



Authenticated Debug Access Control Specification

| | |
|--------------------------|------------------|
| Document number | DEN0101 |
| Document version | 00bet0 |
| Document confidentiality | Non-confidential |
| Date of issue | 2020-10-12 |

Copyright © 2020 Arm Limited or its affiliates. All rights reserved.

BETA

This is a BETA version of the specification. This version is an engineering draft of the full specification. It is meant to obtain feedback from Arm partners and internally within Arm. It includes the majority of features but must not be treated as being complete. It should be treated as a work in progress and is subject to change on the basis of this feedback.

Authenticated Debug Access Control

Release information

| Date | Version | Changes |
|-------------|---------|--|
| 2020/Oct/12 | 00bet0 | <ul style="list-style-type: none">• First public version |

Arm Non-Confidential Document Licence (“Licence”)

This Licence is a legal agreement between you and Arm Limited (“**Arm**”) for the use of Arm’s intellectual property (including, without limitation, any copyright) embodied in the document accompanying this Licence (“**Document**”). Arm licenses its intellectual property in the Document to you on condition that you agree to the terms of this Licence. By using or copying the Document you indicate that you agree to be bound by the terms of this Licence.

“**Subsidiary**” means any company the majority of whose voting shares is now or hereafter owner or controlled, directly or indirectly, by you. A company shall be a Subsidiary only for the period during which such control exists.

This Document is **NON-CONFIDENTIAL** and any use by you and your Subsidiaries (“**Licensee**”) is subject to the terms of this Licence between you and Arm.

Subject to the terms and conditions of this Licence, Arm hereby grants to Licensee under the intellectual property in the Document owned or controlled by Arm, a non-exclusive, non-transferable, non-sub-licensable, royalty-free, worldwide licence to:

- (i) use and copy the Document for the purpose of designing and having designed products that comply with the Document;
- (ii) manufacture and have manufactured products which have been created under the licence granted in (i) above; and
- (iii) sell, supply and distribute products which have been created under the licence granted in (i) above.

Licensee hereby agrees that the licences granted above shall not extend to any portion or function of a product that is not itself compliant with part of the Document.

Except as expressly licensed above, Licensee acquires no right, title or interest in any Arm technology or any intellectual property embodied therein.

THE DOCUMENT IS PROVIDED “AS IS”. ARM PROVIDES NO REPRESENTATIONS AND NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY, SATISFACTORY QUALITY, NON-INFRINGEMENT OR FITNESS FOR A PARTICULAR PURPOSE WITH RESPECT TO THE DOCUMENT. Arm may make changes to the Document at any time and without notice. For the avoidance of doubt, Arm makes no representation with respect to, and has undertaken no analysis to identify or understand the scope and content of, third party patents, copyrights, trade secrets, or other rights.

NOTWITHSTANDING ANYTHING TO THE CONTRARY CONTAINED IN THIS LICENCE, TO THE FULLEST EXTENT PERMITTED BY LAW, IN NO EVENT WILL ARM BE LIABLE FOR ANY DAMAGES, IN CONTRACT, TORT OR OTHERWISE, IN CONNECTION WITH THE SUBJECT MATTER OF THIS LICENCE (INCLUDING WITHOUT LIMITATION) (I) LICENSEE’S USE OF THE DOCUMENT; AND (II) THE IMPLEMENTATION OF THE DOCUMENT IN ANY PRODUCT CREATED BY LICENSEE UNDER THIS LICENCE). THE EXISTENCE OF MORE THAN ONE CLAIM OR SUIT WILL NOT ENLARGE OR EXTEND THE LIMIT. LICENSEE RELEASES ARM FROM ALL OBLIGATIONS, LIABILITY, CLAIMS OR DEMANDS IN EXCESS OF THIS LIMITATION.

This Licence shall remain in force until terminated by Licensee or by Arm. Without prejudice to any of its other rights, if Licensee is in breach of any of the terms and conditions of this Licence then Arm may terminate this Licence immediately upon giving written notice to Licensee. Licensee may terminate this Licence at any time. Upon termination of this Licence by Licensee or by Arm, Licensee shall stop using the Document and destroy all copies of the Document in its possession. Upon termination of this Licence, all terms shall survive except for the licence grants.

Any breach of this Licence by a Subsidiary shall entitle Arm to terminate this Licence as if you were the party in breach. Any termination of this Licence shall be effective in respect of all Subsidiaries. Any rights granted to any Subsidiary hereunder shall automatically terminate upon such Subsidiary ceasing to be a Subsidiary.

The Document consists solely of commercial items. Licensee shall be responsible for ensuring that any use, duplication or disclosure of the Document complies fully with any relevant export laws and regulations to assure that the Document or any portion thereof is not exported, directly or indirectly, in violation of such export laws.

This Licence may be translated into other languages for convenience, and Licensee agrees that if there is any conflict between the English version of this Licence and any translation, the terms of the English version of this Licence shall prevail.

The Arm corporate logo and words marked with ® or ™ are registered trademarks or trademarks of Arm Limited (or its subsidiaries) in the US and/or elsewhere. All rights reserved. Other brands and names mentioned in this document may be the trademarks of their respective owners. No licence, express, implied or otherwise, is granted to Licensee under this

Licence, to use the Arm trade marks in connection with the Document or any products based thereon. Visit Arm's website at <http://www.arm.com/company/policies/trademarks> for more information about Arm's trademarks.

The validity, construction and performance of this Licence shall be governed by English Law.

Copyright © 2020 Arm Limited (or its affiliates). All rights reserved.

Arm Limited. Company 02557590 registered in England.

110 Fulbourn Road, Cambridge, England CB1 9NJ.

LES-PRE-21585 version 4.0

Contents

Authenticated Debug Access Control

| | |
|---|-----|
| Authenticated Debug Access Control | ii |
| Release information | ii |
| Arm Non-Confidential Document Licence ("Licence") | iii |

Preface

| | |
|-------------------------------------|------|
| About this document | viii |
| Conventions | ix |
| Typographical conventions | ix |
| Numbers | ix |
| C source code | ix |
| References | x |
| Feedback | xi |
| Feedback on this document | xi |

Part A Debug Access Control Architecture

Chapter A1

About the Architecture

| | |
|---|----|
| A1.1 About Platform Security Architecture | 14 |
| A1.2 Goals | 15 |
| A1.3 Scope | 16 |

Chapter A2

System Architecture

| | |
|---|----|
| A2.1 System architecture overview | 18 |
| A2.2 Target Architecture | 20 |
| A2.2.1 Access Control | 20 |
| A2.2.2 Target States | 20 |
| A2.2.3 Handling Reset | 20 |
| A2.3 Protocol Architecture | 21 |
| A2.3.1 Link Layer | 21 |
| A2.3.2 Command Protocol | 21 |
| A2.3.3 Authentication Command Set | 21 |

Chapter A3

Security Model

| | |
|---|----|
| A3.1 About the Security Model | 23 |
| A3.2 Threat Model | 24 |
| A3.3 Authentication | 25 |
| A3.4 Trust | 26 |
| A3.5 Constraints | 27 |
| A3.5.1 Scope-limiting constraints | 27 |
| A3.5.2 Permissions | 27 |
| A3.6 Examples | 29 |

Part B Specification

Chapter B1

Common Elements

| | |
|--|----|
| B1.1 Conventions | 32 |
| B1.1.1 Data encoding conventions | 32 |

| | | | |
|-------------------|--------|---|----|
| | B1.1.2 | C language conventions | 32 |
| B1.2 | | Version Numbers | 33 |
| B1.3 | | TLV Data Type | 34 |
| B1.4 | | Type ID Registry | 35 |
| B1.5 | | Key and Signature Types | 37 |
| | B1.5.1 | Cryptosystem Support | 38 |
| Chapter B2 | | Command Protocol | |
| B2.1 | | About the command protocol | 40 |
| B2.2 | | Protocol state machine | 41 |
| B2.3 | | Packets | 42 |
| | B2.3.1 | Request | 42 |
| | B2.3.2 | Response | 42 |
| B2.4 | | Error Handling | 43 |
| Chapter B3 | | Authentication Command Set | |
| B3.1 | | About the authentication commands | 45 |
| | B3.1.1 | Status codes | 45 |
| | B3.1.2 | Authentication response | 45 |
| | B3.1.3 | Authentication command sequence | 47 |
| B3.2 | | Commands | 48 |
| | B3.2.1 | Discovery | 48 |
| | B3.2.2 | Authentication Start | 49 |
| | B3.2.3 | Authentication Response | 49 |
| | B3.2.4 | Close Session | 50 |
| | B3.2.5 | Lock Debug | 51 |
| Chapter B4 | | ADAC Token | |
| B4.1 | | About the ADAC Token | 53 |
| B4.2 | | Format | 54 |
| B4.3 | | Extensions | 56 |
| | B4.3.1 | Allowed Extension Types | 56 |
| B4.4 | | Rules | 57 |
| | B4.4.1 | Construction | 57 |
| | B4.4.2 | Validation | 57 |
| Chapter B5 | | ADAC Certificate | |
| B5.1 | | About the ADAC Certificate | 59 |
| B5.2 | | Format | 60 |
| B5.3 | | Extensions | 63 |
| | B5.3.1 | Allowed Extension Types | 63 |
| B5.4 | | Rules | 64 |
| | B5.4.1 | Construction | 64 |
| | B5.4.2 | Validation | 64 |
| | B5.4.3 | Constraints | 65 |
| Part C | | Appendices | |
| Chapter C1 | | Example System Architectures | |
| C1.1 | | Example Arm Architecture Externally-hosted Target | 68 |
| | C1.1.1 | Debugger Mailbox | 69 |
| Chapter C2 | | Link Layer | |
| C2.1 | | About the link layer | 71 |
| C2.2 | | Link layers | 72 |

| | | |
|--------|--------------------------------------|----|
| C2.2.1 | COM Encapsulation Protocol | 72 |
| C2.2.2 | ACK Token | 72 |
| C2.2.3 | Memory Window | 73 |

Chapter C3

Cryptographic Support

| | | |
|--------|---|----|
| C3.1 | General concepts | 76 |
| C3.1.1 | Algorithm Agility | 76 |
| C3.1.2 | Key Sizes | 76 |
| C3.1.3 | Hash Function | 76 |
| C3.2 | ECDSA | 77 |
| C3.2.1 | P-256 Curve | 77 |
| C3.2.2 | P-521 Curve | 77 |
| C3.3 | RSA | 79 |
| C3.3.1 | RSA 3072-bit keys | 79 |
| C3.3.2 | RSA 4096-bit keys | 79 |
| C3.4 | EdDSA (tentative) | 80 |
| C3.4.1 | Ed25519 Curve | 80 |
| C3.4.2 | Ed448 Curve | 80 |
| C3.5 | ShangMi - SM2 (tentative) | 81 |
| C3.5.1 | SM2 | 81 |
| C3.6 | Secret key algorithms (tentative) | 82 |
| C3.6.1 | CMAC with AES | 82 |
| C3.6.2 | HMAC with SHA-256 | 82 |

Glossary

Preface

About this document

This is a specification document for ADAC.

This document has the following parts:

Part A

Describes the secure debug architecture.

Part B

Provides details of the specification.

Part C

Appendices.

Conventions

Typographical conventions

The typographical conventions are:

italic

Introduces special terminology, and denotes citations.

bold

Denotes signal names, and is used for terms in descriptive lists, where appropriate.

`monospace`

Used for assembler syntax descriptions, pseudocode, and source code examples.

Also used in the main text for instruction mnemonics and for references to other items appearing in assembler syntax descriptions, pseudocode, and source code examples.

SMALL CAPITALS

Used for some common terms such as IMPLEMENTATION DEFINED.

Used for a few terms that have specific technical meanings, and are included in the Glossary.

Red text

Indicates an open issue.

Blue text

Indicates a link. This can be

- A cross-reference to another location within the document
- A URL, for example <http://developer.arm.com>

Numbers

Numbers are normally written in decimal. Binary numbers are preceded by 0b, and hexadecimal numbers by 0x. In both cases, the prefix and the associated value are written in a monospace font, for example 0xFFFF0000. To improve readability, long numbers can be written with an underscore separator between every four characters, for example 0xFFFF_0000_0000_0000. Ignore any underscores when interpreting the value of a number.

C source code

This book contains numerous sections of C source code containing type definitions. These are shown in a monospace font.

References

This section lists publications by Arm and by third parties. See Arm Developer (<http://developer.arm.com>) for access to Arm documentation.

