Ref. Ares(2026)981055 - 28/01/2026

Musical Metaverse made in Europe: An innovation lab for musicians and audiences of the future

D3.2: Report on security and privacy issues and planned measures

European Innovation Council
Funded by the European Union
101184379.

This project has been funded by the European Union as part of the Horizon Europe Framework Program (HORIZON), under the grant agreement

This work was supported by the Swiss State Secretariat for Education, Research and Innovation (SERI) under contract number 25.00026.

Schweizerische Eidgenossenschaft
Confédération suisse
State Secretariat for Education, Research and Innovation SERI

Document Information

Document factsheet

Work Package	WP 3 Musical Networking
Task	T3.2: Report on security & privacy issues and planned measures.
Due date	31/01/2026
Deliverable lead	7Bulls.com
Authors	Katarzyna Materka, Michał Kłosiński
Reviewers	Luca Turchet
Nature of deliverable	R
Dissemination level	PU

Document history

Version	Date	Description
v1.0	26/11/2025	Draft creation
v1.1	12/01/2026	First revision
v1.2	24/01/2026	Second revision
v1.3	28/01/2026	Document finalized

Disclaimer

This project has received funding from The European Union as part of the Horizon Europe Framework Program (HORIZON) under grant agreement No. 101184379. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Innovation Council. Neither the European Union nor the European Innovation Council can be held responsible for them.

This work was also supported by the Swiss State Secretariat for Education, Research and Innovation (SERI) under contract number 25.00026.

Table of contents

LIST OF ACRONYMS **1**

EXECUTIVE SUMMARY **2**

1. INTRODUCTION **3**

1.1 Background: Security & Privacy in the Musical Metaverse 3
1.2 Objectives of the document 3
1.3 Methodology and Data Sources 4
1.3.1 Methodological Approach *4*
1.3.2 Data Sources *5*
1.4 Structure of the Document 5

2. SECURITY- AND PRIVACY-RELEVANT CHARACTERISTICS OF THE MUSMET ECOSYSTEM **6**

2.1 Threat Model Specific to Musical Metaverse 6
2.1.1 Assets and Security Objectives *6*
2.1.2 Adversaries and Threat Actors *7*
2.1.3 Attack Surfaces and System Layers *7*
2.1.4 Constraints and Assumptions *7*
2.1.5 Role of the Threat Model in MUSMET *8*
2.2 Data Types and Sensitivity Levels 8
2.2.1 Real-Time Musical and Audiovisual Data *8*
2.2.2 Expressive Interaction and Motion Data *8*
2.2.3 Biometric and Neurophysiological Data *8*
2.2.4 User Identity, Profile, and Session Metadata *9*
2.2.5 User-Generated Content and Intellectual Property Data *9*
2.2.6 Derived and Inferred Data *9*
2.2.7 Implications for Security and Privacy Design *9*
2.3 System Constraints 10
2.3.1 Ultra-Low-Latency and Tight Synchronization *10*
2.3.2 Continuous, High-Frequency Data Streams *10*
2.3.3 Heterogeneous Devices and Network Technologies *10*
2.3.4 Distributed and Decentralized Architecture *11*
2.3.5 Dynamic Session Management and Participation *11*
2.3.6 Performance-Security Trade-offs *11*
2.3.7 Implications for Security and Privacy Design *11*
2.4 Intellectual Property Rights (IPR) Considerations in MUSMET 11
2.4.1 Nature of Intellectual Assets in the Musical Metaverse *12*
2.4.2 IPR Risks in Real-Time Musical Interaction *12*
2.4.3 Challenges Arising from Collaborative Creation *12*
2.4.4 System-Level Implications for IPR Protection *13*
2.4.5 Relationship Between IPR and Security Measures *13*
2.4.6 Role of IPR Considerations in This Deliverable *13*
2.5 Legal and Ethical Aspects 13
2.5.1 Protection of Personal and Sensitive Data *14*
2.5.2 Consent, Awareness, and User Control *14*
2.5.3 Use of Biometric and Neurophysiological Data *14*
2.5.4 Fairness, Inclusion, and User Safety *14*
2.5.5 Accountability and Trust *14*
2.5.6 Implications for MUSMET Design and This Deliverable *15*

3. RESULTS OF THE MUSMET CONSORTIUM SURVEY **16**

3.1 Methodology of the Survey 16
3.1.1 Framework and Scope *16*
3.1.2 Data Collection and Partner Input *16*
3.2 Reported Perceived Threats 17

3.2.1 Security Threats (STRIDE Perspective) *17*
3.2.2 Privacy Threats (LINDDUN Perspective) *18*
3.3 Risk Prioritization by Stakeholders 18
3.4 Identified Gaps and Recommendations 18

4. SECURITY AND PRIVACY ISSUES IN THE MUSICAL METAVERSE 20

4.1 Network- and Transport-Layer Threats 20
4.1.1 Eavesdropping and Traffic Interception *20*
4.1.2 Man-in-the-Middle and Injection Attacks *20*
4.1.3 Packet Loss, Reordering, and Timing Manipulation *20*
4.1.4 Denial-of-Service and Resource Exhaustion *21*
4.1.5 Session Hijacking and Transport-Level Impersonation *21*
4.1.6 Exposure of Metadata and Traffic Patterns *21*
4.1.7 Implications for the Musical Metaverse *21*
4.2 Application-Layer Threats 22
4.2.1 Avatar Impersonation and Identity Misuse *22*
4.2.2 Unauthorized Access to Virtual Spaces and Sessions *22*
4.2.3 Manipulation of Virtual Environments and Content *22*
4.2.4 Unauthorized Recording, Replication, and Redistribution *23*
4.2.5 Injection of Malicious or Deceptive Content *23*
4.2.6 Session State Manipulation and Logic Abuse *23*
4.2.7 Privacy Leakage Through Application Behaviour *23*
4.2.8 Implications for the Musical Metaverse *24*
4.3 Data and AI/ML Threats 24
4.3.1 Data Exposure and Leakage Across the Processing Pipeline *24*
4.3.2 Inference and Re-Identification from Expressive Data *24*
4.3.3 Neurophysiological and EEG-Specific Threats *25*
4.3.4 Threats in AI/ML Model Training and Use *25*
4.3.5 Centralised vs Distributed Learning Risks *25*
4.3.6 Secondary Use and Long-Term Accumulation Risks *26*
4.3.7 Opacity, Explainability, and Trust *26*
4.3.8 Implications for the Musical Metaverse *26*
4.4 Device- and Sensor-Layer Threats 26
4.4.1 Device Compromise and Unauthorized Control *27*
4.4.2 Sensor Spoofing and Signal Manipulation *27*
4.4.3 Threats Specific to Wearable and Neurophysiological Sensors *27*
4.4.4 Insecure Communication Between Devices and Processing Nodes *27*
4.4.5 Physical Access and Tampering *28*
4.4.6 Heterogeneity and Update Challenges *28*
4.4.7 Side-Channel and Environmental Attacks *28*
4.4.8 Implications for the Musical Metaverse *28*
4.5 IPR and Digital Asset Protection Threats 29
4.5.1 Unauthorized Recording and Capture of Performances *29*
4.5.2 Uncontrolled Replication and Redistribution of Digital Assets *29*
4.5.3 Content Manipulation and Integrity Violations *30*
4.5.4 Impersonation and Fraudulent Attribution *30*
4.5.5 Misuse of Collaborative Creation Mechanisms *30*
4.5.6 Persistence and Reuse of Ephemeral Content *30*
4.5.7 Cross-Platform and Interoperability Risks *31*
4.5.8 Implications for the Musical Metaverse *31*
4.6 User Safety, Harassment, and Social Risks 31
4.6.1 Harassment and Abusive Behaviour in Immersive Spaces *31*
4.6.2 Identity Misrepresentation and Social Engineering *32*
4.6.3 Unwanted Presence and Invasion of Personal Space *32*
4.6.4 Exposure to Disturbing or Manipulative Content *32*
4.6.5 Social Exclusion and Power Imbalances *32*
4.6.6 Psychological and Cognitive Risks *33*
4.6.7 Challenges in Detection and Response *33*
4.6.8 Implications for the Musical Metaverse *33*

5. SUITABILITY ANALYSIS OF EXISTING SECURITY AND PRIVACY PROTOCOLS UNDER MM LATENCY CONSTRAINTS 34

5.1 REQUIREMENTS FOR SECURE ULTRA-LOW-LATENCY COMMUNICATION 34
5.1.1 Latency-Zone Awareness and Differentiation 34
5.1.2 Preservation of Ultra-Low Latency for Musical Interaction 34
5.1.3 Deterministic Timing and Jitter Control 35
5.1.4 Compatibility with Continuous Real-Time Streaming 35
5.1.5 Efficient Session Establishment and Dynamic Participation 35
5.1.6 Robustness Under Network Variability 35
5.1.7 Protection of Payload and Communication Metadata 35
5.1.8 Support for Heterogeneous Devices and Architectures 35
5.1.9 Scalability for Multi-Party and Audience Scenarios 36
5.1.10 Fault Isolation and Graceful Recovery 36
5.2 EVALUATION OF TRANSPORT-LAYER SECURITY PROTOCOLS 36
5.2.1 Evaluation Scope and Methodological Perspective 36
5.2.2 Analysis of Secure Communication and Stream Protection Approaches 37
5.2.3 Comparative Perspective Across Latency Zones 39
5.2.4 Summary 39
5.3 EVALUATION OF NETWORK-LAYER SECURITY PROTOCOLS 39
5.3.1 Role of Network-Layer Security in MUSMET 40
5.3.2 Analysis of Network-Layer Security Approaches 40
5.3.3 Network-Layer Security vs Transport-Layer Security 41
5.3.4 Appropriate Deployment Scenarios in MUSMET 41
5.3.5 Summary 41
5.4 PRIVACY-PRESERVING ARCHITECTURES 42
5.4.1 Data Minimisation and Purpose Limitation by Design 42
5.4.2 Local and Edge-Centric Processing Architectures 42
5.4.3 Distributed and Federated Architectural Patterns 42
5.4.4 Separation of Interaction, Control, and Analytics Planes 43
5.4.5 Architectural Support for Anonymity and Pseudonymity 43
5.4.6 Limitations and Trade-Offs 43
5.5 SUMMARY OF PROTOCOL SUITABILITY FOR MUSMET 43

6. PROPOSED SECURITY & PRIVACY MEASURES FOR MUSMET 45

6.1 POSITIONING OF SECURITY AND PRIVACY WITHIN THE MUSMET PROJECT SCOPE 45
6.2 SECURITY AND PRIVACY ISSUES ADDRESSED THROUGH EXISTING STANDARDS AND TOOLS 46
6.3 MUSMET-SPECIFIC RISKS IDENTIFIED BUT CONSIDERED OUT OF PROJECT SCOPE 46
6.4 SECURITY AND PRIVACY ASPECTS EXPLICITLY ADDRESSED WITHIN MUSMET 47
6.5 SECURITY-BY-DESIGN PRINCIPLES APPLIED IN MUSMET SYSTEM ARCHITECTURE 48
6.6 VALIDATION, MONITORING, AND RISK AWARENESS DURING PROJECT EXECUTION 49
6.7 SUMMARY AND RELATION TO POST-PROJECT EXPLOITATION 50

REFERENCES 51

List of acronyms

AI	Artificial Intelligence
AR	Augmented Reality
BCI	Brain-Computer Interface
DTLS	Datagram Transport Layer Security
GDPR	General Data Protection Regulation
HMD	Head-Mounted Display
IPR	Intellectual Property Rights
ML	Machine Learning
MM	Musical Metaverse
MR	Mixed Reality
NMP	Networked Music Performance
SDN	Software-defined networking
TCP	Transmission Control Protocol
UDP	User Datagram Protocol
UX	User Experience
VR	Virtual Reality
XR	Extended Reality

Executive Summary

This deliverable examines the security and privacy challenges of the MUSMET Musical Metaverse and defines planned measures to address them. It focuses on risks emerging from real-time, immersive, and distributed musical interactions, where ultra-low-latency communication, multimodal data processing, and heterogeneous devices are core requirements. The scope covers technical, legal, and ethical aspects, providing a shared baseline to support the secure and compliant design of MUSMET technologies in line with European regulatory frameworks and project objectives.

The analysis indicates that the Musical Metaverse introduces heightened and domain-specific risks compared to conventional XR or multimedia platforms. Key issues include the protection of real-time audio and audiovisual streams, risks of inference and re-identification from expressive and biometric data, avatar impersonation and unauthorized access, and the misuse or redistribution of user-generated musical content and intellectual property. These risks are amplified by strict latency constraints, continuous high-frequency data flows, and dynamic participation, a concern consistently reflected in the consortium survey results.

To mitigate these challenges, MUSMET follows a security-by-design and privacy-by-design approach. Planned measures include secure low-latency communication mechanisms, strong identity and access control, safeguards for intellectual property, and minimization and controlled processing of personal and biometric data. Where AI/ML components are involved, privacy-preserving techniques and appropriate governance mechanisms are foreseen. Overall, the proposed measures aim to ensure trust, data protection, and regulatory compliance while preserving the performance and responsiveness required for real-time musical collaboration.

1. Introduction

1.1 Background: Security & Privacy in the Musical Metaverse

The emergence of the Metaverse marks a significant shift in the way individuals interact, collaborate, and create within shared virtual environments. Within this broader landscape, the Musical Metaverse (MM) represents a domain with particularly stringent requirements for real-time interaction, expressive communication, and secure exchange of multimodal information (Boem et al., 2025; Wang et al., 2023). As highlighted in recent technological analyses, musical activities - especially collaborative performance - are characterized by extreme sensitivity to latency, jitter, synchronization accuracy, and data integrity, far exceeding those of typical social XR applications (Boem et al., 2025).

In such settings, delays above 30 ms can substantially impair the ability of geographically distributed musicians to perform together, while inconsistent audiovisual or haptic cues can disrupt audience immersion and compromise artistic expression.

The MUSMET project seeks to advance the state of the art by developing an integrated ecosystem of devices, networking techniques, and software services enabling real-time, high-fidelity, multisensory musical interactions in XR environments. Achieving this vision requires not only technical innovation, but also a robust and comprehensive approach to security and privacy, given the unprecedented scale and sensitivity of data involved.

Musical Metaverse interactions rely on continuous streams of audio, video, motion-tracking signals, haptic stimulation, environment sensing, and, in some scenarios, biometric or neurophysiological signals, such as EEG or affective state indicators. Such data can reveal deeply personal information, including identity traits, behavioural patterns, emotional states, or health characteristics. At the same time, MM environments enable creation, manipulation, and distribution of original artistic content, raising important concerns about intellectual property rights (IPR), integrity of creative works, and protection against unauthorized recording or reproduction.

Furthermore, the distributed, multi-vendor, and highly heterogeneous nature of emerging Metaverse infrastructures introduces new vulnerability surfaces: insecure XR devices, compromised IoT sensor nodes, untrusted network paths, malicious avatars, and misuse of real-time data channels (Ali et al., 2024;Wang et al., 2023). Traditional security protocols, often optimized for throughput, reliability, or enterprise scenarios, are not automatically suitable for ultra-low-latency audio-first environments. As a result, the MM domain demands a deep evaluation of existing and emerging security and privacy technologies, including transport-layer encryption, authentication methods, data minimization techniques, and privacy-preserving machine learning, to determine their feasibility under MUSMET's real-time constraints.

1.2 Objectives of the document

The main purpose of this deliverable is twofold. First, to identify and analyse both the security and privacy threats that are common to generic Metaverse environments and those that are specific to the MM. Second, to define the measures that can be undertaken within MUSMET to ensure secure and trustworthy system operation. More specifically, the objectives are:

- To establish the security- and privacy-relevant context of the MUSMET ecosystem, including functional characteristics, latency constraints, data types, user interactions, and IPR considerations that influence risk exposure and protection needs.
- To synthesize threat categories specific to the MM, spanning network-level attacks, device vulnerabilities, identity and access risks, manipulation of XR environments, biometric data misuse, AI/ML-related threats, and infringement of artistic intellectual property (Kang et al., 2024).

