



Testing and Comparing the Performance of Cloud Service Providers Using a Service Broker Architecture

Divyaa Manimaran Elango¹, Frank Fowley¹, and Claus Pahl²(✉)

¹ IC4, Dublin City University, Dublin, Ireland

² SwSE, Free University of Bozen-Bolzano, Bolzano, Italy
Claus.Pahl@unibz.it

Abstract. Service brokers are tools that allow different individual service providers to be integrated. An API can be a mechanism to provide a joint interface. Broker can actually also be use for more than integration. We use a cloud service broker that implements a multi-cloud abstraction API in order to carry out performance comparisons between different cloud services. The broker tool here is a multi-cloud storage API that integrates a number of provided storage services. The library supporting the API is organised into three services, which are a file, a blob and a table service. Using this broker architecture, we developed a performance test scenario to compare the different providers, i.e., to compare a range of storage operations by different providers.

1 Introduction

Integration is a key problem in the cloud services context. A cloud service broker is an intermediary application between a client and cloud provider service that can provide this integration [15]. Brokerage reduces the need for service consumers to analyze different types of services by different providers [1]. This enables a single platform to offer the client a common cloud storage service. This results in cost optimization and reduced level of back-end data management requirements.

For our performance evaluation, we use here a cloud service broker that implements a multi-cloud abstraction API. This multi-cloud storage broker supports GoogleDrive, DropBox, Microsoft Azure and Amazon Web Services as the provided storage services. The API library offers file, blob and table services. The API can facilitates the distribution of different types of cloud provider services [16]. The abstraction library allows the cloud broker to adapt to a rapidly changing marketplace.

Vendor lock-in is often referred to as a critical point in choosing a provider. In order to avoid lock-in, a broker can help. A multi-cloud abstraction library is a suitable mechanism that it makes it easy for the client to switch between cloud providers with different services that are supported by the broker.

Switching or migrating between providers can be driven by quality [27,32]. We use a broker implementation to compare the supported services [8,13,23] from a performance perspective [42]. Service brokers normally remedy interoperability problems [2–5,10]. However, based on this architecture, we look into other service qualities, namely performance which is of key importance for all cloud layers [28–30]. A performance test application was developed here to compare between the services provided through the broker [19]. The performance test scenario was used to compare a range storage operations across different supported providers.

2 Broker – Principles and Supported Services

In Fig. 1 we have outlined the core components of the broker architecture. We also discuss the storage services supported by it in this section.




Fig. 1. Architecture of a multi-cloud storage broker.

2.1 Principle Properties of the Storage Broker

Cloud services are generally provided with specifications, but often constructed in a way that makes them hard to be used as part of a common interface, thus impeding on interoperability. We studied different multi-cloud libraries, including Apache Jclouds, DeltaCloud, Kloudless, SecureBlackBox, and SimpleCloud. The purpose was to adopt a successful solution template.

We decided to adopt an approach similar to the Apache Jclouds library for abstraction, as we will explain now. Jclouds [22] provides cloud-agnostic abstraction. A single instance context approach for the mapping of a user request in jclouds was used in our implementation. The purpose of having each class for each provider across different levels of service was adopted from a similar design in the SecureBlackBox library. Our concept of including a manager interface layer at each component level is adopted from LibCloud, another library.

2.2 Services Supported by the Broker

We have included storage services from Google, Dropbox, Azure and Amazon in our cloud storage broker.

- *Amazon Web Service S3* [35] is a file storage service which is built on REST and SOAP. Their SDK is available in all major development languages.
- *Azure Storage* [37] supports blob, file, queue and table services. The API is built on REST, HTTP and xml, and can be integrated with Microsoft visual studio, eclipse and GIT. The Azure SDK provides a separate API package for each service and has the same code flow across different service APIs.
- *DropBox* [36] is a file hosting service. It also enables synchronised backup and web sharing. The DropBox API is very light-weight and easy for a new user.
- *GoogleDrive* [38] offers a cloud file storage service. The GoogleDrive service includes access to a Google API client library.

