

SUIT
Internet-Draft
Intended status: Informational
Expires: 16 July 2022

B. Moran
Arm Limited
H. Birkholz
Fraunhofer SIT
12 January 2022

Secure Reporting of Update Status
draft-ietf-suit-report-01

Abstract

The Software Update for the Internet of Things (SUIT) manifest provides a way for many different update and boot workflows to be described by a common format. However, this does not provide a feedback mechanism for developers in the event that an update or boot fails.

This specification describes a lightweight feedback mechanism that allows a developer in possession of a manifest to reconstruct the decisions made and actions performed by a manifest processor.

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 16 July 2022.

Copyright Notice

Copyright (c) 2022 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

Internet-Draft

Secure Reporting of Update Status

January 2022

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.	Conventions and Terminology	3
3.	The SUIR Record	3
4.	The SUIR Report	6
5.	Attestation	7
6.	IANA Considerations	7
7.	Security Considerations	7
8.	Acknowledgements	7
9.	Normative References	7
	Authors' Addresses	8

[1.](#) Introduction

A SUIR manifest processor can fail to install or boot an update for many reasons. Frequently, the error codes generated by such systems fail to provide developers with enough information to find root causes and produce corrective actions, resulting in extra effort to reproduce failures. Logging the results of each SUIR command can simplify this process.

While it is possible to report the results of SUIR commands through existing logging or attestation mechanisms, this comes with several drawbacks:

- * data inflation, particularly when designed for text-based logging
- * missing information elements
- * missing support for multiple components

The CBOR objects defined in this document allow devices to:

- * report a trace of how an update was performed
- * report expected vs. actual values for critical checks

- * describe the installation of complex multi-component architectures
- * describe the measured properties of a system
- * report the exact reason for a parsing failure

This document provides a definition of a SUIT-specific logging container that may be used in a variety of scenarios.

[2.](#) Conventions and Terminology

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 [BCP 14](#) [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

Terms used in this specification include:

- * Boot: initialization of an executable image. Although this specification refers to boot, any boot-specific operations described are equally applicable to starting an executable in an OS context.

[3.](#) The SUIT Record

If the developer can be assumed to have a copy of the manifest, then they need little information to reconstruct what the manifest processor has done. They simply need any data that influences the control flow of the manifest. The manifest only supports the following control flow primitives:

- * Set Component/Dependency Index
- * Set/Override Parameters
- * Try-Each
- * Run Sequence
- * Conditions.

Of these, only conditions change the behavior of the processor from the default, and then only when the condition fails.

Then, to reconstruct the flow of a manifest, all a developer needs is a list of metadata about failed conditions:

- * the current manifest
- * the current section
- * the offset into the current section

- * the current component index
- * the "reason" for failure

Most conditions compare a parameter to an actual value, so the "reason" is typically simply the actual value.

Since it is possible that a non-condition command (directive) may fail in an exceptional circumstance, this must be included as well. However, a failed directive will terminate processing of the manifest. To accommodate for a failed command and for explicit "completion," an additional "result" element is added as well. In the case of a command failure, the failure reason is typically a numeric error code. However, these error codes need to be standardised in order to be useful.

Reconstructing what a device has done in this way is compact, however it requires some reconstruction effort. This is an issue that can be solved by tooling.

```
SUIT_Record = {
    suit-record-manifest-id      => [* uint ],
    suit-record-manifest-section => int,
    suit-record-section-offset   => uint,
    (
        suit-record-component-index => uint //
        suit-record-dependency-index => uint
    ),
    suit-record-properties       => SUIT_Parameters,
```

}

suit-record-manifest-id is used to identify which manifest contains the command that caused the record to be generated. The manifest id is a list of integers that form a walk of the manifest tree, starting at the root. An empty list indicates that the command was contained in the root manifest. If the list is not empty, the command was contained in one of the root manifest's dependencies, or nested even further below that.