- To assess the suitability of existing security and privacy protocols - including transport-layer encryption (TLS/DTLS/SRTP), secure media frameworks, QUIC-based approaches, authentication mechanisms, and privacy-preserving data processing methods - under MUSMET's ultra-low-latency and high-fidelity performance requirements.
- To integrate and analyse insights from the MUSMET consortium, capturing partners' perceptions of domain-specific risks, priorities, and unmet needs, thereby grounding the threat model in real stakeholder experience.
- To propose a coherent set of planned security and privacy measures, covering architecture design principles, technology selection, data governance, IPR protection strategies, user safety practices, and integration with MUSMET's orchestration and networking layers.

1.3 Methodology and Data Sources

This deliverable adopts a **multi-source, qualitative methodology** combining state-of-the-art analysis, project-specific contextual investigation, and stakeholder-driven input to identify relevant security and privacy threats and to define appropriate mitigation measures for the Musical Metaverse (Huang et al., 2023, Rahartomo et al., 2025).

1.3.1 Methodological Approach

The methodology followed in D3.2 consists of four complementary steps (Wang et al., 2023; Singh et al., 2025; Dey et al., 2025):

1. **Systematic analysis of existing research and standards:** A structured review of scientific literature, technical reports, and international standards was conducted to identify security and privacy threats commonly associated with Metaverse, XR, IoT, and distributed multimedia systems. Particular attention was paid to works addressing real-time interaction, low-latency communication, biometric data protection, user safety, and privacy-preserving data processing. This step enabled the identification of *baseline threats and countermeasures* applicable to generic Metaverse environments.
2. **Contextual analysis of MUSMET-specific characteristics and requirements:** The general threat landscape was refined by analysing the specific characteristics of the MUSMET ecosystem, including ultra-low-latency musical interaction, continuous multimodal data streams, expressive performance data, XR-based collaboration, and the use of heterogeneous devices and networks. This step allowed the identification of Musical-Metaverse-specific threats that arise due to tight timing constraints, audio-first interaction paradigms, artistic content creation, and the handling of potentially sensitive biometric or neurophysiological signals.
3. **Stakeholder-driven risk identification through consortium input:** To complement the technical analysis, a structured survey was conducted among MUSMET consortium partners to collect their perspectives on perceived security, privacy, and intellectual property risks. The survey results were used to validate assumptions derived from the literature, highlight domain-specific concerns, and prioritise threats based on practical relevance to MUSMET use cases.
4. **Derivation of feasible mitigation measures:** Based on the identified threats and contextual constraints, the deliverable defines a set of practical security and privacy measures that can realistically be implemented within the scope and timeline of the MUSMET project. These measures are evaluated with respect to their compatibility with MUSMET's performance requirements, regulatory obligations, and architectural design principles.

1.3.2 Data Sources

The analysis presented in this deliverable draws upon the following data sources:

- Scientific literature, including surveys, systematic reviews, and peer-reviewed articles on Metaverse security and privacy, XR systems, real-time multimedia networking, biometric data protection, and privacy-preserving machine learning.
- Technical standards and guidelines, such as those related to secure communication protocols, identity and access management, and data protection frameworks relevant to distributed and real-time systems.
- MUSMET project documentation, including architectural concepts, use-case definitions, and performance requirements that influence security and privacy design choices.
- Consortium survey data, capturing expert input from MUSMET partners regarding anticipated risks, domain-specific threats, and expectations for security and privacy safeguards.

By combining these sources, D3.2 ensures that its analysis is both grounded in established research and closely aligned with the practical realities and objectives of the MUSMET project, providing a solid foundation for subsequent design, implementation, and validation activities.

1.4 Structure of the Document

This document is structured as follows. Section 1 introduces the background and objectives of the deliverable, outlines the adopted methodology, and situates security and privacy within the context of the Musical Metaverse. Section 2 describes the security- and privacy-relevant characteristics of the MUSMET ecosystem, including the threat model, data types and sensitivity levels, system constraints, intellectual property considerations, and legal and ethical aspects. Section 3 presents the results of the MUSMET consortium survey, summarising perceived threats, risk prioritisation, and identified gaps. Section 4 provides a detailed analysis of security and privacy issues across network, application, data and AI/ML, and device layers. Section 5 outlines the planned security and privacy measures and mitigation strategies aligned with the identified risks. Finally, Section 6 concludes the document by summarising the main findings and highlighting their implications for subsequent MUSMET design and implementation activities.

2. Security- and Privacy-Relevant Characteristics of the MUSMET Ecosystem

2.1 Threat Model Specific to Musical Metaverse

The Musical Metaverse involves multiple stakeholder groups that generate, process, and interact with data in distinct ways. These stakeholders include composers and performers engaged in musical creation and live collaboration, audience members participating in immersive listening and experiential scenarios, as well as teachers and students involved in educational and training activities. Each group contributes different categories of data, ranging from creative content and real-time performance signals to interaction traces, learning-related information, and audience engagement metrics. Clearly identifying these stakeholder roles is essential for understanding data flows, exposure to threats, and the scope of security and privacy risks addressed by the MUSMET threat model (Wang et al., 2023; Rahartomo et al., 2025).

The MM exhibits a threat landscape that is both broader and more nuanced than that of generic Metaverse environments, due to its audio-first interaction paradigm, stringent real-time constraints, rich multimodal data flows, and creative content lifecycle (Chen et al., 2022). This section describes the threat model adopted in MUSMET, providing a structured yet flexible foundation for analysing security and privacy risks and for identifying appropriate mitigation strategies.

The threat model builds upon well-established security and privacy taxonomies (STRIDE, LINDDUN) commonly used in distributed systems, multimedia networking, and XR research - such as confidentiality, integrity, availability, authenticity, and privacy-related inference risks ((Howard & Lipner, 2006; Deng et al., 2011). However, rather than strictly adhering to a single formal taxonomy, the MUSMET threat model adopts a context-driven perspective, allowing threats to be analysed in terms of their practical impact on musical collaboration, user experience, and artistic integrity.

2.1.1 Assets and Security Objectives

At the core of the MM threat model are the assets whose protection is critical to safe and trustworthy operation. These include:

- Live musical content, such as high-fidelity audio streams, spatialized sound fields, and synchronized audiovisual representations of performances, whose integrity and timing are essential for ensemble coordination.
- Expressive and interaction data, including motion capture, gesture, posture, and timing information, which form an integral part of musical expression and may reveal behavioural patterns or performance signatures.
- User-related data, encompassing identities, profiles, session metadata, and, in some MUSMET scenarios, biometric or neurophysiological signals that can expose highly sensitive personal information.
- Intellectual property assets, such as original compositions, live performances, virtual instruments, digital stages, and user-generated content created within XR environments.
- System resources and services, whose availability and correct operation are prerequisites for uninterrupted, low-latency musical interaction.

The corresponding security objectives include maintaining confidentiality where required, preserving data and system integrity, ensuring continuous availability under real-time constraints, supporting reliable authentication and authorization, and preventing unauthorized inference or misuse of sensitive data.

2.1.2 Adversaries and Threat Actors

The MUSMET threat model considers a diverse set of potential threat actors with different motivations and capabilities:

- External attackers, aiming to intercept, manipulate, or disrupt real-time communication streams, compromise devices, or exploit vulnerabilities in network protocols and services.
- Malicious or misbehaving participants, including users who may attempt impersonation, unauthorized recording or redistribution of performances, content manipulation, or deliberate disruption of collaborative sessions.
- Insider or semi-trusted entities, such as compromised services, misconfigured infrastructure components, or third-party systems integrated into the MM ecosystem.
- Curious or opportunistic actors, whose goal is not necessarily disruption but inference-extracting sensitive information (e.g., identity traits, emotional states, performance styles) from legitimately accessible data streams.

These actors are assumed to operate in a partially trusted environment, where neither users nor infrastructure components can be universally assumed to be fully trustworthy (Sharma et al., 2024).

2.1.3 Attack Surfaces and System Layers

Threats in the MM emerge across multiple, tightly coupled layers:

- Network and transport layers, including real-time audio and XR data channels, synchronization mechanisms, and edge-to-cloud communication paths, where latency, packet manipulation, and traffic analysis are critical concerns.
- Application and XR layers, encompassing shared virtual spaces, avatar interactions, performance tools, and collaborative environments, which may be exploited for impersonation, content spoofing, or behavioural manipulation.
- Device and sensor layers, including XR headsets, musical IoT devices, motion tracking systems, and wearable or neurophysiological sensors, which introduce risks related to device compromise, spoofing, and data leakage.
- Data processing and AI/ML components, where aggregation, analytics, or learning processes can enable privacy violations, model manipulation, or unintended inference.

The real-time coupling of these layers amplifies the potential impact of attacks, as even short-lived disruptions or subtle manipulations can significantly affect musical coordination and user experience (Otoum et al., 2024).

2.1.4 Constraints and Assumptions

The MUSMET threat model explicitly incorporates several domain-specific constraints:

- Ultra-low-latency and tight synchronization requirements, which limit the complexity and overhead of security mechanisms that can be deployed in the critical interaction path.
- Heterogeneous and distributed infrastructure, involving diverse devices, network technologies, and service providers.
- Continuous and high-frequency data generation, increasing the risk of cumulative privacy leakage and long-term profiling.
- Regulatory and ethical obligations, including compliance with data protection regulations (e.g., GDPR) and emerging governance frameworks for AI and biometric data.

2.1.5 Role of the Threat Model in MUSMET

This threat model serves as a guiding framework for the remainder of this deliverable, enabling the differentiation between security and privacy threats common to generic Metaverse platforms and those specific to the MM, the evaluation of the suitability of existing security and privacy protocols under MM constraints, and the identification of feasible, project-relevant mitigation measures that can be progressively implemented within MUSMET. By adopting a descriptive, context-aware threat model informed by established taxonomies but tailored to musical interaction, MUSMET ensures that security and privacy considerations remain closely aligned with technical feasibility and artistic requirements.

2.2 Data Types and Sensitivity Levels

The MM relies on the continuous generation, transmission, and processing of heterogeneous data streams to enable real-time, immersive musical interaction. These data streams differ significantly in terms of content, purpose, temporal characteristics, and sensitivity, and therefore require differentiated security and privacy protections. This section classifies the main data types handled within the MUSMET ecosystem and discusses their associated sensitivity levels and protection implications.

Rather than applying a rigid data classification scheme, the analysis adopts a context-aware perspective, informed by established data protection principles and regulatory frameworks, to highlight how data sensitivity in the MM is shaped by real-time use, expressive content, and potential for inference (Wang et al., 2023; Ali et al., 2024).

2.2.1 Real-Time Musical and Audiovisual Data

At the core of MUSMET are high-fidelity audio streams, often accompanied by synchronized video or spatial rendering data, which carry live musical performances, rehearsals, and collaborative creative activities. While such data may not always constitute personal data in a strict legal sense, they are highly sensitive from both functional and intellectual property perspectives, as unauthorized access, manipulation, or delay can directly degrade musical quality and disrupt ensemble coordination, interception or recording can lead to infringement of performers' rights and unauthorized redistribution of artistic content, and subtle modifications, such as timing shifts or audio artefacts, may not be immediately detectable yet can significantly affect artistic outcomes (Wang et al., 2023). As a result, musical and audiovisual data require strong integrity and confidentiality protections implemented in a manner compatible with ultra-low-latency constraints.

2.2.2 Expressive Interaction and Motion Data

Musical collaboration in XR environments relies heavily on expressive interaction data, including motion capture, gesture trajectories, posture, spatial positioning, and timing cues, which are essential for conveying musical intent, coordination, and non-verbal communication. From a privacy perspective, such data are particularly sensitive because they can reveal behavioural patterns, motor characteristics, or performance signatures unique to an individual, enable long-term profiling of musicians' expressive styles, and support re-identification when combined with other data sources. Although often treated as technical interaction data, expressive motion data may therefore qualify as highly sensitive personal data, warranting careful access control, minimization, and protection against unauthorized inference.

2.2.3 Biometric and Neurophysiological Data

In some MUSMET scenarios, interaction may involve biometric or neurophysiological signals, such as EEG, heart rate, or affective indicators, which represent some of the most sensitive data categories

handled within the MM. Such data can reveal health-related information, emotional or cognitive states, as well as fatigue, stress, or engagement levels during performance. The processing of biometric and neurophysiological data therefore introduces heightened privacy, ethical, and regulatory considerations, requiring strict safeguards including purpose limitation, secure storage and transmission, and compliance with applicable data protection regulations. Even when processed transiently or locally, the potential for misuse or unintended inference necessitates a conservative protection approach.

2.2.4 User Identity, Profile, and Session Metadata

The MUSMET ecosystem also handles identity-related and contextual metadata, including user identifiers, authentication credentials, access rights, session logs, device identifiers, and network-related information. While often considered ancillary, such metadata can be highly sensitive when aggregated over time, as it enables the tracking of user behaviour and participation patterns, correlation across sessions or platforms, and the inference of social relationships and professional activity. Protecting identity and session metadata is therefore essential to prevent impersonation, unauthorized access, profiling, and linkage attacks.

2.2.5 User-Generated Content and Intellectual Property Data

The Musical Metaverse enables the creation of user-generated content, such as compositions, recordings, virtual instruments, digital stages, and collaborative arrangements, which represent significant intellectual and economic value. The sensitivity of such data stems from ownership and authorship considerations, the risks of unauthorized copying, modification, or redistribution, and potential disputes over rights in collaborative creation contexts. Security mechanisms must therefore address not only confidentiality and integrity but also provenance, authenticity, and controlled access, while avoiding intrusive protections that could hinder creative workflows.

2.2.6 Derived and Inferred Data

Beyond raw data streams, MUSMET systems may generate derived or inferred data, such as performance analytics, behavioural models, synchronization metrics, or AI-generated representations. Although often treated as secondary, such data may implicitly encode sensitive information and can in some cases be more revealing than the original inputs, depending on the nature of the inference, the possibility of linking results back to individuals, and the duration and scope of data retention (Huang et al., 2023). Appropriate safeguards are therefore required to prevent unintended disclosure or misuse of inferred information.

2.2.7 Implications for Security and Privacy Design

The diversity of data types and sensitivity levels in the MM highlights the need for differentiated and proportionate protection mechanisms. Security and privacy measures in MUSMET must balance the need for strong protection of sensitive personal and artistic data, the strict performance requirements of real-time musical interaction, and the principles of data minimization, purpose limitation, and user control (Otoum et al., 2024). This classification informs subsequent sections of the deliverable, in particular the analysis of threats, the evaluation of suitable protocols, and the definition of measures to be implemented within the MUSMET project.

2.3 System Constraints

The MUSMET ecosystem operates under a set of stringent technical and architectural constraints that directly influence the feasibility, design, and placement of security and privacy mechanisms. These constraints arise from the real-time nature of musical interaction, the continuous processing of multimodal data, and the distributed, heterogeneous structure of the MM. This section elaborates on the main system constraints that shape security and privacy decisions in MUSMET.

2.3.1 Ultra-Low-Latency and Tight Synchronization

A fundamental constraint of the MMis the requirement for ultra-low end-to-end latency and precise temporal synchronization across all interaction modalities, as distributed musical performance is highly sensitive to timing deviations and even small delays or jitter can disrupt rhythmic alignment, degrade ensemble coherence, and negatively affect expressive performance. From a system perspective, latency constraints apply simultaneously to audio signal transmission and processing, synchronization between audio, visual, and gestural cues, and feedback loops between performers and virtual environments. These requirements significantly restrict the class of security mechanisms that can be deployed in the critical data path, since multi-round protocol handshakes, frequent key renegotiation, or cryptographic operations with unpredictable execution times may introduce delays or variability that are unacceptable in musical contexts (Otoum et al., 2024). Consequently, security mechanisms must be deterministic, lightweight, and tightly integrated with real-time processing pipelines.