This selection of service providers resulted in a grouping of the cloud providers and their services as shown in the table below¹ that summarises the main features of the services:

| Service | Azure | AWS | Google | DropBox |
|---------|---------------------------|--------------------|-------------|---------|
| File | Storage file | - | GoogleDrive | DropBox |
| Blob | Storage blob | AWS S3 | - | - |
| Table | Storage Table, DocumentDB | DynamoDB, SimpleDB | - | - |

3 Broker Architecture

Portability and interoperability are the key objectives of a cloud brokerage tool. Thus the objective of designing and developing an abstraction API is to produce an effective cross-service cloud delivery model [14]. Before describing how we use this to support performance evaluation, we still need to introduce the architectural principles. The main service-oriented functionalities of cloud providers are compute nodes, data volume, load balance, DNS and so on. The advantage of bringing these functionalities to a multi-cloud application provides (1) an easy way of importing and exporting data, (2) choice over price, (3) enhanced SLA, and (4) the elimination of vendor lock-in.

As concept and function integration is the key difficulty in constructing the broker, this broker implementation is based on an ontology that at conceptual level defines the integration. This Storage Abstraction Ontology describes the common naming and meaning approach of the abstraction API [25, 26, 34]. The model consists of four main layers, namely Service, Provider, Level-2 (composite storage objects) and Level-1 (core storage objects).

- *Level-4 Service*: The Service layer is the top layer and is directly integrated to the user interface layer. This layer basically describes the services that the multi-cloud storage abstraction API supports. There are three services currently supported. They are Blob, Table and File service.
- *Level-3 Provider*: The Provider layer is the second layer, which is one of the context object parameters mapped to the service layer. The multi-cloud storage abstraction supports four main providers, namely Microsoft Azure,

¹ <https://www.google.com/drive/>;
<https://www.dropbox.com/>;
<https://aws.amazon.com/s3/>;
<https://azure.microsoft.com/en-us/services/storage/>;

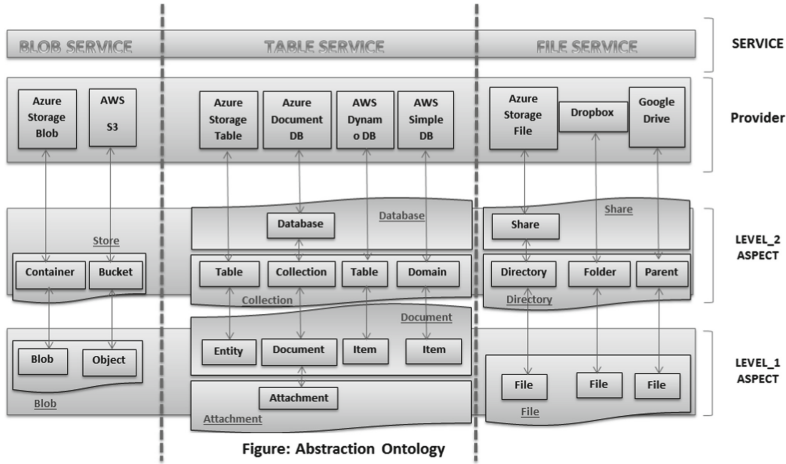


Fig. 2. Ontology-based layered broker architecture.

Amazon Web Services, GoogleDrive and DropBox. The corresponding services supported by the providers are shown below:

| Service | Provider |
|---------|---|
| Blob | Azure storage blob; AWS S3 |
| Table | Azure storage table; Azure DocumentDB; AWS DynamoDB; AWS SimpleDB |
| File | Azure storage file; DropBox and GoogleDrive |

- *Level-2 Composite:* Level-2 is the next layer. This layer represents the first level or higher level of composite object abstraction. This layer is service-neutral and brings out the common naming across the providers specific functionalities. Each layer is abstracted based on the common operations and aspect of how the main function is applied in that particular service. Common naming is represented to easily categorise storage resources and group them to make the development of the coding easier.

