

OTRA: A Risk Management Ontology for Transparent Service Level Agreements in Federated Cloud Environment

Giulia Biagioni¹ ^a, Resul Serkan Keskin¹ ^b, Lawrence Cook² ^c, Edwin Harmsma¹ ^d and Erik Langius¹ ^e

¹Netherlands Organisation for Applied Scientific Research (TNO), The Hague, The Netherlands

²BIT, Ede, The Netherlands

{giulia.biagioni, serkan.keskin, edwin.harmsma, erik.langius}@tno.nl, lawrence@bit.nl

Keywords: Cloud Federation, Service Level Agreement, Risk Management Ontology, Bow-Tie RDF-Based Implementation.

Abstract: Cloud Federation, a model where multiple cloud providers collaborate to offer interconnected services, has become increasingly renowned in cloud computing. In this environment, Service Level Agreements (SLAs) play a critical role by formalising service expectations and fostering transparency between providers and consumers. Although existing SLA ontologies provide a semantic foundation for standardised SLA management, they lack components for detailed risk management and representation within SLA. This paper presents our work on developing OTRA - Ontology for Transparent Agreement - specifically designed for risk management within SLAs in federated cloud environments. By integrating structured risk identification, prevention and mitigation strategies, our ontology extends the Gaia-X Core Ontology. We validate our ontology through a controlled simulation that demonstrates its effectiveness in structuring SLA risk information, offering a foundation for transparent, accountable service agreements.


1 INTRODUCTION


In the evolving landscape of cloud computing, cloud federation - a model where multiple cloud service providers (CSPs) collaborate to deliver an interconnected network of services (Kurze et al., 2011) - is gaining increasing popularity. This environment offers significant benefits such as data sovereignty, which allows organisations to store and manage their data within specific jurisdictions, ensuring compliance with regional regulations (Holfelder et al., 2022)(De Filippi and McCarthy, 2012). Additionally, cloud federation helps avoid vendor lock-in, giving organisations the flexibility to choose services from different providers without being tied to a single vendor's ecosystem (Bermbach et al., 2013). This flexibility ensures better resource optimisation, cost management, and control over critical infrastructure choices (Bonfiglio, 2021).


Within this model, Service Level Agreements (SLAs) play a crucial role. SLAs are formal contracts that define the expected level of service between providers and consumers. They ensure transparency and accountability, helping to maximise value for all parties by clearly defining expectations and mitigating legal risks (Lê and Nguyen, 2020).


Several ontologies have been developed to represent SLAs in the context of cloud federation, such as CSLAOnto (Labidi et al., 2016) and the generic SLA ontological model of SLAVIDES (Karamanlioglu and Alpaslan, 2018). These ontologies address the need for standardised, automated, and flexible SLA management. They offer a semantic-based format that enables automatic monitoring, understanding and enforcement of SLAs, ensuring that service agreements are clearly defined and adhered to. When integrated into expert systems, they enable the detection of SLA violations across multiple sectors, potentially leading towards improving compliance and performance management in federated cloud ecosystems.


Despite the considerable efforts made in the landscape of SLA ontology development, semantic models specifically designed to delineate risks in SLAs for cloud federation have not been developed to the same

^a  <https://orcid.org/0000-0002-9005-7945>

^b  <https://orcid.org/0000-0002-4844-7982>

^c  <https://orcid.org/0009-0008-0536-0511>

^d  <https://orcid.org/0009-0009-8892-4593>

^e  <https://orcid.org/0009-0006-2765-6781>

extent. Semantic-based formats that could be fed into expert systems to trigger strategies aimed at preventing the occurrence of hazardous events, or mitigating their impact are still under development. This gap represents an opportunity for further advancements in the integration of risk management and prevention mechanisms into the SLA frameworks used in cloud federation.

In this article, we introduce our work on developing OTRA, a specialised ontology for risk management within SLAs in federated cloud environments. This ontology is designed to fill existing gaps in SLA frameworks by explicitly modelling and representing risks, which are essential in cloud federation scenarios where multiple Cloud Service Providers (CSPs) interact and collaborate across varying regulatory environments. The ontology aims to provide a structured and semantic foundation for risk identification, prevention, and mitigation within SLAs, thereby enhancing transparency and accountability for both providers and consumers.

