
Make the most out of available cloud offerings with CloudPerfect Tools

White Paper

1 Introduction

The cloud market includes a vast number of cloud providers, ranging from the big players, such as Amazon Web Services, Google Cloud or Microsoft Azure, to smaller cloud providers or combination of private server clusters, along with a large variety of services and offerings.

Companies are attracted by cloud services because of their multidimensional value, including agility, scalability and cost benefits. Most cloud services fall into three main categories: 1) infrastructure as a service (IaaS) - offering virtualized compute resources ranging from processing power and storage space to complete servers, accessible over telecommunication networks; 2) platform as a service (PaaS) – supplying an on-demand environment for developing, testing, delivering and managing software applications; 3) software as a service (SaaS) – delivering software applications accessible over telecommunication networks, on demand.

Cloud infrastructure as a service or IaaS¹ is considered the fastest-growing segment of the cloud market as confirmed by a great number of studies e.g. according to the “RightScale 2018 state of the cloud report” over the 997 technical professionals interviewed, 96% are using cloud (public and/or private) and 81% are using multi cloud strategies. According to the results of a survey commissioned by Forbes, more than half of the enterprises have already applied cloud-first² strategies (52%) and many more (77%) are planning to become cloud-first in the near future (within two years)³.

Despite the rising popularity of IaaS, cloud IaaS adopters still face a number of challenges. In this white paper we focus on those cloud adopters that host or run their applications on cloud IaaS e.g. owners of on demand applications offering them as services to third parties (SaaS) or owners of applications of high computational demand.

In the following sections we follow the cloud adopters, as defined above, in their journey towards the use of cloud IaaS: from the assessment and comparison of available cloud offerings (Section 2), to the deployment of an application especially in a multiple-cloud environment (Section 3) and finally to the monitoring of the quality of service (QoS) that is assured through Service Level Agreements⁴ (SLAs) signed between the cloud adopter and the provider (Section 4).

The white paper looks into some of the challenges that adopters of cloud IaaS might face during their transition towards cloud and presents a set of solutions being developed in the context of the EU Horizon 2020 CloudPerfect⁵ project (<http://cloudperfect.eu/>) that can support cloud adopters throughout their cloud adoption process. It is noted that CloudPerfect delivers also tools for cloud providers, but this is out of scope of this white paper.

¹ Forecast Analysis: Public Cloud Services, Worldwide, 4Q17 Update, Gartner 2018

² New apps and services hosted in the Cloud as opposed to on infrastructure that the enterprise owns and manages.

³ Allan Leinwand, What The Rise Of Cloud Computing Means For IT Pros, Forbes Feb. 2017, <https://www.forbes.com/sites/forbestechcouncil/2017/02/28/what-the-rise-of-cloud-computing-means-for-it-pros/#6a9414be5518>

⁴ A Service Level Agreement (SLA) is the contract signed between the Cloud provider and the Cloud Adopter serving as a mutual agreement upon outline of the services being offered by the provider and that the adopter is entitled to.

⁵ The CloudPerfect project is co-funded by the European Union (EU) Horizon 2020 program under Grant number 732258.

2 Selecting the right cloud

The selection of the most appropriate cloud provider can be quite challenging. Cloud providers' offerings differ in many aspects and cloud adopters need to evaluate several factors, i.e., from compliance with regulatory requirements to certifications, from security options to customer support. To make things even more complicated, potential cloud adopters need to comprehend the technology stack (size, performance and resource type) offered by each cloud and to find their way through different pricing models. As a result, it is not unusual that many cloud adopters struggle and often do not make the best choice; thus not reaching the benefits envisaged in the use of cloud.

Several cloud comparison and recommendation tools come to the rescue of 'frustrated' cloud adopters supporting them in understanding which cloud provider better fits their business characteristics. Existing tools usually list various available providers, presenting their resource/services offerings and costs or include reputation systems, where users are encouraged to rate particular features of the services.

A crucial aspect, usually not taken under consideration by existing cloud comparison tools, is to understand how the application will perform on a specific cloud. Performance of an application is heavily affected by several factors like the actual hardware equipment of the cloud infrastructure, its co-location with other applications, the configuration of the Virtual Machines⁶ (VMs). Unfortunately, most of this information is not available to the cloud adopter. For example, despite the fact that information on the amount of resources and on the size of VMs is usually available, information on the respective physical hardware or other VMs running on the same nodes is usually unknown.

