

From Consistent Data to Trustworthy Shared Context: AI-Assisted XR for Data-Center Operations

Dawid Lizurek^{1,*†}, Bosen Zhou^{1,*†}, Özlem Mis^{1,†}, Wanja Pasternak^{1,†} and Frank Feinbube^{1,*†}

¹SAP SE

Abstract

A persistent challenge in XR-supported operations is not the availability of digital information, but operators' uncertainty regarding its validity and alignment with the physical environment. Building on our earlier proposal combining AAS, XR and AI-assisted interaction, this position paper examines how shared digital representations can support reliable operational decisions. Six interviews with data-center professionals highlighted discrepancies between digital records and physical conditions that require on-site verification and manual data maintenance. We connect these findings with our AI context-awareness and XR interaction research to propose requirements for communicating temporal validity, provenance, verification status, and conflicting or incomplete evidence. A hypothetical rack-verification scenario explores how operators and AI could recognize conflicting information and request verification before action.

Keywords

Extended Reality (XR), Trust Aware XR, Digital Twin, Asset Administration Shell (AAS), Critical infrastructure, Datacenter, Challenging Environments, Data verification

1. Introduction

Modern datacenters are becoming increasingly complex, requiring a combination of monitoring systems, asset management platforms, digital twins, and operational tools to support infrastructure management and decision-making.

While existing research has investigated uncertainty management, synchronization, and interoperability within digital twins [1, 2, 3] and cyber-physical systems [4, 5], limited attention has been given to how these challenges affect operator trust in XR-assisted workflows [6] within mission-critical datacenter environments. In particular, there remains a lack of guidance regarding how information provenance, temporal validity, verification status, conflicting evidence, and confidence indicators should be represented within XR interfaces to support operational decision-making.

Existing XAI trust frameworks primarily focus on fostering appropriate trust in AI systems through transparency, interpretability, explanation quality, uncertainty communication, and trust calibration while questioning the reliability and credibility of AI model predictions and recommendations [7, 8, 9, 10]. While these approaches help users understand and evaluate AI-generated decisions, they provide limited support for assessing the reliability of the underlying information on which operational decisions depend [8, 9].

The proposed framework therefore complements rather than replaces established XAI trust constructs by introducing reliability-oriented indicators centered on the information ecosystem itself rather than the AI model. Our previous CHI26 workshop paper proposed combining Extended Reality (XR), Asset Administration Shell (AAS) [11] information models and AI assistance for datacenter operations [12]. This paper builds on that proposal by examining how operators and AI can assess conflicting evidence and coordinate verification. [13, 14, 15]

EuroXR 2026 & VWP Annual Forum 2026: 2nd International Workshop on XR for Challenging Environments, November 3, 2026, Valencia, Spain

*Corresponding author.

† These authors contributed equally.

✉ dawid.lizurek@sap.com (D. Lizurek); bosen.zhou@sap.com (B. Zhou); ozlem.mis@sap.com (Ö. Mis);
wanja.pasternak@sap.com (W. Pasternak); frank.feinbube@sap.com (F. Feinbube)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

The six interviews conducted for this work revealed recurring concerns regarding information reliability, validation effort, and trust in digital representations. While exploratory in nature and limited in sample size, these interviews suggest that operator trust may represent a critical but insufficiently studied factor in successful XR adoption within complex operational environments.

This paper therefore addresses this gap by exploring practitioner perspectives from datacenter professionals and arguing for a shift from visualization-centric XR systems toward trust-aware XR environments, where the reliability and contextual information such as provenance, update status, and confidence should be surfaced as prominently as the underlying content. Rather than asking how more information can be displayed to operators, we investigate what information operators require in order to trust the information they are shown.

2. Interview Findings from Datacenter Practitioners

Six semi-structured interviews were conducted with datacenter professionals occupying both operational and supervisory roles. Half of the participants had prior experience with XR technologies, specifically the Microsoft HoloLens to render information, place devices and identify equipment.

As summarized in Table 1, the most frequently reported challenges were monitoring and information quality, device identification, and software integration. Monitoring and information quality emerged as the dominant concern, being raised by all interviewees and encompassing fragmented tools, inconsistent data, and significant effort associated with information gathering and validation.

Table 1
Identified challenges.

