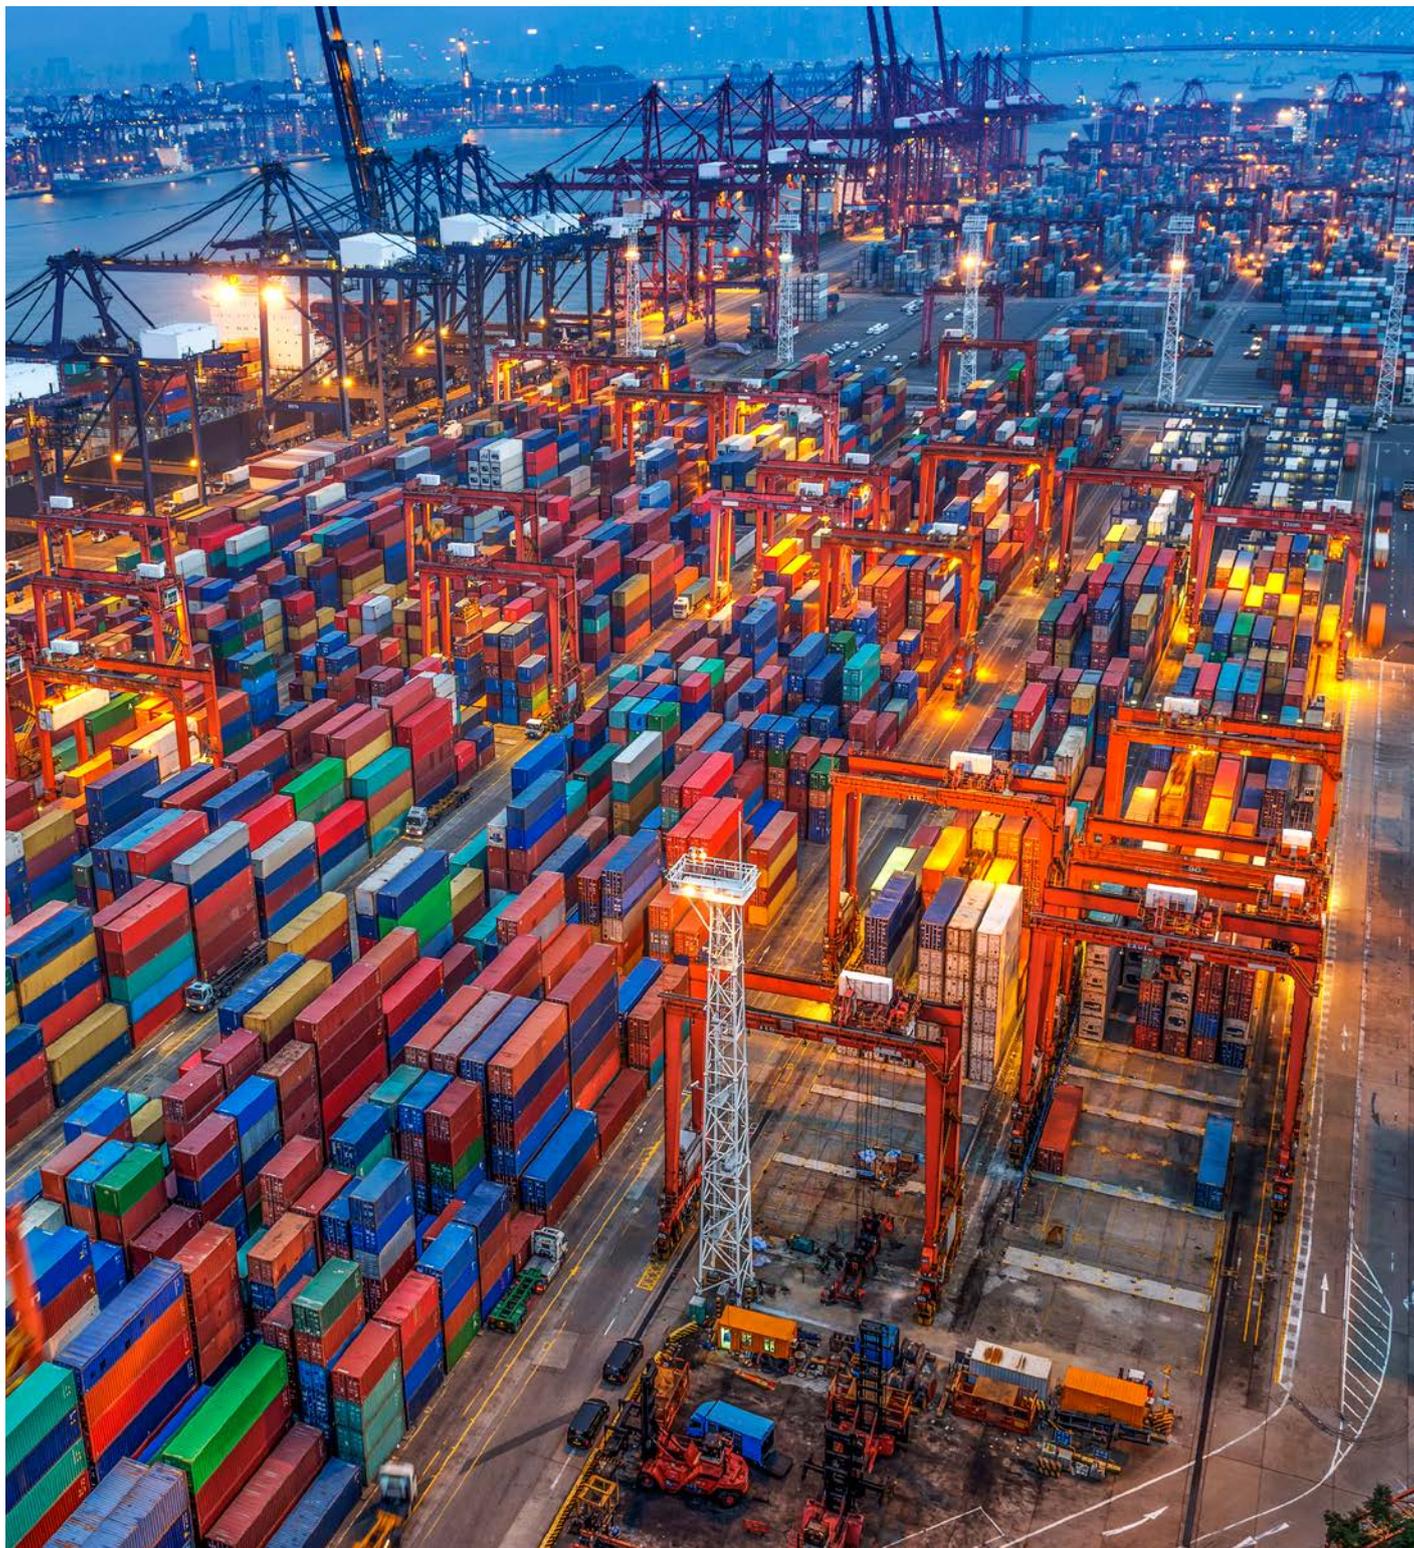


the cloud report



03-2019

Container Technology

Additional:

Kubernetes - GC Build - Death of the Developer - Deploying Tests

T · · Systems ·

**GET IT NOW.
BUT GUIDED.**

MANAGED CLOUD SERVICES

Take control with our intelligent Managed Cloud Services – secure and reliable.





Container Technology

Cloud technology not only changes the way data is stored. Collaboration models, tools, availability, business models, ... are changing as well. DevOps is almost an old hat. Everywhere it is said that processes have to change and the individual has to adapt. But people are creatures of habit.

Every developer set up their workplace, computers, monitors, programs they need and want to use, ... And then suddenly you have to develop for cloud? In the cloud? Directly together with others? Who then perhaps also want to make operations?

Meeting these challenges can be scary, but it's definitely worth the effort to learn something new, to grow with it, find new tools, cool tricks. And still from home, the computer remains the same, the monitors, the silence in the private room or the chosen noise. Even the tools can stay the same or be set up to feel familiar. But results come faster, feedback comes more purposefully and brings immediate progress. And so, the sense of achievement comes faster.

The new doesn't have to be frightening, because it is above all fun. And how this can look in practice is shown here in this issue. As in every issue, we have again tested cloud vendors, such as AWS, Google Cloud Platform, Azure, IBM Cloud and OTC. We test cases, time, costs in various areas like compute, network, security, as a Service provisions, and many more. Find more details at the end of this issue. In addition, Kubernetes is already 5 years old, that has to be celebrated! This will also be celebrated at the Container Days. We are media partners of the Container Days and with you on site! Have fun reading, listening, talking and good days in Hamburg.

HAVE A GOOD TIME AT THE CONTAINER DAYS!

Best wishes,
Friederike
Editor in Chief



the cloud report
IMPRINT

Publisher	Cloudical GmbH, Edisonstr. 63, 12459 Berlin
Managing directors	Michael Dombek and Karsten Samaschke
Publishing director / Editor in chief	Friederike Zelke
Editorial	Julia Hahn, Emelie Gustafsson
Online editorial	Stefan Klose
Sales	Emelie Gustafsson
Artdirection and Layout	Anna Bakalovic
Production	Andreas Merkert
Editorial office contact	presse@cloudical.io
Sales contact	marketing@cloudical.io / sales@cloudical.io

Copyright © Cloudical GmbH

the cloud report	published by Cloudical GmbH Edisonstr. 63, 12459 Berlin
Managing directors	Michael Dombek and Karsten Samaschke Mail: info@cloudical.io

the-report.cloud

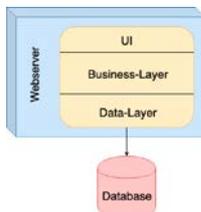
ISSN 2626-1200

The Cloud Report is published quarterly at the beginning of January, April, July and October. The Cloud Report is available in two versions: the online edition which can be accessed online or via downloaded from the homepage, and the printed edition which can be subscribed at a price of 20 Euros a year via the personalized customer portal in which a personal account is set up. When you register, please enter the information to specify the execution in which you want to obtain the report. The subscription can be cancelled at any time via the personal access of the subscriber, by mail at the latest two weeks before the publication of the new edition via the e-mail address: sales@cloudical.io. We collect relevant personal customer data for the subscription. Individual issues can also be purchased without a subscription at a price of: 5 Euro, in this case too, relevant personal data will be used to fulfil the purchase contract. Further information can be found at: <http://the-report.cloud/privacy-policy>

the cloud report
TABLE OF CONTENT

EDITORIAL

Container Technology 1



FOCUS

- Building the Clouds 4
- Dynamic Systems and Their Testing 8
- Kubernetes, the essential ingredient for every enterprise 12
- The Death of the Developer 14



INTERVIEW

Interview with Nils Brauckman,
CEO, SUSE 20



PROJECT LEADERSHIP

Digital Transformation in
Medical Technology 24

CONFERENCES

- Collaboration without Boundaries 28
- Happy 5th Birthday, Kubernetes! 31
- SUSECON 2019 34
- On the Go - GoDays 2019 in Berlin 37



TESTS

- We are testing clouds 41
- And the winners are ... 42
- Test Results 44

Building the Clouds

Google Cloud has a lot of impressive things to discover – one of the interesting and least known aspects is its integrated build system, Cloud Build. Let's set it up for hosting a GIT-repository, for building a Microservice and for storing the generated artifacts in a Docker registry inside Google Cloud.

First of all, there are some prerequisites to fulfill before using Google Cloud Build:

- › Google Cloud account with payment functionality enabled needs to be present
- › A Project has to be created inside Google Cloud – our sample project is called *cloud-report-project*
- › The Google Cloud SDK is to be installed on your local machine

If all of these steps are taken, we should have a Microservice we want to build. In this article, we use a Spring Boot based Microservice written in Java (you can find the source code at [1]), but you can of course use a language of your liking – Google Cloud Build actually does not care about a programming language, it refers to builders being supported on the platform.

How does Google Cloud Build work?

Before building anything, we should understand the basic idea behind Google Cloud Build, since it is different to the approaches AWS or Azure have. Essentially, Google Cloud Build does exactly, what the name supposes: It imports sources, builds them, creates container images, stores and – if requested to – deploys them inside of Google's own cloud environment.

At the center of it, there is the Google Cloud Build configuration file, a YAML-file being expected to be present inside the project. Builds are reflected as steps in that YAML-file, each step is represented by a builder being supported by the platform or the community.

Already existing builders are tools such as Maven, Gradle, NPM, Bazel, Gulp or wget, that are able to be executed

as some sort of command inside the environment. A list of builders supported by the platform is available at [2], but there are way more builders provided by the community – and you can even create and submit your very own builder.

Each build step is executed inside a container. These containers run inside a local Docker network called *cloudbuild*. This allows for communication between the different builders and build steps being referenced in the build process.

Generated artifacts of a build step can be containerized and stored in a container registry, such as Docker Hub or (preferably) a private container registry inside the Google Cloud environment.

The build process can be triggered manually using the *gcloud builds*-command on your local machine or by utilizing hooks in GIT-repositories.

Generally, the lifecycle of a Cloud Build build is like this:

- › We create our application
- › A Cloud Build config-file is added to our application
- › The build is submitted to the Cloud Build service
- › The build process is executed inside Cloud Build
- › Generated container images are pushed to a container registry

Google Cloud Build preparations

Before building any software project, we need to activate Google Cloud Build and create a GIT-repository for hosting the sources.

Hosting a GIT-repo

First, point your browser to <https://console.cloud.google.com> and log in there with your account. Then use the Hamburger menu and select *Cloud Source Repositories*

```

Microservice --- -bash --- 138x41
Step #1: Removing intermediate container 0959d9431e50
Step #1: ----> d946dc4b6dad
Step #1: Step 3/7 : VOLUME /tmp
Step #1: ----> Running in 44c16fe1fd46
Step #1: Removing intermediate container 44c16fe1fd46
Step #1: ----> 96224eef62e2
Step #1: Step 4/7 : EXPOSE 80
Step #1: ----> Running in 3acc6aafb24
Step #1: Removing intermediate container 3acc6aafb24
Step #1: ----> 57397a4a689c
Step #1: Step 5/7 : ARG JAR_FILE=target/CGPMicroservice-0.0.1-SNAPSHOT.jar
Step #1: ----> Running in c1698e41bbe9
Step #1: Removing intermediate container c1698e41bbe9
Step #1: ----> 6dc922415e93
Step #1: Step 6/7 : ADD ${JAR_FILE} service.jar
Step #1: ----> f24e24a849a1
Step #1: Step 7/7 : ENTRYPOINT ["java","-Djava.security.egd=file:/dev/./urandom","-jar","/service.jar"]
Step #1: ----> Running in 8a34bdab0c1d
Step #1: Removing intermediate container 8a34bdab0c1d
Step #1: ----> 04d235ea340a
Step #1: Successfully built 04d235ea340a
Step #1: Successfully tagged gcr.io/cloud-report-project/cloud-report-sample:latest
Finished Step #1
PUSH
Pushing gcr.io/cloud-report-project/cloud-report-sample:latest
The push refers to repository [gcr.io/cloud-report-project/cloud-report-sample]
ff2ed6976da: Preparing
88a9dbfd3d5c: Preparing
f1b5933fe4b5: Preparing
f1b5933fe4b5: Layer already exists
ff2ed6976da: Pushed
88a9dbfd3d5c: Pushed
latest: digest: sha256:3cac50d46c36a9b8d92e794f1c55be6cbe1b3be8edc1d5d1bf28cb4cfff16147 size: 953
DONE
-----
ID                CREATE_TIME          DURATION  SOURCE                                STATUS
IMAGES
911b8cb6-91c5-4332-b85f-10f5736838fc  2019-06-17T05:12:09+00:00  1M31S    gs://cloud-report-project-cloudbuild/source/1560748327.09-86ec61s28179416e808f4655f7881741.tgz  gcr.io/cloud-report-project/cloud-report-sample (+1 more)  SUCCESS
Karstens-MacBook-Air:Microservice ksamaschke$

```

Figure 1: Output of a Cloud Build operation

or head over to <https://sources.cloud.google.com>. Create a new GIT-repository here by pressing the *Add repository*-button, name it *cloud-report-source*. This just takes some seconds.

Enable Cloud Build

Next, Cloud Build needs to be enabled. Head over to <https://console.cloud.google.com/cloud-build/builds> and click on the *Enable Cloud Build API*-button.

The software project

Obviously, in a software project, some build- and configuration files need to be present. Usually, these are located in the root folder of the project.

In our example (a Java-based Spring Boot Microservice), we have three configuration files (*pom.xml*, *Dockerfile* and *cloudbuild.yaml*) being present for building, containerization and Google Cloud Build execution.

The Maven build configuration

As we are building a Spring Boot based Microservice with Java, we need to provide a Maven POM-file for building the Microservice. This is a standard POM-file, so there is no need to look into it here.

The Dockerfile

Then, we need to provide a Docker configuration file, since the generated artifact from the Maven build step needs to be dockerized to run in a cloud-native environment. Let's have a look at it:

```

# Start with a base image containing Java runtime
FROM openjdk:13-jdk-alpine

# Add Maintainer Info
LABEL maintainer="karsten.samaschke@cloudical.io"

# Add a volume pointing to /tmp
VOLUME /tmp

# Make port 80 available to the world outside this container
EXPOSE 80

# The application's jar file
ARG JAR_FILE=target/CGPMicroservice-0.0.1-SNAPSHOT.jar

# Add the application's jar to the container
ADD ${JAR_FILE} service.jar

# Run the jar file
ENTRYPOINT ["java","-Djava.security.egd=file:/dev/./urandom","-jar","/service.jar"]

```

This Dockerfile is very straightforward: As base-image is the current version of the *openjdk*-Docker-image based on Alpine Linux referenced. The container exposes port 80. The generated JAR-file containing the Microservice is copied and renamed to *service.jar*. Finally, the command line entry-point for starting the Microservice inside this container is defined.

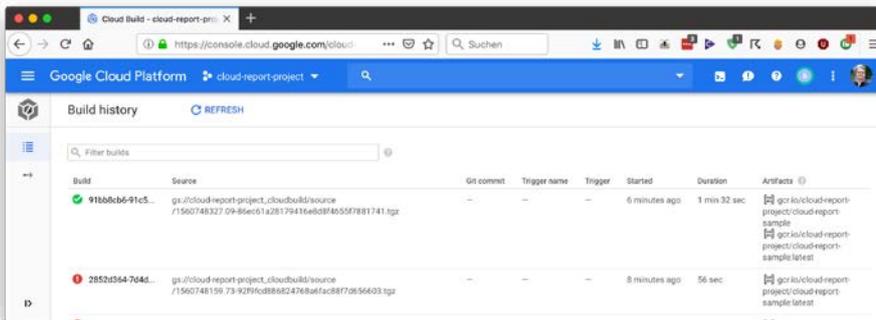


Figure 2: Build history in the Cloud Build console

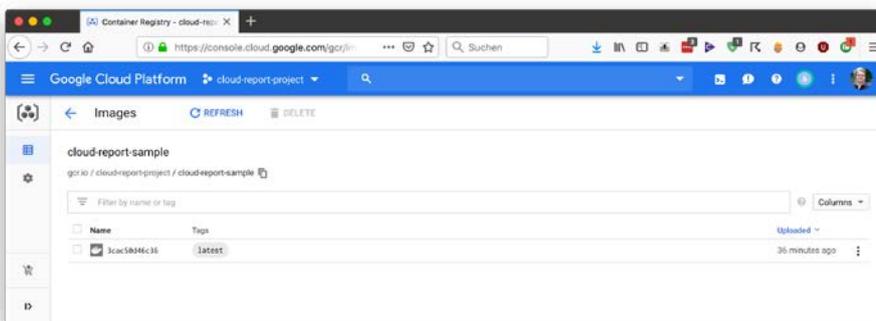


Figure 3: Generated image in the Container Registry

The Google Cloud Build configuration

Now, let's have a look at the Google Cloud Build configuration file *cloudbuild.yaml*. As stated above, each build (and deployment) step is noted down here. Additionally, the name of the image to be created from the build steps are depicted as well:

steps:

- name: 'gcr.io/cloud-builders/mvn'
 args: ['install']
- name: 'gcr.io/cloud-builders/docker'
 args: ['build', '--tag=gcr.io/\$PROJECT_ID/cloud-report-sample:latest', '.']

```
images: [
  'gcr.io/$PROJECT_ID/cloud-report-sample:latest']
```

When executing this file, the sources are built using Maven. The generated artifact is dockerized using the Dockerfile provided within the project. The generated docker image is made available in the local container registry associated with the project having the tag *latest*.

Note: The name *cloud-report-sample* references the internal name of the project. It can be selected to your liking, but it should be a simple and unique name within your Google Cloud project.