2.3.2 Continuous, High-Frequency Data Streams

MUSMET systems rely on continuous, high-frequency data streams, including raw audio samples, motion capture data, spatial positioning information, and sensor readings, which are produced and consumed continuously, often at fixed sampling rates, rather than following traditional request–response patterns. This streaming paradigm constrains security and privacy approaches in several ways, as there is limited opportunity for buffering, batching, or deferred processing, inline processing must be completed within strict time bounds, and interruptions or delays in data flow have immediate perceptual consequences (Wang et al., 2023). As a result, security mechanisms that depend on extensive inspection, accumulation of data over time, or offline validation are difficult to apply, and protection strategies must instead operate incrementally and transparently while preserving the continuity and temporal structure of data streams.

2.3.3 Heterogeneous Devices and Network Technologies

The MUSMET ecosystem integrates a wide variety of devices and communication technologies, ranging from high-performance edge and cloud servers to resource-constrained XR headsets, wearable sensors, and musical IoT components, which differ substantially in computational capacity, memory availability, energy constraints, and supported networking protocols. This heterogeneity introduces system-level constraints, as not all devices can support the same cryptographic algorithms or protection mechanisms, security capabilities may vary across system components leading to uneven protection, and upgrading or patching devices may be constrained by hardware limitations or deployment conditions. Therefore, security and privacy mechanisms must be adaptable and modular, allowing protections to be tailored to device capabilities while maintaining coherent end-to-end guarantees where possible.

2.3.4 Distributed and Decentralized Architecture

MUSMET adopts a distributed architecture in which computation and data processing are spread across multiple edge and cloud nodes, with interactions occurring across different network domains. While this architectural choice is essential for scalability and responsiveness, it imposes additional constraints on security design, as there is no single point with complete visibility or control over all interactions, trust relationships must be established dynamically across components, and the coordination of security policies across administrative boundaries is complex. These factors limit the effectiveness of centralized enforcement mechanisms and require security solutions that function reliably in partially decentralized and federated environments, often relying on local decision-making and peer coordination.

2.3.5 Dynamic Session Management and Participation

MM scenarios often involve dynamic session composition, with participants joining, leaving, or reconfiguring groups during runtime, and with session duration, group size, and interaction patterns changing rapidly in response to artistic or collaborative needs. This dynamism constrains session establishment and teardown mechanisms, key distribution and renewal strategies, and access control models that assume stable membership. Security mechanisms must therefore support fast onboarding, low-overhead reconfiguration, and seamless adaptation to changing session contexts without interrupting ongoing musical interaction.

2.3.6 Performance-Security Trade-offs

Across all layers of the MUSMET system, there exists an inherent trade-off between performance and security, as the strict requirements on latency, jitter, and computational overhead mean that the strongest or most comprehensive security mechanisms are not always feasible in time-critical paths. System design must therefore consider the selective application of security controls based on data criticality, the prioritization of protections for the most sensitive or impactful data streams, and the placement of heavier security mechanisms outside real-time interaction loops where possible. Rather than seeking maximal protection in all cases, MUSMET must pursue balanced, risk-aware security designs that preserve musical performance while providing meaningful protection against realistic threats.

2.3.7 Implications for Security and Privacy Design

Taken together, these system constraints highlight that security and privacy in the MM must be performance-aware, adaptive, and context-sensitive, as conventional security solutions designed for transactional or delay-tolerant systems are often unsuitable without modification. The constraints identified in this section provide essential input for evaluating the suitability of existing security and privacy protocols, designing MUSMET-specific protection strategies, and selecting mitigation measures that can be realistically implemented within the project lifecycle. They also frame the analysis presented in subsequent sections, ensuring that proposed solutions remain aligned with the operational realities of real-time musical interaction.

2.4 Intellectual Property Rights (IPR) Considerations in MUSMET

Intellectual Property Rights (IPR) constitute a central consideration in the MM, as MUSMET systems enable the creation, transformation, and real-time exchange of original artistic content within immersive and collaborative XR environments. Unlike many generic Metaverse applications, musical

interaction inherently involves copyrighted material, performance rights, and creative authorship, making IPR protection a critical aspect of secure and trustworthy system operation.

In MUSMET, IPR considerations intersect closely with security and privacy requirements, as unauthorized access, interception, manipulation, or replication of content directly translate into potential IPR violations. This section outlines the main IPR-related characteristics and risks specific to the MUSMET context and their implications for system design.

2.4.1 Nature of Intellectual Assets in the Musical Metaverse

The MUSMET ecosystem handles a diverse set of intellectual assets, including:

- Live musical performances, transmitted as high-fidelity audio streams and spatial representation;
- Original compositions and arrangements, created or modified collaboratively in XR environments;
- User-generated content, such as virtual instruments, sound designs, digital stages, and performance artefacts;
- Expressive performance data, including gesture and timing information, which may themselves constitute creative expression;
- Derived digital artefacts, such as recordings, remixes, or AI-generated representations of musical content.

These assets may be created individually or collaboratively, often in real time, complicating traditional notions of authorship, ownership, and control.

2.4.2 IPR Risks in Real-Time Musical Interaction

The real-time and distributed nature of MUSMET introduces IPR-related risks that differ in both form and severity from those encountered in conventional digital media platforms. Key risks include:

- Unauthorized recording or capture of live performances, potentially without the knowledge or consent of performers.
- Interception and redistribution of musical content during transmission, particularly in shared or public XR environments.
- Content manipulation or misattribution, including subtle alterations that affect artistic integrity while remaining difficult to detect.
- Impersonation or identity misuse, leading to fraudulent creation or attribution of musical works.
- Uncontrolled replication of user-generated assets, facilitated by the ease of digital duplication in immersive environments.

Such risks are amplified by the tight coupling between interaction, performance, and content creation in the MM.

2.4.3 Challenges Arising from Collaborative Creation

MUSMET supports collaborative and co-creative workflows in which multiple participants jointly contribute to musical output, which is a core strength of the platform but also introduces technical challenges related to IPR protection. These challenges include distinguishing individual contributions within a shared performance, tracking the provenance and evolution of creative artefacts, and managing access and modification rights dynamically during live sessions. Without appropriate technical safeguards, such collaborative environments may blur ownership boundaries and increase the likelihood of disputes or misuse.

2.4.4 System-Level Implications for IPR Protection

From a system perspective, IPR considerations impose several requirements on MUSMET architectures and security mechanisms:

- Controlled access to content, ensuring that only authorized participants can receive, modify, or reproduce musical assets;
- Protection of content integrity, preventing undetected modification or degradation of artistic works;
- Support for provenance and traceability, enabling attribution and tracking of creative contributions over time;
- Compatibility with real-time constraints, as IPR protection mechanisms must not introduce latency or disrupt musical interaction.

Importantly, IPR protection mechanisms must be designed to operate in harmony with performance requirements, avoiding intrusive controls that hinder creativity or collaboration.

2.4.5 Relationship Between IPR and Security Measures

In MUSMET, IPR protection cannot be treated as a separate concern from security and privacy. Many technical safeguards-such as encryption, authentication, secure streaming, and access control-serve a dual role in protecting both personal data and intellectual assets. Conversely, weaknesses in security mechanisms often translate directly into IPR vulnerabilities (Huang et al., 2023).

For this reason, IPR considerations are integrated into the broader threat analysis and mitigation strategy presented in this deliverable. Rather than relying solely on external enforcement mechanisms, MUSMET emphasises technical protections embedded within the system design, aligned with the platform's real-time and collaborative nature.

2.4.6 Role of IPR Considerations in This Deliverable

Within D3.2, IPR considerations inform the identification of Musical-Metaverse-specific threats, the evaluation of suitable security protocols and architectures, and the definition of measures that can be realistically implemented during the MUSMET project. By explicitly addressing IPR at the system level, MUSMET aims to support creative freedom and collaboration while safeguarding the rights and interests of content creators participating in the Musical Metaverse (Rahartomo et al., 2025).

2.5 Legal and Ethical Aspects

The MM raises a set of legal and ethical considerations that are closely intertwined with security and privacy risks, yet extend beyond purely technical constraints. MUSMET systems support immersive, real-time, and expressive interaction, involving sensitive personal data, creative content, and collaborative social environments. As a result, legal and ethical aspects must be considered as design drivers that inform how security and privacy measures are defined and implemented throughout the project.

Rather than providing a comprehensive legal analysis, this section outlines the key legal and ethical dimensions relevant to MUSMET and highlights their implications for system design and risk mitigation.

2.5.1 Protection of Personal and Sensitive Data

MUSMET environments process multiple categories of personal data, including identity information, behavioural patterns, interaction data, and, in some scenarios, biometric or neurophysiological signals, which from a legal and ethical perspective require heightened protection due to their potential impact on individual autonomy, dignity, and well-being. Ethical considerations include respecting user expectations regarding the use and visibility of their data, avoiding unnecessary data collection or retention, and preventing secondary use of data beyond the original interaction context. These considerations reinforce the need for data minimization, purpose limitation, and transparency, which must be reflected in the design of security and privacy mechanisms.

2.5.2 Consent, Awareness, and User Control

Immersive environments can obscure traditional cues that signal data collection or monitoring, meaning that in the MM users may not always be aware of which data are being captured, processed, or shared in real time (Rahartomo et al., 2025). From an ethical standpoint, MUSMET systems should therefore support meaningful user awareness of data flows, mechanisms for expressing and managing consent, and user control over participation, visibility, and data sharing where feasible. Security and privacy measures must accordingly be designed to enable user agency, rather than merely enforcing protection in a passive or opaque manner.

2.5.3 Use of Biometric and Neurophysiological Data

Where MUSMET scenarios involve biometric or neurophysiological data, additional ethical concerns arise due to the intimate nature of such signals, as they can reveal information about health, cognitive state, or emotional responses, making their misuse particularly harmful (Wang et al., 2023). Ethical considerations therefore include avoiding intrusive or excessive sensing, ensuring proportionality between the data collected and the intended functionality, and preventing the inference of sensitive attributes that are not strictly required for musical interaction. These aspects necessitate conservative handling strategies and reinforce the importance of privacy-preserving processing approaches.

2.5.4 Fairness, Inclusion, and User Safety

The MM is a social and collaborative space, and ethical design must therefore consider not only data protection but also fairness, inclusivity, and user safety. Relevant concerns include preventing harassment, impersonation, or exclusion within virtual musical spaces, avoiding system designs that disadvantage certain user groups due to accessibility, device limitations, or interaction styles, and ensuring that security mechanisms do not disproportionately burden or exclude participants (Sharma et al., 2024). Ethical considerations thus extend to how security and privacy measures affect social dynamics and the overall user experience.

2.5.5 Accountability and Trust

Trust in the MM depends on the ability to assign responsibility and understand system behaviour when incidents occur. From both a legal and ethical perspective, MUSMET systems should support a degree of accountability, including the traceability of actions within collaborative environments, clarity regarding the roles of different system components, and mechanisms to address misuse or harmful behaviour. While MUSMET does not aim to enforce legal accountability mechanisms directly, security and privacy design choices can either support or hinder accountability and trust (Huang et al., 2023).

2.5.6 Implications for MUSMET Design and This Deliverable

Legal and ethical aspects influence MUSMET primarily by shaping requirements and constraints on acceptable system behaviour, informing the classification of data sensitivity, the prioritization of threats and mitigation measures, and the evaluation of suitable security and privacy protocols (Singh et al., 2025). Within D3.2, these considerations provide a normative backdrop against which technical solutions are assessed. By acknowledging legal and ethical dimensions alongside technical constraints, MUSMET aims to promote a MM that is not only secure and privacy-preserving, but also respectful of users, creators, and collaborative practices.

3. Results of the MUSMET Consortium Survey

This section presents the results of the MUSMET Consortium Security & Privacy Anticipation Survey and provides a stakeholder-driven perspective on the security, privacy, and intellectual property risks anticipated within the Musical Metaverse ecosystem. The analysis synthesizes responses from 14 distinct respondents representing 13 consortium organisations, including 8 industrial partners and 5 research organisations, and aligns them with the methodological framework introduced earlier in this deliverable. By doing so, it grounds the theoretical threat model in the practical realities, design constraints, and expectations expressed by participating members of the MUSMET consortium.

The MUSMET ecosystem is characterized by a complex interplay of high-fidelity, ultra-low-latency audio streams, immersive XR environments, heterogeneous networking infrastructures, and - in some scenarios - highly sensitive biometric and neurophysiological sensors. These characteristics significantly amplify the potential impact of security and privacy incidents. Consequently, capturing partner anticipation of risks is essential for prioritizing threats and defining feasible mitigation measures within the scope of the project.

3.1 Methodology of the Survey

To complement the literature review and architectural analysis, a consortium-wide survey was conducted to collect forward-looking insights into anticipated security and privacy challenges. The survey was explicitly designed to elicit expectations, concerns, and design trade-offs rather than to document already implemented safeguards.

3.1.1 Framework and Scope

The survey adopted a structured qualitative methodology based on two widely recognised frameworks:

- **STRIDE** (Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, Elevation of Privilege) to systematically identify anticipated *security threats (Howard & Lipner, 2006)*.
- **LINDDUN** (Linkability, Identifiability, Non-repudiation, Detectability, Disclosure of Information, Unawareness, Non-compliance) to explore *privacy risks* arising from immersive, sensor-rich environments (Deng et al., 2011).

This dual-framework approach ensured balanced coverage of technical, organisational, and user-centric risks, while remaining compatible with MUSMET's distributed and heterogeneous architecture.

3.1.2 Data Collection and Partner Input

Partners were first asked to identify the types of data their components are expected to collect, transmit, or process. Reported data categories included:

- real-time audio and audiovisual streams;
- 3D spatial and interaction data;
- identifiable user and avatar information;
- network and session metadata;
- biometric and neurophysiological signals, such as EEG or affective-state indicators.

Subsequently, partners reflected on how their respective components-such as the XR Music Box, Head-Mounted Displays (HMDs), Brain-Computer Interfaces (BCIs), AI/ML modules, and 5G/edge/cloud infrastructures-might be exposed to STRIDE and LINDDUN threat categories. Particular attention was paid to trade-offs between cryptographic protection, computational overhead, and ultra-low-latency performance, which are central to MMapplications.

The collected responses thus provide a context-aware, MUSMET-specific risk landscape, complementing generic Metaverse security analyses.

3.2 Reported Perceived Threats

The survey responses reveal a diverse but convergent set of concerns, reflecting both the technical fragility of real-time musical interaction and the ethical sensitivity of the data involved. Reported threats can be grouped into security-related risks (STRIDE perspective) and privacy-related risks (LINDDUN perspective).

3.2.1 Security Threats (STRIDE Perspective)

From a security standpoint, partners consistently highlighted risks that could directly compromise musical integrity, system trustworthiness, and collaborative safety.

Identity Spoofing (Spoofing). A significant concern was raised regarding the impersonation of legitimate users, avatars, devices, or infrastructure components. Examples include fraudulent musician avatars entering private rehearsal spaces or malicious entities masquerading as trusted 5G nodes or XR devices. Such attacks could lead to unauthorized access, reputational damage, fraudulent attribution of creative works, and erosion of trust within collaborative musical environments.

Data Tampering (Tampering). The integrity of real-time data streams-particularly audio, control signals, and EEG data-was identified as paramount. Partners emphasized that even subtle manipulations, such as micro-timing shifts or minor signal distortions, could disrupt synchronization, degrade artistic expression, or bias adaptive system behaviour. In MUSMET's ultra-low-latency context, these attacks may be difficult to distinguish from benign network anomalies, increasing their practical risk.

System Availability and Disruption (Denial of Service). Real-time streaming pipelines, orchestration services, and 5G connectivity were identified as highly vulnerable to resource exhaustion or denial-of-service attacks. Unlike delay-tolerant applications, MMservices cannot gracefully degrade: even short-lived disruptions may irreversibly break live performances or rehearsals. Availability was therefore perceived as a **core security objective**, tightly coupled with user experience.