- [1] Arm Ltd, “Arm® Platform Security Architecture Security Model,” *DEN 0079*, Feb. 2019, [Online]. Available: <https://developer.arm.com/architectures/security-architectures/platform-security-architecture/documentation>.
- [2] Arm Ltd, “Advanced Communications Channel Architecture Specification,” *Arm IHI 0076A*, May 2018, [Online]. Available: <https://developer.arm.com/documentation/ih0076/a>.
- [3] National Institute of Standards and Technology, “Digital Signature Standard (DSS),” *Federal Information Processing Standards Publication (FIPS PUB) 186-4*, Jul. 2013, doi: [10.6028/NIST.FIPS.186-4](https://doi.org/10.6028/NIST.FIPS.186-4).
- [4] American National Standards Institute, “Public Key Cryptography for the Financial Services Industry: The Elliptic Curve Digital Signature Algorithm (ECDSA),” *ANSI X9.62-2005*, Nov. 2005, [Online]. Available: <https://standards.globalspec.com/std/1955141/ANSI%20X9.62>.
- [5] Standards for Efficient Cryptography Group, “Recommended Elliptic Curve Domain Parameters,” *SEC2*, [Online]. Available: <https://www.secg.org/sec2-v2.pdf>.
- [6] IETF, “Deterministic Usage of the Digital Signature Algorithm (DSA) and Elliptic Curve Digital Signature Algorithm (ECDSA),” *RFC6979*. <https://tools.ietf.org/html/rfc6979>.
- [7] IETF, “PKCS #1: RSA Cryptography Specifications Version 2.2,” *RFC8017*. <https://tools.ietf.org/html/rfc8017>.
- [8] IETF, “Edwards-Curve Digital Signature Algorithm (EdDSA),” *RFC8032*. <https://tools.ietf.org/html/rfc8032>.
- [9] International Organization for Standardization, “IT Security techniques – Digital signatures with appendix – Part 3: Discrete logarithm based mechanisms,” *ISO ISO/IEC 14888-3:2018*, Nov. 2018, [Online]. Available: <https://www.iso.org/standard/76382.html>.
- [10] International Organization for Standardization, “IT Security techniques – Hash-functions – Part 3: Dedicated hash-functions,” *ISO ISO/IEC 10118-3:2018*, Oct. 2018, [Online]. Available: <https://www.iso.org/standard/67116.html>.
- [11] National Institute of Standards and Technology, “Recommendation for Block Cipher Modes of Operation: The CMAC Mode for Authentication,” *NIST Special Publication 800-38B*, May 2005, doi: [10.6028/NIST.SP.800-38B](https://doi.org/10.6028/NIST.SP.800-38B).
- [12] IETF, “The AES-CMAC Algorithm,” *RFC4493*. <https://tools.ietf.org/html/rfc4493>.
- [13] IETF, “HMAC: Keyed-Hashing for Message Authentication,” *RFC2104*. <https://tools.ietf.org/html/rfc2104>.
- [14] IETF, “US Secure Hash Algorithms (SHA and SHA-based HMAC and HKDF),” *RFC6234*. <https://tools.ietf.org/html/rfc6234>.
- [15] National Institute of Standards and Technology, “Recommendation for Applications Using Approved Hash Algorithms,” *NIST Special Publication 800-107r1*, Aug. 2012, doi: [10.6028/NIST.SP.800-107r1](https://doi.org/10.6028/NIST.SP.800-107r1).

Feedback

Arm values the feedback of its partners.

Feedback on this document

If you have comments on the content of this document, send an e-mail to arm.psa-feedback@arm.com. Give:

- The title (Authenticated Debug Access Control).
- The number (DEN0101 00bet0).
- The page numbers to which your comments apply.
- A concise explanation of your comments.

Arm also welcomes general suggestions for additions and improvements.

Part A

Debug Access Control Architecture

Chapter A1

About the Architecture

This chapter describes the goals and scope of the ADAC architecture. It contains the following sections:

A1.1 About Platform Security Architecture

A1.2 Goals

A1.3 Scope

A1.1 About Platform Security Architecture

Arm's Platform Security Architecture is a holistic set of threat models, security analyses, hardware and firmware architecture specifications, an open source firmware reference implementation, and an independent evaluation and certification scheme. Platform Security Architecture provides a recipe, based on industry best practice, that allows security to be consistently designed in, at both a hardware and firmware level.

A1.2 Goals

Introducing security in debug is about making sure that only authorized people have access to select parts of firmware and hardware.

ADAC aims at making sure that debug capabilities do not become attack vectors. Debug security cannot be an afterthought when designing an SoC and the kind of debug solution needed is driven by the threat models for the device use case.

The ADAC architecture is designed to be flexible to meet varying vendor needs, adaptable to work with many different hardware and software components, and scalable from small embedded or IoT systems to complex server environments. At the same time, it strives to be simple and resilient against attack.

This requires:

- Strong authentication
- Partitioning firmware and hardware into fine-grained domains
- Enforcing debug limitations

ADAC needs to offer fine-grained access control of system resources, accessed through the debug port, across device lifecycle states.

This specification targets software-only layers that sit above the physical debug link. As such, any security aspects of the physical layer, including confidentiality and integrity, are out of scope for this document.

A1.3 Scope

The ADAC specification covers the following topics:

- Command protocol
- Debug authentication commands
- Certificate format
- Token format

The specification considers both device-level and application or software-level debug.

Chapter A2

System Architecture

This chapter describes the overall architecture of ADAC. The contains the following sections:

A2.1 System architecture overview

A2.2 Target Architecture

A2.3 Protocol Architecture

A2.1 System architecture overview

The high level ADAC system architecture is described by [Figure A2.1](#).

It is composed of these two systems:

- *Debug target*: The system or subsystem containing the resources for which permission to debug is requested through the secure debug protocol.
- *Debug host*: The device that initiates the debug session. It drives the authentication sequence and provides credentials to the debug target.

The debug host and debug target are connected via a debug link. This is any interface from which the debug host can perform debug operations on the target. This can be debug probe hardware driving a wire protocol such as SWD or JTAG, self-hosted debug capability between cores in a multi-core device, the self-hosted debug capability of a single processor debugging itself, or any similar debug link.

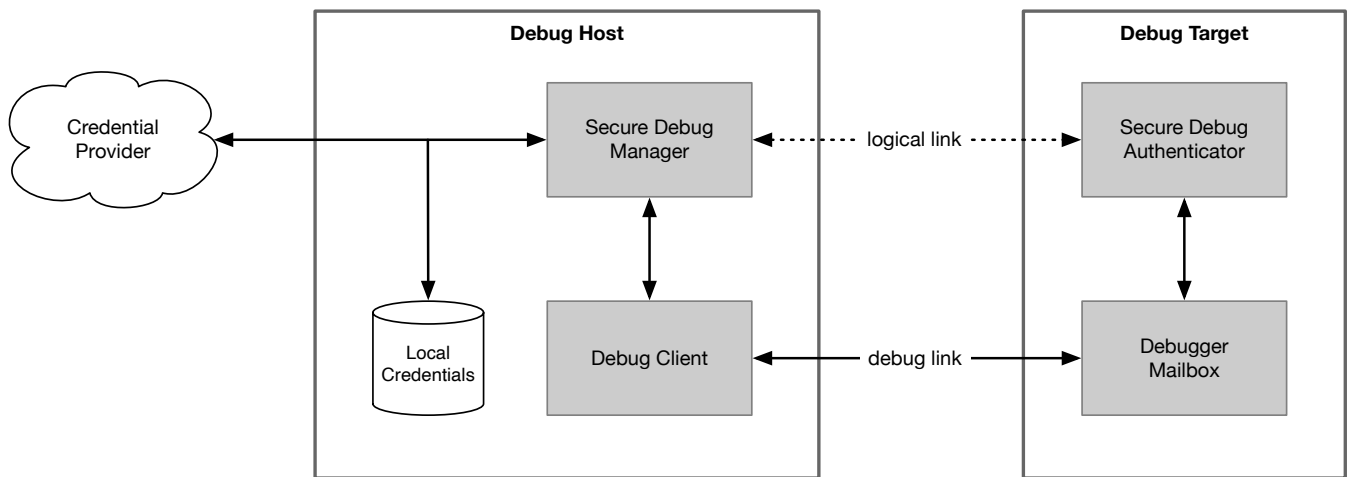


Figure A2.1: High level system block diagram.

The following components are used to establish the communications channel over which authentication is performed.

- *Debug client*: The component that drives the debug link between host and target (e.g., debugger software).
- *Debugger mailbox*: Generic term for a communications endpoint, inside the debug target, where the Secure Debug Authenticator (see definition below) can receive commands. The key attribute is that a debugger mailbox is available even when debug access is otherwise restricted due to the PSA security lifecycle. However, there can be other temporal restrictions on when the Secure Debug Authenticator is available to respond to requests.

The system architecture avoids placing requirements on the debug link and communications channel, except to define the debugger mailbox as the endpoint through which commands and responses are transferred. This allows a common architecture to support both externally-hosted debug and self-hosted debug.

For externally hosted debug, the debug link is typically in the form of a debug probe connected to the host with a high-bandwidth protocol such as USB or Ethernet.

The two primary components that implement this architecture are:

- *Secure Debug Manager (SDM)*: Initiates the logical communications link with the Secure Debug Authenticator on behalf of the debug client and manages the secure debug protocol. It receives credentials from either a local or remote credential provider, and passes those credentials to the Secure Debug Authenticator.

The debug client and SDM can communicate to allow a user to choose requested permissions and credentials, or the selection of permissions and credentials can be fully automated.

- *Secure Debug Authenticator (SDA)*: Accepts commands sent by the SDM via the debugger mailbox. It issues an authentication challenge and validates the credentials provided in response. Upon successful authentication, the SDA handles the hardware or software aspects of enabling debug access to target resources.

In addition, upon request it can provide the SDM with information about the debug target. This information can be used for discovery and identification purposes, to pass to the credential provider, to present the user with data to make an informed decision when choosing credentials, or other purposes.

The SDM resides on the debug host, and the SDA resides on the debug target. The two components establish a logical communications link between themselves. This logical link is routed through the debug client, debug link, and debugger mailbox.

The SDA must be contained within a trusted domain. It can run from several possible execution contexts:

- Immutable Root of Trust (e.g., boot ROM)
- Mutable Root of Trust (e.g., updatable bootloader)
- Trusted runtime service

Each execution context has distinct properties and capabilities that can influence the debug authentication sequence.

The execution context of the SDA is not required to be in the same system or PE over which it controls access. For instance, an SDA can execute from a secure enclave and control access only for the system outside the enclave.

A2.2 Target Architecture

A2.2.1 Access Control

For discussion of the target architecture, these terms are defined:

- *Access control signals*: Logical signals that control debug access to hardware or software components, or modify the functionality of components during the time when debug access is enabled. The signals themselves are IMPLEMENTATION DEFINED and can be implemented in either hardware or software.
- *Access control domain*: The system or subsystem over which the SDA controls debug access via the access control signals. An access control domain can be wholly software-defined, and only exist after a particular point in the boot process, such as in the case where the SDA is a trusted runtime service that only exists after the Application RoT is running.

There can be more than one access control domain in a debug target, and thus more than one SDA. It is acceptable for a single SDA to control more than one access control domain, but the method of selecting the access control domain for which the SDM is authenticating is IMPLEMENTATION DEFINED.

These are several examples of functions that access control signals can serve:

- Control debug access to a CPU or a security or privilege level within a CPU.
- Control access to a system memory bus.
- Control the availability of certain cryptographic keys in a hardware cryptography accelerator.
- Force the hiding of secrets in the root of trust.

The method used by the SDA to apply access control signals is IMPLEMENTATION DEFINED.

A2.2.2 Target States

An access control domain can be in one of the following states:

- Fully Locked
- Partially Unlocked
- Fully Unlocked

A2.2.3 Handling Reset

Note

Whether this section applies to an access control domain that only exists after the system boots to a certain stage is determined by the system architecture of the debug target.

There are two types of reset that impact the ADAC architecture:

- Cold Reset, often called a Power-On Reset.
- Warm Reset.

It is not possible to have a Cold reset without also having a Warm reset.

Access control signals must be reset to defined state, dependent on the device's security lifecycle state, on a Cold Reset. When in the Secured lifecycle state, a Cold reset should place the target in the Locked state.