Execute the build

Now, that we have all components in place, we can execute the build process. Open your command-line, switch into

the directory of your sources and execute the following command (including the trailing dot):

```
gcloud builds submit -config cloudbuild.yaml.
```

Once you submitted this command, Google Cloud SDK will create an archive of your project, upload it to the Cloud Build infrastructure and execute all defined build steps there. While doing this, you are informed about the process as depicted in Figure 1.

In the Cloud Build-console at <https://console.cloud.google.com/cloud-build/builds>, you can now see the result of all build processes as depicted in Figure 2.

Head over to the Container Registry using the Hamburger menu to the right (or via <https://console.cloud.google.com/gcr>) and notice the generated Docker container as depicted in Figure 3.

Let's automate it!

Now, there is only very few steps left for creating an automated build process. Remember the GIT-repository created earlier? Just commit and push your code there and then head over to the Triggers section of Cloud Build configuration page in the Google Cloud Console at <https://console.cloud.google.com/cloud-build/triggers>.

Here we can create a new build trigger. Click on the button *Create Trigger* to start. Now select *Cloud Source Repository* (or any of the other available options, if applicable) and click *Continue*. Select the repository you created earlier on and click on *Continue*. Now define the trigger itself - name a specific branch (or leave it with the default option pointing to all branches, but I highly rec-

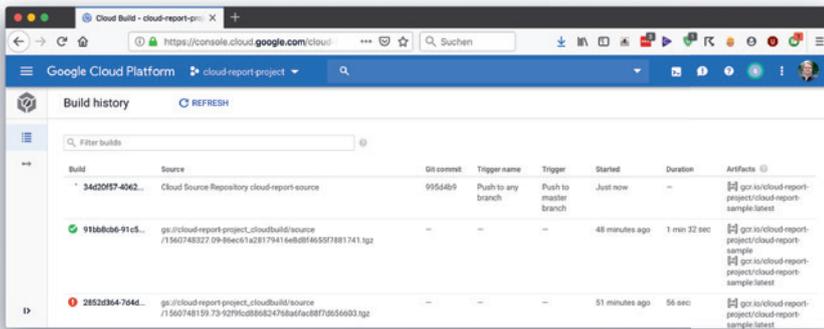


Figure 4: The current build is triggered automatically

ommend pointing it to specific branch only) and select the option *Cloud Build configuration file* in the Build configuration section. Finally, click on *Create Trigger*, and you are all set!

Finally, change something in your sources, commit and push them – and then head over to the Cloud Build console and see the magic happen (Figure 4).

Summary

Google Cloud Build is a very interesting and easy to use approach to build and build-automation. It does not require any local builds anymore, since every build step is executed in a cloud environment. We as developers gain a lot of power with very few effort – as of now, there is no

such easy to set up and use mechanism available on other cloud platforms.

And we did not even speak about publishing everything to Kubernetes, which we will cover in the next issue...

References

- ▶ [1] Source code of the sample project – <https://github.com/cloudical-io/google-cloud-build-sample>
- ▶ [2] List of supported builders – <https://cloud.google.com/cloud-build/docs/cloud-builders>

Karsten Samaschke



CONTAINER DAYS SPONSORING



Technological boost for the core business of SMBs

Resulting from the merge of 1&1 Internet and IaaS expert ProfitBricks, 1&1 IONOS is a leading cloud service provider for small and medium-sized businesses (SMBs), and the largest hosting company in Europe with a personalized service. It manages more than 8 million customer contracts and hosts over 12 million domains in its own data centers in the US and Europe. A best-in-class server and cloud portfolio empowers customers to continue running, and develop a successful business. 1&1 IONOS offers a cutting-edge cloud and dedicated server platform for those looking for high-performance, security and flexibility. This includes extras like pay-per-use billing, full API control, unlimited traffic and load balancing. The easy to use Cloud Panel helps customers to manage, configure and scale their infrastructure. The interface also allows them to configure their external firewall, use load balancing, and monitor their system. In addition to easy administration of their server and cloud infrastructure, the Cloud Panel also allows the setup, (auto)scale and update of Kubernetes clusters. A broad choice of pre-installable solutions from Helm to Fabric8



are also available so that the customers can be up and running in no time. By providing its customers with the latest technology, 1&1 IONOS gives them the foundations they need to be successful online and to dedicate themselves to their business.

- ▶ Learn more about 1&1 IONOS at www.ionos.com

Dynamic Systems and Their Testing

Before software is handed over, it is tested. Before you run it, you want to be sure it's doing what it's supposed to. But the more agile the process becomes and the more automated it becomes, the faster software is deployed, the less space there is in the process for testing. This article shows how testing can be integrated and planned into the agile process from the beginning, so that high-quality results can be delivered.

Software systems are usually constructed so that their function is spread over a set of coupled components. Then, that set of software components forms a logically dependant, but technically isolated, composite, which in its entirety forms the application. Such a modular construction of software systems has been an accepted industry standard for some years now and has been grown to new complexity, some might say extremes, with the advent of micro-service architectures.

Adding to the complexity of (multi-tiered) software, application software relies on a number of backend systems behind the application's "actual backend", i.e. auxiliary software (message brokers, backup, logging, monitoring) and infrastructure components (that stuff that the application actually runs on, viz. computers, networks, storage systems); and it is – by and large – undisputed that application software cannot run without it¹. For these, distribution of components is de-facto mandatory.

Distributed Systems

Such distributed systems are, while not exactly fragile, still complicated and leave copious space for errors.² To reduce the probability of such errors hitting production, software development processes stress the importance of quality control, often called testing. Before a given piece of software will reach production, it has been subjected to a multitude of unit, integration and acceptance tests – all at different abstraction levels. These tests are, because developers are very fast at producing errors, themselves software. And, as we all know, software can, as anything computerized, run faster, more often, at times when developers sleep, etc. So, in theory, the testing stays ahead of the developers at all times. In practice, it doesn't, and software development has coined a less aggressive sounding word for such errors, "bugs", to save themselves from the embarrassment.

¹ The usual suspects have been waiting to divert the discussion to λ all along, but neither kidnapping Alonso Church nor enlisting the help of μ , ν or members of other alphabets may bend the fact that irrespectively, that stuff still needs to run somewhere.

² Some call them beasts, and while this sometimes may be true, more often than not the beasts are long dead and their decomposing carcasses form a swamp which may kill the unsuspecting just as efficiently, albeit slower, as the beasts themselves could, when they still were alive.

Administrators, operators or whatever the current term may be, do not err when deploying or maintaining; accordingly, they do not need to resort to insectoid imagery and certainly do not need to automate away the testing.³

Classical, manually administered environments

Those that work and do not break with any new release – distinct themselves from software development and its methods of verification and testing, and the speed of producing errors anyway. To guarantee the correct setup and accordingly, the specified functionality of a computing environment, so-called “changes”, i.e. alterations to the systems, are “transported” stepwise over a number of increasingly more production-like pre- production “stages”, where any errors, should they occur, can be caught. Correctness is given when testers run and analyze a set of tests and they return without error – the tests, not the testers.

This practice has spilled over to automated, “neo-classical” environments, where, even if the mechanics of deployment have been passed to an automaton such as Puppet or Ansible, the proven practices of performing changes have stayed more or less the same.⁴

With that reasoning understood

According to that, unsurprisingly, testing and software testing alike is all but common practice for infrastructure systems and their guardians. Although the “underlying” infrastructure or auxiliary software and systems comprise non-trivial portions of application function, when verifying and testing an application, they are largely ignored.⁵

However, such an approach to operations is generally and increasingly unfit for a dynamic setting: When automaton rapidly perform scaling and replication or re-provision software systems to other computers autonomously, entirely without human intervention, the approach of provisioning and testing over pre-production stages cannot keep up with the speed, the agility of change itself. Because operating untested software is a Bad Thing™ – and prohibited in regulated environments such as finance anyway – automation should generally not be trusted with the operation of professional IT.⁶

Unfortunately, they re-invented MULTICS lately, called in air-mobile cavalry of supreme marketing, christened that “the cloud” and, alas, managers aspiring to get honorable mentions in Gartner reports or whatever gossip & glamour magazines the so-called working, so-called professionals consume, demand their IT organizations to use and to operate cloud services.⁷

Speed in dynamic settings

In this inherently dynamic setting, resources are provisioned rapidly and on demand and are decommissioned when no longer required, and because of the speed at which these operations need to be performed to be useful, they require a high degree of automation, all of which is code in some aspect or the other. Administrators are thus forced to look into the cloudy skies instead of keeping their focus firmly on the ground and now, they produce errors at the same dangerous rate developers do and are used to do.

To add insult to injury, they have lately introduced orchestration middleware such as Kubernetes to IT operations. By this cunning scheme, cloud wobbliness has taken over in-house operations. Why build roofs, when you like rain so much? In this context, administrators might, even with some remaining nausea, look at how developers have managed to live with change, which has led – with varying degrees of success and satisfaction⁸ – to the DevOps idea of actually introducing development methods into the operation of IT systems.⁹

Automated Testing

When automated or even autonomous systems, which are themselves defined by code, provision and deploy, then the deployed infrastructure itself is defined by code. Such code will then be developed like any other software component or system, and development methodology may – should – be applied. This includes testing, and, more specifically, such infrastructure code should then be submitted to the same automatic tests (software by itself) to demonstrate it fit to the specification.

3 At 120kg and just a flight of stairs away from a halting cardiac, administrators are a dignified folk and not so fast at producing anything, which includes errors, mind you, and have, by the way, understood the recursion problem with so-called software testing finding software errors and, accordingly, do not even try to accomplish the impossible.

4 I.e., find some machine in the development stage, perform the change, prod it to find out if it is strike though effect still alive, resurrect it when dead by fate of chance and gradually reduce the prod-resurrect iterations to zero till production.

5 We may now debate if the term “underlying” still fits, when the contribution to an applications’ function is non-trivial, which we won’t, because IT professionals are too expensive as it is, which is a problem in itself and won’t be helped by declaring the janitors and cleaners important parts of an enterprise. Issue closed, thank you very much.

6 Anything remotely resembling opposition is quickly killed off casually pointing out the algorithmic difficulties for any automaton to correctly select the proper type of ITIL change and actually get it approved in this eon. QED.

7 Oh my gosh! Blimey!

8 Und wenn du lange in einen Abgrund blickst, blickt der Abgrund auch in dich hinein. (And if thou gaze long into an abyss, the abyss will also gaze into thee.)

9 If the DevOps idea of increased collaboration between developers and operators has led to successful advances or even is on the right path requires large amounts of beer and even larger amounts of time, which in turn requires more beer, albeit without significantly increasing the chances of timely resolution ...

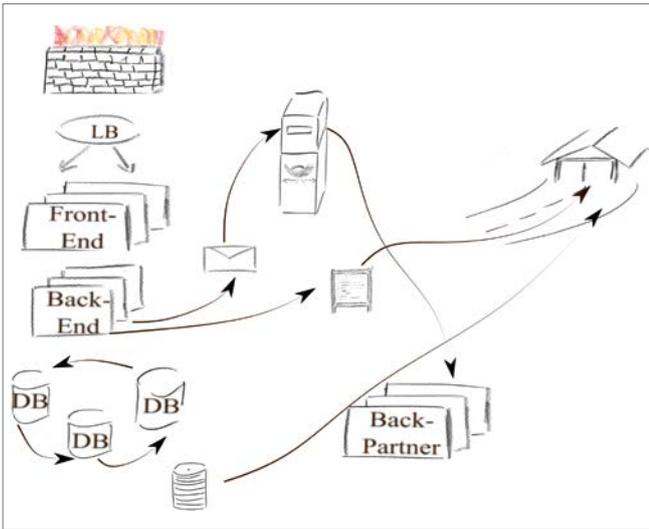


Figure 1: Distributed Systems

Reasoning thus, when operating rapidly changing infrastructure and auxiliary systems, either cloud or cloud-like in-house middleware, methods of development and operations converge, and automated testing should be introduced to the operators *modus operandi*.¹⁰ Having thus been convinced – or coerced – that dividing between infrastructure, auxiliary and application systems is rather artificial, it may be time to revisit infrastructure from an application’s point of view.

Persistent Storage

Most applications rely on storage systems, often databases, but, depending on the application domain, also large object stores or buckets of files. No-one actually would argue against persistently stored state being integral to the function of many applications. Cloud or no cloud, professional storage systems can reasonably and well be trusted in so far as they do not loose data, but their correct configuration cannot be blindly assumed.

The permissions governing the access to data could well be wrong.¹¹ Performance may rely on correctly spreading the load over a number of participating systems, which nobody really checks. Resilience relies on redundancy,¹² so distributed databases claiming to survive the loss of one or even many nodes should really be distributed instead of being strike through smeared over only a subset of ma-

chines with the rest of ‘em happily idling – contributing to the power bill, but not to the required failover properties of the system.

Often, applications write journals of logs.

Logs are often of a technical nature, sometimes even purely so, but sometimes, they already have or later gain a business relevance up to being part of a mediation system. Then, their retention and deletion may well fall under different laws governing trade and taxation, making their correct handling much more compulsory instead of errors being a mere nuisance.¹³

Conversely, under European data protection laws, certain data may well be subject to a maximum retention time and accordingly, must be periodically deleted. Failure to do so not only may but will result in severe and punitive penalties. Even without GDPR-effects, failure to timely delete some data may be even more punitive: No law enforcement agency can sub-poena data to be handed over which, legitimately and in compliance with the relevant laws, has been deleted with their retention period long past.¹⁴

Even so, retention and deletion in databases and archive systems seldom are of concern to the application, seldom are they tested as relevant parts of application function, instead, the corresponding tasks are delegated to operations as the cleaners and janitors of IT.

¹⁰ When advocating testing, we implicitly admit the possibility of erroneous infrastructure provisioning, which we earlier have convincingly ruled out. Combatting the resulting cognitive dissonance with even more beer – settling at a bottle per page – and assuming enough experience in the industry, we may reach the unsettling conclusion that the time- and battle-proven approach of testing deployments over stages has some faults indeed and, in fact, some errors reach production, but ...

¹¹ In 2017, the Swedish Transport Agency or their contractors, it does not matter in terms of responsibility or the lack thereof, mistakenly leaked the witness relocation programme from publicly accessible AWS S3 storage buckets. Don’t drink and strike though deploy!

¹² Never ever misconfigure a RAID. Never ever!

¹³ A premature data loss possibly leading to an estimation of undeclared revenue von Amts wegen is entirely different in terms of gravity indeed when compared to a loss of some random logs.

¹⁴ Lessons well learned by the financial services and automotive industry with agro-sciences yet to join the party.

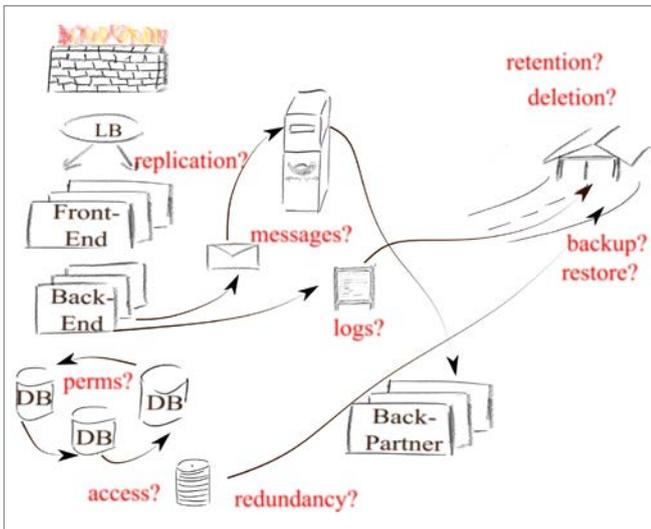


Figure 2: Automated Testing of Infrastructure Components Considered Harmful.

A possible propagation, no, proliferation, of data to a set of (on- and off-site) backup systems should be, by the same reasoning, tightly controlled. Yet, most are blissfully oblivious to the toxic waste lurking in the depths of the enterprise data vaults, not just waiting but rather begging to be discovered by law enforcement at most untimely moments.¹⁵ Yet again, punitive as it may be to mess up retention or deletion, where and in whose stewardship data may end up before it is forgotten usually is not part of application testing.

Dynamic Operations

In this respect, dynamic operations – whether cloud or non-cloud does not matter – are a game changer! While we may – successfully or not – have sat out the necessity to prove or at least verify the correctness of infrastructure and auxiliary systems regularly, when dealing with highly distributed, highly interdependent and rapidly moving systems, we simply cannot.

With automatically¹⁶ redeploying, redistributing, re-provisioning infrastructure components, transporting over stages does not work and adaptations to production are made without testing in pre-production first. The frequency of changes will be high, and errors will be made, to err is human, after all, and the frequency of errors will initially correspond to the frequency of provisioning.

So, principally, after every change the automaton introduces, the stack should be tested, and with much sighing and occasionally, fainting, we accept the inevitable of auto-

mated testing. We may speculate on the wee hours in the night when we, for a short time, may skip some tests, but conceptual, testing event driven requires more adaptation than just counting the ticks of a clock, it may be easier just to test periodically. Then, conceptually, testing and monitoring start to converge.

The outrage (“No! We do not test in production!”) starts to become bearable (“Yes, certainly we do monitor our production, mind you, it would be insane not to!”) and suddenly, we realize, that, even if the factors driving development and operations may be different, the aims are not and, in the end, similar tasks are performed by very similar means.

Then, the engineers operating dynamic infrastructure – may they call themselves Platform, Site Reliability or Operations Engineers – will soonish have monitoring or testing – again, the naming does not matter – suites in place, which will close the gap of application testing to include the infrastructure and the auxiliaries.

Only this I cannot imagine: Calling it “Test Driven Operations” – that would probably have problems to be accepted.¹⁷

Christopher J. Ruwe

Dipl.-Kfm. u. M.Comp.Sc.

Systemanalyse und -beratung

Christopher J. Ruwe (1982), having studied Business Administration, decided afterwards that learning something actually useful might have been more sensible. Having then added Computer Science at FernUni Hagen to the mix and consulting as DevOps-Engineer or rather System’s Archaeologist, professional gardening might prove an even more sensible choice. One could grow roses ...

<https://www.cruwe.de>

¹⁵ The more technical versed are untouched by such perils: Any connoisseur of enterprise IT well knows that backup may be tested, but data restore procedures almost certainly never are. The probability to successfully restore data is a function of time, decay rates and the fact that after a certain amount of time, the archaeologists of IT are required to decipher whether these glyphs over there have been a legitimate file format at some time, or just droppings left behind by marauding hordes of rodents.

¹⁶ Walle, walle, manche Strecke ... (Quote: Der Zauberlehrling Johann Wolfgang von Goethe)

¹⁷ In DevOps, we gamify operations. Two outages, and Ops-Team is out. (DevOps Borat)

Kubernetes

The essential Ingredient for every Enterprise

Less than five years old, Kubernetes has already become the de facto container management system worldwide. In 2018 Forrester's cloud predictions declared Kubernetes to be the victor in the "war for container orchestration dominance." Since then its popularity has only continued to grow and CIOs across every sector now consider it to be the gold standard for container management. In particular when it comes to supporting DevOps within their business.

This popularity is no surprise given the benefits of the technology. Kubernetes groups application containers into logical "packages" for simple, fast management and discovery. It also automates the deployment and scaling of containerised applications. Not quite a true platform in itself, Kubernetes can be combined with additional elements to provide the ease of use of Platform-as-a-Service for developers, with the adaptability of Infrastructure-as-a-Service to make it easier to move workloads across infrastructure providers.

Modern enterprise use containers as an increasingly important aspect of their business. But there are still those implementing container technology without Kubernetes. CIOs should consider the following three key reasons to embrace the industry standard:

Tech industry's influence on Kubernetes

Developed as an open source project by Google in 2014, Kubernetes immediately benefitted from its creator's heritage of employing the world's best talent and its legacy of experience around providing software services at a large scale.

But it wasn't just Google's legacy that led to the rapid growth of the system. Kubernetes brought with it major benefits including its accessibility through major cloud providers like Azure, Amazon Web Services, and Google Cloud Platform. Being hosted and supported by the Cloud Native Computing Foundation (CNCF) and The Linux Foundation also proved valuable. Association with both foundations helped the system gain credibility through its ability to continually improve through best practices and contributions from members, including SAP, AWS, SUSE, Google and more.

