

NETWORK WORKING GROUP

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: 07 November 2026

C.P. Ross
Independent
07 May 2026

Mercurius Window System (MWS)
draft-ross-mercurius-02

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This document is an individual submission to the IETF.
Distribution of this document is unlimited.

The latest version of this draft can be found at:
<https://mercurius.tebibyte.org/draft-ross-mercurius-02.txt>

Author's Address

Christopher Ross
Independent
Email: chris@tebibyte.org
Project Website: <https://mercurius.tebibyte.org>

Abstract

The Mercurius Window System (MWS) is a networknative, serverside rendering system that enables graphical sessions to be accessed remotely with explicit semantics for windows, input, audio, and session state. MWS allows a user to interact with a workstation from untrusted or resourceconstrained client devices without exposing application data, GPU workloads, or compositor state to those devices. The protocol defines a zerotrust client model, a structured session and window architecture, and distinct planes for rendering, video, audio, and input over an SCTP multistream transport profile. This document specifies the MWS architecture, message formats, transport requirements, and security model.

Executive Summary (NonNormative)

The Mercurius Window System (MWS) is a secure, highperformance window system for environments where applications run on one machine but users work from another. Whether the host is a personal workstation, a shared server, a large deployment, or a private cloud, MWS provides responsive access from any location without assuming that "local" networks or client devices are trustworthy.

MWS is designed for modern mobility patterns: consultants, remote workers, and digital nomads who move between client sites, hotels, airports, and home offices. These users often rely on

lightweight laptops or tablets that may be lost, stolen, or compromised. MWS ensures that possession of a client device is never sufficient to access the workstation. Authentication requires explicit user presence, and no sensitive data, credentials, or GPU workloads reside on the client.

MWS preserves the network transparency that made X11 valuable, whilst replacing its implicit trust and CPUbound rendering with a modern, zerotrust, GPUaccelerated architecture. The result is a window system that feels local even when the GPU is in another room, another building, or another country, supporting workflows ranging from office productivity to highrefreshrate interactive applications on modern networks.

Table of Contents

1. Introduction
 - 1.1. Scope and Applicability
 - 1.2. Design Rationale
 - 1.3. Cloud and Distributed Computing Context
 - 1.4. HighPerformance Rendering and Gaming
 - 1.5. Terminal Requirements and Wireless Considerations
 - 1.6. Session Mobility and Detachable Operation
2. Conventions Used in This Document
3. System Architecture
 - 3.1. Architectural Principles
 - 3.2. Major Components
 - 3.3. WorkstationCentric Model
 - 3.4. Rendering and Surface Model
 - 3.5. Session Model
 - 3.6. ZeroTrust Client Model
 - 3.7. Network Considerations
 - 3.8. Transport Requirements
4. Detailed Architecture
 - 4.1. Sessions
 - 4.2. Seats
 - 4.3. Windows
 - 4.4. Compositor Model
 - 4.5. Rendering Model
 - 4.6. Audio Model
 - 4.7. Stream Allocation
 - 4.7.1. Session Identity and Message Routing
 - 4.8. Session Lifecycle
 - 4.9. Session and Seat Model
 - 4.10. Local Transport Profile (NonNormative)
 - 4.11. Security Model
5. Protocol Specification
 - 5.1. Message Framing
 - 5.2. Control Messages (Stream 0)
 - 5.2.1. Initial Handshake (001099)
 - 5.2.1.1. Session Identifier Semantics
 - 5.2.2. Session Management (100199)
 - 5.2.2.1. Resume Semantics
 - 5.2.3. Window Lifecycle (200299)
 - 5.2.3.1. Window Identifier Scope
 - 5.3. Rendering Commands (300399) — Stream 1
 - 5.4. Input Events (400499) — Stream 2
 - 5.4.1. Pointer Motion Events
 - 5.4.2. Input Scoping and Session Isolation
 - 5.5. Video Fallback (500599) — Stream 3
 - 5.6. Audio Plane (600699) — Stream 4

- 5.7. Protocol State Machine
 - 5.7.1. Initial Connection
 - 5.7.2. Session Resume
- 5.8. WSI Extension (Surface Creation)
 - 5.8.1. Surface Binding and Session Validation
 - 5.8.2. Required Enums
- 5.9. Error Handling (700799)
 - 5.9.1. Session and Resource Validation Errors
- 6. Reference Implementation
 - 6.1. Server Components
 - 6.2. Client Components
 - 6.3. Demonstration Clients
 - 6.3.1. Vulkan Demonstration Clients
 - 6.3.2. Simple Game Demonstration
 - 6.3.3. Desktop Environment Support
 - 6.4. Dependencies
 - 6.5. Bootstrap Example
- 7. Implementation Requirements and Validation
 - 7.1. Test Matrix
 - 7.1.1. Core Validation Tests
 - 7.1.2. Reference Implementation Commands (NonNormative)
 - 7.2. GPU Isolation Requirements
 - 7.3. Bandwidth and Transport Isolation Requirements
- 8. Performance Considerations
- 9. Security Considerations
 - 9.1. DANE Deployment (NonNormative)
- 10. IANA Considerations
- 11. Acknowledgements
- 12. References
 - 12.1. Normative References
 - 12.2. Informative References
- 13. Copyright
- Appendix A. MWS Opcode Registry
 - A.1. Handshake and Authentication (000099)
 - A.2. Session Management (100199)
 - A.3. Window Lifecycle (200299)
 - A.4. Rendering Commands (300399)
 - A.5. Input Events (400499)
 - A.6. Video Plane (500599)
 - A.7. Audio Plane (600699)
 - A.8. Error Reporting (700799)
 - A.9. Reserved for Future Extensions (800899)
 - A.10. Experimental and VendorSpecific (900999)
- Appendix B. Authentication Mechanism Registry
 - B.1. Standard Mechanisms
 - B.2. Extensible Mechanisms
 - B.3. Private and Experimental Mechanisms
 - B.4. Registration Policy
- Appendix C. SCTP Stream Usage Summary
 - C.1. Stream 0 — Control Plane
 - C.2. Stream 1 — Rendering Commands
 - C.3. Stream 2 — Input Events
 - C.4. Stream 3 — Video Plane
 - C.5. Stream 4 — Audio Plane
 - C.6. Additional Streams

Appendix D. Protocol State Machine Diagrams

- D.1. Initial Connection State Machine
- D.2. Session Resume State Machine
- D.3. Error Handling State Machine
- D.4. Stream Interaction Summary

13. Copyright

1. Introduction

The Mercurius Window System (MWS), named for Mercurius, the Roman messenger god of swift communication, is a secure window system for both local and remote use. A user may work directly at the console of a workstation as on a conventional Unixlike desktop, with full access to its GPU, input devices, audio devices, and local display. The same session may also be accessed from lightweight, mobile, or untrusted client devices elsewhere, without replicating the workstation's software environment or exposing its data or GPU resources. Compute, storage, rendering, and audio processing remain on the workstation; clients act solely as authenticated display and input/audio endpoints.

MWS is intended to let a workstation remain itself while being reached from elsewhere. The workstation is treated as a longlived environment that accumulates tools, history, and identity; remote devices are simply places from which the user inhabits that environment. A user may begin work at a powerful machine in the office and later continue the same session from a laptop, thin client, or secondary desktop in another location, without maintaining multiple environments or synchronising state. Remote access is an extension of the local workstation rather than a separate mode of operation.

MWS is not a local display protocol, nor a pixelstreaming system. It defines a structured, messageoriented protocol for presence, input, audio, and rendering state on a workstation. Rendering is serverresident and GPUaccelerated; audio capture and playback are explicitly negotiated streams; and the protocol transmits structured commands and media units rather than raw framebuffers. The transport is agnostic but optimised for SCTP's multistream, messageoriented semantics [RFC9260], with separate streams for control, rendering commands, input, video, and audio.

This architecture continues the lineage of early Unix window systems such as X11, which supported networktransparent interaction with applications running on central servers, while applying modern zerotrust security [NIST800207], authenticated multistream transport [RFC9260][RFC4895], and GPU isolation. Earlier systems such as NeWS explored serverside rendering but lacked the transport and security mechanisms required for contemporary workloads. Wayland, by contrast, is intentionally scoped to trusted local compositing and does not address remote GPUs, untrusted clients, or relocatable sessions. MWS occupies a distinct design space: secure, zerotrust remote presence for serverresident graphical environments, without compromising firstclass local console use.

1.1. Scope and Applicability

This document specifies the Mercurius Window System (MWS) protocol, the transportlevel protocol used by MWS to establish, authenticate, and maintain a user's graphical presence on a workstation. The MWS protocol defines message framing, capability negotiation, session attachment, and input/output semantics,

including structured handling of video and multichannel audio, over a secure transport profile based on TLS 1.3 [RFC8446] and SCTP [RFC9260].

MWS is intended for environments where:

- applications execute on a central workstation or server
- users may work locally at the console or remotely from other devices
- clients may be untrusted, mobile, or ephemeral
- users may relocate sessions across devices
- GPUaccelerated workloads must remain serverresident
- audio capture and playback must remain serverresident or explicitly brokered
- network transparency is a firstclass requirement
- loss or theft of a device must not compromise workstation security.

MWS does not replace local display protocols such as Wayland, nor does it extend them. It provides a complementary mechanism for secure local and remote presence in multiuser and distributed environments where local display protocols do not apply.

1.2. Design Rationale

Early Unix window systems, including X11, were explicitly designed for network transparency: applications executed on powerful central servers while users interacted from remote terminals. This model proved valuable in multiuser and distributed environments, but X11' s permissive trust model, unrestricted client capabilities, and CPUbound rendering are incompatible with modern zerotrust requirements and GPUaccelerated workloads [NIST800207].

Wayland addresses these issues by assuming a single trusted local compositor, a local GPU, and a singleuser environment. This model provides excellent performance on personal workstations but does not support remote GPUs, relocatable sessions, multiuser deployments, or untrusted clients. These use cases lie outside Wayland' s design goals.

MWS intentionally revives and modernises the network transparent workstation model. It retains the architectural advantages of centralised execution and remote interaction while adopting a zerotrust security model based on mutual TLS [RFC8446], authenticated SCTP [RFC9260][RFC4895] streams, and perclient GPU isolation. Rendering is serverresident and GPUaccelerated; clients transmit structured rendering commands rather than framebuffers, and audio is carried as explicit timestamped streams rather than devicelocal side effects. All compositor policy, input routing, and window management occur on the server, ensuring multiuser correctness and preventing privilege escalation.

Crucially, possession of a client device is never sufficient to access the workstation. Client certificates must be protected by the platform, session resume requires fresh authentication, and no sensitive data or credentials are stored on the client. This ensures that a lost or stolen laptop cannot be used to compromise the workstation.

The result is a window system that provides deterministic semantics, strong isolation, and relocatable sessions, enabling users to inhabit remote workstations with the performance and responsiveness of a local environment, whilst preserving

firstclass local console operation.

1.3. Cloud and Distributed Computing Context

Many organisations operate private cloud or workstationcluster environments where users access centralised compute and GPU resources from thin clients or mobile devices. Public cloud deployments exhibit similar characteristics: applications execute on remote servers while clients roam across untrusted networks.

MWS aligns with this model by centralising execution and distributing only the user interface. This avoids the inefficiencies of distributed compute systems whilst preserving the benefits of remote access, session mobility, and strong isolation between users. Because clients are untrusted, MWS ensures that compromise or loss of a client device does not grant access to the workstation.

1.4. HighPerformance Rendering and Gaming

MWS is primarily intended for workstation and privatecloud deployments in which terminals connect over wellprovisioned LANs and VPNs, typically with DANE [RFC6698][RFC7671] securing mutually authenticated transport over TLS 1.3 [RFC8446] carried over SCTP [RFC3436]. In these environments, modern GPUs provide hardwareaccelerated AV1 encoding with extremely low latency, enabling highresolution and highrefreshrate streaming for everyday workstation workloads. All workstation rendering in MWS is performed using the Vulkan [VK14] API.

Nevertheless, the same architecture could accommodate high performance remote rendering workloads, including interactive 3D applications and games. Support for such workloads is a stretch goal rather than a primary target, but these use cases inform the design of the transport, security, and rendering model to ensure that MWS remains viable for demanding graphical applications.

1.5. Terminal Requirements and Wireless Considerations

MWS clients are treated as untrusted endpoints. Practical deployments assume a minimum level of capability. A typical terminal is expected to provide a modern CPU, a hardware accelerated GPU capable of AV1 decoding, and at least gigabit class network connectivity. Higher resolutions or refresh rates benefit from greater bandwidth, but MWS remains usable at reduced quality on lowercapacity links.

Emerging wireless standards such as WiFi 8 (IEEE 802.11bn) are expected to deliver multigigabit throughput and submillisecond airinterface latency, enabling MWS terminals to operate over wireless links without compromising interactive performance.

1.6. Session Mobility and Detachable Operation

Because sessions in MWS are serverresident and independent of client connections, the system naturally supports detachable operation. A user may disconnect from one terminal and later resume the same session from another device, with all windows, GPU state, and compositor context preserved.

This model is conceptually similar to terminal multiplexers such as screen or tmux, but applied to a full GPUaccelerated graphical environment. Session mobility is a core design goal of MWS and informs its authentication, transport, and rendering architecture.

