European International Journal of Multidisciplinary Research and Management Studies (ISSN: 2750-8587)

Received: 01 January 2026 | Accepted: 15 January 2026 | Published: 31 January 2026

RESEARCH ARTICLE

Reliability-Aware Service Level Governance and Error Budget Economics in Large-Scale Cloud and Internet of Things Ecosystems

Julian T. Krause

Technical University of Munich, Germany

VOLUME: Vol.06 Issue 01 2026

PAGE: 146-152

Abstract

The accelerating convergence of large-scale cloud platforms with Internet of Things infrastructures has introduced unprecedented complexity in the governance of service reliability, performance, and contractual accountability. As cloud-native architectures expand toward highly distributed, latency-sensitive, and mission-critical IoT workloads, the traditional interpretation of Service Level Agreements and Quality of Service frameworks becomes increasingly inadequate. This article develops a comprehensive theoretical and methodological foundation for reliability-aware service governance through the lens of Site Reliability Engineering and error budget economics. The research is grounded in the systematic integration of formal SLA specification languages, autonomic cloud management, and QoS-driven orchestration frameworks, while being critically informed by recent advances in SRE practices for error budget management in large-scale systems (Dasari, 2025). By synthesizing decades of SLA theory, distributed systems monitoring, and performance engineering with contemporary reliability-driven operations models, this study constructs a unified conceptual framework that bridges contractual service guarantees and real-time operational risk control.

The central thesis of this article is that error budgets provide a missing economic and operational link between abstract service contracts and the lived reality of cloud and IoT platforms. While SLA languages define what must be delivered, error budgets define how much risk can be taken in delivering it. Through extensive theoretical elaboration, this study demonstrates that reliability engineering, when embedded into SLA-aware orchestration layers, enables a shift from static compliance checking to dynamic risk-adjusted service governance. The methodology employs a multi-layer analytical synthesis combining formal specification theory, cloud and fog computing architectures, and reliability engineering economics. Results are presented as an interpretive mapping of how SLA metrics, QoS indicators, and reliability objectives can be operationalized through continuous monitoring, adaptive control, and release engineering.

KEYWORDS

Site Reliability Engineering, Error Budget Management, Service Level Agreements, Cloud Computing, Internet of Things, Quality of Service, Reliability Governance

INTRODUCTION

The evolution of distributed computing from centralized data centers to globally federated cloud platforms and Internet of Things ecosystems has fundamentally altered the meaning of service reliability, performance, and accountability. In early enterprise computing, service provision was typically governed by tightly controlled environments where infrastructure,

applications, and users were co-located within a single organizational domain. Service Level Agreements emerged in this context as contractual instruments that codified expectations about availability, response time, and fault tolerance, forming the basis for trust between providers and consumers (Sturm et al., 2000). However, as computing architectures have become increasingly virtualized, decentralized, and automated, the adequacy of traditional SLA paradigms has been called into question by both scholars and practitioners (Skene et al., 2010).

The modern cloud computing landscape is characterized by elastic resource provisioning, multi-tenant platforms, and globally distributed services that operate at scales previously unimaginable. At the same time, the proliferation of Internet of Things devices has extended digital services into physical environments, where sensors, actuators, and embedded systems continuously generate and consume data under stringent latency and reliability constraints (Fok et al., 2011). These developments have created a new class of socio-technical systems in which failures are no longer isolated events but emergent properties of complex, interdependent networks (Mahmud et al., 2016). In this context, ensuring service quality is not merely a technical challenge but a governance problem that spans organizational, contractual, and operational domains (Maarouf et al., 2015).

Traditional SLA frameworks, rooted in static specification languages and periodic compliance reporting, struggle to cope with the dynamic and probabilistic nature of cloud and IoT services (Uriarte et al., 2014). Formal SLA languages such as WSLA, GXLA, and domain-specific SLA specifications were developed to express service obligations in machine-readable form, enabling automated monitoring and enforcement (Ludwig et al., 2003; Tebbani and Aib, 2006; Vaderna et al., 2015). While these languages provide a rigorous foundation for specifying service metrics, they often lack a direct connection to the operational realities of large-scale distributed systems, where failures, load fluctuations, and deployment changes are continuous and unavoidable (Okanovic et al., 2013).