In addition to the security threats most frequently emphasized by consortium partners, several STRIDE-related risks were mentioned only sporadically or implicitly and are therefore less represented in the aggregated analysis. In particular, repudiation and elevation of privilege threats were not consistently highlighted across partners' inputs, they are perceived as secondary to real-time performance and data protection concerns. Repudiation-related risks, such as disputes over responsibility for disruptive actions or misuse of shared environments, were occasionally referenced but did not emerge as a dominant concern, possibly reflecting an assumption that logging and governance mechanisms will be addressed at later stages of system integration. Similarly, privilege escalation risks, including unauthorized access to administrative functions or cross-component control interfaces, were acknowledged by a limited subset of respondents, typically in relation to platform management or orchestration layers rather than end-user interactions. While these threats were less frequently raised, they remain relevant in a distributed, multi-role ecosystem such as MUSMET and warrant consideration in the overall security design, particularly as system complexity and interoperability increase.

3.2.2 Privacy Threats (LINDDUN Perspective)

Privacy-related concerns were particularly pronounced due to the granularity and intimacy of data processed in MUSMET environments.

Linkability and Identifiability. Partners expressed concern that expressive performance data, motion patterns, or neurophysiological signals could function as quasi-biometric identifiers. Even when pseudonymisation is applied, such data may allow re-identification of individuals across sessions or platforms, enabling long-term profiling of musicians' styles, behaviours, or cognitive states.

Unawareness and Information Disclosure. Immersive XR environments may obscure traditional cues indicating data collection. Partners noted a risk that users-both performers and audiences-may be unaware of the extent to which emotional, cognitive, or physiological data are being captured and processed, particularly in scenarios involving BCIs or affective sensing.

Regulatory Non-compliance. Emerging regulatory frameworks, notably the GDPR and the EU AI Act, were cited as major challenges. Neurophysiological data were consistently recognised as falling into highly sensitive categories, requiring explicit consent, strict purpose limitation, and robust safeguards. Partners highlighted the risk that insufficiently transparent or poorly governed data flows could lead to non-compliance, with legal and ethical consequences.

3.3 Risk Prioritization by Stakeholders

Based on partner input, risks were implicitly prioritised according to their anticipated impact on musical integrity, user trust, and regulatory exposure. Table 1 summarises the resulting prioritisation.

Priority	Risk Category	Key Concern
High	Neurophysiological data leakage	Irreversible privacy harm due to exposure of EEG or affective signals
High	Real-time stream disruption	Violation of the 30 ms latency threshold critical for musical interaction
Medium	IPR infringement	Unauthorized recording or redistribution of collaborative performances or collaborative composition sessions
Medium	Avatar impersonation	Social engineering and fraudulent attribution in virtual studios

Table 1. Main risks identified by stakeholders.

This prioritisation underscores that privacy breaches involving biometric data and availability attacks affecting real-time performance are perceived as the most critical threats for MUSMET, surpassing even traditional concerns such as repudiation or auditability.

3.4 Identified Gaps and Recommendations

Beyond identifying threats, the survey revealed several systemic gaps that inform the definition of planned security and privacy measures in subsequent sections of this deliverable.

Performance-Aware Security Design. Partners repeatedly stressed the need for security mechanisms that are explicitly designed for *audio-first, ultra-low-latency interaction*. Conventional cryptographic solutions may introduce jitter or processing delays that are unacceptable in musical contexts, motivating the exploration of lightweight, performance-aware protection schemes.

Integrated Identity and Zero-Trust Principles. The lack of a unified identity and access management model tailored to distributed MMenvironments was identified as a key gap. Partners recommended adopting Zero-Trust principles and identity federation to manage roles and permissions across heterogeneous devices, platforms, and administrative domains.

Coordinated Governance Across Partners. Given strong interdependencies between components developed by different partners, respondents emphasised the need for *joint security controls and shared assumptions*, for example between networking, cloud orchestration, and ethical governance activities. Fragmented approaches were seen as a potential source of systemic vulnerability.

User Transparency and Awareness Mechanisms. To mitigate risks related to unawareness, partners recommended explicit transparency mechanisms, such as clear indicators when sensitive sensors (e.g., BCIs) are active and transmitting data, as well as user-accessible controls for consent and participation.

In summary, the MUSMET consortium survey confirms that security and privacy are foundational enablers of the MM rather than auxiliary concerns. The results validate and refine the threat model presented earlier in this deliverable, while directly informing the design principles and concrete measures proposed in the following sections of D3.2.

4. Security and Privacy Issues in the Musical Metaverse

4.1 Network- and Transport-Layer Threats

Network- and transport-layer threats represent a critical risk category in the MM, as MUSMET relies on continuous, low-latency, high-fidelity data exchange to support real-time musical interaction. Unlike delay-tolerant applications, MM services are highly sensitive to even minor disruptions, making network- and transport-layer vulnerabilities particularly impactful (Ali et al., 2024; Boem et al., 2025). This section describes the main classes of threats affecting these layers and their implications for MUSMET.

4.1.1 Eavesdropping and Traffic Interception

The transmission of real-time audio streams, motion data, and synchronization signals across distributed networks exposes MUSMET systems to the risk of eavesdropping and traffic interception. Attackers positioned on untrusted or shared network paths may capture data packets without necessarily altering them. In the MMcontext, such interception can lead to:

- unauthorized access to live musical performances;
- leakage of expressive or behavioural data;
- inference of participation patterns or collaboration structures.

Even when intercepted data are not intelligible in isolation, traffic analysis techniques may still reveal sensitive information based on timing, volume, or communication patterns.

4.1.2 Man-in-the-Middle and Injection Attacks

Network- and transport-layer vulnerabilities may allow adversaries to perform man-in-the-middle attacks, enabling them to intercept, modify, or inject packets into ongoing communication streams. In real-time musical interaction, even subtle packet manipulation can have disproportionate effects. Potential consequences include:

- degradation of audio quality or timing precision,
- disruption of synchronization between participants,
- injection of malicious or misleading control signals.

Such attacks may remain difficult to detect, as their effects can resemble benign network anomalies rather than overt failures.

4.1.3 Packet Loss, Reordering, and Timing Manipulation

MUSMET systems are highly sensitive to packet loss, reordering, and timing variation, whether caused by malicious activity or exploited through intentional manipulation. Adversaries may deliberately introduce delays or reorder packets to disrupt musical coherence without fully interrupting communication. These threats are particularly relevant because:

- small timing perturbations can impair rhythmic alignment;
- effects may accumulate over time, degrading performance quality;

- distinguishing between malicious interference and normal network congestion can be challenging.

As a result, timing manipulation represents a subtle yet effective attack vector in MM environments.

4.1.4 Denial-of-Service and Resource Exhaustion

Network-level denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks pose a significant threat to MUSMET, given its reliance on continuous connectivity and real-time processing. Attackers may attempt to overwhelm network links, transport endpoints, or session management mechanisms. In the MM, the impact of such attacks extends beyond service unavailability:

- interrupted sessions may irreversibly disrupt live performances;
- repeated disruptions may erode user trust and willingness to participate;
- resource exhaustion at transport endpoints may cascade to other system components.

Even short-lived attacks can have severe consequences due to the time-critical nature of musical interaction.

4.1.5 Session Hijacking and Transport-Level Impersonation

Transport-layer vulnerabilities can enable session hijacking or impersonation, allowing attackers to take control of an existing communication session or masquerade as a legitimate participant. In MUSMET, this can have particularly harmful effects. Potential impacts include:

- unauthorized participation in private musical sessions;
- disruption or manipulation of performances;
- misuse of established trust relationships between participants.

Because MM sessions often involve dynamic group composition, session hijacking may go unnoticed for extended periods.

4.1.6 Exposure of Metadata and Traffic Patterns

Even when payloads are protected, network- and transport-layer communication may expose metadata, such as packet timing, size, frequency, and endpoint identifiers. In MUSMET, such metadata can reveal sensitive contextual information. Examples include:

- inference of rehearsal schedules or performance duration;
- identification of prominent performers or collaborators;
- reconstruction of interaction intensity and group dynamics.

Metadata exposure therefore represents a privacy-relevant threat, particularly when combined with other data sources.

4.1.7 Implications for the Musical Metaverse

Network- and transport-layer threats are especially critical in MUSMET due to the tight coupling between communication quality and musical performance (Ali et al., 2024). Attacks at these layers may not need to fully disrupt connectivity to be effective; minor degradations or subtle manipulations can suffice to undermine musical interaction and user experience.

These threats highlight the need for careful selection and configuration of transport protocols, robust session management, and protection strategies that account for both data content and communication patterns, while remaining compatible with MUSMET's stringent real-time constraints.

4.2 Application-Layer Threats

Application-layer threats arise from vulnerabilities and misuse of the software components that enable interaction, collaboration, and content creation in the MM. In MUSMET, the application layer encompasses XR environments, avatar representations, collaborative musical tools, session logic, and content management functions. Because these components directly mediate user experience and artistic interaction, threats at this layer can have immediate and highly perceptible effects on both system operation and creative outcomes (Otoum et al., 2024).

4.2.1 Avatar Impersonation and Identity Misuse

In immersive environments, avatars serve as the primary representation of users. Application-layer weaknesses may enable avatar impersonation, allowing attackers to assume the identity of legitimate participants or to create misleading representations. Such threats can lead to:

- unauthorized participation in musical sessions;
- reputational harm to performers or creators;
- fraudulent attribution of musical actions or content.

Impersonation at the application layer can be particularly damaging in collaborative musical contexts, where trust between participants is essential for effective interaction.

4.2.2 Unauthorized Access to Virtual Spaces and Sessions

MUSMET applications manage access to virtual stages, rehearsal rooms, and collaborative environments. Weaknesses in application-layer access control can permit unauthorized entry into restricted spaces. Potential consequences include:

- exposure of private rehearsals or creative processes;
- disruption of live performances;
- leakage of sensitive interaction or performance data.

Given the dynamic nature of MM sessions, ensuring consistent and context-aware access control is especially challenging.

4.2.3 Manipulation of Virtual Environments and Content

Application-layer threats may involve the manipulation of virtual environments, musical tools, or shared content. Attackers may alter parameters, inject malicious content, or subtly modify the virtual setting in ways that affect musical interaction. Examples include:

- altering spatial audio parameters or instrument settings;
- injecting misleading visual or auditory cues;
- modifying shared scores or performance artefacts.

Such manipulations may not always be immediately detected, yet they can significantly degrade artistic integrity and collaborative effectiveness.

4.2.4 Unauthorized Recording, Replication, and Redistribution

At the application layer, MUSMET systems may expose functionality for capturing, exporting, or sharing content. If insufficiently controlled, these mechanisms can be exploited for unauthorized recording or replication of live performances and user-generated content. This threat is particularly relevant in the MM, where:

- live interaction often coincides with content creation;
- recording capabilities may be embedded in XR applications;
- collaborative settings blur the boundary between private and public performance.

Application-layer vulnerabilities may therefore directly enable IPR violations and loss of control over creative output (Zhao et al., 2023).

4.2.5 Injection of Malicious or Deceptive Content

Attackers may exploit application-layer interfaces to inject malicious or deceptive content, such as altered audio streams, misleading visual elements, or corrupted interaction data. Such injections can:

- mislead participants about the state of the performance;
- disrupt synchronization and coordination;
- degrade trust in the virtual environment.

Because these attacks target application-level semantics rather than raw network traffic, they may bypass lower-layer protections.

4.2.6 Session State Manipulation and Logic Abuse

MM applications maintain complex session state, including participant roles, synchronization status, and shared resources. Application-layer vulnerabilities may allow attackers to manipulate this state or abuse application logic. Examples include:

- forcing premature session termination;
- altering participant permissions during a session;
- exploiting race conditions or inconsistencies in session management.

Such attacks can disrupt performances and undermine confidence in system reliability.

4.2.7 Privacy Leakage Through Application Behaviour

Beyond explicit data exposure, application-layer behaviour may unintentionally leak sensitive information. For example:

- interface elements may reveal participation patterns or user roles;
- application responses may vary in ways that enable inference;
- logging or debugging features may expose sensitive data.

In the MM, such leakage can reveal creative processes, collaboration structures, or individual performance characteristics.

4.2.8 Implications for the Musical Metaverse

Application-layer threats are particularly impactful in MUSMET because they directly affect user interaction, artistic expression, and collaborative trust. Unlike lower-layer attacks, which may manifest as technical anomalies, application-layer attacks often exploit domain semantics and may therefore be harder to detect and diagnose. These threats underscore the need for robust application-level security design, including secure session management, careful control of recording and sharing functions, and consistent enforcement of access and identity policies. Their identification provides essential input for the mitigation strategies and architectural measures discussed in later sections of this deliverable.

4.3 Data and AI/ML Threats

Data-driven processing and AI/ML-based components play an increasingly important role in the MM, supporting functionalities such as real-time synchronization, adaptive interaction, performance analysis, expressive augmentation, and user experience optimisation. While these techniques enable advanced musical interaction, they also introduce a distinct and complex class of data-centric and learning-related threats. In MUSMET, such threats are amplified by the continuous, expressive, and highly personal nature of musical interaction data, as well as by the real-time constraints under which processing must occur. This section analyses the principal threats associated with data handling and AI/ML usage in the MM, with particular attention to high-sensitivity data types, including biometric and neurophysiological signals such as EEG.

4.3.1 Data Exposure and Leakage Across the Processing Pipeline

MUSMET systems rely on distributed data pipelines spanning end devices, edge nodes, and cloud services. At each stage, data may be exposed to unauthorized access, leakage, or unintended disclosure due to misconfiguration, compromised components, or insufficient isolation between services. Data exposure risks affect:

- raw input streams (e.g. audio, motion, EEG);
- intermediate representations generated during preprocessing;
- stored datasets used for analysis or training;
- outputs produced by AI/ML components.

Because MM data streams are continuous and high-resolution, even brief exposure events can result in significant leakage, with cumulative effects over time. Leakage of expressive or performance-related data may also enable imitation, profiling, or misuse of artistic output.

4.3.2 Inference and Re-Identification from Expressive Data

A critical threat category in MUSMET arises from inference attacks, whereby sensitive information is derived from data that may appear innocuous in isolation. Expressive musical interaction data-such as timing patterns, gesture dynamics, or performance variability-can act as strong behavioural identifiers. In the MM, inference attacks may enable:

- re-identification of users across sessions or platforms;
- reconstruction of individual performance styles or signatures;
- inference of emotional, cognitive, or physical states.

These threats are particularly difficult to mitigate because they do not necessarily rely on explicit breaches or unauthorized access, but on the aggregation and analysis of legitimately available data.

4.3.3 Neurophysiological and EEG-Specific Threats

Where MUSMET scenarios involve EEG or other neurophysiological signals, the threat landscape becomes significantly more severe. EEG data represent one of the most sensitive forms of personal data, as they may encode information related to cognition, affect, neurological conditions, and individual neural signatures. Key EEG-related threats include:

- Irreversible privacy loss, as neural patterns are difficult or impossible to fully anonymize;
- Identity linkage, where EEG features enable persistent re-identification across contexts;
- Health and cognitive inference, extending far beyond the original musical interaction purpose;
- Function creep, where EEG-derived models or features are reused for unintended analysis.

Unlike many other data types, EEG signals retain sensitivity even after transformation, feature extraction, or aggregation, making both raw data and learned representations high-risk assets.

4.3.4 Threats in AI/ML Model Training and Use

AI/ML models trained on MUSMET data introduce additional attack vectors, particularly when learning is performed on expressive or neurophysiological inputs. Adversaries may attempt to extract information about training data through:

- model inversion, reconstructing approximate input data from trained models;
- membership inference, determining whether specific users or sessions contributed to training.

In MUSMET, such attacks could expose sensitive performance data or confirm participation in particular sessions, even when raw data are no longer accessible.

Data Poisoning and Manipulation. In online or adaptive learning scenarios, MUSMET systems may be vulnerable to data poisoning attacks, where malicious inputs are injected to distort model behaviour. Even small amounts of poisoned data can bias models or destabilize adaptive components. Potential impacts include:

- degraded synchronization or interaction quality;
- systematic bias in adaptive behaviour;
- subtle manipulation of collaborative dynamics.

These attacks may be particularly difficult to detect in real-time systems, where performance fluctuations are common.

