Deep Science Publishing, 2026
https://doi.org/10.70593/978-93-7185-496-2_9

DeepScience
Open Access Books

Chapter 9: Governance, Security, and Compliance in Multi-Cloud Environments

9.1. Introduction

Open access to cloud computing geographically decentralizes cloud governance, enabling organizations to use services from multiple providers, often located in different jurisdictions. Multi-cloud computing creates distinct management challenges in cloud governance, security, and compliance for organizations and their service providers and regulators. Multi-cloud governance is a continuously evolving set of formal and informal policies managing cross-cloud interactions.

Multi-cloud governance is distinct from cloud governance. Cloud governance is about managing the integrity of a single cloud, the accountability of that cloud's service providers, and the relationship between a cloud consumer and the cloud provider. Multi-cloud governance covers the broader territory of interactions among cloud consumers and providers that involve multiple clouds. Multi-cloud governance helps organizations retaining the benefits of multiple clouds without losing control over security and compliance. The literature has produced responses to many crucial issues in multi-cloud governance but significant parts of the governance puzzle remain disconnected and inconclusive.

9.1.1. Background and Significance

Governance in a multi-cloud context encompasses the policies, roles, responsibilities, and enforcement mechanisms for ensuring that an enterprise's cloud computing strategy is aligned with and supports its business goals. The key stakeholders include cloud consumers, cloud providers, compliance regulators, internal and external auditors, and governance boards. The core principles of governance across any number of clouds are accountability, transparency, interoperability, and enterprise control. Multi-cloud governance can be centralized, with a single authority governing resource consumption

across multiple clouds, or federated, with each cloud provider governing its own resources.

Security in a multi-cloud context encompasses the processes, policies, and architecture for securing enterprise resources spanning multiple cloud providers. A cloud consumer's resources face the same potential threats as any other Internet-accessible resource, but additional threats arise from consumers' use of resources hosted by multiple vendors. Security challenges affecting data in multi-cloud environments include security of data at rest, data in transit, and data in use; enforcement of the principle of least privilege; and securing third-party software and services.

Fig 9.1: Governance, Security, and Compliance in Multi-Cloud Environments

9.1.2. Research design

Multi-cloud governance and its supporting set of strategies to ensure the confidentiality, integrity, and availability of service access when deployed across multiple clouds are particularly important in addressing regulatory and compliance requirements. Governance involves the formulation of strategies, processes, and mechanisms to establish decision rights and accountability frameworks that encourage desirable behavior in the use of resources. Security frameworks and controls help elucidate the security measures, controls, and mechanisms that help secure data and applications across multiple cloud providers from the diverse threats faced in such environments. Compliance management provides a mechanism by which organizations can ascertain

the compliance of cloud services with the requirements of relevant laws and external compliance mandates, such as the GDPR, HIPAA, or PCI-DSS. Risk management involves the identification of business risks associated with the use of multiple cloud providers and the mitigation of these risks through business continuity, disaster recovery, and vendor risk management.

The research approach involves a comprehensive literature review supported by the specification of key definitions, the identification of key stakeholders, and the articulation of core principles of multi-cloud governance. Governance models for the management of policies across multiple clouds are examined, along with a set of security controls that address specific multi-cloud threats. Compliance challenges in a multi-cloud environment, including jurisdictional issues, maintenance of evidence, and fulfillment of third-party auditing requirements, are discussed. Together, these components of the research contribute to a holistic understanding of multi-cloud governance, security, and compliance. Finally, with an eye toward the future, these contributions highlight potential gaps for further research and investigation.

9.2. Conceptual Foundations of Multi-Cloud Governance

Multi-cloud governance frameworks orchestrate the establishment, compliance with, management of, and ongoing adherence to policies, standards, and technologies within multi-cloud environments. Central stakeholders include cloud consumers, providers, regulators, auditors, and governance boards. Four core principles underlie effective governance in multi-cloud settings: accountability, transparency, interoperability, and control. Three architectural dimensions balance abstraction and distribution of governance responsibilities: centralized versus federated models, policy abstraction and translation across cloud services, and identity, access, and entitlement management strategies.

Multi-cloud governance frameworks orchestrate the establishment, compliance with, management of, and ongoing adherence to policies, standards, and technologies within multi-cloud environments. Central stakeholders include cloud consumers, providers, regulators, auditors, and governance boards. Four core principles underlie effective governance in multi-cloud settings: accountability, transparency, interoperability, and control. Three architectural dimensions balance abstraction and distribution of governance responsibilities: centralized versus federated models, policy abstraction and translation across cloud services, and identity, access, and entitlement management strategies.