Deployments are expected to configure a reconnection grace period so that brief network outages or short unscheduled breaks such as to move or charge a client device do not cause the user's session to be lost.

2. Conventions Used in This Document

The key words MUST, MUST NOT, REQUIRED, SHALL, SHALL NOT, SHOULD, SHOULD NOT, RECOMMENDED, NOT RECOMMENDED, MAY, and OPTIONAL in this document are to be interpreted as described in RFC 2119 and RFC 8174 when, and only when, they appear in all capitals.

Terminology relating to Vulkan follows the definitions and naming conventions of the Vulkan 1.4 specification [VK14].

Terminology relating to SCTP follows the Stream Control Transmission Protocol specification [RFC9260]. Terminology relating to TLS 1.3 follows the Transport Layer Security specification [RFC8446].

Unless otherwise stated:

“workstation” refers to the system on which applications execute, and where all rendering, compositing, and session management occur.

“server” refers to the MWS server daemon (mwsd) running on the workstation. In this document, “workstation” and “server” refer to the same system at different levels of abstraction.

“client device” refers to the physical device used by a user to access the workstation. Examples include laptops, tablets, phones, thin clients, and embedded devices.

“client” refers to the MWS client component (mwsc) running on the client device. The client is the protocol endpoint that communicates with the server.

“terminal” refers to the logical role a client assumes once connected: a seatproviding endpoint that delivers input and receives rendered output.

“session” refers to a persistent graphical environment maintained on the workstation independently of client connections.

“seat” refers to a set of input devices and output mappings associated with a session and bound to a terminal.

“surface” refers to a drawable region managed by the compositor and rendered by the workstation's GPU.

“command stream” refers to the structured, Mercurius protocol messages exchanged on SCTP stream 0 (control stream).

“video surface” refers to a highmotion region encoded using a hardwareaccelerated codec such as AV1.

The names Alice, Bob, Eve, and Mallory are used in their standard roles from security literature. Alice and Bob denote honest users, Eve denotes a passive eavesdropper, and Mallory denotes an active attacker. These names are used solely for threatmodel examples and do not correspond to real users or implementation artefacts.

All multibyte integers are transmitted in network byte order unless explicitly specified otherwise.

3. System Architecture

This section provides a highlevel overview of the Mercurius Window System (MWS). It describes the conceptual model, major components, and architectural principles that inform the detailed design in Section 4 and the protocol specification in Section 5.

MWS is designed around a workstationcentric model in which all rendering, compositing, audio processing, session management, and windowmanagement policy reside on a central server. Client devices act solely as authenticated display, audio, and input endpoints. This model preserves the semantics of a local workstation while enabling secure remote presence across modern networks.

MWS assumes client devices with at least gigabitclass connectivity, including modern WiFi networks that routinely exceed 1Gb/s. The protocol is optimised for 10GbE LANs, where uncompressed or lightly compressed surfaces, highmotion content, and lowlatency audio can be delivered with minimal delay. Devices with substantially lower bandwidth may operate at reduced quality but are not a primary design target.

MWS also provides a realtime audio transport suitable for workstationclass media workloads. The protocol supports fullduplex, timestamped PCM audio with deterministic latency, enabling remote use of digital audio workstations, conferencing applications, and highfidelity monitoring. Audio transport is integrated into the same multistream SCTP association as rendering, input, and video, but operates on dedicated streams to ensure isolation from congestion and headofline blocking.

3.1. Architectural Principles

The design of MWS is guided by the following principles:

- Applications execute on a central workstation or server.
- Local and remote interaction share identical session semantics.
- Client devices may be untrusted, mobile, or ephemeral.
- Users may relocate sessions across devices without restarting applications.
- GPUaccelerated workloads remain serverresident.
- Realtime audio is a firstclass subsystem with strict latency and ordering requirements.
- Network transparency is a firstclass requirement.
- Loss or theft of a client device must not compromise workstation security.

These principles reflect the goal of treating the workstation as a longlived environment with continuity of storage, configuration, and identity.

Alternative transports such as QUIC were considered. However, SCTP' s native multistreaming, messageoriented delivery, and support for partial reliability align directly with the requirements of MWS. QUIC' s multiplexed bytestream model, together with the absence of partially reliable streams, would require additional framing and scheduling logic to emulate SCTP semantics. For these reasons, SCTP is the primary transport for MWS.

The protocol' s guarantees depend on transport properties that SCTP provides natively, including independent ordered streams, preservation of message boundaries, optional partial reliability, avoidance of crossstream headofline blocking, and stable associations. These properties are required to ensure deterministic compositor behaviour, responsive input under load,

support for highmotion video surfaces, lowlatency audio transport, relocatable sessions, and multiseat concurrency.

TCP does not provide these properties without substantial additional protocol machinery. A TCPbased transport would therefore be unable to meet the latency, isolation, and concurrency requirements of MWS as defined in this document, and is out of scope for this specification.

3.2. Major Components

MWS consists of the following major components, corresponding to the object model defined in Appendix A:

The compositor, which manages windows, surfaces, focus, input routing, and presentation.

The renderer, which executes GPUaccelerated drawing and performs surface composition on the workstation.

The audio subsystem, which manages audio device and audio stream objects, and provides fullduplex, timestamped PCM transport for playback and capture.

The session manager, which maintains Session objects independently of client connections and supports detachable operation.

The transport layer, which provides a secure, multistream, messageoriented channel between workstation and client.

The input subsystem, which manages Seat and InputDevice objects and forwards input events to the compositor.

The client device, which decodes surfaces, presents them to the user, plays audio, and forwards input events to the workstation.

These components interact to provide a deterministic, structured, zerotrust window system suitable for both local console use and remote presence.

3.3. WorkstationCentric Model

The workstation is the authoritative environment. It owns all GPU resources, audio devices, compositor state, and inputrouting policy. Applications run exclusively on the workstation, and all rendering and audio processing are performed on workstationresident hardware.

Client devices own their physical input hardware. Input events are generated on the client and forwarded to the workstation, which applies focus, routing, and seat semantics. The workstation never interacts with the client' s physical devices directly; it operates only on the logical input events they produce.

GPU resources are exposed to MWS exclusively through the Vulkan API. The workstation enumerates all available GPUs using `vkEnumeratePhysicalDevices()` and creates one compositor instance per physical device. All rendering, composition, and presentation operations are defined in terms of Vulkan objects and capabilities.

Audio resources are exposed to MWS through audio devices, each of which represents a physical or virtual playback or capture endpoint. Audio streams are created dynamically to transport PCM audio between the workstation and client devices.

A user may interact with the workstation in two ways:

Local console mode, using the workstation's own keyboard, pointer, display, and audio hardware.

Remote presence mode, using an authenticated client device elsewhere on the network.

Local and remote interaction share the same compositor, window tree, audio devices, and session state. Remote presence is an extension of the local workstation, not a separate mode of operation. Sessions persist across transient disconnections, but longterm persistence requires explicit detachment.

3.4. Rendering and Surface Model

MWS supports two classes of graphical output:

structured rendering commands, representing deterministic drawing operations for lowmotion or vectororiented surfaces

video surfaces, representing highmotion content encoded using hardwareaccelerated codecs such as AV1

The compositor selects the appropriate representation based on surface characteristics and available bandwidth. This allows MWS to operate efficiently on both 1GbE and 10GbE networks, as well as on modern WiFi links.

Surface objects represent drawable regions within the compositor. Window objects reference one or more surfaces, and the compositor determines whether a surface is transmitted as structured commands or as a video stream. The renderer produces Vulkan command streams for structured surfaces, while video surfaces are encoded using hardware acceleration when available.

3.5. Session Model

Sessions are serverresident and persist independently of client connections, but longterm persistence requires explicit detachment. MWS tolerates transient network interruptions; if a client reconnects within the configured grace period, the session continues without interruption. If a client disappears without detaching, the session is preserved only for the duration of this grace period. Once the period expires, the session is closed, and applications terminate in the same manner as a workstation session without an active seat. Because all application, window, and session state resides on the workstation, failure, loss, or destruction of a client device does not by itself risk loss of inprogress work; the session remains intact on the server and may be resumed from another device, subject only to the reconnection grace period.

A user may explicitly detach a session to preserve it beyond the reconnection grace period and may later resume it from any authorised device. A single transport association may provide multiple seats for a session when permitted by policy. The precise rules governing how many transport associations may attach to a session, and under what conditions, are defined in Section 4.7.

Sessions contain seats, each of which aggregates input devices, audio streams, and presentation state for a particular user interaction context. Multiple seats may be active concurrently, enabling multiuser or multiterminal operation.

This model enables mobility across devices while preserving the semantics of a traditional workstation and avoiding longlived orphaned sessions.

3.6. ZeroTrust Client Model

All client devices are treated as untrusted endpoints, even on local LANs. Trust is established exclusively through cryptographic identity and explicit authorisation rather than network location. A client device is not an identity and is not authorised to access a session by virtue of its presence on the network; only the user is authorised.

Historically, X11 enabled a powerful and flexible model in which users could inhabit remote workstations as naturally as local ones. Even on dedicated X terminals, users authenticated as themselves and could not access another user's data. The flaw was not the login model but the assumption that any client on the network was inherently trustworthy. Wayland addressed this by eliminating remote clients entirely. Mercurius instead removes the assumption of trust: client devices may be anywhere, but trust is derived solely from what the user can cryptographically prove.

Device identity is established during the TLS 1.3 handshake carried over SCTP. Deployments SHOULD use DNSBased Authentication of Named Entities (DANE) [RFC6698][RFC7671] to bind the server's certificate to DNSSECprotected TLSA records, allowing clients to verify that they are communicating with the correct workstation without relying on public certificate authorities or assumptions about local network topology. Typically only the server requires a TLSA record; client devices may obtain certificates through local provisioning mechanisms. DANE prevents an attacker on the same network from impersonating a Mercurius server, without depending on public CA infrastructure.

User authentication is performed at the application layer using the mechanismagnostic model defined in Section 5.2.1. The server advertises supported mechanisms (for example, "PAM", "FIDO2"), and the client selects one. This allows deployments to integrate passwordbased, hardwaretoken, federated, or certificatebased user authentication without modifying the protocol.

A client device is assumed to be mobile and at risk of loss or theft. To limit the impact of device compromise, a client stores no confidential data, longterm session state, or reusable credentials. Because all application and session state resides exclusively on the workstation, loss or compromise of a client device does not expose user data or permit modification of workstationresident files. Authentication requires both possession of the device's private key (validated via mutual TLS and, when deployed, DANE) and successful user authentication via one of the advertised mechanisms. Mallory stealing Bob's laptop gives him no more access to the workstation than if he had purchased a brandnew laptop; the device alone is insufficient to access or resume a session.

In contrast to traditional window systems such as X11, client devices in MWS are not part of the trusted computing base; they are merely display, audio, and input conduits for a workstationresident session.

3.7. Network Considerations

MWS is designed to operate over untrusted IP networks, including public networks and variable-quality wireless links. The protocol does not assume that terminals are located on the same LAN as the workstation, nor that any network segment provides meaningful security. A terminal on a local LAN and a terminal on a remote network are treated identically by the workstation.

MWS is designed for modern networks:

10GbE provides optimal performance and headroom for multiple high-resolution seats on a single workstation.

WiFi 6/6E/7 provides multi-gigabit throughput with variable jitter and is fully supported for single-seat clients.

1GbE provides a usable baseline for typical desktop workloads and a small number of seats on a workstation.

Sub-gigabit links are outside the primary design envelope and are not expected to provide an acceptable experience for high-resolution, high-refresh workloads or low-latency audio.

The transport layer adapts to available bandwidth through dynamic surface encoding, selective use of video surfaces, adaptive refresh rates, and prioritised input, audio, and control streams.

The detailed architecture is specified in Section 4, and the wire protocol is defined in Section 5.

3.8. Transport Requirements

MWS requires a transport that provides structured, message-oriented delivery with support for multiple independently ordered channels. The transport **MUST** preserve message boundaries, **MUST** support concurrent streams with independent ordering, and **SHOULD** provide mechanisms for partial reliability to avoid retransmission of stale high-volume data such as video surfaces and real-time audio frames.

The transport **MUST** avoid cross-stream head-of-line blocking. Input events, control messages, rendering commands, audio streams, and video surfaces are logically independent flows, and the correctness of compositor behaviour depends on their timely and ordered delivery within their respective channels. A transport that enforces global ordering across all data would introduce latency coupling between these flows and would not meet the responsiveness requirements of MWS.

The transport **MUST** support stable associations that survive transient network changes, including client mobility across networks. Session attachment, reconnection semantics, and multi-seat operation rely on the ability to maintain a consistent transport-level association identity.

SCTP satisfies these requirements through its native multi-streaming model, message-oriented delivery, optional partial reliability, and support for multi-homing. These properties align directly with the architectural principles defined in Section 3.1 and are required for deterministic compositor behaviour, responsive input under load, support for high-motion video surfaces, low-latency audio transport, relocatable sessions, and multi-seat concurrency.

