

CTO SUMMIT REPORT EU 2022

RESILIENCY IN MULTI-CLOUD

August 2022

Foreword by Amr Abdelhalem, *Fidelity*;
Arun Gupta, *Intel*; Ricardo Torres, *Boeing*;
Pratik Wadher, *Intuit*

Authored by Dr. Udo Seidel,
Speakeasy Strategies

With Hilary Carter, *Linux Foundation Research* and
Priyanka Sharma, *Cloud Native Computing Foundation*

In partnership with





Contents

Foreword	3
Executive Summary	4
Key Findings	5
Discussion Topics	7
Managing the Cloud Native Ecosystem	8
TECHNOLOGIES	8
PEOPLE AND PROCESSES	9
GAPS TO BE CLOSED	10
Multi-Cloud Availability	10
TECHNOLOGIES	10
PEOPLE AND PROCESSES	11
GAPS TO BE CLOSED	12
The Collaboration Work Continues	13
Appendices	14
Appendix A: Summit Process	14
Appendix B: Sub-topics	14
MULTI-CLOUD MANAGEMENT	14
MANAGING THE CLOUD NATIVE ECOSYSTEM IN YOUR ORGANIZATION	14
WHAT TENETS OF MULTI-CLOUD ARE ESSENTIAL FOR BUSINESS?	15
MULTI-CLOUD AVAILABILITY	15
Appendix C: Discussion Framework	15
Appendix D: Glossary	15
Acknowledgments	16
About the Author	17
Disclaimer	17



Foreword

The goal of the Linux Foundation, and Cloud Native Computing Foundation (CNCF) is to share the best software innovations and practices to move technologies, and thereby companies, economies, and societies, forward.

The rise of cloud computing, and cloud native computing, comprises one of the largest shifts that the computing industry has seen in decades.

There are several benefits of cloud computing. Companies get the computing power that they need when they need it and can scale applications and workloads faster. Additionally, the cloud democratizes resources and makes them more accessible to more companies and entities of all sizes.

“Resiliency in a cloud-defined world means being able to maintain operations under all circumstances and adjust accordingly, especially in the face of unexpected events.”

The idea years ago that companies would move all applications to one cloud provider has not panned out, with good reason. Companies cannot afford to have their operations disrupted by cloud or network outages, which often occur.

Resiliency in a cloud-defined world means being able to maintain operations under all circumstances and adjust accordingly, especially in the face of unexpected events.

Multi-Cloud is the practice of using services (e.g., computing, databases, and networking) from more than one cloud service provider (CSP) at the same time. Multi-Cloud can include public CSPs, such as Amazon Web Services (AWS) and Microsoft Azure, private clouds, or a combination of the two.

Despite the advantages of Multi-Cloud, companies

struggle to succeed with Multi-Cloud and achieve resilience and optimal operations under all conditions.

This emerging arena inspired the Linux Foundation and the CNCF to host the first-ever Chief Technology Officer (CTO) Summit about Multi-Cloud resiliency and how to achieve it.

The May 18, 2022 summit included 21 participants from six business verticals. The participants represented diverse industry sectors and functions, including aeronautics, automotive, semiconductor, insurance, telecommunication, healthcare, business services, technology, banking, fintech and finance, e-commerce, social media, and audio streaming. Generally speaking, the participants held leading information technology roles in their companies, ranging from the chief Kubernetes architects to the CTO level.

They all experience Multi-Cloud differently. Some companies deal with it within their organization, while others consult and support their customers on the subject. It is often a combination of both.

The following criteria were critical:

- Experience in Multi-Cloud and associated technologies
- Decision-making and budget accountability
- Cloud native knowledge

In addition, all participants are consumers of cloud services rather than providers, which the CNCF and the technology industry broadly define as the **end-user community**.

The following report captures the significant findings of the Summit participants, the questions raised, and the concerns that must be addressed.

Amr Abdelhalem, Fidelity
Arun Gupta, Intel
Ricardo Torres, Boeing
Pratik Wadher, Intuit
CTO Summit Committee



Executive Summary

Today, Multi-Cloud is a reality for many companies and organizations and will become more so as cloud native architectures mature. The rise of Multi-Cloud is driven by several factors, including mergers and acquisitions that bring different clouds together into one entity and the desire to avoid widespread outages by spreading the workload and risk across multiple clouds.

Currently, more than 90% enterprises have a Multi-Cloud **strategy**, 82% have a hybrid cloud strategy, and more of the same is to come. The COVID-19 pandemic has caused organizations to accelerate their migration to the cloud.

Multi-Cloud resiliency means avoiding or mitigating an adverse event's impact and being ready for unexpected outcomes. Achieving Multi-Cloud resiliency requires different approaches than those that are used within a single cloud, no-cloud environments, or even hybrid cloud environments. Finding a path to federate Multi-Cloud architectures is a growing concern for many organizations.

To paraphrase the Summit participants: Multi-Cloud is an inevitable and significant challenge.

In general, the Summit participants agreed that there is a clear need for a Multi-Cloud strategy to govern a rollout. There is also a need for proven best practices to better

deal with multiple cloud providers. Best practices should include a framework to exchange ideas and experiences.

Perhaps most importantly, the Summit participants largely agreed that while there is the potential for better team performance, service availability, and lower operating costs with a Multi-Cloud environment, it has not materialized for many companies. This is because they have varying levels of knowledge and experience in managing multiple cloud environments, building something that is cloud native, handling the exponentially increasing risks that are inherent in Multi-Cloud environments, connecting clouds, and getting data out of a cloud once it is in without a hefty price tag.

The Summit participants pointed out that challenges grow exponentially when organizations add multiple cloud providers. To achieve Multi-Cloud resiliency, entities need best practices when leveraging people, processes, and technologies. The Summit participants geared their discussion toward offering insights and best practices to help companies and organizations to achieve Multi-Cloud Resiliency.

We thank all participants for sharing their experiences, knowledge, and insight.



Key Findings

“No organization is alone in facing challenges to become a Multi-Cloud operator.”

Multi-Cloud is here to stay.