If the target architecture allows, access control signals should remain unmodified through a Warm reset. This reduces the need for frequent authentication during target debug sessions. It can also allow for debug of the boot process. A target in the Partially or Fully Unlocked state can temporarily disable debug access during boot following a Warm reset to ensure that execution of the boot ROM is protected.

A2.3 Protocol Architecture

The ADAC specification defines the protocol used by the SDM to request debug access from the SDA. Several protocol layers are defined, as shown in [Figure A2.2](#).




Figure A2.2: Protocol layer stack.

Note that API layers are not included in this diagram.

The sections that follow describe each protocol layer in turn, from bottom up.

A2.3.1 Link Layer

The bottom layer in the stack is the debug link, which was defined earlier. The link layer sits atop the debug link and provides guarantees upon which the Command Protocol layer can be built.

Because the link layer is dependent upon the specific debug link, this specification does not require any one link layer. Several example link layers are documented in [Chapter C2 Link Layer](#).

A2.3.2 Command Protocol

The purpose of the Command Protocol is to abstract the communication channel between host and target.

The Command Protocol provides a simple request/response mechanism suitable for implementing the Authentication Command Set atop arbitrary link layers.

Vendors can extend the usefulness of the debugger mailbox by designing vendor-specific command sets that reuse the Command Protocol.

The Command Protocol is documented in [Chapter B2 Command Protocol](#).

A2.3.3 Authentication Command Set

The Authentication Command Set specifies a number of commands that implement ADAC. It defines these aspects of the protocol:

- The command sequence used to perform authentication.
- Status and error codes.
- Common data types.
- The set of supported data formats that implement the trust mechanism (see [Chapter A3 Security Model](#)).

The principal elements of the trust mechanism are the [ADAC Token](#) and [ADAC Certificate](#) formats. Other authentication and trust mechanisms or formats, either standard extensions or proprietary, can also be used.

The Authentication Command Set is documented in [Chapter B3 Authentication Command Set](#).

Chapter A3

Security Model

This chapter describes the security model for ADAC. The contains the following sections:

A3.1 About the Security Model

A3.2 Threat Model

A3.3 Authentication

A3.4 Trust

A3.5 Constraints

A3.6 Examples

A3.1 About the Security Model

Physical access to the target device being required for debug link operation, the main objective is for the target to securely authenticate the host. Authentication of the target by the host is not in scope for this document. Security of communications over the debug link, including confidentiality and integrity, are also not in scope for this document.

In order to simplify key management and improve security this specification focuses on asymmetric key cryptography.

A3.2 Threat Model

The goal of this specification is to offer an access control mechanism to prevent attackers with access to debug functionality from gaining access to devices resources.

The threat model does not consider the following attack vectors:

- Attacks on the physical debug interface.
- Other physical attacks on the device.
- Denial of service attacks.

A3.3 Authentication

The default mechanism for authentication (see [Chapter B4 ADAC Token](#)) relies on a challenge-response protocol. The challenge is used to protect against replay attacks.

Note

The challenge vector can be:

- a cryptographically random value,
- a random value (low entropy),
- a combination of different elements that will make the challenge to be different when performing authentication on distinct devices and when performing multiple authentication on the same device. Those elements can be (non-exhaustive list):
 - device unique constant values (e.g. serial number),
 - non-repeatable values (e.g. monotonic counter, secure time),
 - device constant values with some entropy (e.g. MAC address, manufacturing date),
 - non-constant values (e.g. high-precision counter or clock).

The randomness, non-malleability and unpredictability of the challenge vector is important to avoid the re-use of tokens for a given device or across devices. The options above are listed in a generic order of preference that might not apply to all cases.

The response to the challenge is a signed authentication token, also called the debug token in this specification. The key used to verify the signature must be trusted (see [A3.4 Trust](#)).

A3.4 Trust

In order to verify the signature of the authentication token, a (public) key is needed.

The simplest option has this key present on each device directly. This yields two extremes in terms of management options (or a combination of the two):

- Use the same key for all devices.
- Use a different key for each device.

A chain of certificates links the key used to sign to the authentication token to a set of one or more trusted anchors (roots of trust). Based on the vendor needs, the chain can be of arbitrary length, ending in a key directly linked to a programmed root authority (e.g. hash of public key(s) programmed into OTP).

Adding intermediate steps in the certificate chain adds some overhead to the authentication process in the form of extra verification operations and increased data size. However, intermediate certificates also limit the exposure of the most sensitive keys, allowing that those keys can be used less often and protected with added security (e.g., stored on offline/air-gapped systems or other physical protections).

This specification defines a certificate format (see [Chapter B5 ADAC Certificate](#)) to build trust chains and offer the flexibility to deal with complex scenarios (see [A3.5 Constraints](#)).

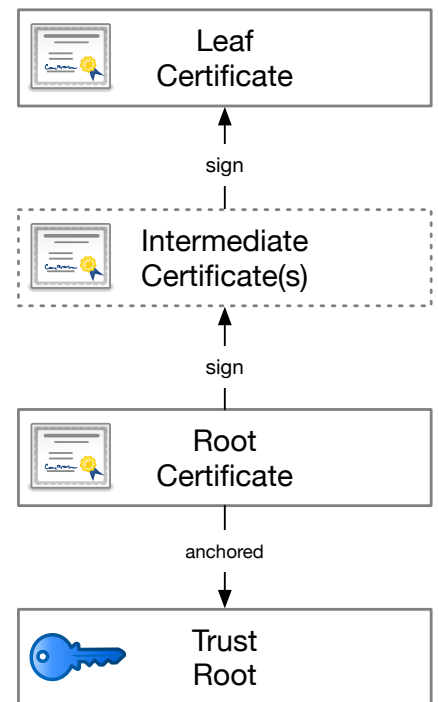


Figure A3.1: Example chain of trust.

The diagram in [Figure A3.1](#) shows an example trust chain composed of a leaf certificate, zero or more intermediate certificates, ending in a root certificate. The certificate chain is anchored to the device's root of trust.

A3.5 Constraints

Each certificate in the chain can add constraints to the authentication process in order to limit the scope of authentication and restrict permissions that its holder can unlock. This allows granting a specific set of privileges to a specific set of targets to either exercise or optionally to delegate further by issuing sub-certificates.

Two types of constraints are supported:

- Scope-limiting constraints. Described below.
- Permission-limiting constraints. See [A3.5.2 Permissions](#).

A3.5.1 Scope-limiting constraints

Several scope-limiting features are included in each certificate:

- `role`: Controls the ability of a certificate to sign other certificates.
- `lifecycle`: Limits to a specific lifecycle state.
- `soc_class`: Limits to specific vendor defined family of devices.
- `soc_id`: Limits to specific device.
- `custom_constraint`: Limits to devices with a matching customizable constraint value.

Using certificate extensions it is possible to add other types of constraints (e.g. `target_identity` or `sw_partition_id`).

All scope-limiting constraints have a neutral value to indicate no further restriction. Values for these constraints in certificates need to be consistent with the target configuration and between all certificates in the certificate chain:

- If its configuration or state does not match the expected value the target will reject the certificate.
- If two or more distinct non-neutral values are present in the certificate chain, a failure is triggered.

A3.5.2 Permissions

This specification supports two different permissions models:

- Logical permissions bits.
- Software partition permissions.

A3.5.2.1 Permissions bitmap

For fine-grained access-control, this specification defines a standard mechanism to associate certificates in the chain with a bitmap of logical permissions.

The debug host requests access to a set of permissions via the authentication token. The effective permissions are computed by masking requested permissions with permission-limiting constraints from the certificate chain. This mechanism allows for certification authorities to restrict permissions of issued certificates.

The exact semantics for the permissions is an IMPLEMENTATION DEFINED combination of SoC-specific access control signals.

Logical permissions do not necessarily map 1:1 to debug access signals. The Secure Debug Authenticator implementation can apply an IMPLEMENTATION DEFINED mapping operation to convert logical effective permissions to the debug access signals programmed then into system control registers. This allows for compression of debug access signals into a simpler set of permissions, as well as to ensure a consistent security configuration.

A value of 1 for a logical permission bit means that access is granted, while a value of 0 means that access is denied.

A3.5.2.1.1 Computing effective permissions

Definitions:

- `Perm_req`: Permissions vector requested by the debug host.

- **Perm_mask**: Mask of permissions allowed by the certificate chain to be enabled.
- **Perm_eff**: Final effective permissions after masking the requested permissions.
- **Soc_mask**: Static permissions mask value provided to the Secure Debug Authenticator.
- **Soc_value**: Static value used for permissions that, due to **Soc_mask**, are not allowed to be controlled via debug authentication.
- **Crt_count**: Number of certificates in the chain. The leaf certificate is at index 0, and the root certificate is at index **Crt_count** - 1.

Soc_mask and **Soc_value** are static values passed to the Secure Debug Authenticator from an IMPLEMENTATION DEFINED source. The intended use case is to allow the provisioning process of the target to set permanent restrictions on permissions requested by the host, and further, to set the values for those permissions that are restricted. These values can be programmed into an SoC's OTP memory or another trusted memory, fixed in hardware, set at software development time, or can simply be set to 0 if unavailable or not desired.

Steps to compute the effective permissions:

- The host requests a set of debug signals (**Perm_req**).
- The target combines the permission masks present in the certificates of the trust chain (**Perm_mask**).

```
let Perm_mask = ~0 // Initialize to all 1s.  
for n in 0..(Crt_count - 1)  
    let Perm_mask = Perm_mask & crt[n].permissions_mask
```

- Finally, the target computes the effective permissions (**Perm_eff**), merging with the SoC-programmed permission constraints (**Soc_mask** and **Soc_value**).

```
let Perm_eff = (Perm_req & Perm_mask & Soc_mask) | (Soc_value & ~Soc_mask)
```

The effective debug signals (**Sig_eff**) will then be used by the Secure Debug Authenticator to alter the debug configuration of the target system in an IMPLEMENTATION DEFINED manner.

A3.5.2.2 Software partition permissions

Access to software partitions can also be requested in the authentication token, and restricted by the certificate chain.

Each software partition is identified by a unique ID. The authentication token includes a list of zero or more software partition IDs for which access is requested.

Certificates also include a list of zero or more software partition IDs. The rules for interpretation are as follows:

- If a partition ID appears in a certificate, then the certificate is declaring that debug access to that software partition can be granted if requested in the authentication token.
- If at least one software partition ID is listed anywhere in the certificate chain, then access to only those software partitions whose IDs are listed in the certificate chain can be granted.
- If no software partition IDs are listed in the certificate chain, then debug access to software partitions is not constrained, and access to any partition whose ID is listed in the authentication token can be granted.

Other factors can influence whether permission to access a given software partition is granted. For instance, if the SoC-programmed permission constraints disallow any debugging of the SPE, then a request for debug access to an Application RoT partition must be disallowed. The interpretation of these additional factors and how they are applied (for instance, by hardware or software) is IMPLEMENTATION DEFINED.

The definition of software partition and the definition and size of software partition IDs is outside the scope of this specification.

A3.6 Examples

The following scenarios illustrate some of the flexibility of the authentication mechanism combined with certificates:

- Manufacturing equipment with a device-class certificate (and matching key stored in a hardware security module) is able to authenticate itself to the devices on the production line to initialize them (flash, credentials, root of trust. . .).
- A developer uses a device-locked certificate with a local key to debug their application on the device.
- A technician connects diagnostics equipment to a device, the authentication token is generated in the cloud to unlock access and perform maintenance.

Part B

Specification

Chapter B1

Common Elements

This chapter covers common definitions and conventions used throughout the ADAC specification. It contains the following sections:

[*B1.1 Conventions*](#)

[*B1.2 Version Numbers*](#)

[*B1.3 TLV Data Type*](#)

[*B1.4 Type ID Registry*](#)

[*B1.5 Key and Signature Types*](#)

B1.1 Conventions

B1.1.1 Data encoding conventions

The size of all outer-level aggregate data types must be 32-bit word aligned.

All multi-byte scalar values must be encoded in little-endian byte order.

Non-scalar, untyped, multi-byte arrays are encoded as little-endian byte arrays. As an example, take the value “DDCCBBAA8877665544332211”, written as a big-endian hex string. The most-significant byte, bits [95:88], has the value 0xDD. It is encoded as the byte sequence:

```
[ 11 22 33 44 55 66 77 88 99 AA BB CC DD ]
```

This table shows the same encoded value in several formats:

| Word Number | Bytes | 32-bit Integer |
|-------------|-----------------|----------------|
| 0 | [11 22 33 44] | 0x44332211 |
| 1 | [55 66 77 88] | 0x88776655 |
| 2 | [AA BB CC DD] | 0xDDCCBBAA |

B1.1.2 C language conventions

C structure definitions are used in this specification as a convenient syntax method for defining data types. However, the definitions should be considered pseudocode, as they can include non-conformant struct member definitions that are intended to better describe data sizes and relationships. In addition, all structures are defined as if packed (i.e., alignment is set to 1 byte), and include explicit padding to ensure the desired alignment of members.

