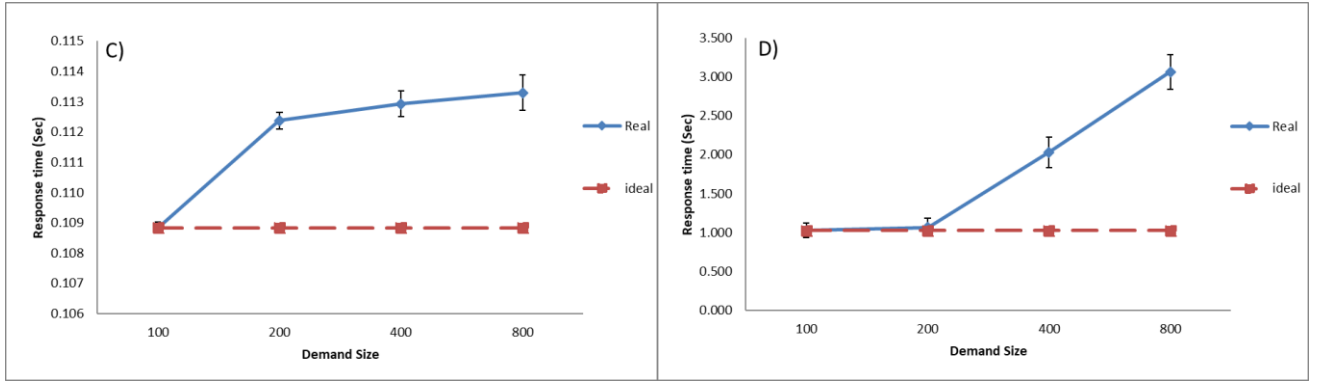
447 opportunities for further investigation of the technical match between a software system and
448 the cloud platforms on which it may run.

449 **4.2.2 Results for different cloud-based software systems on EC2**

450 We used different software configurations, hardware settings, and workload generator in
451 this set of experiments to measure the scalability of the two scenarios for both cloud-based
452 software services that have been hosted in EC2. We changed the instance type and the
453 workload generator in order to see the changes in scalability performance when using different
454 and larger experimental settings. The purpose of this kind of comparison is to see the effects
455 on the scalability performance using the same cloud platform while using different types of
456 instances and workload generators. The average number of OrangeHRM instances for both
457 scenarios and for the four demand workload levels are shown in Fig. 6A and Fig. 6C. The
458 average numbers of MediaWiki instances for both scenarios and for the four workload levels
459 are shown in Fig. 8A and Fig. 8B. The average response times of OrangeHRM for both
460 scenarios and four demand workload levels are shown in Fig. 7A and Fig. 7C. The average
461 response times of MediaWiki for both scenarios and for the four workload levels are shown in
462 Fig. 8C and Fig. 6D.

463





464

465 Fig. 8. The average response times and number of software instances for MediaWiki in EC2. A,B) Average
 466 number of software instances- Steady rise and fall of demand scenario, Series of step-wise increases and
 467 decreases of demand scenario respectively. C,D) Average response times – Steady rise and fall of demand
 468 scenario, Series of step-wise increases and decreases of demand scenario respectively.

469 We note that in the case of the MediaWiki we found a case of over-provisioning of
 470 software instances, i.e. when the measured average number of software instances is larger than
 471 what would be expected as ideal performance according to equation (1) – see Fig. 8B. Given
 472 that we found this for the scenario with many stepwise up and down changes of the demand, a
 473 possible reason for this is the slow or delayed down-elastic response of the cloud platform.
 474 Our volume performance metric does not account for over-provision as it assumes by default
 475 under-provision. Consequently, the over-provision, in this case, distorts somewhat the
 476 performance metric (increases it). One way to correct for this is to include a penalty for over-
 477 provisioning. Considering the symmetric nature of the deviation from the idea (downward or
 478 upward) in terms of its impact on the performance and on the geometric calculations in
 479 equation (5), we can modify this equation as follows:

$$480 \quad A = \sum_{k=1, \dots, n} (D_k - D_{k-1}) \cdot (I_k - 2 \cdot [I_k - I_k^*]_+ + I_{k-1} - 2 \cdot [I_{k+1} - I_{k+1}^*]_+) / 2 \quad (10)$$

481 where $[x]_+$ represents the value of x if it is positive and 0 otherwise. This change of the
 482 calculation avoids the distortion of the metric caused by potential over-provision.

