Cloud Price/Performance Evaluation

Results of a provider comparison in cooperation with 1&1 IONOS Cloud GmbH

by René Büst, Maximilian Hille and Björn Böttcher (freelance IT consultant)

April 2015

crisp RESEARCH

A current research project performed in cooperation with 1&1 IONOS Cloud GmbH

Executive Summary

- The current supply situation makes it hard for CIOs and CTOs to distinguish between individual public infrastructure-as-a-service providers. Comparing the almost infinite number of packages, different billing models and the underlying infrastructures takes a great deal of effort.
- IT decision-makers need a resilient decision-making basis in order to compare the costs and services of individual providers. Before decision-makers outsource large workloads, they should find out about the implications of price and performance.
- This price/performance test compares the four relevant public-cloud providers Amazon Web Services, Google Cloud Platform, Microsoft Azure and 1&1 IONOS. For performance measurement, the average response time was taken into account. To ensure an objective test scenario, a real and standardized workload was used that covers the requirements of a majority of users.
- With its outstanding performance and competitive price, the Enterprise Cloud by 1&1 IONOS (93 percent of the reference value) was the price/performance winner. Google came next (87 percent), followed by Microsoft and Amazon AWS (73 percent each).
- This shows that size is not a decisive factor in a good price/performance ratio. Small and local providers promise a good combination of price and performance compared with the large ones.

Contents

1. Background and Objectives

When planning an IT infrastructure, it is constantly necessary to find the right balance of resources in order to operate all applications with sufficient performance. At the same time, the challenge is to avoid oversizing the infrastructure thus wasting capacity. Excessive deployment of infrastructure leads to further costs. If too few resources are factored in, there is a risk of performance restrictions or failures.

The golden path to correct sizing of the infrastructure therefore consists of managing capacity and performance. What had to be heeded in previous years when planning the in-house on-premise infrastructure now has to be fully taken into account when using public-cloud infrastructures. A powerful IT infrastructure is essential to ensuring the customer experience of modern digital services.

The public infrastructure-as-a-service (IaaS) market is growing incessantly. According to current forecasts, German companies will invest around EUR 2 billion in public-cloud infrastructure resources in 2018. Instead of relying solely on the statements of public-cloud providers, decision-makers should use performance management and further control mechanisms to constantly verify the services provided by the providers. The price/performance ratio is a killer criterion for CIOs here.

Price/performance ratio

Both the costs and the performance of infrastructure resources from the public cloud are hard to estimate, especially when attempting to compare them with multiple providers. To get a rough idea of what resources are needed within a public cloud and what costs arise during use, a requirements analysis with regard to capacity planning and expected performance is first required in order to ensure stable operation of the applications and systems.

→ Definition Performance

Performance is a measure of how well IT systems work. It depends on the chosen infrastructure configuration, the interaction between the individual systems and the most optimum possible mapping of the respective workloads. Within these tests, the average response time of the infrastructure per unit of time was chosen to measure performance. A shorter response time implies higher performance.

However, the performance properties vary depending on the type and architecture of the application and its use. These factors make it difficult to determine performance in relation to price across cloud infrastructures.

One prominent evaluation factor is the fact that all public-cloud infrastructures differ considerably in terms of how they are built, their size and their detailed pricing structure.

Every public cloud is based on an individual physical infrastructure that consists of different components with regard to the processors used, the network infrastructure, the type of memory and various virtualization stacks. In each case, performance and price can only be determined from the individual infrastructure components such as the virtual machines (VMs), storage space and bandwidth. However, this cannot be mapped directly to the price/performance ratio of an application.

In addition, every public-cloud infrastructure has to contend with highly dynamic environmental characteristics that differ significantly in terms of capacity utilization. Virtualization management should not be ignored here: it may result in discrepancies in performance depending on the public cloud and how and when a VM is provided. These performance differences are possible between different public clouds and, over time, within a public cloud.

To ensure a certain degree of transparency regarding the various VM types and their provisioning processes when determining the performance of an application, it is necessary to apply exactly the same application to various public-cloud infrastructures and to conduct a performance test on each of these infrastructures.

Transparency of IaaS costs

In addition to the variations of the technical infrastructure, the public IaaS services particularly differ in terms of packages and pricing models for virtual machines and further resources. This results in little transparency for users and makes comparison confusing. For example, an instance type with the designation “large” from one public-cloud provider is hard or even impossible to compare with the same “large” type from another provider, as it is not certain whether the virtual CPU delivers the same performance for both of them.

Furthermore, it is apparent that most providers dimension their VMs in a specific way, meaning that the next machine up is much more expensive, but oversized for the purpose in mind. The customer should therefore pay less for the actual purpose if there were a more varied selection. The same applies to the promised flexibility, which is not delivered in many cases as many providers have different bases for their billing, from per hour to per minute.

Another level of complexity is introduced if infrastructure-related services are used in addition to virtual servers, storage and network resources. Although they help to develop the application more quickly and efficiently on the public-cloud infrastructure, they also incur further, occasionally unpredictable, costs, as their level of use by the application fluctuates.