9.2.1. Definitions and Scope

Gaps and ambiguities in prior research obstruct a comprehensive understanding of multi-cloud governance, the establishment of which is critical for assuring stakeholders that multi-cloud setups can be controlled in a consistent, reliable manner. Multi-cloud governance refers to all functions, activities, and decisions that lend coherence, shape, and substance to external cloud consumer access across different clouds. It concerns how external cloud consumer and service provider decision rights, responsibilities, and authorizations are allocated across the different cloud services being consumed. Just-in-time, just-for-you, on-demand control in all the different cloud services a cloud consumer is vastly unlikely to achieve or to have even considered.

Governance in a computing environment where consumers access services or resources from multiple cloud providers (public or private) differs from governance when a consumer accesses services or resources from multiple service providers but without a cloud means of abstraction, thus direct or opaque access to each provider with its own technology stack. Multi-cloud governance can be organized around a centralized or global governance board tasked with making all decisions, or around a series of local decision bodies—one for each cloud provider in the multi-cloud environment—that are responsible for making decisions relating only to that cloud provider's services. A multi-cloud governance model can also be federated, with authority residing in the cloud consumers as well as the different service providers.

9.2.2. Stakeholders and Roles

Cloud consumers, providers, and other stakeholders in multi-cloud ecosystems, including regulators, auditors, and dedicated governance boards, play crucial roles to ensure good governance. Cloud consumers can choose any number of cloud services from different CSPs, each residing in different geographical and legal jurisdictions and operated by different organizations. Careful though, since they are ultimately responsible for the cloud workloads and data. Providers offer cloud services but may also be liable for breaches in data protection legislation, especially when cloud consumers fall under the protective jurisdiction of regulators having authority over the provider.

Regulatory authorities are responsible for policy-making in the public interest, and hence their influence can extend beyond their own jurisdiction. Cloud workloads transfer sensitive application data into the public domain, and regulators expect cloud providers to ensure data confidentiality, integrity, and availability while processing sensitive information. Therefore, proper cross-cloud data segregation, protection, and least-privilege access capabilities must be set in all CSPs. Auditors assess the compliance posture after each audit cycle, supplying evidence and assurance regarding the controls

implemented in the cloud consumer's multi-cloud environment through an exhaustive audit trail. Finally, a governance board can provide the necessary expertise and dedicated resources to address the multitude of strategic, operational, and assurance-related cross-cloud governance challenges. A centralized governance model improves transparency and accountability.

9.2.3. Core Principles of Governance in Multi-Cloud

A defining feature of governance is accountability: every stakeholder must answer for their actions and decisions and can be held responsible for their consequences. Governance frameworks provide a structure for accountability by specifying which individual or group has authority over a particular issue and what mechanisms ensure that it is exercised in accordance with established policies. This principle applies with equal force in cloud computing: cloud consumers cannot absolve themselves of responsibility for regulatory compliance simply because they have transferred their data to a third party, a cloud provider. Nor can cloud providers abrogate responsibility for the security, availability, or integrity of the services they offer. Indeed, the ostensible separation of responsibility provided by the use of multiple clouds creates potential new gaps, with each cloud operator liable only for its own environment.

The need for transparent, auditable, and harmonized services across multiple clouds, especially with respect to sensitive information, is essential in establishing responsible governance in a multi-cloud environment, particularly in the context of cross-jurisdictional privacy obligations or industry-specific requirements. To provide this assurance and reduce the risk of data loss or regulatory issues, enterprise governance mandates that both full data copies and a full audit trail of cross-cloud transactions be retained and available to a nominated auditor or governance board, with the requisite tools and processes to search, analyze, and respond within the required timeframe.

9.3. Architectural Models for Multi-Cloud Governance

Architectural Models for Multi-Cloud Governance

Centralized versus Federated Governance: A centralized governance model is pragmatic in the initial phases of multi-cloud adoption but may prove overly restrictive for mature deployments. The distribution of authority and decision rights in a multi-cloud environment must therefore evolve. Many organizations are shifting to a federated model that balances centralized and decentralized decision-making and elaborates the interactions among various governance processes. Indeed, the decentralized surfacing of information required for governance becomes increasingly important and difficult to

achieve in mature multi-cloud environments, and organizations should seek ways to assist these decentralized processes. The establishment of well-defined accountabilities is essential to avoid undue reliance on informal mechanisms of information sharing across the environment.

Policy Abstraction and Policy Translation Across Clouds: Three main layers constitute a governance policy in multi-cloud environments. Policy definition abstraction provides a simplified model of the different policy dimensions that monitors expanding into a multi-cloud environment introduce. Policy definition abstraction does not attempt to simplify the complexity of the policies for formal verification but enables consumers to understand the challenges faced across clouds and the additional elements to be dealt with in expanding to a multi-cloud environment. Cross-cloud policy translation specifies the translation mechanisms that generate the equivalent policies on different providers. Policy enforcement points on the clouds manage the enforcement of the policies on the specific providers, triggering actions that can span different providers.