Based on the Abstraction Ontology Fig. 2, the Blob service has Store which groups Container from Azure Storage Blob and Bucket from AWS S3. The Table service has two different sub-layers - where Database belongs to Azure DocumentDB Database, and where Collection groups Table from Azure Storage Table, collection from Azure DocumentDB Collection, Table from AWS DynamoDB and Domain from AWS SimpleDB. The File service has two different sub-layers where Share belongs to Azure Storage File and Directory groups Directory from Azure Storage File, Folder from DropBox and Parent from the GoogleDrive service.

- *Level-1 Core:* Level-1 represents the lower level of core object abstraction. This layer contains the core functionalities of a particular service across different providers. The classes in this level are extended from an abstract class called AbstractConnector. The class implements the abstract methods defined in the AbstractConnector class. The mapping from Level-2 to Level-1

is performed by an interface class called *Manager*. This *Manager* identifies the provider class by its key. Basic CRUD operations on the storage resources are included as core methods. In order to achieve these functions, each operation request should pass through the Level-2 mappings and are then mapped across the service and providers.

The *Blob* service has *Blob* which groups *Blob* from Azure Storage *Blob* and *Object* from AWS *S3*. The *Table* service has two different sub layers. It has *Item* which groups *Entity* from Azure Storage *Table*, *Document* from Azure *DocumentDB* *Document*, *Item* from AWS *DynamoDB* and *Item* from AWS *SimpleDB*. Also, the second sub layer *Attachment* belongs to Azure *DocumentDB*. The *File* service has *File* which groups *File* from Azure Storage *File*, *File* from *DropBox* and *File* from *GoogleDrive*.

4 Performance Testing and Provider Comparison

We have used the broker to compare performance values for the four providers selected. The broker is instrumented to provide the response time results.

In this section, we describe the performance test set-up and the results for the three service types *blob*, *file* and *table* across the different providers. We organise this section based on the storage types *blob*, *file* and *table*.

Not all providers support each of the storage types. So, the number of compared services provided varies between two and four. We report on the time consumed for a number of standard operations at the two important levels 1 and 2 of the layer architecture. In this way, we cover individual objects (items) and composites (collections) and a range of standard operations on them such as creating or deleting.

4.1 Blob Service Performance Test

The *Blob Service Performance Test* was performed on two providers, namely Azure Storage *Blob* and AWS *S3*. This performance test includes two object levels. Level-2 represents *Store* (which includes *container* and *Bucket*). Level-1 represents *Blob* (which includes *Blob* and *Object*).

- The total number of tests performed was 27 to fully cover the respective core and composite objects and the different relevant operations on them. The performance test compares the operations across the service providers.
- Each operation was run 10 times in order to avoid any accidental performance irregularities due to external factors, and the corresponding process time for each request from T1 to T10 was calculated.
- Each request was processed with the same blob size of 10.2 MB, which resembles a standard object size.

The result includes start time, end time, average time and total duration – see Figs. 3 and 4.




Fig. 3. Blob service Level-2 composite object.




Fig. 4. Blob service Level-1 core object.

4.2 File Service Performance Test

The *File Service Performance Tests* were performed on Azure Storage File, GoogleDrive and DropBox. The tests include two object levels. Level-2 represents Share and Directory. Level-1 represents Files.

- The total number of tests performed was 26 to cover the combinations of different object types and different operations on them.
- The performance tests compare the operations across the service providers. Each operation was run 10 times to eliminate irregular single behaviour, and the corresponding process time for each request from T1 to T10 was calculated.
- Each request was processed with same file size of 10.2 MB as a common size for the object type in question.