This article is structured as follows: following the introduction, Section 2 provides an overview of the state of the art in ontology development for risk management and SLAs, summarising existing approaches and identifying gaps that motivate our work. Section 3 details the ontology design of OTRA, presenting its conceptual underpinnings through First-Order Logic (FOL) axioms. This approach ensures that our conceptualisation remains technologically agnostic, allowing for flexible implementation across different languages and evolving standardised semantics.

In Section 4, we describe the evaluation process of the ontology, including a simulation designed to validate its effectiveness in a controlled operational environment, achieving a Technology Readiness Level (TRL) of 5. Section 5 discusses the results of this evaluation, interpreting the ontology's performance and insights gained from the simulation. Finally, Section 6 outlines future work, exploring potential extensions to enhance the ontology's applicability and expressiveness in dynamic cloud federation contexts.

2 STATE OF THE ART

Ontologies, in the context of RDF, are formal representations of knowledge that define concepts, entities, and the relationship between them within a specific domain. Regarding their technical application, they can be used for enabling interoperability and connecting data silos (Roussey et al., 2011)(Allemang et al., 2020)(Blumauer et al., 2020), feeding expert systems with structured knowledge (Prapakorn and Chittaya-

sothorn, 2009), where example of such applications can be found in domains like health (Mustafa et al., 2023) and energy (Pruvost et al., 2023), allowing machines to understand the meaning of information (Hogan, 2020), and enhancing AI trustworthiness by providing explainable reasoning when integrated with machine learning-driven methods (Pan et al., 2023).

In risk management, the Open Risk community¹ has contributed significantly to developing ontological models, promoting open-source risk analysis tools for transparency within financial sectors. Their frameworks, such as the Description of a Model Ontology (DOAM)² and the Risk Function Ontology (RFO)³, offer structured means to annotate and categorise key concepts in risk modelling and management roles.

DOAM, for example, facilitates the semantic organisation of risk model components, while RFO structures essential skills, occupations, and functions to support interoperability. Though valuable for analysis, both ontologies prioritise metadata-level annotation, which limits their direct application in expert systems or within service-level agreements (SLAs) between providers and consumers in cloud federations.

Incorporating detailed risk-related information within SLAs between providers and consumers is critical for: reducing ambiguity, ensuring awareness of conditions, clarifying mitigation strategies, and making implicit assumptions explicit (Uriarte et al., 2016) (Alkhamees, 2022). If DOAM and RFO lack the capability to directly map risk frameworks onto data due to their analytical focus and financial sector orientation, SLA-specific ontologies like SLAVIDES and CSLAOnto similarly miss essential components for representing risk and management strategies within service agreements.

CSLAOnto is specifically designed for cloud computing SLAs, focusing on standardising terms related to service quality, performance, and responsibilities (Labidi et al., 2016), thus enhancing semantic interoperability but not directly addressing risk management. SLAVIDES (Karamanlioglu and Alpaslan, 2018), while formalising SLA terms, similarly lacks built-in risk representation and management components. Consequently, these ontologies offer valuable structure but are not yet at a readiness level for comprehensive, risk-sensitive applications in operational

¹Regarding the Open Risk Community, cf: <https://www.openriskmanagement.com>

²<https://www.openriskmanual.org/ns/doam/index-en.html>

³<https://www.openriskmanual.org/ns/rfo/index-en.html>

environments.

Regarding ontological efforts aimed at fostering interoperability in cloud federations at an operational and industry-driven level, the Gaia-X initiative offers a notable example. Gaia-X is an initiative dedicated to developing digital governance that can be applied to foster transparency, controllability, portability, and interoperability in the federated cloud ecosystem and data spaces⁴. Among its various activities, Gaia-X supports the collaborative development of ontological modules that enable seamless data and service integration across cloud providers. One of its foundational ontological modules is the Gaia-X Core Ontology, which structures key concepts such as Participants (e.g., Provider, Consumer), Resources (e.g., Data, Software Asset), and Service Offerings to enhance compatibility across cloud systems. While the Gaia-X ontologies effectively establish foundational roles and service interactions, they currently lack the necessary components to define a comprehensive risk management framework within SLAs between providers and consumers.

To address these identified gaps, we have developed a new risk management ontology specifically designed for SLAs, aiming to comprehensively represent risk and management strategies within service agreements. This ontology extends the Gaia-X Core Ontology module by facilitating structured, actionable risk information, allowing its seamless integration within SLAs. The following section presents our framework in detail, highlighting its potential to enhance transparency and risk management within federated cloud environments.