C99 standard scalar types such as `uint32_t` are used for all integer values.

The following macros are used in definitions.

```
// Round a size *n* up to the nearest *r* multiple.
#define ROUND_UP(n, r) (((n) + (r) - 1) / (r) * (r))

// Round a size *n* up to the nearest 32-bit word.
#define ROUND_TO_WORD(n) (ROUND_UP((n), sizeof(uint32_t)))
```

B1.2 Version Numbers

Version numbers are used throughout the specification in order to allow for future changes while remaining backwards compatible. The authentication protocol itself has a version. And each major aggregate data type has a version number as its first member.

Version numbers consist of these two components:

- *Major version*: Must only be incremented when the data type is entirely redefined.
- *Minor version*: Must be incremented due to added or changed data type members, where the majority of the data type remains compatible.

Version numbers are encoded as a sequence of two 8-bit integers, with the major version first and minor version following. A C structure definition follows.

```
struct adac_version {  
    uint8_t major;  
    uint8_t minor;  
};
```

For example, version 1.2 is represented by the byte sequence [0x01 0x02].

Version number counting starts at version 1.0.

B1.3 TLV Data Type

A simple Type-Length-Value (TLV) data type is used in multiple places throughout the specification.

Each value consists of a 32-bit header with type ID and length, followed by the value data. The size of the entire TLV instance must be rounded up to the nearest 32-bit word.

Each TLV instance consists of the following fields:

| Name | Bytes | Description |
|---------|----------|--|
| Type ID | 2 | Unique identifier for the value type. |
| Length | 2 | Length of value in bytes. Does not include the size of any required padding. |
| Value | <i>n</i> | Value data. Must be padded with 0 to align on a 32-bit boundary. |

The C type definition for a TLV instance follows.

```
struct adac_tlv {  
    uint16_t type_id;  
    uint16_t length_in_bytes;  
    uint8_t value[ROUND_TO_WORD(length_in_bytes)];  
};
```

In standard usage within other data types, multiple TLV instances are placed back to back in a variable-length array. This construction is called a “TLV sequence”. Because the size of every TLV instance is rounded up to be 32-bit aligned, all TLV instances naturally start on 32-bit boundaries. The total size of the TLV sequence is specified by another member of the parent data type.

Nested TLV types are not used in this specification.

B1.4 Type ID Registry

The type ID for values in TLV instances is a 16-bit value. Any type ID with bit 15 set is vendor-specific. This leaves room for 32768 possible standard type IDs. However, related type IDs are grouped into consecutive ranges, so the ID space is sparsely populated.

The following table contains the complete list of type IDs used for TLV instances in this specification.

| ID | Name | Bytes | Description |
|--------|----------------------|--------------|---|
| 0x0000 | null_type | 0 | Indicates “no data”. |
| 0x0001 | adac_auth_version | 2 | Major and minor versions for ADAC. See Version Numbers . |
| 0x0002 | vendor_id | 2 | Vendor JEP106 ID. |
| 0x0003 | soc_class | 4 | SoC class. |
| 0x0004 | soc_id | 16 | SoC unique identifier. |
| 0x0005 | target_identity | <i>n</i> | Cryptographic identity for the target. |
| 0x0006 | hw_permissions_fixed | 16 | Value of hardware-fixed permissions. |
| 0x0007 | hw_permissions_mask | 16 | Mask of permissions bits allowed to be modified. |
| 0x0008 | psa_lifecycle | 2 | Current lifecycle state. |
| 0x0009 | sw_partition_id | <i>n</i> | Persistent, unique ID for a software partition. |
| 0x0100 | token_formats | 2 * <i>n</i> | Array of supported debug token formats. |
| 0x0101 | cert_formats | 2 * <i>n</i> | Array of token and certificate formats supported by the target, with each format being represented by a 16-bit type ID. |
| 0x0102 | cryptosystems | 1 * <i>n</i> | List of supported algorithms plus keys sizes. |
| 0x0200 | token_adac | <i>n</i> | ADAC Token format. |
| 0x0201 | cert_adac | <i>n</i> | ADAC Certificate format. |
| 0x0202 | rot_meta | <i>n</i> | Platform-specific metadata about the root of trust. |
| 0x8000 | | | Start of vendor-specific value type IDs. |

Note

Not all type IDs are accepted in all situations. The specification of each data type that uses TLV will contain a list of accepted type IDs.

Detailed information of each type ID follows.

- `null_type` : The ID 0x0000 is reserved and is used to represent “no data” or list termination in special cases.
- `adac_auth_version` : The [version number](#) for the authentication protocol command set. Currently defined as version 1.0.
- `vendor_id` : JEDEC JEP106 vendor ID. Bits [6:0] hold the “Identity Code” value; bits [15:7] contain the count of 0x7F Continuation Codes. For example, Arm’s `vendor_id` is 0x023B.

- `soc_class` : Vendor-unique identifier for the SoC part number and revision. The layout and meaning of any individual fields within `soc_class` is the responsibility of the vendor.
- `soc_id` : 128-bit identifier for the SoC. For a given root of trust, this value should be unique. This value is not sensitive. It is tied to the hardware and does not change for the life of the physical SoC. Often, the ID is a serial number composed of die and wafer coordinates plus a random number programmed into OTP memory by the silicon vendor.
- `target_identity` : Cryptographic identity for the target. The length of this value depends upon the cryptosystem used for its construction. Might not be available in all circumstances. For instance, a boot ROM might not have access to the information required to construct this value.
- `hw_permissions_fixed` : Bitfield setting the fixed value of any permissions bits whose corresponding bit in `hw_permissions_mask` is cleared.
- `hw_permissions_mask` : Hardware-defined permissions mask. When a permission bit is set to 0 in this field, it indicates that the hardware disallows modification of that permission via the authentication protocol—the permission is always fixed to the value of the corresponding bit in `hw_permissions_fixed`.
- `psa_lifecycle` : Represents the current lifecycle state of the PSA RoT. The state is represented by an integer that is divided to convey a major state and a minor state. A major state is defined by PSA-SM[1]. A minor state is optional, and is IMPLEMENTATION DEFINED. The PSA security lifecycle state and implementation state are encoded in Trusted Firmware-M (<https://www.trustedfirmware.org>) as follows:
 - `version[15:8]` – PSA security lifecycle state
 - * `PSA_LIFECYCLE_UNKNOWN` (0x0000)
 - * `PSA_LIFECYCLE_ASSEMBLY_AND_TEST` (0x1000)
 - * `PSA_LIFECYCLE_PSA_ROT_PROVISIONING` (0x2000)
 - * `PSA_LIFECYCLE_SECURED` (0x3000)
 - * `PSA_LIFECYCLE_NON_PSA_ROT_DEBUG` (0x4000)
 - * `PSA_LIFECYCLE_RECOVERABLE_PSA_ROT_DEBUG` (0x5000)
 - * `PSA_LIFECYCLE_DECOMMISSIONED` (0x6000)
 - `version[7:0]` – IMPLEMENTATION DEFINED state.
- `sw_partition_id` : This value uniquely identifies a software partition. The length and semantics of a partition ID value are not determined by this specification.
- `token_formats` : Array of 16-bit type IDs indicating which token formats are supported by the debug target.
- `cert_formats` : Array of 16-bit type IDs indicating which certificate formats are supported by the debug target.
- `cryptosystems` : Array of 8-bit [cryptosystem IDs](#) indicating those cryptosystems supported by the debug target.
- `token_adac` : The ADAC Token format.
- `cert_adac` : The ADAC Certificate format.
- `rot_meta` : Optional vendor-defined data required by a Secure Debug Authenticator to verify the provided public root key matches the hardware Root of Trust Public Key (ROTPK).

B1.5 Key and Signature Types

The specification currently supports the following key types:

- ECDSA P-256 (see [C3.2.1 P-256 Curve](#))
- ECDSA P-521 (see [C3.2.2 P-521 Curve](#))
- RSA 3072 (see [C3.3.1 RSA 3072-bit keys](#))
- RSA 4096 (see [C3.3.2 RSA 4096-bit keys](#))
- Ed25519 (see [C3.4.1 Ed25519 Curve](#))
- Ed448 (see [C3.4.2 Ed448 Curve](#))
- SM2 (see [C3.5.1 SM2](#))
- CMAC (see [C3.6.1 CMAC with AES](#))
- HMAC (see [C3.6.2 HMAC with SHA-256](#))

The specification currently supports the following signature types:

- ECDSA P-256 with SHA-256
- ECDSA P-521 with SHA-512
- RSA 3072 with SHA-256
- RSA 4096 with SHA-256
- Ed25519 with SHA-512
- Ed448 with SHAKE256
- SM2 with SM3
- CMAC with AES
- HMAC with SHA-256

Currently a given key type only supports one signature algorithm.

The list of accepted cryptosystem combinations and the unique constants for identifying each follows.

A cryptosystem ID is an 8-bit value. Vendor-defined cryptosystems are allowed by setting the MSB (bit 7) of the ID.

| ID | Name | Public Key | Signature Algorithm |
|------|---------------------------------------|----------------------|----------------------|
| 0x01 | ECDSA_P256_SHA256 | Elliptic Curve P-256 | ECDSA with SHA-256 |
| 0x02 | ECDSA_P521_SHA512 | Elliptic Curve P-521 | ECDSA with SHA-512 |
| 0x03 | RSA_3072_SHA256 | RSA (3072-bit key) | RSA-PSS with SHA-256 |
| 0x04 | RSA_4096_SHA256 | RSA (4096-bit key) | RSA-PSS with SHA-256 |
| 0x05 | ED_25519_SHA512 | Ed25519 | EdDSA with SHA-512 |
| 0x06 | ED_448_SHAKE256 | Ed448 | EdDSA with SHAKE256 |
| 0x07 | SM_SM2_SM3 | SM2 | SM2 with SM3 |
| 0x08 | CMAC_AES | Nonce | CMAC with AES |
| 0x09 | HMAC_SHA256 | Nonce | HMAC with SHA-256 |
| 0x80 | Start of vendor-defined cryptosystems | | |

The IDs listed above are used in several places:

- The `cryptosystems` value (see [above](#)).
- Signature algorithm ID in the ADAC Token header.
- Key and signature algorithm ID in the ADAC Certificate header.

B1.5.1 Cryptosystem Support

Debug targets are expected to support only one or a small number of related cryptosystems. For instance, a debug target might support only `ECDSA_P256_SHA256`, or could support both `RSA_3072_SHA256` and `RSA_4096_SHA256`.

It is likely that the debug host supports more cryptosystems than the target. On the other hand, vendor-specific implementations of the Secure Debug Manager can choose to implement only those cryptosystems known to be supported by that vendor's target-side implementation.

Chapter B2

Command Protocol

This chapter specifies the generic high level command protocol. It contains the following sections:

B2.1 About the command protocol

B2.2 Protocol state machine

B2.3 Packets

B2.3.1 Request

B2.3.2 Response

B2.4 Error Handling

B2.1 About the command protocol

The Debug Mailbox functionality can and has been implemented in many different ways. In order to abstract this implementation aspect, this document defines a command protocol. The link layer and wire protocol are left to be implementation specific.

In order to support most existing debug mailbox solutions and not create additional requirements for future ones, the goal is for the command protocol to be extremely simple. This facilitates simple target-side software, which is important for constrained systems and execution environments such as a boot ROM. Critically, it also makes reasoning about the security aspects of an implementation much easier.

Key attributes of the command protocol:

- Variable size, word aligned packets.
- Simple request/response model.

Non-requirements:


- Error correction.
- Flow control.
- Multiple in-flight commands.
- Out of order packets.
- Fixed size packets.

The link layer is assumed to provide error correction and flow control where necessary. Some link layers additionally provide their own packetization method.

All packets are 32-bit word aligned. This easily supports debugger mailboxes that transfer either byte or word sized data. It is assumed that commands can easily arrange for word aligned parameters.

B2.2 Protocol state machine

The following diagram depicts the state machine for the command protocol.



The states are labeled from the point of view of the debug host. For the debug target, the sending and receiving roles are reversed.