Topic	Key Issues	IDs
Monitoring & Information Quality	Fragmented tools, inconsistent data, costly information gathering and validation	1–6
Device Identification	Configuration errors, location tracking, maintenance of up to 50,000 devices, validation effort	1,4,5,6
Software Integration	Vendor software interoperability challenges	1,4,6
Hardware Integration	Integration of heterogeneous vendor hardware	4,5
Communication & Human Error	Coordination issues between onsite and administrative teams; human error incidents	4
Downtime	Service interruptions caused by system changes	3
Capacity Planning	Balancing spare capacity and cost across global datacenters	6
XR Adoption Barriers	Setup overhead, compatibility issues, device availability, platform EOL concerns	2,6

Table 2 indicates that participants primarily prioritized enhanced information access, AI-assisted decision support, and XR-enabled verification and remote support. These themes reflect a common objective of reducing manual effort and improving operational decision-making through timely access to information, intelligent assistance, and context-aware interaction mechanisms. Together, they suggest a vision in which digital technologies augment operators’ ability to access, validate, and act upon infrastructure information efficiently.

Table 2

Future vision and desired capabilities.

Topic	Desired Capability	IDs
Information Access	Reduce reliance on spreadsheets and paperwork, improve visibility, and provide information at hand	1,2,3
Central Information Repository	Provide a single source of truth for operational information	3,5
Information Sharing	Enable consistent information sharing across teams and regions	6
Digital Twin Synchronization	Maintain real-time alignment between physical assets and digital representations	1
Trusted Digital Verification	Reduce onsite verification through trusted digital models	3
Simulation & Planning	Support GPU deployment planning and datacenter retrofit simulations	5,6
AI Decision Support	Assist operators through hardware analysis, recommendations, and efficiency optimization	1,3,6
XR Visualization	Provide contextual overlays on racks, cables, ports, and assets	1,2
XR Collaboration	Reduce communication delays and improve operational efficiency	4,5
XR Verification & Remote Support	Enable AR-based verification, remote expert support, and RFID-assisted inventory management	1,4,6
Infrastructure Monitoring	Continuously monitor datacenter capacity and hardware health	4,5

Overall, participants viewed XR as a promising technology for maintenance support, infrastructure validation, and collaboration. However, successful deployment was consistently linked to reliable, current, and trustworthy information. Across both technical and managerial perspectives, confidence in underlying data was viewed as a prerequisite for practical XR adoption.

3. Toward Trust-Aware XR in Mission-Critical Datacenter Operations

The interviews indicate that information access and validation are a greater challenge than information availability. Operators routinely validate information from monitoring systems, documentation, maintenance records, and digital twins through cross-checking, physical inspections, and consultation with colleagues. Consequently, uncertainty about information accuracy often increases cognitive effort and delays decision-making. For XR systems, simply presenting information closer to physical assets does not automatically improve trust. Operators must also understand where information originates, whether it is current, and whether it has been verified. Similarly, AI should primarily act as an integration layer that aggregates information from fragmented enterprise systems while preserving provenance and verification records.

These findings suggest that future XR systems should support trust assessment alongside information presentation. To address this requirement, we suggest a preliminary framework consisting of five design dimensions derived from our empirical findings and research context:

- *Source Transparency* enables users to identify whether information originates from sensors, documentation systems, asset repositories, digital twins, or AI-generated interpretations.
- *Temporal Validity* communicates when information was last updated, allowing operators to distinguish between current and potentially outdated records.
- *Verification Status* indicates whether information has been confirmed through automated monitoring, procedural checks, human review, maintenance activities, or physical inspection.
- *Conflict Awareness* explicitly highlights inconsistencies between information sources, enabling operators to recognize uncertainty and perform additional validation when necessary.
- *Confidence Indication* provides an evidence-based assessment of information reliability derived from factors such as source credibility, verification history, information age, corroborating evidence, and conflicting records. Confidence should be dynamic, increasing with additional verification and decreasing when information remains unverified over time.

To support practical use, XR systems should present these trust attributes through a layered interface. A simple confidence indicator can provide an immediate overview during routine tasks, while additional interaction layers reveal verification history, information recency, source provenance, detected conflicts, and the evidence underlying confidence assessments. This approach supports both efficient operation and deeper investigation when reliability becomes critical, although its proposed benefits remain subject to future empirical investigation and validation.

4. Illustrative Scenario: Trust-Aware XR Decision Support

In a future datacenter environment, physical assets such as racks, servers, ports, cables, and power equipment can be linked to digital representations through machine-readable identifiers such as QR codes or RFID tags. Using an XR application, operators can immediately access configuration information, maintenance records, provenance data, verification history, and confidence indicators associated with an asset from a centralized AAS-source, reducing manual information retrieval and strengthening the connection between physical and digital system representations.

