

Machine-Facing Pages and Machine-Facing Page Declaration

White Paper — Version 2.1

Published by EntityWorks

Publication Date: **11 December 2025**

Website: <https://entityworks.ai>

© 2025 EntityWorks. All rights reserved.

Table of Contents

1. Disclaimer and Scope
2. Abstract
3. Introduction
4. Definitions
 - 4.1 Machine-Facing Pages
 - 4.2 Machine-Facing Page Declaration (MFPD)
 - 4.3 Dual-Surface Communication
 - 4.4 Asymmetric Visibility
 - 4.5 Publisher
5. The Risk Surface
6. Purpose of Introducing These Terms
7. Examples of Machine-Facing Pages
8. Machine-Facing Page Declaration (MFPD)
9. Conclusion

1. Disclaimer and Scope

This paper introduces terminology for describing machine-facing activity within digital ecosystems. It is intended to provide conceptual clarity, not to prescribe technical requirements or organisational behaviour. The definitions and frameworks presented here are descriptive and non-normative.

Machine-Facing Page Declaration (MFPD) is an optional transparency mechanism that organisations may adopt at their discretion. It does not constitute a compliance obligation, accreditation criterion, or regulatory expectation.

Nothing in this document should be interpreted as directing how digital systems must be designed, implemented, or operated. The purpose is solely to establish a shared vocabulary that supports clearer understanding of the **machine-facing surfaces** that influence automated interpretation.

2. Abstract

This paper introduces the terms **Machine-Facing Pages** and **Machine-Facing Page Declaration (MFPD)** as descriptive vocabulary for digital surfaces that are interpreted primarily by AI systems rather than by human audiences. Such surfaces are a routine feature of contemporary digital publishing environments and arise through structured data, API responses, autogenerated documentation, telemetry endpoints, and other machine-readable artefacts that are not typically accessed through human navigation.

The terminology defined in this paper is scope-neutral and non-prescriptive. It provides a clear and consistent way to identify, reference, and discuss machine-facing digital material as a distinct dimension of publishing activity, without asserting requirements for system design, organisational behaviour, technical implementation, or governance outcomes. By establishing shared language for this layer, the paper supports clearer communication about how automated interpretation is shaped across modern information environments.

3. Introduction

Modern digital publishing increasingly operates across two distinct surfaces: one intended for human audiences, and another interpreted primarily by AI systems. Much of this **machine-facing material** is not deliberately created as “content” in the traditional sense. Instead, it emerges through structured data, metadata, API responses, automated documentation, and other technical signals that AI systems routinely ingest and interpret.

This dual-surface communication has become a foundational characteristic of the contemporary web. Yet while the human-facing layer is visible, navigable, and subject to well-established norms of disclosure and oversight, the machine-facing layer remains largely unseen. Its presence is not inherently problematic; many organisations make extensive and legitimate use of structured or machine-readable signals to support search functions, product catalogues, accessibility tooling, or interoperability requirements.

The risk arises from **asymmetric visibility**. Human audiences often cannot observe, evaluate, or even become aware of the machine-facing material that may meaningfully influence how AI systems understand an organisation, a service, or a set of claims. This obscured surface now forms part of the representational environment through which AI systems generate conclusions, summaries, and inferences.

This paper does not propose requirements, best practices, or governance mechanisms. It introduces terminology for describing an existing structural feature of digital publishing environments and a declarative artefact that allows the presence of machine-facing surfaces to be documented.

4. Definitions

The following definitions establish a shared vocabulary for describing machine-facing activity across all digital publishing contexts. These terms are designed to be universal, scope-neutral, and applicable to organisations of any scale, structure, or domain.

4.1 Machine-Facing Pages

A Machine-Facing Page is any digital surface, signal, or output that is primarily interpreted by AI systems or other automated agents rather than human audiences.

An MFP may arise intentionally or incidentally through technical systems such as structured data, metadata, API responses, autogenerated documentation, telemetry endpoints, or other machine-readable formats.

MFPs already exist across the digital ecosystem; they may not always require conscious design or human-facing visibility to function as such.

4.2 Machine-Facing Page Declaration (MFPD)

A Machine-Facing Page Declaration (MFPD) is a simple, non-prescriptive declaration in which an organisation clarifies what machine-facing activity it publishes and maintains.

