

# The Evolution of Cloud Computing Towards a Vendor Agnostic Market Place Using the SKY CONTROL Framework

Henry-Norbert Cocos<sup>a</sup>, Christian Baun<sup>b</sup> and Martin Kappes<sup>c</sup>

Frankfurt University of Applied Sciences, Nibelungeplatz 1, Frankfurt am Main D-60318, Germany  
{cocos, christianbaun, kappes}@fb2.fra-uas.de

Keywords: Sky Computing, Multi-Cloud, Vendor-Agnostic Framework, Workload Placement.

Abstract: Multi-cloud environments offer benefits like vendor diversification and resilience but pose challenges such as increased management complexity, lack of cost transparency, and compliance. This concept paper introduces SKY CONTROL, a vendor-agnostic framework for small and medium-sized enterprises (SMEs). SKY CONTROL integrates cost control and risk management into multi-cloud setups, providing static and dynamic resource analyses, a cost calculator, and risk assessment tools. By leveraging the Sky Computing paradigm, SKY CONTROL simplifies resource orchestration and enhances security. This novel framework is the first implementation of the innovative Sky Computing concept. It aims to improve cost efficiency, regulatory compliance, and strategic IT planning for SMEs, offering a unified approach to managing hybrid infrastructures.

## 1 INTRODUCTION

Cloud Computing, based on distributed systems (van Steen and Tanenbaum, 2017), integrates virtualization and modern web technologies to provide scalable IT infrastructure, platforms, and applications as on-demand services. The National Institute of Standards and Technology (NIST) defines Cloud Computing by five key properties (Mell et al., 2011):

- **On-demand Self-Service:** Automatic resource provisioning.
- **Broad Network Access:** Accessibility via standard network interfaces.
- **Resource Pooling:** Shared, scalable resources.
- **Rapid Elasticity:** Dynamic scalability.
- **Measured Service:** Usage-based billing.

These features facilitate seamless cloud integration, reducing costs and effort. Figure 1 illustrates the NIST model, covering service properties, models, and deployment types.

Since the late 2000s, IT infrastructures have shifted from on-premise to multi- and hybrid cloud systems (Hong et al., 2019; Gundu et al., 2020).

By the mid-2010s, multi-cloud strategies, combining services from multiple providers, became common (Jamshidi et al., 2016). Benefits include (Hong et al., 2019; Georgios et al., 2021; Petcu, 2013):

- Reduced provider dependency
- Cost optimization
- Load balancing
- Business continuity through redundancy
- Free service selection
- Enhanced security via data diversification

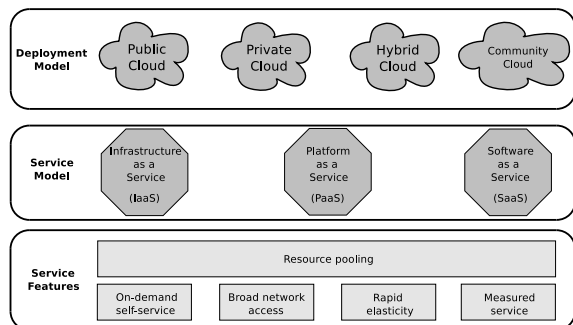


Figure 1: NIST Definition of Cloud Computing.

Multi-cloud strategies foster innovation by integrating diverse platforms. Hybrid models use infrastructure-as-a-service (IaaS) to minimize administrative tasks and software-as-a-service (SaaS) for more straightforward configuration and integration.

<sup>a</sup> <https://orcid.org/0009-0001-7573-0361>

<sup>b</sup> <https://orcid.org/0009-0004-9955-3752>

<sup>c</sup> <https://orcid.org/0000-0002-8768-8359>

However, managing multi-cloud environments remains complex, with interoperability a significant challenge. Sky Computing aims to simplify this by abstracting cloud resources across providers, enabling optimal service selection. Though still being researched, it has the potential to enhance multi-cloud adoption.

This paper explores Sky Computing and its role in advancing Cloud Computing.