It is within this gap between contractual specification and operational reality that Site Reliability Engineering has emerged as a transformative paradigm. Originating in large technology organizations managing planetary-scale infrastructures, SRE reframes reliability not as an absolute requirement but as a quantifiable objective that must be balanced against innovation, cost, and system complexity (Beyer et al., 2016). Central to this paradigm is the concept of the error budget, which defines the allowable level of unreliability a service can experience while still meeting its reliability targets. Rather than striving for perfect uptime, SRE acknowledges that a certain amount of failure is both inevitable and economically rational, provided it remains within agreed limits.

The significance of this shift has been rigorously articulated in recent work on error budget management in large-scale systems, which demonstrates how reliability targets can be translated into actionable operational policies (Dasari, 2025). By formalizing the relationship between service objectives and acceptable risk, error budgets enable engineering teams to make informed decisions about deployments, feature releases, and architectural changes. This approach transforms reliability from a binary compliance metric into a continuous control variable that can be optimized in real time. However, despite its practical success, SRE has largely evolved outside the formal SLA and QoS research traditions, resulting in a conceptual disconnect between contractual service governance and reliability engineering practice (Dasari, 2025).

The literature on SLA specification and QoS management provides a rich theoretical foundation for understanding how service obligations can be modeled, monitored, and enforced. Research on cloud SLAs has produced formal languages that capture availability, performance, and cost metrics in machine-processable formats, enabling automated negotiation and verification (Uriarte et al., 2014; Andrieux et al., 2007). In parallel, IoT-oriented SLA frameworks such as WSN-SLA and utility metrics specifications have sought to address the unique challenges of sensor networks and cyber-physical systems, where energy consumption, data freshness, and network reliability are critical concerns (Gaillard et al., 2014; Calbimonte et al., 2014). Yet these frameworks often remain detached from the operational strategies that determine whether service targets are actually met in practice (Jayaraman et al., 2015).

A growing body of research in performance engineering and autonomic computing has attempted to bridge this gap by developing monitoring frameworks and adaptive control mechanisms that respond to SLA violations in real time (van Hoorn et al., 2012; Walter et al., 2017). SLA-driven workload

management, adaptive monitoring, and declarative performance engineering all aim to create feedback loops between service objectives and system behavior (Stamatakis and Papaemmanouil, 2014; Walter et al., 2016). However, these approaches often focus on technical metrics without explicitly incorporating the economic and organizational dimensions of reliability management that are central to SRE (Dasari, 2025).

The literature gap that motivates this research lies in the absence of an integrated theoretical framework that connects formal SLA specification, QoS management, and SRE-based error budget governance. While SLA languages define what services should deliver and monitoring frameworks measure what they actually deliver, SRE defines how engineering teams should act in the face of uncertainty and failure. Without a unified model, these domains remain siloed, leading to fragmented governance structures in which contractual compliance, technical monitoring, and operational decision-making are misaligned (Uriarte, 2015; Dasari, 2025).

This article addresses this gap by proposing a reliability-aware service governance framework that embeds error budget economics into SLA-driven cloud and IoT architectures. Drawing on a comprehensive synthesis of cloud computing, IoT QoS, SLA specification, and SRE literature, it argues that error budgets can serve as a unifying abstraction that links contractual obligations to operational practices. By treating reliability as a consumable resource governed by explicit budgets, service providers can align deployment velocity, fault tolerance, and customer expectations in a transparent and quantifiable manner (Dasari, 2025; Beyer et al., 2016).

The importance of this integration becomes particularly evident in the context of fog and edge computing, where services must operate under highly variable network conditions and limited computational resources (Mahmud et al., 2016). In such environments, rigid SLA thresholds can lead to either excessive overprovisioning or chronic violations, both of which undermine the sustainability of service ecosystems. Error budget–aware SLA management, by contrast, allows for adaptive trade-offs between performance, availability, and cost, reflecting the probabilistic nature of distributed systems (Duan et al., 2011; Dasari, 2025).

Furthermore, the increasing prevalence of zero-downtime release engineering in large-scale web services highlights the need for reliability-aware deployment strategies that are explicitly linked to service guarantees (Naseer et al., 2020). Continuous delivery pipelines introduce frequent changes that inherently carry risk, yet traditional SLA frameworks offer little guidance on how much risk is acceptable at any given time. Error budgets provide a mechanism for governing this risk, enabling organizations to pace innovation according to their remaining reliability capacity (Dasari, 2025).