To simulate how an application would behave on different clouds, cloud adopters use benchmarking tools specialized in generating application-specific workloads and in running performance tests on different clouds. However, to have realistic results the benchmark tests need to be executed repeatedly and their results need to be compared. When taking into account the multitude of cloud providers, in conjunction with the different resource types and sizes that are offered and that need to be analysed, it becomes obvious that the simulation of an application behaviour is a time consuming and not trivial task.

CLOUDPERFECT APPROACH

CloudPerfect proposes a cloud comparison and recommendation tool, namely the QoE Portal, that offers support to cloud adopters by ranking the various cloud offerings taking into account the workloads that are going to be moved into the cloud.

Through the CloudPerfect QoE Portal the adopter can choose the workload category among a set of existing ones and CloudPerfect runs the performance assessment, for the selected workload categories, for a number of clouds (provided by different cloud providers). This is done by executing benchmarking workloads (generated by existing specialized tools) repeatedly on these accessible clouds. Based on average performance metrics, as well as their deviation, the QoE Portal provides an estimation of the application performance (on average conditions) on a per cloud (provider) basis. Once configured, the entire process is completely automated and repeatable. Furthermore, thanks

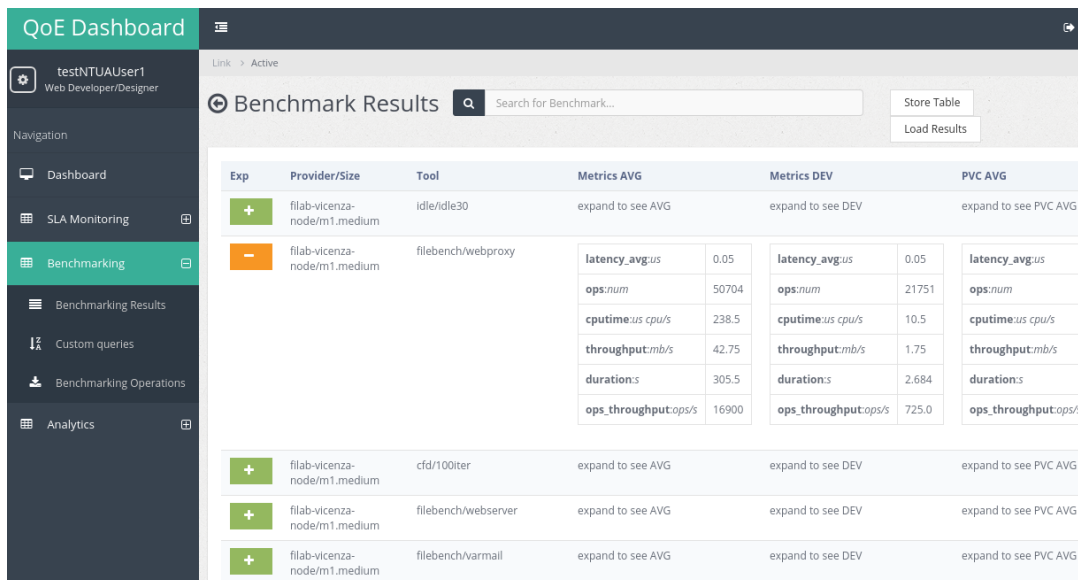
⁶ A Virtual Machine (VM) is an emulation of a computer system which provides the same functionalities as a physical computer.

to a plug-in architecture and the usage of abstraction libraries, CloudPerfect solution is extensible by design with the possibility to add support for new workloads and clouds.

Identifying the workload of an application deployment is a very challenging task. Especially for cloud adopters this is tightly related to requesting the suitable amount of cloud resources, avoiding overdimensioning and minimising cost. To support cloud adopters in this task, CloudPerfect is experimenting a tool providing insight on the behaviour of applications (resource usage i.e. cpu, memory, network) under a specific range of input workloads – that correspond to the specific deployment of the application - at runtime. For a specific application deployment the application is, then, classified into a set of pre-defined generic workload categories that behave in a similar way during runtime.