483 Table 5 shows the calculated values for the scalability metrics η_I and η_t for the two demand
 484 scenarios for both OrangeHRM and MediaWiki cloud-based systems. The corrected volume
 485 scalability performance metric, according to equation (10), for the MediaWiki for the second
 486 scenario is reported in Table 5 in italics.

487

TABLE 5: SCALABILITY METRICS VALUES

Cloud-Based System	Scenario	Metric	
		η_I	η_t
OrangeHRM	Steady rise and fall	0.5687	0.9041
	Step-wise increase and decrease	0.5882	0.5201
MediaWiki	Steady rise and fall	0.7556	0.9664
	Step-wise increase and decrease	0.7421 <i>0.7183</i>	0.5012

488 The calculated metrics show that in terms of volume scaling the two scenarios give similar
 489 performance metrics for both systems. The scaling is slightly better in the context of the
 490 scenario with step-wise increase and decrease of demand for OrangeHRM. In contrast, for
 491 MediaWiki, the performance metrics indicate that the software performs slightly better in the
 492 first scenario, steady rise and fall of demand than in the second scenario. In terms of quality
 493 scalability, both systems scale much better in the context of the first scenario, steady rise and
 494 fall of demand, than in the case of the second scenario with step-wise increase and decrease of
 495 demand.

496 Comparing the two software systems running on the EC2, the metrics show that the
 497 MediaWiki runs at a considerably higher volume scalability performance than the
 498 OrangeHRM in both demand scenarios. The quality scalability metrics show at the MediaWiki
 499 has higher performance than the OrangeHRM in this respect in the first scenario and the
 500 performances are relatively close in this sense in the case of the second scenario. One possible

501 factor behind the different volume scalability performance is that we ran the MediaWiki on
502 t2.medium virtual machines, while the OrangeHRM was run on t2.micro virtual machines.
503 Interestingly this difference in the virtual machines made no major difference to the quality
504 scaling of the two software systems. In principle, the difference in the volume scalability
505 performance may point to the possibility that technical solutions in the MediaWiki system
506 support more the volume scaling of the system than the corresponding solutions in the
507 OrangeHRM. A deeper insight and investigation into the components of these systems
508 responsible for the performance difference could deliver potentially significant improvements
509 to the system with the weaker scalability performance metrics.