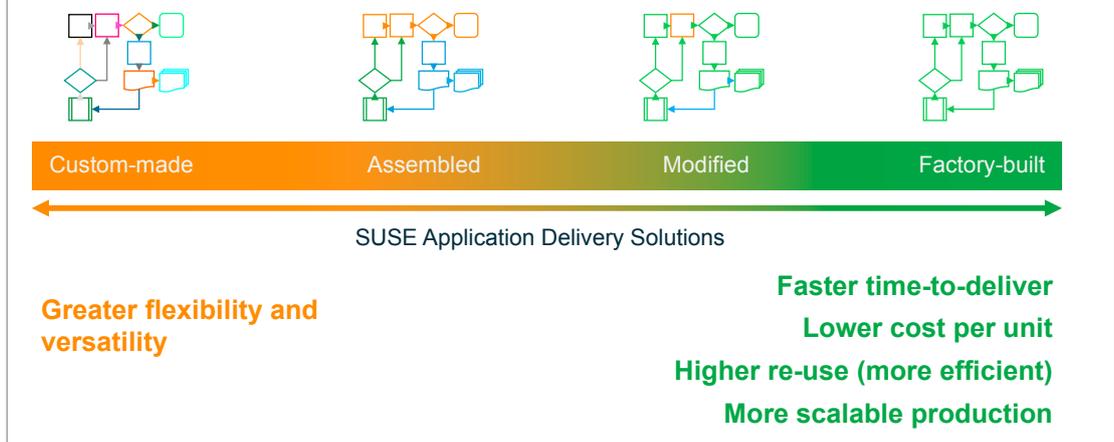
With such widespread support across the industry, Kubernetes eventually sparked the interest of enterprise vendors, which ultimately led it to where it is today – being a central component of innovation within the space.

Flexibility is the key

There is hardly an organisation today that does not deploy a combination of cloud and on-premises data centres. Thus IT needs the ability to easily and securely run and move workloads as well as share data between the various environments.

DevOps Workflows for Kubernetes

A spectrum of automation options



The unrivalled flexibility of Kubernetes makes it ideal for this hybrid environments: it can be implemented across cloud and on-premises, or even run as a service. As the system does not rely on a particular infrastructure, DevOps teams can build applications and then use Kubernetes to deploy them just as easily in the organisation's data centre as in the public cloud.

From the latest cloud-native micro-services to a legacy application which has been migrated into containers, Kubernetes can handle practically any application. In a nutshell: if an application works in containers, it will work smoothly on Kubernetes. Designed to support a wide range of workloads, Kubernetes also works well with state-of-the-art development technology, such as serverless architecture.

DevOps innovation at its best

Achieving business agility requires more than software and IT, yet applications are necessary to all companies' success. The DevOps method is revolutionising application development and helping organisations to achieve business agility – in part because of the flexibility of containers.

Luckily, Kubernetes simplifies DevOps as it allows developers to package their applications into containers and be certain that the apps will process as well in production environments as they would anywhere else. The platform removes a level of work for developers, specifically removing the need for them to spend time on scripting specific

deployment, scaling and updating workflows. Because this is all done automatically – especially when put together with application platform technologies such as Cloud Foundry – there's no need for developer input.

Developers across the world have already started to see tangible success metrics when it comes to Kubernetes. One great example is in the case of Ancestry, the global leader in family history and consumer genomics. Faced with a website that had become quite cumbersome and time-consuming in its processing, the company needed a way to accelerate and become more agile in delivering its products and solution. Fast forward to today and it is working with rapid speeds, enabling it to make customer discoveries faster than ever before.

Kubernetes is a critical component of application development and delivery solutions today. This ground-breaking technology will drive enterprise innovation now and in years to come. CIOs in need of achieving business agility should embrace Kubernetes for their businesses now to dramatically accelerate application delivery with containerised and cloud-native workloads.



Thomas Di Giacomo,
 president of engineering,
 product and innovation at SUSE,
 considers why organisations using containers, but not yet
 using Kubernetes, should employ the technology.

The Death of the Developer

Cloud, especially when understood as PaaS (Platform-as-a-Service)-approach changes a lot for developers, technologically and in regard to mindset. In fact, it changes so much, that we as developers need to reinvent ourselves – and actually need to bury our old incarnations. This new series will cover all of these aspects and starts by giving a high-level overview on the growing complexities and challenges, developers are confronted with.

When moving into a cloud environment, we could try to execute as we did in the past. This is embraced by *lift & shift*-approaches, that basically recreate environments existing in traditional datacenters in clouds.

Although this is a viable first step, it should actually be considered as the first step only, since opportunities and challenges in cloud-environments demand completely new approaches in regard to software, operations and integrations. If adjusted to these approaches, the total-cost-of-ownership (TCO) of a software will dramatically reduce while quality of service increases, and time-to-market will lower significantly.

To understand these correlations, we need to see the bigger picture first.

New requirements to software

Software gets more and more into the center of attention of organizations, since it is understood as a crucial aspect of execution and strategy. Therefore, developers need to deliver software faster and in a continuous way, ensuring

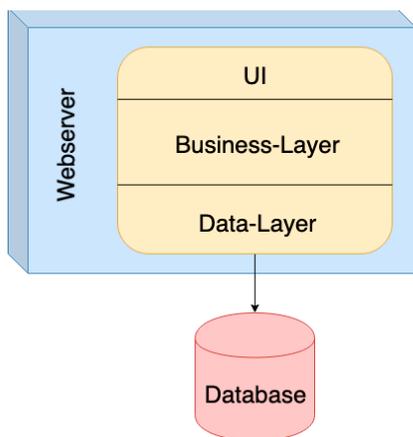
a constant stream of updates and ever-evolving quality. Software needs to run stable, it needs to allow for better testing, and it needs to integrate deeply with infrastructure, allowing for more automated operations and lower operational expenses. And, software needs to be future-proven by design, preventing from technology- and vendor-lock-ins.

These requirements cannot be fulfilled by using old, traditional approaches. Software running in cloud environments needs to be written differently, needs to be set up differently and needs to be orchestrated differently, ultimately forcing the developer to think and to act differently and change our perspectives and mindsets.

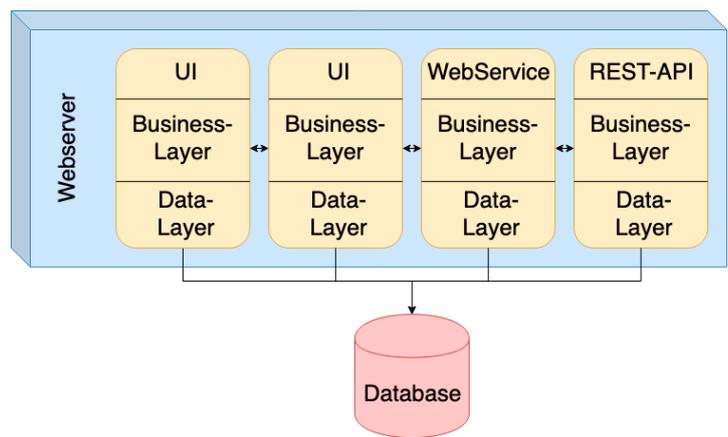
Also, the architecture of software needs to be different. This is where Microservices come into play, since they promise to deliver on these requirements.

Microservices

Traditional, monolithic software was easy to develop, but it did not allow for scalability, maintainability and man-



Monolithic Application



Application consisting of Verticals

ageability. All parts of a software were strongly coupled with each other, ensuring high performance by utilizing in-memory-communication.

This was an appropriate approach to desktop-applications, but it lacked a lot of today's requirements, as testability, simple deployment, fault-tolerance and scalability-requirements are not met at all. Furthermore, all parts of the software were written in one programming language, leading to *historically grown*, even harder to maintain, software.

This was tackled by the introduction of verticals, spreading an application over some not-so-big-services. While this was a step into the right direction, it is not enough for our modern kind of applications and environments.

This is, where Microservices, depicting a radically different way of building software, appear to be a solution. They emphasize on loose coupling of components. Communication between application components is handled via HTTP/HTTPS-protocols.

Each Microservice is completely self-contained, exposing a REST-api to the outside world, having its own datastore and is executed on a dedicated webserver. This allows for greater independence of each server from any other service, for better fail-over-behavior, for easier scaling and for using the best technology and programming language for the service's purpose.

Each Microservice is usually packed as container-image, allowing for simple installation and independence of the surrounding environment.

The usage of Microservices changes the whole game not only for developers, but also for application architects, since it implies a switch to a completely decoupled application architecture. This kind of architecture raises a lot of questions, in different areas of interest: Orchestration, deployment, fail-over-approaches, transaction management, CI/CD, automation, etc.

Kubernetes as Container Orchestrator

Since containerized microservices can be tough to operate (as there could be literally hundreds of instances running at the same time), it is important to have some sort of orchestration tool in place.

Currently, the orchestrator of choice is Kubernetes for most enterprises. As a result, developers need to familiarize themselves with it and need to understand where, when and how to use it.

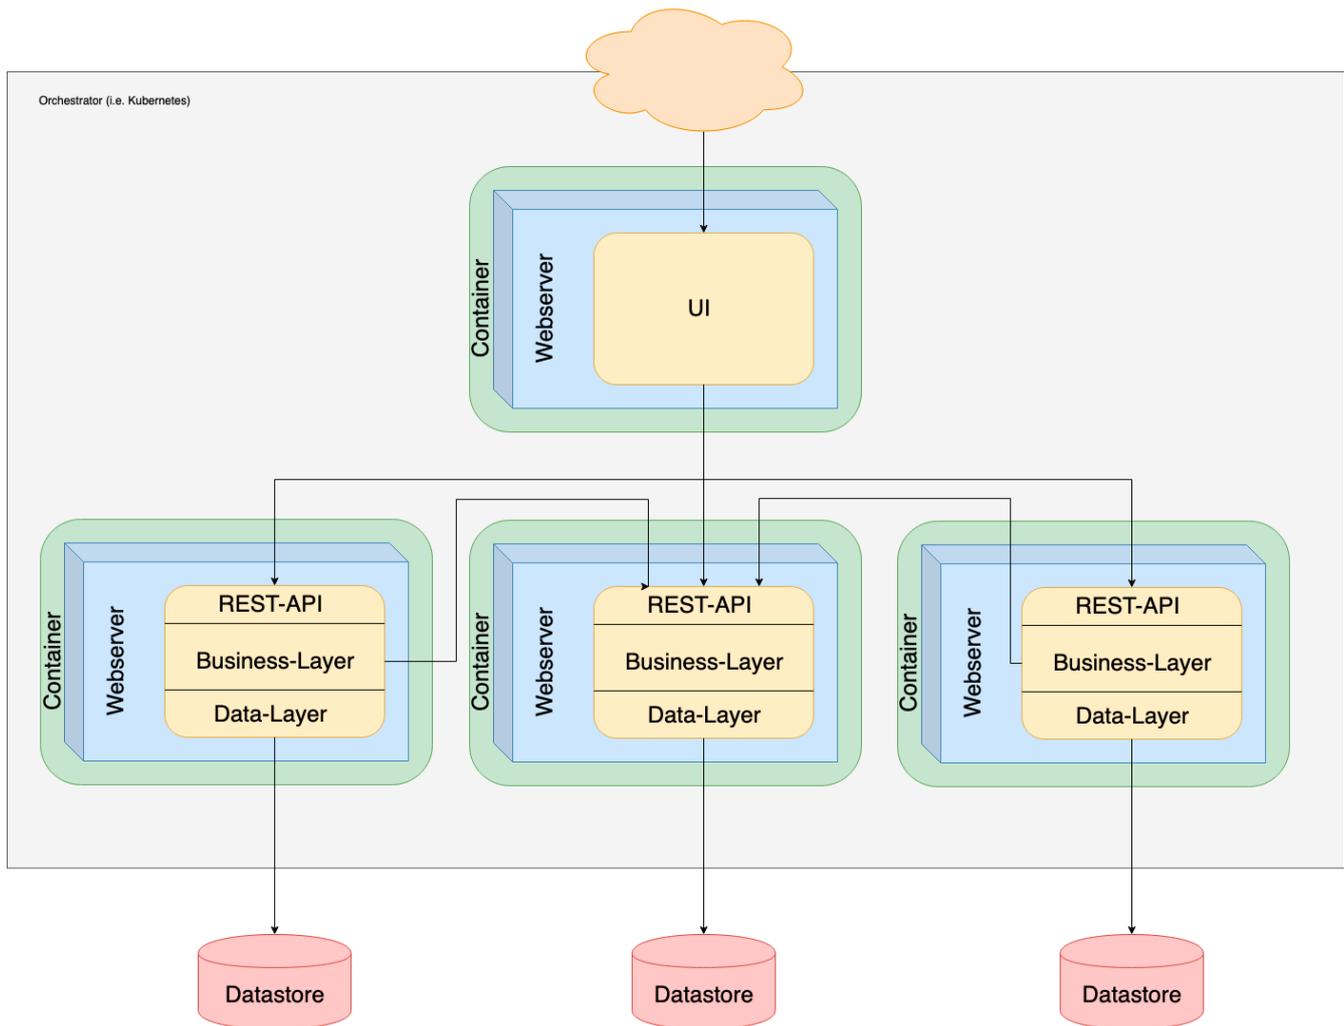
As there are different flavors of Kubernetes available, developers, operators and other stakeholders need to ensure to understand their specific pros and cons. Especially enterprise-grade variants, such as RedHat OpenShift or SUSE CaaS / CAP, are opinionated in their approaches – allowing for deeper integration with build- and deployment-processes and enforcing a tool-level knowledge beyond a default Kubernetes variant.

Adding different storage-technologies (such as Rook) and network-overlays to the game, the decision for a specific solution cannot be done by a development team alone, since this kind of technology tends to get strategic on the longer term, is complex by itself and is by no means easy to set up, to integrate and to operate in production grade environments.

(Secure) DevOps

As the complexity of applications and environments increases, the complexity of operations increases as well. Actually, scaling and automatically handling applications implies a lot of additional work to operation teams, if application and environment are not integrated deeply with each other.

Such an integration is not set up easily, since a lot of requirements, approaches and ideas from both sides of



Microservice-Architecture

the fence need to be coordinated and discussed upon. Therefore, it is essential to set up a DevOps-process by constantly involving developers and operators with each other. Each application-related aspect and requirement of operations needs to be known to developers as early as possible, and each operations-related aspect in regard to infrastructure and processes needs to be known to and understood by the operations team.

And that still is not enough.

Security of applications and environments needs to be targeted as early as possible as well. If not approached early, security requirements will get really expensive, in terms of costs and efforts as well. Application developers and -operators need to be aware of security requirements and constraints, applications need to be written with security in mind, infrastructures need to be set up accordingly. Since setting up a robust DevOps-process involving security aspects is already often ignored in traditional environments – and ignoring this in cloud environments, with services

constantly scaling up and down and new releases being deployed every few minutes, implies even bigger problems.

(Secure) CI/CD

Speaking of deployments: Since cloud-native applications mainly consist of dozens of Microservices (or serverless functions) and are supposed to be deployed by an orchestration environment such as Kubernetes automatically, the process of creating those deployments and providing the images needs to be automated as well. And – perhaps even more importantly – this process needs to run in a secure environment. This basically implies: Base container images are downloaded from a trustworthy and controlled source, external libraries need to be downloaded from a trustworthy and controlled source as well (perhaps they even need to be acknowledged by a security team), everything is build inside a controlled environment and is to be deployed from that controlled environment as well.



DevOps (Photo by Marvin Meyer on Unsplash)

Or in different words: No binaries as libraries, no base images from Docker hub and no manual installation or deployment of artifacts or images. This implies not only to think of automation, but also to set up a controlled and secure build- and deployment chain, which is understood to be a vital and essential part of any development- and operations process.

Automation and transparency

In cloud-environments, automation is a crucial aspect of lowering costs and ensuring steady and ongoing deployments of new versions. For operators, there is no such thing as SSH anymore, at least not in their tool chain. For developers, everything needs to be automated – including creating their test infrastructure and pushing out each deployment. Surely, there can be manual steps involved, but these should only be understood as quality gates. Every single aspect of building and running an application, needs to be automated.



Cloud-native Orchestrators and Tools



Grafana Dashboard (Picture taken from <https://github.com/grafana/grafana>)

Another very important aspect is transparency: In the past, developers simply dropped some output in a log file, and when something happened, an operator (or the developer) would look into the log file and tried to understand the error message or the traces being visible in there. Basically, this approach would be valid for cloud-native applications and environments as well – with one caveat: With cloud-native environments, not only one instance of a service is available at a given time, but perhaps dozens. And when an instance crashes, the orchestration software will simply remove that instance and start a new one. This leads to major problems when trying to solve issues, if not tackled by utilizing centralized logging and centralized tracing approaches.

Integration

The ultimate goal when writing cloud-native applications, is to bring all of the things previously discussed together – and to add even more aspects on top, allowing for deep integration with cloud-native infrastructure and orchestrators. This integration ensures the environment to be aware of the state of any application component and to be able to act accordingly on its own without manual interaction. This implies scaling services up and down, replacing faulty instances, allowing for A/B-testing and Blue/Green-deployment-approaches without any downtimes.

The death of the developer (as we understood it)

All of these aspects are achievable, when we as developers (and operators) bury our old approaches and our traditional mindsets. We need to understand ourselves as a team, being responsible for an application end-to-end, involving security-experts and other stakeholders from the very beginning.

In this sense, the old understanding of developers working on their own, of monolithic applications and unchanged environments is a thing of the past. This new type and new generation of developers succeeds their predecessors in many aspects, having a wider picture and living a broader collaboration than ever before. The mindset is fundamentally different to that of past generations, ultimately implying the death of these old-fashioned developers.

In the next issues of this magazine, we will dig deeper into how to create a Microservice and how to deploy it to Kubernetes.



Karsten Samaschke
Co-Founder und CEO der Cloudical
karsten.samaschke@cloudical.io
cloudical.io / cloudexcellence.io



PRODYNA is your Partner in the Digital Space

PRODYNA is a Pan-European IT consultancy creating innovative custom software applications to solve the business challenges of enterprise customers. We design, implement and operate digital solutions for mid- to large size enterprises all over Europe.

Driven by the major cloud providers, our industry currently undergoes a tremendous shift in how we build and run software. Speed, agility, and scalability are at the core of this new "Cloud Native" approach. With the "Cloud Native Computing Foundation" (CNCF), the industry now converges towards an open and standardized set of technologies with software containers and Kubernetes at its center. As PRODYNA, we are committed to offer our customers innovative and future prove solutions which require us to continuously adapt to such developments. This has lead us to join the CNCF and to become one of the first European Certified Kubernetes Service Provider (CKSP) and Kubernetes Training Partner (KTP). Cloud Native technologies



are the foundation on which your digital business can thrive. Let us help you with your cloud native transformation.

▶ Please visit www.prodyna.com for customer references and project examples.

Extend your cloud capabilities.



Interview with Nils Brauckman, CEO, SUSE

At the SUSECON 2019 in Nashville, Tennessee, I got the possibility to talk with the CEO of SUSE Nils Brauckman about open source, cloud technology, possibilities and opportunities in the IT world.

The IT world is changing rapidly and offers incredible opportunities. We are undergoing a digital transformation, but it can be challenging to be a pioneer or even keep up with the changes. SUSE is also in the midst of changes – are these easier after its independence? And is the company moving in a new direction?

At SUSE, we listen closely to our customers to understand their needs. Together with our partners, our primary mission is to enable them to succeed in today's rapidly changing business environment. Our new status as an independent business allows us to focus on our customers and partners' goals more than ever, while being the open, open source company that provides freedom of choice.

By putting our customers' needs front and centre, we help them to succeed in their digital transformation journeys, without being locked in a vendor stack. Since launching SUSE as an independent business in March, we have actually faced the same challenges as our customers. For example, to ensure agility and keep pace with both market changes and customer requirements, we had to create our own IT stack that supports our employees across different functions. This stack needs to be able to unlock the value of our data and provide insight into how our customers use SUSE's services.

Industry leaders around the world are being driven by the same business challenges and all

require the necessary tools to adapt, scale and achieve agility. Many of us have to deal with competitors – in some cases, disruptive competitors – and we want to stay one step ahead of them and the application stack. Engaging effectively with the outside world by servicing, supporting and learning from customers and partners is extremely important if an organization is to maintain a competitive edge. As the largest independent truly open, open source company,





we are reinforcing our efforts to make our customers' journeys into digitization a success.

Digitisation and open source bring great opportunities. How does SUSE live and breathe open source and what digitisation strategies is it pursuing?

Our open, open source culture is one of the key pillars of our business. Open source is the main driver of digitization and innovation.

At SUSE, our end goal is really about ensuring our customers can be innovative, agile and competitive through our industry-hardened enterprise open source solutions. We get there by providing our customers with best of breed enterprise-ready solutions that enable them to create, deploy and manage workloads anywhere – whether that be on-premises or in hybrid and multi-cloud environments.

Our recent independence and position as a long-term trusted advisor in the open source industry allows us to deliver expertise effectively. Moreover, working alongside innovative open source communities, as well as our partners and broader ecosystem enables us to take innova-

tion and provide our customers with valuable consumer technology.

SUSE is still mainly seen as a Linux distributor rather than a cloud technology provider. But the latter is the direction the new offerings are taking. With what strategy do you want to show the world that you stand for cloud? Where do you see yourself regarding cloud, machine learning, AI, IoT, ... ?