In summary, the introduction establishes that the convergence of cloud computing, IoT, and continuous deployment has created a pressing need for new models of service governance that transcend static SLA compliance. By integrating SRE principles and error budget management into the theoretical and operational fabric of SLA and QoS frameworks, this research seeks to redefine how reliability is specified, monitored, and controlled in large-scale digital ecosystems (Dasari, 2025; Uriarte et al., 2014). The following sections develop this argument through a detailed methodological synthesis, interpretive results, and an extensive discussion of theoretical and practical implications.

METHODOLOGY

The methodological foundation of this research is based on an integrative analytical synthesis that combines formal SLA theory, cloud and IoT service management frameworks, and Site Reliability Engineering principles into a unified conceptual model. Rather than employing empirical experimentation or simulation, the methodology is grounded in a systematic and critical engagement with the scholarly and technical literature, reflecting the epistemological orientation of design-oriented and theory-building research in information systems and distributed computing (Skene et al., 2010). This approach is particularly appropriate given the complexity and scale of the systems under study, which defy simple experimental isolation and instead require multi-layered conceptual integration (Uriarte, 2015).

The first methodological pillar involves the structured analysis of SLA specification languages and QoS frameworks. Formal SLA languages such as WSLA, SLAC, GXLA, and domain-specific SLA models provide a rigorous vocabulary for expressing service obligations, measurement points, and violation conditions (Ludwig et al., 2003; Uriarte et al., 2014; Tebbani and Aib, 2006; Vaderna et al., 2015). By examining these languages, the study identifies the underlying assumptions about service stability, metric observability, and enforcement mechanisms that shape how reliability is

conceptualized in contractual terms (Maarouf et al., 2015). This analysis reveals that most SLA languages implicitly treat reliability as a fixed threshold to be met, rather than as a dynamic resource to be managed over time.

The second methodological pillar draws on cloud and IoT service management architectures, including autonomic management frameworks, fog computing models, and QoS orchestration systems. Research on autonomic cloud management emphasizes the need for self-monitoring, self-healing, and self-optimizing capabilities that can respond to changing workloads and failure conditions without human intervention (Uriarte, 2015). In the IoT domain, QoS architectures address the heterogeneity of devices, networks, and applications, seeking to balance competing stakeholder requirements under resource constraints (Fok et al., 2011; Duan et al., 2011). By synthesizing these perspectives, the methodology constructs a layered view of service management in which contractual objectives, monitoring systems, and adaptive controls interact continuously (Jayaraman et al., 2015).

The third and most critical methodological pillar is the incorporation of SRE and error budget theory into this layered model. SRE literature conceptualizes reliability as a probabilistic target expressed through Service Level Objectives and operationalized through error budgets (Beyer et al., 2016). Recent work has further formalized this relationship by demonstrating how error budgets can be used to guide release engineering, incident response, and capacity planning in large-scale systems (Dasari, 2025). Methodologically, this study treats error budgets as a meta-layer that governs how SLA-defined objectives are pursued in practice, providing a bridge between formal specification and operational decision-making.

The integration of these three pillars is achieved through a comparative conceptual mapping. SLA languages are mapped onto SRE constructs to identify correspondences between service metrics and reliability objectives. For example, availability percentages specified in an SLA are translated into error budgets that define allowable downtime over a given period (Dasari, 2025). QoS metrics such as latency and throughput are similarly mapped to reliability targets that reflect user experience rather than raw technical performance (Kritikos et al., 2013). This mapping enables the construction of a unified governance model in which contractual and operational perspectives are mutually reinforcing rather than contradictory.

A key methodological challenge in this synthesis lies in reconciling the deterministic logic of formal SLA specifications with the probabilistic and adaptive nature of SRE. SLA languages are designed to express precise conditions under which violations occur, often triggering penalties or remediation actions (Skene et al., 2010). SRE, by contrast, accepts that violations will occur and focuses on managing their frequency and impact within acceptable bounds (Beyer et al., 2016). The methodology addresses this tension by conceptualizing SLA compliance not as a binary state but as a distribution of outcomes governed by an error budget envelope (Dasari, 2025).