3 ONTOLOGY DESIGN FOR RISK MANAGEMENT

Our risk management ontology, OTRA, is grounded in the bow-tie method⁵. The bow-tie method is a structured risk assessment model used in industries to visualise risks by organising both preventive measures and mitigation strategies around a central hazardous event (Omidvar et al., 2024) (Kluwer, 2017) (Kluwer, 2019). Resembling a bow tie, the model consists of two branches extending from the "knot" or core event: the left side represents potential threats and preventive actions, while the right side outlines

potential consequences and mitigation measures. The bow-tie model is displayed in Figure 1.




Figure 1: The bow-tie diagram for risk assessment.

Industries such as oil and gas (Kbah et al., 2020), healthcare (McLeod and Bowie, 2020), and aviation (Aust and Pons, 2020) use the bow-tie method to develop clear action plans for preventing and addressing the outcomes of hazardous events.

Our ontology structures each *Hazardous Event* to include essential elements like *Preconditions*, *Preventive Mechanisms*, *Postconditions*, and *Remedies*, closely mirroring the bow-tie method's approach to capture potential triggers and responses. This alignment allows our ontology to systematically map risks and mitigation strategies that apply to SLAs in federated cloud environments. Key elements include:

- **Hazardous Event:** Identifies risks that may impact SLA parties, linked to *Preconditions* (factors that may trigger the event) and *Postconditions* (the resulting outcomes or states after the event).
- **Prevention Mechanism:** Outlines preventative actions aimed at reducing identified risks.
- **Remedy:** Details corrective measures to address adverse events if they occur.
- **Agreement:** Serves as the core SLA element, connecting Gaia-X Consumer and Gaia-X Provider roles and embedding risk considerations within the SLA framework.

In our formal delineation of the ontology, which subsequently guided its technical implementation, each *Hazardous Event* x , representing a risk within SLA agreements, is linked to a *Precondition* y that may act as a trigger. This relationship is expressed in First-Order Logic (FOL) as:

$$\forall x (\text{HazardousEvent}(x) \rightarrow \exists y (\text{Precondition}(y) \wedge \text{hasPrecondition}(x, y)))$$

To reduce the likelihood of hazardous events occurring, each *Hazardous Event* is linked to a *Prevention Mechanism* w , which is specifically designed to counteract its occurrence:

⁴<https://gaia-x.eu/what-is-gaia-x/about-gaia-x/>

⁵The bow-tie method reflects conceptual structure outlines in the ISO 31000:2018. This standard identifies eight core concepts critical for effective risk management: Risk, Risk Management, Stakeholder, Risk Source, Event, Consequence, Likelihood and Control.

$$\begin{aligned} \forall x (\text{HazardousEvent}(x) \rightarrow \\ \exists w (\text{PreventionMechanism}(w) \\ \wedge \text{hasPreventionMechanism}(x, w))) \end{aligned}$$

Each *Precondition* y associated with a *Hazardous Event* is also assigned a *Prevention Mechanism* w to address conditions that could lead to potential risks. We represent this relationship in First-Order Logic as follows:

$$\begin{aligned} \forall y (\text{Precondition}(y) \rightarrow \\ \exists w (\text{PreventionMechanism}(w) \\ \wedge \text{hasPreventionMechanism}(y, w))) \end{aligned}$$

The `hasPreventionMechanism` relationship includes an OR operator in its domain, applying to either a *Precondition* or a *HazardousEvent*, with a range on *PreventionMechanism*.

When a *HazardousEvent* x occurs, it produces an outcome represented by the *Postcondition* z , indicating the resulting state:

$$\begin{aligned} \forall x (\text{HazardousEvent}(x) \rightarrow \\ \exists z (\text{Postcondition}(z) \wedge \text{hasPostcondition}(x, z))) \end{aligned}$$

Lastly, each *HazardousEvent* is connected to a *Remedy* v , which serves as a corrective action to prevent the resulting *Postcondition* from manifesting. This relationship can be formalised as:

$$\begin{aligned} \forall x (\text{HazardousEvent}(x) \rightarrow \\ \exists v (\text{Remedy}(v) \wedge \text{hasRemedy}(x, v))) \end{aligned}$$

Each *Postcondition* z associated with a *HazardousEvent* is similarly assigned a *Remedy* v , ensuring proactive response measures. This relationship is represented in First-Order Logic as:

$$\begin{aligned} \forall z (\text{Postcondition}(z) \rightarrow \\ \exists v (\text{Remedy}(v) \wedge \text{hasRemedy}(z, v))) \end{aligned}$$

Ultimately, similar to the nature of the `hasPreventionMechanism` relationship, the `hasRemedy` relationship includes an OR operator in its domain, indicating it applies to either a

Postcondition or a *HazardousEvent*, with a defined range on *Remedy*.