Technology doesn't emerge on its own but in response to the world around it. Many of our customers are exploring hybrid computing, cloud and multi-cloud computing and Software-as-a-Service (SaaS). These options provide multiple benefits including scalability, speed to market and economic flexibility.

Companies need the ability to move to new computing models and paradigms while simultaneously fulfilling customization requirements for technology that they may potentially already have on-premises. A lift and shift approach to the cloud will not suffice here. Crucially, organizations need to consider integration between the technologies in place and new systems. For

many enterprises the concept of hybrid is not just a buzzword, it's a reality. To embrace emerging technologies, you still have to get the most out of on-premises assets – you can't have 100 per cent one or the other.

When we consider current deep and machine learning (ML) technologies – or more generally, the future of artificial intelligence (AI) – we are already seeing the business benefits. AI is used for a variety of purposes including facilitating automation, simulation in healthcare research and helping businesses understand customers. Current AI and ML capabilities are now possible at a larger scale than ever before, thanks to recent innovation in high performance computing (HPC), specific software and hardware, hybrid and multi-cloud application and data management. I would argue that all new technologies – AI, HPC and IoT to name a few – are based on open source innovation.

Many of SUSE's customers and partners are still more classically positioned, so how do you intend to change mind-sets around transitioning to the cloud? Are you looking for new partners, or do you want to enable your existing partners and customers?

From listening and talking to our customers and partners, we know that all business leaders are driven by the same need to adapt and scale and be agile. Hybrid cloud and multi-cloud go hand in hand. Businesses want to achieve the

degree of flexibility that only best of breed can give them. Therefore, they may decide to work with several SaaS providers, and also several cloud providers. Multi-cloud strategies need to be managed, orchestrated, and fitted within an IT landscape that has been growing for 20-30 years.

At SUSE we see ourselves in the role of building bridges and opening new possibilities while catering to the existing infrastructure and ensuring everything works together seamlessly. Our Kubernetes-based container solution is a good example of this. Most businesses use a combination of cloud and on-premises data centres. IT requires the ability to easily and securely run and move workloads, as well as share data between various locations. Kubernetes is ideal for this hybrid environment as its flexibility means it can be implemented across cloud and on-premises, or even run as a service.

Our customers need partners at their side that are willing to really understand their requirements. Ultimately, partners need to be deeply engaged, dependable and act as trusted advisors. They must understand the rapidly transforming business space and the current demand for mixed IT environments. SUSE works extremely closely with existing partners. But we are also always searching for new partners who are driven by the same motivation – working togeth-



er to help our customers successfully transform their business is our core goal.

Why should customers choose SUSE products? What is the value for customers to move to the cloud with SUSE?

Open source fuels digital transformation. SUSE is committed to working within the spirit of open source technology. By collaborating with our partners and wider ecosystem to bring valuable technology to our customers, we are also making our efforts available to open source communities.

Our open, open source approach ensures that our customers are not locked into a vendor stack, but have the freedom of choice to select the best solution for their environments. SUSE connects the familiar stability and functionality of enterprise systems with an expanded, innovative, industry-hardened open source system. This makes it possible for companies to implement agile working methods using modern IT infrastructure and open source software. In a nutshell, SUSE enables customers to create, orchestrate, deploy and manage workloads anywhere – on-premises and in hybrid or multi-cloud environments.

At SUSECON, you can feel a lot of SUSE's own spirit. How does your company understand culture? What is important to you?

SUSE's culture is an integral part of our business. Being the open, open source company drives the way we work to achieve goals and solve problems together with our partners, customers, community and employees.

It was incredibly exciting to see how the SUSE teams across the globe gathered together on our independence day. Our chameleon and the colour green was everywhere, SUSE bands played, and there were mega cakes – the Chinese team had the largest cake I've ever seen! It is astonishing to see this level of pride, passion and belief in the future of SUSE. I am very proud that we are now at around 1,750 employees from 75 different nationalities in over 34 different countries. We are all motivated by the same mantra of true openness – and this is essential to SUSE.

The interview was conducted by Friederike Zelke.



CONTAINER DAYS SPONSORING



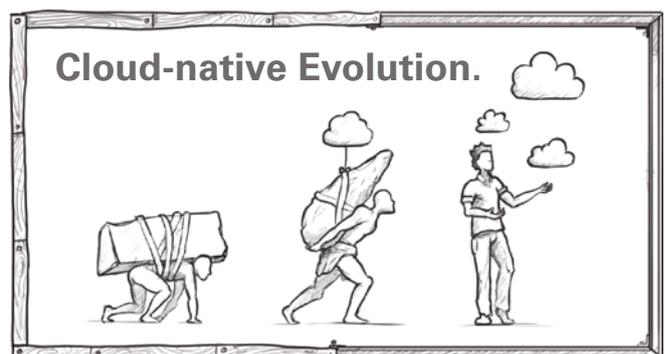
Making Cloud Native Work

QAware is an independent software manufacturer and IT consultancy. We create value by means of high-end computer science. Amongst our customers are BMW, Allianz and Deutsche Telekom.

We are there to help in technical crises. We renovate systems and take over end-to-end responsibility. We build state-of-the-art software systems that provide long-term benefit to our customers. We maximise benefit and revenue through continuous innovation in software engineering. We minimise costs through high quality and productivity.

Cloud native applications help to cope with an ever-growing number of users, devices and data volume, they cater for short release cycles and antifragility. But with the inherent paradigm shift in development and operations comes significant complexity that is increasingly difficult to handle by individual teams.

Enterprise Cloud Native is our approach to reducing the vertical range of manufacture of modern systems and to boost the cloud-native developers' productivity and happiness.



DevOps and **Continuous Delivery** change the software development process sustainably. An integrated tool chain and a high degree of automation are required to efficiently and continuously deliver high quality software: our software assembly line. It is the backbone for successful CI/CD.

At the ContainerDays we want to learn from other experts and share our field-tested knowledge of cloud native software engineering.

Visit us also at www.qaware.de.

Digital Transformation in Medical Technology

Interview with Jörg Haist, Dentsply Sirona



With the merger of Dentsply and Sirona in 2016, the company has become the world's largest manufacturer of dental products and technologies in the field of dentistry and dental technology. During the past 3 years, Dentsply Sirona managed the multiple transformation of merging over 16.000 employees in over 40 countries into one single company and realizing an unprecedented digitization journey centers around connecting practitioners, patients' needs and equipment. Julia Hahn and Felix Evert discussed with Jörg Haist, Vice President of Platform Management Equipment & Instruments, how Dentsply Sirona leveraged the opportunities of digital transformation in medical equipment.

Dentistry is undergoing an enormous upheaval. The keyword Digital Dentistry moves the entire industry, especially the equipment manufacturers. How do you observe this transformation process?

10 years ago, X-ray systems were the start of digitization. It all started on a small scale: Initially, the devices were equipped with a digital interface, but were not yet networked with each

other. Soon the desire arose in the practices that the images should also be available in other rooms than just the X-ray room. The dentists now wanted access to the images, for example, in the treatment room, at the reception desk or in the patient files. At that time, the images were still available in 2D; today, they are of course available in 3D. The practice management systems followed the same path. Here we faced the same challenge: Patient data should not only be available at the reception desk, but also in the treatment room. These two systems have greatly advanced digitization and networking in the dental practice and continue to act as catalysts for other systems. Today, all devices at Dentsply Sirona are networked. The devices have various interfaces at which they collect and pass on information; patients can communicate with the management systems, for example. On the operating side, for example, cleaning or maintenance cycles can be documented. The technical challenge lies in the growing volume of information. Sensor technology is becoming faster and more precise. This results in a large amount of data that we have to process sensibly and securely.

The transformation to digital dentistry means a great effort for all involved, which can only be justified by better treatment results for patients. What does Dentsply Sirona's promise "Inspired by your needs" mean in this context? What needs do you see now and in the medium term?

One of the major tasks of digital networking is that it must be made suitable for users. Our motto for this year's IDS, "Inspired by your needs", starts right there: What is the problem, why does this problem exist in practice and how can we solve it?

It is currently becoming apparent that the virtual patient is the focus of attention: All information about a patient is transferred together to a virtual filing system. There it is made accessible for the treating dentist. Dentists today no

About Jörg Haist

Jörg Haist is Vice President of Platform Management Equipment & Instruments at Dentsply Sirona. Prior to the merger of the two companies, he was Product Manager Imaging Systems at Sirona. There he was responsible for the X-ray machine Orthophos SL, Galileos Comfort Plus and the Xios XG sensors.

Before joining Sirona, he worked at Fraunhofer Institut für Graphische Datenverarbeitung (IGD). He has also taught graphical data processing and medical informatics at the universities of Heilbronn, Heidelberg and Mainz and at the TU Darmstadt.

longer want "disparate information bites". A good example of this is implant planning: the restoration does not consist solely of the implant itself. During X-rays, the doctor can also see what it looks like under the crown. Without an x-ray, it is only possible to look at the restoration from above.

The patient does not initially care which technology is used. The patient would like to receive a complete treatment from his dentist. This is why we brought the various technologies together several years ago.

Dentsply Sirona deliberately works directly with dentists to incorporate their needs and wishes into product planning. There is a close exchange here, both in development and in the treatment practice, i.e. where dentists use our systems. Especially in software development, there is an increased risk that a lot of money will be invested in something that will not be used afterwards. Therefore, it is essential to analyze and consider the needs of users. Not only the dentists, but also we as a manufacturer are satisfied if everything works smoothly and without complications.

Dentsply Sirona actively promoted a digitization strategy more than 10 years ago. How do you rate and document the progress made so far? What path does Dentsply Sirona still have ahead of it?

We are constantly developing our products and systems and are always looking for ways to improve them. Especially in the field of digitization we see great potential for optimization

of the intranet of dental practices. Above all, we, as a manufacturer, focus on cross-practice networking via the cloud in order to facilitate the exchange of electronic patient data and X-ray images between the attending dentists. Further developments in artificial intelligence (AI) also appear promising in this context. Such AI systems are also increasingly being used in human medicine - for diagnostics, for example. In the future, these systems will probably even be able to suggest findings or automate entire processes. Incidentally, this technology is strongly cloud based.

As these systems are very maintenance-intensive, digital networking is gaining in importance here in particular. This allows the systems to draw attention to themselves independently and communicate with the user. The systems can thus draw the attention of their users to pending maintenance and care work with "Attention, I need maintenance", "Salt" or "There is a problem", for example, and thus support smooth work processes.

In addition to technical innovations, it is also worth taking a look at developments in dental practices in Germany: more and more large practice chains with centralized purchasing or technical maintenance are emerging. Networked and automated communication is particularly worthwhile for these practice chains, in order for them to simplify and accelerate organizational processes.

Dentsply Sirona is rapidly bringing innovation to the ecosystem. How do

you develop this power of innovation? Where does the drive come from?

In dentistry, basic technological research is less pronounced. This is more likely to happen in other disciplines. The art lies in adaptation, i.e. in bringing an idea into the dental environment that already exists elsewhere. In imaging, for example, there is already very good basic work from human medicine. Our Orthophos SL X-ray machine was one such development. It enables panoramic images with the best resolution at the lowest possible radiation dose. Innovation is therefore often achieved through the adaptation and new composition of existing approaches. The merger of Dentsply and Sirona has brought together first-class specialists in a wide range of fields. We learn a lot from each other and can now focus more eyes on the dentist from different perspectives at the same time. Our perspective has widened.

At IDS, Dentsply Sirona also emphasizes its “simpler and clinically safe”

solutions. How is this approach integrated into the innovation process?

Of course, there are comprehensive guidelines on how to develop medical devices. These are, among other things, in the areas of readiness for use and usability. Dentistry, in particular, is a discipline that is highly technological compared to other medical fields. As a result, the many different technical products must be easy to use. No user can master 50 devices down to the smallest detail. Therefore, we try to make our products as user-friendly and “simple” as possible. Many physicians are therefore closely involved in our development processes, be it in test installations, concepts or with personal feedback. We want to prove that products and solutions fulfil their purpose: Planning therapies, diagnostics, carrying out treatments.

At the end a short outlook. Dentsply Sirona has just announced its cooperation with exocad. The digital workflow between practice and

laboratory is to be promoted and improved. What will it look like?

As I have already mentioned, all our products are becoming increasingly networked. This concerns the solutions from our company, but of course also the networking with the devices of our competitors. exocad is such a cooperation in which we work with validated workflows. The manufacturers exchange information in the interests of the doctors: How do the products actually interact and how do we ensure that this interaction functions smoothly? - We deliberately do not say: it only works with Dentsply Sirona. The physician should have and retain freedom of choice. The networking within the practice should nevertheless remain guaranteed, regardless of the product choice.

Mr Haist, thank you for this interview.

The Interview was conducted by Julia Hahn and Felix Evert.



CONTAINER DAYS SPONSORING



Fedora CoreOS & Fedora Silverblue

Become fully cloud-native by adopting best practices across all your professional computing – from your desktop and server needs to container deployment at scale. For your desktop, Fedora Silverblue means container-focused workflows and failsafe upgrades. This variant of Fedora Workstation is in high demand in developer communities. Its image-based nature and git-like structure is familiar from OSTree-based servers.

Server-side, Fedora CoreOS is an automatically updating, minimal, monolithic, container-focused operating system. It is designed for clusters, optimized for Kubernetes, but also operable standalone. It aims to combine the best of both CoreOS Container Linux and Fedora Atomic Host, integrating technology like Ignition from Container Linux with rpm-ostree and SELinux hardening from Project Atomic. Its goal is to provide the best container host to run containerized workloads securely and at scale.

Built around a core of OCI container packaging and Kubernetes container cluster management, the next project



to check out is OKD, the Origin Community Distribution of Kubernetes that powers Red Hat OpenShift and provides a complete open source container application platform.

Visit <https://coreos.fedoraproject.org/> to learn more.



CLOUDICAL

Delivering CloudExcellence

At the earliest possible date we are looking for a

Cloud Engineer for OpenStack (f/m/d)

TASKS

- Design and implementation of OpenStack
- Building cloud setups on public, private or hybrid OpenStack infrastructures
- Competent contact person for our customers
- Technical consulting to customers
- Conducting technical trainings
- Sustainable and consistent automation of recurring tasks
- "Hands on" Troubleshooting
- „Containerization“ of existing solutions
- Monitoring and operation of customer infrastructures
- Participation in community meetings and events

SKILLS

- Experience in building, operation and troubleshooting of midrange and large OpenStack environments
- Advanced knowledge in at least three of the areas:
 - Software defined Storage (Ceph, Quobyte, NetApp)
 - Software defined Network (Juniper Contrail, NSX-T/V, OpenDaylight, OVS, Tungsten Fabric)
 - Container (Docker, Kubernetes)
 - Configuration Management (Puppet, Chef, SaltStack, Ansible, Terraform)
 - Monitoring (Prometheus, Grafana)
 - Logging (Elasticsearch, Logstash, Fluentd, Kibana)
 - Continuous Integration/Delivery (Jenkins, Gitlab CI, DroneCI)
 - Source Code Management (Git, Github, Gitlab, Gerrit)
 - Databases (MySQL, PostgreSQL, CouchDB, MongoDB, Cassandra, CockroachDB)
 - Programming languages (Golang, Python)

ADDITIONAL

- Certifications in the above areas are beneficial
- High willingness to learn
- Quick grasp of complex IT landscapes
- Willingness to travel
- Very good spoken and written English

WE ARE

Our aim is to provide our customers with the best possible advice, to accompany them through the challenges and upheavals of the cloud revolution and at the same time to always be at the cutting edge of technology and processes. We pass on this knowledge and experience to our customers in the form of enablements, coaching and workshops and accompany them as reliable partners on their way into and within the cloud. We, Cloudical, act manufacturer and platform independent.

Please send us your complete application including your salary expectations and your earliest possible starting date to the following address: tom@cloudical.io

Come to us!

Contact: career@cloudical.io
cloudical.io

Collaboration without Boundaries

As announced in November in Berlin, the OpenStack Summit is officially re-named Open Infrastructure Summit and had its kickoff in Denver on April 30th. But this is not only a change of name, but also a change of strategic orientation. As of now, not only OpenStack projects will be supported, but other software and infrastructure projects as well. This is an important step in a rapidly growing and changing ecosystem. Jonathan Bryce, executive director of the OpenStack Foundation stated, while opening the keynotes:

“OpenStack is one of the most successful projects in history. That’s because of the community, and it’s also because of the problems we’re solving and the bigger trends in the world. Technology drives progress and change, and technology today always includes software. Cloud is how software moves faster, and software is how humans move faster.” Thus, it only seems a logical consequence to open up the Summit to more and other infrastructure and software projects and to collaborate more – without boundaries.

The Four Opens

Open Source

We do not produce “open core” software.

We are committed to creating truly open source software that is usable and scalable. Truly open source software is not feature or performance limited. There will be no “Enterprise Edition”.

We use the Apache License, 2.0.

OSI approved

GPLv3 compatible

DFSG compatible

Open Design

We are committed to an open design process. With every development cycle, the OpenStack community holds face-to-face events to gather requirements and write specifications for the upcoming release. Those events are open to anyone, and include users, developers, and upstream projects. We gather requirements, define priorities and flesh out technical design to guide development for the next development cycle.

The community controls the design process. You can help make this software meet your needs.

Open Development

We maintain a publicly available source code repository through the entire development process. We do public code reviews. We have public roadmaps. This makes participation simpler, and allows users to follow the development process and participate in QA at an early stage.

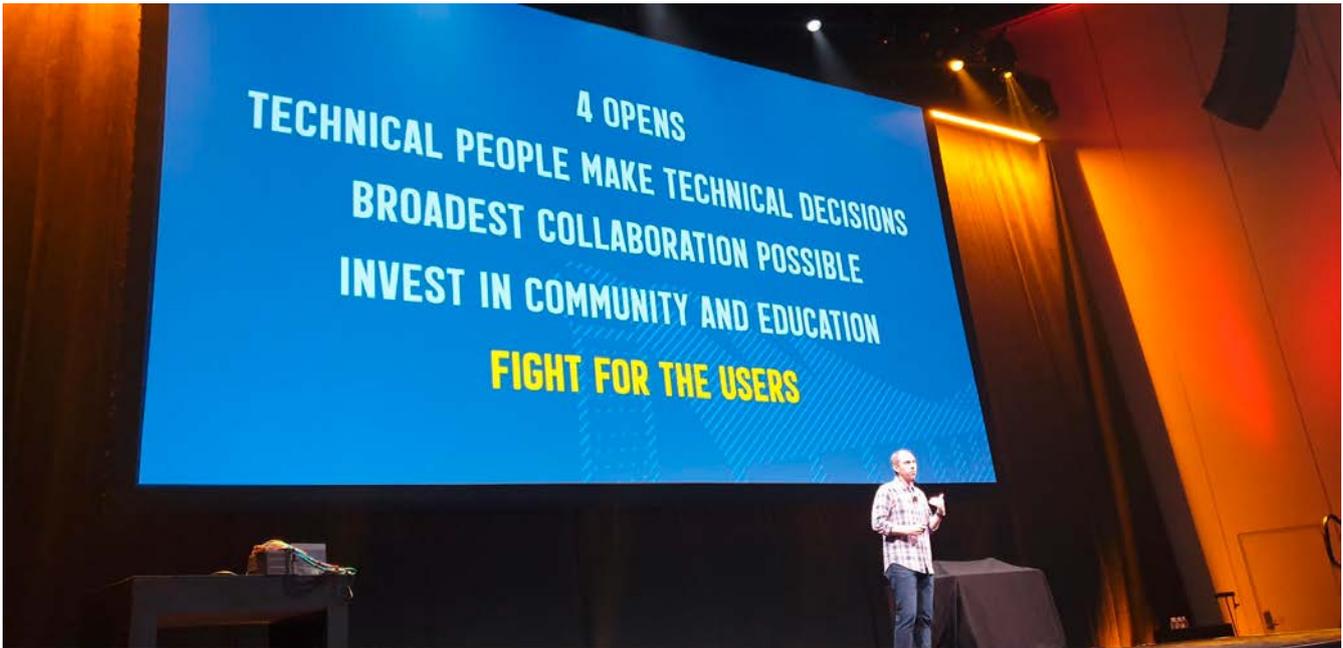
Open Community

One of our core goals is to maintain a healthy, vibrant developer and user community. Most decisions are made using a lazy consensus model. All processes are documented, open and transparent.

The technical governance of the project is provided by the community itself, with contributors electing team leads and members of the Technical Committee.

All project meetings are held in public IRC channels and recorded. Additional technical communication goes through public mailing lists and is archived.