## 2 BACKGROUND AND RELATED WORKS

Cloud computing has expanded significantly in recent years, leading to diverse public and private offerings. Based on the hybrid cloud model, multi-cloud environments have become common, enabling companies to avoid vendor lock-in and adopt flexible strategies (Mulder, 2020). These environments distribute services across multiple providers, offering economic and organizational benefits, such as flexibility and increased resource availability. However, challenges include service transparency, interoperability issues between APIs, and the complexity of managing diverse cloud services (Ardagna, 2015; Barker et al., 2015).

Sky computing aims to address these challenges by introducing an abstraction layer over CSP offerings, enabling uniform service provisioning regardless of location (Fortes, 2010; Monteiro et al., 2014; Yang et al., 2023). The SkyPilot project (Wei et al., 2021) at Berkeley is a pioneering effort in this domain. It features an intercloud broker optimized for machine learning workloads across multiple providers. Despite its innovative approach, SkyPilot remains experimental, with limited cloud service diversity and a focus on IaaS environments.

This raises questions about applying Sky Computing to diverse workloads and its potential to generalize across services. We explore Sky Computing opportunities in section 3 and critically assess its demands on the current cloud computing landscape in section 4.

## 3 SKY COMPUTING

Sky Computing is a new conceptual vision for Cloud Computing that introduces an intercloud broker for the mediation of services. The Sky Computing concept (Stoica and Shenker, 2021) proposes an additional abstraction layer between the cloud services from the providers (Amazon Web Services, Google

Cloud Platform, etc.) and the workloads of the end users. Sky Computing aims to complete the abstraction of cloud resources from different providers so that applications and users can access these resources without worrying about where the resources or services are located in the individual clouds. For this reason, the term "*cloud of clouds*" is also used, as the additional abstraction of resources leads to creating a uniform and interoperable cloud and thus includes several individual cloud providers. This abstraction layer is called an intercloud broker (Yang et al., 2023).

One idea for solving the problem of heterogeneity in cloud services and their respective APIs would be an effort for standardization across cloud vendors. As Chasins et al. (Chasins et al., 2022) describe in their work, introducing standardization in the field of Cloud Computing poses risks to future developments since it would slow down or even stop innovation by locking the ecosystems into a set of interfaces that may not be suitable for future applications. Many examples in the past have shown that standardization efforts have failed. One example is the introduction of the Common Object Request Broker Architecture (CORBA) by the Object Management Group (OMG) (Siegel, 1999). In an attempt to standardize the communication of different distributed objects on various operating systems, this effort posed strict limitations on the implementation of distributed systems since it was designed as an interfacing technology for other programming languages and imposed harsh conventions on the usage of the standard. Therefore, the standard adoption in the industry was poorly met, and it was doomed to become a niche technology. This example demonstrates that standardization is not the solution to the interoperability problem. Sky Computing aims to alleviate this challenge by introducing the intercloud broker, which acts as a common platform for different cloud services. It focuses on the workloads and types and places them in the appropriate service on the vendor side.

Sky Computing redefines cloud usage by routing jobs through an intercloud broker, which selects and manages cloud services for execution. This creates a two-sided market, linking user-submitted jobs with cloud services. Some services are multi-cloud (e.g., Kubernetes, Apache Spark), while others are cloud-specific (e.g., AWS Inferentia) (Chasins et al., 2022). Unlike traditional multi-cloud setups, Sky Computing allows workloads to run on multiple clouds or split across them, offering greater flexibility and benefits. It focuses on partial compatibility, enabling many—but not all—jobs to run across clouds, with compatibility improving over time.