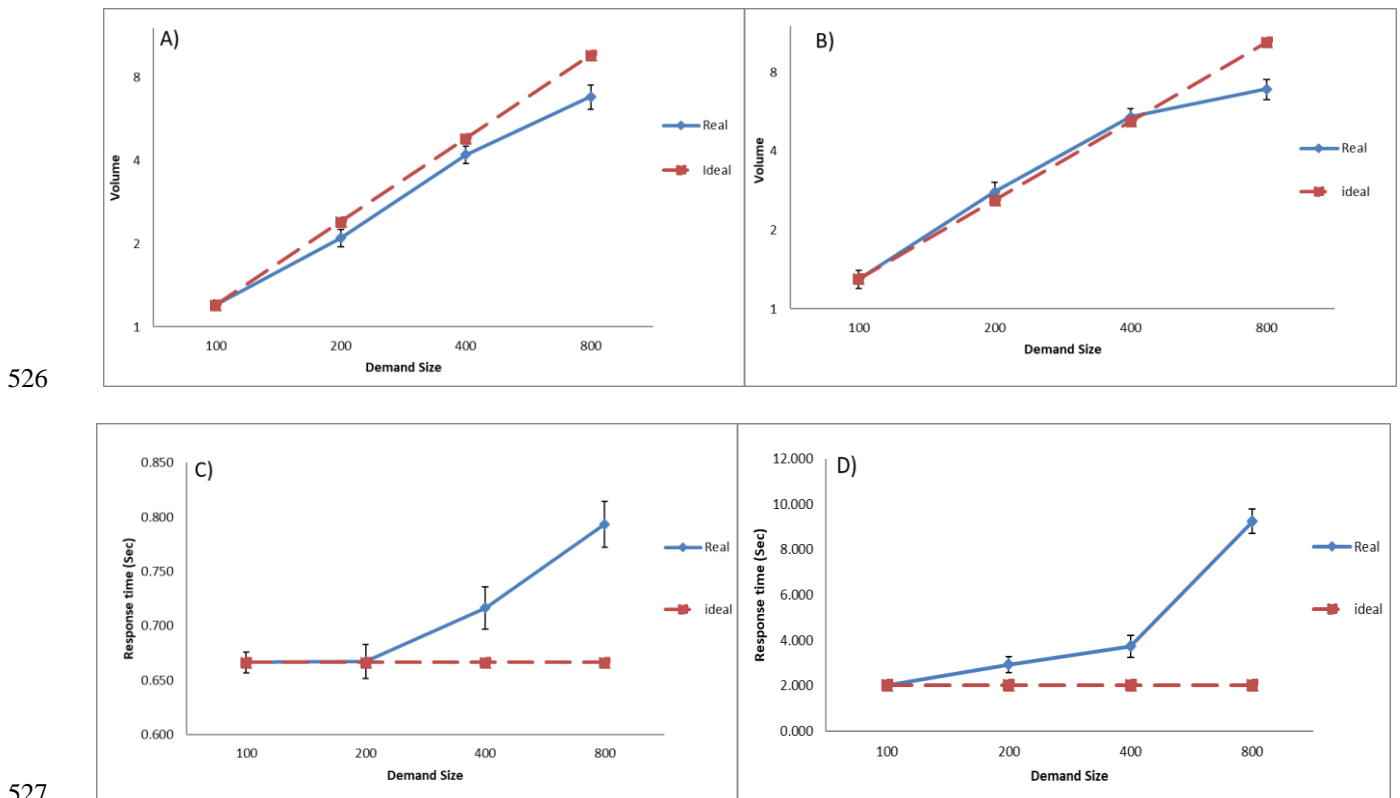
510 **4.2.3 Results for the same cloud-based software system on EC2 with different Auto-**
511 **scaling policies**

512 We used the same software configurations, hardware settings, and workload generator in
513 this set of experiments to measure the scalability of the two scenarios for the same cloud-
514 based software services that have been hosted in EC2, with different Auto-Scaling policies.
515 The first set of policies are the default policies that are provided by EC2 cloud when setting
516 up an Auto-Scaling group (option 1). We pick out random scaling policies for the second set
517 of experiments (option 2). The Auto-scaling policies that have been used for this set of
518 experiments are given in Table 6.

519 TABLE 6: AUTO-SCALING POLICES

Auto-Scaling Policies		
Option 1	Add Instance	When 80% \geq CPUUtilization $<$ +infinity
	Remove Instance	When 30% \leq CPUUtilization $>$ -infinity
Option 2	Add Instance	When 70% \geq CPUUtilization $<$ +infinity
	Remove Instance	When 10% \leq CPUUtilization $>$ -infinity

520 The purpose of this kind of comparison is to see the effects on the scalability performance
 521 using the same cloud platform while using same types of instances and workload generators,
 522 with different auto-scaling policies. The average number of MediaWiki instances (Option 2)
 523 for both scenarios are shown in Fig. 9.A,B. The average response times of MediaWiki (Option
 524 2) for both scenarios shown in Fig. 9.C,D. The average response times and number of software
 525 instances for MediaWiki in EC2 (Option 1) - see Fig. 8.



527
 528 Fig. 9. The average response times and number of software instances for MediaWiki in EC2 (Option 2). A,B)
 529 Average number of software instances- Steady rise and fall of demand scenario, Series of step-wise increases
 530 and decreases of demand scenario respectively. C,D) Average response times – Steady rise and fall of demand
 531 scenario, Series of step-wise increases and decreases of demand scenario respectively.

532 We note two cases of over-provisioning of MediaWiki software instances for both 200 and
 533 400 demand size, when we used new set of auto-scaling policies – see Fig. 8B. Table 7 shows
 534 the calculated values for the scalability metrics η_I and η_t for the two demand scenarios for
 535 MediaWiki cloud-based systems for both auto-scaling policies options. The corrected volume

536 scalability performance metric, according to equation (10), for the second scenario is reported
 537 in Table 7 in italics.