Source: <https://governance.openstack.org/tc/reference/opens.html>

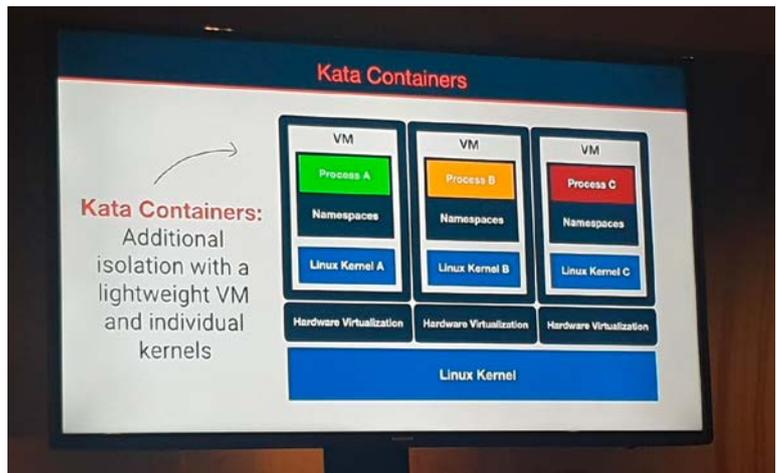


Jonathan Bryce talks about “the four open”.

Those collaborations were to be listened to and worked on in over 300 sessions and workshops on Public Cloud, Private and Hybrid Cloud, Security, Container Infrastructure, CI/CD, and diverse members of open source communities. And of course there were some new members introduced during the event. Zuul and Kata-Container, which were pilot projects so far, were announced as leaving their incubating status earlier than the max. of 18 months and are from now on graduated projects. This means, that they are fully governed by the rules of “The Four Opens” as of now (see below).

Also unchanged was the diversity of attendants at the Summit. There were 2.000 registered participants from all continents. North Americans were the leading participant group with 66% of all attendants, followed by Europeans (20%), APAC (10%), Middle East (2%) and Middle and South America (2%). 11 percent of all attendees were women.

The overall feedback of the conference was positive: “I really enjoyed the





conference and think it was a good idea to have three whole days followed by the Project Teams Gathering (PTG) Days afterwards.” stated Stacey Fletcher from St. Louis. And David Sha liked the most, that one could “actually meet people that you normally work with virtually in person. The event was a good possibility of networking. I also went to a lot of sessions and they were all really good. I mainly enjoyed the product updates.”

Join the people building and operating open infrastructure, like Airship, Ansible, Ceph, Docker, Kata Containers, Kubernetes, ONAP, OpenStack, Open vSwitch, OPNFV, StarlingX, Zuul and more.

Julia Hahn



CONTAINER DAYS SPONSORING

SUSE CaaS Platform

KUBERNETES, READY FOR THE ENTERPRISE

SUSE CaaS Platform is an enterprise class container management solution that enables IT and DevOps professionals to more easily deploy, manage, and scale container-based applications and services. It includes Kubernetes to automate lifecycle management of modern applications, and surrounding technologies that enrich Kubernetes and make the platform itself easy to operate. As a result, enterprises that use SUSE CaaS Platform can reduce application delivery cycle times and improve business agility.

Accelerate modern application delivery with Kubernetes, today’s leading container management platform.

Simplify Kubernetes administration with an exceptional platform operator experience.

Maximize return on investment with a flexible, no lock-in solution.

Key Features

A Cloud Native Computing Foundation (CNCF) certified Kubernetes distribution, SUSE CaaS Platform automates the orchestration and management of your containerized applications and services with powerful Kubernetes capabilities, including:



- » Workload scheduling places containers according to their needs while improving resource utilization
- » Service discovery and load balancing provides an IP address for your service, and distributes load behind the scenes
- » Application scaling up and down, accommodates changing load
- » Non-disruptive Rollout/Rollback of new applications and updates enables frequent change without downtime
- » Health monitoring and management supports application self-healing and ensures application availability

Happy 5th Birthday, Kubernetes!



Bryan Liles, Senior Staff Engineer, VMware, with the „CNCF project update“



Dan Kohn, Executive Director, Cloud Native Computing Foundation, holding his Keynote „Stitching things together“ to open KubeCon

And a big party it was! 7.700 guests gathered on May 20th, at KubeCon CloudNativeCon in Barcelona, to celebrate the 5th birthday of Kubernetes. This was the biggest conference on K8 and CloudNative technologies in Europe hosted by the CNCF so far (Copenhagen 2018: 4.300 attendants, Berlin 2017: 1.500 attendants). Moreover, there are more than 400

members in CNCF right now, and a community of 88 End-User-companies. The whole community is growing rapidly: As of now, Kubernetes has more than 2,66 Million contributions and 26.214 contributors. Thus, the overall theme of the conference – “Empower the Community” – seemed only logic. As a result, this year’s conference offered a lot of different work-

shops on the whole ecosystem, with networking and sharing of ideas on how to apply cloud native techniques in different areas and sectors.

Traditionally, the Keynote introduced several new releases and projects, such as Rook, Helm, Open Telemetry, Harbor, fluentd, CRI-O or linkerd et al. (see box below). Moreover, it tried to give a glance not only

Brief Summary

Rook: version 1.0 released before KubeCon; is a storage manager for the easy enablement of simple and uniform usage of storage implementations such as Ceph or NFS on Kubernetes (Source: <https://rook.io>).

Helm: first Alpha release of Version 3.0 is available now; is a K8 package manager, new feature: Tiller is not necessary any more, which increases security and decreases complexity (Source: <https://helm.sh>)

OpenTelemetry: created through a merger of OpenTracing and OpenCensus; aims to provide effective observability and "is made up of an integrated set of API's and libraries as well as a collection mechanism via an agent and collector." (Source: <https://opentelemetry.io>)

Harbor: version 1.8. released; is an "open source cloud native registry that stores, signs and scans container images for vulnerabilities" (SIC! Source: <https://goharbor.io>)

Fluentd: now officially graduated CNCF project; it is collecting data for unified logging layers (Source: <https://www.fluentd.org>)

CRI-O: incubating project of CNCF; offers lower runtime for containers (Source: <https://cri-o.io>)

TiKV: incubation project of CNCF; "distributed transactional key-value database" (Source: <https://tikv.org>)

OpenEBS: sandbox project of CNCF; storage provider for containers (Source: <https://openebs.io>)

Linkerd: service mesh for Kubernetes and an alternative to Istio (Source: <https://linkerd.io>)





at Kubernetes users and contributors but also at the younger generation, when Nikhita Raghunath and Lucas Käldestrom entered the stage. 17 year old Lucas, CNCF Ambassador, and Nikhita, young college graduate and software engineer, were talking about their very own experiences with Kubernetes: “Inclusive is better than exclusive. Anyone is welcome. Each contribution is valued”, Lucas said. And Nikhita, confirming this, told the audience how much she already learned from the community and how supportive everyone was to her.

The importance of supporting each other was confirmed later that day by Aparna Sinha, director of product management for Kubernetes and Anthos at Google. She talked about the three reasons why Kubernetes is

so successful today: “First, the early on release cadence of three months was of great importance. It gave the community a lot of confidence. Second, the open source community itself, the feedback of a lot of early users and the many developments evolving from it helped to make Kubernetes so big and successful. Last but not least, it helped a lot, that there was a foundation established behind it.”

And Abby Kearns, Executive Director of Cloud Foundry Foundation stated: “Kubernetes is only one technology here, there are many different technologies covered here at KubeCon CloudNativeCon as well. It is all about the cloudnative technologies and the work in the community. Next to Kubernetes, there are other really great projects one can get ex-

cited about, like Prometheus, Containerd, Envoy, Linkerd and others. CloudFoundry, for example, is working on a project called Eirini at the moment. Eirini is offering a pluggable scaler inside CloudFoundry. Now, operators can use Kubernetes to orchestrate application container instances. At the moment, Eirini is in Alpha version and is being tested by IBM.”

To sum up, CNCF really got the message of the community-spirit across at this KubeCon in Barcelona: Inclusive is truly better than exclusive, especially in a vast growing community as this one.

Julia Hahn



SUSECON 2019

Celebrating Independency and New Technology

In the first week of April the open source community met in Nashville, Tennessee, to learn more about the opportunities of open source cloud technology. These five Days of open source technology attracted 1100 attendees from 45 different countries. Three full packed conference days with hundreds of different talks followed combined with celebration, communication, events all around open source.

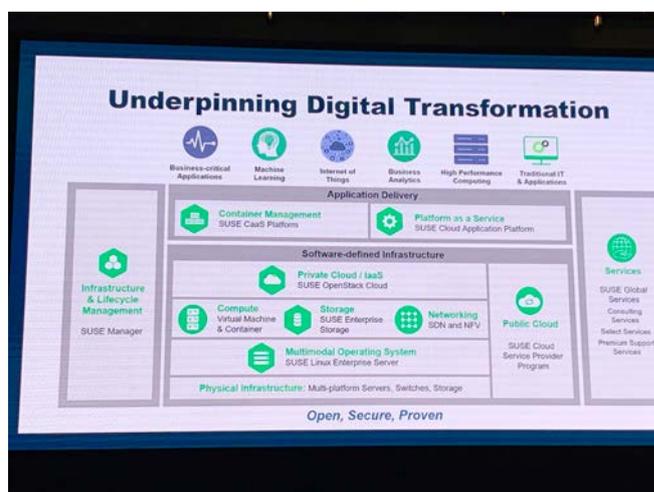
SUSE as an independent open source software company invited to celebrate their new independency and the new course of SUSE. With the title *Open open source! My Kind of OPEN* SUSE wanted to bring over their new approaches based on their

new solutions for a new market. SUSE understands themselves to be much more than just an enterprise Linux distributor. Their vision is: Provide a smart way with balanced approaches for their customers – organic and inorganic development, expertise, deliver the customer’s strategy faster into the market, as SUSE’s chairman Nils Brauckmann pointed out.

Therefore, SUSE has developed different products such as *Container Application Platform (CAP)*, a CloudFoundry solution based on Kubernetes, allowing to run containers on CloudFoundry. Applications and their developers can utilize the power of CloudFoundry to ensure faster time-to-market and more automation, al-

lowing them to focus more on their needs. CAP is SUSE’s Platform as a Service-solution. As per SUSE, Diego containers in CloudFoundry are unnecessary now, since CAP allows containers to run in a native Kubernetes environment.

Another product, SUSE wants to focus on, is *CaaS Platform*, a managed container platform. It includes Kubernetes to automate lifecycle management of modern applications, and surrounding technologies that enrich Kubernetes and make the platform itself easy to operate. As a result, enterprises using SUSE CaaS Platform could reduce application delivery cycle times and improve business agility.





Both solutions are flexible and run in several clouds, private, public, hybrid, multi, ... As per SUSE, technology should be an innovator for new approaches, be a partner for the change, empower the community to explore the opportunities of software. The vendor sees a lot of change in IT roles and the way they are understood within enterprises, since change is not only driven by customer's demand, but also from within the organization. As SUSE's CTO for Europe, Dr. Gerald Pfeifer pointed out: "Software and technology are becoming the center of organizations".

As most of the attendees came from Linux operating system, they learned about the new technological possibilities of digital transformation. In the conference talks were given deep insides in cloud technology and practical success stories, questions of how business can change using cloud technology, implementation of new processes, deeper engagement of



Dev- and Ops-roles, and the creation of new development experiences by the help of containerization took the most room and were discussed openly and intensively.

SUSECON 2019 was an interesting conference – not only by the number of SUSE executives being present and approachable, but by the emphasis being put on cloud as the major thriving force for SUSE’s future development.

Deep insights into a changing organization were given, the commitment to Open Source was even strengthened and important milestones in regard to cloud and cloud technologies were presented. The event worked on a multitude of levels – business, personal and networking aspects were covered, and as the size of the conference was still manageable, a great sense of optimism and trust in SUSE’s future prospects was perceptible.

The last evening rounded this week of celebrating technology with a party in the famous Wildhorse Saloon in Nashville where the SUSE Band performed live! So, we left with music in our heart and technology in our heads. We are looking forward to next year.

Friederike Zelke



CONTAINER DAYS SPONSORING



Migration from PaaS to Managed Kubernetes in 4 weeks

At the beginning of 2019, the German provider of smart project management software TeamGrid succeeded in migrating into a managed Kubernetes solution from ScaleUp Technologies that offers dedicated cloud resources and intelligent load balancing.

The move from PaaS to the new Kubernetes cluster environment, including the necessary adaptations and testing, took only four weeks.

Designed for stability and reliability, the solution provided by ScaleUp consists of a Managed Kubernetes Cluster combined with a managed MongoDB database running in OpenStack. „Thanks to the heavy exchange of expertise, we were able to build up the necessary Kubernetes basic knowledge very quickly. Today I am convinced: It was ab-



solutely the right decision to switch from PaaS to Managed Kubernetes”, says Tobias Hieb, CEO of TeamGrid.

TeamGrid is now currently working on an enterprise solution which will offer TeamGrid App customers a dedicated Kubernetes cluster environment. The cluster images and the virtual infrastructure will then simply be reproduced.

» <https://www.scaleuptech.com>

On the Go – GoDays 2019 in Berlin



The programming language Go has already entered the cloud and fans all over the world are pushing it forward. That's why the vivid community recognized happily the very first GoDays powered by Loodse and located in the Innovation Factory in Berlin. Where they talk, discuss, reveal the news about Go and make the event rounded successful.

Julian Hansert, CEO and Co-Founder of Loodse, about the conference: *We've been working with Go for some time now and got to know a very cool community there that gave us a lot of input. We are very grateful and wanted to give something back. That's why*

we started the GoDays and are happy that so many have registered.

The first day was filled with practical workshops and hand-ons about the possibilities of Go and an evening Meetup with the Women who Go. The second days was the conference part of the GoDays. As often I could only visit some of the talks, there were two in parallel, but I give a short overview about what I listened.

Natalie Pistunovich - Fraugster started as a very active member of the Go community. As founder of Women Who Go she brings in a fresh point of view and new input into old problems.

WWG meet regularly in Berlin and is open to new members and new approaches in GO.

Another topic at this day was security with Go, Elena Grahovac - N26 and Baruch Sadogursky - JFrog showed some approaches to develop secure with Go. Date security can be achieved with several apps like Fluentd, SQL injection, GoSec... But one important point while using the open source language Go is to be aware, that it is open source. Everybody can see, copy, use it as well. So, you need continuously monitoring, no automated updates without a check, enough maintainer of Go, the newest ver-

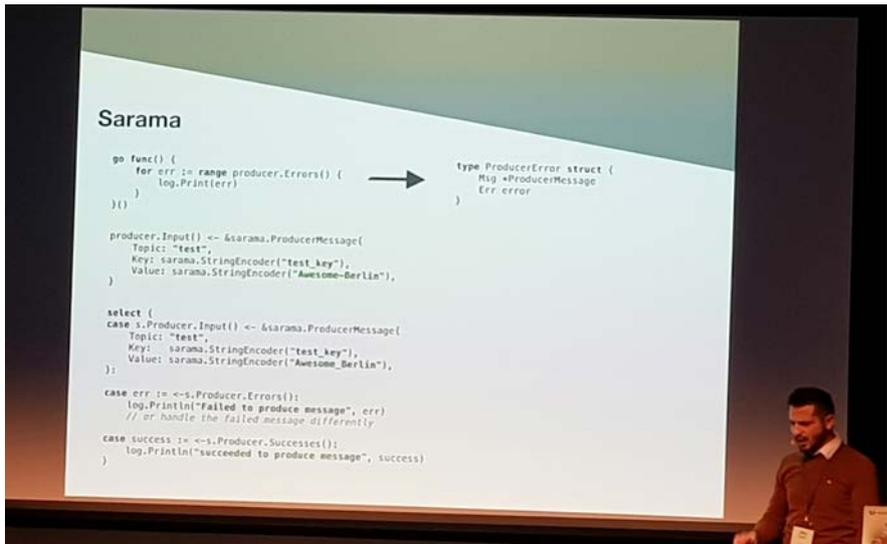
sions, immutable dependencies, and inhouse repositories. Go.mod, for example, is a monitoring open source library for Go, but you find more about at GitHub and GoCentre.io.

Moath Qasim – Ligatus GmbH showed how Apache Kafka can be used with Go. Therefore he introduced Sarama. Package Sarama is a pure Go client library for dealing with Apache Kafka.

A highlight of the GoDays was the short presentation of Tobias Kohlhaas. He was able to program a cleaning robot into a web radio without touching the device just by hacking the software and using Go.

All in one the days were filled with cool stuff around Go and showed the vivid, playful and interested open source community with news, tricks, and practical tips. And I'm happily looking forward to the next year to the second GoDays.

Friederike Zelke



CONTAINER DAYS SPONSORING 

Looking for a seamless cloud experience?

Do you rely on container technology and operate more than one Kubernetes cluster? IBM is committed in helping their customers become a best-run business. With IBM Cloud Private (ICP) IBM has developed a solution which helps the cloud world run better. With the IBM Multicloud Manager as part of ICP you can manage different cluster easily. With the IBM Multicloud Manager, IBM is addressing the need for greater container flexibility and portability. This allows a consistent management of containerized workloads distributed across different clouds - without a vendor lock-in. IBM makes things easy manageable - Operate Kubernetes clusters and containers where they make sense, depending on the requirements - in the public cloud, the private cloud or both. IBM Cloud Private incorporates open source tools like k8s, istio, helm, terraform and many more tools out of the box. A proven solution for a seamless cloud experience regardless where it resides.

We are TechData - one of the world's largest technology distributors. We help companies like IBM to bring their products to market, and we offer a wide range of technical



and business support services. We certify, train, instruct and support our reseller customers, and help configure, install and finance their purchases. Many of the hardware devices and software applications that you use in your daily life have passed through our logistics centers on their way to market.

Find out more about IBM Cloud Private - Contact us today.



CLOUDICAL

Delivering CloudExcellence

At the earliest possible date we are looking for a

Automation Engineer (m/f/d)

TASKS

- Building cloud setups on public, private or hybrid cloud infrastructures
- Design and implementation of IaC (Infrastructure as Code) deployments
- Participation in community meetings and events
- Competent contact person for our customers
- Technical consulting to customers
- Conducting technical trainings
- Sustainable and consistent automation of recurring tasks
- "Hands on" Troubleshooting
- „Containerization“ of existing solutions
- Monitoring and operation of customer infrastructure

SKILLS

- Advanced knowledge in at least three of the areas:
 - Public Cloud (Alibaba, AWS, Azure, GCE, Huawei Cloud, Open Telekom Cloud)
 - Private Cloud (OpenStack, VMware)
 - Infrastructure Management (ManageIQ, Morpheus, RedHat CloudForms, RedHat Satellite, Scalr, Spacewalk, SUSE Manager)
 - Container (Docker, Kubernetes)
 - Configuration Management (Ansible, Ansible Tower, AWX, Chef, Puppet, SaltStack, Terraform)
 - Continuous Integration/Delivery (Jenkins, Gitlab CI, DroneCI)
 - Source Code Management (Git, Github, Gitlab, Gerrit)
 - Enterprise Linux Distributions (RedHat, SUSE)

ADDITIONAL

- Certifications in the above areas are beneficial
- High willingness to learn
- Quick grasp of complex IT landscapes
- Willingness to travel
- Very good spoken and written English

WE ARE

Our aim is to provide our customers with the best possible advice, to accompany them through the challenges and upheavals of the cloud revolution and at the same time to always be at the cutting edge of technology and processes. We pass on this knowledge and experience to our customers in the form of enablements, coaching and workshops and accompany them as reliable partners on their way into and within the cloud. We, Cloudical, act manufacturer and platform independent.

Please send us your complete application including your salary expectations and your earliest possible starting date to the following address: tom@cloudical.io

Come to us!

Contact: career@cloudical.io
cloudical.io

We are testing clouds

Cloud computing offerings are changing rapidly. Even the offerings of the individual providers are regularly being further developed. This makes it almost impossible to keep track of things. We, Cloudical, would like to remedy this situation and gradually examine the offers and evaluate them from an objective point of view.

Our technicians have developed tests for this purpose. We test general information on onboarding, availability, SLAs, data centers, compute, storage, network, limitations, scaling, technologies, but also more internal information such as backup, security, image service, patch management, monitoring, CI/CD, as a Service offerings and, of course, the cost factor. This results in rankings and tables help customers to inform themselves independently and to find the right provider for themselves. But not only readers of the Report receive comprehensive, independent data, providers can also find out about their market, see where they stand and where their strengths lie in comparison. They can also identify their possible weaknesses and potentials, see possible pent-up demand or discover approaches for further specialization and improvement. And of course, they present themselves to interested readers and potential customers.

Currently we have tested the providers AWS, Azure, Google Cloud, IBM Cloud and the Open Telekom Cloud. On the following pages you will find the evaluations sorted by individual topics. You will find the complete evaluations here. We will gradually add more clouds, so that in the next issues only exemplary test evaluations will be shown, you will find the detailed tables online.