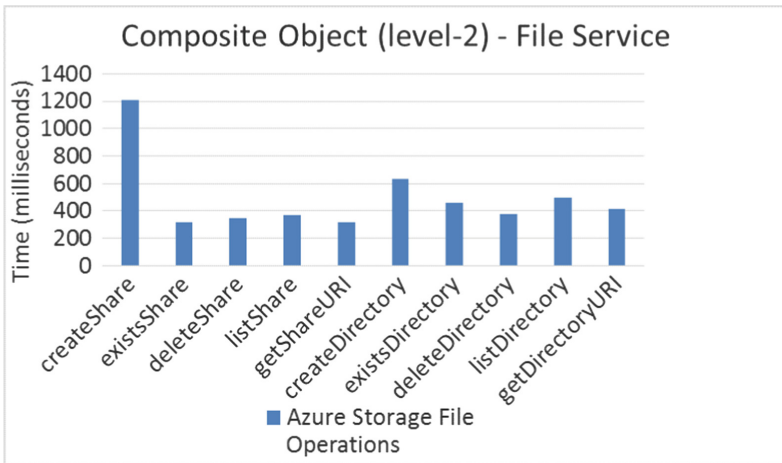


Fig. 5. File service Level-2 composite object.




Fig. 6. File service Level-1 core object.

The results include start time, end time, average time and total duration – see Figs. 5 and 6.

4.3 Table Service Performance Test

The *Table Service Performance Tests* were performed on Azure Storage Table, Azure DocumentDB, AWS DynamoDB and AWS SimpleDB. The Tests include two object levels. Level-2 represents Database and Collections (which includes Table, Collections, Table and Domain). Level-1 represents Item (which includes Entity, Document, Table Item, Domain Item) and Attachment.

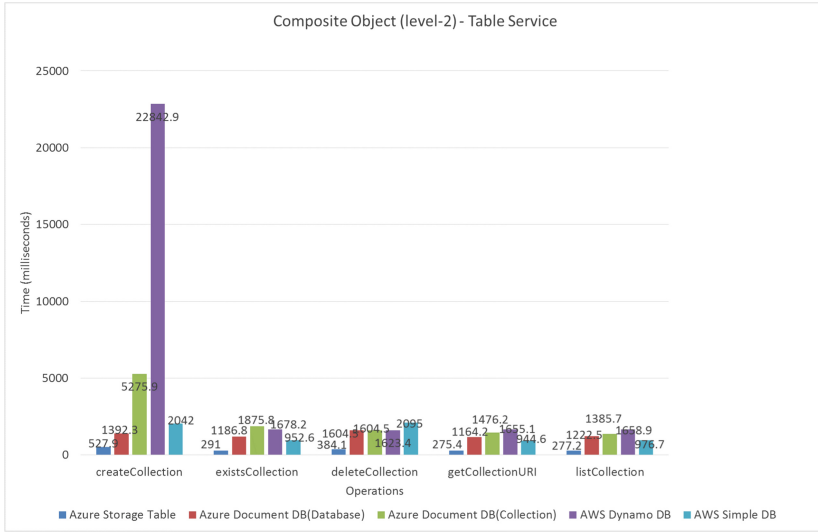


Fig. 7. Table service Level-2 composite object.