538 In the term of average response time, we note that there are big differences in the average
 539 of response times for the second scenario as it gradually from 2.035 seconds for demand size
 540 100 to 9.24 seconds for demand size 800. While it graduates from 1.02 seconds for demand
 541 size 100 to 3.06 seconds for demand size 800, for the second scenario- Step-wise increase
 542 and decrease.

543 TABLE 7: SCALABILITY METRICS VALUES

Cloud-Based System	Scenario	Metric	
		η_r	η_t
MediaWiki (Auto-Scaling policies option 1)	Steady rise and fall	0.7556	0.9664
	Step-wise increase and decrease	0.7421 <i>0.7183</i>	0.5012
MediaWiki (Auto-Scaling policies option 2)	Steady rise and fall	0.7923	0.9202
	Step-wise increase and decrease	0.8510 <i>0.8217</i>	0.4060

544 We note in term of volume scaling that the experiments of MediaWiki with the second
 545 option of auto-scaling policies, increased 4% and 11% for the first and second scenarios
 546 respectively. While in term of quality scaling the the values has decreased 4.5% and 10% for
 547 the first and second scenarios respectively. If we draw a comparison between the two options
 548 of auto-scaling policies, we note that efficiency is increased when we used the default auto-
 549 scaling policies (option 1).

550 **5. DISCUSSION**

551 The scalability metrics [12] address both volume and quality scaling of cloud-based
552 software services and provide a practical measure of these features of such systems. This is
553 important in order to support effective measurement and testing the scalability of cloud-based
554 software systems. These metrics are distinct from elasticity oriented metrics [4].

555 We used two demand scenarios to demonstrate the effect of demands patterns on scaling
556 metrics. Using more than one scenario can be used to improve cloud-based software services
557 to fit specified demand scenario expectations. This can be useful, to track changes in such
558 scenarios that trigger interventions in terms of systems upgrade or maintenance or direct
559 investment of software engineering resources in the development of focused upgrades for the
560 system. Demand scenarios combined with multi-aspects of quality scaling metric can also be
561 used to determine rational QoS expectations and likely variations depending on changes in
562 demand scenarios.

563 Here we use the quality scalability metric defined by considering the system average
564 response time. Alternative quality scaling metrics may be defined by considering other quality
565 aspects of the system such as system throughput or recovery rate [11]. Expanding the range of
566 quality measurements provides a multiple factor view of quality scalability to support the
567 trade-off options in the context of QoS offerings in the case of service scaling.

568 We understand the importance and need for utility-perspective scalability metric and
569 measurements. Therefore, our proposed metrics can be integrated into the utility-oriented
570 scalability metric proposed by Hwang et al. [11], by combining our metrics as the performance
571 and/or quality components of their utility-oriented scalability metric. This will allow the
572 analysis of the scalability of cloud-based software services from both technical and

573 production-driven perspectives. The utility oriented productivity metric ($P(\Lambda)$) is given as
 574 [11]:

$$575 \quad P(\Lambda) = p(\Lambda) \cdot \omega(\Lambda) / c(\Lambda) \quad (11)$$

576 where Λ is the system configuration, $p(\Lambda)$ is the performance component of the metric – in our
 577 case this is the volume scalability metric, $\omega(\Lambda)$ is the quality component of the metric – in our
 578 case this is the quality scaling metric, and $c(\Lambda)$ is the cost component of the metric. This leads
 579 to a re-definition of the utility-oriented metric as:

$$580 \quad P(\Lambda) = \eta_l(\Lambda) \cdot \eta_t(\Lambda) / c(\Lambda) \quad (12)$$

581 We calculated the integrated scalability metric (see costs in Table 1) for the two demand
 582 scenarios for all cloud-based applications for both cloud platforms. The values of the utility-
 583 oriented scalability metrics are shown in Table 8 – note that the MediaWiki experiments used
 584 more powerful and more expensive virtual machines than the experiments with the
 585 OrangeHRM on the EC2. Our utility oriented scalability calculations show that in the case of
 586 the systems that we compared the best choice is to use smaller and cheaper virtual machines
 587 on the EC2. The corrected integrated scalability metric, based on equation (10), for the
 588 MediaWiki for the second scenario, is reported in Table 8 in italics.