TCP does not provide these properties without substantial additional protocol machinery. TCP offers only a single in-order byte stream, lacks message boundaries, enforces global

headofline blocking, and provides no support for partial reliability or multistreaming. A TCPbased transport would therefore be unable to meet the latency, isolation, and concurrency requirements of MWS as defined in this document, and is out of scope for this specification.

4. Detailed Architecture

The Mercurius Window System (MWS) is structured around a central server (mwsd) that owns all GPU resources, audio devices, input routing, and compositor state, and a set of untrusted client devices that connect over a secure, messageoriented transport. Once connected, a client device acts as a terminal providing a seat. This section describes the architectural model of sessions, seats, windows, rendering, and compositor behaviour. The wire protocol and message formats are defined in Section 5.

4.1. Sessions

A session represents the complete graphical environment associated with a single authenticated user, including windows, workspaces, GPU resources, audio devices, and compositor state. Sessions are serverresident and MAY persist independently of client connections when explicitly detached.

User identity is established during the secure transport handshake. The authenticated identity (for example, the client certificate subject or a federated identity token) is mapped to a local user account via the system's authentication framework (such as PAM). All windows, seats, and compositor state created over that association belong to the resulting session.

4.2. Seats

A seat represents a set of input devices, audio streams, and an output binding for a session. A session MAY have multiple seats simultaneously. Each seat corresponds to a particular terminal, whether that terminal is the local console or a remote client device acting in the terminal role.

Input events are tagged with a seat_id, and the compositor routes them according to seatspecific focus and pointer state. Output mappings (for example, which windows appear on which displays) and audio routing may also be seatspecific.

4.3. Windows

Windows are servermanaged objects representing toplevel application surfaces. Each window belongs to exactly one session and is associated with one or more rendering surfaces (structured swapchains or video surfaces) depending on compositor policy.

Window identifiers are scoped to a session. A terminal MUST NOT reference or interact with windows belonging to any other session. The server MUST enforce this isolation and MUST reject or ignore any protocol message that attempts to target a window outside the authenticated session.

4.4. Compositor Model

The compositor maintains the global window tree, stacking order, focus, workspaces, and output mappings for each session. It is responsible for:

- applying windowmanagement policy

- routing input events based on seat and focus
- managing swapchains and presentation timing
- selecting between structured rendering and video fallback
- revoking or reconfiguring windows according to policy

The compositor SHOULD expose a uservisible mechanism to forcibly terminate an unresponsive window. This mechanism is implementation defined (for example, a “kill window” gesture similar to CtrlAltEsc in KDE).

The compositor MAY revoke swapchains, reconfigure windows, or migrate them between outputs according to local policy, resource constraints, or security requirements. When a swapchain is revoked, the server notifies the terminal and MAY substitute a placeholder or video surface.

Each compositor instance is bound to a single GPU or output pipeline. It owns the devicelevel rendering resources for that GPU (device context, queues, swapchains, and associated GPU buffers), and no other component may submit rendering work directly to that device.

4.5. Rendering Model

Rendering in MWS is serverside. Applications submit rendering commands to the server, which validates and executes them on the GPU. Client devices do not access GPU resources directly.

The compositor selects the appropriate representation for each surface:

- structured rendering for lowmotion or interactive content
- video surfaces for highmotion or bandwidthsensitive content

Because rendering is serverresident, a stalled or misbehaving terminal cannot block the compositor. The server MAY revoke a window’s rendering resources, substitute a placeholder surface, or terminate the client if rendering deadlines are repeatedly missed.

The MWS specification assumes a modern explicit GPU API for rendering and composition (for example, Vulkan [VK14]) and requires that all rendering and presentation operations be performed through the compositor’s devicelevel abstraction. Client devices are not required to implement any graphics API.

4.6. Audio Model

The audio subsystem manages audio devices and audio streams and provides fullduplex, timestamped PCM transport between the workstation and client devices. Audio is treated as a firstclass subsystem with strict latency and ordering requirements; audio traffic is logically independent of rendering and control traffic but shares the same secure, multistream transport.

An audio device represents a physical or virtual playback or capture endpoint on the workstation, such as speakers, headphones, microphones, instrument inputs, multichannel mixers, loopback devices, and virtual sinks. Devices are enumerated and managed on the server; client devices do not own or configure audio hardware directly.

Logical audio streams are created dynamically to carry PCM samples between the workstation and the client. Each audio stream is bound to a specific audio device and seat, and is direction

specific (playback or capture). Streams are timestamped at the server, and the client maintains playout buffers that honour these timestamps while minimising latency and jitter.

Playback streams carry audio from applications on the workstation to the client device for presentation. Capture streams carry audio from clientattached input devices to the workstation, where it is injected into the appropriate session and applications according to policy. The server MAY apply policy to limit or redirect capture streams (for example, to prevent inadvertent capture in shared environments, or to restrict which multichannel devices are exposed to a given session).

In the base profile, audio messages use a dedicated SCTP stream (stream 4), separate from control, rendering, input, and the Video Plane, to avoid headofline blocking and to ensure predictable latency. Implementations MAY allocate additional SCTP streams for audio as an optimisation, but MUST NOT multiplex audio opcodes with control, input, or video opcodes on the same SCTP stream. The wirelevel definition of the Audio Plane, including opcodes and negotiation of playback and capture streams, is specified in Section 5.6.

4.7. Stream Allocation

MWS uses SCTP multistreaming to separate control traffic from rendering traffic and to prevent headofline blocking between independent windows. Stream allocation is defined as follows:

Stream 0 is reserved for control messages and MUST NOT carry rendering data.

Each window is associated with exactly one rendering stream. All structured rendering commands and videosurface updates for that window are sent on its assigned stream.

Multimonitor configurations do not affect stream allocation. A window that spans multiple outputs continues to use a single rendering stream.

The compositor MAY allocate additional streams for specialised rendering contexts (for example, offscreen surfaces or auxiliary swapchains), but these MUST be explicitly negotiated during window creation and are scoped to the window that requested them.

Streams are not reused across windows unless the compositor has explicitly revoked the prior window and returned the stream to the allocator.

This model ensures that rendering for one window cannot block or delay rendering for another, while avoiding unnecessary proliferation of streams in multimonitor environments.

4.7.1. Session Identity and Message Routing

Each SCTP association corresponds to exactly one client session. The server MUST treat the SCTP association identifier (assoc_id) as the authoritative session identity. No clientsupplied field may select, reference, or influence the session to which a message is delivered.

A session is created only after successful completion of the handshake defined in Section 5. Until the handshake completes, the server MUST ignore all messages received on streams other than 0.

For any message received on a nonzero stream, the server MUST:

- identify the session associated with the SCTP association

- verify that the session is active

- dispatch the message to the subsystem corresponding to the stream

- reject or ignore the message if it is malformed or references resources outside the session

Messages referencing windows, seats, or other resources not owned by the session MUST be rejected with `MWS_ERROR(type=700, fatal=0)`.

Messages referencing a session other than the one implied by the SCTP association MUST be ignored.

If a message is received for an association that has no active handshake or no active session, or whose session has been closed, the server MUST silently discard the message.

4.8. Session Lifecycle

A session is a longlived serverside construct that persists independently of any particular network connection. A session becomes ACTIVE when a client completes the handshake defined in Section 5 and remains ACTIVE until it is explicitly terminated or reclaimed by policy.

A new client device connection MUST create a new session in the ACTIVE state unless the user explicitly requests to resume an existing session. The server MUST NOT automatically reattach a client device to a prior session solely on the basis of matching user identity.

A client MAY explicitly detach from an ACTIVE session. Detach transitions the session from ACTIVE to DETACHED. In the DETACHED state, the session continues to run without any attached transport association; its windows, compositor state, audio streams, and GPU resources remain serverresident. DETACHED sessions MAY be resumed by any authenticated client device belonging to the same user, subject to server policy.

Loss of the SCTP association (for example, network failure, timeout, client crash) while a session is ACTIVE does not immediately terminate the session. Instead, the server MUST transition the session to a GRACE state and start a reconnection grace timer. In the GRACE state, the session remains active but has no attached client. If a client reconnects and successfully resumes the session before the grace timer expires, the server MUST transition the session back to ACTIVE and the session continues without loss of state.

If the reconnection grace period expires without a successful resume, the server MUST treat the session as ABANDONED unless the user has explicitly detached it. ABANDONED sessions MUST be terminated and all associated resources reclaimed. Implementations MUST provide a configurable reconnection grace interval and SHOULD allow values sufficient to tolerate brief network outages on typical WiFi and WAN links. Servers SHOULD return a specific error status when a resume request targets an expired session.

Longterm persistence is an explicit, optin behaviour: a session

continues to exist beyond the reconnection grace period only if the user has explicitly detached it or otherwise marked it for later resumption. Implementations MUST NOT retain ABANDONED or stale sessions indefinitely. The server SHOULD reclaim resources associated with inactive sessions according to local policy (for example, idle timeout, logout event, or administrative limits).

Only one transport association (client instance) MUST be attached to a given session at a time in the base protocol. That association MAY provide one or more seats for the session, subject to server policy. An implementation MAY provide a mechanism that allows additional associations to attach to the same session (for example, for technical support), but such behaviour is outside the scope of this specification and MUST NOT alter the semantics defined for the singleassociation model above.

4.9. Session and Seat Model

Sessions MAY persist independently of client connections. When a client device disconnects, the associated session and its windows MAY remain active in either the GRACE or DETACHED state. The compositor MAY blank or lock the session's outputs according to local policy while no seat is attached.

When a user resumes a session (from GRACE or DETACHED), the server:

1. Binds a new seat_id to the resumed session.
2. Sends the current window list, geometry, and focus state.
3. Associates the new seat's outputs and audio routing with the session.

MWS supports both independent sessions and multiseat attachment within a single session. A user may maintain multiple concurrent sessions (for example, one on the local console and another accessed remotely), or may attach multiple seats to the same session via a single transport association, subject to the singleassociation rule in Section 4.7.

MWS also supports explicit session detachment. A user may detach a running session, leaving its windows, compositor state, audio streams, and GPU resources active on the server without any attached seats. The user may then initiate a new session on the same client device (for example, to perform unrelated work) and later resume the detached session exactly where it was left. This behaviour is directly analogous to detaching and reattaching a GNU Screen or tmux session, but applied to a full graphical desktop environment spanning one or more seats.

4.10. Local Transport Profile (NonNormative)

Although MWS treats all client devices as untrusted endpoints and applies the same protocol semantics regardless of network location, implementations MAY apply transportlayer optimisations when the client device and server reside on the same physical host. These optimisations MUST NOT alter protocol semantics, message ordering, authentication requirements, or session isolation, and MUST remain transparent to the terminal.

Permitted implementationlevel optimisations include:

- loopbackspecific SCTP acceleration
- reduced cryptographic overhead
- sharedmemory fast paths

GPUDirect resource sharing where supported

These optimisations MUST NOT:

- grant additional privileges to local client devices
- bypass certificate validation or user authentication
- modify the behaviour of control, input, audio, or rendering streams
- introduce protocol features unavailable to remote client devices

MWS remains a transportagnostic, networktransparent window system. Local optimisations exist solely to ensure that client devices running on the same host as the server achieve performance comparable to traditional localonly systems without compromising the zerotrust security model.

4.11. Security Model

All client devices are treated as untrusted. Trust is established exclusively through cryptographic identity and explicit authorisation rather than network location. The server enforces strict isolation between users, sessions, seats, and windows. In particular:

- each session is bound to a single SCTP association, and the association identifier serves as the authoritative session identity

- window identifiers are scoped to a session and cannot be referenced by other sessions

- input events are scoped to a seat and session, and cannot target windows outside that session

- clients cannot observe, enumerate, or reference resources belonging to other sessions

- all clientoriginated messages are validated before being processed

Transport security is provided by mutually authenticated TLS 1.3 carried over SCTP. Deployments SHOULD use DNSBased Authentication of Named Entities (DANE) to bind the server's certificate to DNSSECprotected TLSA records, allowing clients to verify that they are communicating with the correct workstation even in the presence of compromised or misissued CA certificates. Device authentication alone does not grant access to a user session.

MWS does not mandate a specific encryption mechanism beyond requiring confidentiality, integrity, and mutual authentication of endpoints. Deployments SHOULD use a transport that provides forward secrecy, such as TLS 1.3, WireGuard, or IPsec with PFSenabled cipher suites. The choice of transportlayer security does not affect protocol semantics, and MWS remains agnostic to whether encryption is provided directly by TLS or by an external secure tunnel.

User authentication is performed at the application layer using the mechanismagnostic model defined in Section 5.2.1. The server advertises supported mechanisms (for example, "PAM" and "FIDO2"), and the client selects one. This separation of device and user identity ensures that compromise of a device does not grant access to a user's session without the corresponding user credential.

The server validates all clientoriginated messages, including

window creation, input events, and rendering commands. A client may not reference windows, sessions, or resources outside its authenticated session. Attempts to do so are rejected with `MWS_ERROR(type=700, fatal=0)`. Malformed or semantically invalid messages are ignored, and the session continues unless the error is marked fatal.

The client is not part of the trusted computing base. It stores no confidential data, longterm session state, or reusable credentials. If a client disconnects unexpectedly, the session persists only for the duration of the reconnection grace period unless the user has explicitly detached. After this period, the session is closed and applications terminate.