- Multi-Cloud challenges all organizations, regardless of the industry sector, business model, or technological maturity.
- Multi-cloud architecture requires new thinking models, workflows, and cloud native projects.
- Implementing a resilient Multi-Cloud solution is where we will see the most diversity in adoption and maintenance.
- The basic questions and challenges are the same, regardless of your organization's size or business vertical.
- A clear architecture with blueprints that reflect your organization and use cases encourage adoption and compliance.
- The most prevalent concern with Multi-Cloud availability is network federation across an organization.
- It is worth checking to see how you can use it to increase or improve your operations and achieve high availability, which is neither easy nor obvious.
- No organization is alone in facing challenges to become a Multi-Cloud operator.

People and processes are crucial for success.

- Multi-Cloud is not a technology-only topic. People and processes should not be an afterthought.
- Success requires a culture of going the extra mile.
- Organizational setup and people management are essential, regardless of the technology used.
- Managing access can be difficult, and not all workflows must migrate entirely into the cloud.
- Cloud environments are new and unique, so processes and governance will have to change and evolve significantly.
- Education and training are necessary to increase existing talent's skill levels or bring in new talent to a cloud native environment.
- To help organizations master new cloud computing challenges, sharing information via forums like the Summit and sharing best practices or other valuable information is a great start.
- Community is part of the solution. Bringing developers together in the community and learning from the open source community is critical.

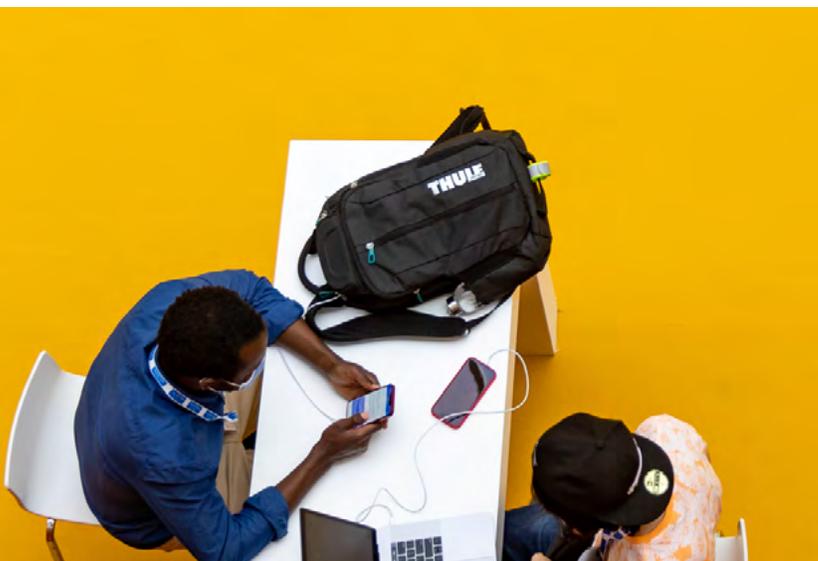
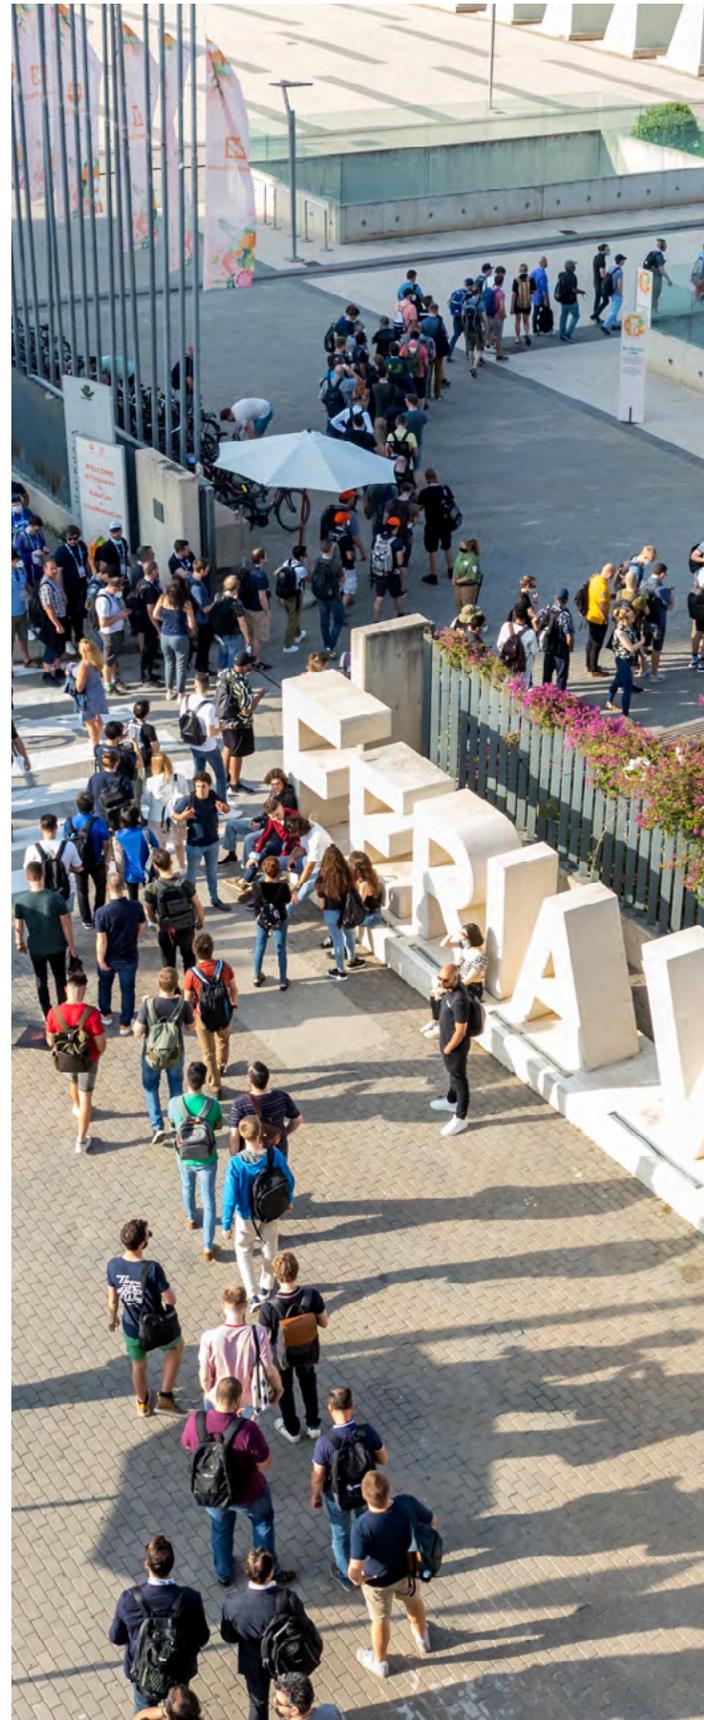
“Community is part of the solution. Bringing developers together in the community and learning from the open source community is critical.”