589 TABLE 8: INTEGRATED SCALABILITY METRIC VALUES

Cloud-Based System / Cloud provider	Scenario	Integrated Metrics
OrangeHRM / EC2	Steady rise and fall	38.95
	Step-wise increase and decrease	23.18
OrangeHRM / Azure	Steady rise and fall	4.93
	Step-wise increase and decrease	2.21
MediaWiki (Auto-Scaling policies option 1)	Steady rise and fall	14.04

Cloud-Based System / Cloud provider	Scenario	Integrated Metrics	
	Step-wise increase and decrease	7.15	6.92
MediaWiki (Auto-Scaling policies option 2)	Steady rise and fall	14.02	
	Step-wise increase and decrease	6.64	6.42

590 The technical scalability metrics that we used in this paper allow exploring in more detail
591 the contribution to the system scalability of various components and techniques used in
592 software systems. By instrumenting the software system [40] it becomes possible to determine
593 these contributions and using this information to improve the system. Potentially, different
594 components, technologies or technical solutions may fit different degree with the cloud
595 platform’s provisions. The technical scalability metrics that we used here combined with
596 instrumentation could allow the identification of best matches that can improve the system
597 scalability.

598 6. CONCLUSIONS AND FUTURE WORK

599 In this paper, we demonstrate the use of two technical scalability metrics for cloud-based
600 software services for the comparison of software services running on the same and also on
601 different cloud platforms. The underlying principles of the metrics are conceptually very
602 simple and they address both the volume and quality scaling performance and are defined
603 using the differences between the real and ideal scaling curves. We used two demand
604 scenarios, two cloud-based open source software services (OrangeHRM and MediaWiki) and
605 two public cloud platforms (Amazon AWS and Microsoft Azure). Our experimental results
606 and analysis show that the metrics allow clear assessments of the impact of demand scenarios
607 on the systems, and quantify explicitly the technical scalability performance of the cloud-
608 based software services. The results show that the metrics can be used effectively to compare

609 the scalability of software on cloud environments and consequently to support deployment
610 decisions with technical arguments.

611 Some interesting scalability behavior has been noted through the analysis, such as big
612 variations in average response time for similar experimental settings hosted in different clouds.
613 A case of over provision state has been accrued when using higher capacity hardware
614 configurations in the EC2 cloud.

615 We believe that the technical-based scalability metrics can be used in designing and
616 performing scalability testing of cloud-based software systems, in order to identify system
617 components that critically contribute to the technical scaling performance. We have shown the
618 integration of our technical scalability metrics into a previously proposed utility oriented
619 metric. Our metrics can also be extended, by considering multiple service quality aspects and
620 combined with a range of demand scenarios to support the fine-tuning of the system. Such
621 things can help the identification of QoS trade-offs, and estimation of genuine scalability
622 performance expectations about the system depending on demand scenarios.

623 Future work will include the consideration of other cloud platforms (e.g. Google Cloud,
624 IBM), demand workload generators, and other cloud-based software services, in order to
625 extend the practical validity of the work. We also aim to consider further demand patterns
626 (such as variable width sudden peaks in demand, seasonal demand) to see the impact of these
627 scenarios on the scalability performance of cloud-based software services. Another aspect of
628 future work will focus on using whole code instrumentation technique in order to identify the
629 software system or cloud platform components that contribute critically to variations in
630 average response times for the same cloud-based application with the similar experimental
631 settings in different clouds.

632

633 **DECLARATIONS**

634 *Abbreviations*

635 SLA: Service level agreement.

636 QoS: Quality of service.

637 SaaS: Software as services.

638 VMs: Virtual machines.

639 η_l : Volume scalability metric.

640 η_t : Quality scalability metric.

641 D and D': Service demand volumes.

642 I and I': The corresponding number of software instances.

643 t_r and t'_r : The corresponding average response times.

644 *Availability of data and materials*

645 Not applicable.

646 *Competing interests*

647 The authors declare that they have no competing interests.

648 *Funding*

649 Not applicable.

650 *Authors' contributions*

651 The core of this paper is based on work developed for Al-Said Ahmad's PhD project at the
652 University of Keele, supervised by Peter Andras. Both authors read, edited, and approved the
653 final manuscript.