Loss of the transport association for any reason (network failure, timeout, endpoint crash) is treated as a fatal transport error. The session MAY persist according to the rules in Section 4.7 and MAY be resumed from another client device subject to policy.

5. Protocol Specification

5.1. Message Framing

All MWS messages consist of a fixed-size header followed by an optional payload. Messages are carried within a single SCTP user message and MUST NOT be fragmented across multiple SCTP user messages.

The header format is:

```
struct MwsHeader {
    uint32_t magic;           // MWS_MAGIC_VALUE
    uint16_t type;           // MWS_* opcode
    uint16_t reserved;       // MUST be zero
    uint32_t length;         // payload length in bytes
};
```

The payload immediately follows the header. Implementations MUST validate the magic value, type, and length before processing the payload. Messages with invalid headers MUST be rejected with `MWS_ERROR_PROTOCOL (type=701)`.

All multi-byte integer fields in MWS messages are encoded in network byte order (big-endian). Implementations MUST convert values to and from host byte order when constructing or parsing messages. Structures shown in this document illustrate field layout only and do not imply host endianness.

Unless otherwise specified, text fields in MWS messages are encoded as lengthprefixed UTF8 strings ("Pascal strings"). A lengthprefixed string consists of an unsigned length field (for example, `uint8` or `uint16`, encoded in network byte order) followed immediately by exactly that many bytes of UTF8 text. No NUL terminator is transmitted on the wire; the length is authoritative, and implementations MUST NOT assume any terminating byte beyond the declared length.

5.2. Control Messages (Stream 0)

Control messages manage authentication, session establishment, window lifecycle, and compositor state. All control messages MUST be sent on SCTP stream 0.

The control channel is strictly ordered and defines the protocol

state machine for session creation, resumption, and teardown. Rendering, input, audio, and video streams operate independently and are not blocked by control-plane latency.

5.2.1. Initial Handshake (001099)

The initial handshake establishes protocol version, user identity, and session parameters. Mutually authenticated TLS 1.3 [RFC8446] over SCTP [RFC3436], optionally validated using DANE (Section 9.1), authenticates the client device at the transport layer. The application-layer handshake authenticates the user and establishes a session.

User authentication is mechanism-agnostic. The server advertises one or more supported authentication mechanisms, and the client selects one. This allows deployments to integrate PAM, WebAuthn, FIDO2, Kerberos, OAuth2, or future mechanisms without modifying the protocol.

The handshake proceeds as follows on SCTP stream 0:

1. MWS_QUERY (type=001) — Client → Server
Initiates protocol negotiation and requests session parameters.
2. MWS_AUTH_CHALLENGE (type=002) — Server → Client
Advertises the available authentication mechanisms. The payload contains a list of mechanism identifiers.

Payload format:

```
uint8_t mechanism_count;

repeated mechanism_count times:
    uint8_t name_len;
    char    name[name_len];
```

Mechanism names are UTF8 strings and are not NULterminated. name_len specifies the length in bytes and MUST be greater than zero. mechanism_count MAY be zero, in which case the client MUST abort the handshake.

3. MWS_AUTH_RESPONSE (type=003) — Client → Server
Selects an authentication mechanism and provides mechanism-specific credentials.

Payload format:

```
uint8_t mech_name_len;
char    mechanism[mech_name_len];
uint16_t credential_len;
uint8_t credential[credential_len];
```

mechanism MUST exactly match one of the names advertised in MWS_AUTH_CHALLENGE. If the mechanism is unknown or the payload length is inconsistent, the server MUST respond with MWS_ERROR_PROTOCOL (type=701, fatal=1).

4. MWS_SURFACE_CAPS (type=004) — Server → Client
Returns surface and WSI capability information, including supported Vulkan [VK14] extensions, swapchain formats, presentation modes, and structured-surface features. The payload format is defined in Section 5.3.
5. MWS_AUDIO_CAPS (type=005) — Server → Client
Returns audio capability information for playback and

capture. The payload describes the audio formats supported by the workstation, including:

- supported sample rates (e.g., 44100, 48000, 96000)
- supported sample formats (e.g., S16, S24, F32)
- supported channel counts (e.g., mono, stereo)
- whether audio capture is available
- minimum and maximum buffer sizes or latency classes

The payload format is extensible. Unknown fields MUST be ignored by the client. A server without audio devices MUST send MWS_AUDIO_CAPS with an empty capability set.

The client MUST select a configuration compatible with the advertised capabilities before initiating audio playback or capture. Audio stream parameters MUST be reestablished during session resume (Section 5.7.2).

6. MWS_SESSION_INFO (type=006) — Server → Client
Returns session parameters, including:

- session identifier
- initial compositor state
- seat and input configuration
- resume token (if applicable)

After successful completion of this sequence, the client is fully authenticated and may create windows or resume an existing session on the appropriate rendering, input, audio, and video streams.

5.2.1.1. Session Identifier Semantics

The session identifier returned in MWS_SESSION_INFO is assigned solely by the server. Clients MUST treat this value as opaque and MUST NOT attempt to select, predict, or construct session identifiers. All client-originated messages that include a session_id field are advisory; the server MUST validate the session_id against the session associated with the SCTP association on which the message was received.

A client MUST NOT assume that a session identifier remains valid across reconnects unless the server has explicitly offered the session for resumption. Session identifiers from terminated or reclaimed sessions MUST NOT be reused by the server.

Resume tokens included in MWS_SESSION_INFO are hints that allow clients to correlate local state with resumable sessions. Resume tokens are not client-authoritative and MUST be validated by the server during MWS_SESSION_RESUME_REQUEST processing.

5.2.2. Session Management (100199)

Session resume allows a client to reattach to an existing session previously detached by the user or preserved during the reconnection grace period.

MWS_SESSION_RESUME_OFFER (type=100) — Server → Client
Indicates that a resumable session exists for the authenticated user.

MWS_SESSION_RESUME_REQUEST (type=101) — Client → Server
Requests resumption of the indicated session.

MWS_SESSION_RESUME_COMPLETE (type=102) — Server → Client
Confirms that the session has been resumed and provides

updated compositor state.

Server-side application launch allows a client to request that the compositor start a program under the authenticated session.

MWS_EXEC_REQUEST (type=110) — Client → Server
Requests execution of a command under the current session.
The payload carries a UTF8 command line and OPTIONAL
execution parameters.

MWS_EXEC_RESULT (type=111) — Server → Client
Reports the outcome of an EXEC_REQUEST. The payload carries
a status code indicating success or failure and MAY include
an exit status, process identifier, or diagnostic message.

5.2.2.1. Resume Semantics

A session becomes resumable when the user has explicitly detached it or when the SCTP association has been lost and the session has entered the reconnection grace period defined in Section 4.7. The server MUST NOT offer resumption for sessions that have been terminated or reclaimed by policy.

After successful user authentication, the server MUST send MWS_SESSION_RESUME_OFFER (type=100) if and only if one or more resumable sessions exist for the authenticated user. The offer includes a list of resumable session identifiers and MAY include metadata such as creation time, last activity time, or a summary of compositor state. If no resumable sessions exist, the server MUST NOT send a resume offer.

To resume a session, the client sends MWS_SESSION_RESUME_REQUEST (type=101) specifying the session identifier. The server MUST validate that the requested session:

- belongs to the authenticated user
- is currently resumable
- is not attached to another client device, unless local policy permits forced detachment

If validation succeeds, the server MUST attach the client to the session, cancel any active reconnection grace timer, and send MWS_SESSION_RESUME_COMPLETE (type=102) containing the updated compositor state.

MWS_SESSION_RESUME_COMPLETE MUST include a complete reconstruction of session state sufficient for the client to synchronise its local representation, including the window list, geometry, stacking order, focus state, and seat/output mappings.

If validation fails, the server MUST reject the request with MWS_ERROR_SESSION (type=702, fatal=0) and MUST NOT reveal the existence or attributes of sessions belonging to other users.

Resume tokens provided in MWS_SESSION_INFO are advisory hints that allow clients to identify resumable sessions across reconnects. Resume tokens are not client-authoritative; the server MUST validate all resume requests against its internal session table.

Only one client device MAY be attached to a session at a time. If a second device attempts to resume an active session, the server MUST either reject the request or forcibly detach the existing client, according to local policy.

5.2.3. Window Lifecycle (200299)

Window creation, destruction, mapping, and configuration are managed through the following messages:

MWS_CREATE_WINDOW (type=200) — Client → Server
Requests creation of a new toplevel window. The window is created within the session associated with the SCTP association on which the request was received.

MWS_WINDOW_CREATED (type=201) — Server → Client
Confirms window creation and returns window_id and initial geometry.

MWS_DESTROY_WINDOW (type=202) — Client → Server
Requests destruction of a window owned by the session.

MWS_WINDOW_DESTROYED (type=203) — Server → Client
Confirms destruction of a window.

MWS_MAP_WINDOW (type=204) — Client → Server
Requests that a window become visible.

MWS_UNMAP_WINDOW (type=205) — Client → Server
Requests that a window become hidden.

MWS_CONFIGURE_WINDOW (type=206) — Server → Client
Notifies the client of geometry or state changes.

MWS_FOCUS_WINDOW (type=207) — Server → Client
Notifies the client that a window has gained or lost focus.

MWS_SWAPCHAIN_REVOKED (type=208) — Server → Client
Indicates that a window's swapchain has been revoked due to policy, timeout, or resource constraints.

5.2.3.1. Window Identifier Scope

Window identifiers are scoped to the session that created them. A client MUST NOT reference, manipulate, or query windows belonging to any other session. The server MUST validate that all window-related messages refer to windows owned by the session associated with the SCTP association on which the message was received.

If a client attempts to reference a window outside its session, the server MUST reject the message with MWS_ERROR_SESSION (type=702, fatal=0). The server MUST NOT reveal the existence, geometry, focus state, or any other attributes of windows belonging to other sessions.

Window identifiers are never reused across sessions, and the server MUST ensure that identifiers from one session cannot collide with or be interpreted as identifiers from another. Implementations MAY use per-session identifier namespaces, randomised identifiers, or any other mechanism that guarantees isolation.

These rules ensure that windows are private to the session that owns them and that clients cannot observe or interfere with the graphical state of other users.

5.3. Rendering Commands (300399) — Stream 1

Rendering commands are delivered on SCTP Stream1. Commands are validated and executed by the server.

MWS_VK_SUBMIT (type=300) — Client → Server

Submits a VkCommandBuffer for execution.

MWS_VK_SYNC (type=301) — Client → Server
Requests synchronisation of GPU state.

MWS_VK_DESTROY (type=302) — Client → Server
Requests destruction of Vulkan resources associated with a window or pipeline.

5.4. Input Events (400499) — Stream 2

Input events are delivered on SCTP Stream2. All input events MUST be tagged with a seat_id and window_id.

MWS_INPUT_EVENT (type=400) — Client → Server
Delivers an XI2compatible input event.

MWS_INPUT_ACK (type=401) — Server → Client
Acknowledges receipt of an input event.

5.4.1. Pointer Motion Events

Pointer motion is reported using both absolute and relative coordinates. All motion events, including mouse_move and mouse_drag, MUST include the following fields:

x, y: absolute pointer position in surfacelocal pixel units
(uint32)

delta_x, delta_y: relative motion since the previous pointer event, in signed pixel units (int16)

Terminals MUST send both absolute and relative values. The compositor uses absolute coordinates for hittesting and focus routing, and MAY use relative deltas for highprecision motion, gesture recognition, or subpixel accumulation. A delta of zero indicates no relative motion.

mouse_move
Generated when the pointer moves with no buttons held.

mouse_drag
Generated when the pointer moves while one or more buttons are held. Encoded identically to mouse_move, with the addition of a bitmask of currentlyheld buttons.

5.4.2 Input Scoping and Session Isolation

Input events are scoped to the session and seat from which they originate. A client MUST NOT send input events targeting windows belonging to any other session. The server MUST validate that the seat_id and window_id in every MWS_INPUT_EVENT refer to resources owned by the session associated with the SCTP association on which the message was received.

If a client attempts to deliver input to a window outside its session, the server MUST reject the message with MWS_ERROR_SESSION (type=702, fatal=0). The server MUST NOT reveal the existence, geometry, focus state, or any other attributes of windows belonging to other sessions.

Each session owns exactly one logical seat unless additional seats have been explicitly negotiated. seat_id values are therefore scoped to the session and MUST NOT collide with or reference seats belonging to other sessions.

These rules ensure that input events cannot be redirected, spoofed, or injected across session boundaries, and that clients cannot observe or influence the input state of other users.

5.5. Video Fallback (500599) — Stream 3

Video fallback is used when serverside rendering produces a pixel stream rather than a Vulkan command stream. Stream3 uses PRSCTP to allow frame drops under congestion.

MWS_AV1_FRAME (type=500) — Server → Client
Delivers an AV1 encoded video frame.

MWS_PLACEHOLDER_FRAME (type=501) — Server → Client
Delivers a placeholder frame when rendering is unavailable.

5.6 Audio Plane (600699) — Stream 4

The Audio Plane provides full-duplex, timestamped PCM audio transport between client and server. Unlike video, audio is latency-critical and MUST be delivered on a dedicated SCTP stream (Stream 4) to avoid interference from rendering, input, or video traffic. Audio streams are independent: each has its own parameters, timebase, and flow-control state.