Limitations of this methodological approach include its reliance on secondary sources and conceptual reasoning rather than primary empirical data. While this allows for a broad and theoretically rich analysis, it also means that the proposed framework must be validated in future work through case studies, simulations, or experimental deployments (Walter et al., 2017). Nonetheless, given the objective of developing a foundational theory of reliability-aware service governance, this integrative methodology provides a robust and coherent basis for advancing the field (Uriarte, 2015; Dasari, 2025).

RESULTS

The analytical synthesis undertaken in this study yields a set of conceptual results that redefine how service reliability, QoS, and SLA compliance can be understood and managed in large-scale cloud and IoT ecosystems. These results are not numerical but interpretive, grounded in the alignment of formal service specifications with SRE-driven operational practices (Dasari, 2025). One of the most significant findings is that error budgets can be systematically derived from SLA-defined service levels, transforming abstract contractual commitments into actionable reliability constraints that guide day-to-day engineering decisions.

In traditional SLA frameworks, availability or response time targets are typically expressed as fixed thresholds, such as a service being available 99.9 percent of the time over a month (Skene et al., 2010). The synthesis demonstrates that this target can be directly translated into an error budget that quantifies the allowable downtime in minutes or hours, providing a tangible measure of how much failure the system

can tolerate while remaining compliant (Dasari, 2025). This transformation has profound implications for how services are operated, as it enables teams to budget their risk in the same way they budget financial or computational resources.

Another key result is the identification of a functional alignment between QoS orchestration mechanisms and SRE-based reliability controls. QoS frameworks in cloud and IoT systems dynamically allocate resources, prioritize traffic, and adjust service configurations to meet performance objectives (Jayaraman et al., 2015; Duan et al., 2011). When these mechanisms are informed by error budget consumption, they can prioritize stability when reliability is threatened and enable more aggressive optimization when reliability reserves are high (Dasari, 2025). This creates a feedback loop in which operational controls are directly linked to contractual service guarantees.

The synthesis also reveals that SLA-driven monitoring frameworks, such as those supported by Kieker and other performance engineering tools, can serve as the data foundation for error budget accounting (van Hoorn et al., 2012; Walter et al., 2017). By continuously measuring service metrics and mapping them to reliability objectives, these tools can provide real-time visibility into error budget consumption, enabling proactive intervention before violations occur (Okanovic et al., 2013; Dasari, 2025). This shifts the role of monitoring from retrospective reporting to predictive risk management.

A further result concerns the governance of continuous deployment and zero-downtime release processes. The integration of error budgets into release engineering frameworks allows organizations to modulate their deployment velocity based on their remaining reliability capacity (Naseer et al., 2020). When error budgets are healthy, teams can release new features more frequently, accepting a higher risk of minor failures. When budgets are depleted, release pipelines can be automatically throttled or halted to protect service stability (Dasari, 2025). This mechanism provides a formal and transparent way to balance innovation and reliability, addressing a long-standing tension in software engineering.

Finally, the synthesis indicates that error budget–aware SLA management is particularly well-suited to fog and IoT environments, where variability and uncertainty are inherent (Mahmud et al., 2016; Fok et al., 2011). By replacing rigid compliance thresholds with probabilistic reliability envelopes, service providers can accommodate transient network disruptions, device failures, and workload spikes without triggering disproportionate contractual penalties (Gaillard et al., 2014; Dasari, 2025). This fosters a more resilient and cooperative service ecosystem.

DISCUSSION

The theoretical and practical implications of integrating Site Reliability Engineering and error budget management into SLA-driven cloud and IoT governance are profound, reshaping long-standing assumptions about service quality, contractual accountability, and operational control (Dasari, 2025). At a theoretical level, this integration challenges the dominant paradigm of SLA compliance as a static, binary state, replacing it with a dynamic, probabilistic conception of reliability that better reflects the realities of distributed systems (Skene et al., 2010). This shift aligns with broader trends in performance engineering and autonomic computing, which increasingly emphasize adaptive, feedback-driven control over rigid specification (Walter et al., 2016).

One of the central debates in the SLA and QoS literature concerns the tension between expressiveness and enforceability. Highly expressive SLA languages can capture a wide range of service attributes, but their complexity often makes them difficult to monitor and enforce in practice (Maarouf et al., 2015; Uriarte et al., 2014). Error budgets offer a way to abstract this complexity into a single, actionable metric that summarizes the overall reliability posture of a service (Dasari, 2025). By doing so, they enable both providers and consumers to reason about service quality in a more intuitive and economically meaningful way.