› the-report.cloud

If you have any suggestions for supplementing the questions, please write to us at: presse@cloudical.io.

In this issue we tested the IBM Cloud on the classic infrastructure which offers nearly the same services but on different conditions compared to the newest console of IBM Cloud (see: <https://www.ibm.com/cloud/blog/permanent-redirect-to-cloud-ibm-com-from-console-bluemix-net>). To show the differences between classic and newest we represent the actual results from the classic infrastructure and the deviant results from the new console of the last issue of The Cloud Report. This is only a showcase and not a real comparison due to the testing dates: we were only able to test on the new infrastructure in March and

on the classic infrastructure in May. But we think this shows the general differences within IBM Cloud provisions. As a customer you can choose between the both different infrastructures and our showcase give some insights where the both offerings have their different opportunities.

Note: Three virtual machines of different sizes are used in the evaluations:

Small means:

- › OS Ubuntu 16.04
- › 2vCPUs
- › 8GB RAM
- › min. 50GB HDD
- › Location: Germany, if not Western Europe, if not Europe

Medium means:

- › OS Ubuntu 16.04
- › 4vCPUs
- › 16GB RAM
- › min. 50GB HDD
- › Location: Germany, if not Western Europe, if not Europe

Large means:

- › OS Ubuntu 16.04
- › 8vCPUs
- › 32GB RAM
- › min. 50GB HDD
- › Location: Germany, if not Western Europe, if not Europe

And the winners are ...

As with every issue of this magazine, we also have our extensive tests in place. This time we changed some aspects, ensuring higher test quality and adding a new scoring scheme, developed by Michael Dombek, CTO of Cloudical.

Let's first of all address, how we ensured a higher test quality: During the last months, we incorporated several changes into our testing process. We added new members to our testing team, maxing out on a number of seven Cloud Engineers and Architects being involved at the peak. We also added new questions and ensured the quality of answers by introducing an expert-driven quality assurance process. Ultimately, we ensured a steady and ongoing process by introducing

a dedicated product manager to the team. More on this in the next issue of this magazine.

Besides changing the way we execute when conducting the tests, we also asked one of the most in-depth experts to develop a new scoring scheme, allowing for better comparability of outcomes and for introducing different weights to different aspects of each cloud environment. This scoring scheme will also be set up as an application later this year, allowing you to

decide, which aspects to emphasize upon and to create your very own scoring. So, stay tuned for this as well.

Winners in the categories

As always, our tests are broken down into categories. Each category has subcategories, that are weighted against each other.

Here is the overview of each category's winners, weightings and winners of the respective subcategories:

Category	Winner	Weighting	Winners of the subcategories
Storage	IBM Cloud	Features: 40% Provisioning Duration: 15% Costs: 5% Speed: 40%	Features: Tie (AWS, Google, IBM, OTC) Fastest Provisioning: OTC Lowest Costs: Google Cloud Fastest IOPS: IBM Cloud
Network	Google Cloud	Features: 40% Cost: 10% Performance: 50%	Features: Tie between every contender Cost: AWS Performance: Google Cloud
Backup	IBM Cloud	Features: 55% Cost: 30% Multiple Locations: 15%	Features: Azure Cost: IBM Locations: IBM
Compute	Google Cloud	Features: 25% Start Times: 10% Usage / Console / Formats: 5% Performance: 35% Cost: 10% Bare-Metal: 5% Patch-Management: 10%	Features: Tie between every contender Start Times: OTC Usage / Console / Formats: Azure Performance: AWS Cost: Google Cloud Bare-Metal: Tie between AWS, Google Cloud, IBM Cloud, OTC Patch-Management: Tie between Azure, Google, IBM Cloud
DBaaS	AWS	MySQL: 50% Postgres: 50%	MySQL: AWS Postgres: AWS
CaaS	Azure	Kubernetes / Features: 90% Mesosphere: 10%	Kubernetes / Features: Azure Mesosphere: Tie between Azure and Google Cloud
Images	Azure, Google Cloud, IBM, OTC	Custom Images: 30% Custom Licenses: 10% Image Build Service: 20% Cloud VM Images Creation: 20% Patch level: 20%	Custom Images: Tie between every contender Custom Licenses: Tie between every contender Image Build Service: Tie between Azure, Google, IBM, OTC Cloud VM Images Creation: Tie between every contender Patch Level: Tie between every contender

Security	Azure	Platform Pen Test: 4% Tools: 30% Features: 66%	Platform Pen Test: Azure Tools: Tie between AWS, Azure, Google Cloud and IBM Features: Tie between every contender
SaaS	IBM Cloud	Mobile Office: 20% Mobile App Services: 40% Application Environments: 40%	Mobile Office: Tie between Google Cloud and IBM Cloud Mobile App Services: Tie between Azure, Google and IBM Application Environments: Tie between AWS, Azure and IBM
Monitoring	AWS, Azure, Google Cloud, IBM	Dashboard: 10% VMs: 20% APPs: 20% Network: 20% LB: 8% Storage: 8% External Monitoring: 15%	Dashboard: Tie between every contender VMs: Tie between every contender APPs: Tie between every contender Network: Tie between every contender LB: Tie between every contender Storage: Tie between every contender External Monitoring: Tie between AWS, Azure, Google and IBM

And the overall winner is...

Just by taking the count of the winning offerings, one might find an overall winner.

But this would perhaps not reflect reality, since not every aspect is emphasized equally to each other.

Therefore, we weighted the categories against each other, added our internal scores for each category and ultimately found an overall winner.

Here are the details (higher numbers imply better results):

Category	Winner	Reason
Backup, Recovery and Availability	Azure	Azure is the winner in this case due to having more “strategies” for backups, e.g., full, differential, incremental, available.
Compute	OTC	Compute performance wise OTC is the winner, though looking at other factors such as costs and startup times Google Cloud Platform is a good choice, too.
Databases (DBaaS)	AWS	AWS through their pure variability in big data and databases is the winner here. Ranging from their in house Amazon DynamoDB to Amazon Kinesis, common relational databases such as MySQL and PostgreSQL, but also including non-relational databases Redis and Memcached, has such a vast amount to choose from, to get the perfect one for the project/application.
IaaS, PaaS and SaaS - Patch Management	Google Cloud Platform	Google Cloud Platform is the winner here, because they offer OSes for the most common use cases, such as Debian and Red Hat but also specialized ones like container optimized OS from Google and Red Hat Enterprise Linux for SAP applications.
Network	Azure and Google Cloud Platform	If you need high bandwidth in the same and different availability zone, go for Google Cloud Platform. When needing especially high bandwidth between different regions, go for Azure.
Security	Azure	Azure is knowingly running penetration tests against their cloud platform to ensure safety for users. Though the other cloud providers are also providing the “standard” set of features, such as Intrusion Detection and Prevention System. Only OTC is missing features in that regard.
Storage	AWS for object, IBM Cloud for block and file	Looking for fast block and filesystem storage? IBM Cloud has you covered with measured write speeds up to around 240 Megabyte per second. Price wise OTC is the winner, though together with Google Cloud Platform they are a good amount slower than their competition. Fast object storage? AWS is the “king” of S3 in regard to capabilities and price as they “invented” the protocol. In aspect to maximum object size AWS and Google Cloud Platform are on the same level.

Congratulations to the Google Cloud Team!

Overall, **Google Cloud** takes the crown this time with a slight advantage to Azure.

But notice, how close the contenders end up to each other – there is ob-

viously no such thing as a bad cloud offering, just different emphasizes on different aspects. So, it is still your job to find the right offering to your needs. And we’re happy to rightfully state this.

Compute

Questions	AWS	Azure	
Small VM: OS Ubuntu 16.04; 2vCPUs; 8GB RAM; min. 50GB HDD; Location: Germany, if unavailable: Western Europe, if unavailable: Europe	yes	yes	
Medium VM: OS Ubuntu 16.04; 4vCPUs; 16GB RAM; min. 50GB HDD; Location: Germany, if unavailable: Western Europe, if unavailable: Europe	yes	yes	
Large VM: OS Ubuntu 16.04; 8vCPUs; 32GB RAM; min. 50GB HDD; Location: Germany, if unavailable: Western Europe, if unavailable: Europe	yes	yes	
GPU support for the VM?	yes	yes	
AutoScaling for VM?	yes	yes	
Availability Zones (i.e Availability set) possible	yes	yes	
Startup-time (till time of availability) - Small - Medium - Large	50 sec 51 sec 52 sec	105 sec 108 sec 163 sec	
Count of steps until VM is created	7 steps	7 Steps	
RAM throughput (sysbench, Block size 1k) - Read - Write	813.27 MB/sec 785.38 MB/sec	4510.18 MB/sec 3517.93 MB/sec	
CPU speed (geekbench) - Small Single Core - Small Multi Core - Medium Single Core - Medium Multi Core - Large Single Core - Large Multi Core	3203 6104 3372 11696 3407 21761	3017 3541 3089 6559 2925 11200	
VM accessible via Console	no	yes	
Total cost of VM per month (732hrs) - Small - Medium - Large	€ 69.91 / \$ 78.48 € 139.80 / \$ 156.95 € 279.60 / \$ 313.89	€ 75.29 / \$ 99.28 € 150.58 / \$ 198.56 € 301.16 / \$ 350.40	
Supported disk formats / images	- OVA - VMDK - RAW - VHD/VHDX	- VHD - VMDK - VHDX - QCOW2 - RAW	
Are there any limitations per VM?	Amount CPUs: 128 RAM size: 1952 GB Disk size: 2048 GB	Amount CPUs: 128 RAM size: 3892 GB Disk size: 4096 GB Amount Disk: 64	
Can bare-metal servers be deployed via the cloud?	yes	no	
Which hypervisor is used?	- KVM - Xen	- Hyper-V	
Is autorecovery available	yes	yes	

	Google Cloud Platform	IBM Cloud	OTC
	yes	yes	yes
	yes	yes	yes
	yes	yes	yes
	yes	yes	no
	yes	yes	yes
	yes	yes	yes
	42 sec 44 sec 42 sec	385 sec 792 sec 903 sec	35 sec 40 sec 42 sec
	2 Steps	4 Steps	3 Steps
	4295.24 MB/sec 3514.16 MB/sec	640.75 MB/sec 613.16 MB/sec	4262.33 MB/sec 3469.08 MB/sec
	3002 3831 3058 7079 2968 13063	2426 4495 2391 8282 2442 15133	2910 5375 2928 9623 2957 16790
	yes	yes	yes
	€ 56.01 / \$ 62.54 € 112.14 / \$ 125.09 € 224.28 / \$ 250.18	€ 80.63 / \$ 89.94 € 154.54 / \$ 172.39 € 302.19 / \$ 337.10	€ 74.75 / \$ 83.86 € 150.28 / \$ 168.60 € 292.42 / \$ 328.20
	- VMDK - VDH - RAW	- VMDK - AKI - ARI - AMI - QCOW2 - RAW	- VMDK - QCOW2 - RAW - VHD - VHDX
	Amount CPUs: 160 RAM size: 3844 GB Disk size: 64 TB Amount Disk: 128	Amount CPUs: 64 RAM size: 512 GB Disk size: 12 TB	Amount CPUs: 60 RAM size: 940 GB
	yes	yes	yes
	- KVM	- PowerVM - VMware ESX Server - Xen - KVM - z/VM	- XEN - KVM
	no	yes	yes

Container as a Service

Questions	AWS	Azure	
Which technologies are being provided/supported?	<ul style="list-style-type: none"> - Kubernetes - Docker 	<ul style="list-style-type: none"> - Kubernetes - Mesosphere - Docker 	
Is a managed container service available?	yes (EKS)	yes (AKS)	
Can worker nodes be accessed directly by customers?	yes	yes	
Can master nodes be accessed directly by customers?	no	yes	
Which version of the technologies/Kubernetes is being offered?	1.12,1.11,1.10	1.12.6, 1.12.5, 1.11.8, 1.11.7, 1.10.13, 1.10.12, 1.9.11, 1.9.10	
How much time does it take to provide the container service for four nodes	10 min	<2min	
Costs <ul style="list-style-type: none"> - Managed service - 732hrs per month - small flavor - hosted in Frankfurt or Western Europe - Storage / IPs not included * Prices in USD have been converted to EUR 	€ 201.19 / \$ 227.88	€ 344.74 / \$ 383.82	
Shared or dedicated Container Engine Cluster?	dedicated	shared	
Do predefined StorageClasses exist in Kubernetes? Name - Provisioner	gp2 - kubernetes.io/aws-efs	default (default) - kubernetes.io/azure-disk managed-premium - kubernetes.io/azure-disk	
Limitations - What is the maximum cluster size?	max. 50	max. 100	
Do you have full access to all K8s resources (no RBAC restriction)?	yes	no	
Does the Container Service provide a Load-Balancer Service?	yes	yes	
Is a Storage Class useable?	yes	yes	

Backup, Recovery and Availability

Questions	AWS	Azure	
Are managed backups offered (Provider is responsible to take backups)	yes	yes	
Which types of backups are supported for VMs?	<ul style="list-style-type: none"> - Snapshots - Incremental Backups 	<ul style="list-style-type: none"> - Full Backups - Differential Backups - Incremental Backups - Snapshots 	
Where will the backup be stored?	<ul style="list-style-type: none"> - Amazon S3 - Amazon Glacier - Different datacenter - Storage-Cluster 	<ul style="list-style-type: none"> - Recovery Services Vault - Different Datacenter 	
Can backups be scheduled?	yes	yes	
Usage costs per month <ul style="list-style-type: none"> - 500 GB Backup Storage - Western Europe 	€ 22.13 / \$ 25.00	€ 10.12 / \$ 12.00	
Is a managed Backup-Service for a VM provided?	no	yes	

	Google Cloud Platform	IBM Cloud	OTC
	<ul style="list-style-type: none"> - Kubernetes - Mesosphere 	<ul style="list-style-type: none"> - Kubernetes 	<ul style="list-style-type: none"> - Kubernetes - Docker - Cloud Container Engine
	yes	yes	yes
	yes	yes	yes
	yes	yes	yes
	1.12.5-gke.5, 1.11.7-gke.6, 1.11.7-gke.4, 1.10.12-gke.7	1.14.1, 1.13.6, 1.12.8, 1.11.3, 1.10.8, 1.9.8	1.9.2-r2, 1.11.3
	<3min	15 min	15 min
	€ 222.53 / \$ 247.78	€ 247.68 / \$ 275.78	€ 285.69 / \$ 318.07
	shared	dedicated	shared
	standard (default) - kubernetes.io/gce-pd	ibmc-file-bronze (default) - ibm.io/ibmc-file ibmc-file-custom - ibm.io/ibmc-file ibmc-file-gold - ibm.io/ibmc-file ibmc-file-retain-bronze - ibm.io/ibmc-file ibmc-file-retain-custom - ibm.io/ibmc-file ibmc-file-retain-gold - ibm.io/ibmc-file ibmc-file-retain-silver - ibm.io/ibmc-file ibmc-file-silver - ibm.io/ibmc-file	cce-evs - cce-evs cce-sfs - cce-sfs
	max. 5000	n/a	Max nodes/cluster: 1000
	yes	no	yes
	yes	yes	yes
	yes	yes	yes

	Google Cloud Platform	IBM Cloud	OTC
	no	yes	yes
	<ul style="list-style-type: none"> - Snapshots - Incremental Backups 	<ul style="list-style-type: none"> - Snapshot - Full Backups - Incremental Backups 	<ul style="list-style-type: none"> - Snapshot - Full Backups - Incremental Backups
	<ul style="list-style-type: none"> - Google Cloud Storage - Storage Cluster 	<ul style="list-style-type: none"> - Evault - dedicated backup space - IBM Cloud Backup, - IBM Object Storage archive 	Different Data-Centers
	yes	yes	yes
	€ 4.46 / \$ 5.00	€ 0.88 / \$ 1.00	€ 5.00 / \$ 5.63
	no	yes	yes

Databases (DBaaS)

Questions	AWS	Azure
Which DB engines are offered?	<p>Relational DB</p> <ul style="list-style-type: none"> - MySQL - PostgreSQL - MariaDB - Oracle - Microsoft SQL Server - Amazon Aurora <p>Non-Relational DB</p> <ul style="list-style-type: none"> - Amazon DynamoDB - Amazon ElastiCache - Amazon Neptune - Redis - MemCached <p>Data Warehouse / Big Data</p> <ul style="list-style-type: none"> - Amazon Redshift - Amazon Athena - Amazon EMR (Hadoop, Spark, HBase, Presto, etc.) - Amazon Kinesis - Amazon Elasticsearch Service - Amazon Quicksight 	<p>Relational DB</p> <ul style="list-style-type: none"> - Azure SQL Database - Azure Database for MySQL - Azure Database for PostgreSQL - Azure Database for Maria DB - Microsoft SQL Server <p>Non-Relational DB</p> <ul style="list-style-type: none"> - Azure Cosmos DB - Azure Table Storage - Redis <p>Data Warehouse / Big Data</p> <ul style="list-style-type: none"> - SQL Data Warehouse - HDInsight (Hadoop, Spark, Hive, LLAP, Kafka, Storm, R) - Azure Databricks (Spark) - Azure Data Factory - Azure Stream Analytics
Performance of MySQL (MySQL Sysbench, table-size (row data): 1000000, Threads: 16)	<p>- Read</p> <p>- Write</p> <p>- Read / Write</p>	<p>Transactions: 50246 (837.13 / sec)</p> <p>Transactions: 67802 (1129.82 / sec)</p> <p>Transactions: 30248 (503.90 / sec)</p>
Provisioning time for a MySQL instance	235 sec	198 sec
Performance of PostgreSQL	<p>Transactions: 1089280 (18152.17 / sec)</p> <p>Transactions: 415225 (6919.63 / sec)</p> <p>Transactions: 35151 (585.64 / sec)</p>	<p>Transactions: 337927 (5630.91 / sec)</p> <p>Transactions: 84989 (1415.90 / sec)</p> <p>Transactions: 12088 (201.15 / sec)</p>
Provisioning time for a PostgreSQL instance	262 sec	132 sec
Supported DB Versions	<ul style="list-style-type: none"> - MySQL 8.0,5.7, 5.6, 5.5 - MariaDB 10.3,10.2,10.1,10.0 - Microsoft SQL Server 2017 RTM, 2016 SP1, 2014 SP2, 2012 SP4, 2008 R2 SP3 - Oracle 12c (12.1.0.2, 12.1.0.1), Oracle 11g (11.2.0.4, 11.2.0.3, 11.2.0.2) - PostgreSQL 11.2, 11.1,10.6,10.5, 10.4, 10.3, 10.1, 9.6.x, 9.5.x, 9.4.x, 9.3.x,9.2.x - Amazon Aurora - compatible with MySQL 5.6.10a 	<ul style="list-style-type: none"> - MySQL 5.7, 5.6 - MariaDB 10.2 - Azure SQL Database: Microsoft SQL Server 2017 - Microsoft SQL Server 2017, 2016 SP1, 2014 SP2, 2012 SP4, 2008 R2 SP3 - PostgreSQL 10.3, 9.6.x, 9.5.x - Azure Cosmos DB
Troubleshooting as a Service	yes	yes
- Rollback	yes	yes
- Support		
Total price for the database per month	€ 114.13 / \$ 128.13	€ 142.29 / \$ 159.50
- MySQL		
- 2 vCores		
- 100 GB Storage		
- Frankfurt / Western Europe		
- 100% active per month		
- No dedicated backup		
Total price for the database per month	€ 121.64 / \$ 136.34	€ 142.29 / \$ 159.50
- PostgreSQL		
- 2 vCores		
- 100 GB Storage		
- Frankfurt / Western Europe		
- 100% active per month		
- No dedicated backup		