A characteristic of Sky Computing is its dis-

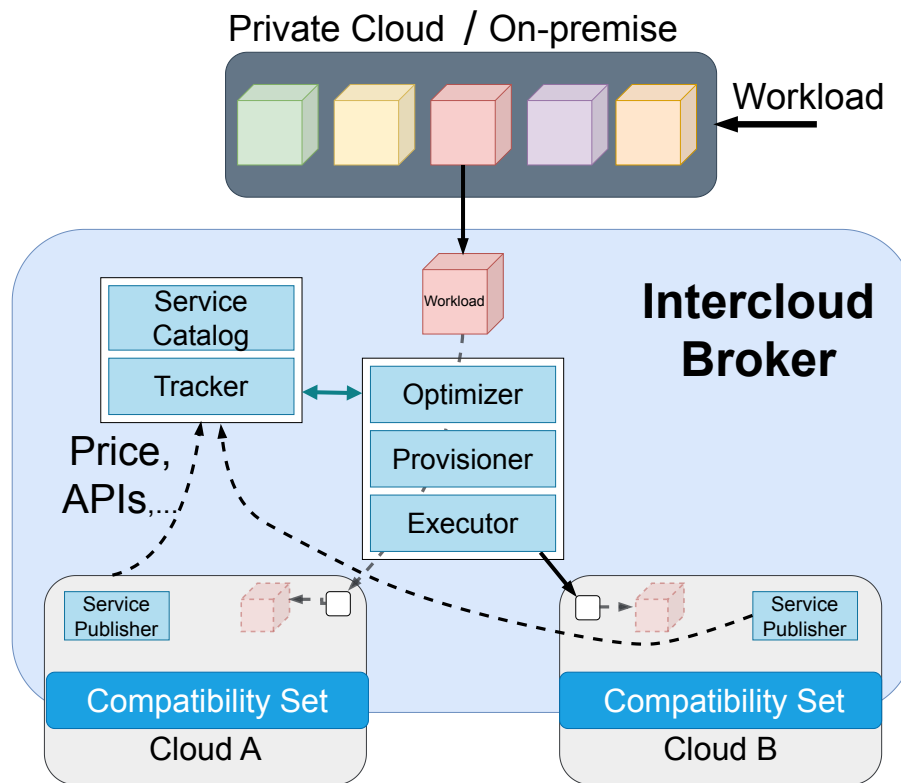


Figure 2: Overview of the components of the inter-cloud broker (Yang et al., 2023).

tributed infrastructure, which dynamically allocates workloads across providers. This ensures scalability and independence from workload type or provider, maximizing resource utilization. Figure 2 visually represents the Sky Computing framework, illustrating how it integrates cloud resources into a cohesive and scalable ecosystem.

The Service Catalog captures the instances and services available in each cloud, detailed information about locations that offer them, and the APIs for allocating and accessing them. It also stores the long-term pricing for on-demand virtual machines, data storage, egress, and services (typically, these prices stay the same for months). The term refers to all prices for a fixed time horizon. Cloud computing services are highly time-dependent, meaning that the time interval of the service offered is specified in advance (Ibrahimi, 2017). The Service Catalog can provide filtering and search functionality based on information published by the cloud providers, listed by a third party, or collected by the broker.

The tracker tracks spot prices (which may change more frequently, e.g., hourly or daily) and the availability of resources across different cloud providers and locations. This module is of central importance as the prices of cloud providers and the associated ser-

vices are a decision metric for service placement.

The Optimizer processes the workload requirements and checks the availability of instances and services and their prices, which the Service Catalog and Tracker provide. This module then calculates the optimal placement of the services. In a broader context, the optimality of workload placement is heavily dependent on the used metrics. An algorithm using multiple parameters for calculating optimality criteria and decision-making is needed (see section 7). If resource availability and/or price change, the Optimizer can perform a new optimization.

The Provisioner module manages the resources by allocating the resources required for execution. The Optimizer’s execution plan allocates the resources accordingly and releases them when each task is completed. The executor manages the application by aggregating the functions of each workload and executing them based on the resources allocated by the cloud provider and service provider.

Compatibility sets are an essential feature of Sky Computing. The focus is on using existing services and APIs from all cloud providers, intended to offer transparent and standardized options for connecting services without reimplementing them.

Sky Computing provides a reasonable basis for

implementing our proposed framework since it adds an abstraction layer between users and cloud service providers. However, having a tool or component in the Sky Computing concept is still essential for analyzing the costs drawn by the services used.