Critics might argue that error budgets oversimplify the multidimensional nature of QoS, reducing diverse performance attributes to a single reliability figure. However, the synthesis presented here demonstrates that error budgets can be defined across multiple SLOs, each corresponding to a different QoS dimension such as latency, throughput, or availability (Beyer et al., 2016; Dasari, 2025). In this sense, error budgets do not replace QoS metrics but rather provide a unifying framework for managing their trade-offs over time.

Another important scholarly debate concerns the role of automation and human judgment in service management. Autonomic computing frameworks aim to automate as many

control decisions as possible, reducing the need for manual intervention (Uriarte, 2015). SRE, while heavily reliant on automation, also emphasizes the importance of human expertise in interpreting reliability data and making strategic trade-offs (Beyer et al., 2016). The integration of error budgets into SLA governance supports a hybrid model in which automated systems enforce reliability constraints while human operators decide how to allocate risk within those constraints (Dasari, 2025).

From a practical perspective, the adoption of error budget–aware SLA management has significant implications for organizational culture and incentive structures. Traditional SLA regimes often create adversarial relationships between providers and consumers, as each side seeks to minimize penalties or maximize compensation (Sturm et al., 2000). By framing reliability as a shared budget that must be jointly managed, error budgets encourage collaboration and transparency, aligning the interests of engineering teams, business stakeholders, and customers (Dasari, 2025; Beyer et al., 2016).

Nevertheless, the framework proposed in this study is not without limitations. One challenge lies in the accurate measurement and attribution of reliability metrics in highly distributed and multi-tenant environments (Kim et al., 2000; Calbimonte et al., 2014). Noise, partial observability, and complex dependency chains can obscure the true sources of failure, complicating error budget accounting (Okanovic et al., 2013). Future research must therefore focus on improving monitoring fidelity and causal analysis in cloud and IoT systems (Walter et al., 2017).

Another limitation concerns the legal and contractual implications of shifting from rigid SLA thresholds to probabilistic error budgets. While error budgets align well with engineering practice, they may require new forms of contract language and regulatory acceptance to be fully effective (Maarouf et al., 2015; Dasari, 2025). Developing standardized frameworks for expressing error budgets in SLA documents is an important area for future work (Uriarte et al., 2014).

Despite these challenges, the integration of SRE and SLA represents a promising path toward more resilient, adaptive, and economically sustainable digital infrastructures. By grounding service governance in both formal specification and operational reality, error budget–aware frameworks can support the continued growth and diversification of cloud and IoT ecosystems (Mahmud et al., 2016; Dasari, 2025).

CONCLUSION

This research has advanced a comprehensive theoretical framework for reliability-aware service governance in large-scale cloud and Internet of Things environments by integrating Site Reliability Engineering and error budget management into the foundations of SLA and QoS theory. Building on formal SLA languages, autonomic cloud management, and contemporary SRE practices, the study demonstrates that error budgets provide a powerful abstraction for linking contractual service guarantees to operational decision-making (Dasari, 2025). By reconceptualizing reliability as a consumable and governable resource, this framework enables more transparent, adaptive, and economically rational management of complex digital services.

The implications of this integration extend beyond technical operations to encompass organizational governance, contractual design, and the future evolution of service ecosystems. As cloud and IoT infrastructures continue to expand in scale and complexity, the need for reliability-aware, risk-sensitive service governance will only become more acute. Error budget–driven SLA management offers a path toward meeting this need, fostering a new generation of resilient and trustworthy digital platforms (Beyer et al., 2016; Dasari, 2025).

REFERENCES

1. Walter, J., Stier, C., Koziolek, H., and Kounev, S. An Expandable Extraction Framework for Architectural Performance Models. Proceedings of the International Workshop on Quality-Aware DevOps, 2017.
2. Maarouf, A., Marzouk, A., and Haqiq, A. A review of SLA specification languages in the cloud computing. Proceedings of the International Conference on Intelligent Systems: Theories and Applications, 2015.
3. Dasari, H. Site Reliability Engineering Practices for Error Budget Management in Large-Scale Systems. International Journal of Applied Mathematics, 38(5s), 991–1001, 2025.
4. Beyer, B., Jones, C., Petoff, J., and Murphy, N. R. Site Reliability Engineering: How Google Runs Production Systems. O'Reilly Media, 2016.
5. Uriarte, R. B. Supporting Autonomic Management of

Clouds: Service-Level-Agreement, Cloud Monitoring and Similarity Learning. IMT School for Advanced Studies Lucca, 2015.