Therefore, native infrastructure-related services should be ignored in order to ensure a uniform comparison of the price/performance ratio from different public-cloud providers.

Aim of the study

The public IaaS market has developed rapidly in recent years. Because there are so many public IaaS providers with different fixed or flexible infrastructure configurations and pricing models, there is a lack of transparency regarding costs, which are very hard or even impossible to compare. Consequently, companies do not have a sound decision-making basis that provides specific information on the price/performance ratio of a public IaaS service. Furthermore, due to generally inflexible packages, companies overpay for the actual application that they operate on the infrastructures.

In view of this, the standardized benchmark set out here is intended to show IT decision-makers which public-cloud provider has the best price/performance ratio.

2. Methods and Procedure

Configuration of the test infrastructures and technical execution of the tests were carried out by the independent IT consultant Björn Böttcher, who specializes in development and integration of cloud and mobile applications in the enterprise sector. The price/performance test compared a total of four public-cloud infrastructures: three from global providers, and one from 1&1 IONOS, a leading local provider in the German market.

1. Amazon Web Services (AWS)
2. Google Cloud Platform (GCP)
3. Microsoft Azure
4. 1&1 IONOS

A real workload was the priority during the tests. The UnixBench[1] Suite was used in most performance tests. This is a collection of tools that generates a UnixBench score at the end, providing information on the overall performance of a Unix-like system (e.g. Linux). Although solutions such as UnixBench are very widely used in the industry and capable of comparing the local performance of systems, the test sequences are predetermined and do not reflect reality.

To provide information on the real performance of a public-cloud infrastructure, it is necessary to take a genuine standard workload into account, i.e. either a standard application or a known use case. This is the only way in which companies and decision-makers can attain a meaningful basis for deciding in favor of or against a specific public-cloud infrastructure.

The test does not take into account the amount of work or the configuration costs for set-up and deployment of the infrastructures. Likewise, the complexity necessarily involved in setting up and operating a public-cloud application is ignored.

Test configuration and prices

A common basis is the essential requirement when testing different public-cloud infrastructures. To guarantee optimum comparability of the services, the initial configuration of a virtual machine (VM) that reflects the lowest common denominator across all providers was normalized. For this test, the basic configuration of a Magento web shop constituted the normalized initial configuration.

1 UnixBench, https://code.google.com/p/byte-unixbench/

Normalized initial configuration:
1 core/at least 3.5 GB RAM/at least 4 GB HDD

	Amazon Web Services	1&1 IONOS	Microsoft Azure	Google GCP
VM name	m3.medium	VDC	Virtual Linux computers, standard, D1	n1-standard-1
Core	1 Core	1 Core	1 Core	1 Core
Main memory (RAM)	3,7 GB	4 GB	3,5 GB	3,8 GB
Hard drive (HDD)	4GB	4 GB	50 GB	Not included

SOURCE: Crisp Research AG, 2015

The biggest difference among the selectable VMs is apparent here from the local hard-drive size (HDD) in the VM configuration. At least 4 GB of local instance memory is required in order to operate a shop system suitable for the test scenario. If the 4 GB HDD specified for the test was not directly available from a provider, the difference in gigabytes (GB) was purchased at extra cost. As a result, 4 GB of local HDD memory [2] had to be added for the Google Cloud Platform.

	Amazon Web Services	1&1 IONOS	Microsoft Azure	Google GCP
VM Name	m3.medium	VDC	Virtuelle Linux Computer, Standard, D1	n1-standard-1
Core	1 Core	1 Core	1 Core	1 Core
Main memory (RAM)	3,7 GB	4 GB	3,5 GB	3,8 GB
Hard drive (HDD)	4GB	4 GB	50 GB	Not included
		Prices (in EUR*)		
VM**	€0,07	€0,04	€0,07	€0,06
HDD***	Included in the VM price	Included in the VM price	Included in the VM price	€0,60
Price per month	€48,96	€28,80	€50,47	€44,52
Price per hour; utilization: 24h, 30 days	€0,07	€0,04	€0,07	€0,06

* Currency conversion, ** Costs per hour, *** Costs per month

SOURCE: Crisp Research AG, 2015

The price per hour was calculated for a VM with the defined initial configuration based on 24-hour utilization 30 days per month. The prices for each provider were obtained on January 30, 2015. The prices for the Europe region were selected for all providers. Furthermore, in the case of Amazon Web Services and the Google Cloud Platform, the prices were converted from USD to EUR.

2 Google Compute Engine - Local SSD, https://cloud.google.com/compute/docs/disks/local-ssd

Price/performance evaluation

During the price/performance evaluation (PPE), the prices and performance of the respective providers were determined as a ratio. The price/performance ratios of the providers were then compared with each other, taking a reference value into consideration.

Classification of prices

The prices were assigned a score from 1 to 10 points, in 1-cent bands. A low price per hour resulted in a higher score.