There are gaps to be closed.

- Multi-Cloud exponentially increases known challenges and adds new ones.
- There are still some open topics without a solid solution or set of approaches.
- Most of the gaps are related to seamlessly connecting clouds from different providers.
- Entities should think beyond CSPs with Multi-Cloud and include other global infrastructure services, such as hostname resolution or content management.
- Use cases are better than blueprints. Things like reference architectures are too theoretical and difficult to adapt to different industries. Use cases are easier to consume and more practical. By design, solutions must be topic-based and tailored to the industry sector.
- New user groups (UGs) are necessary. Special interest groups (SIGs) in the CNCF currently represent a deep focus on a technology topic or workflow. Many companies are not staffed or do not have the time to attend SIG meetings regularly. More UGs that are focused on solutions within a specific business vertical will be beneficial alongside the SIGs. Companies can engage with SIGs as needs arise and UGs for the long term.

“Use cases are better than blueprints.”





Discussion Topics

The Summit participants chose to focus on the following two primary topics:

- 1. Managing the cloud native ecosystem in one's organization.** This is critical because as cloud native and microservices environments grow, entities face challenges in managing them. Tools and services may need to change, along with developer access to such things as Kubernetes clusters, namespaces, or applications.
- 2. Multi-Cloud availability.** This topic covered the services that are necessary to achieve a resilient Multi-Cloud platform that enables regular iteration and improvement.

The Summit participants were broken into three smaller sub-groups to facilitate deeper discussion. They were asked to concentrate on how leveraging people, processes, and technology in both arenas should—and could—change to increase Multi-Cloud resiliency and to identify what the priorities should be.

BREAKOUT SESSION TOPIC 1

Managing the Cloud Native Ecosystem

The introductory statement was as follows:

Managing Kubernetes clusters, compute, and applications across multiple clouds requires a particular set of tools and comprehensive workflows.

This discussion aims to understand which kinds of workflows teams should implement to allow for massive infrastructure scaling without massively scaling team sizes or budgets.

BEST PRACTICES

Technologies

- **Strong opinions, loosely held.** To increase resiliency, accept that the solutions that work today may not be the solutions that work tomorrow, and solutions at the scale that you have today may not be the solutions when you scale tomorrow.
- **Reduce complexity.** Automation can significantly help, i.e., integrating governance checks or security tests in the delivery pipelines. With Multi-Cloud, different approaches and ways of working can cause friction among teams, and the presence of inconsistent processes hinders teams from doing their best.
- **Eliminate the “not invented here” mindset.** There is no right or wrong approach with self-hosted/made vs. managed. It depends on your particular setup, organization, and business. Decisions must be transparent and based on the available skill sets and current challenges, including project time frames. The right choice today might not be the best tomorrow. The solution to the business problem should be the driver and not a particular technical interest.
- **Embrace regulatory reality.** We all exist in regulated industries, as the General Data Protection Regulation (GDPR) impacts everyone. In the future, regulation may require a Multi-Cloud approach to lessen the impacts of outages and breaches.
- **Select where you are going to standardize.** To do this, you must define your goal of what you want from a cloud and your standardization to achieve that goal.

“The products you have ...
might need to change as you
migrate to the cloud.”



People and Processes

- **People make the difference.** Technology will support master challenges, but people will implement solutions and drive the necessary change. Things like end-to-end responsibility are done on a human level. Technology cannot “go the extra mile,” but people can. Company culture—in general, and in particular—plays a key role. Take empowerment and enablement seriously, i.e. leave technical decisions to the technicians, and trust the judgment of your staff with first-hand experience and knowledge. Recognize that some legacy teams will hinder the progress of people doing the Kubernetes or cloud-based work because they are worried about losing their jobs and fear change.
- **Craft an open source strategy and engage with projects.** Participate in open source projects to reflect your use cases and needs. This can also influence the roadmap of vendors. Be mindful when combining proprietary technology and free and open source software in a solution or software stack.
- **Retain and upskill talent.** This entails attracting new team members¹ and keeping the existing talent engaged and empowered. Multi-Cloud adds the additional dimension of needing even more skills. Furthermore, it includes upskilling and training and detecting unknown talents within your organization and developing them.
- **Use standards and governance** as frameworks without limiting innovation. A means of control is needed regarding where the budget is spent. That is even more true in a Multi-Cloud setup. However, these frameworks should be seen as guidelines rather than rigid cages. There should be a place for experimentation and trying out new things.

“You must scale up your people...Give them what they need.”