The extension of the cloud by constructing an intercloud broker poses new challenges that need to be tackled, especially in networking and security. In Sky Computing, tackling the more ambitious task of automated assurance of suitable security levels for all assets is a considerable challenge. A few years ago, this would have been an almost impossible task in heterogeneous environments. However, the concept of SASE - Secure Access Service Edge - was introduced in 2019 and has since evolved into a solid technology that offers new possibilities for our purposes (Islam et al., 2021; MacDonald et al., 2019). SASE is a synergistic combination of techniques for controlling security - i.e., enforcing security policies - in complex, heterogeneous infrastructures.

SASE generally includes the following functionalities:

- **SD-WAN**(Yang et al., 2019) optimizes wide-area networks by allowing organizations to use multiple transport services (MPLS, LTE, 5G, broadband) for secure connectivity(Islam et al., 2021).
- **Secure Web Gateway (SWG)** filters and monitors web traffic to protect users from threats and ensure compliance.
- **Cloud Access Security Broker (CASB)** enforces security policies between cloud users and providers.
- **Firewall-as-a-Service (FWaaS)** provides scalable, cloud-based firewall functionalities.
- **Zero Trust Network Access (ZTNA)** follows the *never trust, always verify* principle, securing user sessions inside and outside corporate networks.

SASE enables centralized policy management with distributed enforcement points, ensuring local decision-making when needed (van der Walt and Venter, 2022). For example, policies can be enforced locally via Customer Premises Equipment (CPE) or managed device agents, enhancing security flexibility.

Integrating SASE into Sky Computing addresses security challenges such as asset and risk-based management, though practical implementation is still evolving. While these concepts promise seamless workload distribution and cloud interoperability, their adoption—especially for SMEs—requires further research. The following section explores industry use cases, potential developments, and open questions.

## 4 PROPOSITION FOR THE TRANSFORMATION OF THE CLOUD COMPUTING MARKET PLACE

Sky Computing offers significant potential for the industry, mainly by providing end-users with more options for optimizing workload operations. The intercloud broker could shift the balance from large cloud vendors to companies, benefiting small to medium-sized enterprises (SMEs). However, Sky Computing must become an attractive, incentive-driven technology for widespread adoption.

A key focus is addressing end-users' needs, especially when deploying multi-cloud environments. Multi-cloud setups offer benefits like preventing vendor lock-in and enabling the best services from various providers. However, managing resources and costs across multiple providers and meeting governance requirements remains a challenge, particularly for smaller companies.

The complexity of multi-cloud environments increases for SMEs, and understanding their challenges is essential for developing a framework that solves these issues. Although research on multi-cloud architectures exists (Baryannis et al., 2013; Kavitha and Radha, 2022), introducing an intercloud broker can help reduce complexity by creating an abstraction layer between end-users and cloud vendors.

In summary, we propose adding an abstraction layer to create a new marketplace for end-users, enabling interaction with vendors through a new platform for negotiation rather than changing the vendors' services or APIs.

## 5 OPEN QUESTIONS

The proposition in Section 4 raises key questions requiring further research. One major challenge is establishing an open and user-friendly marketplace for end-users. While an intercloud broker is a step forward, practical implementation remains unclear. Adoption by cloud vendors is not critical, but end-user acceptance is essential, necessitating a standardized framework for multi-cloud setups.

Another challenge is designing an attractive framework for stakeholders, including executives, technicians, and security auditors. C-level executives require cost insights, technicians need performance data and easy configurations, and auditors prioritize governance and risk visualization. Although the SkyPilot project (Stoica and Shenker, 2021; Yang

et al., 2023) has demonstrated Sky Computing’s viability for ML pipelines, broader industry-driven use cases must be explored. We, therefore, want to start inspecting the placement of IaaS (Infrastructure as a Service) workloads since many companies use this type of service model for migrating legacy applications to the cloud (so-called “Lift-and-Shift” operation) (Ahmad et al., 2018). A key research question is: *How can Sky Computing enhance workload placement efficiency for SMEs compared to multi-cloud setups?*

Efficient workload placement is crucial, requiring automation to minimize manual intervention, especially as workload volumes grow. SMEs with limited budgets and resources would particularly benefit from productivity and cost improvements through Sky Computing. Technologies such as SD-WAN and SASE introduce security and network configuration automation, but integrating computing resources like VMs and container runtimes remains an open research area.