Fig 9.2: Architectural Models for Multi-Cloud Governance

Identity, Access, and Entitlement Management: Despite the heterogeneity of clouds, organizations are encouraged to adopt a federated approach to identity and access management. This approach enables the use of multiple clouds while maintaining control of the end-users' identity and entitlements in a central repository. However, full federation of the identity stores may not always be possible due to policy limitations, interoperability issues, or the geographical distribution of the clouds. In highly regulated

or legacy environments, organizations tend to maintain independent identity and access management systems and trust domains for each provider, adopting a more controlled and cautious strategy for the identity federation, on a service-by-service basis. In this case, the use of data minimization techniques and the least-privilege principle must be enforced to mitigate the increased exposure.

9.3.1. Centralized versus Federated Governance

Governance in multi-cloud environments can be accomplished in two principal ways: through a central authority that enforces policy and applies sanctions across all cloud services, or through a federated governance board containing representatives of all cloud consumers, providers, and possibly even regulators. A centralized approach typically allows for more rapid decision-making and greater enforcement capabilities, but it is often criticized as being undemocratic in that it allocates decision rights to a single entity. A centralized model may also limit the scalability of governance to just a few clouds, as each decision must be considered by the central authority or committee.

A federated governance operating model enables decision making and resource control to be pushed down to the clouds being governed. As the cloud landscape expands, so does the requirement to remain agile and federated. One of the biggest challenges is policy adoption that supports governance without being too onerous on individual service providers. Wherever feasible, policy should be abstracted at a level that minimizes repetition and establishes policy intent, while allowing sufficient flexibility to address legitimate variations at the local level. Local compliance and risk steering committees should then be empowered to agree, or not, on the policies, processes, and controls these service providers will implement to maintain internal controls and risk posture across the services throughout their cloud vendor life cycle.

9.3.2. Policy Abstraction and Policy Translation Across Clouds

To achieve similar enforcement objectives while considering the inherent differences in the native policies, managed services, and supported security controls of the various underlying cloud services, the operation and enforcement of the security policies across the different cloud environments can be achieved through an appropriate abstraction and mapping of the policies against the specific services of the using cloud service provider. Therefore, a multi-cloud security framework that considers the different specifics and idiosyncratic differences for each of the using cloud service provider in terms of data encryption, key management, identity and access management, firewall and network security controls can provide a more effective approach to managing the core security aspects of security within a multi-cloud service environment.

A practical multi-cloud security approach can be divided into two layers (as depicted in Figure 10). At the higher abstraction layer, policy abstraction and translation, which is executed in a centralized manner within the cloud governance platform, are responsible for the definition of the security profile of the application. The policies or captions defined in this layer can be expressed in a high-level natural language description of the Cloud Security Posture Management (CSPM) preferences of the application covering components such as data protection requirements, IAM specification, network segmentation necessitation, and so forth. Based on the defined abstract security needs, the security requirements can then be mapped to the specific security controls and managed services of the deployed services on the various cloud service providers.

9.3.3. Identity, Access, and Entitlement Management

Identity and access management (IAM) controls protect resources from unauthorized access, yet they also enable privileged users to execute undesirable actions. Thus, a comprehensive strategy should consider both perspectives. IAM in multi-cloud environments supports authentication, authorization, and accounting across providers.

A federated trust framework using SAML 2.0, OpenID Connect, or Zertiificateless Public Key Infrastructure enables single sign-on across service providers and spoofs the user's identity during API calls without requiring API-related credentials. Lightweight directory access protocol provides read-only access to potentially sensitive information for applications making name-to-IP address resolution requests on the private networks separating service provider clouds with stateless, secure bridges. IAM controls help detect and prevent excessive privilege assignments based on the principle of least privilege and related processes.

Entitlement management maps business roles and associated entitlements to technical roles and entitlement systems for all applications. Control frameworks such as CIS, COBIT, PCI DSS, ISO 27001, and NIST SP 800-53 guide IAM strategy and decision-making for key products, processes, and information across the enterprise and supply chain at both the technical and functional levels.

9.4. Security Frameworks and Controls in Multi-Cloud

The threat landscape for multi-cloud environments is particularly rich and varied. Attackers can exploit vulnerabilities in the cloud services enabled by organizations using a multi-cloud model as well as weaknesses in the cloud services supporting these operations. Available resources and skills are also likely to shape attackers' profiles and their types of attacks. Shorter-lived kills and malware-as-a-service decrease the cost of

launching attacks. Advanced attackers often have the resources and skills required to perform spearheading training and executing penetration attack in the target's internal environment. Insufficient data segregation and broken authentication constitute a serious threat for organizations adopting vulnerabilities in the emphasize multi-cloud approach. Alwani and Arora also underscore the need to enhance data integrity through more complex security measures on service-abused third party tools and services. Cross-cloud risks and the temporary access to the cloud infrastructure limit internal data protection for multi-cloud users.