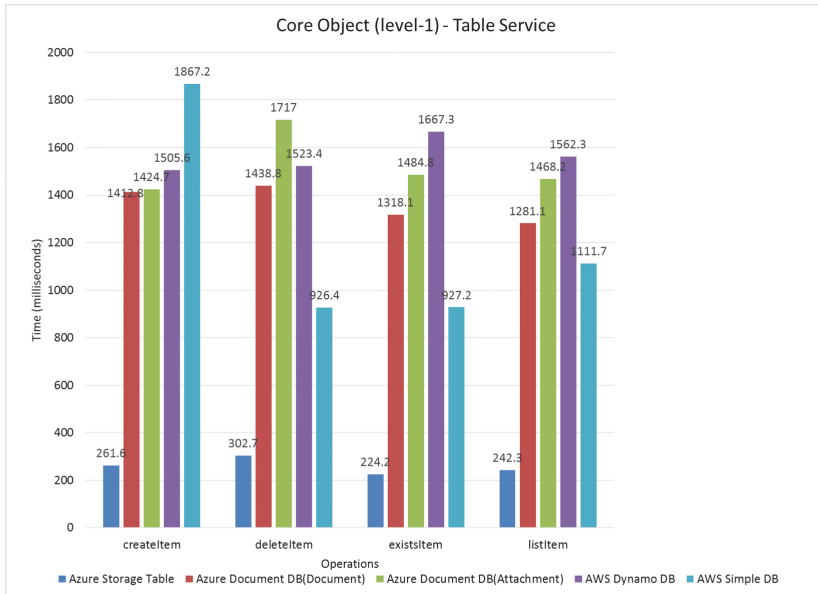


Fig. 8. Table service Level-1 core object.

- The total number of tests was 45. As already explained for the blob tests, this number covers the combination of different objects and the different operations on them. The performance tests compare the operations then across the service providers.

5 Discussion of Results and Conclusions

The aim of cloud service brokerage is customising or integrating existing services or making them interoperable. We have developed what based on common classification schemes in [11, 12, 39] is categorised as an integration broker. The purpose of a broker is intermediation between consumers and providers to provide advanced capabilities (interoperability and portability [33]) that builds up on an intermediary/broker platform to provide for instance a marketplace to bring providers and customers together and automatically facilitate multi-provider usage or portability across providers. The broker for cloud storage service providers implement a joint interface to allow

- easy portability for the user and
- easy extensibility for the broker provider.

This broker solution enables through the joint API also the opportunity for a cloud storage user to easily migrate between service providers and evolve the systems [9, 21], without having sufficient standards [6, 7, 20].

We investigated here the usage of the broker to carry out comparative performance tests across the providers in order to support the user with the decision which provider to choose, if this is taken based on a performance criterion.

Our observations from the performance tests we described earlier are the following:

- (a) Core Objects: We can observe that the performance of core object storage operations varies significantly across Cloud providers. Azure outperforms AWS S3 by a factor of between 4 and 5 in our test scenario. For individual object operations, Azure is also up to 5 times faster in terms of access speed. For example, the common function of UploadBlob takes approximately 4 seconds on Azure and 10 seconds on AWS S3 for a 10.2 MB file.
- (b) Composite Objects: The tests of composite object operations [31] that relate to collections show that Azure has significantly more access performance than other providers. In particular, AWS DynamoDB has a unusually long access time for its CollectionCreate operation. The tests on individual table entity operations show Azure to be the fastest by a considerable margin with over 5 to 6 times lesser access speeds on average.
- (c) Upload and Download: The average of the combined file upload and download speeds do not vary considerably across the providers tested.

We have defined some parameters, such as object size, in a specific way. Other choices might result in different observations. Our aim here was not to recommend a particular provider. The aim was to demonstrate the usefulness of instrumenting brokers for either decision making or as an ongoing monitoring approach. Any selection can anyway not happen without considering other properties such as security.

In the future, we plan to consider more storage services. Furthermore, the impact of different architectures in terms on IaaS or PaaS with and without the use of container technologies [17, 18, 24, 40, 41] shall be explored.

Acknowledgements. This work was partly supported by IC4 (Irish Centre for Cloud Computing and Commerce), funded by EI and the IDA.