Adversarial Inputs and Evasion. AI/ML components may also be susceptible to adversarial inputs crafted to trigger incorrect or unstable outputs. In the MM, adversarial manipulation of audio features, motion signals, or sensor inputs may cause learning-based components to misinterpret interaction states or produce unintended responses. Such attacks exploit the gap between model assumptions and real-world variability, and may remain unnoticed until their cumulative effects become apparent.

4.3.5 Centralised vs Distributed Learning Risks

The architecture chosen for model training significantly affects the threat surface. Centralised learning approaches, where raw data are collected and stored in a central repository, create high-value targets and amplify the consequences of breaches or misuse-particularly for EEG and biometric data.

While distributed approaches, such as federated learning, reduce exposure of raw data, they do not eliminate threats entirely. Residual risks include:

- leakage through shared gradients or model updates;
- poisoning or manipulation of distributed training processes;
- inference attacks targeting shared model parameters.

From a threat perspective, learning architecture selection represents a critical design decision that shapes both risk exposure and mitigation options.

4.3.6 Secondary Use and Long-Term Accumulation Risks

Data and models developed within MUSMET may persist beyond individual sessions, enabling long-term accumulation and secondary use. Over time, aggregated datasets and evolving models may reveal patterns not apparent in short-term interaction. This creates risks related to:

- longitudinal profiling of users or performers;
- unintended cross-context inference;
- gradual erosion of privacy guarantees.

In the MM, where creative interaction may span repeated sessions and collaborations, such cumulative risks are particularly pronounced.

4.3.7 Opacity, Explainability, and Trust

AI/ML-driven processing can introduce opacity into system behaviour, making it difficult for users and operators to understand how data influence outcomes. This lack of transparency complicates detection of malicious manipulation, assessment of privacy risks, and attribution of errors. In collaborative musical environments, unexplained changes in system behaviour, such as altered synchronization or adaptive responses, can undermine trust between participants and in the platform itself.

4.3.8 Implications for the Musical Metaverse

Data and AI/ML threats represent a core challenge for MUSMET, as they often operate silently, cumulatively, and indirectly, without triggering clear system failures. The combination of expressive interaction data, neurophysiological signals, and learning-based processing creates a threat profile that is both high-impact and difficult to observe. The identification of these threats motivates:

- careful classification and handling of sensitive data;
- cautious deployment of learning-based components;
- and evaluation of privacy-preserving and robust learning strategies compatible with real-time operation.

These considerations directly inform the mitigation measures and architectural decisions discussed in subsequent sections of this deliverable.

4.4 Device- and Sensor-Layer Threats

Device- and sensor-layer threats arise from vulnerabilities associated with the physical and embedded components that enable interaction in the MM. MUSMET environments rely on a diverse ecosystem of XR headsets, motion tracking systems, musical IoT devices, wearable sensors, and, in some scenarios, neurophysiological sensing equipment. These components form the primary interface between users and the virtual environment and are therefore critical to both system functionality and user safety.

Because device- and sensor-layer components operate close to the physical world, threats at this layer can compromise not only data confidentiality and integrity, but also interaction fidelity, trust in system behaviour, and, in extreme cases, user well-being.

4.4.1 Device Compromise and Unauthorized Control

XR headsets, wearable devices, and musical IoT components often run complex firmware and software stacks that may contain vulnerabilities exploitable by attackers. Device compromise can occur through malware installation, exploitation of unpatched vulnerabilities, or malicious firmware updates.
Once compromised, devices may:

- leak sensitive sensor data or credentials;
- manipulate outgoing data streams;
- inject falsified or misleading interaction signals;
- serve as entry points for attacks on other system components.

Given the central role of end devices in MUSMET, such compromises can have cascading effects across the system.

4.4.2 Sensor Spoofing and Signal Manipulation

Many MUSMET interactions rely on sensor data-such as motion tracking, spatial positioning, haptic input, or physiological signals-to drive real-time behaviour. Sensor spoofing attacks aim to inject false or manipulated signals into the system, either by tampering with hardware or by exploiting software interfaces. In the MM, spoofed sensor data may:

- disrupt timing and coordination between performers;
- alter perceived gestures or posture;
- manipulate expressive cues critical to musical interaction.

Because sensor data are often trusted implicitly by higher-level components, spoofing attacks can be difficult to detect and may persist undetected for extended periods.

4.4.3 Threats Specific to Wearable and Neurophysiological Sensors

Wearable sensors and neurophysiological devices, such as EEG headsets, introduce additional threat vectors due to their proximity to the human body and the sensitivity of the data they collect. Threats at this layer include:

- unauthorized access to raw physiological signals;
- manipulation of sensor calibration or signal processing pipelines;
- injection of noise or artefacts affecting downstream analysis.

In addition to privacy risks, manipulation of neurophysiological data may influence adaptive system behaviour in unintended ways, potentially altering user experience or system responses.

4.4.4 Insecure Communication Between Devices and Processing Nodes

Device- and sensor-layer components typically communicate with edge or cloud processing nodes over wireless or short-range links. Weaknesses in these communication channels can expose systems to eavesdropping, replay attacks, or message injection.

In MUSMET, insecure device-to-node communication may result in:

- leakage of raw sensor or audio data;
- desynchronization between physical and virtual actions;
- injection of malicious control messages.

Because these links often operate under strict latency and energy constraints, they may lack robust protection by default.

4.4.5 Physical Access and Tampering

Unlike purely virtual components, MUSMET devices are susceptible to physical access and tampering. Attackers with physical access may manipulate hardware, extract stored data, or alter firmware and sensor configurations. Physical tampering may enable:

- long-term compromise of device integrity;
- covert manipulation of sensor outputs;
- bypassing of software-level security mechanisms.

These threats are particularly relevant in shared or public environments, such as rehearsal spaces, studios, or live performance venues.

4.4.6 Heterogeneity and Update Challenges

The MUSMET device ecosystem is inherently heterogeneous, comprising devices from different manufacturers with varying security capabilities, update mechanisms, and lifecycle management practices. This heterogeneity introduces systemic risks, including:

- inconsistent patching and update policies;
- uneven security guarantees across devices;
- reliance on legacy components with limited protection options.

Attackers may exploit the weakest devices in the ecosystem as entry points, undermining overall system security.

4.4.7 Side-Channel and Environmental Attacks

Certain device- and sensor-layer threats may arise from side-channel leakage or environmental manipulation. For example:

- electromagnetic emissions or power consumption patterns may reveal information about device activity;
- environmental interference may affect sensor accuracy or reliability.

In the MM, such attacks could indirectly expose interaction patterns or disrupt sensor-driven functionality without direct system intrusion.

4.4.8 Implications for the Musical Metaverse

Device- and sensor-layer threats are particularly impactful in MUSMET because they operate at the boundary between the physical and virtual worlds. Compromise at this layer can undermine trust in the

authenticity and integrity of musical interaction, propagate misleading information throughout the system, and amplify the effects of higher-layer attacks. The identification of these threats highlights the importance of:

- robust device-level security practices;
- secure integration of sensors into real-time pipelines;
- and careful consideration of device heterogeneity and lifecycle management.

These considerations directly inform the selection of mitigation measures and architectural safeguards discussed in later sections of this deliverable.

4.5 IPR and Digital Asset Protection Threats

IPR and digital asset protection threats constitute a distinct and critical category within the MM. MUSMET environments are inherently creative and content-centric, enabling real-time musical performance, collaborative composition, and the generation of rich user-created digital artefacts. As a result, threats affecting intellectual assets directly undermine not only economic value, but also artistic integrity, trust between collaborators, and long-term sustainability of the platform.

Unlike generic Metaverse applications, where digital assets may be primarily visual or transactional, MUSMET systems handle time-sensitive, expressive, and often ephemeral musical content, which is particularly vulnerable to misuse, unauthorized replication, and manipulation.

4.5.1 Unauthorized Recording and Capture of Performances

One of the most significant IPR-related threats in the MM is the unauthorized recording or capture of live musical performances. XR environments may provide built-in capabilities for audio or audiovisual capture, screen recording, or data export, which can be exploited without the knowledge or consent of performers. Such threats can lead to:

- loss of control over unpublished or experimental material;
- unauthorized redistribution of performances;
- erosion of performers' trust in collaborative environments.

The real-time nature of MUSMET exacerbates this risk, as performances often occur before formal rights management or protection mechanisms can be applied.

4.5.2 Uncontrolled Replication and Redistribution of Digital Assets

Digital assets created or used within MUSMET, such as compositions, sound designs, virtual instruments, spatial audio scenes, or digital stages, can be replicated at negligible cost once accessed. Application- or system-level weaknesses may allow uncontrolled copying and redistribution of these assets. This threat is particularly acute in collaborative contexts, where:

- access rights may change dynamically;
- assets may be shared transiently for creative purposes;
- boundaries between private and public use are blurred.

Uncontrolled replication undermines creators' ability to manage how and where their work is used.

4.5.3 Content Manipulation and Integrity Violations

Beyond unauthorized copying, attackers may attempt manipulation of musical or digital assets, altering content while preserving its apparent authenticity. Such integrity violations may include:

- subtle modification of audio streams;
- alteration of compositions or arrangements;
- changes to expressive parameters embedded in performances.

These manipulations can damage artistic intent, misrepresent creators, and create disputes over authorship or responsibility.

4.5.4 Impersonation and Fraudulent Attribution

IPR threats in MUSMET are closely linked to identity misuse and impersonation. Attackers may present themselves as legitimate artists, performers, or collaborators, creating or distributing content under false identities. This can result in:

- fraudulent attribution of works,
- reputational harm to legitimate creators,
- confusion regarding ownership and rights.

In immersive environments, where identity is mediated by avatars and virtual representations, such threats are particularly difficult to detect.

4.5.5 Misuse of Collaborative Creation Mechanisms

MUSMET supports collaborative and co-creative workflows, enabling multiple participants to jointly shape musical output. While central to the project's vision, these mechanisms also introduce IPR-related vulnerabilities. Threats include:

- unauthorized modification of shared assets;
- extraction of partial contributions without consent;
- disputes arising from unclear provenance or contribution tracking.

Without appropriate safeguards, collaborative creation environments may inadvertently facilitate IPR violations.

4.5.6 Persistence and Reuse of Ephemeral Content

MM interactions often involve ephemeral performances or transient creative processes that participants may not intend to persist beyond a session. However, system-level logging, caching, or recording mechanisms may inadvertently enable persistence and later reuse. Such persistence may:

- conflict with creators' expectations;
- enable retrospective misuse of content;
- complicate control over lifecycle and deletion of assets.

This threat highlights the tension between system functionality (e.g. debugging, optimisation) and creative intent.

4.5.7 Cross-Platform and Interoperability Risks

As MUSMET aims to interoperate with diverse tools, platforms, and services, cross-platform movement of digital assets introduces additional IPR risks. Differences in access control models, metadata handling, or rights representation may lead to loss of protection when assets cross system boundaries. Attackers may exploit these gaps to:

- strip attribution or provenance information;
- bypass usage restrictions;
- redistribute assets in less protected environments.

4.5.8 Implications for the Musical Metaverse

IPR and digital asset protection threats in MUSMET are deeply intertwined with security, identity, and data protection challenges. Unlike traditional digital media platforms, MUSMET must protect creative assets while they are being created and performed in real time, under strict latency and usability constraints. The identification of these threats underscores the need for:

- integrated protection mechanisms that operate seamlessly during interaction;
- safeguards that preserve artistic integrity without disrupting creativity;
- and alignment between technical security measures and the unique workflows of musical collaboration.

These threats inform the design of mitigation measures and protection strategies discussed in subsequent sections of this deliverable, particularly those addressing access control, secure streaming, provenance, and rights-aware system design.

4.6 User Safety, Harassment, and Social Risks

Beyond technical vulnerabilities, the MM introduces a range of user safety, harassment, and social risks that arise from immersive, persistent, and highly interactive environments. In MUSMET, these risks are closely intertwined with security and privacy considerations, as misuse of system capabilities, identity mechanisms, or data flows can directly impact user well-being, trust, and willingness to participate in musical collaboration.

Musical interaction in XR is inherently expressive and social, often involving close virtual proximity, shared creative spaces, and emotionally engaging experiences. As a result, social risks in the MMmay have stronger psychological and behavioural effects than those observed in conventional online platforms.

4.6.1 Harassment and Abusive Behaviour in Immersive Spaces

Immersive environments enable new forms of harassment that extend beyond text or voice-based abuse. In the MM, harassment may manifest through:

- verbal abuse or disruptive audio behaviour during rehearsals or performances;
- intrusive or aggressive avatar positioning and gestures;
- deliberate disruption of musical interaction to undermine performers or collaborators.

Because musical performance relies on concentration, trust, and emotional engagement, such behaviours can have a disproportionate impact on participants, discouraging creative expression and collaboration.

4.6.2 Identity Misrepresentation and Social Engineering

As discussed in earlier sections, identity in the Musical Metaverse is mediated through avatars and application-level representations. Identity misrepresentation can therefore become a social risk when used to deceive, manipulate, or exploit other participants. Examples include:

- impersonation of trusted collaborators or recognized performers,
- creation of deceptive avatars to gain access to private sessions,
- manipulation of social dynamics through false identity cues.

Such behaviour undermines trust within collaborative musical environments and can lead to both social harm and technical security incidents.

4.6.3 Unwanted Presence and Invasion of Personal Space

XR environments blur traditional notions of personal space. In the Musical Metaverse, users may experience unwanted presence, where other participants intrude into virtual proximity or persistently engage without consent. In musical contexts, such intrusions may:

- disrupt concentration and performance flow;
- create discomfort or stress;
- impair the sense of safety within shared creative spaces.

These risks are amplified by the immersive and embodied nature of XR interaction, where virtual actions can feel psychologically immediate.

4.6.4 Exposure to Disturbing or Manipulative Content

Users may be exposed to disturbing, misleading, or manipulative content, either intentionally injected by malicious actors or inadvertently introduced through misuse of creative tools. In MUSMET, this may include:

- sudden injection of loud, disorienting, or harmful audio signals;
- misleading visual or spatial cues affecting perception;
- manipulation of collaborative tools to create confusion or distress.

Such exposures may have both psychological and physical implications, particularly in prolonged immersive sessions.

4.6.5 Social Exclusion and Power Imbalances

Collaborative musical environments may unintentionally reproduce or amplify social exclusion and power imbalances. Certain participants may dominate interaction, control shared resources, or marginalize others through technical or social means. These dynamics can be reinforced by:

- unequal access to moderation or control tools;
- disparities in technical proficiency or device capabilities;
- lack of transparent rules governing interaction.

While not malicious in intent, such risks can still undermine fairness and inclusivity in MUSMET environments.

4.6.6 Psychological and Cognitive Risks

The immersive and emotionally engaging nature of the MM may intensify psychological and cognitive risks, particularly during extended sessions or high-pressure performances. Potential risks include:

- increased stress or fatigue;
- sensory overload due to complex audiovisual stimulation;
- reduced awareness of physical surroundings.

When combined with harassment or disruptive behaviour, these factors may significantly affect user well-being.

4.6.7 Challenges in Detection and Response

User safety and social risks are often difficult to detect and address using purely technical means. Unlike network or application-layer attacks, harmful social behaviour may be subtle, context-dependent, and subjective. Challenges include:

- distinguishing intentional harassment from artistic expression;
- balancing freedom of creative interaction with protection from harm;
- responding to incidents in real time without disrupting musical flow.

These challenges complicate the design of effective safeguards and highlight the need for integrated approaches.

4.6.8 Implications for the Musical Metaverse

User safety, harassment, and social risks are integral to the overall security and trustworthiness of the Musical Metaverse. In MUSMET, these risks can directly affect participation, collaboration quality, and long-term adoption of the platform. Identifying these threats underscores the importance of:

- designing interaction mechanisms that support safe and respectful collaboration;
- integrating security, privacy, and social safeguards coherently;
- and considering user well-being as a core system requirement rather than an afterthought.

The mitigation of these risks is addressed in subsequent sections, where technical, architectural, and organisational measures are discussed in the context of MUSMET's real-time and collaborative constraints.