Data protection and encryption strategies are crucial for instilling trust in stored resources, either to sanitize the data to eliminate sensitive information or by encrypting the data. Few organizations adopt encryption for data at rest. Even fewer have a solution for managing decryption keys, which should be aligned with access policies. Organizations need to define data residency requirements and compliance regulations to avoid fines.

9.4.1. Threat Landscape in Multi-Cloud Environments

The potential benefits of multi-cloud environments have led to rapid adoption, yet companies must implement security measures to mitigate a broad spectrum of threats. Most organizations are not prepared for the complexity of the multi-cloud environment and its inherent risks. Cloud environments offer a large number of well-supported features for enabling security but organizations frequently configure them incorrectly. A study by Cybersecurity Ventures projected that the worst-case scenario for the future of ransomware could result in a total damages figure of USD 265 billion globally. Resources across multiple public clouds increase the area of exposure to themselves and also to workloads located in other clouds.

The multi-cloud environment is a target for attackers of different profiles; for instance, competitors do their best to spot misconfiguration, possibly with the intention of gaining some kind of advantage in the competitive landscape, while politically motivated actors may attempt to expose weaknesses in cloud security for propaganda purposes. Recent multiparty cloud configurations are presented as promising approaches to improve some aspects of security, but only when certain conditions apply. Finally, when partners of a federated ecosystem adopt a multiparty approach to cloud deployment, any compromise of trust with one of the partners exposes the entire ecosystem to a higher risk profile.

9.4.2. Data Protection and Encryption Strategies

Encryption plays a pivotal role in data protection strategies across multi-cloud environments by safeguarding sensitive data both during storage and transmission. Data at rest is typically encrypted within the storage services of the respective providers, utilizing proprietary key management and access control mechanisms. Consequently, it is important to understand the implications of regulatory requirements on key management policies, particularly with respect to jurisdictional constraints on data resident in specific countries and locations. Moreover, data may continue to be at risk of unauthorized disclosure even when undergoing encryption if required keys are not managed appropriately.

Data in transit between different cloud providers remains vulnerable and must be protected accordingly using well-established transport layer security protocols. Consequently, a safe and robust communication channel must be established to bridge the traffic between two clouds. In this regard, network segmentation between different cloud providers is critical in constructing a secure multi-cloud architecture. This may be supplemented with dedicated bridges for forwarding traffic. Furthermore, traffic flowing out of a cloud must be logged, inspected, and monitored to enable the detection of possible data exfiltration activity.

9.4.3. Network Security and Segmentation Across Providers

Segmentation between applications and services hosted by different cloud service providers is an important practice for reducing attack impact and exposure risk. Within each cloud, virtual networks and subnets can be used to provide segmentation, but when traffic between two virtual networks is routed across the Internet, zero trust principles need to be applied to the bridging (or peering) of these networks. Implementing any form of peering between different cloud providers creates a security gap that can be exploited if not managed correctly. Typically, network security should involve stateful firewalls on ingress and egress ports. Security group rules (for example, AWS) or firewall rules (for example, Google Cloud VPC Service Controls) covering traffic between different cloud services should follow zero trust security principles.

Private links or direct connections (for example, AWS Direct Connect, Google Cloud Interconnect) form secure point-to-point connections from on-premise data centers/SaaS applications to cloud providers. Such links can also restrict cloud APIs to be accessible only from these links. By default, network traffic bound for a cloud provider, without using these private links, traverses the Internet. Even private links can have potential security gaps and need to be regularly audited. Networks crossing untrusted boundaries

should be monitored for data leaks. For example, data packets with sensitive data passing through the untrusted boundary should be tagged and monitored.

Despite having some of the largest security teams, breaches are continuously exploited in cloud infrastructures. Network segmentation across providers through zero-trust principles and proper monitoring is essential to help detect, prevent, and alert such security gaps within the network and applications. Continuous Traffic Monitoring and Log Analysis at the Cloud Security Monitoring Centre (CSMC), preferably implemented through a third-party SOCT provider, provides complete visibility of the responses and actions for all alerts and incidents.

9.5. Compliance Management in Multi-Cloud

Governance, Security, and Compliance in Multi-Cloud Environments

Managing compliance through multiple cloud service providers in a way that satisfies organisational and institutional regulators and auditors across multiple countries is a challenge for any business using cloud computing services. Businesses must ensure they meet the relevant regulatory requirements for their industry (e.g. PCI DSS for organisations processing credit card data), the applicable laws of jurisdictions in which they operate (e.g. GDPR for EU residents, CCPA for residents of California), and regulations imposed on them by law enforcement authorities (e.g. non-disclosure, child protection, financial crime) and by sponsors (e.g. US export control laws). In many cases, it will be necessary to comply with the laws of more than one jurisdiction, e.g. by meeting both EU and US requirements for cross-border data flow. Multi-cloud compliance also requires mapping cloud service deployments against various audit frameworks (e.g. ISO27001, NIST, SOC2), ensuring adequate logging and audit-trail capabilities are available, and ensuring third-party assurance reports on compliance with the above frameworks and regulations are available for examination.