Price/points system

Price in euro cents	0–2 cents	2–3 cents	3–4 cents	4–5 cents	5–6 cents	6–7 cents	7–8 cents	8–9 cents	9–10 cents	>10 cents
Score in points	10	9	8	7	6	5	4	3	2	1

SOURCE: Crisp Research AG, 2015

Classification of performance

Performance was measured in milliseconds and assigned a score from 1 to 10 points. A low milliseconds range indicates higher performance, resulting in a higher score.

Performance/points system

Performance in ms	< 3.000	3.000 – < 5.000	5.000 – < 7.000	7.000 – < 9.000	9.000 – < 11.000	11.000 – < 13.000	13.000 – < 15.000	15.000 – < 17.000	17.000 – < 19.000	ab 19.000
Score in points	10	9	8	7	6	5	4	3	2	1

SOURCE: Crisp Research AG, 2015

Price/performance ratio (PPR)

The price/performance ratio was determined by adding up the respective scores and then dividing the result by two. The two factors were not weighted, as price or performance has greater significance depending on the subjective assessment.

→ PPR = (price score + performance score)/2

PPR example

If the VM costs amount to EUR 0.086 per hour, the provider receives 3 points. For a performance of between 7,000 milliseconds and 9,000 milliseconds, the provider receives 7 points. This results in the following calculation basis:

→ PPR = (3 price + 7 performance)/2 = 5 PPR

The providers are measured against a reference value made up of the mean value of the best price rating and the best performance rating.

For example, if provider A scores highest in the price benchmark with 8 points and provider B scores highest in the performance benchmark with 7 points, this produces a PPR reference value of 7.5 points.

The provider whose total score is closest to the reference value comes out on top in the cloud price/performance evaluation.

3. Test and Results

For the public-cloud performance test, typical use of a web shop in the cloud-hosting environment was evaluated. The e-commerce software used for the test was Magento. It is a commonly used platform in the professional sector. The solution is based on PHP and needs a LAMP stack (LINUX/Apache/MySQL/PHP).

The actual Magento framework used was not optimized, and set up as a Magento demo shop in order to ensure a uniform impression for all platforms. Furthermore, optimizations, e.g. based on use of Amazon RDS or Google Cloud SQL, were not taken into account. However, what is more important is the fact that not all providers have corresponding dedicated higher-level services in their range, which would allow a comparison at this level as well. Therefore, only the virtualized computers were used and compared with each other for this test.

This specific use case usually requires customers. In addition, there are generally peaks in the load behavior over time. BlazeMeter [3] was used to simulate this case. BlazeMeter is a software-as-a-service product for load tests, and is compatible with Apache JMeter. Via the test definition in JMeter, it is possible to access several pages within the web shop and even simulate a purchase transaction.

Based on the response times of the system and other parameters, conclusions can be drawn regarding technical bottlenecks or, as in the case of this study, the performance of this system can be compared with that of others. Metrics for performance measurement include the average response time of the system. The average response time indicates how long a page needs in order to load. This measurement is performed from the perspective of the client making this request, i.e. from the user's perspective.

Scripts for this analysis were created in JMeter by means of a graphical user interface. They were then executed in parallel by several computers, for example in order to simulate a larger number of users or emulate different locations. Furthermore, it is possible to record navigation through a web shop by a user, for example, and to get the tool to automatically simulate the behavior afterwards. The aim of the test should always be to simulate user behavior as realistically as possible.

The tests were performed with a simulation of 50 users acting in parallel, and were repeated several times at irregular intervals. The purpose of this is to filter out potential external influences or regional disruptions at the provider end. However, only one test cycle is reproduced below, as there were no significant differences between the relative positions of the providers. This test was not performed simultaneously, but was scheduled at the same time of day (between 1:30 p.m. and 4:30 p.m., GMT+1) to allow all providers to use the same conditions.

3 Blazemeter, http://blazemeter.com/

- Amazon Web Services: 1:30 p.m. – 2:15 p.m.
- 1&1 IONOS: 2:30 p.m. – 3:15 p.m.
- Google Cloud Platform: 3:30 p.m. – 4:15 p.m.
- Microsoft Azure: 4:30 p.m. – 5:15 p.m.

Infrastructure size is no guarantee of performance

The performance of the four assessed public IaaS services was evaluated on the basis of the calculated response times. Providers that have low values throughout this test run have a correspondingly higher performance.