In this protocol, the debug host is always the master and initiator of commands.

B2.3 Packets

B2.3.1 Request

A request consists of a single-word header word containing the command ID and parameter length in words, followed by the optional parameter words. The complete request size must be word aligned.

A request must only be sent from the Idle protocol state.

The high bit (bit 15) of the command ID indicates a vendor-specific command.

| Word | Byte 0 | Byte 1 | Byte 2 | Byte 3 |
|------|--------------|---------------|-----------------|------------------|
| 0 | command[7:0] | command[15:8] | data_count[7:0] | data_count[15:8] |
| 1 | data | ... | ... | ... |

The C structure definition for a request is as follows:

```
struct request_packet {  
    uint16_t command;  
    uint16_t data_count;  
    uint32_t data[];  
};
```

B2.3.2 Response

Similar to a request, a response packet has a single header word containing the command status and response data word count. Following the header is the optional response data. As with requests, the total size must be word aligned.

A response packet is always a reply to the most recent request. Because the protocol does not allow for out of order or asynchronous commands, there is no need to include a copy of the command ID in the header or a sequence number.

The response to a command must be sent only in the Response Pending protocol state by the receiver of the request. No intervening packets, sent by either endpoint, are allowed between the request and response. Once the response is received, the protocol transitions back to the Idle state, and another request can be sent.

| Word | Byte 0 | Byte 1 | Byte 2 | Byte 3 |
|------|-------------|--------------|-----------------|------------------|
| 0 | status[7:0] | status[15:8] | data_count[7:0] | data_count[15:8] |
| 1 | data | ... | ... | ... |

The C structure definition for a response is as follows:

```
struct response_packet {  
    uint16_t status;  
    uint16_t data_count;  
    uint32_t data[];  
};
```

B2.4 Error Handling

Error handling capabilities of the command protocol are intentionally limited. It is assumed that either a method specific to the link layer or a system reset can be used to restore the communications channel functionality.

The status value of 0x7FFF is reserved for unrecognized commands. If a request with an unrecognized command ID is received by the target, it must reply with a response packet consisting of a status value of 0x7FFF and no data words. As a byte sequence, this error packet is [0xFF 0x7F 0x00 0x00].

Chapter B3

Authentication Command Set

This chapter specifies the debug authentication commands and their parameters. It contains the following sections:

B3.1 About the authentication commands

B3.1.1 Status codes

B3.1.2 Authentication response

B3.1.3 Authentication command sequence

B3.2 Commands

B3.2.1 Discovery

B3.2.2 Authentication Start

B3.2.3 Authentication Response

B3.2.4 Close Session

B3.2.5 Lock Debug

B3.1 About the authentication commands

ADAC defines the following set of commands for performing debug authentication. These commands are defined as a layer building on the [Command Protocol](#).

All commands below must be implemented by the debug target. However, some commands can return a status code indicating they are unsupported on the target or in the execution environment. This is documented per command.

Detailed specifications for each command follow in this chapter.

| ID | Constant | Command | Description |
|--------|-----------------------|---|---|
| 0x0001 | SDP_DISCOVERY_CMD | Discovery | Query target properties. |
| 0x0002 | SDP_AUTH_START_CMD | Authentication Start | Initiate authentication sequence; receive challenge. |
| 0x0003 | SDP_AUTH_RESPONSE_CMD | Authentication Response | Send authentication data. |
| 0x0004 | SDP_CLOSE_SESSION_CMD | Close Session | Terminate command session. |
| 0x0005 | SDP_LOCK_DEBUG_CMD | Lock Debug | Lock debug access; restore system to locked state. |

Additional command sets can be defined in further specifications.

Vendor-specific commands can be added by the integrator. All vendor-specific commands must have bit 15 set in the command ID.

Currently defined versions of the ADAC protocol are as follows. The version is reported via the `auth_version` value returned by the [Discovery](#) command.

| Version | Description |
|---------|------------------|
| 1.0 | Initial version. |

B3.1.1 Status codes

The following table lists the complete set of status codes returned by the authentication commands.

| Status | Constant | Description |
|--------|----------------------|---|
| 0x0000 | SDP_SUCCESS | The command has succeeded without error. |
| 0x0001 | SDP_FAILURE | The command has failed. |
| 0x0002 | SDP_MORE_DATA_NEEDED | More data is required to complete the authentication. |
| 0x0003 | SDP_UNSUPPORTED | The command is not supported by the target. |
| 0x7FFF | SDP_INVALID_COMMAND | The command ID is unrecognized. |

The status code 0x7FFF is special in that it is not returned by a specific command but instead indicates an unrecognized command ID. See the [Error Handling](#) section of the command protocol for more information.

B3.1.2 Authentication response

A complete authentication response consists of a sequence of separate cryptographically signed data structures including one or more certificates and the debug token. It can additionally include vendor-specific credentials or cryptographic material required for the target to verify the certificate chain or debug token.

The certificate chain is required to be provided in order from root to leaf.

B3.1.2.1 Response fragments

A response fragment is defined as one of the data structures composing the authentication response. The primary classes of response fragment are the debug token and certificate. Each class of response fragment has a set of valid formats in which it can be represented. Every supported response fragment format has a unique [type ID](#).

The debug target is not required to support all possible response fragment formats, and indeed is expected to only support a small subset. The debug host can use the [Discovery](#) command to query the target for supported response fragment formats. This process is intended to be used as additional validation rather than for protocol negotiation.

Type IDs for accepted response fragment formats are listed in the following table.

| Value | Name | Description |
|--------|----------------------------|------------------------------|
| 0x0200 | token_adac | ADAC Token format |
| 0x0201 | cert_adac | ADAC Certificate format |
| 0x0203 | rot_meta | Vendor-specific RoT metadata |

B3.1.2.2 Example authentication response

As an example, an authentication response can be constructed from this sequence of response fragments:

1. ADAC Certificate: Root certificate.
2. ADAC Certificate: Leaf certificate.
3. ADAC Token

The following diagram shows the relationships between the response fragments in this example.

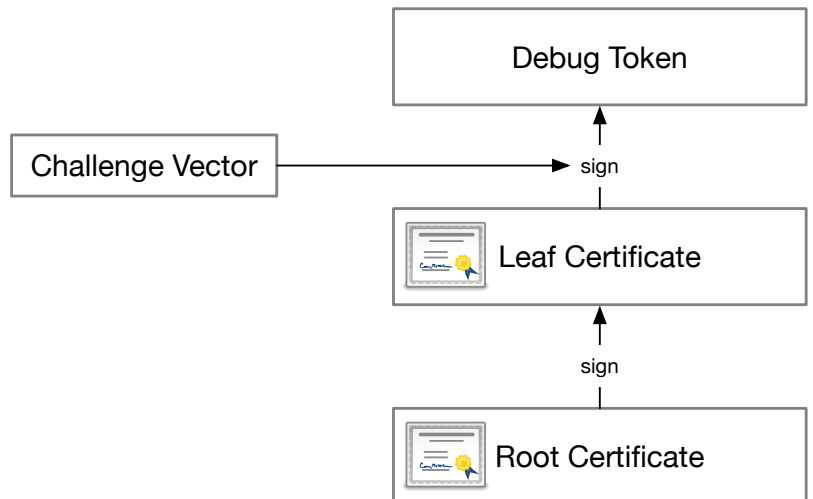


Figure B3.1: Example Authentication Response

The number of certificates used is a decision made by the customer to trade off security versus processing time and management complexity. This example uses two certificates; three is also common. Use of a single certificate is possible but not recommended.

As can be seen, the root certificate signs the leaf certificate. The leaf certificate then signs the debug token and challenge vector.

Note that the root certificate is tied to the target's hardware root of trust. Normally the root certificate contains the target's ROTPK.

B3.1.3 Authentication command sequence

The process of enabling debug access to system resources is accomplished by sending a sequence of the commands defined in this chapter. Only some of the commands defined in this chapter are required as part of the authentication sequence; the others are optionally used for gathering information about the target or requesting a change in system state that does not require authentication.

The required sequence of commands for authentication is as follows.

1. **Authentication Start**: Host requests challenge from target.
2. **Authentication Response** (one or more): Host sends certificate chain (possibly other response fragments) and debug token to target.

The number of **Authentication Response** commands sent is determined by the number of response fragments from which the complete authentication response is constructed.

If **Authentication Response** return the `SDP_MORE_DATA_NEEDED` status code, then the target expects another **Authentication Response** command to be sent in order to continue or complete the authentication sequence.

Multiple authentication sequences are supported by the specification, but whether allowed this is allowed is IMPLEMENTATION DEFINED. This can be used for various purposes, such as to unlock more than one access control domain, potentially with multiple roots of trust, to gain access to more than one software partition, or other purposes.

More than one authentication sequence in parallel is not allowed. This means that once an authentication sequence is started, no intervening commands are allowed to be issued until the sequence completes successfully or fails due to an error.

After all intended authentication sequences are complete, the **Close Session** command must be sent in order to terminate communications.

Note

The benefits of issuing one command per response fragment are two-fold:

1. Reduced target memory requirements by supporting incremental processing.
 2. Allows for the possibility of early error termination.
-

B3.2 Commands

B3.2.1 Discovery

The debug host requests information about the debug target using with this command.

The debug host can optionally include a list of [type IDs](#) for values requested from the target. The requested ID list is discretionary; the target can reply with any set of values that is equal to or a super set of the requested values, excluding any values unavailable on the target or not recognized by the target.

If a requested ID list is not included, the target must reply with all available values.

Use of this command is optional and is not part of the required authentication command sequence. More than one [Discovery](#) command is allowed to be sent by the host.

B3.2.1.1 Request

Command ID SDP_DISCOVERY_CMD (0x0001)

Request data Array of requested [type IDs](#).

The request data is an optional array of requested [type IDs](#), each a 16-bit half-word.

If the number of requested IDs is not even, then an extra `null_type` ID with value 0x0000 is appended to round up to the request data length to the next 32-bit word as required by the [command protocol](#). If the `null_type` ID is present in the requested ID array at any position other than the last, it terminates processing of the array early.

The request sequence must be ordered by increasing ID value.

B3.2.1.2 Response

Response data [TLV sequence](#).

Possible Status Codes

SDP_SUCCESS (0x0000) A valid discovery response follows.

SDP_FAILURE (0x0001) Discovery failed.

The response consists of a [TLV sequence](#). This provides a simple mechanism for managing variable length values. It also allows vendors to extend the response with custom information if needed. Because this command is not critical to security of the overall protocol, the required parsing is not a concern. In addition, all response parsing happens on the host.

Possible response value [type IDs](#):

| Value | Name |
|--------|--------------------------------------|
| 0x0001 | adac_auth_version |
| 0x0002 | vendor_id |
| 0x0003 | soc_class |
| 0x0004 | soc_id |
| 0x0005 | target_identity |
| 0x0006 | hw_permissions_fixed |
| 0x0007 | hw_permissions_mask |
| 0x0008 | psa_lifecycle |
| 0x0100 | token_formats |

| Value | Name |
|--------|-------------------------------|
| 0x0101 | cert_formats |
| 0x0102 | cryptosystems |

The response sequence must be ordered by increasing ID value.

B3.2.2 Authentication Start

The debug host sends this commands to start the authentication sequence. Its primary purpose is for the target to provide a random 256-bit challenge vector used to prevent replay attacks.

B3.2.2.1 Request

Command ID SDP_AUTH_START_CMD (0x0002)

Request data None

B3.2.2.2 Response

Response data adac_auth_challenge struct.

Possible Status Codes

SDP_SUCCESS (0x0000) A valid challenge structure follows.

SDP_FAILURE (0x0001) Target is unable to send a challenge.

The response data consists of the following versioned structure.

```
struct adac_auth_challenge_v1_0 {  
    struct adac_version format_version;  
    uint16_t _reserved;  
    uint8_t challenge_vector[32];  
};
```

Note that the `challenge_vector` field is encoded as a multi-byte array, as described under [B1.1.1 Data encoding conventions](#).

Currently defined versions of this structure are as follows.

| Version | Description |
|---------|------------------|
| 1.0 | Initial version. |

B3.2.3 Authentication Response

This command is used to provide debug token and additional credentials as part of a complete authentication response to the target. One or more [Authentication Response](#) commands must occur to form the complete authentication response

This command will return an `SDP_MORE_DATA_NEEDED` status code to indicate that further data is required to complete the authentication and thus other [Authentication Response](#) must be sent. This sequence will continue until all required data has been provided to allow the target to validate the trust chain and debug token.