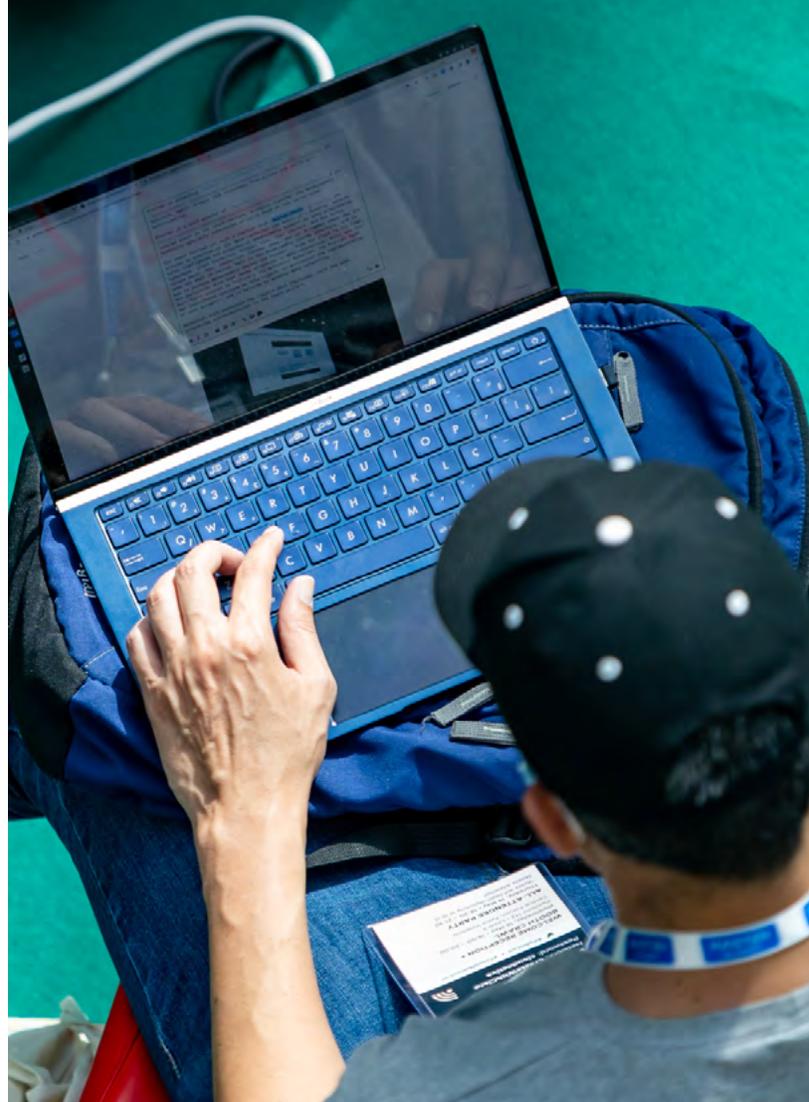
¹ <https://training.linuxfoundation.org/resources/2022-open-source-jobs-report>



Gaps to be Closed

- **Who owns the intellectual property** that was created? An important part of the cloud native ecosystem is contributing to projects. Often, the fastest path is to pay somebody to do that. The code needs to be maintained, and bugs must be fixed. This becomes a problem if that upstream code is not picked up by the commercial offerings of vendors.
- **What about non-cloud native workloads?** Although container and container orchestration have solved many problems, they have not yet addressed the entire picture. One example is real-time systems that are embedded in firmware.

“Not everything is a suitable [container] candidate.”



BREAKOUT SESSION TOPIC 2

Multi-Cloud Availability

The introductory statement was as follows:

Many services like domain name systems (DNSs), content distribution networks (CDNs), and artifact storage can be instantiated in public or private clouds.

This session covered which services are requirements for achieving a resilient Multi-Cloud platform that enables regular iteration and improvement.

BEST PRACTICES

Technologies

- **Consider architecture and design as key.** Multi-Cloud comes with new challenges and tasks. An example is the cross-CSP transfer of information, such as DNS. The dependency graphs are becoming very complicated to read and manage. The foundation of a solution starts at the drawing board. Although seasoned experience with a single CSP is handy, it is not sufficient. In-depth knowledge of each involved cloud is mandatory.

“Kubernetes is part of the solution.”

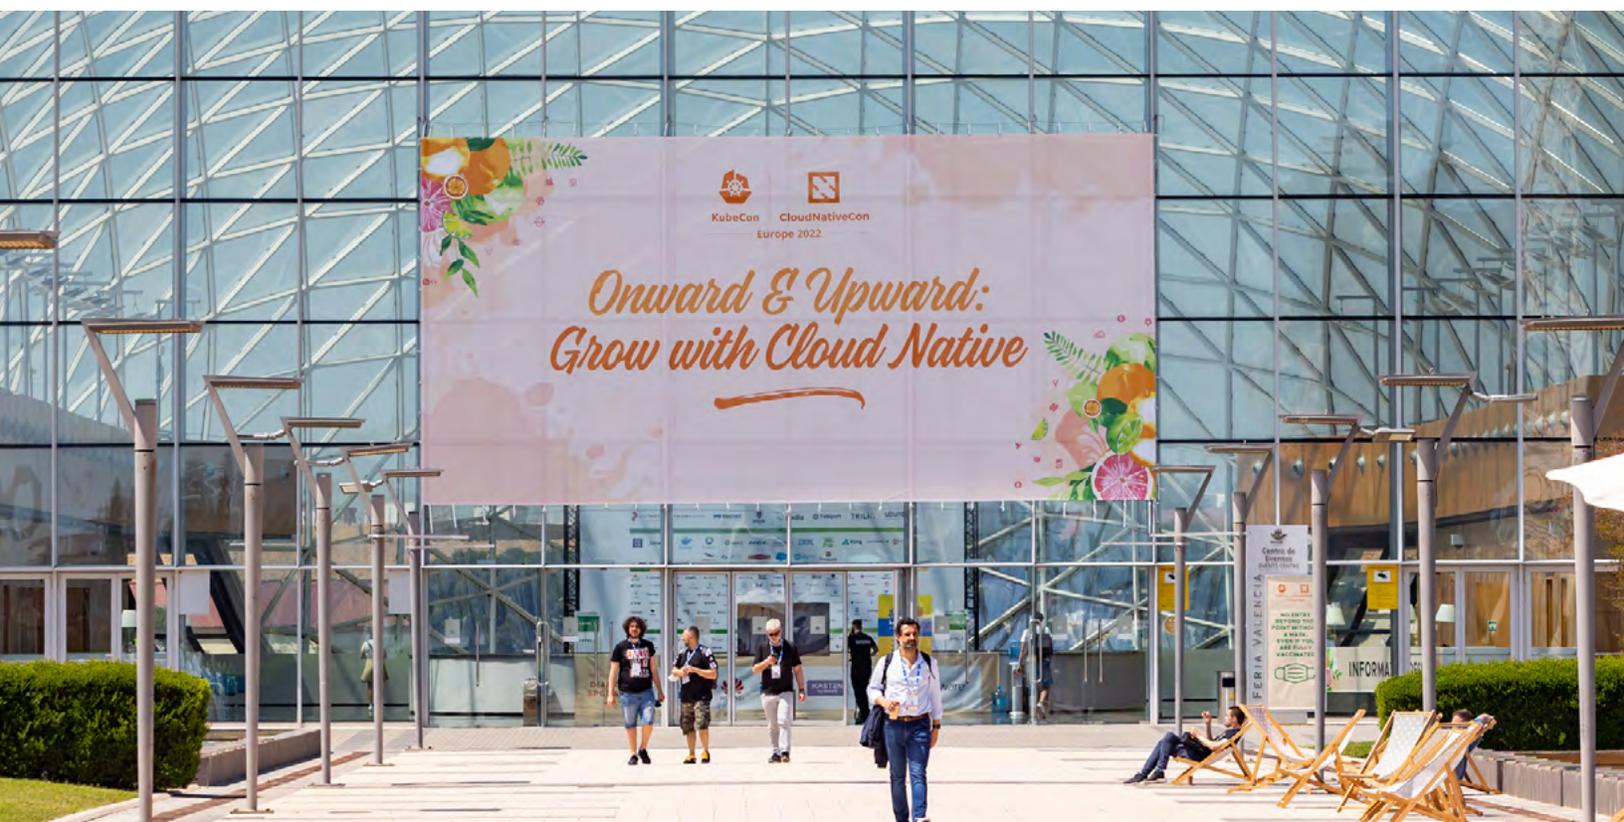
- **Recognize that data is both easy and complicated.** Using Multi-Cloud to duplicate your data on different clouds is a positive and so-called easy win. The current pricing model supports that. However, you should carefully review the ingress and egress costs for your use case. The data location policy is more complicated. Due to a missing cross-CSP abstraction layer, data may end up in the wrong place if the specifics of a contributing cloud are not fully understood.
- **Manage vendor lock-in to your advantage.** To minimize egress costs, consider grouping the same types of business into different clouds. For example, direct all mobile connectivity into one cloud and all webpage and front end into another cloud. With little interaction between the two, the lock-in becomes less of an issue, and you can achieve the benefits of leveraging the price advantages.