The performance assessment provided by CloudPerfect does not rely only on average scores, it also uses more sophisticated metrics. Through the QoE portal, in fact, the cloud adopter can check the percentage of the performance fluctuation over time (i.e. stability level of the observed performance) according to the benchmarking scores obtained⁷.

Figure 1: CloudPerfect QoE Portal



The screenshot shows the 'QoE Dashboard' for a user named 'testNTUAUser1'. The main section is titled 'Benchmark Results' and contains a table with columns for 'Exp', 'Provider/Size', 'Tool', 'Metrics AVG', 'Metrics DEV', and 'PVC AVG'. The table lists several experiments, including 'idle/idle30', 'filebench/webproxy', 'cfd/100iter', 'filebench/webserver', and 'filebench/varmail'. The 'filebench/webproxy' experiment is expanded to show detailed metrics for 'latency_avg:us', 'ops:num', 'cputime:us cpu/s', 'throughput:mb/s', 'duration:s', and 'ops_throughput:ops/s' for both 'DEV' and 'PVC AVG' scenarios.

Exp	Provider/Size	Tool	Metrics AVG	Metrics DEV	PVC AVG
+	filab-vicenza-node/m1.medium	idle/idle30	expand to see AVG	expand to see DEV	expand to see PVC AVG
-	filab-vicenza-node/m1.medium	filebench/webproxy	latency_avg:us ops:num cputime:us cpu/s throughput:mb/s duration:s ops_throughput:ops/s	latency_avg:us ops:num cputime:us cpu/s throughput:mb/s duration:s ops_throughput:ops/s	latency_avg:us ops:num cputime:us cpu/s throughput:mb/s duration:s ops_throughput:ops/s
+	filab-vicenza-node/m1.medium	cfd/100iter	expand to see AVG	expand to see DEV	expand to see PVC AVG
+	filab-vicenza-node/m1.medium	filebench/webserver	expand to see AVG	expand to see DEV	expand to see PVC AVG
+	filab-vicenza-node/m1.medium	filebench/varmail	expand to see AVG	expand to see DEV	expand to see PVC AVG

The cloud adopter can exploit the data available in the QoE Portal to get a recommendation for his/her application and create rankings of clouds (cloud providers) based on the performance results for the specific application workload and the related pricing parameters, collected while executing the benchmark test. Thanks to a customizable ratio between average performance and costs, the cloud adopter can request a recommendation, based on a trade-off between the best performance expected for a specific workload type and the cheapest solution available.

Using CloudPerfect QoE Portal, the time spent by cloud adopters to select the cloud provider that better fits their needs is considerably reduced. By accessing the QoE Portal the adopter can run automatically the benchmarking tests configured for the chosen application category and on the selected clouds. Results can be viewed in a few minutes instead of spending multiple hours or in

⁷ The metric used in CloudPerfect is the Performance of Virtual Cores (PVC). It was developed in the context of <http://slalom-project.eu/> and the Standard Performance Evaluation Corporation (SPEC) Cloud Working Group. The metric abstracts the measurement from a simple absolute value (of any given metric) and includes aspects deviation in a series of measurements.

more complex cases even a couple of days to run the tests with the various benchmarking tools. The QoE portal can also be configured to periodically perform all the tests for all the providers in order to have updated information but also to assess the stability/variability over time of the performance results.

Last but not least, CloudPerfect enables the ranking of cloud providers based on SLAs. This is done on the basis of specific measurable QoS parameters (to be guaranteed through SLAs); following the ISO 19086-2 standard⁸. The ranking basis is adjustable to the service and the performance requirements of each cloud adopter. In other words, through the QoE portal, cloud adopters can choose the parameters which are considered more relevant for their application or business (e.g. minimum downtime limits) and obtain cloud ranking results according to their needs.


⁸ ISO 19086-2 is the family of standards for Cloud computing SLAs, used to create a strict schema for strictly describing the SLA of a public provider, in order to be evaluated and become comparable.

3 Deploying on the cloud

Deploying an application on a cloud is not simple and hardly do any two providers offer the same interfaces or configuration details. In all cases, the cloud adopter cannot avoid to read the deployment manual in detail, provide the required information and generate an according script before actually deploying the service(s).