5. Suitability Analysis of Existing Security and Privacy Protocols under MM Latency Constraints

5.1 Requirements for Secure Ultra-Low-Latency Communication

Secure communication in the Musical Metaverse must reconcile stringent performance constraints with strong security and privacy guarantees. MUSMET scenarios involve heterogeneous communication flows with distinct latency requirements, ranging from ultra-low-latency, tightly synchronized interaction between musicians to low-latency, large-scale dissemination from musicians to audiences. These differing interaction modes impose fundamentally different constraints on the design and deployment of security mechanisms.

This section defines the key requirements that secure communication solutions must satisfy to be considered suitable for MUSMET, serving as the foundation for the protocol evaluation presented in subsequent sections.

5.1.1 Latency-Zone Awareness and Differentiation

MUSMET communication spans at least two latency zones that must be explicitly recognised and supported by security mechanisms.

Musician-to-musician communication requires ultra-low end-to-end latency, on the order of tens of milliseconds, minimal jitter, and highly deterministic timing to preserve rhythmic alignment and ensemble coherence.

Musician-to-audience communication can tolerate low latency, typically in the range of 200–300 ms, allowing greater buffering, aggregation, and processing flexibility while prioritising scalability and robustness.

In addition, **audience-to-audience communication** also falls within the low-latency category. Such interactions include social exchange, shared experiences, and collaborative audience participation, where strict musical synchronisation is not required but timely and responsive interaction remains important. While latency constraints are less stringent than for musician-to-musician interaction, audience-to-audience communication may involve larger groups and dynamic participation, placing greater emphasis on scalability, moderation, and consistent user experience.

Security mechanisms must therefore be capable of differentiating between communication paths and applying protection strategies appropriate to each latency zone and interaction context. Uniform security treatment across all flows risks either compromising ultra-low-latency musical interaction or underutilising available latency margins in audience-related communication, where stronger or more scalable protection mechanisms could be applied without perceptible impact on user experience.

5.1.2 Preservation of Ultra-Low Latency for Musical Interaction

For musician-to-musician communication, the primary requirement is that security mechanisms preserve ultra-low latency and do not introduce delays that impair musical performance, meaning that cryptographic processing, authentication checks, and integrity verification must be lightweight, predictable, and carefully integrated into real-time pipelines. This requirement excludes security mechanisms that introduce additional buffering or retransmission delays, multi-round handshake procedures executed during active interaction, and cryptographic operations with highly variable execution times. In contrast, musician-to-audience communication can accommodate stronger processing and buffering, provided that perceived latency remains within acceptable bounds.

5.1.3 Deterministic Timing and Jitter Control

Musical interaction depends not only on low average latency but also on stable and deterministic timing, as even small fluctuations in delay (jitter) can degrade synchronization between performers. Secure communication mechanisms must therefore minimise timing variability introduced by encryption, authentication, or integrity checks, avoid dynamic reconfiguration that alters processing paths during live sessions, and maintain consistent timing behaviour across different data modalities. While deterministic timing remains important for audience-facing communication, moderate buffering and adaptive delivery techniques are more acceptable in this latency zone.

5.1.4 Compatibility with Continuous Real-Time Streaming

MUSMET communication is dominated by continuous, high-frequency data streams, including audio, visual, and interaction signals. Security mechanisms must therefore operate natively on streaming data without interrupting flow or requiring session restarts. This entails packet- or frame-level protection compatible with real-time transport, seamless operation across long-lived sessions, and preservation of synchronization across multiple data streams. Security approaches designed primarily for transactional or batch-oriented communication are thus unsuitable for MUSMET's interaction paths.

5.1.5 Efficient Session Establishment and Dynamic Participation

MM sessions are inherently dynamic, with participants joining, leaving, or reconfiguring groups during runtime, and secure communication must therefore support rapid session establishment and adaptation without disrupting ongoing interaction. This includes fast authentication and key exchange for musician-to-musician sessions, efficient rekeying mechanisms that avoid audible or perceptible disruption, and scalable onboarding for audience-facing communication. Security overhead that delays session start or interrupts live performance is incompatible with MUSMET use cases.

5.1.6 Robustness Under Network Variability

MUSMET systems operate over heterogeneous networks with variable conditions. Secure communication mechanisms must maintain protection under packet loss, reordering, and fluctuating bandwidth, without amplifying network instability. For ultra-low-latency musician-to-musician interaction, continuity of interaction takes precedence over perfect reliability. For musician-to-audience communication, stronger reliability and recovery mechanisms may be employed, provided they do not significantly increase perceived latency.

5.1.7 Protection of Payload and Communication Metadata

In addition to protecting payload confidentiality and integrity, secure communication mechanisms must consider the exposure of metadata, such as timing, packet rates, and session identifiers. In the MM, such metadata can reveal sensitive information about collaboration patterns, rehearsal schedules, or performance dynamics. Where feasible, security mechanisms should minimize metadata leakage, particularly for musician-to-musician communication, while balancing scalability requirements for audience-facing dissemination.

5.1.8 Support for Heterogeneous Devices and Architectures

Secure communication in MUSMET must operate across a diverse ecosystem of devices and infrastructures, including resource-constrained XR headsets, edge computing nodes, and cloud-based

distribution platforms. Security mechanisms must therefore adapt to varying computational and energy constraints, support both peer-to-peer and mediated communication models, and remain interoperable across distributed and decentralised architectures. A single, monolithic security solution is unlikely to be suitable across all MUSMET components.

5.1.9 Scalability for Multi-Party and Audience Scenarios

Musician-to-audience communication introduces scalability requirements that differ markedly from small-group musical interaction. Secure communication mechanisms must support efficient dissemination to large audiences without linear increases in latency or processing overhead. At the same time, scalability mechanisms must be designed so that they do not interfere with ultra-low-latency musician-to-musician interaction, which must remain isolated from audience-related processing.

5.1.10 Fault Isolation and Graceful Recovery

Finally, secure communication mechanisms must support fault isolation and graceful recovery. Failures, attacks, or degradations affecting one communication path should not cascade to other parts of the system. In ultra-low-latency musical interaction, recovery mechanisms must be rapid and minimally disruptive, while audience-facing communication may tolerate more visible recovery procedures.

5.2 Evaluation of Transport-Layer Security Protocols

Transport-layer security and stream protection mechanisms are a key building block for ensuring confidentiality, integrity, and authenticity of communication in the Musical Metaverse. In MUSMET, these mechanisms must operate under heterogeneous latency constraints, support continuous real-time streaming, and scale across distributed, multi-party interaction scenarios.

This section evaluates the suitability of established transport-layer security protocols and lightweight stream protection techniques, including scrambling approaches inspired by digital television and broadcast streaming systems. The evaluation is guided by the requirements defined in Section 5.1 and explicitly differentiates between ultra-low-latency musician-to-musician communication and low-latency musician-to-audience communication.

5.2.1 Evaluation Scope and Methodological Perspective

The evaluation presented in this section does not aim to identify a single universally optimal protocol. Instead, it adopts a context-aware perspective, analysing tradeoffs between security guarantees, latency impact, scalability, and operational complexity. Each protocol or technique is assessed with respect to:

- its impact on latency and jitter;
- its suitability for continuous real-time streaming;
- session establishment and rekeying overhead;
- robustness under network variability;
- scalability to multi-party and audience-facing scenarios.

This approach reflects the fact that MUSMET requires differentiated security solutions aligned with distinct communication paths and latency zones.

5.2.2 Analysis of Secure Communication and Stream Protection Approaches

5.2.2.1 TLS 1.3 over TCP

Transport Layer Security (TLS) 1.3 is the dominant standard for securing reliable, connection-oriented transport protocols such as TCP.

- **Musician-to-musician communication:** TLS over TCP is generally unsuitable for ultra-low-latency musical interaction. TCP's congestion control, retransmission mechanisms, and head-of-line blocking introduce latency and jitter that conflict with the tight synchronization requirements of ensemble performance. Even with TLS 1.3's reduced handshake latency, connection setup and recovery remain too heavyweight for time-critical interaction.
- **Musician-to-audience communication:** TLS over TCP is technically feasible for musician-to-audience communication, as this interaction path can tolerate higher end-to-end latency. It may therefore be used for non-time-critical channels such as control signalling, metadata exchange, or management interfaces. However, for live musical streaming, TLS over TCP remains sub-optimal due to latency variability and unpredictable behaviour under adverse network conditions, which can negatively affect perceived quality of experience.
- **Audience-to-audience communication:** Audience-to-audience interaction also falls within the low-latency category and can, in principle, be supported by TLS over TCP, particularly for social interaction, messaging, or shared experience features. Nevertheless, similar to musician-to-audience scenarios, TLS over TCP is sub-optimal for interactive or media-rich audience communication, as TCP-level retransmissions and congestion control may introduce unnecessary delay and jitter when compared to UDP-based alternatives.
- **Assessment:** Feasible for low-latency, non-time-critical communication paths, but sub-optimal for real-time or interactive media exchange. Overall applicability is limited, primarily outside MUSMET's real-time musical interaction paths.

5.2.2.2 DTLS 1.3 over UDP

Datagram Transport Layer Security (DTLS) 1.3 extends TLS-equivalent security guarantees to unreliable, datagram-oriented transports such as UDP.

- **Musician-to-musician communication:** DTLS is a strong candidate for ultra-low-latency interaction, as it avoids TCP head-of-line blocking and supports packet-oriented streaming. Improvements in DTLS 1.3 reduce handshake overhead, but session establishment and rekeying must still be carefully managed to avoid disrupting musical timing.
- **Musician-to-audience communication:** DTLS is also suitable for low-latency audience-facing communication, although scaling to large audiences may require additional architectural support, such as relays or edge-based distribution.
- **Audience-to-audience communication:** DTLS is suitable for audience-to-audience interaction in low-latency scenarios, supporting secure, packet-oriented communication without TCP-related blocking. It can be effectively used for interactive or media-enabled audience features, provided that session management is kept lightweight.
- **Assessment:** High suitability for MUSMET real-time communication when optimised for minimal handshake and rekeying overhead, including musician-to-musician interaction and low-latency audience-facing communication, such as musician-to-audience and audience-to-audience paths.

5.2.2.3 Secure Real-Time Transport Protocol (SRTP)

The Secure Real-Time Transport Protocol (SRTP) is specifically designed to protect real-time audio and video streams, providing encryption, integrity protection, and replay prevention with minimal per-packet overhead.

- **Musician-to-musician communication:** SRTP is particularly well aligned with ultra-low-latency musical interaction. Its lightweight design, streaming-oriented operation, and predictable processing characteristics support tight synchronization and continuous audio exchange.
- **Musician-to-audience communication:** SRTP also scales well to one-to-many scenarios and is widely used in live media distribution. It remains effective across both latency zones, provided that key management is handled efficiently.
- **Audience-to-audience communication:** SRTP is also suitable for audience-to-audience interaction involving real-time audio or audiovisual exchange. Its low per-packet overhead and predictable behaviour support low-latency audience interaction without introducing unnecessary processing delays.
- **Assessment:** Very high suitability for MUSMET, particularly for audio-centric real-time interaction, including musician-to-musician, musician-to-audience, and audience-to-audience communication.

5.2.2.4 QUIC and HTTP/3-Based Secure Transport

QUIC is a modern transport protocol built on UDP, integrating security based on TLS 1.3 and supporting multiplexed streams without head-of-line blocking.

- **Musician-to-musician communication:** While QUIC offers reduced connection setup latency and improved resilience to packet loss, its internal complexity and stream multiplexing mechanisms may introduce timing variability that is problematic for the strictest ultra-low-latency musical interaction paths.
- **Musician-to-audience communication:** For musician-to-audience communication, QUIC and HTTP/3 are strong candidates. They support efficient large-scale distribution, adaptive delivery, and robust performance under variable network conditions, fitting well within the relaxed latency budget.
- **Audience-to-audience communication:** Audience-to-audience interaction falls within the low-latency category and aligns well with the characteristics of QUIC-based transport. QUIC and HTTP/3 can support interactive and media-rich audience communication by providing efficient connection establishment, multiplexed streams, and resilience to packet loss. These properties are particularly beneficial in scenarios involving dynamic participation and varying network conditions, where responsiveness and scalability are important but strict musical synchronisation is not required.
- **Assessment:** Moderate suitability for ultra-low-latency musician-to-musician interaction; high suitability for audience-facing communication, including both musician-to-audience and audience-to-audience interaction. QUIC-based approaches offer a strong balance between security, robustness, and scalability for low-latency interaction paths in MUSMET.

5.2.2.5 Lightweight Scrambling Techniques Inspired by Digital TV Streaming

Scrambling techniques widely used in digital television and broadcast streaming systems provide *lightweight stream obfuscation* with minimal computational overhead and near-zero added latency. These techniques are designed to deter casual or opportunistic access rather than to provide strong cryptographic confidentiality.

- **Musician-to-musician communication:** In ultra-low-latency interaction, scrambling may serve as a *complementary or fallback protection mechanism, particularly during session setup or transient peer-to-peer communication. However,* scrambling alone does not provide sufficient protection for sensitive data and cannot replace cryptographic security.
- **Musician-to-audience communication:** Scrambling techniques are *particularly well suited* to musician-to-audience scenarios, which closely resemble broadcast and IPTV use cases.

They enable scalable, low-latency dissemination while protecting artistic and commercial value against unauthorized access.

- **Audience-to-audience communication:** Audience-to-audience interaction also fits well with scrambling-based protection. In social or shared-experience scenarios, scrambling can deter casual interception while preserving low latency and supporting large numbers of participants. As with musician-to-audience communication, this level of protection is generally sufficient where the primary goal is content access control rather than strong confidentiality guarantees.
- **Assessment:** High applicability for audience-facing communication, including both musician-to-audience and audience-to-audience interaction. Scrambling techniques offer effective, low-overhead protection for scalable dissemination, while remaining only a complementary measure for ultra-low-latency musician-to-musician interaction.

5.2.3 Comparative Perspective Across Latency Zones

The evaluation confirms that no single transport-layer protection mechanism satisfies all MUSMET requirements. Instead:

- Musician-to-musician communication benefits most from real-time media-oriented solutions with minimal overhead (e.g. SRTP, carefully configured DTLS), with lightweight scrambling as a supplementary measure;
- Musician-to-audience communication can leverage a broader range of techniques, including SRTP, QUIC/HTTP/3, and scrambling, balancing security, scalability, and latency.

This differentiated approach aligns with MUSMET's multi-zone latency model and supports a risk- and performance-aware security architecture.

5.2.4 Summary

This section demonstrates that effective security in the Musical Metaverse requires context-sensitive selection of transport-layer protection mechanisms. Lightweight scrambling techniques, long established in digital TV streaming, emerge as a valuable complementary approach in MUSMET, particularly for audience-facing communication paths where scalability and latency constraints limit the applicability of heavyweight cryptographic solutions.

The insights from this evaluation directly inform the definition of MUSMET-specific security measures and architectural decisions discussed in subsequent sections of this deliverable.

5.3 Evaluation of Network-Layer Security Protocols

Network-layer security protocols provide protection at a lower level of the communication stack than transport-layer mechanisms, typically securing traffic between hosts, subnets, or network segments independently of application semantics. In MUSMET, network-layer security can play a supporting role in protecting infrastructure-level communication and establishing baseline trust between system components. However, its suitability for real-time musical interaction is strongly constrained by latency, jitter, and lack of application awareness.

This section evaluates the applicability of network-layer security protocols in the context of the Musical Metaverse, considering their strengths, limitations, and appropriate deployment scenarios.

5.3.1 Role of Network-Layer Security in MUSMET

Network-layer security mechanisms are designed to be transparent to applications, offering generic protection for all traffic traversing a given network path, which in MUSMET can be both an advantage and a limitation. At a conceptual level, network-layer security may protect communication between trusted infrastructure components such as edge–cloud links, reduce exposure on untrusted or shared networks, and provide baseline confidentiality and integrity guarantees. However, because these mechanisms lack awareness of interaction context, latency zones, and data semantics, they are ill-suited as the primary protection mechanism for time-critical musical streams.

5.3.2 Analysis of Network-Layer Security Approaches

5.3.2.1 IPsec (Tunnel and Transport Modes)

IPsec is the most widely deployed network-layer security framework, providing encryption, authentication, and integrity protection at the IP level.