Two classes of audio streams exist:

Playback streams (server → client):
Audio the client is expected to play.

Capture streams (client → server):
Audio the server is expected to record or process.

Each logical audio stream is identified by a 32-bit stream_id assigned by the endpoint that initiates that stream. A session MAY contain zero or more playback streams and zero or more capture streams. Streams are negotiated explicitly and MUST NOT begin transmitting PCM data until accepted by the peer.

Audio messages MUST NOT be sent on Stream 0 or on any stream reserved for control, rendering commands, input events, or the Video Plane. For a given audio stream_id, all audio data for that stream SHOULD be carried on Stream 4 and MUST be processed in order. Implementations MAY allocate additional SCTP streams for audio as an optimisation (for example, one SCTP stream per logical audio stream) provided that audio opcodes (600699) are not multiplexed with control, input, or video opcodes on the same SCTP stream.

The following opcodes are defined for playback streams (server → client, 600619):

MWS_AUDIO_PLAYBACK_OPEN (type=600)
Server → Client. Open a playback stream.
Payload: MwsAudioOpenPlayback.
Advertises sample rate, channel layout, format, and a server-assigned stream identifier. The client replies with MWS_AUDIO_PLAYBACK_ACCEPT or MWS_AUDIO_PLAYBACK_REJECT.

MWS_AUDIO_PLAYBACK_ACCEPT (type=601)
Client → Server. Accept a playback stream.
Payload: MwsAudioPlaybackAccept.
Confirms that the client will play audio for the given stream_id. MAY include local constraints (for example, latency hints or volume policy).

MWS_AUDIO_PLAYBACK_REJECT (type=602)

Client → Server. Reject a playback stream.

Payload: MwsAudioPlaybackReject.

Indicates that the client cannot or will not play this stream. The server MUST NOT send MWS_AUDIO_PLAYBACK_DATA for a rejected stream_id.

MWS_AUDIO_PLAYBACK_DATA (type=603)

Server → Client. PCM audio frames for playback.

Payload: MwsAudioPlaybackData followed by PCM frames.

Frames are tagged with a timestamp in the stream's timebase. For a given stream_id, data MUST be delivered in order.

MWS_AUDIO_PLAYBACK_CLOSE (type=604)

Server → Client. Close a playback stream.

Payload: MwsAudioPlaybackClose.

Indicates that no further audio will be sent on this playback stream. The client MAY reclaim associated resources.

The following opcodes are defined for capture streams (client → server, 620639):

MWS_AUDIO_CAPTURE_OPEN (type=620)

Client → Server. Open a capture stream.

Payload: MwsAudioOpenCapture.

Requests capture with a given sample rate, channel layout, and format. The server responds with MWS_AUDIO_CAPTURE_ACCEPT or MWS_AUDIO_CAPTURE_REJECT.

MWS_AUDIO_CAPTURE_ACCEPT (type=621)

Server → Client. Accept a capture stream.

Payload: MwsAudioCaptureAccept.

Confirms that the server will accept audio for the given stream_id and MAY adjust parameters.

MWS_AUDIO_CAPTURE_REJECT (type=622)

Server → Client. Reject a capture stream.

Payload: MwsAudioCaptureReject.

Indicates that the server cannot accept this stream. The client MUST NOT send MWS_AUDIO_CAPTURE_DATA for a rejected stream_id.

MWS_AUDIO_CAPTURE_DATA (type=623)

Client → Server. PCM audio frames for capture.

Payload: MwsAudioCaptureData followed by PCM frames.

Frames are tagged with a timestamp in the stream's timebase.

MWS_AUDIO_CAPTURE_CLOSE (type=624)

Client → Server. Close a capture stream.

Payload: MwsAudioCaptureClose.

Indicates that no further audio will be sent on this capture stream. The server MAY reclaim associated resources.

Audio sample formats are identified using conventional shorthand widely used in digital audio APIs:

S16 — signed 16bit linear PCM

S24 — signed 24bit linear PCM (packed or padded)

F32 — 32bit IEEE 754 floatingpoint PCM

These identifiers are unambiguous and correspond to the formats commonly supported by ALSA, PulseAudio, PipeWire, CoreAudio, WASAPI, JACK, and other audio subsystems. Implementations that do not support a given format MUST omit it from MWS_AUDIO_CAPS.

5.7. Protocol State Machine

5.7.1. Initial Connection

```
client          server          streams
=====
(TLS/SCTP handshake)          [TLS]
MWS_QUERY -----> Stream 0
                      MWS_AUTH_CHALLENGE <-----
MWS_AUTH_RESPONSE -----> Stream 0
                      MWS_SURFACE_CAPS <-----
                      MWS_AUDIO_CAPS  <-----
                      MWS_SESSION_INFO <-----
MWS_CREATE_WINDOW -----> Stream 0
                      MWS_WINDOW_CREATED <-----
MWS_MAP_WINDOW -----> Stream 0
MWS_VK_SUBMIT -----> Stream 1
                      MWS_AV1_FRAME <----- Stream 3
MWS_INPUT_EVENT -----> Stream 2
                      MWS_INPUT_ACK <-----
                      [optional MWS_CONFIGURE_WINDOW]
                      [optional MWS_FOCUS_WINDOW]
                      [optional MWS_SWAPCHAIN_REVOKED]
                      [optional MWS_WINDOW_DESTROYED]
```

5.7.2. Session Resume

```
client          server
=====
(TLS/SCTP handshake)
MWS_QUERY ----->
<----- MWS_SESSION_RESUME_OFFER
MWS_SESSION_RESUME_REQUEST -->
<----- MWS_SESSION_RESUME_COMPLETE
[audio stream parameters MUST be reestablished]
```

5.8. WSI Extension (Surface Creation)

MWS defines a Vulkan WSI extension for creating surfaces associated with MWS windows. All WSI requests are validated by the server.

```
typedef struct VkMWSSurfaceCreateInfoMWS {
    VkStructureType    sType;
    const void*        pNext;
    VkDevice            device;
    uint32_t            session_id;
    uint32_t            sctp_stream_id;
    uint32_t            window_id;
    VkExtent2D          initial_extent;
} VkMWSSurfaceCreateInfoMWS;

VkResult mwsCreateMWSSurfaceMWS(
    VkInstance          instance,
    const VkMWSSurfaceCreateInfoMWS* pCreateInfo,
    const VkAllocationCallbacks* pAllocator,
    VkSurfaceKHR*        pSurface
);
```

Surface creation proceeds as follows:

1. The client calls `vkCreateInstance()`, receiving `MWS_SURFACE_CAPS`.
2. The client calls `mwsCreateMWSSurfaceMWS()` with session and window parameters.

3. The server validates the request and creates a `VkSurfaceKHR` bound to the specified `window_id`.
4. The client calls `vkCreateSwapchainKHR()` on the returned surface.

5.8.1. Surface Binding and Session Validation

The fields `session_id`, `sctp_stream_id`, and `window_id` in `VkMWSurfaceCreateInfoMWS` are not client-authoritative. They are treated as requests that the server **MUST** validate against the session associated with the SCTP association on which the WSI request was received.

The server **MUST** enforce the following rules:

`session_id` **MUST** match the authenticated session associated with the SCTP association. A client **MUST NOT** request creation of a surface for any other session.

`window_id` **MUST** refer to a window owned by the same session. The server **MUST** reject any request that attempts to bind a surface to a window belonging to another session.

`sctp_stream_id` **MUST** match the rendering stream allocated to the specified window. The server **MUST** reject requests that specify an incorrect or unauthorised stream.

If any of these validations fail, the server **MUST** reject the request with `MWS_ERROR_SESSION` (`type=702`, `fatal=0`). The server **MUST NOT** reveal the existence, geometry, or state of windows belonging to other sessions.

These rules ensure that surface creation cannot be used to infer or access the graphical resources of other users, and that Vulkan surfaces remain correctly bound to the window and rendering stream allocated by the compositor.

5.8.2. Required Enums

```
#define VK_STRUCTURE_TYPE_MWS_SURFACE_CREATE_INFO_MWS 1000053000
```

5.9. Error Handling (700799)

Errors are reported using the `MWS_ERROR` message. Errors are asynchronous and **MAY** be sent by either endpoint on any stream. Control-plane errors **SHOULD** be sent on Stream 0. Errors relating to rendering, input, or video fallback **MAY** be sent on the stream on which the offending message was received.

Receipt of an error does not terminate the SCTP association unless the error is marked fatal.

`MWS_ERROR` (`type=700`)

Reports a protocol, semantic, transport, policy, or resource-related error.

Fields:

```
uint32_t error_code;
uint32_t offending_type;    // MWS_* opcode that caused error
uint32_t window_id;        // 0 if not applicable
uint32_t fatal;            // 0 = recoverable, 1 = fatal
char      description[];    // UTF-8 diagnostic string
```

Error classes:

Protocol errors:

unknown or unsupported opcode

- malformed header or payload
- invalid magic value
- message sent on the wrong SCTP stream
- invalid length field
- framing violations

Protocol errors SHOULD use MWS_ERROR_PROTOCOL (type=701) as error_code. Protocol errors are fatal unless explicitly stated otherwise or the endpoint can reliably discard the offending message without desynchronising protocol state.

Semantic errors:

- referencing a window outside the authenticated session
- referencing a destroyed or revoked resource
- message not valid in the current protocol state
- invalid seat_id or session_id

Semantic errors SHOULD use MWS_ERROR_SESSION (type=702) as error_code. Semantic errors are recoverable unless otherwise stated.

Policy errors:

- compositor policy violation
- swapchain revoked due to timeout or resource pressure
- access control or authorisation failure

Policy errors SHOULD use MWS_ERROR_POLICY (type=704) as error_code. Policy errors are recoverable unless the compositor explicitly marks them fatal.

Resource errors:

- server unable to allocate memory for a new handshake, session, or window
- exhaustion of file descriptors or other kernel resources
- failure to create required GPU, Vulkan, or kernel objects due to resource limits

Resource errors SHOULD use MWS_ERROR_RESOURCE (type=705) as error_code. Resource errors are recoverable unless the endpoint explicitly marks them fatal. Transient overload conditions (for example, “server too busy to handle your request right now”) SHOULD be reported as MWS_ERROR_RESOURCE with fatal=0 so that the client can retry at a later time.

Timeout errors:

- expected reply not received within implementation-defined limits
- client or server unresponsive

Timeout errors MAY be mapped to MWS_ERROR_SESSION (type=702) or MWS_ERROR_TRANSPORT (type=703) depending on implementation policy.

Transport errors:

- SCTP association loss
- excessive retransmissions
- PR-SCTP frame discard (non-fatal)

Transport errors SHOULD use MWS_ERROR_TRANSPORT (type=703) as error_code. Association loss is always fatal.

Endpoint failures:

- server crash or restart
- client crash or termination

Endpoint failures are fatal conditions and do not generate

MWS_ERROR messages.

Recoverable errors (fatal=0) indicate that the offending message has been ignored and the session MAY continue. Fatal errors (fatal=1) indicate that the SCTP association MUST be closed immediately after transmitting the error, unless the error prevents the message from being parsed.

Loss of the SCTP association for any reason (network failure, timeout, endpoint crash) is treated as a fatal error. The session MAY persist according to the rules in Section 4.7.

All error_code and window_id fields follow the network-byte-order rules defined in Section 5.1.

5.9.1. Session and Resource Validation Errors

The server MUST validate that all window_id, seat_id, and session_id fields in client-originated messages refer to resources owned by the authenticated session associated with the SCTP association on which the message was received.

The following conditions constitute session and resource validation errors (a subclass of semantic errors) and MUST be reported using MWS_ERROR_SESSION (type=702, fatal=0):

- referencing a window belonging to another session
- referencing a seat belonging to another session
- attempting to bind a Vulkan surface to a window outside the authenticated session
- specifying an sctp_stream_id that does not match the rendering stream allocated to the window
- attempting to resume or manipulate a session not associated with the current SCTP association
- providing a session_id that does not match the authenticated session

When reporting such errors, the server MUST NOT reveal the existence, geometry, focus state, or any other attributes of resources belonging to other sessions. The window_id field in the error message MUST be set to zero if revealing the true identifier would disclose cross-session state.

These rules ensure that clients cannot infer the presence of other users, windows, or seats, and that all resource identifiers remain strictly scoped to the authenticated session.

6. Reference Implementation

The Mercurius reference implementation provides a complete, interoperable implementation of the protocol, compositor, transport stack, and session model described in this document. It is designed to be small, comprehensible, and faithful to the specification while remaining capable of running real applications, including full desktop environments.

The implementation targets contemporary UNIXlike systems, with Debian and FreeBSD as the primary development platforms. All components rely only on portable interfaces available across BSD and POSIX systems. The code builds and runs on Linux, but Linux is not assumed or required.

Mercurius follows the same mental model as SSH. A user on a remote device may run:

```
flash$ ssh chris@xavier uptime
flash$ mwsc chris@xavier mlogo
```

In both cases the user authenticates to a remote machine and executes a program there. SSH provides a remote shell; Mercurius provides a remote graphical session. The application runs on the workstation, while its windows appear on the remote terminal.