The target validates the provided debug token. If validation fails, the `SDP_FAILURE` status is returned.

B3.2.3.1 Request

Command ID `SDP_AUTH_RESPONSE_CMD` (0x0003)

Request data `adac_auth_fragment_header` plus response fragment

The data for **Authentication Response** request phase consists of a 32-bit header specifying the type ID of the included response fragment followed by the entire response fragment for the debug token.

The C definition of the response is as follows:

```

struct adac_auth_fragment_header {
    uint16_t fragment_type_id;
    uint16_t _reserved;
};

struct adac_auth_fragment {
    struct adac_auth_fragment_header header;
    uint8_t data[];
};

```

B3.2.3.2 Response

Response data None

Possible Status Codes

| | |
|--|--|
| <code>SDP_SUCCESS</code> (0x0000) | Authentication has succeeded. The authentication sequence is complete. |
| <code>SDP_FAILURE</code> (0x0001) | Authentication failed. |
| <code>SDP_MORE_DATA_NEEDED</code> (0x0003) | More data is required to complete the authentication sequence. |

B3.2.4 Close Session

This command requests that the communications session between the Secure Debug Manager and the Secure Debug Authenticator be terminated.

Depending on the execution environment of the Secure Debug Authenticator, closing the session can cause system boot to continue. Any means of resuming execution is acceptable, including performing a system reset. Note that if a system reset is used to resume execution, it should be performed after sending the response. Implementations should take link layer specifics into account to ensure a good likelihood of the response being delivered successfully.

B3.2.4.1 Request

Command ID `SDP_CLOSE_SESSION_CMD` (0x0004)

Request data None

B3.2.4.2 Response

Response data None

Possible Status Codes

| | |
|-----------------------------------|--|
| <code>SDP_SUCCESS</code> (0x0000) | The communications session was closed. |
|-----------------------------------|--|

B3.2.5 Lock Debug

The **Lock Debug** command restores the the device's debug access controls to the Fully Locked state, given its current lifecycle state.

Not all targets will have the capability to lock debug access after it is unlocked without going through a Power-On Reset.

B3.2.5.1 Request

Command ID SDP_LOCK_DEBUG_CMD (0x0005)

Request data None

B3.2.5.2 Response

Response data None

Possible Status Codes

SDP_SUCCESS (0x0000) Debug access is now locked.

SDP_UNSUPPORTED (0x0004) Debug access cannot be locked.

Chapter B4

ADAC Token

This chapter specifies the structure of the binary debug authentication token. It contains the following sections:

[*B4.1 About the ADAC Token*](#)

[*B4.2 Format*](#)

[*B4.3 Extensions*](#)

[*B4.3.1 Allowed Extension Types*](#)

[*B4.4 Rules*](#)

B4.1 About the ADAC Token

The ADAC Token is part of the authentication mechanism defined for ADAC. A token is sent by the debug host in response to a challenge vector (sometimes call nonce) sent by the target, and is cryptographically linked to the challenge vector and Root of Trust.

The ADAC Token also contains debug access permissions requested by the debug host.

A proprietary binary token format is used in order simplify requirements for parsing.

B4.2 Format

The components of a ADAC Token are the following:

- **Header:** The header contains all mandatory fields and the length of extensions in 32-bit words. The size of the header remains constant for different cryptosystems.
- **Extensions hash:** Hash of Extensions sequence. The algorithm and length depend on the Signature Type field in Header. Value is all zeros if extensions length is zero.
- **Signature:** The signature is performed over Header and Extension Hash parts of the Token as well as the challenge sent by the target.
- **Extensions:** [TLV](#) sequence of non-mandatory and vendor-specific fields.

Mandatory fields contained in the header:

- **format_version:** See [Version](#).
- **signature_type:** [Cryptosystem ID](#) specifying the algorithm used to generate the token's signature, as well as the extensions hash.
- **requested_permissions:** Bitfield of requested debug permissions. A set bit indicates a request for the permission corresponding to that bit to be granted.

The data layout for the token header is shown in the following table.

| Word | Byte 0 | Byte 1 | Byte 2 | Byte 3 |
|------|--|----------------------|------------------|--------|
| 0 | format_version.major | format_version.minor | signature_type | (0) |
| 1 | (0) | (0) | extensions_words | |
| 2 | requested_permissions[31:0] ¹ | | | |
| 3 | requested_permissions[63:32] ¹ | | | |
| 4 | requested_permissions[95:64] ¹ | | | |
| 5 | requested_permissions[127:96] ¹ | | | |

¹ Encoded as a multi-byte array, detailed in [B1.1.1 Data encoding conventions](#).

The following C structures describe the token header.

```

struct adac_debug_auth_token_header_v1_0 {
    struct adac_version format_version;
    uint8_t signature_type;
    uint8_t _reserved1;
    uint16_t _reserved2;
    uint16_t extensions_words;
    uint8_t requested_permissions[16];
};

struct adac_debug_token_v1_0 {
    struct adac_debug_auth_token_header_v1_0 header;
    uint8_t extension_hash[HASH_LENGTH];
    uint8_t signature[SIGNATURE_LENGTH];
    uint32_t extension_data[header.extensions_words]; // array of variable-length
                                                    // adac_tlv structs
};

```

Currently defined versions of the `adac_debug_token_v1_0` structure are as follows.

| Version | Description |
|---------|------------------|
| 1.0 | Initial version. |

B4.3 Extensions

The extensions for a ADAC Token are structured as a [TLV sequence](#). In the `adac_debug_token_v1_0` struct, the extensions TLV sequence is placed in the `extensions_data` member.

The size of the extensions data is specified in the certificate header `extensions_words` member as a number of 32-bit words.

B4.3.1 Allowed Extension Types

The following table lists those [Type IDs](#) that are accepted in the extensions data for a ADAC Certificate. Any extension value with a Type ID not included within this list will be ignored.

| Type ID | Name | Description |
|---------|---------------------------------|-----------------------|
| 0x0005 | target_identity | Target identity |
| 0x0009 | sw_partition_id | Software partition ID |

The [sw_partition_id](#) extension specifies a software partition to which access is granted. This extension value can be included more than once, in which case access is granted to all listed software partitions.

B4.4 Rules

B4.4.1 Construction

The hash of the token extensions is computed as follows, using the hash algorithm identified by the `header.signature_type` cryptosystem.


```
extensions_hash = Hash(extensions_data)
```

The signature over the token is computed as follows. The signature algorithm used is specified by the `header.signature_type` cryptosystem.

```
signature = Sign(leaf_cert.private_key, header || extensions_hash ||  
challenge_vector)
```

In the above expression, the symbol `leaf_cert` refers to the leaf certificate that signs the token.

The following diagram shows the inputs to the token signature algorithm.



B4.4.2 Validation

The `header.format_version` member can be used to select an appropriate struct definition for parsing the entire header.

These members of the `adac_debug_token_v1_0` struct must be validated before any further processing of the certificate is performed:

- `header.format_version`
- `header.signature_type`

Chapter B5

ADAC Certificate

This chapter specifies the structure of binary certificates used in ADAC. It contains the following sections:

B5.1 About the ADAC Certificate

B5.2 Format

B5.3 Extensions

B5.3.1 Allowed Extension Types

B5.4 Rules

B5.1 About the ADAC Certificate

The high complexity of X.509v3 certificates, in addition to being costly (effort, code size, RAM, etc.), has been the source of bugs and security issues.

Due to the low bandwidth of some of the underlying transports and the potential resource constraints of the target, this specification recommends the use of the alternative, purpose-built certificate format described in this chapter.

A ADAC Certificate is an element of the chain of trust. It also contains a set of optional constraints applied to debug authentication and debug permissions.

B5.2 Format

The components of a ADAC Certificate are the following:

- *Header*: The header contains all mandatory fields and the length of the extensions in 32-bit words. The size of the header remains constant for all cryptosystems.
- *Extensions hash*: Hash over extensions data. Algorithm and length depend on Signature Type field in Header. Value is all zeros if extensions length is zero.
- *Public key*: Content and length depend on Key Type field in Header.
- *Signature*: The signature is performed over Header, Extension Hash, and Public Key.
- *Extensions*: TLV sequence of non-mandatory and vendor-specific fields.

Mandatory fields contained in the header:

- `format_version`: See [Version](#).
- `signature_type`: [Cryptosystem ID](#) for the algorithm used to generate the certificate's signature and extensions hash.
- `key_type`: [Cryptosystem ID](#) indicating the algorithm and key size for the public key contained in the certificate.
- `role`: Certificate role. Whether the certificate is a root, intermediate, or leaf certificate. The table below lists accepted values.
- `usage`: Certificate usage. Specifies additional operational usage expressed by the certificate, if any. See the table below for defined usage values.
- `lifecycle`: Restricts authentication to a particular PSA lifecycle state. Encoding is the same as the [psa_lifecycle](#) data type.
- `custom_constraint`: Customizable constraint bitfield. Semantics are defined by the integrator and used to apply additional constraints on authentication.
- `soc_id`: Device unique ID. See the [soc_id](#) data type.
- `soc_class`: Vendor-defined device family ID. See the [soc_class](#) data type.
- `permissions_mask`: Bit mask limiting the permissions that can be requested in the ADAC Token. A set bit indicates that the permission corresponding to that bit position is allowed by the containing certificate to be requested in the token. The full certificate chain and hardware-defined permissions mask must be taken into account to determine the final permissions mask.

The data layout for the certificate header is shown in the following table.

| Word | Byte 0 | Byte 1 | Byte 2 | Byte 3 |
|------|---------------------------------------|----------------------|------------------|----------|
| 0 | format_version.major | format_version.minor | signature_type | key_type |
| 1 | role | usage | lifecycle | |
| 2 | customer_constraint | | extensions_words | |
| 3 | soc_class | | | |
| 4 | soc_id[31:0] ¹ | | | |
| 5 | soc_id[63:32] ¹ | | | |
| 6 | soc_id[95:64] ¹ | | | |
| 7 | soc_id[127:96] ¹ | | | |
| 8 | permissions_mask[31:0] ¹ | | | |
| 9 | permissions_mask[63:32] ¹ | | | |
| 10 | permissions_mask[95:64] ¹ | | | |
| 11 | permissions_mask[127:96] ¹ | | | |

¹ Encoded as a multi-byte array, detailed in [B1.1.1 Data encoding conventions](#).

The following C structures describe the certificate header:

```

struct adac_certificate_header_v1_0 {
    struct adac_version format_version;
    uint8_t signature_type;
    uint8_t key_type;
    uint8_t role;
    uint8_t usage;
    uint16_t lifecycle;
    uint16_t custom_constraint;
    uint16_t extensions_words;
    uint32_t soc_class;
    uint8_t soc_id[16];
    uint8_t permissions_mask[16];
};

struct adac_certificate_v1_0 {
    struct adac_certificate_header_v1_0 header;
    uint8_t extension_hash[HASH_LENGTH];
    uint8_t public_key[KEY_LENGTH];
    uint8_t signature[SIGNATURE_LENGTH];
    uint32_t extension_data[header.extensions_words]; // array of variable-length
                                                    // adac_tlv structs
};

```

Currently defined versions of the `adac_certificate_v1_0` structure are as follows.

| Version | Description |
|---------|------------------|
| 1.0 | Initial version. |

The `role` member of the header must have one of the following values.

| Value | Constant | Description |
|-------|-----------------------|--|
| 0x1 | SDP_ROLE_ROOT | The certificate is a root certificate. |
| 0x2 | SDP_ROLE_INTERMEDIATE | The certificate is intermediate. |
| 0x3 | SDP_ROLE_LEAF | The certificate is a leaf certificate. |

The `usage` member of the header must have one of the following values.

| Value | Constant | Description |
|-------|--------------------|--|
| 0x1 | SDP_USAGE_STANDARD | The certificate has no special usage. |
| 0x2 | SDP_USAGE_RMA | The certificate moves the device to the RMA lifecycle state. |

B5.3 Extensions

The extensions for a ADAC Certificate are structured as a [TLV sequence](#). In the `adac_certificate_v1_0` struct, the extensions TLV sequence is placed in the `extensions_data` member.

The size of the extensions data is specified in the certificate header as a number of 32-bit words.

B5.3.1 Allowed Extension Types

The following table lists those [Type IDs](#) that are accepted in the extensions data for a ADAC Certificate. Any extension value with a Type ID not included within this list will be ignored.

| Type ID | Name | Description |
|---------|---------------------------------|-----------------------|
| 0x0005 | target_identity | Target identity |
| 0x0009 | sw_partition_id | Software partition ID |

The [sw_partition_id](#) extension specifies a software partition to which access is granted. This extension value can be included more than once, in which case access is granted to all listed software partitions.