The *Hazardous Events*, representing risks in SLAs, are inherently linked to *Agreements*. The *Agreement* concept formalises the binding between specific risks and the involved parties. An agreement a connects to a hazardous event x via the `hasRisk` relationship, embedding risk considerations within the contractual terms:

$$\begin{aligned} \forall a (\text{Agreement}(a) \rightarrow \\ \exists x (\text{HazardousEvent}(x) \wedge \text{hasRisk}(a, x))) \end{aligned}$$

Each agreement a also incorporates both a consumer c and provider p , represented by the `involvesParty` relationship. This connects the entities directly accountable for the risk terms outlined within the SLA. Formally, we represent this as:

$$\begin{aligned} \forall a (\text{Agreement}(a) \rightarrow \\ \exists c, p (\text{Consumer}(c) \wedge \text{Provider}(p) \\ \wedge \text{involvesParty}(a, c) \wedge \text{involvesParty}(a, p))) \end{aligned}$$

The ontology, as formally described above, has been technically implemented using OWL and RDFs standardised semantics, capturing the relationships between key elements such as hazardous events, preconditions, prevention mechanisms, postconditions, and remedies. This structured semantic model can be visualised as shown in Figure 2, where rectangles represent classes, and arrows illustrate object property relationships.

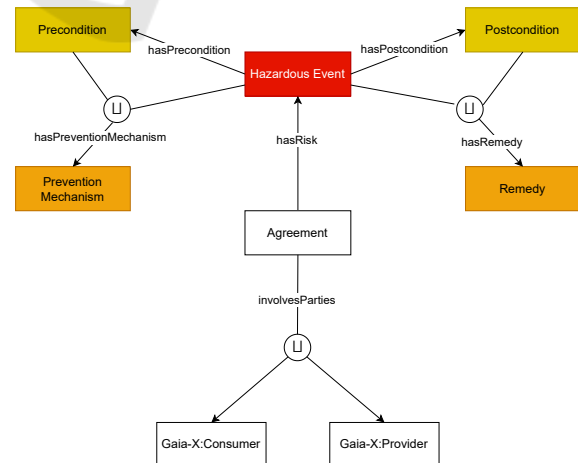


Figure 2: Visual representation of the ontology.

The OWL and RDF-based ontology is accessible for reference and collaborative contributions at the following repository: <https://github.com/OTRA-anon/OTRA>.

4 EVALUATION PROCESS

OTRA ontology was validated to ensure both structural consistency and relevance in representing domain knowledge effectively. For consistency, we applied the Hermit 1.4.3.456 reasoner to check the ontology's logic and coherence with and without instance data, confirming its integrity.

To further test the ontology's effectiveness in organising SLA risk information, we developed a simulation of a digital marketplace environment during the "Tech-X Conference and Hackathon 2024" event. This conceptual marketplace was designed to enable cloud service providers (CSPs) to offer services such as IaaS, PaaS, and SaaS, with simulated consumers capable of selecting and combining services across multiple providers.

For seamless service interoperability within the simulated cloud federation, we proposed Ligo⁶, an open-source Kubernetes extension, as a mechanism to facilitate peering between providers (Iorio et al., 2022). Although the simulation did not involve actual users or infrastructure, our conceptual pipeline leveraged Ligo's capability to create "virtual nodes" that represent the capacity of a remote cluster within the local cluster's environment. This approach enabled simulated clusters to expand their resource pool dynamically, supporting resource sharing and flexible workload distribution across Kubernetes clusters in a federated cloud context.

In this model, the consumer cluster, depicted on the left part of Figure 3, initiates an outgoing peering relationship with a provider cluster on the right, thereby gaining access to additional resources from the provider.