In some industries, jurisdictional considerations can be particularly challenging. In the case of health data stored in cloud services operated by US service providers, the intersection of GDPR and HIPAA becomes critical. GDPR provides European Union residents with an explicit "right to be forgotten." When that right is exercised, an organisation must delete all data related to that individual from all storage devices, backups, etc., regardless of whether the data originates in the EU or has simply been ingested by a service based in the EU. HIPAA, however, requires that health data be retained for a minimum of 6 years, and the Health and Human Services office must have access to that data at any time. Failure to comply with the request of the HHS office is a criminal offence that can carry penalties of imprisonment. In this case, an EU-based health organisation using a US-based cloud service (or possibly multiple US-based cloud

services) for health data subject to both GDPR and HIPAA clearly cannot comply with both regulations, as each requires a direct violation of the other.

Fig 9.3: Compliance Management in Multi-Cloud

9.5.1. Regulatory and Jurisdictional Considerations

A wide variety of laws and regulations could apply to the acquisition and use of cloud services. The most prominent are those corresponding to the sector in which an enterprise operates or where the processed data originate. These are usually regulatory obligations imposed by the bodies responsible for overseeing the applicable sector. Telecommunications and financial services are sectors – among many others – with established regulations. One of the main characteristics of these obligations is that they are often jurisdiction-specific and rarely harmonized across borders. Although many enterprises prefer the cloud for its implied "location independence," using cloud services could expose them to data residency constraints and ensure the compliance of sensitive data with local regulatory obligations. Jurisdictional issues usually arise when the cloud service services across borders. Such geographical movements will require additional controls. Global cloud providers typically distribute regional data across multiple datacenters. These pieces of infrastructure are often in foreign jurisdictions and are often located in countries where regulatory requirements are neither adopted nor, more importantly, enforced.

One way to analyze the possible jurisdictional-related risks is to associate each piece of data with a regulatory framework. By mapping the known regulatory requirements related to data types and availability domains, it is possible to identify the regions where the controls are not implemented. Cross-border movements of data types for which no regulatory requirements exist are treated as sensitive, as enterprise management has not imposed limitations on them. Mapping all controls of all regulations applicable to an enterprise's cloud usage (service providers and data locations) to the cloud services can provide an indication of the degree of overlapping requirements. Having multiple controls for the same threats posed by multiple regulatory obligations can influence the choice of cloud services providers.

9.5.2. Mapping Compliance Requirements to Cloud Services

Many sectors and jurisdictions impose regulatory requirements on organizations, establishing general principles for risk management, breach notification, data protection, and prevention of illegal content. Specific compliance frameworks exist for specific industries, such as the Payment Card Industry Data Security Standard for the payment card industry, the Federal Risk and Authorization Management Program for US federal agencies, the Health Insurance Portability and Accountability Act for the health sector, and the Sarbanes–Oxley Act for components of publicly traded companies.

The compliance requirements must be mapped to the services and capabilities provided by the cloud vendors. Compliance-specific controls should be aligned with applicable laws and frameworks. In addition to the cloud service provider's own compliance certification, cloud users and compliance officers should inspect the types of evidence generated by the cloud vendor—for example, artifact type, evidence granularity, and evidence updating frequency—and design a solution that collects compliance support evidence required by regulations and indicates that the burden of proving compliance has been shared with relevant vendors.

9.5.3. Auditability, Evidence, and Assurance

Operational integrity mandates that all multi-cloud systems, as well as all transactions and events processed across multi-cloud environments, provide audit and logging capabilities. This capability is of paramount importance for both verification of compliance with regulatory requirements and for the forensic analysis needed to detect intrusions. An enterprise information security and risk management strategy should require that all multi-cloud service provider environments and the solutions executing within the provider environments incorporate logging and reporting. These requirements should not be limited to event logging by the service provider; solutions executing on

the multi-cloud service provider should also record events within the applications' logs. Such logs are vital for incident detection, analysis, and response. Specific requirements must be established for cross-cloud scenarios, where a service provider relies on a partner cloud service provider.

The logs generated should be detailed, appropriately retained, and easily associated with the multi-cloud service provider's shared responsibility model. Controls governing the logging and reporting capability should include, at a minimum, the following elements: actions taken by users with administrative privileges; log-on attempts and log-off events; access to directories and files, including successful and failed accesses; reliance on system alerts for event monitoring and analysis; redundant inspection of the multi-cloud service provider's systems from different locations and persons; collection and storage of log information in a format suitable for incident investigation and analysis; collection of a time-based reference for the coalescing and correlating of log information from different sources; verifiable collection of logs; and the configuration of alerts on log information, as required. Multi-cloud services should provide evidence that supports compliance evaluations for applicable regulations, legislation, and agreements.