6. Jayaraman, P. P., Mitra, K., Saguna, S., Shah, T., Georgakopoulos, D., and Ranjan, R. Orchestrating quality of service in the cloud of things ecosystem. IEEE International Symposium on Nanoelectronic and Information Systems, 2015.
7. Ludwig, H., Keller, A., Dan, A., King, R. P., and Franck, R. Web Service Level Agreement Language Specification, 2003.
8. Mahmud, R., Kotagiri, R., and Buyya, R. Fog computing: a taxonomy, survey and future directions. Internet of Everything, Springer, 2016.
9. Gaillard, G., Barthel, D., Theoleyre, F., and Valois, F. SLA specification for IoT operation: the WSN-SLA framework, 2014.
10. Okanovic, D., van Hoorn, A., Konjovic, Z., and Vidakovic, M. SLA-driven adaptive monitoring of distributed applications. Computer Science and Information Systems, 10(1), 2013.
11. Skene, J., Raimondi, F., and Emmerich, W. Service-level agreements for electronic services. IEEE Transactions on Software Engineering, 36(2), 2010.
12. Fok, C. L., Julien, C., Roman, G. C., and Lu, C. Challenges of satisfying multiple stakeholders: Quality of service in the internet of things. Workshop on Software Engineering for Sensor Network Applications, 2011.
13. Duan, R., Chen, X., and Xing, T. A QoS architecture for IoT. International Conference on Internet of Things, 2011.
14. Uriarte, R. B., Tiezzi, F., and De Nicola, R. SLAC: A formal service-level-agreement language for cloud computing. IEEE/ACM International Conference on Utility and Cloud Computing, 2014.
15. Kritikos, K., Pernici, B., Plebani, P., Cappiello, C., Comuzzi, M., Benrernou, S., Brandic, I., Kertesz, A., Parkin, M., and Carro, M. A survey on service quality description. ACM Computing Surveys, 46(1), 2013.
16. Naseer, U., Niccolini, L., Pant, U., Frindell, A., Dasineni, R., and Benson, T. A. Zero Downtime Release: Disruption-free Load Balancing of a Multi-Billion User Website. ACM SIGCOMM, 2020.
17. Kim, E. C., Song, J. G., and Hong, C. S. An integrated CNM architecture for multi-layer networks with simple SLA monitoring and reporting mechanism. Network Operations and Management Symposium, 2000.
18. Calbimonte, J. P., Riahi, M., Kefalakis, N., Soldatos, J., and Zaslavsky, A. Utility metrics specifications. OpenIoT Deliverable D4.2.2, 2014.
19. Tebbani, B., and Aib, I. GXLA a language for the specification of service level agreements. IFIP Conference on Autonomic Networking, 2006.
20. Vaderna, R., Vukovic, Z., Okanovic, D., and Dejanovic, I. A domain-specific language for service level agreement specification. International Conference on Information Technology, 2015.
21. van Hoorn, A., Waller, J., and Hasselbring, W. Kieker: A framework for application performance monitoring and dynamic software analysis. ACM/SPEC International Conference on Performance Engineering, 2012.
22. Walter, J., van Hoorn, A., Koziolek, H., Okanovic, D., and Kounev, S. Asking "What?", Automating the "How?": The Vision of Declarative Performance Engineering. ACM/SPEC International Conference on Performance Engineering, 2016.
23. Andrieux, A., Czajkowski, K., Dan, A., et al. Web services for management specification. Open Grid Forum, 2007.
24. Stamatakis, D., and Papaemmanouil, O. SLA-driven workload management for cloud databases. IEEE International Conference on Data Engineering Workshops, 2014.
25. Li, B., and Yu, J. Research and application on the smart home based on component technologies and internet of things. Procedia Engineering, 15, 2011.
26. Practical guide to cloud service agreements version 2.0. Cloud Standards Customer Council, 2015.
27. Bhuyan, B., Sarma, H. K. D., Sarma, N., Kar, A., and Mall, R. Quality of service provisions in wireless sensor networks and related challenges. Wireless Sensor Networks, 2(11), 2010.