Today more and more enterprises move to multi-cloud strategies and sometimes want to make use also of their own resources to contribute to the service capabilities, protect vital data better etc.; in other words they try to get the best of all worlds (cf. Figure 2). Using different clouds (possibly also by different cloud providers) supports the cloud adopter in meeting specific workload or cost requirements and mitigates the overall concern on cloud availability. To this end, once a cloud adopter decides to opt for a multi-cloud strategy to better fit his business or applications needs, deployment gets even more complicated, because the aforementioned process has to be repeated for every separate cloud provider.

Figure 2: Cloud Strategy for big Enterprises, according to RightScale 2018 report



CLOUDPERFECT APPROACH

CloudPerfect offers a specific deployment tool called CLOUDIATOR which allows cloud adopters to deploy the same type of image/service multiple times in completely different infrastructure environments – including own ones.

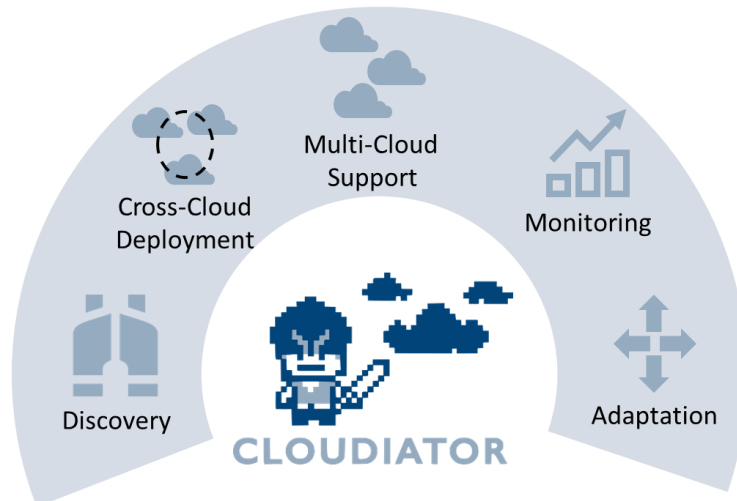
To make this happen, CLOUDIATOR relies on 3 types of information that govern every deployment:

- (1) the Cloud Model which specifies the destination provider (simple selection from a list) which allows CLOUDIATOR to decide how to execute the deployment. In order to perform the deployment, the necessary credentials also need to be provided at initial deployment. By assigning a name to that cloud instance later reference to it is facilitated.
- (2) the Application Model that describes the specific application and allows the user to perform powerful application behaviour rules and link multiple instances together by dedicated communication links. An application model can consist of multiple, reusable application components that altogether define the full application, so that e.g. different database components can be easily specified and reused in different applications, such as for a Wiki or a Customer Relationship Management (CRM) application.
- (3) the Resource Demands, finally, which specifies on which resources the application should be deployed. CLOUDIATOR allows its user to query the resources offered by a cloud provider so

as to select the best options from the respective list, but it is similarly possible to just (re)use an entry from the database.

With CLOUDIATOR the same cloud setup can equally be used for different applications and vice versa (the same application model for different target cloud providers).

Figure 3: CLOUDIATOR main features



The CLOUDIATOR toolkit can reduce the administrator's effort for multi-cloud deployment considerably; respectively also for recurring deployment. Complex deployments can take expert administrators multiple hours and if the according effort has to be done multiple times, this sums up to multiple days and weeks over the usage period of one specific service or application. CLOUDIATOR requires the one-time-investment into the application specification which can take in complex cases up to a day, but from that time on, any further deployment on any provider is just a question of seconds, also removing the need for an administrator's time to perform the deployment process.

In addition to the deployment support, CLOUDIATOR also offers integrated capabilities to:

- Monitor the application and its performance at run-time
- Adapt (migrate, scale) instances according to pre-defined rules
- Deploy application components across multiple clouds.

4 Maintaining Record of the Quality of Service

Following the deployment of the application on a specific cloud, the focus of the cloud adopter moves to monitoring the service performance guaranteed by the Service Level Agreement (SLA) between a cloud provider and the cloud adopter.