Consider a scenario in which a remote operator reviewing a digital twin identifies a discrepancy involving a network cable. The inventory system indicates that the cable is connected to Port A1, while a recently uploaded photograph suggests that the cable is actually attached to Port A2. Rather than concealing this inconsistency, the XR system makes it explicit. The inventory record is shown to originate from the AAS and has not been verified for six months, while the photographic evidence is more recent but partially obscured and therefore uncertain. Based on these factors, the system reports a low-confidence assessment.

Instead of producing a definitive recommendation, the AI assistant explains [16] the available evidence and its limitations. It reports that inventory records indicate Port A1, visual evidence suggests Port A2, recent physical verification is unavailable, and additional validation is required. Operators can directly inspect the evidence sources and review the reasoning process that produced the assessment. The purpose of the AI is not to replace human judgment, but to help operators manage large numbers of interconnected assets and identify inconsistencies more efficiently.

To resolve the discrepancy, the system initiates a verification workflow. An onsite technician inspects the connection, confirms that the cable is connected to Port A2, uploads new photographs, updates the documentation, and records supporting test results. As a result, confidence increases because the connection state is now supported by recent physical inspection and additional corroborating evidence. Operators can also examine exactly why the confidence assessment changed, including the age of the original inventory record, the uncertainty associated with the initial photograph, and the contribution of the new verification activities.

Even after establishing that the cable is physically connected to Port A2, a further question remains: whether Port A2 represents the intended design state. Comparison with engineering documentation may reveal that the inventory record was outdated, or alternatively that the cable was connected incorrectly and requires corrective maintenance. Throughout this process, confidence remains a dynamic property that evolves as evidence changes. Confidence decreases as information becomes outdated and increases when new inspections, verifications, and supporting evidence become available.

5. Conclusion and Open Research Questions

This paper explored trust requirements for XR-supported operational tasks through six semi-structured interviews with data center operators. Contrasting our findings with existing trust and confidence indicators proposed in previous XR and AI literature suggests that participants were primarily concerned not with the AI system itself, but with the trustworthiness of the information communicated by the system. In particular, trust was closely associated with the ability to assess the origin, validity, and reliability of underlying data sources. Since AI systems operate on information retrieved from enterprise databases and monitoring infrastructures, the central challenge appears to be establishing confidence

in the data supplied to the AI and in the transparency of its transformation and presentation processes, rather than in the AI component alone.

The findings should also be interpreted within the context of the investigated domain. Data centers represent highly structured and tightly controlled operational environments with the opportunity of standardized equipment, labeling conventions, identification procedures, and the AAS as a centralized datasource. In contrast, transferring the proposed framework to more dynamic and unpredictable environments, such as emergency response or firefighting scenarios, is likely to introduce additional challenges. In such settings, rapidly changing conditions, incomplete information, and limited opportunities for verification may affect both the feasibility and the effectiveness of confidence indicators.

Further limitations must be acknowledged. The empirical basis of this work consists of six interviews and is therefore exploratory rather than representative. Consequently, the five dimensions presented in this paper should be understood as a preliminary conceptual framework rather than a validated model of trust in XR-AI systems. Future work should extend the empirical foundation of the framework through larger and more diverse studies across different operational domains. Furthermore, controlled user studies are required to investigate whether the proposed confidence indicators produce the anticipated benefits in practice, including improvements in data reliability assessment, decision-making efficiency, and operator trust. Such investigations will be essential for validating, refining, and generalizing the framework for trustworthy XR-AI systems.

Ultimately, as XR becomes increasingly integrated into mission-critical environments, the key challenge may no longer be how much information can be displayed, but how effectively systems communicate what is known, what remains uncertain, and why. Addressing this challenge will be essential for enabling trustworthy AI- and XR-assisted workflows in datacenters and other industrial environments.

Acknowledgments

- Acknowledgement: "This research work was carried out as part of the project "Apeiro Reference Architecture", in short ApeiroRA, under the funding ID 13IPC007, financed by the European Union - NextGenerationEU and funded by the German Federal Ministry for Economic Affairs and Energy."
- Disclaimer: "Financed by the European Union - NextGenerationEU. The views and opinions expressed are solely those of the author(s) and do not necessarily reflect the views of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them."

Declaration on Generative AI

During the preparation of this work, the author(s) used Microsoft Copilot solely for language-related improvements such as grammar correction, wording refinement, and stylistic suggestions. After using this tool, the author(s) carefully reviewed, validated, and edited all content and take full responsibility for the final publication.