References

1. Ried, S.: Cloud Broker - A New Business Model Paradigm. Forrester (2011)
2. Elango, D.M., Fowley, F., Pahl, C.: An ontology-based architecture for an adaptable cloud storage broker. In: *Advances in Service-Oriented and Cloud Computing*. Springer CCIS (2018, to appear)
3. Benslimane, D., Dustdar, S., Sheth, A.: Services mashups - the new generation of web applications. *Internet Comput.* **12**(5), 13–15 (2008)
4. Bernstein, D., Ludvigson, E., Sankar, K., Diamond, S., Morrow, M.: Blueprint for the inter-cloud: protocols and formats for cloud computing interoperability. In: *International Conference on Internet and Web Applications and Services* (2009)
5. Buyya, R., Ranjan, R., Calheiros, R.N.: InterCloud: utility-oriented federation of cloud computing environments for scaling of application services. In: Hsu, C.-H., Yang, L.T., Park, J.H., Yeo, S.-S. (eds.) *ICA3PP 2010*. LNCS, vol. 6081, pp. 13–31. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13119-6_2
6. Cloud Standards (2017). <http://cloud-standards.org/>
7. ETSI Cloud Standards (2017). <http://www.etsi.org/newsevents/news/734-2013-12-press-release-report-on-cloudcomputing-standards>
8. Fehling, C., Mietzner, R.: Composite as a service: cloud application structures, provisioning, and management. *Inf. Technol.* **53**(4), 188–194 (2011)
9. Pahl, C., Jamshidi, P., Weyns, D.: Cloud architecture continuity: change models and change rules for sustainable cloud software architectures. *J. Softw. Evol. Process* **29**(2) (2017)
10. Pahl, C., Jamshidi, P., Zimmermann, O.: Architectural principles for cloud software. In: *ACM Transactions on Internet Technology*. (2018, to appear)
11. Fowley, F., Pahl, C., Zhang, L.: A comparison framework and review of service brokerage solutions for cloud architectures. In: *1st International Workshop on Cloud Service Brokerage* (2013)
12. Fowley, F., Pahl, C., Jamshidi, P., Fang, D., Liu, X.: A classification and comparison framework for cloud service brokerage architectures. *IEEE Trans. Cloud Comput.* (2017). <https://doi.org/10.1109/TCC.2016.2537333>. <http://ieeexplore.ieee.org/document/7423741/>
13. Garcia-Gomez, S., et al.: Challenges for the comprehensive management of cloud services in a PaaS framework. *Scalable Comput. Pract. Experience* **13**(3), 201–213 (2012)
14. Elango, D.M., Fowley, F., Pahl, C.: Pattern-driven architecting of an adaptable ontology-driven cloud storage broker. In: *University of Oslo, Department of Informatics, Research report 471*, pp. 33–47 (2017)
15. Gartner: Cloud Services Brokerage. Gartner Research (2013). <http://www.gartner.com/it-glossary/cloud-servicesbrokerage-csb>
16. Grozev, N., Buyya, R.: InterCloud architectures and application brokering: taxonomy and survey. *Softw. Pract. Experience* **44**(3), 369–390 (2012)
17. Pahl, C., Jamshidi, P.: Microservices: a systematic mapping study. In: *Proceedings CLOSER Conference*, pp. 137–146 (2016)
18. Taibi, D., Lenarduzzi, V., Pahl, C.: Processes, motivations and issues for migrating to microservices architectures: an empirical investigation. *IEEE Cloud Comput.* (2018). Accepted for publication