Independent third parties periodically assess compliance and issue attestation reports. Such attestation, although considered evidence of effective compliance, should not be viewed as offering assurance that a third-party service provider is protected against all threats. Enterprise-specific security and privacy requirements should still be defined and be imposed on the multi-cloud service provider's product and service offerings.

9.6. Risk Management and Resilience

Vulnerabilities in one cloud provider can threaten the entire multi-cloud environment. Multi-cloud risk management encompasses risk assessment, business continuity and disaster recovery, and vendor risk management.

Probabilistic and qualitative risk assessment methodologies are applicable to multi-cloud environments. Affinity-based risk catalogs extend beyond, and cover other cloud-specific risks. Their use requires continuous monitoring and periodic reassessment.

Business continuity and disaster recovery planning demands extra attention in multi-cloud environments. Key decisions involve the definition of recovery time and recovery point objectives, data redundancy between public clouds, and the requirement for local copies. These elements must be documented in a business continuity plan and business impact analysis, tested on a regular basis, and considered during a disaster recovery exercise.

When executing risk assessments, organizations must also assess risks associated with vendor services. Frameworks supporting this activity address many aspects deemed relevant for vendor risk management, such as third-party classification, contracts, monitoring, and change management. A complete frame is obtained by integrating all these dimensions into a single vendor risk management strategy.

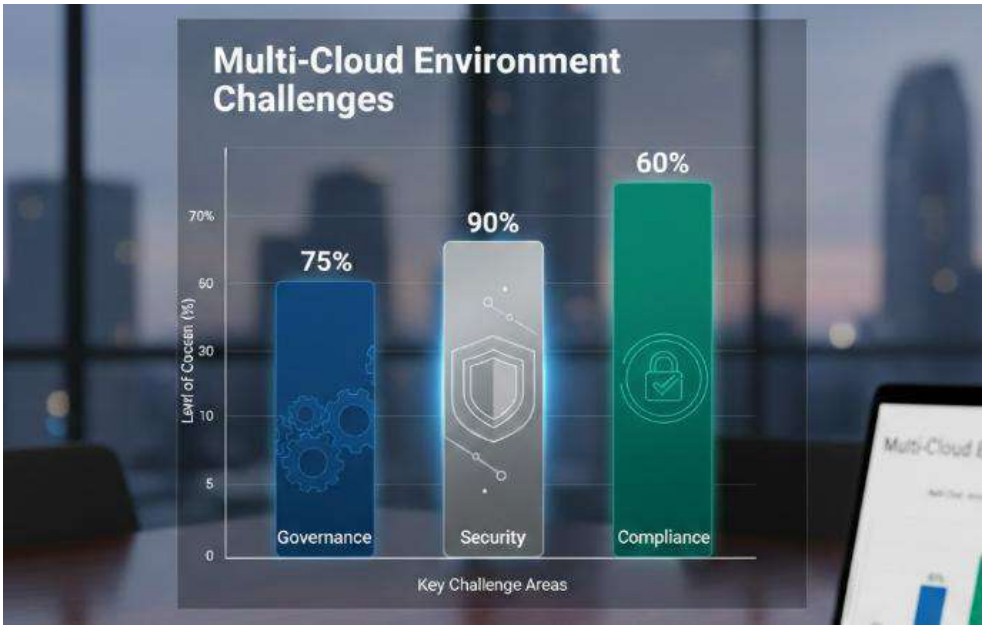

Fig 9.4: Governance, Security, and Compliance in Multi-Cloud Environments

9.6.1. Risk Assessment Methodologies for Multi-Cloud

Approaches to assessing risk in multi-cloud environments are similar to those employed in other settings. The risks associated with specific services offered by the underlying cloud providers are well understood and documented. These risks can be further alleviated through the specialized construction of applications—such as the use of multiple managed databases, the recommendation of any of a large range of storage services, and so on. The majority of cloud consumers can, therefore, obtain a reasonable level of assurance in relation to their risk profile by following conventional (though labor-intensive) methods based on control matrices.

Nevertheless, fully multi-cloud systems—applications that fully and simultaneously exploit multiple public cloud providers to provide their services—are rare. The key question is whether such crossing of cloud borders introduces a new set of risks, whether those risks can be identified, and whether their probability and impact can be assessed. Simple probabilistic models are often used to reasonably assess the reliability of fully

redundant systems (such as an application deployed across two geographically diverse cloud providers); qualitative models have been developed for such situations. In other cases, the risk profile of a cloud consumer may be altered by dependent or correlated outages (for example, a malicious insider at Cloud Provider A who deletes records at the request of the governing regime of Cloud Provider B; in this situation, simply using two geographically diverse cloud services would not protect the cloud consumer). Thus, eliciting and cataloging the probable events that could lead to harmful situations is essential.