654 *Acknowledgements*

655 This research is supported by a PhD scholarship from Philadelphia University – Jordan for
656 Amro Al-Said Ahmad.

657 *Authors' information*

658 Amro Al-Siad Ahmad has a PhD in scalability analysis of cloud-based systems (2019) from
659 Keele University, UK. Prior to his PhD, he obtained bachelor degree in Software Engineering
660 (2009) from Philadelphia University in Jordan, and a Master Degree in Computer Science
661 (2014) with distinction from Amman Arab University, Jordan. He works in the areas of
662 scalability of cloud computing and software engineering.

663 Peter Andras has a BSc in computer science (1995), an MSc in artificial intelligence (1996)
664 and a PhD in mathematical analysis of neural networks (2000), all from the Babes-Bolyai
665 University, Cluj, Romania. He is a Professor in the School of Computing and Mathematics,
666 Keele University, UK. He has published 2 books and over 100 papers. He works in the areas
667 of complex systems, computational intelligence and computational neuroscience. Dr. Andras
668 is Senior Member of IEEE, member of the International Neural Network Society, of the
669 Society for Artificial Intelligence and Simulation of Behaviour, and Fellow of the Royal
670 Society of Biology.

671

672 **REFERENCES**

- 673 1. Liu HH (2011) Software performance and scalability: a quantitative approach. John Wiley &
674 Sons, Hoboken, N.J
- 675 2. Atmaca T, Begin T, Brandwajn A, Castel-Taleb H (2016) Performance Evaluation of Cloud
676 Computing Centers with General Arrivals and Service. *IEEE Trans Parallel Distrib Syst*
677 27:2341–2348 . doi: 10.1109/TPDS.2015.2499749
- 678 3. Becker M, Lehrig S, Becker S (2015) Systematically Deriving Quality Metrics for Cloud
679 Computing Systems. In: *Proceedings of the 6th ACM/SPEC International Conference on*
680 *Performance Engineering - ICPE '15*. ACM, New York, NY, USA, pp 169–174
- 681 4. Herbst NR, Kounev S, Reussner R (2013) Elasticity in Cloud Computing: What It Is , and
682 What It Is Not. In: *Presented as part of the 10th International Conference on Autonomic*
683 *Computing*. USENIX, San Jose, CA, pp 23–27
- 684 5. Lehrig S, Eikerling H, Becker S (2015) Scalability, elasticity, and efficiency in cloud
685 computing: A systematic literature review of definitions and metrics. In: *Proceedings of the*
686 *11th International ACM SIGSOFT Conference on Quality of Software Architectures - QoSA*
687 *'15*. pp 83–92
- 688 6. Buyya R, Ranjan R, Calheiros RN (2010) InterCloud : Utility-Oriented Federation of Cloud
689 Computing Environments for Scaling of. In: Hsu C-H, Yang LT, Park JH, Yeo S-S (eds)
690 *Algorithms and Architectures for Parallel Processing (10th International Conference, ICA3PP*
691 *20)*. Springer Berlin Heidelberg, Berlin, Heidelberg, pp 13–31
- 692 7. Hwang K, Shi Y, Bai X (2015) Scale-out vs. scale-up techniques for cloud performance and
693 productivity. In: *Proceedings of the International Conference on Cloud Computing*
694 *Technology and Science, CloudCom*. pp 763–768
- 695 8. Blokland K, Mengerink J, Pol M (2013) Testing Cloud Services: How to Test SaaS, PaaS &
696 IaaS. Rocky Nook
- 697 9. Aljahdali H, Albatli A, Garraghan P, et al (2014) Multi-tenancy in cloud computing. In:
698 *Proceedings - IEEE 8th International Symposium on Service Oriented System Engineering,*