19. Hofer, C.N., Karagiannis, G.: Cloud computing services: taxonomy and comparison. *J. Internet Serv. Appl.* **2**(2), 81–94 (2011)
20. IEEE Cloud Standards (2015). <http://cloudcomputing.ieee.org/standards>
21. Jamshidi, P., Ahmad, A., Pahl, C.: Cloud migration research: a systematic review. *IEEE Trans. Cloud Comput.* **1**(2), 142–157 (2013)
22. jclouds: jclouds Java and Clojure Cloud API (2015). <http://www.jclouds.org/>
23. Ferrer, A.J., et al.: OPTIMIS: a holistic approach to cloud service provisioning. *Future Gener. Comput. Syst.* **28**(1), 66–77 (2012)
24. Gacitua-Decar, V., Pahl, C.: Structural process pattern matching based on graph morphism detection. *Int. J. Softw. Eng. Knowl. Eng.* **27**(2), 153–189 (2017)
25. Pahl, C.: Layered ontological modelling for web service-oriented model-driven architecture. In: Hartman, A., Kreische, D. (eds.) *ECMDA-FA 2005*. LNCS, vol. 3748, pp. 88–102. Springer, Heidelberg (2005). https://doi.org/10.1007/11581741_8
26. Pahl, C., Giesecke, S., Hasselbring, W.: Ontology-based modelling of architectural styles. *Inf. Softw. Technol.* **51**(12), 1739–1749 (2009)
27. Pahl, C., Xiong, H.: Migration to PaaS clouds - migration process and architectural concerns. In: *IEEE 7th International Symposium on the Maintenance and Evolution of Service-Oriented and Cloud-Based Systems MESOCA* (2013)
28. Konstantinou, A.V., Eilam, T., Kalantar, M., Totok, A.A., Arnold, W., Snibler, E.: An architecture for virtual solution composition and deployment in infrastructure clouds. In: *International Workshop on Virtualization Technologies in Distributed Computing* (2009)
29. Jamshidi, P., Sharifloo, A., Pahl, C., Arabnejad, H., Metzger, A., Estrada, G.: Fuzzy self-learning controllers for elasticity management in dynamic cloud architectures. In: *12th International ACM SIGSOFT Conference on Quality of Software Architectures QoSA* (2016)
30. Arabnejad, H., Jamshidi, P., Estrada, G., El Ioini, N., Pahl, C.: An auto-scaling cloud controller using fuzzy Q-learning - implementation in openstack. In: Aiello, M., Johnsen, E.B., Dustdar, S., Georgievski, I. (eds.) *ESOCC 2016*. LNCS, vol. 9846, pp. 152–167. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-44482-6_10
31. Mietzner, R., Leymann, F., Papazoglou, M.: Defining composite configurable SaaS application packages using SCA. In: *International Conference on Internet and Web Applications and Services, Variability Descriptors and Multi-tenancy Patterns* (2008)
32. Pahl, C., Xiong, H., Walshe, R.: A comparison of on-premise to cloud migration approaches. In: Lau, K.-K., Lamersdorf, W., Pimentel, E. (eds.) *ESOCC 2013*. LNCS, vol. 8135, pp. 212–226. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40651-5_18
33. Petcu, D., et al.: Portable cloud applications - from theory to practice. *Future Gen. Comput. Syst.* **29**(6), 1417–1430 (2013)
34. Javed, M., Abgaz, Y.M., Pahl, C.: Ontology change management and identification of change patterns. *J. Data Semant.* **2**(2–3), 119–143 (2013)
35. Amazon Simple Storage Service (S3) Cloud Storage AWS. <https://aws.amazon.com/s3/>
36. Dropbox. <https://www.dropbox.com/>
37. Azure Storage - Secure cloud storage. <https://azure.microsoft.com/en-us/services/storage/>
38. Google Drive - Cloud Storage & File Backup. <https://www.google.com/drive/>
39. Jamshidi, P., Pahl, C., Mendonca, N.C.: Pattern-based multi-cloud architecture migration. *Softw. Pract. Experience* **47**(9), 1159–1184 (2017)

40. Pahl, C., Brogi, A., Soldani, J., Jamshidi, P.: Cloud container technologies: a state-of-the-art review. *IEEE Trans. Cloud Comput.* (2017). <https://doi.org/10.1109/TCC.2017.2702586>. <http://ieeexplore.ieee.org/document/7922500/>
41. Aderaldo, C.M., Mendonca, N.C., Pahl, C., Jamshidi, P.: Benchmark requirements for microservices architecture research. In: 1st International Workshop on Establishing the Community-Wide Infrastructure for Architecture-Based Software Engineering. *IEEE* (2017)
42. Heinrich, R., van Hoorn, A., Knoche, H., Li, F., Lwakatare, L.E., Pahl, C., Schulte, S., Wettinger, J.: Performance engineering for microservices: research challenges and directions. In: *Proceedings of the 8th ACM/SPEC on International Conference on Performance Engineering Companion* (2017)