Adoption risks include resistance from cloud service providers (CSPs), who may restrict API access to maintain market control, hindering interoperability. Addressing this requires careful research into business and technical strategies.

This section’s open questions guide the conceptualization of the SKY CONTROL framework, which aims to optimize multi-cloud environments for end-users. As an initial step, SKY CONTROL will be applied to IaaS workloads in SMEs, evaluating its performance against deployments using major CSPs like AWS and Azure. Later, higher-abstraction models like PaaS will be explored. The next section details the SKY CONTROL framework and its core concepts.

## 6 CURRENT DEVELOPMENT IN THIS FIELD

As an outlook and initial attempt, we present the conceptual frameworks to address the challenges discussed in the previous sections. We introduce SKY CONTROL, a framework for SMEs to control and optimize multi-cloud setups. This provides a common ground for companies to onboard Sky Computing. Another concept discussed in this section is the analysis of methods and technologies for attaching workloads to the Sky Computing ecosystem (section 8). In this project, the geographic distribution of workloads across WAN boundaries will be investigated to identify suitable solutions for the technological implementation of an intercloud broker.

## 7 SKY CONTROL: A FRAMEWORK FOR SMEs

The distributed nature of multi-cloud environments offers many benefits for SMEs, but it also brings significant challenges in managing workloads across different CSP platforms. The overview of the costs of the companies’ services is challenging, and there are substantial drawbacks in choosing such architectures. Another critical aspect is the overview and analysis of potential security risks and the governance of assets. SKY CONTROL addresses the challenges SMEs face when using multi-cloud deployments. The following sections present the architecture of SKY CONTROL, along with the desired functionalities needed to fulfill the requirements of SMEs.

### 7.1 Architecture of SKY CONTROL

This section describes the key components of the SKY CONTROL framework:

- **Cost Control:** Analyzes, calculates, and visualizes costs for both on-premise and cloud resources. It performs static analysis (e.g., resource IDs, hardware specs) and dynamic analysis (e.g., CPU/memory usage, network bandwidth). Pricing trends and predictions are derived from this data, considering the complexity of cloud service integration. A control and planning tool provides insights into resource usage across multiple cloud providers, with visualization for easier comprehension.
- **Risk Management:** Manages customer assets, collects detailed insights, and performs risk analysis. It evaluates risks based on asset criticality, data sensitivity, and compliance standards like C5 (Cloud Computing Compliance Criteria Catalogue) (Di Giulio et al., 2017) for German SMEs. This helps businesses meet governance requirements and enhances compatibility with larger enterprises. The module also provides risk and asset visualization for CIOs, aiding audits and risk mitigation.

Figure 3 illustrates the architecture of the proposed framework.

This section highlights the necessity of analyzing pricing, workload criticality, and governance information to define decision criteria for optimal workload placement. To achieve this, an algorithm must be developed to calculate and formalize these decisions within the Optimizer module (as defined in section 3). Given that “optimality” varies across different factors, the SKY CONTROL framework must first establish a