Within this simulated federation, SLAs played a foundational role in formalising service agreements. We embedded our risk management ontology in the SLAs to clearly outline potential risks, preventive measures, and remedies, thereby creating a structured, transparent framework for risk management. This approach aimed to foster informed decision-making and trust within federated cloud environments.

The SLA template from BIT⁷ was used as a case




Figure 3: Illustration of Ligo-enabled resource sharing in a federated cloud environment. In the figure, the directed arrow signifies the peering connection, which allows the consumer to offload workloads to the provider, mirroring the provider's capacity within its own infrastructure through a "virtual node." This virtual node in the consumer cluster represents the resources of the provider cluster, enabling workloads or pods to run on the remote infrastructure while remaining under the consumer's management framework.

study in our simulation to demonstrate how the ontology could support digital, legally binding agreements. This allowed all parties (providers and consumers) to view potential risks and planned response strategies.

In the simulated SLA, an agreement labelled *Agreement34* was formed between the consumer, *Alice*, and the provider, *Company X*. This agreement contained a termination clause specifying conditions under which *Company X* could end the contract. According to this clause, *Company X* reserved the right to terminate the SLA if *Alice* engaged in harmful or non-compliant behaviours, such as unauthorised access, denial-of-service attacks, or policy violations. This termination scenario, modelled as the *Hazardous Event risk5678*, represented actions that could lead to contract termination.

The *Hazardous Event risk5678* was directly linked to both *Precondition precondition1245* and *Postcondition Termination89*. Here, *precondition1245* specified behaviours (e.g., data interception, spam, unauthorised access) that might trigger termination risks. If *Alice* engaged in these non-compliant actions, *Termination89* would be activated, resulting in contract termination. This structured linkage created a pathway from risky behaviour (*risk5678*) to its consequence (*Termination89*) within the SLA.

Our ontology encouraged the CSP to incorporate preventive and corrective mechanisms, even though these elements were absent from the original SLA template.

To reduce the likelihood of termination, a *prevention mechanism* labelled *communication89* ("Communication") was recommended. This mechanism, associated with both the hazardous event and the precondition, encouraged proactive man-

⁶<https://docs.liqo.io/en/stable/index.html>

⁷<https://www.bit.nl/en/home-en>

agement of minor issues before escalation. In the event of a violation, the “Compliance as Remedy” (complianceAgreement123) corrective action would allow both parties to restore SLA adherence before triggering termination.

5 DISCUSSION

The ontology developed in this study was designed to enhance risk representation and management in SLAs within federated cloud environments. Through our simulation of a digital marketplace, we validated the ontology’s effectiveness in structuring and conveying complex SLA-related risk information to both providers and consumers. By embedding our risk ontology within SLAs, we enabled participants to view potential risks, preventive measures, and remediation strategies explicitly, fostering transparency and informed decision-making in cloud service interactions.

The structural integrity of the ontology was verified using the Hermit 1.4.3.456 reasoner, ensuring logical consistency both with and without instance data. This technical validation confirmed that the ontology operates coherently within the OWL and RDF frameworks, supporting consistent interpretation across diverse cloud environments.

A key outcome of the simulation was the ontology’s ability to guide CSPs in incorporating preventive and corrective mechanisms, even if these were absent in the original SLA templates. By systematically modelling risk triggers and potential outcomes, the ontology encouraged CSPs to adopt proactive communication and compliance-based remedies. This structured risk pathway not only defines actions and consequences within SLAs but also enhances resilience by establishing actionable approaches for risk mitigation.

We deem that the ontology has the potential to foster trust within the federated cloud environments by establishing a shared, transparent structure for risk representation. Through standardised elements like hazardous events, preconditions, preventive mechanisms, postconditions, and remedies, the ontology provides a clear structure that can reduce ambiguity and improve mutual understanding. This transparency is intended to give both consumers and providers confidence in the reliability of the framework, as each party gains visibility into how risks are identified, managed, and mitigated within the SLA itself prior to signing the contract.

Encouraging CSPs to adopt a structured, replicable approach to risk management, the ontology has the potential to create a consistent standard within

SLAs, making contractual terms and risk mitigation strategies more predictable and accessible across cloud federations. This standardised approach could, over time, support trust-building in federated cloud environments by promoting transparency, accountability, and a shared commitment to proactive risk management.