References

- [1] S. Malihi, Towards uncertainty-aware digital twins in the built environment: A position paper for data-centric engineering, *Digital Engineering* (2026) 100140. doi:doi.org/10.1016/j.dte.2026.100140.
- [2] A. Thelen, X. Zhang, O. Fink, Y. Lu, S. Ghosh, B. D. Youn, M. D. Todd, S. Mahadevan, C. Hu, Z. Hu, A comprehensive review of digital twin – part 2: Roles of uncertainty quantification and optimization, a battery digital twin, and perspectives, 2022. *arXiv:2208.12904*.
- [3] J. Ríos, G. Staudter, M. Weber, R. Anderl, A review, focused on data transfer standards, of the uncertainty representation in the digital twin context, in: C. Fortin, L. Rivest, A. Bernard, A. Bouras (Eds.), *Product Lifecycle Management in the Digital Twin Era*, Springer International Publishing, Cham, 2019, pp. 24–33. doi:10.1007/978-3-030-42250-9_3.
- [4] Y.-P. Chen, Y.-K. Tsai, V. Karkaria, W. Chen, Uncertainty-aware digital twins: Robust model predictive control using time-series deep quantile learning, *Journal of Mechanical Design* 148 (2025). URL: <http://dx.doi.org/10.1115/1.4069104>. doi:10.1115/1.4069104.
- [5] D. C. Villanueva, R. Hickey, A. Markina-Khusid, A. Chaudhuri, Capturing uncertainty in the digital twin state: Addressing data storage and data flow bottlenecks, in: *AIAA SciTech Forum*, 2025. URL: <https://doi.org/10.2514/6.2025-1091>. doi:10.2514/6.2025-1091.
- [6] G. R. Rodrigo Gutierrez-Maquilon, M. Tscheligi, Human-agent trust mediation: Augmented context understanding as a method to dynamically calibrate machine-reliance in emergency response, CHI'26: XR4CE Workshop, April 14, 2026, Barcelona, Spain (2026).
- [7] M. T. Ribeiro, S. Singh, C. Guestrin, Why should i trust you?: Explaining the predictions of any classifier, in: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ACM, 2016, pp. 1135–1144. doi:10.1145/2939672.2939778.
- [8] A. Abdul, J. Vermeulen, D. Wang, B. Y. Lim, M. Kankanhalli, Trends and trajectories for explainable, accountable and intelligible systems: An hci research agenda, *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (2018) 1–18. doi:10.1145/3173574.3174156.
- [9] R. R. Hoffman, S. T. Mueller, G. Klein, J. Litman, Measures for explainable ai: Explanation goodness, user satisfaction, mental models, curiosity, trust, and human-ai performance, *Frontiers in Computer Science* 5 (2023). doi:10.3389/fcomp.2023.1096257.
- [10] C. Newen, D. Bodemer, S. Glantz, E. Müller, M. Wischniewski, L. Schnaubert, Uncertainty awareness and trust in explainable ai: On trust calibration using local and global explanations, *arXiv preprint arXiv:2509.08989* (2025). doi:10.48550/arXiv.2509.08989.
- [11] Industrial Digital Twin Association e.V. (IDTA), Specification of the Asset Administration Shell: Part 1: Metamodel, Industrial Digital Twin Association e.V., 2025. doi:10.62628/IDTA.01001-3-1-2.
- [12] M. M. F. F. Stephan Waltermann, Andre Borrmann, B. Zhou, Xr for complex datacenter environments – why xr needs aas and ai-driven interaction to support operators in next-generation datacenters, in: *CHIXR/2026/04*, 2026.
- [13] A. Mayer, L. Greif, T. M. Häußermann, S. Otto, K. Kastner, S. El Bobbou, J.-R. Chardonnet, J. Reichwald, J. Fleischer, J. Ovtcharova, Digital twins, extended reality, and artificial intelligence in manufacturing reconfiguration: A systematic literature review, *Sustainability* 17 (2025). doi:10.3390/su17052318.
- [14] E. Bozkir, S. Özdel, K. H. C. Lau, M. Wang, H. Gao, E. Kasneci, Embedding large language models into extended reality: Opportunities and challenges for inclusion, engagement, and privacy, in: *ACM Conversational User Interfaces 2024, CUI '24*, ACM, 2024, pp. 1–7. doi:10.1145/3640794.3665563.
- [15] S. Amershi, D. S. Weld, et al., Guidelines for human-ai interaction, in: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*, ACM, New York, NY, USA, 2019, pp. 1–13. doi:10.1145/3290605.3300233.
- [16] F. A. Oliver Hein, Dominik O. W. Hirschberg, P. A. Rauschnabel, Listening between the lines: Exploration of augmented reality supported radio communication for challenging environments, in: *CHIXR/2026/04*, 2026.