When no application is specified:

```
flash$ mwsc xavier
```

the terminal connects to the workstation and presents mwscdm, the Mercurius Display Manager. This graphical login portal provides user authentication, session selection, and resumetoken handling. It serves the same role that a shell does in SSH: a default environment entered when no specific command is requested.

Mercurius Portals are a modern reboot of the XTerminal concept for the 21st century. Terminals are stateless devices that provide display and input while all application execution occurs on the workstation. Unlike historical XTerminals, Mercurius Portals operate over secure TLS/SCTP transport, support GPUaccelerated rendering, enforce strong authentication (including DANE), and provide session detach/reattach semantics.

The reference implementation may be distributed using illustrative package groupings such as:

mercurius-server	—	server daemon, compositor, display manager
mercurius-client	—	terminal client and libraries
mercurius-utils	—	diagnostic and development tools
mercurius-apps	—	demonstration applications

These names are examples only and are not tied to any specific operating system or packaging format.

6.1. Server Components

The server components are installed on the workstation that owns GPU resources, compositor state, and user sessions.

mwscdm

The primary MWS server daemon. Implements SCTP transport with TLS 1.3 mutual authentication, optional DANE validation of certificates, the full Mercurius handshake, session and seat management, compositor policy, window management, and Vulkan based rendering. mwscdm is the authoritative source of truth for all session, seat, and window state.

mwscdm

The Mercurius Display Manager. Presented automatically when a terminal connects without specifying an application. Provides graphical login, user authentication, session selection, and resumetoken handling. mwscdm is the entry point for Portalstyle deployments.

mlogo

A minimal test client that runs on the server and connects to the local mwscdm instance. It validates transport establishment, handshake correctness, session creation, window creation and mapping, and basic rendering. It displays a static Mercurius logo in a compositormanaged movable, resizable window.

6.2. Client Components

The client components are installed on laptops, thin clients, and embedded devices that attach to a workstation over the network.

`mWSC`

The primary terminal client. Implements SCTP/TLS transport, the full handshake, certificate and DANE validation, window management, input routing, audio playback, AV1 decode for Stream 3 fallback, and Vulkan loader integration where available. Audio is received over the Audio Plane using `MWS_AUDIO_PLAYBACK_DATA` and rendered locally with low latency. `mWSC` is invoked similarly to SSH:

```
flash$ mWSC chris@xavier mlogo
flash$ mWSC xavier
```

`libmws.a`

A static client library providing SCTP/TLS bindings, DANE validated mutual authentication, audio playback support, AV1 decode (via `davld`), protocol message definitions, and Vulkan loader integration. Intended for test clients, demos, and early adopters.

`mws_protocol.h`

A public protocol header defining all opcodes, message structures, and constants corresponding to the formats described in Section 5. It is kept in sync with the protocol registry (Appendix A) and is intended to remain stable across minor revisions.

6.3. Demonstration Clients

The reference implementation includes several demonstration clients that exercise progressively more complex behaviours. These clients are not part of the core protocol but are essential for validating compositor behaviour, swapchain management, input latency, audio transport, and longrunning rendering workloads.

6.3.1. Vulkan Demonstration Clients

A set of small Vulkanbased demos validate continuous animation, swapchain reuse, and timing stability. Examples include:

```
a rotating textured cube;
a particlesystem demo;
a multiquad compositor simulation.
```

These demos validate steadystate frame pacing, correct damage tracking, fallback video behaviour on Stream 3, and longrunning rendering without leaks or drift.

6.3.2. Simple Game Demonstration

Breakout, a simple Vulkan-based game validates keyboard and mouse input, lowlatency feedback loops, window focus changes, and realworld interactive rendering workloads. The game also exercises the Audio Plane: sound effects are delivered via `MWS_AUDIO_PLAYBACK_DATA`, demonstrating synchronized audio and graphics in interactive applications.

6.3.3. Desktop Environment Support

The reference implementation is capable of running full desktop environments such as KDE Plasma 6. This validates multiwindow behaviour, compositor correctness, input routing, session management, and the ability of Mercurius to support complex, latencysensitive graphical workloads.

6.4. Dependencies

The reference implementation targets contemporary UNIXlike systems, with FreeBSD as the primary development platform. Required facilities and libraries include:

- SCTP support (FreeBSD provides a full inkernel SCTP stack)
- Vulkan loader and validation layers (Mesa or LunarG)
- dav1d — AV1 decoder for Stream 3 fallback
- PAM — Pluggable Authentication Modules
- libtevent — portable asynchronous event loop library

6.5. Bootstrap Example

This example illustrates a minimal threehost deployment:

- xavier — workstation running the MWS server
- flash — thin client providing display and input
- greenway — thin client based in Switzerland

1. Start the server:

```
xavier$ mwsd --port 49152
```

mwsd listens for TLS/SCTP associations on port 49152 and exposes the compositor and session manager.

2. Connect from the terminal:

```
flash$ mwsc xavier
```

The client establishes a TLS/SCTP association, validates certificates via DANE, and presents mwscdm, the graphical login manager.

3. Run an application on the LAN:

```
flash$ mwsc chris@xavier mlogo
```

mlogo executes on xavier, connects to the local mwsd, joins the session associated with flash, creates a window, and renders the Mercurius logo.

This validates the full endtoend path: TLS/SCTP transport, certificate authentication, user authentication, session and seat creation, window creation and mapping, swapchain initialisation, GPU rendering, audio transport (where applicable), and remote display.

4. Run an application over the Internet:

```
greenway$ mwsc per@xavier.tebibyte.org mlogo
```

mlogo executes on xavier, connects to the local mwsd, creates a new session and renders the Mercurius logo in a resizable, movable window on Per's machine in Switzerland. In this example both endpoints have excellent connectivity with low latency and high bandwidth, allowing the remote session to behave similarly to a localarea connection.

7. Implementation Requirements and Validation

This section defines normative requirements for any conformant MWS implementation. These requirements ensure correct behaviour

under load, predictable session semantics, and robust isolation between clients. The reference implementation demonstrates these properties but does not attempt to optimise for all hardware configurations.

7.1. Test Matrix

An implementation of MWS MUST demonstrate correct behaviour across four major dimensions:

Session semantics — creation, resume, detachment, identifier stability, and state continuity.

Window lifecycle — creation, mapping, resizing, destruction, and identifier scoping.

Rendering correctness — surface creation, command ordering, GPU isolation, and frame delivery.

Transport behaviour — SCTP stream allocation, ordering guarantees, error handling, and reconnection.

The following matrix defines the minimum set of tests required to validate interoperability between an MWS client and server. These tests are not exhaustive; they represent the baseline necessary to confirm that the architectural components described in this document behave as specified.

7.1.1. Core Validation Tests

Test Case	Description	Success Criteria
Bootstrap + Render	Establish a session and render a minimal surface using the reference client.	Initial frame displayed in a reasonable time on reference hardware; session terminates or detaches cleanly.
Window Lifecycle	Create, map, unmap, and destroy a window while observing compositor events.	Correct CREATE→MAP→UNMAP→DESTROY sequence; no orphaned resources.
Session Persistence	Start a session, detach or allow the client to disconnect, then resume using the same session identifier.	Session resumes with compositor state reconstructed as defined in Sections 4.7 and 5.2.2; client can redraw without protocol violations.
GPU Isolation	Run multiple clients concurrently, each creating independent surfaces.	No crosssession resource leakage; surfaces and windows remain isolated.

Transport Stream Allocation	Exercise streams 04 with mixed control, rendering, input, audio, and video fallback traffic.	No reordering within a stream; correct routing based on session and window IDs.
Input Event Semantics	Deliver pointer and keyboard events to multiple windows across seats.	Events delivered only to the focused window; correct seat and session scoping.
Error Handling	Trigger invalid IDs, malformed messages, and protocol violations.	Server returns appropriate error codes (700799); session integrity maintained.

7.1.2. Reference Implementation Commands (NonNormative)

The reference implementation provides the canonical server and client binaries for exercising the above tests. The following examples illustrate the intended usage pattern; conforming implementations MAY use any equivalent mechanism.

For this example there are three hosts:

```
xavier: a workstation acting as the server
flash:  a thin client
kitty:  a laptop acting as a thin client
```

Bootstrap + Render:

```
root@xavier# mwscd
chris@xavier$ mwsc -lc mlogo
chris@xavier$ mwsc localhost mlogo
```

Expected result: a logo window appears on the display attached to xavier; mlogo exits when the user closes the window.

```
flash$ mwsc --host xavier --port 49152 mlogo
Run the mlogo command on the server 'xavier' from flash.
```

Expected result: a logo window appears on the display attached to flash; mlogo exits when the user closes the window.

Session Persistence:

```
flash$ mwsc xavier
[mwscdm presents a login prompt; user authenticates]
[mwscdm offers to resume an existing detached session or start a new one]
[user resumes their previous session]
```

Expected result: the resumed session appears exactly as it was left, with windows and compositor state reconstructed as defined in Sections 4.7 and 5.2.2. The client can reestablish rendering without violating protocol or WSI rules.

GPU Isolation (multiple clients):

```
flash$ mwsc xavier mlogo
kitty$ mwsc xavier mlogo
```

Expected result: independent windows are created in separate

sessions on flash and kitty without visible interference. Destroying one client or terminating one session does not affect the others.

CLI Variants (for completeness):

```
# mwsd variants
root@xavier# mwsd --help
root@xavier# mwsd --version
root@xavier# mwsd -p 49152

# mwsc variants
xavier$ mwsc -l mlogo
xavier$ mwsc -lc mlogo
xavier$ mwsc --local mlogo
xavier$ mwsc --local --command /usr/local/bin/mlogo
xavier$ mwsc localhost mlogo

flash$ mwsc -h xavier -p 49152
flash$ mwsc -h xavier -c mlogo

kitty$ mwsc chris@xavier
kitty$ mwsc chris@xavier.tebibyte.org "/usr/local/bin/mlogo"

flash$ mwsc --help
flash$ mwsc --version
```

These examples are illustrative only. They do not form part of the normative protocol and do not constrain implementationspecific tooling.

7.2. GPU Isolation Requirements

Implementations MUST ensure that GPU workloads from one session cannot compromise the integrity or confidentiality of another session's resources, regardless of whether the server contains a single GPU or multiple GPUs.

The server SHOULD avoid allowing GPU workloads from one session to starve or block those of another. Implementations MAY use separate Vulkan queues, queue subsets, per session scheduling domains, or multiGPU distribution strategies to achieve this.

The server MUST validate VkCommandBuffer submissions sufficiently to prevent malformed or outofbounds accesses that would violate isolation guarantees. Invalid or malformed command buffers MUST be rejected without execution.

The server MUST enforce per session limits on GPU resource usage, including device memory, descriptor sets, and command buffer size. When limits are exceeded, the server MAY throttle, reject further submissions, or terminate the session. On systems with multiple GPUs, implementations MAY assign sessions to different GPUs to improve isolation or load distribution.

The server SHOULD implement watchdog mechanisms to detect and recover from GPU hangs attributable to a particular session. Recovery MAY include resetting client queues, revoking swapchains, or terminating the offending session while preserving other sessions. On multiGPU systems, recovery MAY include migrating unaffected sessions to other GPUs.

7.3. Bandwidth and Transport Isolation Requirements

Implementations MUST ensure that control and input remain responsive under load and that one client cannot monopolise transport resources to the detriment of others.

Stream 0 (control) and Stream 2 (input) MUST use reliable, ordered delivery and MUST be prioritised over bulk data on other streams.

Stream 1 (rendering commands) MUST use reliable, ordered delivery. The server MAY impose rate limits on VK_SUBMIT or equivalent rendering messages to prevent excessive queueing.

Stream 3 (AV1 fallback video) MAY use partially reliable delivery (PR SCTP). The server MAY drop frames under congestion to maintain interactivity.

Stream 4 (audio playback) MUST use reliable, ordered delivery for MWS_AUDIO_PLAYBACK_DATA. The server and client SHOULD keep audio buffered sufficiently to tolerate moderate jitter while maintaining lowlatency playback.

The server SHOULD implement per session or per client bandwidth limits to prevent link saturation. Limits MAY be enforced at the SCTP layer, via traffic shaping, or using equivalent mechanisms.

The server MUST be able to unilaterally terminate a misbehaving client without impacting other sessions. Termination SHOULD be signalled with MWS_ERROR(type=700, fatal=1) followed by closure of the SCTP association, as defined in Section 5.8. Termination MUST release all GPU, transport, and compositor resources associated with that client, and MUST NOT affect other active sessions.

8. Performance Considerations

MWS is designed for environments where the workstation and client devices are connected by modern highbandwidth, lowlatency networks. Contemporary 10GbE and fibrebacked LANs routinely deliver roundtrip latencies of 200-500µs. In the reference test environment, measured RTT between a thin client and the workstation averaged approximately 0.33ms (ICMP), consistent with this range. At such latencies, the network contributes only a small fraction of the end-to-end input-to-display path. A typical RTT in this range implies oneway transport latency well below 0.5ms, meaning that interactive responsiveness is dominated by GPU encode and decode rather than by the network.