SLA assurance is still a major challenge in cloud computing, as there is no common framework – among cloud providers - in the market. Cloud providers employ dedicated tools for monitoring the offered service levels to their customers. Depending on the cloud deployment and the underlying SLAs terms they may also provide cloud adopters with access to online portals or third-party monitoring tools to check whether their SLAs are met. In other cases trusted third parties are involved to perform measurements and resolve possible conflicts.

Availability is the fundamental and most commonly used Key Performance Objective (KPO) in cloud services SLAs; considering that even the smallest outage can cost the cloud adopter money or even impact their reputation. Monitoring the availability from the cloud adopter's side, a usual approach employed by existing frameworks, is to periodically poll cloud infrastructure and record the service uptime. However, this usually corresponds to the general cloud provider infrastructure availability and not to each cloud adopter service. Furthermore, cloud providers measure availability in different ways⁹ (e.g. when occasional downtimes exceed 1 minute, when one or all VMs are not reachable by their end-users), making it even more challenging to assess the actual level of availability on a common basis.

CLOUDPERFECT APPROACH

Going beyond the generic infrastructure availability monitoring –provided by most existing tools- CloudPerfect offers an SLA monitoring and evaluation tool enabling cloud adopters to monitor their own cloud services availability and performance levels, and check them against the thresholds agreed in their SLAs with cloud providers. CloudPerfect calculates the service availability and performance levels of all the VMs of a cloud adopter, by polling the service directly through the cloud adopter's account. To this end, the tool takes into consideration the definition of service availability and performance levels that each cloud provider gives in their SLA. Besides the generic service uptime/downtime which corresponds to availability, the following SLA performance guarantees' monitoring is supported:

- Performance guarantees in terms of Response time: monitoring of the response time of the application and generating notifications when this exceeds a certain limit (e.g. applicable in the MS Azure Blob Storage Service/SLA).
- Performance guarantees in terms of Exceptions received: monitoring of the number or rate of failures and exceptions generating notifications when this exceeds a certain limit; corresponds to a functional evaluation of the service (e.g. applicable in the Google App Engine Service/SLA).

⁹ White paper: Comparing Public Cloud Service Level Agreements - Dimension Data
Availability in the cloud: State of the art, Journal of Network and Computer Applications Volume 60, January 2016, Pages 54-67.

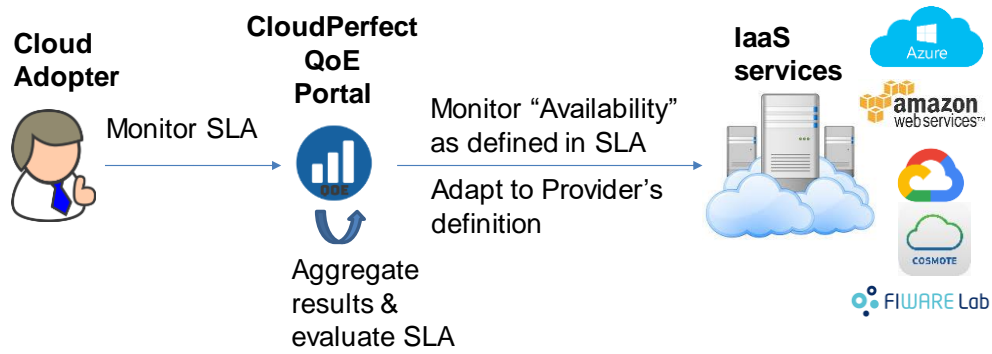


Figure 4: QoS monitoring facility offered by the CloudPerfect QoE portal

Thanks to its abstract software library with specific –per cloud provider/deployment- driver implementations, different interpretations of general service availability levels can be defined and monitored, while a number of metrics and parameters can be included for more complex SLAs' performance level guarantees.

The cloud adopter can access the CloudPerfect SLA monitoring functionalities and view the monitoring results through the QoE Portal, as well as access the detailed SLA monitoring logs to support compensation claims.

CloudPerfect offers the adopters the possibility to monitor the Quality of Service they experience, rather than relying on cloud provider infrastructure services. CloudPerfect also offers an SLA monitoring solution for multi-cloud deployment through a single dashboard.