B5.4 Rules

B5.4.1 Construction

The hash of the certificate extensions is computed as follows, using the hash algorithm identified by the `header`.
→ `signature_type` cryptosystem.


```
extensions_hash = Hash(extensions_data)
```

The signature over the certificate is computed as follows, using the signature algorithm specified by the `header`.
→ `signature_type` cryptosystem.

```
signature = Sign(ca.private_key, header || extensions_hash || public_key)
```

In the above expression, the symbol `ca` refers to the signer certificate. For a root certificate, `ca` is the same as the certificate being signed (self-signing).

The following diagram shows the inputs to the certificate signature algorithm.



B5.4.2 Validation

The `header.format_version` member can be used to select an appropriate struct definition for parsing the entire header.

These members of the `adac_certificate_v1_0` struct must be validated before any further processing of the certificate is performed:

- `header.format_version`
- `header.signature_type`
- `header.key_type`
- `header.role`
- `header.usage`

- `header.lifecycle`

When processing the complete certificate chain, all non-zero values of a given constraint must be equal in every certificate.

B5.4.3 Constraints

When these members of the `adac_certificate_v1_0` struct are set to all zero bits, they are ignored and do not constrain authentication.

- `header.soc_id`
- `header.soc_class`
- `header.lifecycle`
- `header.custom_constraint`

As defined for the `psa_lifecycle` data type, the `header.lifecycle` is composed of a major PSA state and minor IMPLEMENTATION DEFINED state. If `header.lifecycle` is non-zero, the major state must also be non-zero and authentication is restricted to the specified major state. The minor state is always optional; if non-zero, authentication is restricted to the specified minor state in addition to the major state restriction. If the minor state is zero, then any minor state is allowed. A `header.lifecycle` value with a zero major state and non-zero minor state is invalid.

The `header.custom_constraint` field is compared to the static custom constraint value provided to the Secure Debug Authenticator by an IMPLEMENTATION DEFINED mechanism. If the two values match then authentication is not constrained, otherwise authentication fails.

Part C

Appendices

Chapter C1

Example System Architectures

This chapter gives several examples of possible system architectures for ADAC. The contains the following sections:

C1.1 Example Arm Architecture Externally-hosted Target

C1.1 Example Arm Architecture Externally-hosted Target

This section demonstrates an example mapping of this specification to the Arm architecture. The example shown here focuses on externally hosted debug.

Devices based on the Arm architecture vary considerably depending on the size and complexity of the device. However, there are a number of key components that any device with secure debug will have. These are shown in [Figure C1.1](#).

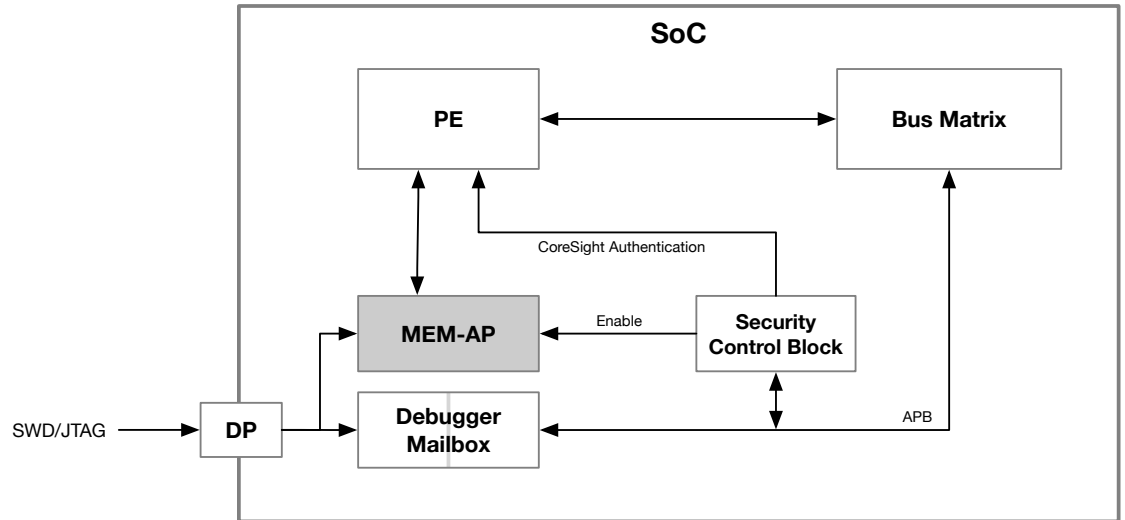


Figure C1.1: Example externally-hosted Arm target block diagram.

As defined by the Arm® Debug Interface Architecture Specification (ADI), the external interface for debugger access is called the Debug Port (DP). The SWJ-DP is an implementation that supports both SWD and JTAG wire protocols. The DP connects to one or more Access Port (AP) components. A device will have at least one MEM-AP that provides the debugger with the ability to perform transactions on the internal bus fabric and control system and PE-level debug logic.

A standard set of four debug access signals called CoreSight authentication signals are available for controlling the level of debug access for each PE:

- **DBGEN**: Invasive Non-secure debug access
- **NIDEN**: Non-invasive Non-secure debug access
- **SPIDEN**: Invasive Secure debug access, **DBGEN** must also be asserted
- **SPNIDEN**: Non-invasive Secure debug access, **NIDEN** or **DBGEN** must also be asserted

Not all systems will have all four authentication signals. For instance, a system that does not implement a Secure PE will not have **SPIDEN** and **SPNIDEN**.

A debug target in the Secured lifecycle state has these four standard authentication signals disabled by default.

In addition to the standard CoreSight authentication signals, each AP also has a its own enable signal. In a Cortex-M system where the AP is routed through the PE, it is not strictly necessary to disable the AP in the Secured lifecycle state. But in larger systems where a MEM-AP is directly connected to the bus fabric, it is a requirement.

Depending on the debug target's system-level architecture, the debug access signals can be connected in different ways. For instance, a multicore device can expose the CoreSight authentication signals for each PE, or they can be shared. In addition, the debug target can define additional, proprietary security control signals. As an example, a cryptographic accelerator peripheral can accept a signal that controls access to and use of device-unique key(s).

The Security Control Block is IP used by the Secure Debug Authenticator to modify the access control signals. Access to the Security Control Block is restricted to trusted software.

C1.1.1 Debugger Mailbox

A type of debugger mailbox that fits well with the ADiv5 architecture is a special type of AP, composed of two sides. One side connects to the DP and is accessible by the debugger as an AP. The other side is a standard APB peripheral. The two sides are themselves connected and can perform byte or word transfers bidirectionally.

A similar debugger mailbox can be built for the ADiv6 architecture, but both sides of the debugger mailbox are APB slaves.

The Arm SDC-600 COM Port is an example of such an debugger mailbox. It is available in versions that support both ADiv5 and ADiv6.

Chapter C2

Link Layer

This appendix documents several common link layer protocols used for ADAC communications channels. It contains the following sections:

C2.1 About the link layer

C2.2 Link layers

C2.2.1 COM Encapsulation Protocol

C2.2.2 ACK Token

C2.2.3 Memory Window

C2.1 About the link layer

The link layer is the protocol layer specific to the communications channel implementation. The primary purpose is to establish a common set of capabilities upon which the higher level protocols can build. It is focused on delivery of packets between the two physical endpoints. Each different communication channel has its own specific link layer protocol.

The link layer provides some or all of these properties for the higher level protocol layers:

- Method to request and initiate communications
- Flow control
- Packetization

Some communication channels have an intrinsic link layer protocol that provides the required properties, and so do not require an additional link layer to be used.

This chapter describes several link layers used for common types of communications channels.

- COM Encapsulation Protocol for Arm SDC-600 COM Port.
- ACK token protocol for common types of proprietary debugger mailbox IP.
- Memory window protocol for devices using a restricted window into system memory.

C2.2 Link layers

C2.2.1 COM Encapsulation Protocol

The COM Encapsulation Protocol is a byte-oriented protocol for transferring data packets across a COM Port interface such as the Arm SDC-600. The protocol is fully defined in the Arm Advanced Communications Channel Architecture Specification[2].

A protocol discovery sequence is defined as part of the COM Encapsulation Protocol. This allows the target to report a unique protocol ID to the host to declare its expected higher level protocol using a method independent of the higher level protocol.

For ADAC, the unique protocol ID is 4 bytes in length, and is defined as shown here:

| Format | Protocol ID |
|---------------|---------------|
| ASCII | 'ADAC' |
| Byte sequence | [41 44 41 43] |

Note

This protocol ID is currently provisional.

Any other protocol ID response sent by the target indicates that a different command protocol is expected, and the host should behave accordingly.

The protocol discovery sequence as byte values are defined as follows. The indicated direction is relative to the host being master.

| Direction | Protocol | Data |
|-----------|---------------|---------------|
| Request | - IDR | [A0] |
| Response | - IDA | [A1] |
| | - Protocol ID | [41 44 41 43] |
| | - END | [AD] |

The COM Encapsulation Protocol fully specifies the method of packetizing the higher level command requests and responses transferred by the ADAC Command Protocol.

Command protocol requests and replies must be sent in least-significant byte-first (LSB-first) order.

C2.2.2 ACK Token

Some SoCs include a simple debugger mailbox that implements transfer of a single 32-bit word at a time in either direction. Most often the hardware does not mandate a particular link layer protocol. For such IP, the ACK token protocol can be used to provide a software-implemented flow control mechanism. An alternative protocol is to make use of status flags in the IP registers, if provided.

In response to each request or reply data word, the other side must send an ACK Token, used for flow control. The request and reply header words do not require an ACK Token to be sent; the reply acts as the ACK for the request.

The upper 16-bits are set by the receiver (the side sending the ACK Token) with number of remaining words

expected to be sent. The lower 16-bits are always set to 0xA5A5. This value should be avoided for a valid command ID or command status value.

| Byte 0 | Byte 1 | Byte 2 | Byte 3 |
|--------|--------|-------------------|--------------------|
| 0xA5 | 0xA5 | remain_count[7:0] | remain_count[15:8] |

The C structure definition for a ACK Token is as follows:

```
struct ack_token {
    uint16_t token;           /* must be set to 0xAFAF */
    uint16_t remain_count;    /* remaining word count */
};
```

C2.2.3 Memory Window

The memory window link layer is designed to be as simple as possible. Performance is not a key concern. The only requirement is that the memory window support both read and write.

```
enum {
    MW_HOST_DONE = 0x12121212,
    MW_TARGET_DONE = 0xEFEFEFEF,
    MW_PATTERN1 = 0xFF00FF00,
    MW_PATTERN2 = 0x00FF00FF
};

struct memory_window {
    uint32_t status[4];
    uint32_t message[];
};
```

C2.2.3.1 Handshake

- The host initiates, writing to the status array the sequence [MW_PATTERN1, MW_PATTERN2, ↪MW_PATTERN1, MW_PATTERN2].
- The target acknowledges, writing to the status array the sequence [MW_PATTERN1, MW_PATTERN2, ↪MW_PATTERN1, MW_PATTERN2].

C2.2.3.2 Message exchange

It is assumed that the [Chapter B2 Command Protocol](#) and its sequence will be used.

- The host writes a [B2.3.1 Request](#) packet in the message member.
- The host writes the MW_HOST_DONE value in the status[0] member.
- Reading the value MW_HOST_DONE in the status[0] member signals to the target that it can read the [B2.3.1 Request](#) packet from the message member.
- Once the [B2.3.1 Request](#) packet is processed, the target writes [B2.3.2 Response](#) packet in the message member.
- The target writes the MW_TARGET_DONE value in the status[0] member.
- Reading the value MW_TARGET_DONE in the status[0] member signals to the host that it can read the [B2.3.2 Response](#) packet from the message member.

Note

Both [B2.3.1 Request](#) and [B2.3.2 Response](#) packets encode their respective lengths.

Chapter C3

Cryptographic Support

This appendix documents cryptographic suites currently defined for ADAC. It contains the following sections:

[*C3.1 General concepts*](#)

[*C3.2 ECDSA*](#)

[*C3.3 RSA*](#)

[*C3.4 EdDSA \(tentative\)*](#)

[*C3.5 ShangMi - SM2 \(tentative\)*](#)

C3.1 General concepts

C3.1.1 Algorithm Agility

Algorithm agility is an important feature of security protocols, which is the reason this specification provides a list of different families of asymmetric key cryptographic algorithms. That list of supported algorithms is extensible.