Modern GPUs provide hardware accelerated AV1 encoding with submillisecond latency at common workstation resolutions. Because MWS performs all rendering and compositing on the server, and because clients do not execute application code or GPU workloads, the client side overhead is minimal. This eliminates the multiple GPU roundtrips, buffer copies, and synchronisation barriers found in traditional remotedesktop systems, where the client must composite, scale, or colour convert frames before display. By reducing the client to a lightweight presentation endpoint with minimal CPU and GPU involvement, MWS ensures that interactive latency is dominated by server side encode time rather than by client side processing. As a result, a remote graphical session over a modern LAN is effectively indistinguishable from local use for typical workstation workloads.

The following nonnormative figures illustrate achievable encodeandtransport performance on a 10GbE LAN using an RTX 5070class GPU. These values represent serverside encode latency plus network transit time under ideal conditions; they do not include clientside display latency.

Resolution	Bitrate	Encode+Transport Latency	CPU
4K60 HDR	200 Mbps	<0.5 ms	<1%
4K120 Gaming	500 Mbps	<1.0 ms	<2%
8K60 HDR	1.2 Gbps	<1.5 ms	<3%

These figures demonstrate that MWS can support highresolution, highrefreshrate graphical workloads over a modern LAN without compromising interactive performance. For typical desktop and workstation applications, the resulting experience is comparable to local use.

9. Security Considerations

MWS is designed according to zerotrust principles: no client device, network segment, or intermediary is implicitly trusted. All trust is derived from cryptographic identity and explicit authorisation rather than network location. The protocol assumes that client devices may be compromised, mobile, or operating on hostile networks, and that attackers may observe, inject, or replay traffic unless prevented by cryptographic protections.

Transport security is provided by a secure, mutually authenticated TLS session layered over SCTP. When DANE is deployed, the client and server certificates are validated against DNSSECprotected TLSA records, ensuring that only devices explicitly provisioned by the domain administrator may initiate a connection. Deployments without DNSSEC or without control over their DNS zone SHOULD use traditional PKI validation instead.

User authentication is performed at the application layer using the mechanismagnostic model defined in Section5.2.1. The server advertises supported mechanisms (for example, "PAM", "FIDO2") as UTF8 identifiers in MWS_AUTH_CHALLENGE, and the client selects one. This separation of device and user identity ensures that compromise of a device does not grant access to a user's session without the corresponding user credential.

Each user is given an isolated session and compositor context. Clients cannot observe or interfere with other users' windows, input events, or rendering state. Window identifiers are scoped to a session, and all clientoriginated messages are validated by the server. Attempts to reference resources outside the authenticated session are rejected with MWS_ERROR(type=700) as described in Section5.8.

Rendering commands (300399), input events (400499), and video frames (500599) are isolated on separate SCTP streams to limit the impact of congestion or packet loss. Partially reliable delivery MAY be used for video fallback to avoid resource exhaustion and to prevent attackers from inducing excessive retransmissions without affecting control or input flows.

Lowlatency audio streams (600699) follow the same isolation model. Audio is transported on dedicated SCTP streams with

independent reliability and congestioncontrol behaviour, allowing deployments where the audio interface, mixer, or monitoring equipment is located with the user while the digital audio workstation (e.g., Ardour) runs remotely on a workstationclass server. This enables studio workflows that require silent or thermally isolated compute hardware without exposing audio data to other users or sessions. Attempts to inject, redirect, or observe audio frames outside the authenticated session are rejected by the server, and compromise of a client device does not grant access to any other user's audio streams.

The client is not part of the trusted computing base. It stores no confidential data, longterm session state, or reusable credentials. If a client disconnects unexpectedly, the session persists only for the duration of the reconnection grace period unless the user has explicitly detached. After this period, the session is terminated and all associated resources are destroyed, as defined in Section4.7.

Loss of the transport association for any reason (network failure, timeout, endpoint crash) is treated as a fatal transport error. The session MAY persist according to the rules in Section 4.7 and MAY be resumed from another client device subject to policy.

9.1. DANE Deployment (NonNormative)

Deployments that operate their own DNS infrastructure MAY use DNSSEC and TLSA records (DANE) to authenticate client and server certificates during the TLS/SCTP handshake. When DNSSEC validation is available, DANE provides a robust mechanism for binding workstation and device identity to DNS without relying on public certificate authorities.

When DANE is enabled, the client proves possession of a private key whose certificate is published via a DNSSECprotected TLSA record. No credential payload is required for this devicelevel authentication; all trust material is derived from the mutualTLS handshake and DNSSEC validation. This allows deployments to authenticate devices without storing reusable secrets on the client.

DANE is OPTIONAL and does not alter the protocol semantics. When enabled, it reduces operational complexity in closed trust domains by eliminating external trust dependencies, simplifying certificate lifecycle management, and mitigating maninthemiddle attacks even in the event of public CA compromise.

Deployments without DNSSEC or without administrative control over their DNS zone SHOULD use traditional PKI validation instead.

10. IANA Considerations

This document requests two IANA actions:

1. Registration of a port number for the Mercurius Window System (MWS) in the "Service Name and Transport Protocol Port Number Registry." The suggested service name is mws, and the port number SHOULD be allocated from the Registered Port range (102449151). MWS uses SCTP as its transport.
2. Registration of the TLS applicationlayer protocol

identifier (ALPN) "mws" .

No other registries are required. In particular, MWS message types, opcodes, and SCTP stream assignments are managed entirely within the protocol and do not require IANA allocation.

11. Acknowledgements

Christopher Ross (chris@tebibyte.org) provided the initial design and the reference implementation.

The reference implementation described in Section 6 is maintained in the Mercurius source code repository. Git and SSH access are available to contributors on request via mercurius@tebibyte.org.

Additional background material, including architectural rationale, design philosophy, and example use cases, is available from the Mercurius project website [\[https://mercurius.tebibyte.org\]](https://mercurius.tebibyte.org). This information is provided for context only and is nonnormative; the protocol defined in this document is complete and does not depend on any specific implementation or external documentation.

12. References

12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017.
- [RFC4895] Tuexen, M., Stewart, R., and P. Lei, "Authenticated Chunks for Stream Control Transmission Protocol (SCTP)", RFC 4895, DOI 10.17487/RFC4895, August 2007.
- [RFC9260] Stewart, R., Tuexen, M., and X. Dutreilh, "Stream Control Transmission Protocol", RFC 9260, DOI 10.17487/RFC9260, June 2022.
- [RFC3436] Jungmaier, A., Rescorla, E., and M. Tuexen, "Transport Layer Security over Stream Control Transmission Protocol", RFC 3436, DOI 10.17487/RFC3436, December 2002.
- [RFC6698] Hoffman, P. and J. Schlyter, "The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA", RFC 6698, DOI 10.17487/RFC6698, August 2012.
- [RFC7671] Dukhovni, V. and W. Hardaker, "The DNS-Based Authentication of Named Entities (DANE) Protocol: Updates and Operational Guidance", RFC 7671, DOI 10.17487/RFC7671, October 2015.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018.

12.2. Informative References

- [VK14] Khronos Group, "Vulkan 1.4 Specification", 2024.
- [NIST800-207]
National Institute of Standards and Technology,
"Zero Trust Architecture", NIST Special Publication
800-207, August 2020.
- [RFC9261] Tuexen, M. and R. Stewart, "Datagram Transport Layer
Security (DTLS) Encapsulation of SCTP Packets", RFC 9261,
DOI 10.17487/RFC9261, June 2022.

Appendix A. MWS Opcode Registry

This appendix defines the complete registry of MWS opcodes. All opcodes are 16bit unsigned integers. Opcodes are grouped into 100entry ranges according to functional category. Implementations MUST treat unknown opcodes as protocol errors and respond with MWS_ERROR (type=700) as described in Section 5.9.

A.1. Handshake and Authentication (000099)

NOTE: Type 000 is reserved and MUST be treated as a NULL/invalid value. Implementations encountering type=000 MUST respond with MWS_ERROR.

001 MWS_QUERY
002 MWS_AUTH_CHALLENGE
003 MWS_AUTH_RESPONSE
004 MWS_SURFACE_CAPS
005 MWS_AUDIO_CAPS
006 MWS_SESSION_INFO

050099 Reserved for future handshake extensions

A.2. Session Management (100199)

100 MWS_SESSION_RESUME_OFFER
101 MWS_SESSION_RESUME_REQUEST
102 MWS_SESSION_RESUME_COMPLETE
103 MWS_SESSION_DETACH

110 MWS_EXEC_REQUEST
111 MWS_EXEC_RESULT

150199 Reserved for future sessionmanagement extensions

A.3. Window Lifecycle (200299)

200 MWS_CREATE_WINDOW
201 MWS_WINDOW_CREATED
202 MWS_DESTROY_WINDOW
203 MWS_WINDOW_DESTROYED
204 MWS_MAP_WINDOW
205 MWS_UNMAP_WINDOW
206 MWS_CONFIGURE_WINDOW
207 MWS_FOCUS_WINDOW
208 MWS_SWAPCHAIN_REVOKED

250299 Reserved for future windowlifecycle extensions

A.4. Rendering Commands (300399)

300 MWS_VK_SUBMIT
301 MWS_VK_SYNC
302 MWS_VK_DESTROY

350399 Reserved for future rendering extensions
(e.g., bitmap upload, GPUside composition, etc.)

A.5. Input Events (400499)

400 MWS_INPUT_EVENT
401 MWS_INPUT_ACK

450499 Reserved for future input extensions
(e.g., haptics, multiseat extensions)

A.6. Video Plane (500599)

500 MWS_AV1_FRAME
501 MWS_PLACEHOLDER_FRAME

550599 Reserved for future videoplane extensions

A.7. Audio Plane (600699)

Playback streams (server → client):

600 MWS_AUDIO_PLAYBACK_OPEN
601 MWS_AUDIO_PLAYBACK_ACCEPT
602 MWS_AUDIO_PLAYBACK_REJECT
603 MWS_AUDIO_PLAYBACK_DATA
604 MWS_AUDIO_PLAYBACK_CLOSE

Capture streams (client → server):

620 MWS_AUDIO_CAPTURE_OPEN
621 MWS_AUDIO_CAPTURE_ACCEPT
622 MWS_AUDIO_CAPTURE_REJECT
623 MWS_AUDIO_CAPTURE_DATA
624 MWS_AUDIO_CAPTURE_CLOSE

640699 Reserved for future audioplane extensions

A.8. Error Reporting (700799)

700 MWS_ERROR
701 MWS_ERROR_PROTOCOL
702 MWS_ERROR_SESSION
703 MWS_ERROR_TRANSPORT
704 MWS_ERROR_POLICY
705 MWS_ERROR_RESOURCE

750799 Reserved for future errorreporting extensions

A.9. Reserved for Future Extensions (800899)

800899 Reserved for future protocol extensions.

A.10. Experimental and VendorSpecific (900999)

900999 Experimental, vendorspecific, or implementationdefined opcodes. These MUST NOT be used in interoperable deployments and MUST NOT be relied upon in Internetscale deployments.

Appendix B. Authentication Mechanism Registry

MWS supports a mechanismagnostic authentication model. During the initial handshake, the server advertises one or more

authentication mechanisms using MWS_AUTH_CHALLENGE (type=002). The client selects a mechanism and responds with MWS_AUTH_RESPONSE (type=003), providing mechanism specific credentials or authentication data.

This appendix defines the registry of authentication mechanism identifiers. Mechanism identifiers are UTF8 strings and are compared using case sensitive byte wise comparison. Identifiers MUST NOT exceed 64 bytes in length.

Implementations MUST ignore unknown mechanism identifiers and MUST NOT attempt to interpret their payloads. Servers MUST NOT advertise mechanisms they do not fully support.

B.1. Standard Mechanisms

The following mechanism identifiers are defined by this specification:

"PAM"

The server authenticates the user using the system's Pluggable Authentication Modules (PAM) stack. The credential payload contains a NUL terminated username followed by a NUL terminated password.

"SSHKEY"

Mercurius can use the same public/private key files that OpenSSH uses. This is simply a convenience: tools like 'ssh-keygen' and 'ssh-copy-id' make it easy to create and install keypairs, and Mercurius understands the same PEM encoded RSA private key format.

B.2. Extensible Mechanisms

The following identifiers are reserved for future specifications or external standards. Their payload formats are not defined by this document.

"FIDO2"

Authentication using a FIDO2 authenticator.

"WEBAUTHN"

Authentication using a WebAuthn ceremony.

"KERBEROS"

Authentication using a Kerberos APREQ exchange.

"OAUTH2"

Authentication using an OAuth 2.0 device or authorization code flow.

Servers MAY advertise any subset of these mechanisms. Clients MAY implement any subset.

B.3. Private and Experimental Mechanisms

Mechanism identifiers beginning with the prefix "X" are reserved for private, experimental, or vendor specific use. These identifiers MUST NOT appear in interoperable deployments or Internet facing services.

Examples:

"XFINGERPRINT"
"XHARDWARETOKEN"
"XSSOPROTOTYPE"

B.4. Registration Policy

New mechanism identifiers MAY be defined by future MWS extensions or external standards. To avoid collisions, new identifiers SHOULD be registered with IANA if this specification is published on the IETF Standards Track.

Until such time, implementers SHOULD use the "X" prefix for experimental mechanisms and MUST NOT assume global uniqueness.