5 Who can benefit from CloudPerfect tools?

Although CloudPerfect tools can be utilized by all cloud adopters, two groups of cloud adopters can benefit most from them:

- *Owners of on demand applications offered to third parties as SaaS on the cloud* – these are companies that need to carefully select the cloud providers taking into consideration costs, cloud performance, application resource requirements and at the same time the quality of the service offered. In fact they need to guarantee certain QoS to their own customers, to make their offering competitive on the market.
- *Owners of applications with high computational demand* – these are companies that usually make use of High Performance Computing (HPC) to run their applications (e.g. perform highly sophisticated simulations) but in the last few years are increasingly using cloud computing as a valid option to HPC. The pay-per-use model suits the often fluctuating need of computation. Also in this case the adopter needs to choose the provider carefully keeping the right balance between costs and performance and often needs to deploy the same application on different cloud providers; thus increasing the deployment costs considerably.

6 Summary and Roadmap

CloudPerfect tools (the QoE Portal and CLOUDIATOR) offer personalized support for the selection of the most appropriate cloud IaaS provider, they support the Service Level Agreement (SLA) monitoring for the service performance guarantee and facilitate the deployment of the applications in cloud (and multi-cloud) environments. More specifically:

The **QoE Portal** supports the cloud adopters in two stages of their cloud journey:

- 1) *In selecting the right cloud* – the QoE Portal supports the cloud adopter in selecting the most appropriate cloud offering (cloud provider/s) according to their application or business needs. It provides the adopters with the possibility to create rankings of clouds (cloud providers) based on the performance results for the specific application workload and the related pricing parameters. Rankings can also be made on the basis of specific measurable QoS parameters that are guaranteed by the cloud providers through their SLAs. The QoE portal provides homogeneous results supporting the adopter in comparing the numerous and different cloud offerings available on the market while reducing the time and the hassle spent on the provider selection phase.
- 2) *In Maintaining Record of the Quality of Service* – through the QoE Portal the cloud adopter can access SLA monitoring functionalities, view monitoring results and access detailed SLA monitoring logs (which can be used to support compensation claims). The monitoring results made available through the QoE Portal go beyond the generic infrastructure availability monitoring, since cloud adopters can monitor their own cloud services availability and performance levels, and check them against the thresholds agreed in their SLAs with cloud providers.

Development Status and Roadmap

The QoE portal is in its final testing and validation phases being performed with the support of companies working in the above targeted domains (in particular: ICON, PHITEC Ingegneria, VALEO Cognito and COSMOTE). CloudPerfect QoE will be offered as a service at the end of the CloudPerfect Project (30th November 2018).

The tools working in the backend of the QoE portal can also be used as single tools to perform specific task (e.g. automated benchmark test or SLA monitoring) or can be integrated in other solutions. They are all released as open source with Apache2.0 license and can be found at the following link: <https://github.com/cloudperfect-project/CloudPerfect/>

CLOUDIATOR supports the cloud adopters in the phase of deploying on the cloud. It offers a common interface and model for describing and configuring applications independently from their actual runtime environment and the infrastructure where they are deployed. This in fact reduces deployment time, since now applications can be seamlessly moved across different cloud provider infrastructures without any additional effort. Normally, this step requires further expertise for the allocation and configuration of cloud provider specific resources. Now this expertise is shifted to the tool itself which allows for cost reduction for provider specific adaptation and expertise. The tool implements a generic and provider independent way of application implementation and management which eases the use of clouds in general and does not bind tenants to specific provider

offers and their platforms (vendor lock-in). It increases flexibility, since now more apps can be handled in parallel on various provider infrastructures and increases the competition between providers in terms of favourable offers and conditions for their resources.

Development Status and Roadmap

CLOUDIATOR is also in its final testing and validation phase. The CLOUDIATOR team aims at creating a spin-off to develop the specification upon request of the adopter. This will specifically focus on small to medium enterprises, for whom the effort for deployment is too high, but experience and time for writing and maintaining own configuration scripts is lacking.

CLOUDIATOR is released as Open Source software with Apache 2.0 license and is available on github at the following link: <https://github.com/cloudiator>

More information on CLOUDIATOR can be found at the following link: <http://cloudiator.org/>