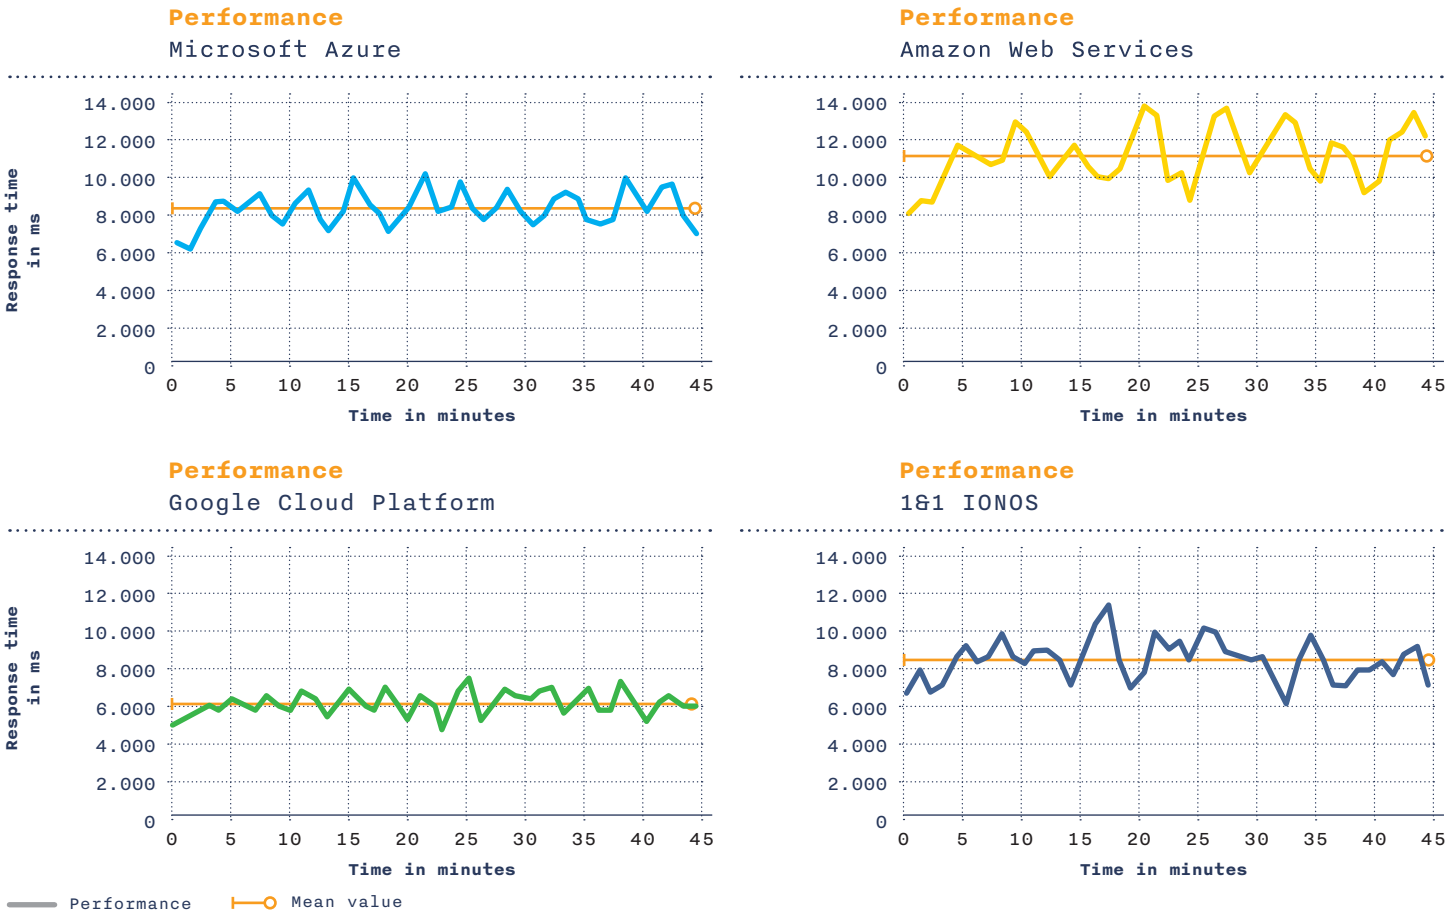

SOURCE:
Crisp Research AG, 2015

The charts show the response times of the virtual machines throughout the entire 45-minute test as well as the average response time per provider in milliseconds. All the tests were performed with identical workloads (see above). This means that exactly the same simulation processes were undertaken in each test, thus ensuring comparability.