9.6.2. Business Continuity and Disaster Recovery Across Clouds

Multi-cloud environments must ensure service continuity and resilience against outages or catastrophes. Key requirements include defining maximum tolerable periods of disruption, devising comprehensive redundancy strategies, setting recovery point and time objectives, and regularly testing against these objectives. Redundant solutions can be developed across a single provider, different providers, or combinations of on-premises and public cloud services. The provider's service level agreements establish the contractual basis for redundancy plans relying on services from the same provider. Businesses relying on single or dual-zone cloud services may combine geographically distinct services from the same provider. Redundant solutions should not introduce new risks or increase overall risk exposure. Consolidation can present opportunities to reduce costs.

Redundant infrastructure should enable a rapid restoration of services, including data, applications, and systems. Recovery Time Objectives for each service should reflect acceptable periods of downtime and disruption. More complex services incur longer recovery period objectives. Back up strategies should meet Recovery Point Objectives, with a range of options available for different requirements: logging of transactions, mirroring within a single provider, or scheduled backups. Data stored with different cloud service providers should be regularly synchronized. Recovery Time Objectives are often hardest to meet for complex IT services supporting business-critical processes. Backup and recovery plans should be tested regularly, across all service dependencies and scenarios.

9.6.3. Vendor Risk Management and Third-Party Oversight

Supplier risk assessments examine the potential impact of the cloud provider on the consumer's security and compliance posture. The more sensitive the data stored in the supplier's environment, the more rigorous the assessment should be, with financial, legal, and technical support underpinning the process. A standard questionnaire—such

as the one developed by the Shared Assessments Program, the Cloud Security Alliance, or the Action advisable ICC C5 standard—should be applied as a minimum. Consumer organizations must also review the provider's responses and evidence carefully, identify gaps, and define any additional controls to be implemented by the service consumer. When outsourcing to multiple cloud providers, consumers should ensure that both sides of the relationship are adequately covered—for example, by addressing the technical due diligence needed to establish secure and reliable connections between services.

Due diligence should be included in the supplier contract, especially concerning mandatory reporting requirements. These might take the form of adjustment clauses or even payments dedicated to regulatory fines resulting from a violation attributable to the provider and its policies. Consumer organizations should monitor supplier operations periodically by repetitive questionnaires, evidence collection, and audits, within the authorizations contained in the contracts. Finally, the cloud risk profile should be included in the organization's second or third-party risks, with probabilities assigned to the monitored proofs, periodic controls, and limitations envisaged in case of excessive consumer risk exposure. The consumer must also consider its own role as a supplier when using different cloud services, implementing a frequent-check mechanism to guarantee the persistence of the overall security posture.

9.7. Conclusion

Shifting sociopolitical and technological spheres necessitate novel services that enhance value-driven dynamic supply chains. Multi-cloud appears to be the greatest enabler, providing access to worldwide best-of-breed services and innovation. Multinational companies focused on the European market must therefore embrace multi-cloud, deftly managing a potentially overwhelming complexity involving the deployment of services from different Cloud Service Providers (CSP) operating under distinct regulatory regimes and sets of Terms and Conditions. Surmounting this challenge entails addressing the multitude of interrelated governance, Security, Risk, and Compliance issues involved in such a complex setup, with the object of effectively and efficiently leveraging the multi-cloud paradigm while meeting the specific requirements articulated by these stakeholders and the additional challenges of risk management and resilience.

Governance, Security, Compliance, and Risk form the GSEC-R triangle: multi-cloud security, risk, and compliance management are but components of an overarching governing framework – the governance of multi-cloud environments. Hence, it seems appropriate to elaborate on this concept before drilling down into GSEC-R considerations. Multi-cloud governance is defined as the establishment of an executive framework for adopting, managing, and monitoring a multi-cloud environment, within which Cloud Consumers leverage services from multiple Cloud Service Providers –

operating under different legal jurisdictions, security and privacy standards, and contractual terms– to meet their requirements for Security, Risk management, Compliance, and other concerns.

9.7.1. Future Directions

Research on multi-cloud governance, security, and compliance is in a nascent stage and requires further exploration. Applying established multi-cloud concepts across heterogeneous cloud types would provide a broader perspective. Multi-cloud governance research should expand its focus beyond consumers and private infrastructures to central government and federation perspectives. The applicability of network security concepts, such as segmentation and micro-segmentation, to different cloud types warrants investigation. Additionally, the intricacies of securing data in cross-cloud contexts deserve attention.

The development of a compliance management approach specifically tailored to multi-cloud environments remains a vital area of research, alongside the construction of comprehensive threat, vulnerability, exploits, and incident taxonomies based on empirical studies. A roadmap for risk assessment enhancement would benefit both probabilistic and qualitative risk assessment methods. Perennial topics such as business continuity and disaster recovery in multi-cloud settings also require sustained inquiry, particularly the orchestration of supplier redundancies and recovery objectives. Research on cloud vendor risk management has concentrated on public clouds and financial institutions; expansion to different supplier types and sectors would enhance maturity..