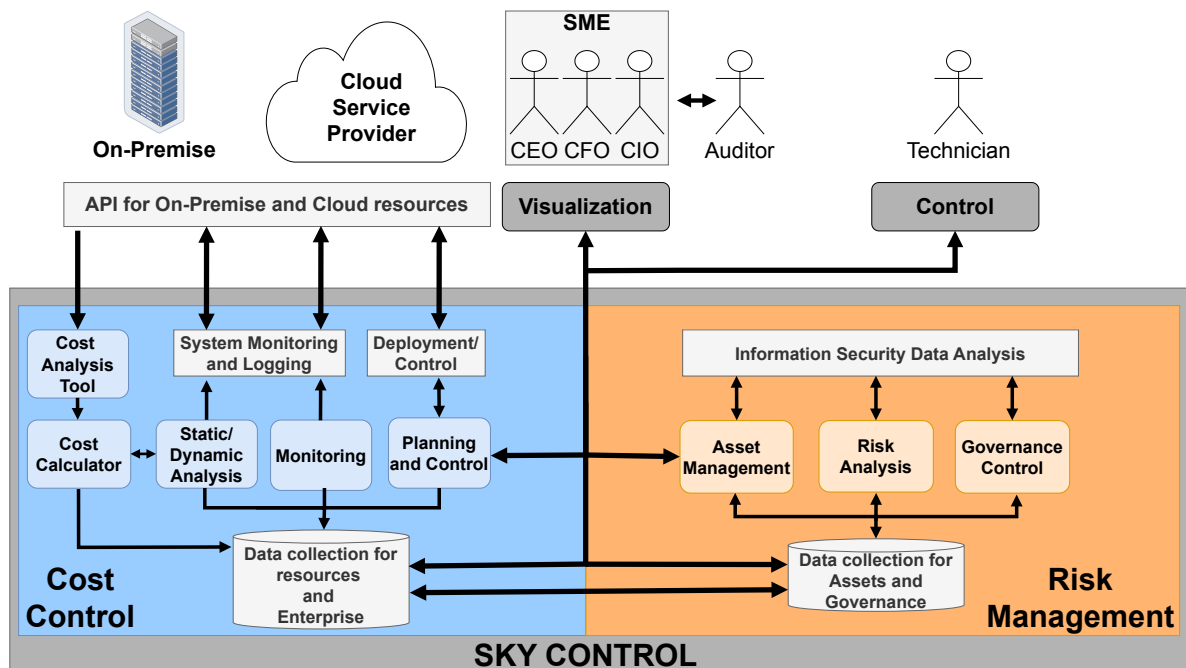


Figure 3: SKY CONTROL architecture.

conceptual foundation for analyzing these criteria before integrating them into the optimization process.

At its core, SKY CONTROL aims to provide transparency into the properties of the entire infrastructure and its components. The project is designed to be highly adaptable, leveraging both existing and emerging technologies. Some of these technologies will inspire SKY CONTROL's development, while others may be directly integrated into its architecture as functional components.

## 8 ATTACHING ON-PREMISES WORKLOADS TO THE SKY

This field of research explores methods for integrating workloads into the Sky Computing ecosystem, focusing on geographic workload distribution across WAN boundaries to develop an intercloud broker. It investigates an offline-first strategy, prioritizing local workload operation while enabling cloud relocation as needed, enhancing service usability and availability.

The emphasis of this sub-project lies in the investigation of geographical scaling, often overlooked in practice, to optimize service proximity to users and enhance performance.

Key resources examined in this strategy include:

- **Computing Capacity** (via VMs, containers)

- **Network Resources** (leveraging SDN virtualization)
- **Storage Resources** (for distributed data synchronization)
- **Software**
- **Services**

The exploration of the outsourcing and cooperation of cloud services with client-side resources, investigating which cloud services can be offloaded to clients and under what conditions is a central point of interest. Unlike the cloud-first approach, an offline-first strategy will prioritize local network service availability, ensuring network-independent cloud services while still enhancing overall functionality through cloud cooperation.

A critical aspect is the scaling of services:

- Vertical scaling (adding resources)
- Horizontal scaling (adding service instances)
- Geographical scaling (placing services closer to users, or "scaling away")

An inter-cloud broker could facilitate seamless interoperability between cloud and on-premise environments using SKY CONTROL mechanisms. This would allow users to access nearby services independently of the cloud, increasing autonomy and resilience without losing cloud benefits.