Google Cloud Platform	IBM Cloud	OTC Cloud
<p>Relational DB</p> <ul style="list-style-type: none"> - PostgreSQL - MySQL - Google Cloud Spanner <p>Non-Relational DB</p> <ul style="list-style-type: none"> - Google Cloud Datastore - Google Cloud BigTable <p>Data Warehouse / Big Data</p> <ul style="list-style-type: none"> - Google Cloud BigQuery - Google Cloud Dataflow - Google Cloud Dataproc (Hadoop / Spark) - Google Cloud Datalab - Google Cloud Dataprep 	<p>Relational DB</p> <ul style="list-style-type: none"> - Db2 on Cloud - PostgreSQL - MySQL <p>Non-Relational DB</p> <ul style="list-style-type: none"> - Cloudant - MongoDB - ScyllaDB - Redis - JanusGraph - etcd - Elasticsearch <p>Data Warehouse / Big Data</p> <ul style="list-style-type: none"> - Db2 Warehouse on Cloud 	<p>Relational DB</p> <ul style="list-style-type: none"> - PostgreSQL - MySQL - Microsoft SQL Server <p>Non-Relational DB</p> <ul style="list-style-type: none"> - MongoDB - Redis
<p>Transactions: 7366 (122.53 / sec) Transactions: 18240 (303.70 / sec) Transactions: 5771 (95.95 / sec)</p>	<p>Transactions: 31026 (516.82 / sec) Transactions: 64777 (1079.35 / sec) Transactions: 22708 (378.18 / sec)</p>	<p>Transactions: 35363 (589.22 / sec) Transactions: 91616 (1526.67 / sec) Transactions: 30347 (505.58 / sec)</p>
194 sec	120 sec	468 sec
<p>Transactions: 117958 (1965.67 / sec) Transactions: 97248 (1620.43 / sec) Transactions: 5628 (93.45 / sec)</p>	<p>Transactions: 837378 (13953.69 / sec) Transactions: 609573 (10158.00 / sec) Transactions: 33815 (563.26 / sec)</p>	<p>Transactions: 1267208 (21116.66 / sec) Transactions: 329666 (5492.69 / sec) Transactions: 23728 (393.46 / sec)</p>
159 sec	195 sec	307 sec
<ul style="list-style-type: none"> - MySQL 5.7, 5.6 - PostgreSQL 9.6.x 	<ul style="list-style-type: none"> - Db2-ge - PostgreSQL 9.6.12, 9.6.10, 9.6.9, 9.5.14, 9.5.13, 9.4.19, 9.4.18 - MySQL 5.7.22 - Cloudant-h7 - MongoDB 3.4.10, 3.2.18, 3.2.11, 3.2.10 - ScyllaDB 2.0.3 - Redis 4.0.10, 3.2.12 - JanusGraph 0.1.1 beta - etcd 3.3.3, 3.2.18 - Elasticsearch 6.2.2, 5.6.9 - Db2 Warehouse-ef 	<p>PostgreSQL 10.0, 9.6.5, 9.6.3, 9.5.5 MySQL 5.7.20, 5.7.17, 5.6.35, 5.6.34, 5.6.33, 5.6.30 Microsoft SQL Server 2016 EE, 2016 SE, 2014 SE</p>
yes yes	yes yes	yes yes
€ 121.43 / \$ 138.75	N/A	€ 298.40 / \$ 335.04
€ 124.21 / \$ 141.81	€ 103.04 / \$ 136.00	€ 312.80 / \$ 350.85

Limitations: How many simultaneous requests to the DB? How much RAM? How many users?	MySQL: - max Connections: 2540 PostgreSQL: - max Connections: 5696	MySQL: - max Connections: 10000 PostgreSQL: - max Connections: 1900	
How does backup/restore work?	Backups: - Automatic Backups. Restore: - Point-in-time restore	Backups: - Automatic Backups. Restore: - Point-in-time restore - Geo-restore	

IaaS, PaaS and SaaS – Patch-Management

Questions	AWS	Azure	
Does the cloud provide a managed patch service?	no	yes (Azure Automation)	
Which operating systems are supported?	Linux: - Red Hat Enterprise Linux (RHEL) 7.0 - 7.4, 6.5 - 6.9 - SUSE Linux Enterprise Server (SLES) 12 - Amazon Linux 2015.03 - 2018.03, 2012.03 - 2017.03 - CentOS 7.1, 6.5 and later - Raspbian Jessie - Raspbian Stretch - Ubuntu Server 18.04, 16.04, 14.04 Windows: - Windows Server 2008 - Windows Server 2012 - Windows Server 2016 including R2 Versions	Linux: - CentOS 6 (x86/x64), 7 (x64) - Red Hat Enterprise 6 (x86/x64), 7 (x64) - SUSE Linux Enterprise Server 11 (x86/x64), 12 (x64) - Ubuntu 14.04, 16.04 (x86/x64) Windows: - Windows Server 2008 - Windows Server 2008 R2 RTM - Windows Server 2008 R2 SP1 and later	
Is the operating system from the deployed VM at a current patch level?	yes	yes	
What is the current available patch level in our sample VM? - Ubuntu 16.04 LTS with latest patches applied	04	04	

Logging as a Service

Questions	AWS	Azure	
Does the cloud platform provide a Logging as a Service functionality?	yes	yes	
Is the data stored in encrypted form?	yes	yes	
Which logging technology is used?	- AWS Cloudwatch - AWS Cloudtrail - AWS VPC flow logs - Amazon Cloudfront access logs - Amazon S3 access logs	- Activity logs - Activity diagnostics logs - Azure AD Reporting - Virtual machines and cloud services - Azure Storage Analytics - Network Security Group (NSG) flow logs - Application insight	

MySQL: - max Connections: 4000 PostgreSQL: - max Connections: 1000	MySQL: - max Connections: 151 PostgreSQL: - max Connections: 1000	MySQL: - max Connections: 151 PostgreSQL: - max Connections: unlimited
Backups: - Automatic Backups. Restore: - On-demand	Backups: - Automatic Backups. Restore: - On-demand	Backups: - Automatic Backups. Restore: - Point-in-time restore

Google Cloud Platform	IBM Cloud	OTC
yes (Google App Engine)	yes (IBM BigFix Patch Management)	no
Linux: - Centos 7, 6 - Container-Optimized OS from Google cos-69-lts, cos-stable, cos-beta, cos-dev - CoreOS coreos-stable, coreos-beta, coreos-alpha - Debian 9 - Red Hat Enterprise Linux (RHEL) 8, 7, 6 - RHEL for SAP, rhel-7-6-sap-ha, rhel-7-4-sap - SUSE Enterprise Linux Server (SLES) 15, 12 - SLES for SAP sles-15-sap, sles-12-sp4-sap, sles-12-sp3-sap, sles-12-sp2-sap, sles-12-sp1-sap - Ubuntu 19.04, 18.10, 16.04, 14.04 Windows: - Windows Server 2019, 2016, 2012 R2, 2008 R2 - Windows Server Core 2019 - Windows Server Core 2019 for containers	Linux: - CentOS 7, 6 - RedHat Enterprise 7, 6 - SUSE Linux Enterprise Server 12, 11 - Ubuntu Minimal 18.04, 16.04, 14.04 Windows: - Windows Server 2016, 2012 R2, 2008 R2	Linux: - openSUSE 42.x, 15.x - CentOS 6.x, 7.x - Debian 9.x, 8.x - Fedora 29, 28, 27, 26 - EulerOS 2.x - Ubuntu 18.04, 16.04, 14.04 - RedHat Enterprise Linux 7.x, 6.x - SUSE 15.x, 12.x, 11.x - Oracle Linux 7.x, 6.x Windows: - Windows Server 2019, 2016, 2012 R2, 2012, 2008
yes	yes	yes
04	04.2	04.06

Google Cloud Platform	IBM Cloud	OTC
yes	yes	yes
yes	yes	yes
- Stackdriver Logging	- IBM Log Analysis with LogDNA - Bluemix UI - Cloud Foundry Line Interface(CLI) - External logging	- cloud trace

Image Service

Questions	AWS	Azure
Which operating systems are offered by the provider with which versions?	<p>Windows:</p> <ul style="list-style-type: none"> - Windows Server 2008, 2012, 2016 Build 1809, 2019 <p>Linux:</p> <ul style="list-style-type: none"> - Amazon Linux 2, 2018.03 - CentOS 6.x, 7.x - Debian 8.x, 9.x - Fedora 26, 27, 28, 29 - Ubuntu 14.04.x, 16.04.x, 18.04.x - SUSE Enterprise Linux 12, 15 - Oracle Linux 6.8, 7.x, - Red Enterprise Linux 6.8, 7.3 	<p>Windows:</p> <ul style="list-style-type: none"> - Windows Server 2008 R2 SP1, 2008 SP2, 2012 R2, 2016, 2019 - Windows Server 2016 Build 1709, 1803, 1809 - Windows 10 <p>Linux:</p> <ul style="list-style-type: none"> - CentOS-based 6.9, 7.4 - ClearLinux - Container Linux - Debian 8, 9 - Red Hat Enterprise Linux 7.x - SLES 11SP4, 12SP3 - Ubuntu 14.04, 16.04, 18.04
Can own images be uploaded?	yes	yes
Can existing licenses be used to minimize costs?	yes	yes
Is there an image build service?	<p>no</p> <p>Supported Formats:</p> <ul style="list-style-type: none"> - OVA File - VMDK - VHD - RAW 	<p>yes</p> <p>Supported formats:</p> <ul style="list-style-type: none"> - VHD - VMDK - VHDX - QCOW2 - RAW
Can images be created from existing cloud instances?	yes	yes
Are different patch levels of images available?	yes	yes

Network

Questions	AWS	Azure
Is network monitoring available?	yes	yes
Is a Content Delivery Network (CDN) available?	yes	yes
<p>Sample Measurements</p> <p>1) Same AZ</p> <p>2) Different AZ</p> <p>3) Different Region</p>	<p>Iperf Result:</p> <p>1)</p> <p>TCP:</p> <p>Bandwidth</p> <p>Sender: 964 Mbit/sec</p> <p>Receiver: 963 Mbit/sec</p> <p>UDP:</p> <p>Bandwidth: 1.94 Gbit/sec</p> <p>2)</p> <p>TCP:</p> <p>Bandwidth</p> <p>Sender: 934 Mbit/sec</p> <p>Receiver: 933 Mbit/sec</p> <p>UDP:</p> <p>Bandwidth: 2.06 Gbit/sec</p> <p>3)</p> <p>TCP:</p> <p>Bandwidth</p> <p>Sender: 162 Mbit/sec</p> <p>Receiver: 161 Mbit/sec</p> <p>UDP:</p> <p>Bandwidth: 868 Mbit/sec</p>	<p>Iperf Result:</p> <p>1)</p> <p>TCP:</p> <p>Bandwidth</p> <p>Sender: 876 Mbit/sec</p> <p>Receiver: 874 Mbit/sec</p> <p>UDP:</p> <p>Bandwidth: 839 Mbit/sec</p> <p>2)</p> <p>TCP:</p> <p>Bandwidth</p> <p>Sender: 910 Mbit/sec</p> <p>Receiver: 908 Mbit/sec</p> <p>UDP:</p> <p>Bandwidth: 927 Mbit/sec</p> <p>3)</p> <p>TCP:</p> <p>Bandwidth</p> <p>Sender: 813 Mbit/sec</p> <p>Receiver: 813 Mbit/sec</p> <p>UDP:</p> <p>Bandwidth: 933 Mbit/sec</p>

Google Cloud Platform	IBM Cloud	OTC
<p>Windows:</p> <ul style="list-style-type: none"> - Windows Server 2008, 2012, 2016, 2019 - Windows Server 2016 Build 1709, 1803, 1809 <p>Linux:</p> <ul style="list-style-type: none"> - CentOS 6.x, 7.x - Container-optimised OS dev, beta, stable, 69-lts - CoreOS alpha, beta, stable - Debian 9.x - Ubuntu 14.04.x, 16.04.x, 17.04.x, 18.04.x - SLES 12, 15 - SLES for SAP 12-sp2-sap, 12-sp3-sap - Oracle Linux 6.8, 7.x, - RedHat Enterprise Linux 6, 7 - RHEL for SAP 7-4-sap, 7-6-sap-ha 	<p>Windows:</p> <ul style="list-style-type: none"> - Windows Server 2012, 2016 <p>Linux:</p> <ul style="list-style-type: none"> - CentOS- Minimal 6.X, 7.x - CentOS-LAMP 6.X, 7.X - Debian Minimal Stable 8.X, 9.x - Debian LAMP Stable 8.X - Red Hat Minimal 6.x, 7.x - Red Hat LAMP 6.x, 7.x - Ubuntu Minimal 14.04, 16.04, 18.04 - Ubuntu LAMP 14.04, 16.04, 18.04 	<p>Windows:</p> <ul style="list-style-type: none"> - Windows Server 2008, 2012, 2016, 2019 <p>Linux:</p> <ul style="list-style-type: none"> - openSUSE 15.x, 42.x - CentOS 6.x, 7.x - Debian 8.x, 9.x - Fedora 26, 27, 28, 29 - EulerOS 2.x - Ubuntu 14.04.x, 16.04.x, 18.04.x - SUSE Enterprise Linux 11, 12, 15 - SUSE SAP 12 - Oracle Linux 6.8, 7.2 - Red Enterprise Linux 6.8, 7.3
yes	yes	yes
yes	yes	yes
<p>yes</p> <p>Supported Formats:</p> <ul style="list-style-type: none"> - VMDK - VHD - VDI - VPC - QCOW2 - RAW 	<p>yes</p> <p>Supported formats:</p> <ul style="list-style-type: none"> - VHD - VMDK - QCOW2 - AKI - ARI - AMI 	<p>yes</p> <p>Supported Formats:</p> <ul style="list-style-type: none"> - VHD - ZVHD - VMDK - VHDX - QCOW - QCOW2 - RAW - ZVHD2 - VDI - QED
yes	yes	yes
yes	yes	yes

Google Cloud Platform	IBM Cloud	OTC
yes	yes	yes
yes	yes	yes
<p>Iperf Result:</p> <p>1)</p> <p>TCP:</p> <p>Bandwidth</p> <p>Sender: 3.65 Gbit/sec</p> <p>Receiver: 3.65 Gbit/sec</p> <p>UDP:</p> <p>Bandwidth: 3.79 Gbit/sec</p> <p>2)</p> <p>TCP:</p> <p>Bandwidth</p> <p>Sender: 3.62 Gbit/sec</p> <p>Receiver: 3.62 Gbit/sec</p> <p>UDP:</p> <p>Bandwidth: 3.79 Gbit/sec</p> <p>3)</p> <p>TCP:</p> <p>Bandwidth</p> <p>Sender: 302 Mbit/sec</p> <p>Receiver: 302 Mbit/sec</p> <p>UDP:</p> <p>Bandwidth: 3.59 Gbit/sec</p>	<p>Iperf Result:</p> <p>1)</p> <p>TCP:</p> <p>Bandwidth</p> <p>Sender: 102 Mbit/sec</p> <p>Receiver: 100 Mbit/sec</p> <p>UDP:</p> <p>Bandwidth: 99 Mbit/sec</p> <p>2)</p> <p>TCP:</p> <p>Bandwidth</p> <p>Sender: 102 Mbit/sec</p> <p>Receiver: 99.8 Mbit/sec</p> <p>UDP:</p> <p>Bandwidth: 98.9 Mbit/sec</p> <p>3)</p> <p>TCP:</p> <p>Bandwidth</p> <p>Sender: 102 Mbit/sec</p> <p>Receiver: 99.8 Mbit/sec</p> <p>UDP:</p> <p>Bandwidth: 99 Mbit/sec</p>	<p>Iperf Result:</p> <p>1)</p> <p>TCP:</p> <p>Bandwidth</p> <p>Sender: 83.3 Mbit/sec</p> <p>Receiver: 81.9 Mbit/sec</p> <p>UDP:</p> <p>Bandwidth: 2.78 Gbit/sec</p> <p>2)</p> <p>TCP:</p> <p>Bandwidth</p> <p>Sender: 80.8 Mbit/sec</p> <p>Receiver: 80.1 Mbit/sec</p> <p>UDP:</p> <p>Bandwidth: 1.53 Gbits/sec</p> <p>3) N/A</p>

Public IPs – Public IPs for VMs? – Available kinds of public IPs for VMs – Public IPs for Load Balancers? – Available kinds of public IPs for Load Balancers	yes floating / static yes static	yes floating / static yes static	
Is a dedicated network connection from data-center to public cloud possible?	yes (AWS Direct Connect)	yes (Azure Express Route)	
Network Security features (Network Traffic analysis, Network Security Groups)	– AWS Web Application Firewall – Network security groups – Network Traffic analysis	– Azure Firewall – Azure Front Door – Azure Network Watcher – Azure Security Center – Azure DDoS protection – Network access control – Network layer control – Network security rules (NSGs)	
VPN as a Service	yes	yes	
Traffic costs per GB	€ 0.13 / \$ 0.15	€ 0.009 / \$ 0.01	

Monitoring

Questions	AWS	Azure	
Dashboard	yes	yes	
Which cloud resources will be monitored?			
VMs	yes	yes	
Apps	yes	yes	
Network	yes	yes	
Load Balancer	yes	yes	
Storage	yes	yes	
Connection/Usage of external monitoring solutions	yes	yes	
Costs per month	€ 49.49 / \$ 55.50	€ 22.31 / \$ 26.45	

Security

Questions	AWS	Azure	
– Integration to a SIEM possible? (Security Information and Event Management)	yes	yes	
– Security Groups	yes	yes	
– Disk Encryption	yes	yes	
– Network Traffic Analyse	yes	yes	
Protection against Denial of Service Attacks	yes	yes	
Firewall - Does the cloud provider provide additional integrated security features i.e. a Next Generation Firewall?	yes	yes	
Does the cloud provider keep an eye on current threats and take action?	yes	yes	
Does the cloud provider support additional integrated security features for cloud resources using 3rd party tools:			
– IDS (Intrusion Detection System)	yes	yes	
– IPS (Intrusion Prevention System)	yes	yes	
– ATP (Advanced Threat Protection)	yes	yes	
Does the provider carry out regular penetration tests against the platform?	No	yes	

	yes floating / static yes static	yes floating/static yes static	yes static yes static
	yes (Google Cloud Interconnect)	yes	yes (Direct Connect - MPLS)
	- Firewall - Network security groups - Network Traffic analysis	- Network Security Groups - Firewalls (Multi VLAN, Single VLAN and Web App) - DDOS mitigation	- Network Security Groups - Firewalls (Multi VLAN, Single VLAN and Web App)
	yes	yes	yes
	€ 0.073 / \$ 0.082	€ 0.078 / \$ 0.087	€ 0.06 / \$ 0.067

	Google Cloud Platform	IBM Cloud	OTC
	yes	yes	yes
	yes yes yes yes yes	yes yes yes yes yes	yes yes yes yes yes
	yes	yes	no
	n/a	n/a	included

	Google Cloud Platform	IBM Cloud	OTC
	yes yes yes yes	yes yes yes yes	yes yes yes yes
	yes	yes	yes
	yes	yes	yes
	yes	yes	yes
	yes yes yes	yes yes yes	no no no
	no	no	no

Software as a Service

Questions	AWS	Azure
Is a mobile office suite offered? Is it deeply integrated with other services?	no n/a	yes no
Managed App Services	<ul style="list-style-type: none"> - AWS Step Functions - Amazon API Gateway - Amazon Elastic Transcoder - Amazon SWF 	<ul style="list-style-type: none"> - Azure Stack - Security and Compliance - Backups and Archives - Disaster Recovery - Cosmos DB - Networks - Active Directory Services - Development and Testing Services - Mobile Services
Mobile App Services <ul style="list-style-type: none"> - Push Notifications - User Management - NoSQL-Datenbase - File Storage - Messaging - Social Networks 	AWS Mobile <ul style="list-style-type: none"> yes yes yes yes yes no 	Azure Mobile App Service <ul style="list-style-type: none"> yes yes yes yes yes yes
Application Environments <ul style="list-style-type: none"> - Websites - Microservices - Messaging - Serverless 	<ul style="list-style-type: none"> yes (AWS Lightsail) yes (AWS Elastic Beanstalk) yes (AWS SQS) yes (AWS Lambda) 	<ul style="list-style-type: none"> yes (Azure Web Sites) yes (Azure Service Fabric) yes (Azure Service Bus) yes (Azure Functions)
Rollback to a previous application version?	yes	yes

Storage

Questions	AWS	Azure
Which kinds of storage are available? <ul style="list-style-type: none"> - Object / Blob Storage - File Storage - Block Storage 	<ul style="list-style-type: none"> yes (S3 / Glacier) yes (EFS) yes (EBS) 	<ul style="list-style-type: none"> yes (Azure Blob Storage) yes (Azure Disk Storage) yes (Azure Files)
Block - Different tier-classes? SATA, SSD, SAS	yes	yes
Which objects storage-engines are offered?	Amazon S3	Azure Blob Storage
File - Accessing file storage via (cluster) file system.	- EFS	<ul style="list-style-type: none"> - GlusterFS - BeeGFS - Luster
Storage capacity limits	Overall size: Unlimited 5 TB per S3 object	Overall size: 500 TB per Storage Account 200 Storage Accounts per Subscriptions
Duration of provisioning?	9 sec	47 sec
Throughput IOPS (only Block- and File-Storage)	<ul style="list-style-type: none"> - Random read: bw = 24.52 MB/s, iops = 3065 - Random write: bw = 128.93 MB/s, iops = 2014 - Random Read and write: <ul style="list-style-type: none"> - read : bw = 44.17 MB/s, iops = 2760 - write: bw = 5.03 MB/s, iops = 314 - Sequential read: bw = 24.55 MB/s, iops = 3068 - Sequential write: bw = 98.08 MB/s, iops = 3064 	<ul style="list-style-type: none"> - Random read: bw = 1.95 MB/s, iops = 244 - Random write: bw = 15.58 MB/s, iops = 243 - Random read and write: <ul style="list-style-type: none"> - read: bw = 3.50 MB/s, iops = 6574; - write: bw = 0.41 MB/s, iops = 25; - Sequential read: bw = 1.95 MB/s, iops = 244 - Sequential write: bw = 7.81 MB/s, iops = 243
Costs per month <ul style="list-style-type: none"> - total price for 50 GB Disk which is mounted to the VM 	€ 5.29 / \$ 5.95	€ 7.65 / \$ 8.58