“Data policies are important.”

People and Processes

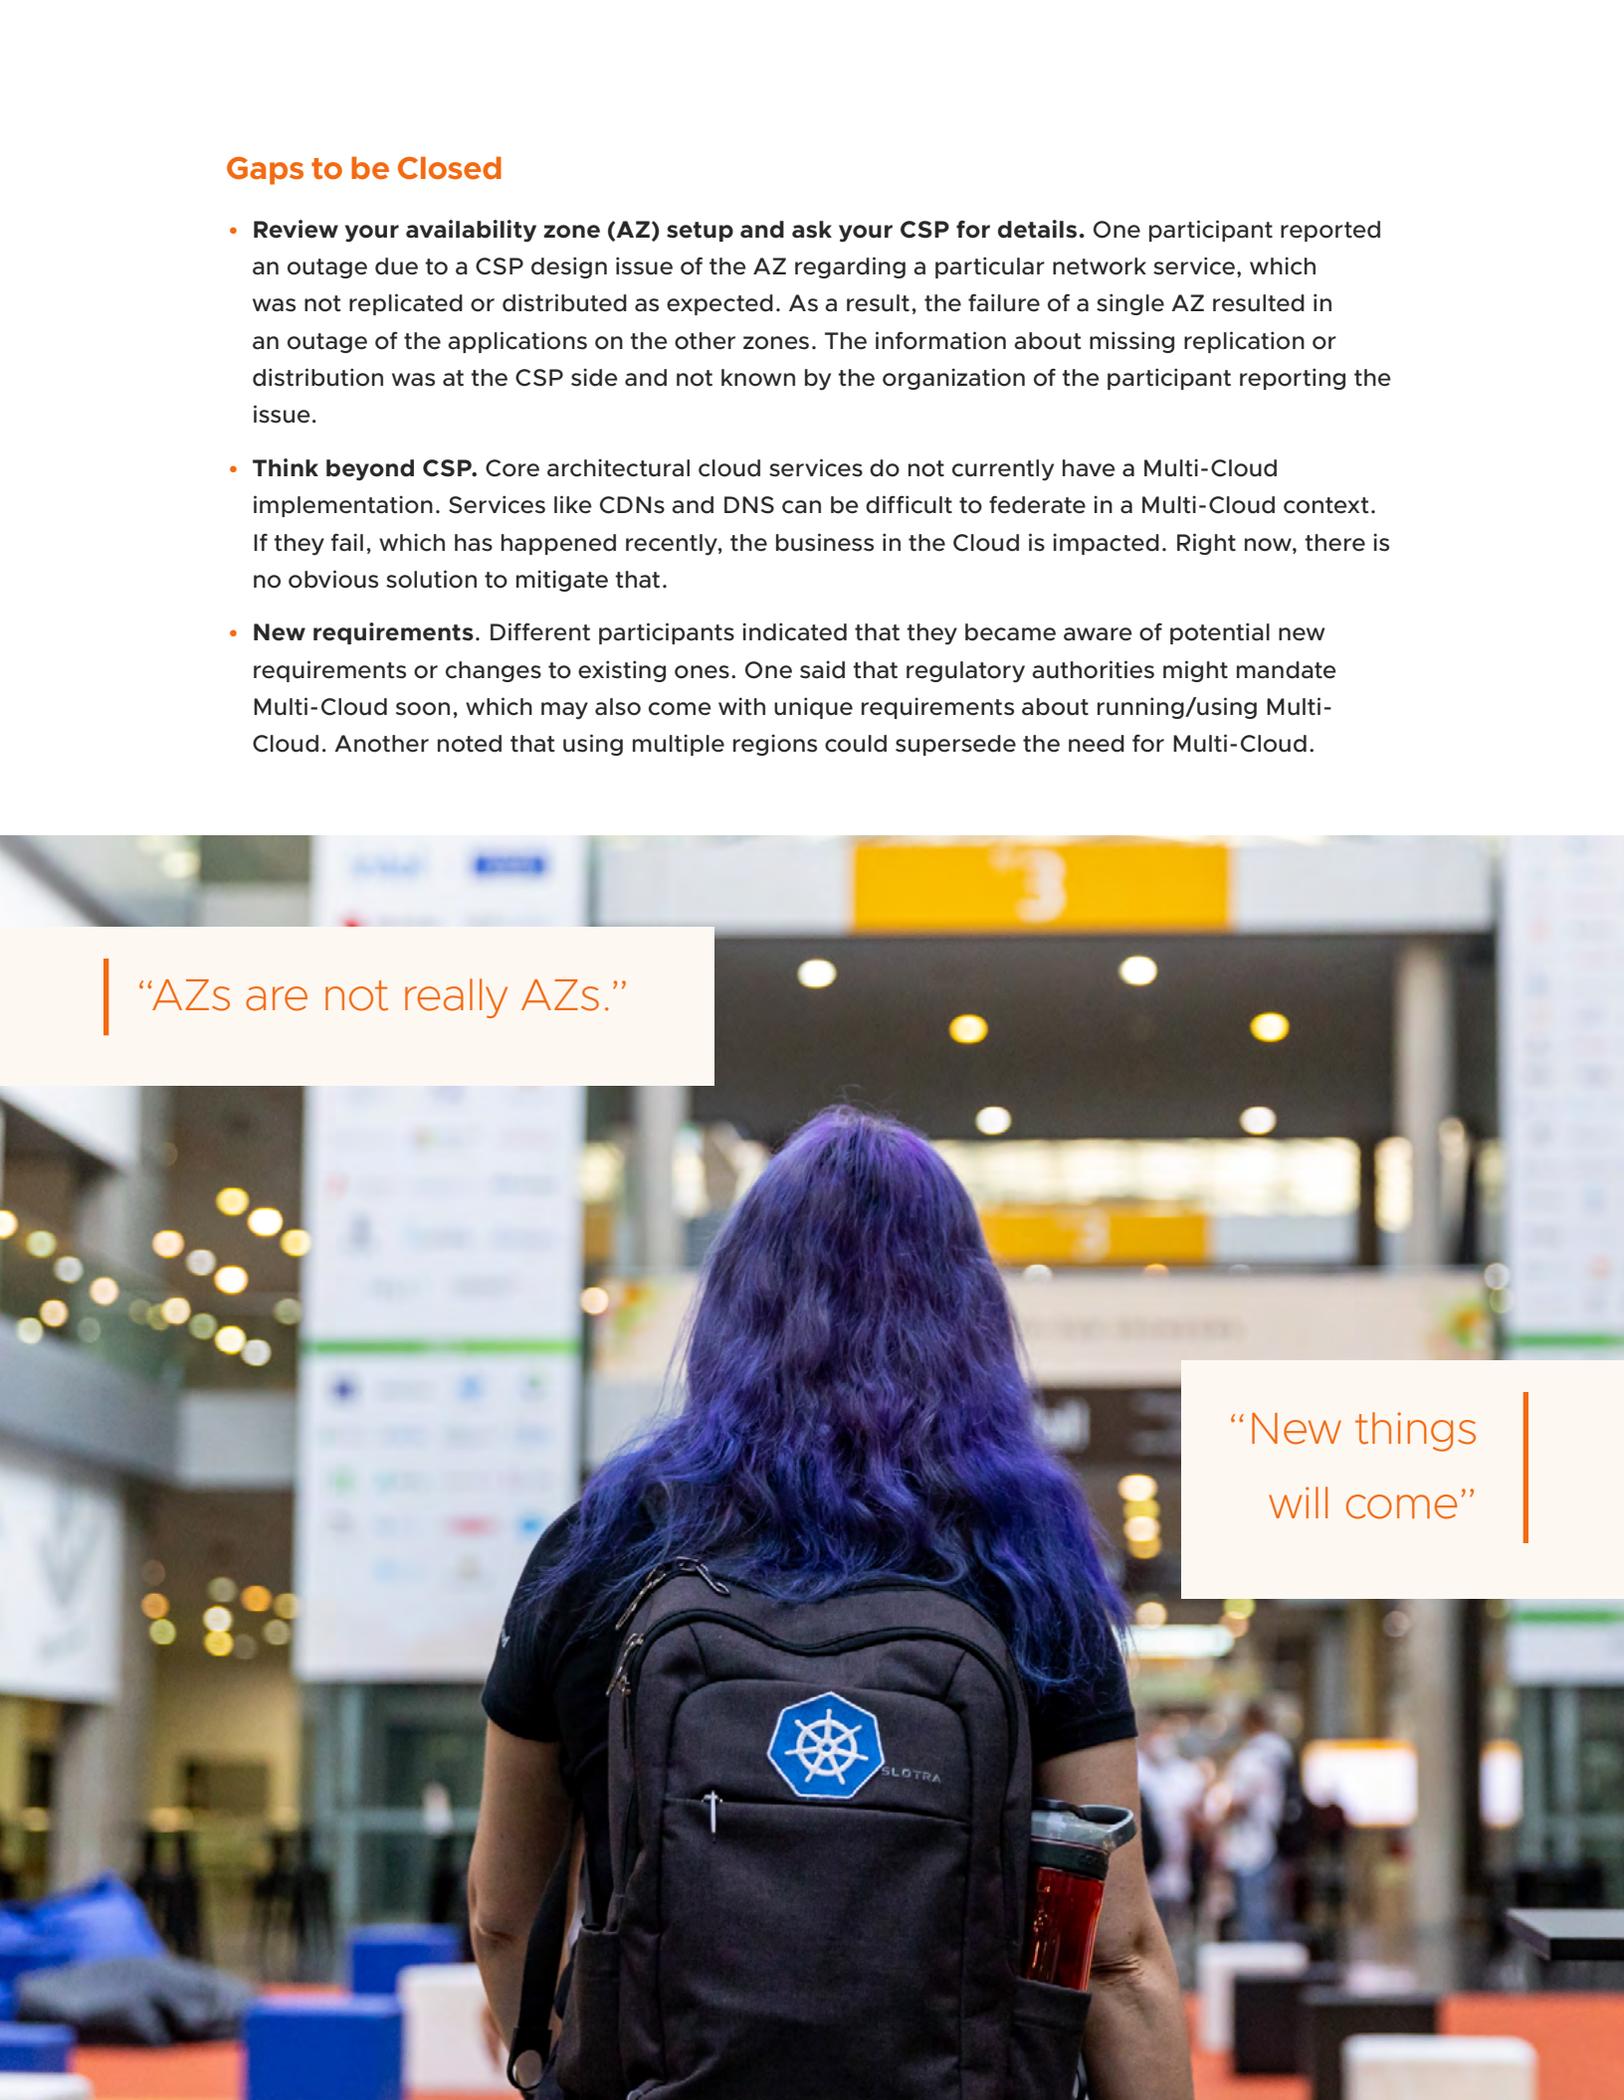
- **Governance is crucial.** Which permissions and access do you grant to whom? How do you make sure that the right locations are selected? What about data policy? How do you effectively deal with the different life cycles? The importance of governance with Multi-Cloud is much higher. Much more forethought is required to cover the different CSPs and their specifics.
- **Take a new approach to scale your teams.** Above all, any new system has to reflect your organization. Hence, you must deal with specialized teams in a single CSP while covering multiple clouds.