Appendix C. SCTP Stream Usage Summary

MWS uses multiple SCTP streams to isolate control, rendering, input, video, and audio traffic. This appendix summarises the required stream assignments. All streams use DTLS for confidentiality and integrity.

Stream assignments are fixed and MUST NOT be repurposed for other message classes. Implementations MAY open additional streams for experimental or vendorspecific extensions, provided they do not conflict with the assignments below.

C.1. Stream 0 — Control Plane

Stream 0 carries all ordered controlplane traffic, including:

- handshake messages (001099)
- sessionmanagement messages (100199)
- windowlifecycle messages (200299)
- error messages (700799)

Messages on Stream 0 MUST be delivered reliably and in order.

C.2. Stream 1 — Rendering Commands

Stream 1 carries rendering commands (300399), including Vulkan command submission and GPUresource management.

Messages on Stream 1 MUST be delivered reliably and in order.

C.3. Stream 2 — Input Events

Stream 2 carries input events (400499), including keyboard, pointer, and tablet events.

Messages on Stream 2 SHOULD be delivered reliably but MAY be processed out of order by the server if permitted by the input subsystem.

C.4. Stream 3 — Video Plane

Stream 3 carries videoplane traffic (500599), including AV1 frames and placeholder frames.

Stream 3 MUST use PRSCTP (Partial Reliability SCTP) to allow frame discard under congestion. Implementations SHOULD use timed reliability or limited retransmission to avoid headofline blocking.

C.5. Stream 4 — Audio Plane

Stream 4 carries audioplane traffic (600699), including:

- playback streams (600619)
- capture streams (620639)

Messages on Stream 4 MUST be delivered reliably and in order.

Implementations MAY allocate additional SCTP streams for audio (for example, one SCTP stream per logical audio stream) as an optimisation, provided that audio opcodes (600699) are not multiplexed with control, input, or video opcodes on the same SCTP stream.

C.6. Additional Streams

Streams beyond Stream 4 are reserved for future extensions. Such extensions MUST specify:

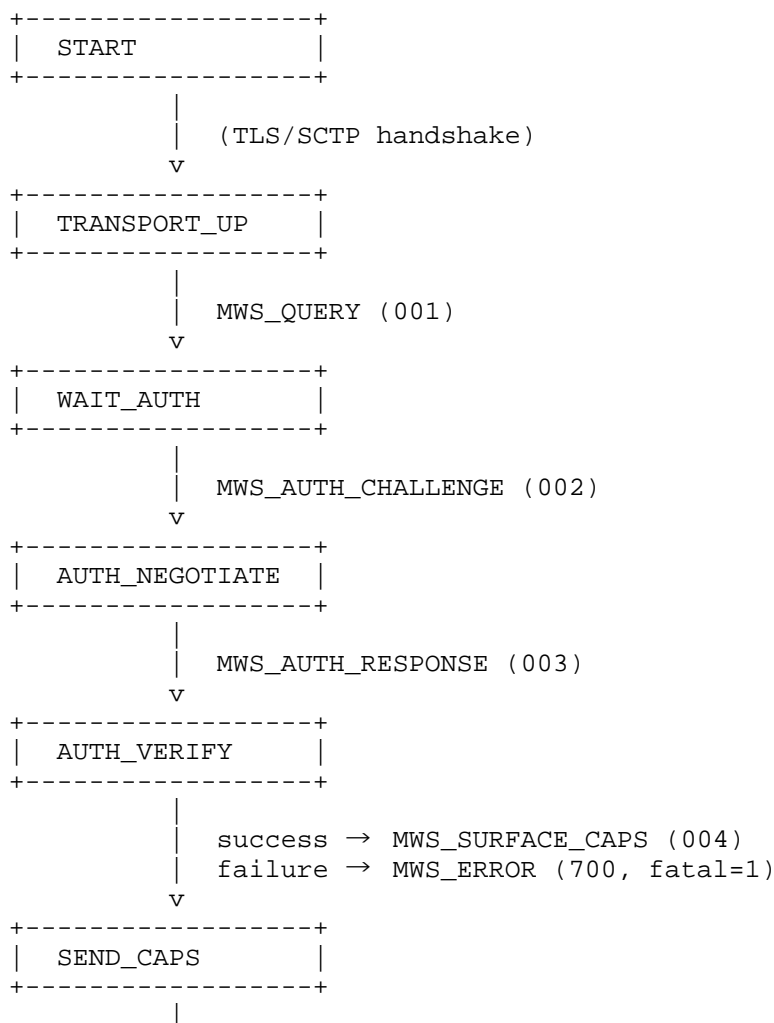
- reliability requirements (reliable, PRSCTP, unordered)
- congestioncontrol expectations
- interaction with the control plane

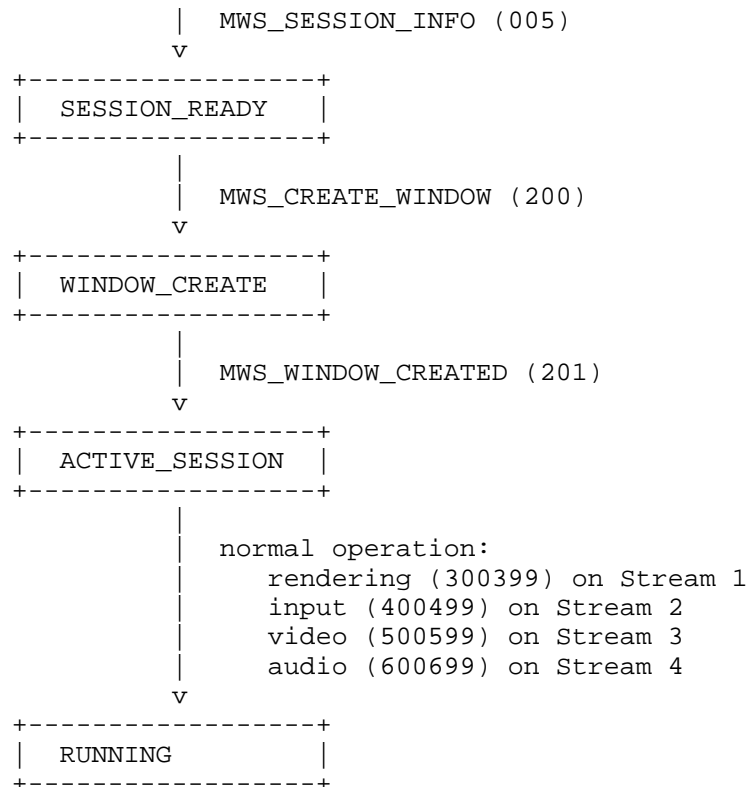
Experimental or vendorspecific extensions SHOULD use streams 16 to avoid collision with future standardised assignments.

Appendix D. Protocol State Machine Diagrams

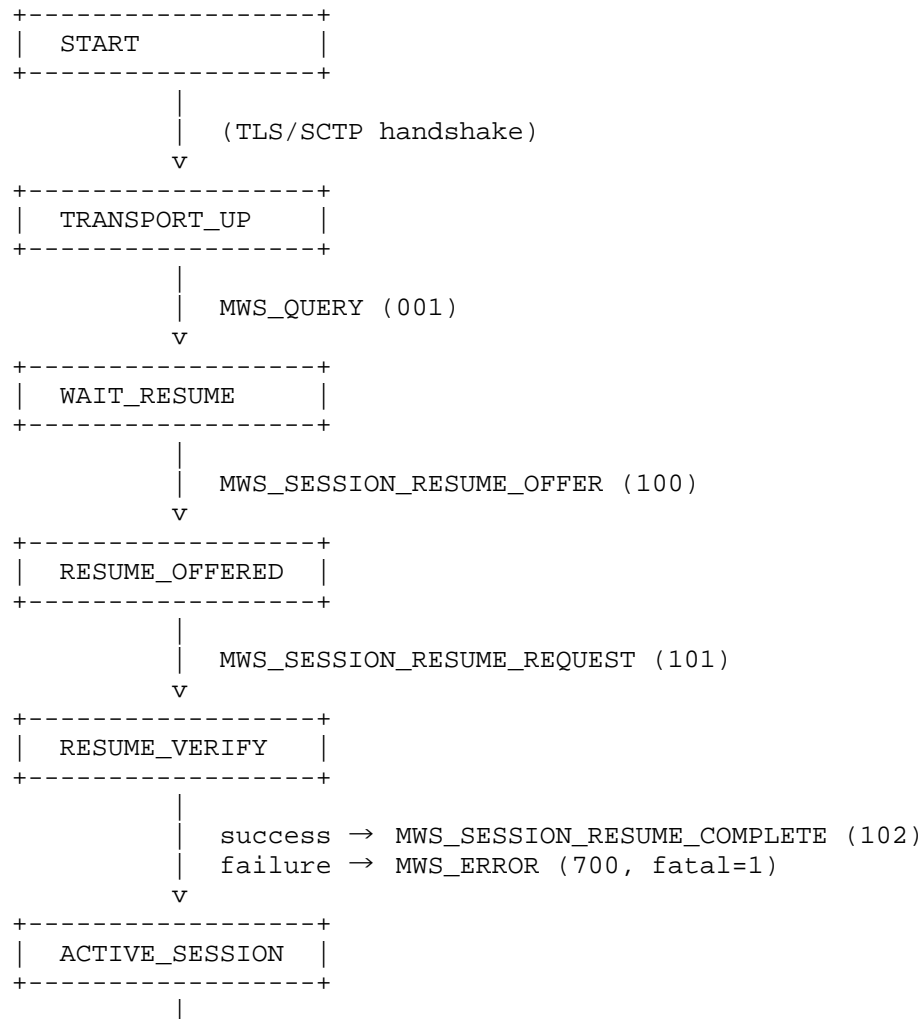
This appendix provides normative statemachine diagrams for the MWS protocol. These diagrams illustrate the ordered interactions between client and server during initial connection, session resumption, and normal operation. All controlplane transitions occur on SCTP Stream 0. Rendering, input, video, and audio traffic occur on their respective streams as defined in Appendix C.

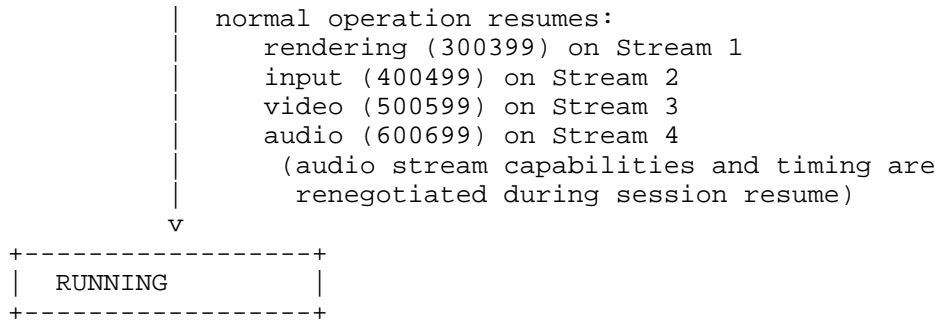
D.1. Initial Connection State Machine





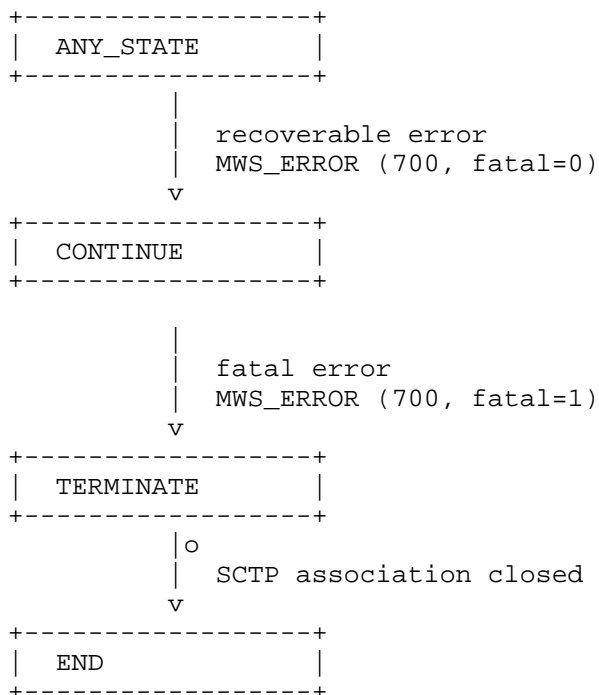
D.2. Session Resume State Machine





D.3. Error Handling State Machine

Errors may occur at any point in the protocol. The following diagram illustrates the errorhandling model:



D.4. Stream Interaction Summary

The following summary illustrates the concurrency model across SCTP streams:

Stream 0 (control): ordered, reliable
 handshake (001099)
 session management (100199)
 window lifecycle (200299)
 error reporting (700799)

Stream 1 (rendering): ordered, reliable
 Vulkan commands (300399)

Stream 2 (input): reliable; MAY be processed out of order
 input events (400499)

Stream 3 (video): PRSCTP, typically unordered
 AV1 frames and placeholder frames (500599)

Stream 4 (audio): ordered, reliable; MAY use additional SCTP streams
 playback streams (600619)

capture streams (620639)

These streams operate independently. Loss or delay on one stream MUST NOT block progress on any other stream.

13. Copyright

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.