Key objectives include:

- Defining criteria for outsourcing computing-intensive tasks, an area with no standardized research methodology.
- Investigating service migration (Rodriguez et al., 2021) to optimize resource distribution across cloud, end devices, and services.
- Exploring vertical migration (between cloud and end devices) and horizontal migration (between end devices), assessing their impact and technical feasibility.

The project will evaluate how local resource utilization enhances service performance and availability.

## 9 CONCLUSION AND OUTLOOK

Sky Computing is a new Cloud Computing paradigm that abstracts end-users from cloud vendors via an intercloud broker, simplifying multi-cloud management by eliminating API complexities.

We analyzed this concept and proposed practical implementation strategies, identifying open questions that highlight the need for new frameworks, methods, and technologies to develop a marketplace from an end-user perspective.

Our proposed SKY CONTROL framework offers cost control and risk management solutions tailored for SMEs. By integrating dynamic analysis, real-time cost tracking, and risk assessment, SKY CONTROL enhances transparency and control over IT infrastructure. Its Sky Computing capabilities improve multi-cloud flexibility and efficiency, giving SMEs a competitive edge.

Additionally, we explored methods for attaching workloads to the Sky Computing ecosystem, focusing on geographic workload distribution across WANs. An offline-first strategy will be investigated to ensure local workload operation while enabling cloud relocation as needed.

This concept paper presents foundational ideas for Sky Computing implementation, acknowledging that many proposed questions remain open. As SKY CONTROL evolves, we aim to provide practical results in the near future.

## ACKNOWLEDGEMENTS

This work was funded by the Federal Ministry for Economic Affairs and Climate Action (Bundesministerium für Wirtschaft und Klimaschutz) under the

Central Innovation Programme for SMEs (Zentrales Innovationsprogramm Mittelstand).

We thank our partners at Systrade GmbH for their support, especially Andreas Schmidt, Thorsten Luft, Martin Schesny, Julian Hofmann, and Abo El Hage. Special thanks to Wei Yin Shing for proofreading and improving the paper's quality.

## REFERENCES

- Ahmad, N., Naveed, Q. N., and Hoda, N. (2018). Strategy and procedures for migration to the cloud computing. In *2018 IEEE 5th International Conference on Engineering Technologies and Applied Sciences (ICETAS)*, pages 1–5.
- Ardagna, D. (2015). Cloud and multi-cloud computing: Current challenges and future applications. In *2015 IEEE/ACM 7th International Workshop on Principles of Engineering Service-Oriented and Cloud Systems*, pages 1–2.
- Barker, A., Varghese, B., and Thai, L. (2015). Cloud services brokerage: A survey and research roadmap. In *2015 IEEE 8th International Conference on Cloud Computing*, pages 1029–1032.
- Baryannis, G., Garefalakis, P., Kritikos, K., Magoutis, K., Papaioannou, A., Plexousakis, D., and Zeginis, C. (2013). Lifecycle management of service-based applications on multi-clouds: a research roadmap. In *Proceedings of the 2013 International Workshop on Multi-Cloud Applications and Federated Clouds, MultiCloud '13*, page 13–20, New York, NY, USA. ACM.
- Chasins, S., Cheung, A., Crooks, N., Ghodsi, A., Goldberg, K., Gonzalez, J. E., Hellerstein, J. M., Jordan, M. I., Joseph, A. D., Mahoney, M. W., Parameswaran, A., Patterson, D., Popa, R. A., Sen, K., Shenker, S., Song, D., and Stoica, I. (2022). The sky above the clouds.
- Di Giulio, C., Sprabery, R., Kamhoua, C., Kwiat, K., Campbell, R. H., and Bashir, M. N. (2017). Cloud standards in comparison: Are new security frameworks improving cloud security? In *2017 IEEE 10th International Conference on Cloud Computing (CLOUD)*, pages 50–57.
- Fortes, J. A. B. (2010). Sky computing: When multiple clouds become one. In *Proceedings of the 2010 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing, CCGRID '10*, page 4, USA. IEEE Computer Society.
- Georgios, C., Evangelia, F., Christos, M., and Maria, N. (2021). Exploring cost-efficient bundling in a multi-cloud environment. *Simulation Modelling Practice and Theory*, 111:102338.
- Gundu, S. R., Panem, C. A., and Thimmapuram, A. (2020). Hybrid it and multi cloud an emerging trend and improved performance in cloud computing. *SN Computer Science*, 1(5).
- Hong, J., Dreiholz, T., Schenkel, J. A., and Hu, J. A. (2019). An Overview of Multi-cloud Computing. In