“Understand the new rules.”



Gaps to be Closed

- **Review your availability zone (AZ) setup and ask your CSP for details.** One participant reported an outage due to a CSP design issue of the AZ regarding a particular network service, which was not replicated or distributed as expected. As a result, the failure of a single AZ resulted in an outage of the applications on the other zones. The information about missing replication or distribution was at the CSP side and not known by the organization of the participant reporting the issue.
- **Think beyond CSP.** Core architectural cloud services do not currently have a Multi-Cloud implementation. Services like CDNs and DNS can be difficult to federate in a Multi-Cloud context. If they fail, which has happened recently, the business in the Cloud is impacted. Right now, there is no obvious solution to mitigate that.
- **New requirements.** Different participants indicated that they became aware of potential new requirements or changes to existing ones. One said that regulatory authorities might mandate Multi-Cloud soon, which may also come with unique requirements about running/using Multi-Cloud. Another noted that using multiple regions could supersede the need for Multi-Cloud.



“AZs are not really AZs.”

“New things
will come”

The Collaboration Work Continues

As the Summit participants noted, Multi-Cloud resiliency is only attainable with changes to the legacy processes, technologies, mindsets, and skillsets. Although it will be a long journey to true resiliency, early results show that the effort will be worth it in terms of an enhanced time to market, faster development cycles, and enhanced capabilities.

Multi-Cloud and its different aspects are still **mostly undiscovered territory**. Over time, we will be able to go deeper in our multi-cloud resiliency skills as we develop our expertise in a collaborative, global environment. Participants left the CTO Summit inspired with fresh perspectives and ideas to bring back to their organizations.

The open discussions and experience exchange, including the logistics around the event, were very much welcomed by the participants. The following requests for continued collaboration and idea exchange were raised:

- **Continue the discussion after the event**
- **Have a framework to easily contact peers or knowledgeable persons on Multi-Cloud topics**
- **Repeat the CTO summit with a new topic at the next conference**

The Linux Foundation and the CNCF are built on the notion that the best ideas come from anywhere and that collaboration is key to achieving the best results. We look forward to continuing this work with others as we work toward Multi-Cloud resiliency to improve efficiency, performance, and innovation.



Appendices

Appendix A: Summit Process

The participants received a brief introduction to the Multi-Cloud topic in general and several core elements in particular. This introduction synchronized the expectations of the Summit and served to broaden the understanding of the subject. Additionally, the participants were allowed to vote on which particular aspects of Multi-Cloud to discuss in breakout sessions.

Three working groups deconstructed two critical topics. After the opening remarks, introductions, and logistics, the participants split up into three designated working groups with dedicated rooms, each separate from the opening plenary discussion area. Given the diversity of the participants, the resulting groups served as a measured representation of the different aspects and viewpoints of Multi-Cloud challenges and experiences across cloud computing.

Each working group explored two topics and evaluated three key areas. Following a vote by a show of hands, the following sub-topics were selected and discussed separately:

- How to best manage the cloud native ecosystem in any organization
- Multi-Cloud availability

To facilitate the discussion, the groups followed a common framework to focus their discussion using the following three key areas:

- Expand on or define the problem statement of the selected sub-topic.
- Solution design
 - » People
 - » Processes
 - » Technology
 - » Priorities

The working group facilitators sparked discussion, managed the flow of ideas and solution design, and recorded the breakout group's key findings. There was one dedicated session per topic for each working group. Afterward, each group reported back on their results to the Summit participants, and then an open discussion ensued.

Appendix B: Sub-topics

Multi-Cloud management

Multi-Cloud is inevitable. It creeps in through best-of-breed services, company acquisitions, and just plain sprawl. In such a world, managing Kubernetes clusters, computing, and applications across multiple clouds require a unique set of tools and comprehensive workflows. Using strategies like GitOps and declarative infrastructure creates consistent environments, which are essential for operating systems at scale. This discussion aims to understand which kinds of workflow teams should implement to allow for massive infrastructure scaling without massively scaling team sizes or budgets.

Managing the cloud native ecosystem in your organization

Do you use a CNCF project by itself? There are unique challenges when running and configuring a project directly rather than through a managed service. The cloud native ecosystem has a vibrant community of particular interest, working, and technical action groups that create boundless updates across the ecosystem that can be difficult to keep up with. During this session, we will cover the practices that high-performance teams implement to stay up to date on project updates without wasting their time.



What tenets of Multi-Cloud are essential for business?