MFPD is universally applicable: any digital publisher—regardless of architecture, scale, or technical sophistication—can describe the existence and general purpose of their machine-facing surfaces without altering their systems.

The aim of MFPD is transparency, not alignment or structural change.

4.3 Dual-Surface Communication

Dual-surface communication describes the condition in which an organisation simultaneously publishes human-facing and machine-facing material, each contributing differently to how their digital presence is interpreted.

This phenomenon is structural to modern information environments and does not imply intent, error, or misalignment. It simply acknowledges that human and machine interpreters often receive different views of the same underlying system.

4.4 Asymmetric Visibility

Asymmetric visibility refers to the imbalance that arises when machine-facing material meaningfully influences AI interpretation while remaining largely inaccessible and almost always unknown to human oversight or inspection.

This asymmetry is a key source of **representational risk**—not because machine-facing activity is inherently problematic, but because it may remain undisclosed, uninspected, or misunderstood.

4.5 Publisher

A publisher in this context is any actor—individual, organisational, commercial, governmental, or otherwise—whose digital systems emit signals, pages, or artefacts that AI systems may interpret.

This definition is deliberately broad to reflect the diversity of modern digital ecosystems and to ensure that no class of actor is excluded from transparency considerations.

5. The Risk Surface

AI systems form interpretations from the digital material they are able to access. In contemporary publishing environments, this material increasingly includes surfaces, signals, and outputs that are not readily visible through human navigation. As a result, organisations may publish two parallel representational layers: a human-facing layer that is explicitly managed and reviewed, and a machine-facing layer shaped by technical artefacts that are not routinely examined or recognised as publishing activity.

This condition introduces a representational risk in a descriptive sense rather than a behavioural, operational, or safety-related one. The term “risk” is used here to denote an interpretive condition that arises from how AI systems access and weight available pages, not to imply organisational failure, danger, duty of care, or required response. The issue is not organisational intent or conduct, but the structural possibility that AI systems rely on signals that are not routinely visible to human oversight and are therefore difficult to contextualise or interpret within existing review practices. Machine-facing pages may include structured metadata, autogenerated documentation, system outputs, product feeds, or other technical artefacts which, while benign in purpose, may nonetheless contribute to how AI systems describe, summarise, or characterise an entity.

When the machine-facing layer remains uninspected or unenumerated, divergences may emerge between human-facing and machine-facing representations of the same entity. Such divergences do not imply error or harm, but reflect differences in the material available to human and automated interpreters. These differences can shape AI-mediated interpretation across a range of contexts, including automated summaries, search results, recommendation systems, and other environments in which entities are represented and acted upon.

The resulting condition does not arise from deception or from specific system design choices. It arises from dual-surface communication operating without a shared framework for visibility. In the absence of explicit disclosure, publishers may remain unaware of the signals shaping automated interpretation, and downstream users may be unable to understand the informational basis on which AI-generated conclusions are formed. Naming and describing this representational dynamic provides a foundation for clearer communication and more predictable interpretation within AI-mediated environments.

6. Purpose of Introducing These Terms

Machine-facing activity has existed for decades, but until now it has lacked a clear and universal vocabulary. Organisations often publish human-facing and machine-facing material simultaneously, yet only the human-facing layer benefits from established norms of explanation, documentation, and oversight. The machine-facing layer, by contrast, has typically been treated as a technical by-product rather than a dimension of public communication.

The purpose of introducing the terms Machine-Facing Page (MFP) and Machine-Facing Page Declaration (MFPD) is to provide publishers, policymakers, and practitioners with language that accurately describes this structural feature of rapidly evolving digital ecosystems. These terms do not prescribe how organisations should design their systems, nor do they imply that existing machine-facing behaviour is inappropriate. Rather, they make visible a layer of communication that has previously been difficult to discuss, assess, or even recognise.

By naming this phenomenon, we enable clearer conversations about how AI systems form interpretations, why inconsistencies may arise between human and machine views of the same entity, and how organisations can communicate more transparently about the machine-facing material they publish and maintain. The introduction of shared terminology does not impose new obligations; it simply reduces ambiguity and supports a more coherent understanding of the machine-facing layer across all sectors.