- Barolli, L., Takizawa, M., Xhafa, F., and Enokido, T., editors, *Web, Artificial Intelligence and Network Applications*, pages 1055–1068, Cham. Springer International Publishing.
- Ibrahimi, A. (2017). Cloud computing: Pricing model. *International Journal of Advanced Computer Science and Applications*, 8(6).
- Islam, M. N., Colomo-Palacios, R., and Chockalingam, S. (2021). Secure access service edge: A multivocal literature review. In *2021 21st International Conference on Computational Science and Its Applications (ICCSA)*, pages 188–194.
- Jamshidi, P., Pahl, C., and Mendonça, N. C. (2016). Pattern-based multi-cloud architecture migration. *Software: Practice and Experience*, 47(9):1159–1184.
- Kavitha, M. G. and Radha, D. (2022). *Quality, Security Issues, and Challenges in Multi-cloud Environment: A Comprehensive Review*, pages 269–285. Springer International Publishing, Cham.
- MacDonald, N., Orans, L., and Skorupa, J. (2019). The future of network security is in the cloud. *Gartner*.
- Mell, P., Grance, T., of Standards, N. I., and Division, T. U. C. S. (2011). *The NIST Definition of Cloud Computing*. NIST special publication. U.S. Department of Commerce, National Institute of Standards and Technology.
- Monteiro, A., Teixeira, C., and Pinto, J. S. (2014). Sky computing: Exploring the aggregated cloud resources — part ii. In *2014 9th Iberian Conference on Information Systems and Technologies (CISTI)*, pages 1–6.
- Mulder, J. (2020). *Multi-Cloud Architecture and Governance: Leverage Azure, AWS, GCP, and VMware vSphere to build effective multi-cloud solutions*. Packt Publishing.
- Petcu, D. (2013). Multi-cloud: expectations and current approaches. In *Proceedings of the 2013 international workshop on Multi-cloud applications and federated clouds*, ICPE'13. ACM.
- Siegel, J. (1999). An overview of corba 3. In Kutvonen, L., König, H., and Tienari, M., editors, *Distributed Applications and Interoperable Systems II*, pages 119–132, Boston, MA. Springer US.
- Stoica, I. and Shenker, S. (2021). From cloud computing to sky computing. In *Proceedings of the Workshop on Hot Topics in Operating Systems, HotOS '21*, page 26–32, New York, NY, USA. ACM.
- van der Walt, S. and Venter, H. (2022). Research gaps and opportunities for secure access service edge. In *International Conference on Cyber Warfare and Security*, volume 17, pages 609–619.
- van Steen, M. and Tanenbaum, A. (2017). *Distributed Systems*. 3rd edition. Self-published, open publication.
- Wei, X., Mohaimenur Rahman, A. B. M., and Wang, Y. (2021). Data placement strategies for data-intensive computing over edge clouds. In *2021 IEEE International Performance, Computing, and Communications Conference (IPCCC)*, pages 1–8.
- Yang, Z., Cui, Y., Li, B., Liu, Y., and Xu, Y. (2019). Software-defined wide area network (sd-wan): Architecture, advances and opportunities. In *2019 28th International Conference on Computer Communication and Networks (ICCCN)*, pages 1–9.
- Yang, Z., Wu, Z., Luo, M., Chiang, W.-L., Bhardwaj, R., Kwon, W., Zhuang, S., Luan, F. S., Mittal, G., Shenker, S., and Stoica, I. (2023). SkyPilot: An intercloud broker for sky computing. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*, pages 437–455, Boston, MA. USENIX Association.