	Google Cloud Platform	IBM Cloud	OTC
	yes yes	yes yes	no no
	- Google App Engine - GSuite	- IBM Cloud API - CI/CD - Database as a Service - Network as a Service - Function as a Service, - Webserver as a Service, - Monitoring as a Service, - Backup as a service, - AI as a service	- BigData MapReduce - Database as a service, - Workspace management, - Backup as a service, - Network as a service, - Monitoring as a service
	Google Firebase / App Engine yes yes yes yes yes yes	IBM Mobile Foundation yes yes yes yes yes yes	yes yes no yes yes no
	no yes (App Engine) yes (Cloud Pub/Sub) yes (Cloud Functions)	yes yes yes (IBM message Hub) yes (Cloud Functions)	yes yes yes no
	yes	yes	no

	Google Cloud Platform	IBM Cloud	OTC
	yes (Google Cloud Storage) yes (Google Drive / Persistent Disk) yes (Google Persistent Disk)	yes (IBM Cloud Object Storage) yes (IBM Cloud file storage) yes (IBM Cloud block storage)	yes (Object Storage Service) yes (Scalable File Service) yes (Elastic Volume Service)
	yes	yes	yes
	Buckets (like S3)	- S3 - Swift	- S3 - OpenStack Swift
	- Google Cloud Storage FUSE - Beta: Google Cloud Filestore	- NFS	- NFS
	Overall size: Unlimited 5 TB per individual object	Overall size: Unlimited 25 GB per month of object storage(Unlimited for standard plan)	50 TB object storage
	22 sec	24 sec	8 sec
	- Random read: bw = 5.86 MB/s, iops = 732 - Random write: bw = 56.29 MB/s, iops = 879 - Random read and write: - read: bw = 7.91 MB/s, iops = 494 - write: bw = 0.93 MB/s, iops = 56 - Sequential read: bw = 88.29 MB/s, iops = 11036 - Sequential write: bw = 122.74 MB/s, iops = 3835	- Random Read: bw = 188.67 MB/s, iops = 23583 - Random Write: bw = 233.98 MB/s, iops = 3655 - Random Read and write: - read: bw = 19.08 MB/s, iops = 11923 - write: bw = 21.74 MB/s, iops = 1358 Sequential Read: bw = 127.40 MB/s, iops = 15924 Sequential Write: bw = 115.70 MB/s, iops = 3615	- Random Read: bw = 8.04 MB/s, iops = 1005 - Random Write: bw = 65.03 MB/s, iops = 1016 - Random Read and write: - read: bw = 14.46 MB/s, iops = 903 - write: bw = 1.65 MB/s, iops = 102 Sequential Read: bw = 11.06 MB/s, iops = 1383 Sequential Write: bw = 58.88 MB/s, iops = 1840
	€ 1.14 / \$ 1.28	€ 8.82 / \$ 9.89	€ 2.30 / \$ 2.58

Comparisons within IBM Cloud Provisions

In this issue we tested the IBM Cloud on the classic infrastructure which offers nearly the same services but on different conditions compared to the newest console of IBM Cloud (see: <https://www.ibm.com/cloud/blog/permanent-redirect-to-cloud-ibm-com-from-console-bluemix-net>). To show the differences between classic and newest we represent the actual results from the classic infrastructure and the deviant results from the new console of the last issue of The Cloud Report. This is only a showcase and not a real comparison due to the testing dates: we were only able to test on the new infrastructure in March and on the classic infrastructure in May. But we think this shows the general differences within IBM Cloud provisions. As a customer you can choose between the both different infrastructures and our showcase give some insides where the both offerings have their different opportunities.

Compute

Questions	IBM Cloud Classic	IBM Cloud latest
Small VM: OS Ubuntu 16.04; 2vCPUs; 8GB RAM; min. 50GB HDD; Location: Germany, if unavailable: Western Europe, if unavailable: Europe	yes	yes
Medium VM: OS Ubuntu 16.04; 4vCPUs; 16GB RAM; min. 50GB HDD; Location: Germany, if unavailable: Western Europe, if unavailable: Europe	yes	yes
Large VM: OS Ubuntu 16.04; 8vCPUs; 32GB RAM; min. 50GB HDD; Location: Germany, if unavailable: Western Europe, if unavailable: Europe	yes	yes
GPU support for the VM?	yes	yes
AutoScaling for VM?	yes	yes
Availability Zones (i.e Availability set) possible	yes	yes
Startup-time (till time of availability)		
- Small	385 sec	125 sec
- Medium	792 sec	160 sec
- Large	903 sec	340 sec
Count of steps until VM is created	4 Steps	4 Steps
RAM throughput (sysbench, Block size 1k)		
- Read	640.75 MB/sec	606.40 MB/sec
- Write	613.16 MB/sec	583.99 MB/sec
CPU speed (geekbench)		
- Small Single Core	2426	2590
- Small Multi Core	4495	4952
- Medium Single Core	2391	2616
- Medium Multi Core	8282	8659
- Large Single Core	2442	2567
- Large Multi Core	15133	15316
VM accessible via Console	yes	yes
Total cost of VM per month (732hrs)		
- Small	€ 80.63 / \$ 89.94	€ 78.77 / \$ 88.37
- Medium	€ 154.54 / \$ 172.39	€ 141.19 / \$ 158.39
- Large	€ 302.19 / \$ 337.10	€ 301.21 / \$ 337.91

Supported disk formats / images	<ul style="list-style-type: none"> - VMDK - AKI - ARI - AMI - QCOW2 - RAW 	<ul style="list-style-type: none"> - VMDK - AKI - ARI - AMI - QCOW2 - RAW
Are there any limitations per VM?	Amount CPUs: 64 RAM size: 512 GB Disk size: 12 TB	Amount CPUs: 64 RAM size: 512 GB Disk size: 12 TB
Can bare-metal servers be deployed via the cloud?	yes	yes
Which hypervisor is used?	<ul style="list-style-type: none"> - PowerVM - VMware ESX Server - Xen - KVM - z/VM 	<ul style="list-style-type: none"> - PowerVM - VMware ESX Server - Xen - KVM - z/VM
Is autorecovery available	yes	n/a

Network

Questions	IBM Cloud Classic	IBM Cloud latest
Is network monitoring available?	yes	yes
Is a Content Delivery Network (CDN) available?	yes	yes
Sample Measurements 1) Same AZ 2) Different AZ 3) Different Region	Iperf Result: 1) TCP: Bandwidth Sender: 102 Mbit/sec Receiver: 100 Mbit/sec UDP: Bandwidth: 99 Mbit/sec 2) TCP: Bandwidth Sender: 102 Mbit/sec Receiver: 99.8 Mbit/sec UDP: Bandwidth: 98.9 Mbit/sec 3) TCP: Bandwidth Sender: 102 Mbit/sec Receiver: 99.8 Mbit/sec UDP: Bandwidth: 99 Mbit/sec	Iperf Result: 1) TCP: Bandwidth Sender: 102 Mbit/sec Receiver: 99.9 Mbit/sec UDP: Bandwidth: 99 Mbit/sec 2) TCP: Bandwidth Sender: 102 Mbit/sec Receiver: 99.9 Mbit/sec UDP: Bandwidth: 98.9 Mbit/sec 3) TCP: Bandwidth Sender: 102 Mbit/sec Receiver: 100 Mbit/sec UDP: Bandwidth: 98.8 Mbit/sec
Public IPs - Public IPs for VMs? - Available kinds of public IPs for VMs - Public IPs for Load Balancers? - Available kinds of public IPs for Load Balancers	yes floating/static yes static	yes floating/static yes static
Is a dedicated network connection from datacenter to public cloud possible?	yes	yes
Network Security features (Network Traffic analysis, Network Security Groups)	<ul style="list-style-type: none"> - Network Security Groups - Firewalls (Multi VLAN, Single VLAN and Web App) - DDOS mitigation 	<ul style="list-style-type: none"> - Network Security Groups - Firewalls (Multi VLAN, Single VLAN and Web App) - DDOS mitigation
VPN as a Service	yes	yes
Traffic costs per GB	€ 0.078 / \$ 0.087	€ 0.078 / \$ 0.087

Storage

Questions	IBM Cloud Classic	IBM Cloud latest
Which kinds of storage are available? - Object / Blob Storage - File Storage - Block Storage	yes (IBM Cloud Object Storage) yes (IBM Cloud file storage) yes (IBM Cloud block storage)	yes (IBM Cloud Object Storage) yes (IBM Cloud file storage) yes (IBM Cloud block storage)
Block - Different tier-classes? SATA, SSD, SAS	yes	yes
Which objects storage-engines are offered?	- S3 - Swift	- S3 - Swift
File - Accessing file storage via (cluster) file system.	- NFS	- NFS
Storage capacity limits	Overall size: Unlimited 25 GB per month of object storage (Unlimited for standard plan)	Overall size: Unlimited 25 GB per month of object storage (Unlimited for standard plan)
Duration of provisioning?	24 sec	25 sec
Throughput IOPS (only Block- and File-Storage)	- Random Read: bw = 188.67 MB/s, iops = 23583 - Random Write: bw = 233.98 MB/s, iops = 3655 - Random Read and write: - read: bw = 19.08 MB/s, iops = 11923 - write: bw = 21.74 MB/s, iops = 1358 Sequential Read: bw = 127.40 MB/s, iops = 15924 Sequential Write: bw = 115.70 MB/s, iops = 3615	- Random Read: bw = 92.88 MB/s, iops = 11610 - Random Write: bw = 24.30 MB/s, iops = 3796 - Random Read and write: - read: bw = 78.24 MB/s, iops = 4889 - write: bw = 8.91 MB/s, iops = 557 Sequential Read: bw = 79.44 MB/s, iops = 9929 Sequential Write: bw = 85.32 MB/s, iops = 2666
Costs per month - total price for 50 GB Disk which is mounted to the VM	€ 8.82 / \$ 9.89	€ 8.82 / \$ 9.89

Software as a Service

Questions	IBM Cloud Classic	IBM Cloud latest
Is a mobile office suite offered? Is it deeply integrated with other services?	yes yes	n/a
Managed App Services	IBM Cloud API CI/CD Database as a Service Network as a Service Function as a Service, Webserver as a Service, Monitoring as a Service, Backup as a service, AI as a service	n/a
Mobile App Services - Push Notifications - User Management - NoSQL-Datenbase - File Storage - Messaging - Social Networks	IBM Mobile Foundation yes yes yes yes yes yes	n/a
Application Environments - Websites - Microservices - Messaging - Serverless	yes yes yes (IBM message Hub) yes (Cloud Functions)	n/a
Rollback to a previous application version?	yes	n/a

CaaS

Questions	IBM Cloud Classic	IBM Cloud latest
Which technologies are being provided/supported?	- Kubernetes	n/a
Is a managed container service available?	yes	n/a
Can worker nodes be accessed directly by customers?	yes	n/a
Can master nodes be accessed directly by customers?	yes	n/a
Which version of the technologies/Kubernetes is being offered?	1.14.1, 1.13.6, 1.12.8, 1.11.3, 1.10.8, 1.9.8	n/a
How much time does it take to provide the container service for four nodes	15 min	n/a
Costs - Managed service - 732hrs per month - small flavor - hosted in Frankfurt or Western Europe - Storage / IPs not included * Prices in USD have been converted to EUR	€247.68 / \$ 275.78	n/a
Shared or dedicated Container Engine Cluster?	dedicated	n/a
Do predefined StorageClasses exist in Kubernetes? Name - Provisioner	ibmc-file-bronze (default) - ibm.io/ ibmc-file ibmc-file-custom - ibm.io/ibmc-file ibmc-file-gold - ibm.io/ibmc-file ibmc-file-retain-bronze - ibm.io/ibmc-file ibmc-file-retain-custom - ibm.io/ibmc-file ibmc-file-retain-gold - ibm.io/ibmc-file ibmc-file-retain-silver - ibm.io/ibmc-file ibmc-file-silver - ibm.io/ibmc-file	n/a
Limitations - What is the maximum cluster size?	n/a	n/a
Do you have full access to all K8s resources (no RBAC restriction)?	no	n/a
Does the Container Service provide a Load-Balancer Service?	yes	n/a
Is a Storage Class useable?	yes	n/a

DBaaS

Questions	IBM Cloud Classic	IBM Cloud latest
Which DB engines are offered?	Relational DB - Db2 on Cloud - PostgreSQL - MySQL Non-Relational DB - Cloudant - MongoDB - ScyllaDB - Redis - JanusGraph - etcd - Elasticsearch Data Warehouse / Big Data - Db2 Warehouse on Cloud	Relational DB - Db2 on Cloud - PostgreSQL - MySQL Non-Relational DB - Cloudant - MongoDB - ScyllaDB - Redis - JanusGraph - etcd - Elasticsearch Data Warehouse / Big Data - Db2 Warehouse on Cloud
Performance of MySQL (MySQL Sysbench, table-size (row data): 1000000, Threads: 16) - Read - Write - Read / Write	Transactions:31026 (516.82 / sec) Transactions:64777 (1079.35 / sec) Transactions:22708 (378.18 / sec)	Transactions:4705 (78.33 / sec) Transactions:11806 (196.52 / sec) Transactions:3691 (61.27 / sec)
Provisioning time for a MySQL instance	120 sec	132 sec

Performance of PostgreSQL	Transactions: 837378 (13953.69 / sec) Transactions: 609573 (10158.00 / sec) Transactions: 33815 (563.26 / sec)	Transactions: 65058 (1084.03 / sec) Transactions: 63624 (1060.11 / sec) Transactions: 2961 (49.14 / sec)
Provisioning time for a PostgreSQL instance	195 sec	180 sec
Supported DB Versions	Db2-ge PostgreSQL 9.6.12,9.6.10,9.6.9,9.5.14,9.5.13, 9.4.19,9.4.18 MySQL 5.7.22 Cloudant-h7 MongoDB 3.4.10,3.2.18,3.2.11,3.2.10 ScyllaDB 2.0.3 Redis 4.0.10,3.2.12 JanusGraph 0.1.1 beta etcd 3.3.3,3.2.18 Elasticsearch 6.2.2, 5.6.9 Db2 Warehouse-ef	Db2-ge PostgreSQL 9.6.10,9.6.9,9.5.14,9.5.13,9.4. 19,9.4.18 MySQL 5.7.22 Cloudant-h7 MongoDB 3.4.10,3.2.18,3.2.11,3.2.10 ScyllaDB 2.0.3 Redis 4.0.10,3.2.12 JanusGraph 0.1.1 beta etcd 3.3.3,3.2.18 Elasticsearch 6.2.2, 5.6.9 Db2 Warehouse-ef
Troubleshooting as a Service - Rollback - Support	yes yes	yes yes
Total price for the database per month - MySQL - 2 vCores - 100 GB Storage - Frankfurt / Western Europe - 100% active per month - No dedicated backup	N/A	N/A
Total price for the database per month - PostgreSQL - 2 vCores - 100 GB Storage - Frankfurt / Western Europe - 100% active per month - No dedicated backup	€103.04 / \$136.00	€103.04 / \$136.00
Limitations: How many simultaneous requests to the DB? How much RAM? How many users?	MySQL: - max Connections: 151 PostgreSQL: - max Connections: 1000	MySQL: - max Connections: 151 PostgreSQL: - max Connections: 1000
How does backup/restore work?	Backups: - Automatic Backups. Restore -On-demand	Backups: - Automatic Backups. Restore -On-demand

Monitoring

Questions	IBM Cloud Classic	IBM Cloud latest
Dashboard	yes	n/a
Which cloud resources will be monitored? VMs Apps Network Load Balancer Storage	yes yes yes yes yes	n/a
Connection/Usage of external monitoring solutions	yes	
Costs per month	n/a	n/a

IaaS, PaaS and SaaS-Patch-Management

Questions	IBM Cloud Classic	IBM Cloud latest
Does the cloud provide a managed patch service?	yes (IBM BigFix Patch Management)	yes (IBM BigFix Patch Management)
Which operating systems are supported?	Linux: <ul style="list-style-type: none"> - CentOS 7, 6 - RedHat Enterprise 7, 6 - SUSE Linux Enterprise Server 12, 11 - Ubuntu Minimal 18.04, 16.04, 14.04 Windows: <ul style="list-style-type: none"> - Windows Server 2016, 2012 R2, 2008 R2 	Linux: <ul style="list-style-type: none"> - CentOS-Minimal 7.X, 6.X - CentOS-LAMP 7.X.6.X - Debian Minimal Stable 9.X, 8.X - Debian LAMP Stable 8.X - Red Hat Minimal 7.x, 6.x - Red Hat LAMP 7.x, 6.x - Ubuntu Minimal 18.04, 16.04, 14.04 - Ubuntu LAMP 18.04, 16.04, 14.04 Windows: <ul style="list-style-type: none"> - Standard 2016 - Standard 2012 - R2 Standard 2012
Is the operating system from the deployed VM at a current patch level?	yes	yes
What is the current available patch level in our sample VM? - Ubuntu 16.04 LTS with latest patches applied	04.2	0.04

Backup, Recovery and Availability

Questions	IBM Cloud Classic	IBM Cloud latest
Are managed backups offered (Provider is responsible to take backups)	yes	yes
Which types of backups are supported for VMs?	<ul style="list-style-type: none"> - Snapshot - Full Backups - Incremental Backups 	<ul style="list-style-type: none"> - Snapshot - Full Backups - Incremental Backups
Where will the backup be stored?	<ul style="list-style-type: none"> - Evault - dedicated backup space - IBM Cloud Backup, - IBM Object Storage archive 	<ul style="list-style-type: none"> - Evault - r1 cdp
Can backups be scheduled?	yes	yes
Usage costs per month - 500 GB Backup Storage - Western Europe	€ 0.88 / \$1.00	€ 0.88 / \$1.00
IaaS - Start backup job with the backup feature of the cloud. Restore the backup and test if VM can be restored. Also try to restore into a new VM.	yes	yes

Logging as a Service

Questions	IBM Cloud Classic	IBM Cloud latest
Does the cloud platform provide a Logging as a Service functionality?	yes	
Is the data stored in encrypted form?	yes	
Which logging technology is used?	<ul style="list-style-type: none"> - IBM Log Analysis with LogDNA - Bluemix UI - Cloud Foundry Line Interface(CLI) - External logging 	

Security

Questions	IBM Cloud Classic	IBM Cloud latest
Integration to a SIEM possible? (Security Information and Event Management)	yes	yes
Security Groups	yes	yes
Disk Encryption	yes	yes
Network Traffic Analyse	yes	yes
Protection against Denial of Service Attacks	yes	yes
Firewall - Does the cloud provider provide additional integrated security features i.e. a Next Generation Firewall?	yes	yes
Does the cloud provider keep an eye on current threats and take action?	yes	yes
Does the cloud provider support additional integrated security features for cloud resources using 3rd party tools:		
IDS (Intrusion Detection System)	yes	yes
IPS (Intrusion Prevention System)	yes	yes
ATP (Advanced Threat Protection)	yes	yes
Does the provider carry out regular penetration tests against the platform?	no	no

Image Service

Questions	IBM Cloud Classic	IBM Cloud latest
Which operating systems are offered by the provider with which versions?	<p>Windows:</p> <ul style="list-style-type: none"> - Windows Server 2012, 2016 <p>Linux:</p> <ul style="list-style-type: none"> - CentOS- Minimal 6.X, 7.x - CentOS-LAMP 6.X, 7.X - Debian Minimal Stable 8.X, 9.x - Debian LAMP Stable 8.X - Red Hat Minimal 6.x, 7.x - Red Hat LAMP 6.x, 7.x - Ubuntu Minimal 14.04, 16.04, 18.04 - Ubuntu LAMP 14.04, 16.04, 18.04 	n/a
Can own images be uploaded?	yes	n/a
Can existing licenses be used to minimize costs?	yes	n/a
Is there an image build service?	<p>yes</p> <p>Supported formats:</p> <ul style="list-style-type: none"> - VHD - VMDK - QCOW2 - AKI - ARI - AMI 	n/a
Can images be created from existing cloud instances?	yes	n/a
Are different patch levels of images available?	yes	n/a



HAMBURG | JUNE 24 - 26, 2019

CONTAINER days

www.containerdays.io