- 699 SOSE 2014. pp 344–351
- 700 10. Islam S, Lee K, Fekete A, Liu A (2012) How a consumer can measure elasticity for cloud
701 platforms. In: Proceedings of the third joint WOSP/SIPEW international conference on
702 Performance Engineering - ICPE '12. ACM, New York, NY, USA, p 85
- 703 11. Hwang K, Bai X, Shi Y, et al (2016) Cloud Performance Modeling with Benchmark
704 Evaluation of Elastic Scaling Strategies. *IEEE Trans Parallel Distrib Syst* 27:130–143 . doi:
705 10.1109/TPDS.2015.2398438
- 706 12. Al-Said Ahmad A, Andras P (2018) Measuring the Scalability of Cloud-Based Software
707 Services. In: 2018 IEEE World Congress on Services (SERVICES). IEEE, San Francisco, CA,
708 pp 5–6. doi: 10.1109/SERVICES.2018.00016
- 709 13. Jayasinghe D, Malkowski S, Wang Q, et al (2011) Variations in performance and scalability
710 when migrating n-tier applications to different clouds. In: Proceedings - 2011 IEEE 4th
711 International Conference on Cloud Computing, CLOUD 2011. pp 73–80
- 712 14. Jayasinghe D, Malkowski S, Li J, et al (2014) Variations in performance and scalability: An
713 experimental study in IaaS clouds using multi-tier workloads. *IEEE Trans Serv Comput*
714 7:293–306 . doi: 10.1109/TSC.2013.46
- 715 15. Gao J, Pattabhiraman P, Bai X, Tsai WT (2011) SaaS performance and scalability evaluation
716 in clouds. In: Proceedings - 6th IEEE International Symposium on Service-Oriented System
717 Engineering, SOSE 2011. IEEE, pp 61–71
- 718 16. Al-Said Ahmad A, Andras P (2018) Measuring and Testing the Scalability of Cloud-based
719 Software Services. In: 2018 Fifth International Symposium on Innovation in Information and
720 Communication Technology (ISIICT). Amman, pp 1–8. doi: 10.1109/ISIICT.2018.8613297
- 721 17. Jennings B, Stadler R (2015) Resource Management in Clouds: Survey and Research
722 Challenges. *J Netw Syst Manag* 23:567–619 . doi: 10.1007/s10922-014-9307-7
- 723 18. Gao J, Bai X, Tsai WT, Uehara T (2013) SaaS testing on clouds - Issues, challenges, and
724 needs. In: Proceedings - 2013 IEEE 7th International Symposium on Service-Oriented System
725 Engineering, SOSE 2013. pp 409–415

- 726 19. Al-Said Ahmad A, Brereton P, Andras P (2017) A Systematic Mapping Study of Empirical
727 Studies on Software Cloud Testing Methods. In: Proceedings 2017 IEEE International
728 Conference on Software Quality, Reliability and Security Companion, QRS-C 2017. pp 555–
729 562. doi: 10.1109/QRS-C.2017.94
- 730 20. Geetha N, Anbarasi MS (2015) Ontology in cloud computing: A survey. International Journal
731 of Applied Engineering Research 10(23):43373-43377. doi: 10.1007/s12243-014-0450-7
- 732 21. Hu Y, Deng B, Peng F, et al (2016) A survey on evaluating elasticity of cloud computing
733 platform. In: World Automation Congress Proceedings. pp 1–4
- 734 22. Herbst NR, Kounev S, Weber A, Groenda H (2015) BUNGEE: An Elasticity Benchmark for
735 Self-Adaptive IaaS Cloud Environments. In: Proceedings - 10th International Symposium on
736 Software Engineering for Adaptive and Self-Managing Systems, SEAMS 2015. pp 46–56
- 737 23. Bauer A, Herbst N, Kounev S (2017) Design and Evaluation of a Proactive, Application-
738 Aware Auto-Scaler. In: Proceedings of the 8th ACM/SPEC on International Conference on
739 Performance Engineering - ICPE '17. ACM, New York, NY, USA, pp 425–428
- 740 24. Beltran M (2016) Defining an Elasticity Metric for Cloud Computing Environments. In:
741 Proceedings of the 9th EAI International Conference on Performance Evaluation
742 Methodologies and Tools. ICST (Institute for Computer Sciences, Social-Informatics and
743 Telecommunications Engineering), ICST, Brussels, Belgium, Belgium, pp 172–179
- 744 25. Kuhlenkamp J, Klems M, Röss O (2014) Benchmarking scalability and elasticity of distributed
745 database systems. Proc VLDB Endow 7:1219–1230 . doi: 10.14778/2732977.2732995
- 746 26. Ilyushkin A, Ali-Eldin A, Herbst N, et al (2017) An Experimental Performance Evaluation of
747 Autoscaling Policies for Complex Workflows. In: Proceedings of the 8th ACM/SPEC on
748 International Conference on Performance Engineering - ICPE '17. ACM, New York, NY,
749 USA, pp 75–86
- 750 27. Hasan Jamal M, Qadeer A, Mahmood W, et al (2009) Virtual machine scalability on multi-
751 core processors based servers for cloud computing workloads. In: Proceedings - 2009 IEEE
752 International Conference on Networking, Architecture, and Storage, NAS 2009. pp 90–97