Furthermore, to address international variations in legal requirements, the ontology’s explicit delineation of risk offers an adaptable structure well-suited for federated environments where CSPs and consumers operate under different jurisdictions. Since SLAs and legal frameworks vary widely between countries, the clear representation of risks, preventive measures, and remedies within the ontology can harmonise understanding across regions. By standardising risk elements within the SLA framework, the ontology facilitates consistent communication of terms, helping to bridge differing legal expectations and regulatory requirements.

This explicit risk delineation is particularly advantageous in federated cloud contexts where cross-border collaborations are common. A transparent, well-defined ontology allows each party to understand applicable risks and mitigation strategies, regardless of jurisdictional differences. The ontology’s flexibility and comprehensiveness provide a foundation for CSPs to maintain compliance with local regulations while fostering mutual understanding.

Ultimately, creating a comprehensive framework for risk management in cloud federations is inherently a collaborative endeavour. While this ontology establishes a foundational structure for representing and managing risks within SLAs, we recognise that the dynamic nature of cloud environments demands ongoing adaptation and enrichment. We encourage, therefore, readers of this article to contribute to expanding the ontology’s expressiveness, ensuring that it continues to meet the evolving needs of federated cloud ecosystems. Collaboration toward this shared objective will be essential in building a robust, reliable setting for cloud federation, one where transparency, accountability, and mutual trust are embedded at the core of cloud service interactions.

6 FUTURE WORK

With the development of our ontology, we have established a framework that allows for the explicit representation of risks associated with SLAs in federated cloud environments, providing a structured view of risks, preventive measures, and remediation strategies before digitally signing agreements. This found-

dation facilitates transparency in risk communication and fosters informed decision-making among cloud service providers (CSPs) and consumers.

Moving forward, one key area of focus will be integrating this ontology into an expert system capable of actively tracking the occurrence of preconditions linked to potential hazards. This system could monitor real-time data to detect risk factors or preconditions as they emerge, triggering preventive mechanisms proactively to avoid the manifestation of hazardous events. By embedding such predictive functionality, we aim to shift risk management from a reactive to a more proactive stance, significantly enhancing the resilience of federated cloud services.

In addition, we plan to develop a complementary system for post-event management. Should a hazardous event occur despite preventive measures, this system would activate remediation strategies, mitigating the impact on associated postconditions. This dual-component approach would provide a comprehensive risk management framework, covering both preventive actions and responsive remedies within SLAs.

Beyond these technical advancements, we envision continuous collaboration with industry stakeholders and researchers to enrich the ontology's expressiveness, ensuring it remains a robust, adaptable resource that reflects emerging needs and complexities within the cloud federation domain.

ACKNOWLEDGEMENTS

This research activity was conducted from the Dutch ECOFED project which is part of the integrated IP-CEI CIS project. The early foundation of this research was established during the Gaia-X Tech-X hackathon in 2024, in which our participation was facilitated by the Dutch Gaia-X hub part of the CoE DSC.