Performance comparison of providers

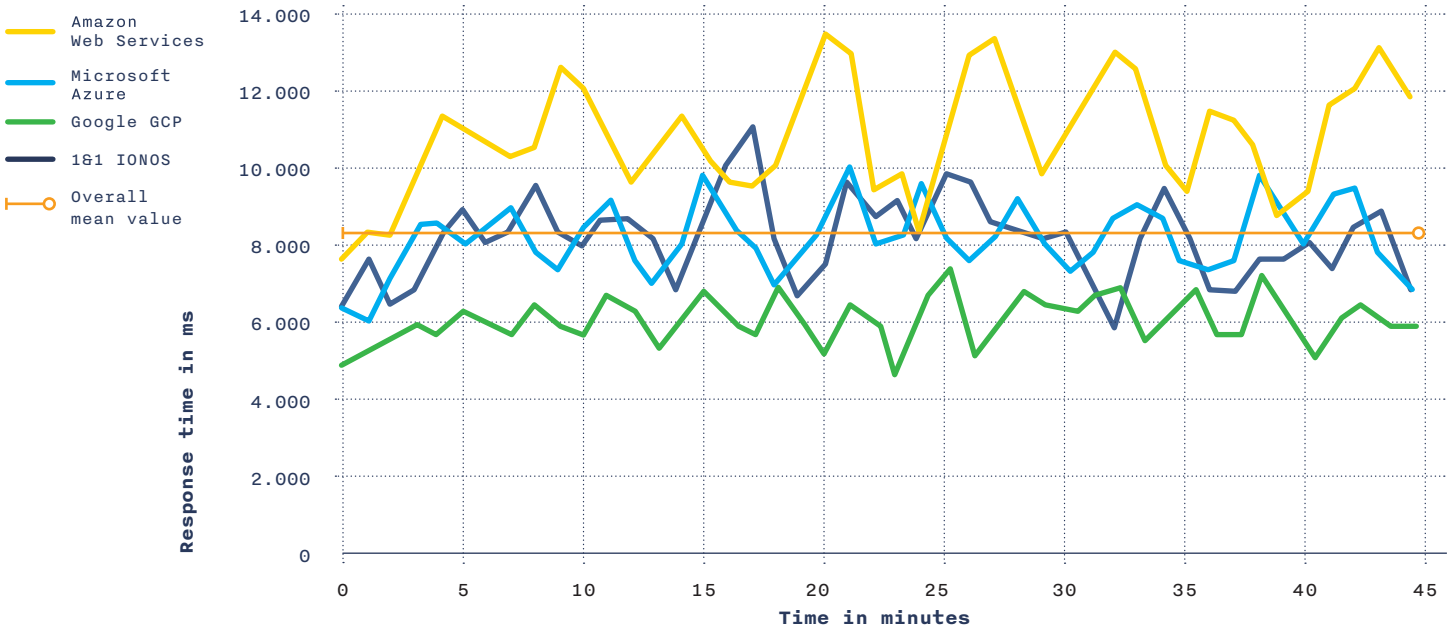

SOURCE: Crisp Research AG, 2015

The individual data tracking shows how effectively the individual configurations can compensate for any load peaks. Google and Microsoft have a very harmonious line, whereas 1&1 IONOS and Amazon AWS show much more irregular patterns.

Average response time in milliseconds

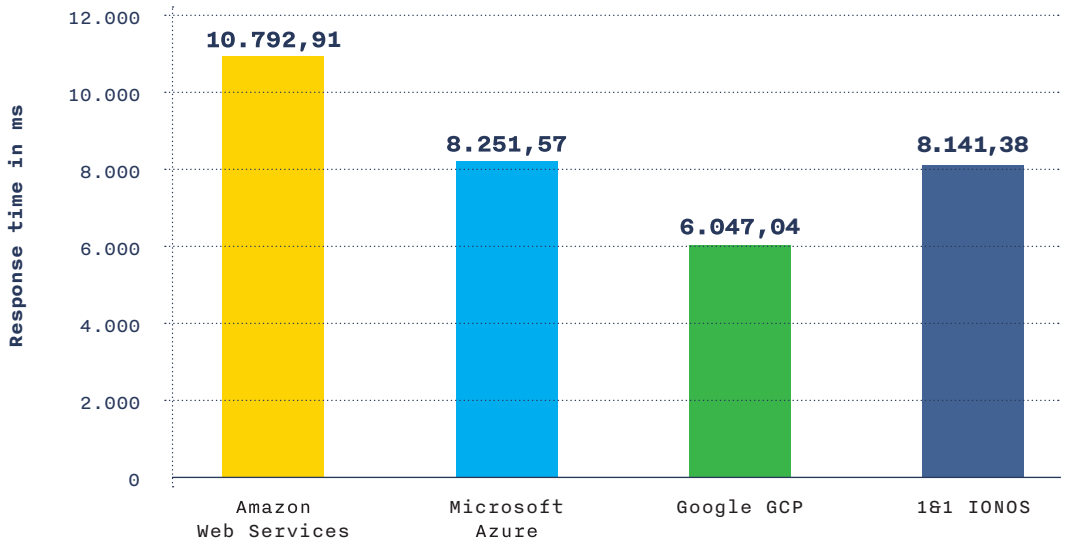

SOURCE: Crisp Research AG, 2015

Compared with the overall performance, a largely clear picture is apparent. Whereas Google shows a high performance with low response times, Amazon AWS tails off noticeably in these test scenarios. Despite having one of the biggest infrastructures and a high market share in the public-cloud sector, Amazon AWS has the lowest performance. The performance of Microsoft Azure and 1&1 IONOS is somewhat better compared with the average across all providers (approx. 8,300 ms). However, in a direct comparison, the differences are more or less marginal.

Even so, on closer inspection, the Enteprise Cloud by 1&1 IONOS can deliver a somewhat higher performance.

Low price as a competitive factor

The price comparison sets out the operating costs of the chosen configuration for one hour based on monthly billing for 24/7 operation.