- 753 28. Lehrig S, Sanders R, Brataas G, et al (2018) CloudStore — towards scalability, elasticity, and
754 efficiency benchmarking and analysis in Cloud computing. *Futur Gener Comput Syst* 78:115–
755 126 . doi: 10.1016/j.future.2017.04.018
- 756 29. Brataas G, Herbst N, Ivansek S, Polutnik J (2017) Scalability Analysis of Cloud Software
757 Services. In: *Proceedings - 2017 IEEE International Conference on Autonomic Computing,*
758 *ICAC 2017.* pp 285–292
- 759 30. Woodside M (2001) Scalability Metrics and Analysis of Mobile Agent Systems. In: Wagner T,
760 Rana OF (eds) *Infrastructure for Agents, Multi-Agent Systems, and Scalable Multi-Agent*
761 *Systems.* Springer Berlin Heidelberg, Berlin, Heidelberg, pp 234–245
- 762 31. Vasar M, Srirama SN, Dumas M (2012) Framework for monitoring and testing web
763 application scalability on the cloud. In: *Proceedings of the WICSA/ECSA 2012 Companion*
764 *Volume on - WICSA/ECSA '12.* p 53
- 765 32. Autili M, Di Ruscio D, Paola I, et al (2011) CHOReOS Dynamic Development Model
766 Definition (D2. 1)
- 767 33. Xiao Z, Chen Q, Luo H (2014) Automatic scaling of internet applications for cloud computing
768 services. *IEEE Trans Comput* 63:1111–1123
- 769 34. Amazon EC2 (2019) What Is Amazon EC2 Auto Scaling?
770 [https://docs.aws.amazon.com/autoscaling/ec2/userguide/what-is-amazon-ec2-auto-](https://docs.aws.amazon.com/autoscaling/ec2/userguide/what-is-amazon-ec2-auto-scaling.html)
771 [scaling.html](https://docs.aws.amazon.com/autoscaling/ec2/userguide/what-is-amazon-ec2-auto-scaling.html). Accessed 23 Jan 2019
- 772 35. OrangeHRM OrangeHRM REST APIS. <https://api.orangehrm.com/?url=/apidoc/index.html>.
773 Accessed 14 Feb 2019
- 774 36. Microsoft Azure (2017) Caching. [https://docs.microsoft.com/en-us/azure/architecture/best-](https://docs.microsoft.com/en-us/azure/architecture/best-practices/caching)
775 [practices/caching](https://docs.microsoft.com/en-us/azure/architecture/best-practices/caching). Accessed 15 Mar 2019
- 776 37. JMeter (2019) JMeter HTTP Request. [https://jmeter.apache.org/usermanual/](https://jmeter.apache.org/usermanual/component_reference.html#HTTP_Request)
777 [component_reference.html#HTTP_Request](https://jmeter.apache.org/usermanual/component_reference.html#HTTP_Request). Accessed 1 Apr 2019
- 778 38. Wang Q, Kanemasa Y, Li J, et al (2012) Response time reliability in cloud environments: An
779 empirical study of n-tier applications at high resource utilization. In: *Proceedings of the IEEE*

- 780 Symposium on Reliable Distributed Systems. pp 378–383
- 781 39. Butler B (2016) Who's got the best cloud latency?
782 [https://www.networkworld.com/article/3095022/cloud-computing/who-s-got-the-best-cloud-
784 latency.html](https://www.networkworld.com/article/3095022/cloud-computing/who-s-got-the-best-cloud-
783 latency.html),. Accessed 19 Mar 2018
- 784 40. Jayathilaka H, Krintz C, Wolski R (2017) Performance Monitoring and Root Cause Analysis
785 for Cloud-hosted Web Applications. In: Proceedings of the 26th International Conference on
786 World Wide Web - WWW '17. International World Wide Web Conferences Steering
787 Committee, Republic and Canton of Geneva, Switzerland, pp 469–478
- 788