REFERENCES

- Alkhamees, S. (2022). Sla negotiation and renegotiation in cloud sla management: Issue and challenges. In Hussain, W. and Jan, M. A., editors, *IoT as a Service*, pages 159–170, Cham. Springer International Publishing.
- Allemang, D., Hendler, J., and Gandon, F. (2020). *Semantic Web for the Working Ontologist: Effective Modeling for Linked Data, RDFS, and OWL*, volume 33. Association for Computing Machinery, New York, NY, USA, 3 edition.
- Aust, J. and Pons, D. J. (2020). A systematic methodology for developing bowtie in risk assessment: Application to borescope inspection. *Aerospace*.
- Bermbach, D., Kurze, T., and Tai, S. (2013). Cloud federation: Effects of federated compute resources on quality of service and cost. In *2013 IEEE International Conference on Cloud Engineering (IC2E)*, pages 31–37.
- Blumauer, A., Nagy, H., and edition mono/monochrom (2020). *The Knowledge Graph Cookbook: Recipes that Work*. edition mono/monochrom.
- Bonfiglio, F. (2021). Gaia-x, vision & strategy. [https://gaia-x.eu/sites/default/files/2021-12/Vision & Strategy.pdf](https://gaia-x.eu/sites/default/files/2021-12/Vision%20&%20Strategy.pdf). CEO Gaia-X, December 16, 2021.
- De Filippi, P. and McCarthy, S. (2012). Cloud computing: Centralization and data sovereignty. *European Journal of Law and Technology*, 3(2).
- Hogan, A. (2020). Web of data. *The Web of Data*.
- Holfelder, W., Mayer, A., and Baumgart, T. (2022). Sovereign cloud technologies for scalable data spaces. In *Designing Data Spaces*, pages 419–436. Springer.
- Iorio, M., Risso, F., Palesandro, A., Camiciotti, L., and Manzalini, A. (2022). Computing without borders: The way towards liquid computing. *IEEE Transactions on Cloud Computing*.
- Karamanlioglu, A. and Alpaslan, F. N. (2018). An ontology-based expert system to detect service level agreement violations. In Shishkov, B., editor, *Business Modeling and Software Design*, pages 362–371, Cham. Springer International Publishing.
- Kbah, Z., Erdil, N. O., and Aqlan, F. (2020). Risk assessment in oil and gas industry using simulation and bowtie analysis. *International Journal of Industrial Engineering: Theory, Applications and Practice*, 27(1).
- Kluwer, W. (2017). The history of bowtie. <https://www.wolterskluwer.com/en/solutions/enablon/bowtie/expert-insights/barrier-based-risk-management-knowledge-base/the-historie-of-bowtie>.
- Kluwer, W. (2019). The bowtie method. <https://www.wolterskluwer.com/en/solutions/enablon/bowtie/expert-insights/barrier-based-risk-management-knowledge-base/the-bowtie-method>.
- Kurze, T., Klems, M., Bermbach, D., Lenk, A., Tai, S., and Kunze, M. (2011). Cloud federation. *The Second International Conference on Cloud Computing, GRIDS, and Virtualization*, pages 32–38.
- Labidi, T., Mtibaa, A., and Brabra, H. (2016). Cslaonto: A comprehensive ontological sla model in cloud computing. *Journal on Data Semantics*, 5.
- Lê, L.-S. and Nguyen, T.-V. (2020). Digitizing service level agreements in service-oriented enterprise architecture. *SN Computer Science*, 1(5):257.
- McLeod, R. W. and Bowie, P. (2020). Guidance on customising bowtie analysis for use in healthcare. In Charles, R. and Golightly, D., editors, *Contemporary Ergonomics and Human Factors 2020*, CIEHF. CIEHF, CIEHF. Based on the CIEHF white paper ‘Human factors in barrier management’.
- Mustafa, E. M., Saad, M. M., and Rizkallah, L. W. (2023). Building an enhanced case-based reasoning and rule-

- based systems for medical diagnosis. *Journal of Engineering and Applied Science*, 70(1):139.
- Omidvar, M., Zarei, E., and Ramavandi, B. (2024). *The Bow-Tie Method: A Hybrid System Safety and Risk Analysis Approach for Safety-Critical Sociotechnical Systems*, pages 123–149. Springer Nature Switzerland, Cham.
- Pan, J. Z., Razniewski, S., Kalo, J.-C., Singhania, S., Chen, J., Dietze, S., Jabeen, H., Omeliyanenko, J., Zhang, W., Lissandrini, M., Biswas, R., de Melo, G., Bonifati, A., Vakaj, E., Dragoni, M., and Graux, D. (2023). Large Language Models and Knowledge Graphs: Opportunities and Challenges. *Transactions on Graph Data and Knowledge*, 1(1):2:1–2:38.
- Prapakorn, N. and Chittayasothorn, S. (2009). An rdf-based distributed expert system. *WSEAS Transactions on Computers*, 8(5):788–798.
- Pruvost, H., Wilde, A., and Enge-Rosenblatt, O. (2023). Ontology-based expert system for automated monitoring of building energy systems. *Journal of Computing in Civil Engineering*, 37.
- Roussey, C., Pinet, F., Kang, M. A., and Corcho, O. (2011). *Ontologies for Interoperability*, pages 39–53. Springer.
- Uriarte, R. B., Tiezzi, F., and Nicola, R. D. (2016). Dynamic slas for clouds. In Villari, M., Zimmermann, W., and Jonas, K., editors, *Service-Oriented and Cloud Computing: 5th European Conference, ESOC 2016, Proceedings*, volume 9846 of *Lecture Notes in Computer Science*, pages 34–49. Springer, Cham. First Online: 25 August 2016.