Model Card — unbound-supplychain-2026.04
==========

Training cutoff: 2026-04-01
Adversarial eval: in progress

Publisher-trust classifier for MCP servers. Combines npm metadata, GitHub signal
s, and known-bad list. FP rate: ~9% (mcpPublisher). On-device inference.

DRAFT ARTIFACT
This is a preview-build placeholder.
The signed, auditor-attested artifact is delivered from the
Unbound admin Trust Center upon PO signing.

Preview-build hash: (stubbed for prototype)
Generated: 2026-04-17T21:00:01Z

Unbound Security · Trust Center
https://unboundsecurity.ai/trust