Price comparison of providers

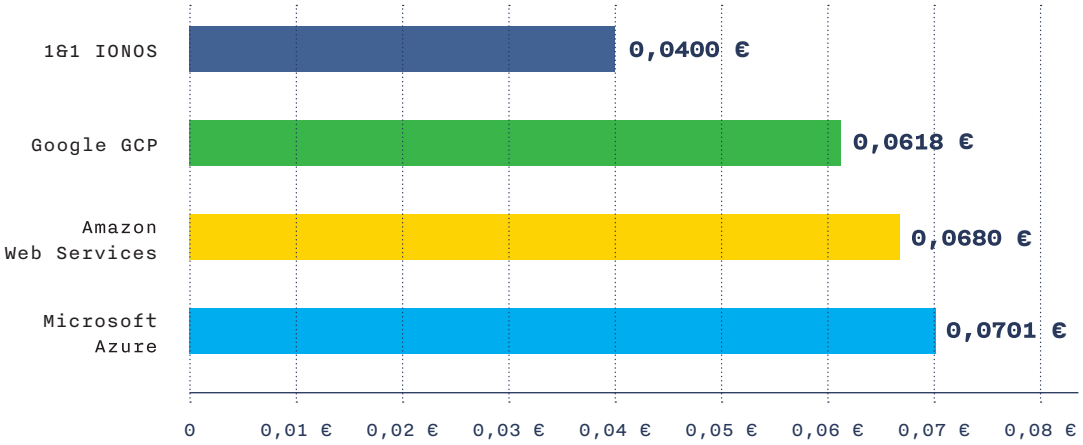

SOURCE:
Crisp Research AG, 2015

It can be seen that the Enteprise Cloud by 1&1 IONOS is by far the most cost-effective provider in the test at €0.04 per hour of operation. Along with the associated storage price, the pricing model also impresses by virtue of flexible billing. For instance, although the price for VM configuration is stated on an hourly basis, billing is to the nearest minute, thus making pricing as use-dependent as possible. Google (€0.062) also has a relatively flexible billing model.

→ **1&1 IONOS „DevOps Central“**

Developers that buy resources via the “DevOps Central” portal receive a further 33-percent discount on US prices.
DevOps Central: https://www.ionos.com/pro/devops/community/

The interval in this case is five minutes. At €0.068, Amazon AWS provides somewhat more expensive configuration and inflexible hourly billing. The Microsoft Azure VM has the highest price per hour in this comparison. However, the approx. €0.07 includes more than enough local instance storage than was required for this test.

What really matters: Price/performance ratio

For a complete picture of the tested public-IaaS configurations, price and performance were assessed in accordance with the procedure described in chapter 2 and compared with the reference value.

Price/performance evaluation

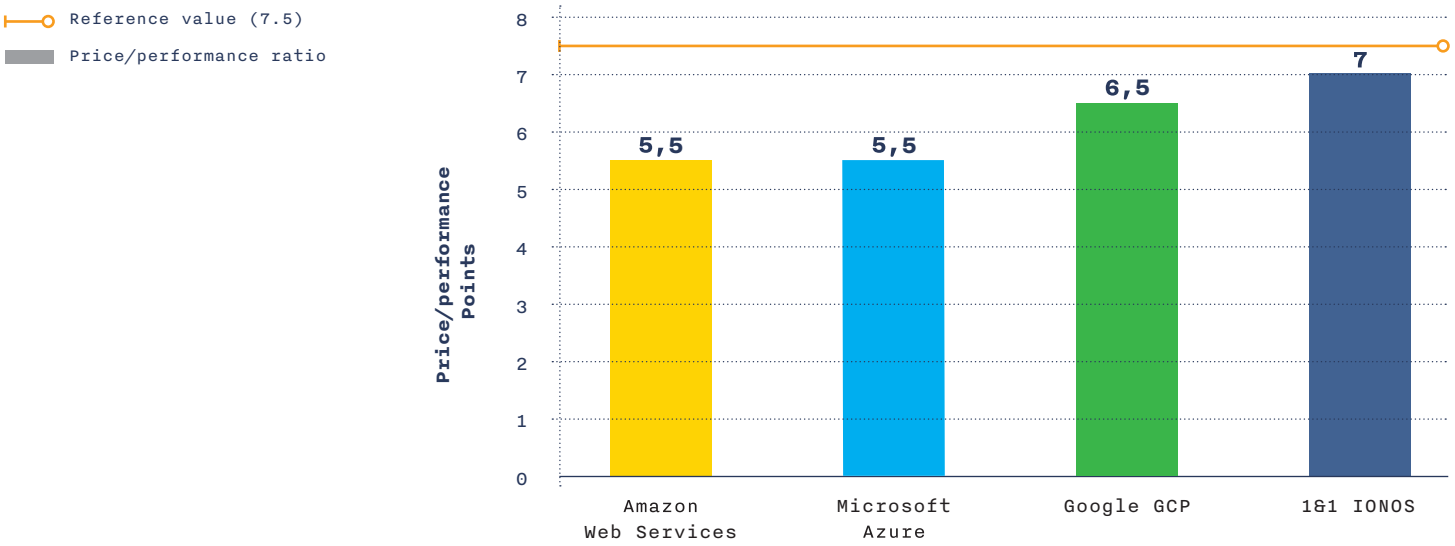

SOURCE:
Crisp Research AG, 2015

In this comparison, the reference value is made up of the performance of Google (8) and the price of 1&1 IONOS (7). Consequently, the reference value is 7.5. All providers in the test must be measured against this total score.