- **Musician-to-musician communication:** IPsec is generally **unsuitable** for ultra-low-latency musical interaction. Its cryptographic processing overhead, packet encapsulation, and limited control over timing behaviour introduce latency and jitter that can disrupt tight musical synchronization. Moreover, IPsec operates independently of application-level timing constraints, making it difficult to align with MUSMET's ultra-low-latency requirements.
- **Musician-to-audience communication:** For audience-facing communication, IPsec may be feasible in limited scenarios, particularly where traffic flows between trusted network segments (e.g. from core infrastructure to content distribution nodes). However, it remains poorly suited for end-to-end protection of live musical streams due to scalability and latency considerations.
- **Audience-to-audience communication:** For audience-to-audience interaction, IPsec may be feasible in constrained scenarios where communication occurs within trusted network segments or managed infrastructure. However, due to its overhead and lack of application-level awareness, it is poorly suited for interactive or media-rich audience communication requiring low and predictable latency.
- **Assessment:** Low suitability for real-time MUSMET communication. IPsec may be used for securing infrastructure-level links or trusted network segments, but remains unsuitable for end-to-end protection of musician-to-musician, musician-to-audience, or audience-to-audience real-time interaction.

5.3.2.2 Virtual Private Networks (VPNs)

VPN technologies, often built on IPsec or similar mechanisms, are commonly used to establish secure tunnels over untrusted networks.

- **Musician-to-musician communication:** VPNs introduce additional encapsulation, routing constraints, and processing overhead that conflict with ultra-low-latency requirements. Their use may significantly increase round-trip times and jitter, making them unsuitable for real-time ensemble interaction.
- **Musician-to-audience communication:** VPNs may be appropriate for administrative access, development environments, or backend services, but they do not scale well to large audiences and are not designed for real-time media distribution.
- **Audience-to-audience communication:** VPN technologies are generally unsuitable for audience-to-audience interaction in MUSMET, particularly where interactive or media-rich communication is expected. The additional tunnelling overhead and routing constraints

introduced by VPNs can negatively affect latency and responsiveness, making them ill-suited for low-latency audience interaction beyond basic administrative or controlled scenarios.

- **Assessment:** Unsuitable for MUSMET real-time interaction, including musician-to-musician, musician-to-audience, and audience-to-audience communication. VPN technologies remain limited to management, administrative access, or backend communication within trusted environments.

5.3.2.3 Software-Defined Networking and Network-Level Policy Enforcement

Software-defined networking (SDN) enables programmable control over network behaviour, including traffic isolation, prioritization, and routing. While SDN is not a security protocol per se, it can support *network-level enforcement of security policies*, such as isolating critical MUSMET traffic from other network flows, prioritising ultra-low-latency musical streams, andlimiting exposure to untrusted network segments.

- **Musician-to-musician communication:** SDN can indirectly enhance security and performance by ensuring predictable network paths and prioritised delivery, but it does not provide confidentiality or integrity protection on its own.
- **Musician-to-audience communication:** SDN may support scalable and controlled distribution paths, particularly in managed network environments.
- **Audience-to-audience communication:** SDN can indirectly support audience-to-audience interaction by enabling traffic isolation and prioritisation in managed networks, helping maintain stable performance without providing direct data protection.
- **Assessment:** Supportive role for MUSMET communication. SDN complements cryptographic security mechanisms, including for audience-to-audience interaction, but does not replace them.

5.3.3 Network-Layer Security vs Transport-Layer Security

Compared to the transport-layer mechanisms evaluated in Section 5.2, network-layer security protocols exhibit several limitations in the MUSMET context, including a lack of awareness of latency zones and interaction semantics, limited ability to differentiate protection strength by data type, and a higher risk of introducing latency and jitter into real-time communication paths. As a result, network-layer security should not be relied upon as the primary mechanism for protecting musical interaction or expressive data streams.

5.3.4 Appropriate Deployment Scenarios in MUSMET

Despite their limitations, network-layer security protocols can play a supporting role in MUSMET when applied selectively and strategically, for example by securing communication between trusted infrastructure components such as edge-to-cloud links, protecting control and management traffic, isolating MUSMET services from unrelated network traffic, and establishing secure administrative access paths. In these scenarios, latency constraints are typically less stringent, and the benefits of network-layer protection outweigh the associated performance costs.

5.3.5 Summary

The evaluation demonstrates that network-layer security protocols have limited applicability for real-time musical interaction in the MM. Their lack of application awareness and potential impact on latency make them unsuitable for protecting ultra-low-latency musician-to-musician communication and suboptimal for live audience-facing streams.

Nevertheless, when deployed judiciously, network-layer security can enhance the overall security posture of MUSMET by protecting infrastructure-level communication and reducing exposure on untrusted networks. Effective MUSMET security architectures should therefore treat network-layer security as a complementary baseline mechanism, while relying on transport- and application-layer solutions for time-critical and content-aware protection.

5.4 Privacy-Preserving Architectures

Privacy-preserving architectures are a critical consideration for the MM, where real-time interaction, immersive embodiment, and expressive performance generate rich and potentially sensitive data. In MUSMET, privacy challenges arise not only from the content of exchanged data, but also from interaction patterns, behavioural signals, and contextual information inherent to musical collaboration in XR environments. This section analyses architectural approaches that can support privacy preservation while remaining compatible with MUSMET's performance, latency, and scalability requirements.

Rather than focusing on individual privacy-enhancing technologies in isolation, the analysis adopts a system-level architectural perspective, examining how data flows, processing placement, and trust assumptions influence privacy outcomes.

5.4.1 Data Minimisation and Purpose Limitation by Design

A foundational principle of privacy-preserving architectures is data minimisation, ensuring that only data strictly necessary for a given function or interaction is processed. In the Musical Metaverse, this principle is particularly important due to the continuous and high-frequency nature of audio, interaction, and sensor data.

Architectures that enforce purpose limitation at design time reduce privacy risks by preventing unnecessary aggregation, replication, or reuse of data across components. This includes separating real-time interaction data from analytics or monitoring pipelines, and avoiding architectural shortcuts that expose expressive or behavioural data to components that do not require it. For MUSMET, such approaches are compatible with both ultra-low-latency and low-latency interaction paths, as minimisation primarily affects architectural structure rather than per-packet processing.

5.4.2 Local and Edge-Centric Processing Architectures

One of the most effective architectural strategies for privacy preservation is the localisation of data processing. By processing sensitive data close to its source, architectures can significantly reduce the risks associated with centralised storage, long-term retention, and cross-context reuse.

In the context of MUSMET, local or near-edge processing is particularly relevant for expressive interaction data and biometric or neurophysiological signals. Architectures that favour local computation enable real-time responsiveness while limiting exposure of raw data to external services or infrastructure components. This approach also aligns well with strict latency requirements for musician-to-musician interaction.

While local processing does not eliminate all privacy risks, it provides a strong architectural baseline that reduces the scope of data governance challenges at the system level.

5.4.3 Distributed and Federated Architectural Patterns

Distributed and federated architectures represent an alternative to fully centralised data processing models. In such architectures, data remains under the control of local nodes, while only aggregated results, model updates, or derived representations are exchanged.

From a privacy perspective, federated approaches can reduce direct data sharing and support collaborative functionality without exposing raw personal data. This is particularly relevant for adaptive or learning-based components that operate on expressive or behavioural data.

However, these architectures introduce additional complexity and coordination overhead. In MUSMET, their applicability must be carefully evaluated against real-time constraints, especially for ultra-low-latency interaction paths. Federated approaches are therefore more naturally aligned with background processing, adaptation, or audience-related scenarios than with tight musical synchronization.

5.4.4 Separation of Interaction, Control, and Analytics Planes

Privacy-preserving architectures often benefit from a clear separation of functional planes within the system. Distinguishing between interaction data, control and signalling information, and analytics or logging streams allows privacy-relevant policies to be applied more precisely.

In the MM, such separation helps prevent secondary use of interaction data beyond its original purpose and reduces the risk of unintended inference. It also supports differentiated privacy treatment for musician-to-musician, musician-to-audience, and audience-to-audience interactions.

This architectural separation complements the latency-aware communication strategies discussed in Sections 5.1 and 5.2 and enables more granular privacy control without introducing excessive overhead.

5.4.5 Architectural Support for Anonymity and Pseudonymity

Where appropriate, privacy-preserving architectures may incorporate anonymity or pseudonymity mechanisms, particularly for audience-facing interaction. By decoupling user identity from interaction data, such architectures can reduce profiling and tracking risks.

In MUSMET, the applicability of anonymity is context-dependent. While performers typically require stable identities for collaboration and attribution, audience participants may benefit from pseudonymous participation, especially in social or exploratory interaction scenarios. Architectures that support multiple identity representations without tightly coupling them to low-level data streams offer greater flexibility in managing privacy expectations.

5.4.6 Limitations and Trade-Offs

Privacy-preserving architectures inevitably involve trade-offs between privacy guarantees, system complexity, and performance. Approaches such as local processing, federated learning, or strict data separation may increase implementation complexity or limit certain types of analytics and optimisation.

In MUSMET, these trade-offs are particularly pronounced due to real-time performance requirements and the expressive nature of musical interaction. As a result, privacy-preserving architectures must be evaluated not only in terms of theoretical privacy benefits, but also with respect to their impact on latency, reliability, and user experience.

5.5 Summary of Protocol Suitability for MUSMET

The analysis presented in Sections 5.1–5.4 demonstrates that the suitability of security and privacy protocols in MUSMET is highly context-dependent and strongly influenced by latency constraints, interaction semantics, and data sensitivity. No single protocol or architectural approach satisfies all requirements of the MM.

For ultra-low-latency musician-to-musician interaction, only lightweight, stream-oriented protection mechanisms are suitable, as they preserve deterministic timing and minimise processing overhead.

Protocols and architectural patterns that introduce buffering, retransmission, or complex session management are incompatible with tight musical synchronisation.

For low-latency communication paths, including musician-to-audience and audience-to-audience interaction, a broader range of security mechanisms can be applied. Transport-layer protocols with stronger security guarantees, scalable delivery mechanisms, and lightweight stream protection techniques are suitable in these contexts, provided that end-to-end latency remains within acceptable bounds.

At the architectural level, privacy-preserving approaches such as data minimisation, local or edge-centric processing, and separation of interaction and analytics functions are well aligned with MUSMET's requirements. More complex privacy-enhancing techniques may be applicable in non-time-critical components but require careful evaluation against performance constraints.

Overall, the findings of this section underline the need for a differentiated and layered approach to security and privacy in MUSMET. Protocols and architectures must be selected and configured based on interaction context rather than applied uniformly. These conclusions directly inform the scoped security and privacy measures presented in Section 6.

6. Proposed Security & Privacy Measures for MUSMET

6.1 Positioning of Security and Privacy within the MUSMET Project Scope

Security and privacy are critical enablers for the MM, yet their effective treatment depends not only on technical solutions but also on a clear understanding of project scope, responsibilities, and realistic outcomes. MUSMET is conceived as a research and innovation project focused on the design, integration, and validation of MM systems, rather than as a dedicated security technology development initiative. Consequently, the role of security and privacy within MUSMET must be positioned accordingly.

The primary contribution of MUSMET in this domain lies in the systematic identification, contextual analysis, and informed application of security and privacy measures that are compatible with the unique requirements of musical interaction in immersive, real-time environments. This includes ultra-low-latency musician-to-musician interaction, low-latency musician-to-audience dissemination, the handling of expressive and potentially sensitive data, and the coexistence of multiple stakeholders within distributed infrastructures. MUSMET does not aim to design new cryptographic primitives, invent novel trust frameworks, or deliver production-grade security infrastructures. Instead, it focuses on understanding how existing and emerging approaches can be selected, configured, combined, or constrained in ways that are appropriate for the MM context.

Based on the analysis presented in Sections 2 to 5, MUSMET adopts a differentiated approach to security and privacy issues. First, a significant subset of identified risks corresponds to well-known challenges already present in other digital environments, such as secure communication, basic identity management, or infrastructure protection. These issues can be addressed through the correct use of mature standards, protocols, and tools, without requiring MUSMET-specific research and development. In such cases, the project's responsibility is to ensure informed selection and appropriate configuration, rather than to propose new solutions.

Second, the analysis reveals a number of risks that are highly specific to the MM, including long-term governance of biometric and neurophysiological data, advanced intellectual property enforcement across interoperable virtual environments, and trust management in open, cross-platform metaverse ecosystems. While these issues are acknowledged as important and, in some cases, critical, they exceed the feasible scope of the MUSMET project. Addressing them adequately would require sustained research, regulatory alignment, or large-scale operational deployment that goes beyond MUSMET's objectives. These areas are therefore treated as out-of-scope but explicitly identified, providing a basis for future research initiatives or post-project exploitation activities.

Third, MUSMET identifies a subset of security and privacy aspects that it actively addresses within the project, at least to a limited or exploratory extent. These include architectural choices influenced by latency constraints, handling of sensitive data in experimental settings, informed selection of communication protection mechanisms, and security-aware orchestration decisions. In these areas, MUSMET aims to contribute practical insights, validated design choices, and evidence-based recommendations rather than complete, end-to-end security solutions.

This scoped and transparent positioning serves two purposes. It ensures that security and privacy considerations are integrated into MUSMET's technical work in a realistic and responsible manner, and it avoids over-claiming outcomes that cannot be delivered within the project's timeframe and scope. The following subsections of Section 6 build on this positioning by clearly distinguishing between issues addressed through existing tools, those acknowledged but left for future work, and those that MUSMET actively incorporates into its system design and experimental activities.

6.2 Security and Privacy Issues Addressed through Existing Standards and Tools

A substantial portion of the security and privacy challenges identified in the MM are not unique to MUSMET, but rather reflect well-understood issues already present in other distributed, real-time, and multimedia systems. In these areas, MUSMET does not aim to introduce new technical solutions or conduct dedicated security research. Instead, the project relies on the informed selection, configuration, and integration of existing standards, protocols, and tools that are already widely accepted and deployed in comparable contexts.

This category includes baseline protection of communication channels, standard identity and authentication mechanisms, and infrastructure-level security controls. As demonstrated in Section 5, a broad ecosystem of mature transport- and network-layer security protocols exists, and their properties, limitations, and suitability for different latency zones are well documented. For communication paths that do not fall under ultra-low-latency constraints, most notably musician-to-audience or audience-to-audience dissemination, and non-time-critical control channels, these protocols provide adequate protection when correctly configured. MUSMET's contribution in this area is therefore not the development of new mechanisms, but the context-aware application of existing ones, informed by the project's analysis of latency, scalability, and data sensitivity requirements.

Similarly, fundamental aspects of identity management and access control are addressed through established approaches that are already common in distributed systems and XR platforms. Authentication, role assignment, and basic access restrictions are treated as standard engineering concerns, supported by existing frameworks and middleware. MUSMET does not attempt to redefine identity models or propose new trust infrastructures, but instead assumes that baseline identity and access management capabilities are provided by the platforms, tools, or services integrated into the MUSMET ecosystem.

At the infrastructure level, security measures such as network isolation, secure configuration of edge and cloud resources, and protection of management interfaces are also considered out of scope for bespoke MUSMET innovation. These aspects are addressed through best practices and standard tooling commonly used in cloud-native and distributed systems. While such measures are essential for safe system operation, they do not require MUSMET-specific research contributions and are therefore treated as prerequisites rather than project outcomes.

Importantly, treating these issues as "addressed through existing standards and tools" does not imply that they are trivial or unimportant. On the contrary, failure to correctly configure or deploy well-known security mechanisms is a frequent source of vulnerabilities in real-world systems. MUSMET explicitly acknowledges this risk and treats correct integration as a necessary condition for meaningful experimentation, but not as a research challenge in its own right.

By clearly separating standard, well-covered security and privacy concerns from MUSMET-specific challenges, the project ensures that its efforts are focused where they can provide the greatest added value. This distinction also avoids duplication of work already performed by the wider research and industrial communities and helps maintain a realistic and credible scope for MUSMET's security and privacy contributions.