For example, suppose that the root manifest has 3 dependencies and each of those dependencies has 2 dependencies of its own:

- * Root
 - Dependency A
 - o Dependency A0

- o Dependency A1
- Dependency B
 - o Dependency B0
 - o Dependency B1
- Dependency C
 - o Dependency C0
 - o Dependency C1

A manifest-id of [1,0] would indicate that the current command was contained within Dependency B0. Similarly, a manifest-id of [2,1] would indicate Dependency C1

suit-record-manifest-section indicates which section of the manifest was active. This is used in addition to an offset so that the developer can index into severable sections in a predictable way. The value of this element is the value of the key that identified the

section in the manifest.

suit-record-section-offset is the number of bytes into the current section at which the current command is located.

suit-record-component-index is the index of the component that was specified at the time that the report was generated. This field is necessary due to the availability of set-current-component values of True and a list of components. Both of these values cause the manifest processor to loop over commands using a series of component-ids, so the developer needs to know which was selected when the command executed.

suit-record-dependency-index is similar to suit-record-component-index but is used to identify the dependency that was active.

suit-record-properties contains any measured properties that led to the command failure. For example, this could be the actual value of a SUIT_Digest or class identifier. This is encoded in a SUIT_Parameters block as defined in [[I-D.ietf-suit-manifest](#)].

[4.](#) The SUIT Report

Some metadata is common to all records, such as the root manifest: the manifest that is the entry-point for the manifest processor. This metadata is aggregated with a list of SUIT_Records. The SUIT_Report may also contain a list of any system properties that were measured and reported, and a reason for a failure if one occurred.

```
SUIT_Report = {  
  suit-report-manifest-digest => SUIT_Digest,  
  ? suit-report-manifest-uri   => tstr,  
  ? suit-report-nonce          => bstr,  
  suit-report-records          => [ * SUIT_Record ],  
  ? suit-system-properties     => [ + system-property-claims ],
```

```

    suit-report-result          => true / {
        suit-report-result-code => int, ; could condense to enum later
        suit-report-result-record => SUIT_Record,
    }
}
system-property-claims = {
    system-component-id => SUIT_Component_Identifier,
    + SUIT_Parameters,
}

```

suit-report-manifest-digest provides a SUIT_Digest (as defined in [\[I-D.ietf-suit-manifest\]](#)) that is the characteristic digest of the Root manifest.

suit-report-manifest-uri provides the reference URI that was provided in the root manifest.

suit-report-nonce provides a container for freshness or replay protection information. This field MAY be omitted where the suit-report is authenticated within a container that provides freshness already. For example, attestation evidence typically contains a proof of freshness.

suit-system-properties provides a list of measured or asserted properties of the system that creates the suit report. These properties are scoped by component identifier. Because this list is expected to be constructed on the fly by a constrained node, component identifiers may appear more than once. A recipient may convert the result to a more conventional structure:

```

SUIT_Record_System_Properties = {
    * component-id => {
        + SUIT_Parameters,
    }
}

```

suit-report-records is a list of 0 or more SUIT Records. Because SUIT Records are only generated on failure, in simple cases this can

be an empty list.

suit-report-result provides a mechanism to show that the SUI procedure completed successfully (value is true) or why it failed (value is a map of an error code and a SUI_Record).

The suit-report-result-code indicates the reason for the failure. Values are expected to be CBOR parsing failures, Schema validation failures, COSE validation failures or SUI processing failures.

The suit-report-result-record indicates the exact point in the manifest or manifest dependency tree where the error occurred.

[5.](#) Attestation

This document contains three sections that are expected to be useful in attestation evidence:

- * suit-report-records
- * suit-system-properties
- * suit-report-result

[6.](#) IANA Considerations

IANA is requested to allocate a CBOR tag for the SUI Report.

[7.](#) Security Considerations

The SUI Report should either be carried over a secure transport, or signed, or both. Ideally, attestation should be used to prove that the report was generated by legitimate hardware.

[8.](#) Acknowledgements

[9.](#) Normative References

Moran, B., Tschofenig, H., Birkholz, H., and K. Zandberg, "A Concise Binary Object Representation (CBOR)-based Serialization Format for the Software Updates for Internet of Things (SUIT) Manifest", Work in Progress, Internet-Draft, [draft-ietf-suit-manifest-16](https://www.ietf.org/archive/id/draft-ietf-suit-manifest-16), 25 October 2021, <<https://www.ietf.org/archive/id/draft-ietf-suit-manifest-16.txt>>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

Authors' Addresses

Brendan Moran
Arm Limited

Email: Brendan.Moran@arm.com

Henk Birkholz
Fraunhofer SIT

Email: henk.birkholz@sit.fraunhofer.de