Amazon AWS and Microsoft came joint third in the price/performance test. Performance and price are well below the average with a combined figure of 5.5. The lower price on the one hand and the better performance on the other cancel each other out between the two providers. Amazon AWS provides somewhat less expensive configuration, while Microsoft Azure is stronger in terms of performance. The performance winner, Google, has the second-best price/performance ratio with a total value of 6.5.
The Enterprise Cloud by 1&1 IONOS beats the competition in the combination of price and performance. Its score of 7 is close to the possible reference value of 7.5 and shows competitors that outstanding performance can also be achieved with low prices.

4. Interpretation of Results

The results show that it is advisable to perform a thorough evaluation before using an IaaS service. Although meeting the individual requirements is the main factor that influences the decision, performance and price are variables that must not be ignored.

In the context of the test, it became clear that significant and sometimes unexpected differences arise for the obvious factors such as price and performance. Google sets the performance benchmark in this test scenario. Measured against this, the Enterprise Cloud by 1&1 IONOS achieves a performance of 74.3 percent, Microsoft Azure 73.3 percent and Amazon AWS 56.0 percent. With regard to price, the Enterprise Cloud by 1&1 IONOS sets the standard with 4 cents per hour of operation. Google costs 54.6 percent more, Amazon AWS 70 percent more and Microsoft Azure as much as 75.3 percent more.

In terms of the price/performance ratio, it is particularly clear that very good performance can be achieved with low costs. The Enterprise Cloud by 1&1 IONOS is the price/performance winner in this test. This shows that size is not always decisive, and that good results can be attained with an infrastructure that is smaller than that of Google, Microsoft and AWS yet still sound. Consequently, 1&1 IONOS has shown in this test that the combination of low prices and a competitive infrastructure leads to overall victory.

Conclusions

The test has shown that the public infrastructure-as-a-service market is nowhere near as homogeneous as people think. Even with a test situation of four different providers and a comparatively simple, small and highly standardized configuration, significant and thoroughly unexpected differences were identified.

A critical look at the performance and pricing models of providers is an absolute must for all infrastructure decision-makers. After all, a high price does not necessarily mean good performance. However, in practice, the balance between the individual requirements and the resources actually used must be right. The performance differences can be even more pronounced, especially in the case of much more complex infrastructure configurations where multiple servers communicate with each other.

This test deliberately ignored all the optimization potential and further services of the respective providers, and focused solely on testing with a standard application. This results in configurations that a large web shop greatly exceeds. Accordingly, the response times are comparatively high at 6 to 12 seconds on average. However, with the infrastructure-related services of individual providers, it is possible to perform fine-tuning and ratchet up performance by adjusting specific parameters. Yet to achieve this, it is necessary to develop an application specifically for the respective provider infrastructure, although this would rule out comparability of performance. Likewise, this would impact on the price, as use of the add-on services results in further costs.

To ensure that a provider delivers the performance required to meet the individual requirements, it is advisable to set up corresponding test systems at the selected providers and evaluate them in advance. This is the only way in which it is possible to objectively evaluate and assess which provider delivers what performance at what price.

A global data-center infrastructure, a large network of partners and an extensive range of value-added services are key factors, and form part of the basis for decision-making. However, the price/performance ratio plays a particularly crucial role. It is apparent here that supposedly small providers can also achieve an influential competitive position.

About Crisp Research AG

Crisp Research AG is an independent IT research and consultancy firm. With a team of experienced analysts, consultants and software developers, Crisp Research assesses current and future technology and market trends. Crisp Research helps companies with the digital transformation of their IT and business processes.

The analyses and commentaries of Crisp Research are published and discussed in a host of business and IT journals and on social-media platforms. As contributing editors at leading IT publications (Computerwoche, CIO, Silicon, etc.), committed BITKOM members and sought-after keynote speakers, our analysts actively contribute to debates surrounding new technologies, standards and market trends and are influential figures in the industry.

Crisp Research was established in 2013 by Steve Janata and Dr. Carlo Velten, and focuses its research and consulting on emerging technologies such as cloud computing, analytics and digital marketing and their strategic and operational implications for CIOs and business decision-makers at companies.