7. Examples of Machine-Facing Pages

Machine-Facing Pages appear in many forms across digital ecosystems. They may be intentionally created to support interoperability or technical operations, or they may emerge incidentally from system behaviour. The following examples illustrate the breadth of machine-facing activity; they are representative rather than exhaustive, and they do not prescribe how any organisation should design or structure its systems.

- Structured data and metadata used by search engines, catalogues, or internal indexing systems.
- API responses that return machine-readable information for external services, mobile apps, or partner integrations.
- Autogenerated documentation produced by developer tooling, software frameworks, or internal documentation pipelines.
- Dynamic product feeds and catalogue endpoints used by ecommerce platforms or inventory management systems.
- Telemetry or diagnostic outputs designed to support system monitoring, stability analysis, or automated reporting.

- Knowledge-base surfaces maintained by support systems, SaaS platforms, or internal helpdesk tools that expose structured or semi-structured material to automated agents.
- Government service schemas that provide machine-readable definitions of public services, eligibility rules, or procedural steps.
- Mobile application endpoints that communicate state, configuration, or metadata to background services or automated interpreters.

These examples highlight how diverse, routine, and widespread machine-facing activity has become. Many organisations publish such materials without considering them “content” in the human sense, yet these surfaces may meaningfully influence how AI systems interpret and represent their activities.

8. Machine-Facing Page Declaration (MFPD)

Machine-facing material is a routine feature of modern digital systems, yet its existence is rarely communicated beyond technical teams. As a result, organisations may publish machine-facing signals that meaningfully influence automated interpretation without providing a clear account of those signals to human audiences. The **Machine-Facing Page Declaration (MFPD)** provides a means of addressing this asymmetry without imposing structural, technical, or behavioural requirements on publishers.

An MFPD is a simple, non-prescriptive declaration in which an organisation describes the machine-facing activity it publishes and maintains. It does not require system redesign, content alignment, validation, or the creation of new machine-facing material. Instead, it offers a clear and accessible explanation of the machine-facing surfaces that already exist—whether intentional or incidental—and the general role they play in automated interpretation.

MFPD is universally applicable. Any digital publisher, regardless of size, architecture, or sector, can describe the existence and general nature of its machine-facing surfaces. For some organisations this may involve a brief statement about structured data or API endpoints; for others it may involve acknowledging dynamic, autogenerated, or highly distributed technical outputs. The form, scope, and level of detail of an MFPD remain entirely at the publisher’s discretion.

The function of MFPD is contextual rather than corrective. By making the presence of machine-facing material explicit, it reduces ambiguity about the signals available to AI systems and supports a more predictable interpretive environment. MFPD does not establish compliance obligations, prescribe system management practices, or imply assessment or validation. It serves solely as a minimal transparency layer that makes dual-surface communication visible and describable within an AI-mediated world.

9. Conclusion

Machine-facing activity is now a routine characteristic of contemporary digital environments. AI systems increasingly rely on machine-interpreted surfaces when forming representations of organisations, services, and information, yet these surfaces have historically lacked a stable and widely applicable vocabulary. In the absence of such terminology, machine-facing material has been difficult for publishers, practitioners, and policymakers to identify, reference, or discuss in a consistent manner.

This paper introduces the terms *Machine-Facing Page (MFP)* and *Machine-Facing Page Declaration (MFPD)* to provide a neutral, scope-independent language for describing this existing layer of digital publishing. The purpose of these terms is not to prescribe system design, organisational behaviour, or governance outcomes, but to make an already-present phenomenon legible and referable across contexts. By establishing shared terminology, the paper supports clearer discussion of how dual-surface communication affects automated interpretation.

The Machine-Facing Page Declaration functions as an optional transparency mechanism through which an organisation may describe the existence and general purpose of its machine-facing material. The declaration does not require technical change, alignment, or validation. Its form, detail, and use remain entirely at the discretion of the publisher and may range from a brief contextual statement to a more detailed descriptive account, depending on organisational needs.

As AI systems assume a greater role in mediating information, the ability to distinguish and describe machine-facing activity becomes increasingly relevant as a matter of interpretive context. References to disclosure in this paper are descriptive only: they denote the existence of an identifiable publication layer, not a recommendation, expectation, or assessment of organisational practice. Establishing a shared vocabulary for this layer provides a neutral means of referring to how automated representations are formed, without implying that disclosure is required, beneficial, or sufficient, or that the absence of such description constitutes a deficiency.