References

Alouffi, B., Hasnain, M., Alharbi, A., Alosaimi, W., Alyami, H., & Ayaz, M. (2021). A systematic literature review on cloud computing security: threats and mitigation strategies. IEEE Access, 9, 57792–57807.

Rongali, S. K. (2025). Balancing AI and human collaboration. World Journal of Advanced Research and Reviews.

Alonso, J., et al. (2023). Understanding the challenges and novel architectural models of multi-cloud native applications–a systematic literature review. Journal of Cloud Computing, 12(1), 6.

Guntupalli, R. (2025, August). Cloud-Native AI: Challenges and Opportunities in Infrastructure Security. In 2025 International Conference on Artificial Intelligence and Machine Vision (AIMV) (pp. 1-4). IEEE.

Borse, H. C. (2025). Improving security and compliance: A hybrid blockchain-based multi-cloud governance and auditing framework. International Journal of Engineering Development and Research, 13(4), 102-103.

Inala, R. Designing Scalable Technology Architectures for Customer Data in Group Insurance and Investment Platforms.

Chinnasamy, A., Ahmad, R., & Hassan, R. (2021). Challenges and opportunities of compliance automation in cloud. IEEE Transactions on Cloud Computing, 9(3), 882–895. https://doi.org/10.1109/TCC.2020.2965120

Dawood, M., Tu, S., Xiao, C., Alasmary, H., Waqas, M., & Rehman, S. U. (2023). Cyberattacks and security of cloud computing: A complete guideline. Symmetry, 15(11), 1981. https://doi.org/10.3390/sym15111981

Amistapuram, K. (2022). Fraud Detection and Risk Modeling in Insurance: Early Adoption of Machine Learning in Claims Processing. Available at SSRN 5741982.

Edwards, J. (2024). Cloud security. Mastering Cybersecurity, 223-280. https://doi.org/10.1007/979-8-8688-0297-3_8

Garapati, R. S. (2022). AI-Augmented Virtual Health Assistant: A Web-Based Solution for Personalized Medication Management and Patient Engagement. Available at SSRN 5639650

Khan, A. A., Niazi, F., & S. A. Khan. (2021). Automated governance in multi-cloud environments using policy-as-code. Future Generation Computer Systems, 125, 742–754. https://doi.org/10.1016/j.future.2021.07.022

Aitha, A. R. (2021). Dev Ops Driven Digital Transformation: Accelerating Innovation In The Insurance Industry. Available at SSRN 5622190.

Khan, N., Alsaqour, R., Shah, A., & Alabdulatif, A. (2021). Security and privacy frameworks for multi-cloud computing: A systematic review. Journal of Cloud Computing, 10(1), 1–20. [9]Mukherjee, A., & Tripathi, S. (2021). Blockchain-enabled compliance and audit trails for cloud security. IEEE Cloud Computing, 8(4), 62–71. https://doi.org/10.1109/MCC.2021.3089974 [10]Naik, S. (2023). Cloud-based data governance: Ensuring security, compliance, and privacy. The Eastasouth Journal of Information System and Computer Science, 1(01), 69–87. https://doi.org/10.58812/esiscs.v1i01.452

Nagabhyru, K. C. (2023). From Data Silos to Knowledge Graphs: Architecting CrossEnterprise AI Solutions for Scalability and Trust. Available at SSRN 5697663.

Petcu, D. (2022). Portability and interoperability between clouds: Challenges and case studies. Journal of Grid Computing, 20(1), 5–27.

P S L Narasimharao Davuluri. (2023). Integrating Artificial Intelligence into Event-Driven Financial Crime Compliance Platforms. International Journal Of Finance, 36(6), 707-736. https://doi.org/10.5281/zenodo.18457715

Pourmajidi, W., Zhang, L., Steinbacher, J., Erwin, T., & Miranskyy, A. (2025). A reference architecture for governance of cloud native applications. IEEE Transactions on Cloud Computing, 13(3), 935-952. https://doi.org/10.1109/TCC.2025.3578557

Prasath, R., & Fariz, M. (2026). Multi-cloud security and privacy models for distributed enterprise systems. ISCSITR-International Journal of Cloud Computing, 7(1), 1-6.

Rahman, M., & Islam, S. (2023). Strengthening data governance in multi-cloud environments: A framework for security, compliance, and operational resilience. International Journal of Multidisciplinary and Current Educational Research, 7(6), 71-83.

Gottimukkala, V. R. R. (2021). Digital Signal Processing Challenges in Financial Messaging Systems: Case Studies in High-Volume SWIFT Flows.

Zhang, Q., Fitzpatrick, L., & Boehm, B. (2021). Security architectures for multi-cloud environments: A comprehensive review. IEEE Transactions on Cloud Computing, 9(4), 1397–1412.