Weissenburgstrasse 10
D-34117 Kassel
TEL +49 561 2207 – 4080
FAX +49 561 2207 – 4081

E-MAIL info@crisp-research.com

WEB crisp-research.com
crisp-analytics.com

TWITTER twitter.com/crisp_research

About 1&1 IONOS

With more than eight million customer contracts, 1&1 IONOS is the leading European provider of cloud infrastructure, cloud services, and hosting services. From VPS and bare-metal servers all the way to high-end IaaS solutions: 1&1 IONOS offers SMEs and large companies all the products they need to set up their hybrid or multi-cloud environment and is the only IaaS cloud computing provider that has its own code stack in Germany. 1&1 IONOS operates one of the world's largest and highest-quality IT infrastructures with over 90,000 servers. In the Cloud Vendor Universe from Crisp Research, 1&1 IONOS has repeatedly been named one of the leading providers of cloud platforms.

The Enterprise Cloud by 1&1 IONOS is the "Cloud – Made in Germany" with a data protection-compliant IaaS platform developed in-house for companies, system vendors/integrators, and managed service providers. It is flexibly scalable and provides free 24/7 support by qualified system administrators. During operation, the capacity of all components can be adapted to current requirements through live vertical upscaling.

1&1 IONOS was established in 2018 after the merger of 1&1 Internet and Berlin-based IaaS provider ProfitBricks and is part of the listed United Internet AG.

Greifswalder Str. 207
10405 Berlin, Germany
TEL +49 30 57700-850

E-MAIL enterprise-cloud@ionos.com

WEB https://www.ionos.com/

TWITTER twitter.com/ionosCLOUD_IAAS

About the Authors

René Büst
Senior Analyst & Cloud Practice Lead
rene.buest@crisp-research.com

René Büst is a Senior Analyst and Cloud Practice Lead at Crisp Research AG, focusing on cloud computing, IT infrastructures, open source and the Internet of Things. Previously, he was Principal Analyst at New Age Disruption and a member of the worldwide Gigaom Research Analyst Network. René Büst is a top cloud computing blogger in Germany and one of the top 50 bloggers in the world in this field. In addition, he is one of the world's top cloud-computing influencers and among the top 100 cloud-computing experts on Twitter and Google+. Since the mid-1990s, René Büst has been focusing on the strategic use of information technology in companies and looking at the influence of IT on our society and disruptive technologies.

René Büst is the author of numerous specialist articles on cloud computing and technology. He regularly writes for prestigious IT publications such as Computerwoche, CIO magazine, LANline and Silicon.de, and is quoted by German and international media – including The New York Times, Forbes Magazine, Handelsblatt, Frankfurter Allgemeine Zeitung, Wirtschaftswoche, Computerwoche, CIO, Manager Magazin and Harvard Business Manager. In addition, René Büst is a speaker and a participant in panel discussions. He is the founder of CloudUser.de and writes about cloud computing, IT infrastructures, technologies, management and strategies. René Büst has a degree in computer engineering from Bremen University of Applied Sciences and an M.Sc. in IT management and information systems from FHDW Paderborn.

Maximilian Hille
Analyst
max.hille@crisp-research.com

Maximilian Hille is an analyst at the IT research and consulting firm Crisp Research AG. Maximilian Hille is responsible for market-research initiatives and consulting projects, primarily in the fields of mobile business and enterprise mobility. Previously, he was Research Manager in the Cloud Computing & Innovation Practice of Experton Group AG.

He is also Product Manager of the research web service Crisp Analytics. His specialist fields are mobile user experience, mobile application performance, mobile development platforms, enterprise mobility and mobile collaboration.

Maximilian Hille was a juror at the 2015 Global Mobile Awards.

Björn Böttcher
Freier IT-Berater
kontakt@parallel-ing.net

Björn Böttcher worked as a Microsoft Student Partner at Hamburg University of Technology (TUHH), and began actively addressing information-technology implementations in many projects in industry as a freelancer while still a student. After graduating in computer engineering, he studied for a doctorate in the cloud and grid environment at the Parallel Computing Group at the TUHH. As well as teaching computer engineering, cloud computing and Big Data at the TUHH, he currently works as an IT consultant on projects in the fields of enterprise mobility, cloud computing and Big Data. In the community sector, he is the founder and organizer of the AWS User Group and CloudHafen Hamburg, a member of the SNIA Cloud Storage Technical Work Group (TWG) and a co-organizer of CloudCamp Hamburg 2010. Furthermore, he writes expert articles on cloud computing, enterprise mobility and Big Data for publications such as Computerwoche. With more than 10 years of professional experience in IT and a scientific background, he focuses on modern solutions with practical benefits.

Copyright

Produced on behalf of:

1&1 IONOS Cloud GmbH
Greifswalder Str. 207
10405 Berlin, Germany

E-MAIL enterprise-cloud@ionos.com
WEB https://www.ionos.com/

Design, layout & infographics:

Hellwig & Buntenbruch
E-MAIL info@hellundbunt.de
WEB hellundbunt.de

Weissenburgstrasse 10
D-34117 Kassel
TEL +49 561 2207 – 4080
FAX +49 561 2207 – 4081

E-MAIL info@crisp-research.com
WEB crisp-research.com clouduser.de
TWITTER twitter.com/crisp_research

crisp RESEARCH