6.3 MUSMET-Specific Risks Identified but Considered Out of Project Scope

The analysis presented in earlier sections identifies a number of security and privacy risks that are highly specific to the MM and that merit careful consideration. At the same time, these risks exceed the realistic scope of the MUSMET project, either due to their systemic nature, long-term implications, or the need for regulatory, organisational, or large-scale technological solutions. In this subsection, such risks are explicitly acknowledged and positioned as out of scope, while clarifying how MUSMET limits its exposure to them.

One prominent category concerns the long-term governance, storage, and secondary use of biometric and neurophysiological data, including EEG signals. As discussed in Sections 2 and 4, such data raise profound privacy, ethical, and trust-related concerns, particularly when collected, aggregated, or reused across sessions and platforms. Addressing these issues comprehensively would require sustained governance frameworks, legal alignment, and long-term data management strategies that go beyond the objectives of MUSMET.

Within MUSMET, this broader problem space is deliberately avoided by design. Neurophysiological and biometric data, where used, are processed locally and transiently, without persistent storage or central aggregation. By limiting processing to local devices or trusted local execution contexts, MUSMET reduces exposure to risks associated with long-term data retention, cross-context reuse, or secondary exploitation. While this approach does not resolve the general challenge of biometric data governance in the MM, it allows MUSMET to explore interaction concepts without engaging in issues that are out of scope for the project.

A second category of out-of-scope risks relates to advanced intellectual property rights enforcement across interoperable Metaverse environments. As identified in Section 4.5, protecting musical works, performances, and collaborative artefacts across multiple platforms raises challenges related to provenance tracking, rights enforcement, and cross-platform consistency. Developing robust, interoperable IPR enforcement mechanisms would require coordination across platforms, stakeholders, and legal regimes, which is beyond the scope of MUSMET. The project therefore limits itself to identifying such risks and considering them in architectural discussions, without attempting to deliver full technical solutions.

Similarly, fully fledged zero-trust frameworks for open, multi-stakeholder MM ecosystems are considered out of scope. While zero-trust principles inform MUSMET's architectural thinking, implementing comprehensive zero-trust infrastructures that span devices, platforms, and organisational boundaries would require extensive integration work and operational deployment that cannot be achieved within the project timeframe.

Additional out-of-scope areas include the development of long-term trust and reputation systems, production-grade digital rights management solutions tightly integrated with artistic workflows, and the enforcement of legal or contractual obligations through technical means. These topics are acknowledged as important for the sustainable evolution of the MM, but they fall outside the scope of MUSMET as a research and integration project.

By explicitly identifying these MUSMET-specific but out-of-scope risks, the project aims to provide a transparent and credible account of its limitations. Rather than ignoring these challenges, MUSMET documents them as areas requiring further research, follow-up projects, or exploitation-phase development. This clear delineation ensures that MUSMET's security and privacy contributions remain focused, realistic, and aligned with its stated objectives.

6.4 Security and Privacy Aspects Explicitly Addressed within MUSMET

Within the boundaries defined in Section 6.1, MUSMET actively addresses a selected set of security and privacy aspects that are both relevant to the MM and feasible within the project scope. These aspects are not treated as standalone security solutions, but as integrated design considerations embedded in system architecture, technology selection, and experimental validation activities.

A central focus of MUSMET is the latency-aware treatment of secure communication. As established in Sections 5.1 and 5.2, the project distinguishes between ultra-low-latency musician-to-musician interaction and low-latency communication paths, including musician-to-audience and audience-to-audience interaction. MUSMET addresses these differences through architectural separation of communication paths and context-aware selection and configuration of protection mechanisms.

Communication channels supporting real-time musical coordination are implemented and evaluated using lightweight, stream-oriented protection approaches that preserve deterministic timing and minimise processing overhead. In contrast, audience-related communication paths leverage their higher

latency tolerance to apply stronger or more scalable protection mechanisms where appropriate. Rather than enforcing a uniform security layer, MUSMET implements and validates differentiated security configurations across interaction contexts, demonstrating how existing mechanisms can be selectively applied without compromising musical performance or user experience.

MUSMET also explicitly addresses the handling of sensitive and expressive data in experimental and pilot settings. While the project does not attempt to establish comprehensive data governance frameworks, it incorporates data minimisation principles in system design and experimentation. Only data strictly necessary for musical interaction and evaluation are processed, and exposure of expressive, biometric, or neurophysiological data is limited by architectural choices, such as local processing and avoidance of persistent storage. This approach allows MUSMET to explore advanced interaction paradigms while reducing privacy risks within the project context.

Another area actively addressed within MUSMET concerns security-aware system orchestration and service placement. The project investigates how orchestration decisions, such as placing processing close to performers or distributing services across edge and cloud resources, influence both latency and security exposure. Security considerations are therefore treated as one of the parameters influencing orchestration strategies, alongside performance and functional requirements. This does not result in a full orchestration security framework, but it provides practical insights into how security and privacy constraints shape architectural trade-offs in the MM.

MUSMET further addresses the informed use of lightweight protection mechanisms, particularly for audience-facing communication. As discussed in Section 5.2, techniques such as real-time media protection protocols and lightweight scrambling approaches are evaluated and applied where appropriate. The project explores how such mechanisms can protect artistic content and reduce opportunistic misuse without introducing prohibitive overhead or disrupting musical experience. These measures are applied selectively and experimentally, rather than as comprehensive rights enforcement solutions.

At the architectural level, MUSMET integrates security-by-design considerations into the definition of system components and interaction flows. This includes explicit consideration of trust boundaries, separation of roles between performers, audiences, and infrastructure components, and avoidance of unnecessary coupling between sensitive data sources and external services. While MUSMET does not implement a complete zero-trust architecture, it adopts a cautious stance towards implicit trust and documents assumptions about trusted components and communication paths.

Finally, MUSMET treats security and privacy not only as objects of validation and reflection, but also as concrete implementation tasks, where existing and mature security and privacy technologies are integrated and deployed within MUSMET prototypes in line with a privacy-by-design approach. During experiments and demonstrations, the project observes how selected security measures interact with musical workflows, latency constraints, and user experience. Identified limitations, residual risks, and trade-offs are documented as part of the project outcomes, contributing evidence-based insights rather than production-ready solutions.

In summary, MUSMET explicitly addresses security and privacy where they intersect with its core technical objectives and experimental activities. The project contributes architectural patterns, configuration choices, and validated design insights that are specific to the MM, while deliberately avoiding claims of delivering comprehensive or deployable security infrastructures. These contributions form a pragmatic and credible basis for future development, exploitation, and follow-up research.

6.5 Security-by-Design Principles Applied in MUSMET System Architecture

In addition to addressing selected security and privacy aspects explicitly, MUSMET adopts a security-by-design perspective in the definition of its system architecture. Rather than treating security and privacy as add-on features, the project incorporates them as design constraints and guiding principles that influence architectural decisions, component boundaries, and interaction flows throughout the system.

A key principle applied in MUSMET is the explicit identification of trust boundaries within the MM architecture. The system is structured to distinguish between performers, audience participants, and infrastructure components, and to avoid unnecessary implicit trust between them. Communication paths, data flows, and processing responsibilities are defined in a way that limits the propagation of sensitive data across components and reduces the impact of potential compromise. This principle is particularly relevant in distributed and heterogeneous environments, where devices and services may operate under different trust assumptions.

Another central design principle is minimisation of exposure, both in terms of data and functionality (Jin et al., 2025). MUSMET architectures are designed to process only the data required for a given interaction context and to keep sensitive data as close as possible to their source. This is reflected, for example, in the preference for local or near-edge processing of expressive, biometric, or neurophysiological data, and in the avoidance of persistent storage where it is not strictly necessary for experimentation or evaluation purposes. By limiting data dissemination and retention, MUSMET reduces the potential attack surface without relying on complex enforcement mechanisms.

Security-by-design in MUSMET is also reflected in the alignment between architectural decisions and latency constraints. Components involved in ultra-low-latency musical interaction are designed to be as simple and predictable as possible, avoiding architectural dependencies that would require complex security processing in time-critical paths. More heavyweight security controls are deliberately placed outside these paths, for example in audience-facing communication or management components, where latency budgets allow greater flexibility. This separation ensures that security measures support, rather than undermine, musical performance requirements.

Furthermore, MUSMET adopts a principle of explicit assumptions and documented limitations. Architectural designs clearly state which components are assumed to be trusted, which threats are mitigated by design choices, and which risks remain unaddressed within the project scope. This transparency is an important element of security-by-design in a research context, as it enables informed interpretation of experimental results and avoids hidden dependencies or unrealistic security claims.

While MUSMET does not implement a full zero-trust architecture, it is informed by zero-trust thinking, particularly the avoidance of broad implicit trust and the preference for scoped, context-dependent access and communication. This approach supports a flexible and realistic balance between security, performance, and experimental feasibility.

Overall, the application of security-by-design principles in MUSMET does not aim to deliver a complete or production-ready security architecture; however, it does include the concrete implementation and integration of established security and privacy mechanisms within the MUSMET ecosystem to address the risks and requirements identified in the project proposal. These implementations serve as a practical foundation for analysis, validation, and experimentation, rather than as final or market-ready security solutions. If MUSMET-specific security or privacy issues arise that cannot be sufficiently addressed using existing solutions, the project will consider whether additional tools are needed; however, at this stage, no such development is expected.

6.6 Validation, Monitoring, and Risk Awareness during Project Execution

Within MUSMET, validation and monitoring activities play an important role in assessing the practical implications of security and privacy design choices, rather than in enforcing operational security controls. As a research and innovation project, MUSMET does not deploy full-scale intrusion detection or incident response infrastructures. Instead, it adopts a risk-aware and reflective approach, using validation activities to observe how selected security measures interact with musical workflows, system performance, and user experience.

During the implementation of prototypes, pilots, and experimental scenarios, MUSMET monitors security- and privacy-relevant aspects as part of broader system evaluation. This includes observing communication behaviour under different protection configurations, assessing the impact of security mechanisms on latency and synchronization, and identifying unintended data exposure or architectural

dependencies. Monitoring is therefore treated as a means of evidence gathering, supporting informed analysis rather than real-time threat mitigation.

Risk awareness is further supported by the explicit documentation of assumptions, limitations, and residual risks identified during experimentation. Where security or privacy trade-offs emerge, for example between protection strength and real-time performance, these are recorded and analysed rather than masked or optimised away. This approach ensures that MUSMET results remain transparent and interpretable, and that conclusions are grounded in observed system behaviour.

MUSMET also places emphasis on contextual validation, recognising that security and privacy implications may differ between rehearsal scenarios, live performances, and audience-facing demonstrations. Validation activities therefore consider multiple interaction contexts and participant roles, allowing the project to identify patterns and sensitivities that are specific to musical interaction rather than generic system behaviour.

Importantly, MUSMET treats security and privacy as cross-cutting concerns throughout project execution. Insights from monitoring and validation activities are fed back into architectural refinement, configuration choices, and experimental design. This iterative process supports incremental improvement and helps avoid rigid or premature security assumptions.

In summary, MUSMET addresses validation, monitoring, and risk awareness as integral components of responsible system experimentation. While the project does not aim to deliver operational security monitoring solutions, it provides empirical insights into how security and privacy measures behave in real-time MM settings, contributing valuable evidence for future development, deployment, and exploitation activities.

6.7 Summary and Relation to Post-Project Exploitation

This section has outlined how security and privacy are addressed within MUSMET in a scoped and realistic manner, aligned with the project's role as a research and innovation initiative. MUSMET integrates security and privacy considerations into system design, technology selection, and experimental validation, while clearly distinguishing between issues addressed through existing solutions, those explicitly tackled within the project, and those that remain out of scope.

Rather than delivering comprehensive security infrastructures, MUSMET contributes architectural insights, configuration strategies, and validated design choices that reflect the specific constraints of real-time musical interaction in immersive environments. This approach ensures that security and privacy are treated as integral system properties without compromising performance or usability.

The identification of residual risks and out-of-scope challenges is intentionally preserved as a key outcome of the project. Topics such as long-term governance of biometric data, cross-platform intellectual property enforcement, and trust management in open Musical Metaverse ecosystems are highlighted as areas requiring further research and development beyond MUSMET.

In this context, MUSMET's results provide a foundation for post-project exploitation activities, including follow-up research projects, industrial uptake, and potential standardisation efforts. By combining analytical assessment with experimental evidence, MUSMET supports informed decision-making for future secure and privacy-aware Musical Metaverse systems.

References

Ali, M., Naeem, F., Kaddoum, G., & Hossain, E. (2024). Metaverse communications, networking, security, and applications: Research issues, state-of-the-art, and future directions. IEEE Communications Surveys & Tutorials, 26(2), 1238–1305. https://doi.org/10.1109/COMST.2023.3347172

Boem, A., Tomasetti, M., & Turchet, L. (2025). Issues and challenges of audio technologies for the Musical Metaverse. Journal of the Audio Engineering Society, 73(3), 94–114. https://doi.org/10.17743/jaes.2022.0193

Chen, Z., Wu, J., Gan, W., & Qi, Z. (2022). Metaverse security and privacy: An overview. In Proceedings of the IEEE International Conference on Big Data (pp. 2950–2957). IEEE. https://doi.org/10.1109/BigData55660.2022.10021112

Deng, M., Wuyts, K., Scandariato, R. et al. A privacy threat analysis framework: supporting the elicitation and fulfillment of privacy requirements. Requirements Eng 16, 3–32 (2011). https://doi.org/10.1007/s00766-010-0115-7

Dey, K., Barradas, D., & Hakak, S. (2025). A systematic literature review on fundamental technologies and security challenges in Metaverse platforms. https://doi.org/10.48550/arXiv.2510.09621

Howard, M., & Lipner, S. (2006). The Security Development Lifecycle. Microsoft Press. https://doi.org/10.1007/s11623-010-0021-7

Huang, Y., Li, Y. J., & Cai, Z. (2023). Security and privacy in Metaverse: A comprehensive survey. Big Data Mining and Analytics, 6(2), 234–247. https://doi.org/10.26599/BDMA.2022.9020047

Jin, Y., Wang, T., Yang, Q., Shi, L., & Zhang, S. (2025). Zero-knowledge federated learning: A new trustworthy and privacy-preserving distributed learning paradigm. IEEE Communications Magazine. https://doi.org/10.48550/arXiv.2503.15550

Kang, G., Koo, J., & Kim, Y.-G. (2024). Security and privacy requirements for the Metaverse: A Metaverse applications perspective. IEEE Communications Magazine, 62(1), 148–154. https://doi.org/10.1109/MCOM.014.2200620

Otoum, Y., Gottimukkala, N., Kumar, N., & Nayak, A. (2024). Machine learning in Metaverse security: Current solutions and future challenges. ACM Computing Surveys, 56(8), Article 215. https://doi.org/10.1145/3654663

Rahartomo, A., Merino, L., & Ghafari, M. (2025). Metaverse security and privacy research: A systematic review. Computers & Security. https://arxiv.org/abs/2507.14985

Sharma, S., Singh, J., Gupta, A., Ali, F., Khan, F., & Kwak, D. (2024). User safety and security in the Metaverse: A critical review. IEEE Open Journal of the Communications Society, 5, 5467–5486. https://doi.org/10.1109/OJCOMS.2024.3397044

Singh, J., Singh, P., Kaur, R., Kaur, A., & Hedabou, M. (2025). Privacy and security in the Metaverse: Trends, challenges, and future directions. IEEE Access, 13, 120209–120235. https://doi.org/10.1109/ACCESS.2025.3587015

Wang, Y., Su, Z., Zhang, N., Xing, R., Liu, D., Luan, T. H., & Shen, X. (2023). A survey on Metaverse: Fundamentals, security, and privacy. IEEE Communications Surveys & Tutorials, 25(1), 319–352. https://doi.org/10.1109/COMST.2022.3202047

Zhao, R., Zhang, Y., Zhu, Y., Lan, R., & Hua, Z. (2023). Metaverse: Security and privacy concerns. Journal of Metaverse, 3(2), 93–99. https://doi.org/10.57019/jmv.1286526