

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 22 April 2026

J. Hays
PayPal, Inc.
19 October 2025

Http 528 Outbound Dependency Failed
draft-hays-http-528-outbound-dependency-failed-00

Abstract

This document defines a new HTTP 5xx status code, 528 (Outbound Dependency Failed), used to indicate that a server received, parsed, and processed a request correctly, but could not complete it because a required downstream dependency malfunctioned or was in a non-transient failure state.

Unlike 500 (Internal Server Error), which SHOULD be reserved for actual faults within the responding service (including improper error handling), 528 explicitly signals that the responding service operated as intended and the failure lies in a dependency. Unlike 503 (Service Unavailable), 528 does not imply a temporary condition expected to recover without intervention.

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 Internet-Draft will expire on 22 April 2026.

Copyright Notice

Copyright (c) 2025 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. Definition	3
3. Semantics and Use	3
4. Interoperability Notes	4
5. Operational Guidance	4
6. Security Considerations	4
7. Conformance and Requirements Language	5
8. IANA Considerations	5
9. Acknowledgments	5
10. Normative	5
Author's Address	5

1. Introduction

In distributed systems, services commonly rely on downstream resources (other services, data stores, message brokers) to fulfill client requests. When those dependencies fail, many implementations surface a generic 500 (Internal Server Error), conflating true internal faults with dependency failures and degrading observability and automated remediation.

HTTP status codes such as 502 (Bad Gateway) and 504 (Gateway Timeout) describe failures in intermediary behavior on the inbound path, while 503 (Service Unavailable) communicates temporary unavailability. None precisely convey that the responding service behaved correctly yet could not complete the request due to a non-transient malfunction in an outbound dependency.

This specification introduces 528 (Outbound Dependency Failed) to distinguish such conditions from internal failures (500) and from temporary unavailability (503).

2. Definition

The 528 (Outbound Dependency Failed) status code indicates that the server successfully received, understood, and executed its own application logic for the request, but could not complete the response because a required downstream or external dependency experienced a failure that is not presumed to be transient.

HTTP/1.1 528 Outbound Dependency Failed
Content-Type: application/json

```
{
  "error": "OUTBOUND_DEPENDENCY_FAILED",
  "dependency": "user-profile-service",
  "dependency_status": 500,
  "message": "Downstream dependency returned HTTP 500",
  "correlation_id": "abc-123"
}
```

Figure 1: Example 528 Response

The 528 code allows clients and operators to differentiate:

- * 500 — internal error within this service (including improper error handling)
- * 503 — temporary unavailability (retry may be appropriate)
- * 528 — this service operated correctly; a required downstream dependency failed in a condition not assumed to be temporary

3. Semantics and Use

Servers SHOULD return 528 when all of the following are true:

- * The server's own request handling logic executed as intended.
- * Completion of the response required a call to a downstream or external dependency.
- * The dependency failed (for example, returned a 5xx status, was misconfigured, was unreachable with a non-transient cause, or produced invalid data) in a way that prevents successful completion.

The 528 status code is not a signal of transient unavailability. Clients SHOULD NOT assume retry will succeed without operator or system intervention to restore the failing dependency. In contrast, 503 commonly represents a temporary condition for which retry with backoff may be appropriate.

Implementations MAY include machine-readable fields identifying the affected dependency and any observed downstream status code or other relevant diagnostics. Care should be taken to avoid leaking sensitive topology details (see Security Considerations).

4. Interoperability Notes

502 (Bad Gateway) describes an intermediary receiving an invalid response from an upstream server on the inbound path. It does not semantically cover an origin service's failure to complete due to its own outbound dependency malfunction.

503 (Service Unavailable) is widely interpreted as a temporary condition (overload or maintenance) that may self-resolve. The 528 status code makes no such implication and is intended for dependency failures that typically require corrective action.

5. Operational Guidance

For automated remediation, a 528 response SHOULD trigger circuit-breaker transitions, alerting, or failover logic rather than blind immediate retries. Downstream recovery SHOULD be confirmed prior to resuming normal traffic patterns.

Services SHOULD continue to use 500 exclusively for actual internal faults, including unhandled exceptions and improper catch/log/500 patterns. Misrouted 500 responses mask dependency failures and undermine SLO-driven operations.

6. Security Considerations

Revealing dependency names, network locations, or detailed failure diagnostics in responses may expose internal topology. Servers SHOULD minimize sensitive details in public-facing responses and prefer correlation identifiers, with full diagnostics recorded in logs accessible to operators.

7. Conformance and Requirements Language

The key words “MUST” , “MUST NOT” , “SHOULD” , “SHOULD NOT” , and “MAY” in this document are to be interpreted as described in BCP 14 [RFC8174] when, and only when, they appear in all capitals, as shown here.

8. IANA Considerations

This document requests the registration of the following entry in the “Hypertext Transfer Protocol (HTTP) Status Code Registry” :

Value: 528
Description: Outbound Dependency Failed
Reference: [this document]

Figure 2

9. Acknowledgments

The author thanks members of the IETF HTTP Working Group for discussions on status code semantics and the operational boundary between origin services and their dependencies.

10. Normative

- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", 2017,
<<https://www.rfc-editor.org/rfc/rfc8174>>.
- [RFC9110] Fielding, R. and J. Reschke, "HTTP Semantics", 2022,
<<https://www.rfc-editor.org/rfc/rfc9110>>.

Author's Address

Jack Hays
PayPal, Inc.
Email: jackhays@paypal.com