Although cloud native technologies can help to support workload, workflow, data, and traffic portability, which tenets of Multi-Cloud are the most beneficial? Do you care about workflows that allow a consistent experience building applications, computing, networks, and storage? Maybe you'd like to write an application consistently and have the platform be dynamically configurable to meet your needs. In this session, we will discover which approaches are most beneficial to teams that operate at a Multi-Cloud scale, including examples of success that have been shared by our program committee from Boeing, Fidelity, Intel, and Intuit.

Multi-Cloud availability

Many services, such as DNS, CDNs, and artifact storage, can be instantiated in public or private clouds. Despite many of these services being offered in major public clouds, relying on managed services does not always make sense when resiliency is the end goal. Which services make the most sense to run as managed services or federate between your various cloud offerings? This session will cover which services are requirements for achieving a resilient Multi-Cloud platform that promotes regular iteration and improvement.

Appendix C: Discussion Framework

To encourage a standard structure and flow for the breakout discussions and guide participants to share intelligence on how existing tactics are implemented either currently or in the future, the participants were asked to address the following points as they relate to each subtopic. This would enable the different discussion groups to be compared.

1. Expand on/define the sub-topic problem statement.
2. Solution design
 - a. How can the following elements be leveraged as they relate to overcoming challenges?
 - b. People
 - i. Team/community leaders, community members, employees, and the qualities of each

- ii. Hard/soft skills
 - iii. Budgeting for people
 - b. Processes
 - i. Operational considerations—what workflows and aspects of workflows need to be implemented to overcome the current challenges?
 - ii. What is the budgeting for workflows?
 - c. Technology/Tooling
 - i. Which tech stacks/hardware/standards can be leveraged/implemented?
 - ii. Other tools?
 - iii. Financial/budget considerations?
3. Priorities
 - a. Where should teams focus their efforts? Choose three opportunities that should be implemented first.

Appendix D: Glossary

- **AWS** - Amazon Web Services
- **AZ** - Availability Zone
- **CDN** - Content Delivery Network
- **CNCF** - Cloud Native Computing Foundation
- **CSP** - Cloud Service Provider
- **CTO** - Chief Technology Officer
- **DNS** - Domain Name System
- **GDPR** - General Data Protection Regulation
- **SIG** - Special Interest Group
- **UG** - User Group



Acknowledgments

Thanks to the CNCF for creating an opportunity for leading end-user experts to convene at KubeCon+CloudNativeCon Europe to collectively discuss the struggles, opportunities, and new architectures for crafting a resilient Multi-Cloud strategy.

Many thanks to the content committee members Pratik Wadher from Intuit, Ricardo Torres from Boeing, and Amr Abdelhalem from Fidelity, who steered the development of the discussion topics for this and future summits. Additional thanks go to Arun Gupta from Intel, who as the governing board chair, was involved in the preparation and execution of the event. Thanks also go to the facilitators Pratik Wadher and Henrik Blixt from Intuit and Ricardo Torres from Boeing, who managed to inspire a discussion on trending cloud native topics while giving participants plenty of room to explore the issues. In addition, this distinguished group of people was instrumental in reviewing this report and raising it to the needed and desired quality.

Special thanks to Orca Security and Tigera for generously sponsoring the reception and dinner that followed the May 18 Summit.

Thanks to the Linux Foundation and CNCF team members Taylor Dolezal, Paige O'Connor, and Kristi Tan for their contributions to the summit and to the events team, who took care of the logistics before, during, and after the event, in particular Vanessa Heric and Wendi West. Thanks also go to Hilary Carter, Linux Foundation Vice President of Research, for her work as the internal and external facilitator—not only related to the event but also for the management and production of this report.

Finally, thanks go to Priyanka Sharma (executive director of the CNCF) and the more than 20 participants of the CTO summit. Where Priyanka's leadership and vision had set the tone and the spirit for the event, the participant's discussions and insights helped craft a rich foundation for this report.

Dr. Udo Seidel
Report author



About the Author



Dr. Udo Seidel would have been a teacher for mathematics and physics if he had not been infected by the open source virus in 1996. After his PhD he worked as Linux/Unix instructor, system administrator, senior solution engineer, architect, digital evangelist, and account CTO. He regularly speaks at conferences and publishes articles in computer magazines.

Disclaimer

This report is provided “as is.” The Linux Foundation and its authors, contributors, and sponsors expressly disclaim any warranties (express, implied, or otherwise), including implied warranties of merchantability, noninfringement, fitness for a particular purpose, or title, related to this report. In no event will the Linux Foundation and its authors, contributors, and sponsors be liable to any other party for lost profits or any form of indirect, special, incidental, or consequential damages of any character from any causes of action of any kind with respect to this report, whether based on breach of contract, tort (including negligence), or otherwise, and whether they have been advised of the possibility of such damage. Sponsorship of the creation of this report does not constitute an endorsement of its findings by any of its sponsors.

If you have any questions or comments about this report, you can get in touch with us at

cto-summit@cncf.io





Founded in 2021, Linux Foundation Research explores the growing scale of open source collaboration, providing insight into emerging technology trends, best practices, and the global impact of open source projects. Through leveraging project databases and networks and a commitment to best practices in quantitative and qualitative methodologies, Linux Foundation Research is creating the go-to library for open source insights for the benefit of organizations across the world.



The Cloud Native Computing Foundation (CNCf) hosts critical components of the global technology infrastructure. It brings together the world's top developers, end users, and vendors and runs the largest open source developer conferences. CNCf is the open source, vendor-neutral hub of cloud native computing and hosts projects like Kubernetes and Prometheus to make cloud native universal and sustainable. The CNCf is part of the Linux Foundation.



August 2022



Copyright © 2022 The Linux Foundation

This report is licensed under the **Creative Commons Attribution-NoDerivatives 4.0 International Public License**.

To reference this work, please cite as follows: Udo Seidel, “CTO Summit Report EU 2022: Achieving Resiliency in Multi-Cloud,” foreword by Amr Abdelhalem, Arun Gupta, Ricardo Torres, Pratik Wadher, The Linux Foundation, August, 2022.