The specification (and reference implementation) recommends and only specifies solutions using the same algorithm and key size. This due to security-sensitive nature of the operations and the constraints of the target-side implementations of the protocols. Supporting multiple cryptographic algorithms and key sizes adds complexity and code size.

C3.1.2 Key Sizes

We have defined for each cryptographic algorithm two public key sizes one that matches the minimum publicly recommended sizes, as well as higher level for high or long term assurance.

C3.1.3 Hash Function

With each public key algorithm and key size is associated a hash function.

C3.2 ECDSA

The Elliptic Curve Digital Signature Algorithm (ECDSA) as defined in FIPS-186-4[3] and X9.62-2005[4] used in conjunction with the following curves:

- NIST P-256 curve (also designated `secp256r1` in SEC2[5] or `prime256v1` in X9.62-2005[4])
- NIST P-521 curve (also designated `secp521r1` in SEC2[5])

The public keys are encoded in uncompressed format (without the `0x04` typically used in formats that allow compressed representations).

The signatures do not include the ASN.1 DER encoding.

We recommend using the deterministic variant of ECDSA (see [6]) for generating signature, this is transparent to the implementation on the target.

C3.2.1 P-256 Curve

```
#define ECDSA_P256_PUBLIC_KEY_SIZE 64
#define ECDSA_P256_SIGNATURE_SIZE 64
#define ECDSA_P256_HASH_SIZE 32
#define ECDSA_P256_HASH_ALGORITHM PSA_ALG_SHA_256
#define ECDSA_P256_SIGN_ALGORITHM PSA_ALG_DETERMINISTIC_ECDSA(PSA_ALG_SHA_256)

typedef struct {
    certificate_header_t header;
    uint8_t pubkey[ECDSA_P256_PUBLIC_KEY_SIZE]; // P-256 public key
    uint8_t extensions_hash[ECDSA_P256_HASH_SIZE]; // SHA-256 hash
    uint8_t signature[ECDSA_P256_SIGNATURE_SIZE]; // P-256 with SHA-256 signature
    uint32_t extensions[];
} certificate_p256_t;

typedef struct {
    token_header_t header;
    uint8_t extensions_hash[ECDSA_P256_HASH_SIZE]; // SHA-256 hash
    uint8_t signature[ECDSA_P256_SIGNATURE_SIZE]; // P-256 with SHA-256 signature
    uint32_t extensions[];
} token_p256_t;
```

C3.2.2 P-521 Curve

```
#define ECDSA_P521_PUBLIC_KEY_SIZE 132
#define ECDSA_P521_SIGNATURE_SIZE 132
#define ECDSA_P521_HASH_SIZE 64
#define ECDSA_P521_HASH_ALGORITHM PSA_ALG_SHA_512
#define ECDSA_P521_SIGN_ALGORITHM PSA_ALG_DETERMINISTIC_ECDSA(PSA_ALG_SHA_512)

typedef struct {
    certificate_header_t header;
    uint8_t pubkey[ROUND_TO_WORD(ECDSA_P521_PUBLIC_KEY_SIZE)]; // P-521 public key
    uint8_t extensions_hash[ECDSA_P521_HASH_SIZE]; // SHA-512 hash
    uint8_t signature[ROUND_TO_WORD(ECDSA_P521_SIGNATURE_SIZE)]; // P-521 with SHA
    ↪-512 signature
    uint32_t extensions[];
} certificate_p521_t;

typedef struct {
    token_header_t header;
```

```
uint8_t extensions_hash[ECDSA_P521_HASH_SIZE]; // SHA-512 hash
uint8_t signature[ECDSA_P521_SIGNATURE_SIZE]; // P-521 with SHA-512 signature
uint32_t extensions[];
} token_p521_t;
```

C3.3 RSA

RSA (see 7) is included in specification.

This specification forces the use of the value F_4 (65537 in decimal, 0x10001 in hexadecimal) for exponent.

The public keys are encoded as the raw value of the modulus (without the leading zero mandated by ASN.1 DER encoding).

Signatures use the Probabilistic Signature Scheme.

C3.3.1 RSA 3072-bit keys

```
#define RSA_3072_PUBLIC_KEY_SIZE 384
#define RSA_3072_SIGNATURE_SIZE 384
#define RSA_3072_HASH_SIZE 32
#define RSA_3072_HASH_ALGORITHM PSA_ALG_SHA_256
#define RSA_3072_SIGN_ALGORITHM PSA_ALG_RSA_PSS(PSA_ALG_SHA_256)

typedef struct {
    certificate_header_t header;
    uint8_t pubkey[RSA_3072_PUBLIC_KEY_SIZE]; // RSA 3072-bit public key
    uint8_t extensions_hash[RSA_3072_HASH_SIZE]; // SHA-256 hash
    uint8_t signature[RSA_3072_SIGNATURE_SIZE]; // RSA with SHA-256 signature
    uint32_t extensions[];
} certificate_rsa3072_rsa3072_t;

typedef struct {
    token_header_t header;
    uint8_t extensions_hash[RSA_3072_HASH_SIZE]; // SHA-256 hash
    uint8_t signature[RSA_3072_SIGNATURE_SIZE]; // RSA with SHA-256 signature
    uint32_t extensions[];
} token_rsa3072_t;
```

C3.3.2 RSA 4096-bit keys

```
#define RSA_4096_PUBLIC_KEY_SIZE 512
#define RSA_4096_SIGNATURE_SIZE 512
#define RSA_4096_HASH_SIZE 32
#define RSA_4096_HASH_ALGORITHM PSA_ALG_SHA_256
#define RSA_4096_SIGN_ALGORITHM PSA_ALG_RSA_PKCS1V15_SIGN(PSA_ALG_SHA_256)

typedef struct {
    certificate_header_t header;
    uint8_t pubkey[RSA_4096_PUBLIC_KEY_SIZE]; // RSA 4096-bit public key
    uint8_t extensions_hash[RSA_4096_HASH_SIZE]; // SHA-256 hash
    uint8_t signature[RSA_4096_SIGNATURE_SIZE]; // RSA with SHA-256 signature
    uint32_t extensions[];
} certificate_rsa4096_rsa4096_t;

typedef struct {
    token_header_t header;
    uint8_t extensions_hash[RSA_4096_HASH_SIZE]; // SHA-256 hash
    uint8_t signature[RSA_4096_SIGNATURE_SIZE]; // RSA with SHA-256 signature
    uint32_t extensions[];
} token_rsa4096_t;
```

C3.4 EdDSA (tentative)

Edwards-Curve Digital Signature Algorithm (EdDSA), see 8.

C3.4.1 Ed25519 Curve

```
#define EDDSA_ED25519_PUBLIC_KEY_SIZE 32
#define EDDSA_ED25519_SIGNATURE_SIZE 64
#define EDDSA_ED25519_HASH_SIZE 64
#define EDDSA_ED25519_HASH_ALGORITHM PSA_ALG_SHA_512
#define EDDSA_ED25519_SIGN_ALGORITHM PSA_ALG_EDDSA_PH(PSA_ALG_SHA_512) // Not
    ↪defined yet

typedef struct {
    certificate_header_t header;
    uint8_t pubkey[ROUND_TO_WORD(EDDSA_ED25519_PUBLIC_KEY_SIZE)];
    uint8_t extensions_hash[EDDSA_ED25519_HASH_SIZE];
    uint8_t signature[ROUND_TO_WORD(EDDSA_ED25519_SIGNATURE_SIZE)];
    uint32_t extensions[];
} certificate_ed25519_t;

typedef struct {
    token_header_t header;
    uint8_t extensions_hash[EDDSA_ED25519_HASH_SIZE]; // SHA-512 hash
    uint8_t signature[EDDSA_ED25519_SIGNATURE_SIZE]; // Ed25519 signature
    uint32_t extensions[];
} token_ed25519_t;
```

C3.4.2 Ed448 Curve

```
#define EDDSA_ED448_PUBLIC_KEY_SIZE 57
#define EDDSA_ED448_SIGNATURE_SIZE 114
#define EDDSA_ED448_HASH_SIZE 64
#define EDDSA_ED448_HASH_ALGORITHM PSA_ALG_SHAKE256 // Not defined yet
#define EDDSA_ED448_SIGN_ALGORITHM PSA_ALG_EDDSA_PH(PSA_ALG_SHAKE256) // Not
    ↪defined yet

typedef struct {
    certificate_header_t header;
    uint8_t pubkey[ROUND_TO_WORD(EDDSA_ED448_PUBLIC_KEY_SIZE)];
    uint8_t extensions_hash[EDDSA_ED448_HASH_SIZE];
    uint8_t signature[ROUND_TO_WORD(EDDSA_ED448_SIGNATURE_SIZE)];
    uint32_t extensions[];
} certificate_ed448_t;

typedef struct {
    token_header_t header;
    uint8_t extensions_hash[EDDSA_ED448_HASH_SIZE]; // SHAKE256 hash
    uint8_t signature[EDDSA_ED448_SIGNATURE_SIZE]; // Ed448 signature
    uint32_t extensions[];
} token_ed448_t;
```

C3.5 ShangMi - SM2 (tentative)

SM2 is a set of elliptic curve based cryptographic algorithms including digital signature (see 9). SM2 is used with the SM3 hash function (see 10).

C3.5.1 SM2

```
#define SM2_SM3_PUBLIC_KEY_SIZE 64
#define SM2_SM3_SIGNATURE_SIZE 64
#define SM2_SM3_HASH_SIZE 32
#define SM2_SM3_HASH_ALGORITHM PSA_ALG_SM3 // Not defined yet
#define SM2_SM3_SIGN_ALGORITHM PSA_ALG_SM2 // Not defined yet

typedef struct {
    certificate_header_t header;
    uint8_t pubkey[SM2_SM3_PUBLIC_KEY_SIZE]; // SM2 public key
    uint8_t extensions_hash[SM2_SM3_HASH_SIZE]; // SM3 hash
    uint8_t signature[SM2_SM3_SIGNATURE_SIZE]; // SM2 with SM3 signature
    uint32_t extensions[];
} certificate_sm2sm3_sm2sm3_t;

typedef struct {
    token_header_t header;
    uint8_t extensions_hash[SM2_SM3_HASH_SIZE]; // SM3 hash
    uint8_t signature[SM2_SM3_SIGNATURE_SIZE]; // SM2 with SM3 signature
    uint32_t extensions[];
} token_sm2sm3_t;
```

C3.6 Secret key algorithms (tentative)

C3.6.1 CMAC with AES

CMAC with AES is authentication algorithm based on CMAC with the 128-bit Advanced Encryption Standard (AES), see [11](#) and [12](#).

```
#define CMAC_PUBLIC_KEY_SIZE 16
#define CMAC_SIGNATURE_SIZE 16
#define CMAC_HASH_SIZE 16
#define CMAC_HASH_ALGORITHM PSA_ALG_CMAC
#define CMAC_SIGN_ALGORITHM PSA_ALG_CMAC

typedef struct {
    certificate_header_t header;
    uint8_t pubkey[CMAC_PUBLIC_KEY_SIZE]; // Nonce
    uint8_t extensions_hash[CMAC_HASH_SIZE]; // CMAC
    uint8_t signature[CMAC_SIGNATURE_SIZE]; // CMAC
    uint32_t extensions[];
} certificate_cmac_cmac_t;

typedef struct {
    token_header_t header;
    uint8_t extensions_hash[CMAC_HASH_SIZE]; // CMAC
    uint8_t signature[CMAC_SIGNATURE_SIZE]; // CMAC
    uint32_t extensions[];
} token_cmac_t;
```

C3.6.2 HMAC with SHA-256

HMAC a mechanism for message authentication using cryptographic hash functions. See [13](#), [14](#) and [15](#)

```
#define HMAC_PUBLIC_KEY_SIZE 32
#define HMAC_SIGNATURE_SIZE 32
#define HMAC_HASH_SIZE 32
#define HMAC_HASH_ALGORITHM PSA_ALG_SHA_256
#define HMAC_SIGN_ALGORITHM PSA_ALG_HMAC(PSA_ALG_SHA_256)

typedef struct {
    certificate_header_t header;
    uint8_t pubkey[HMAC_PUBLIC_KEY_SIZE]; // Nonce
    uint8_t extensions_hash[HMAC_HASH_SIZE]; // SHA-256 hash
    uint8_t signature[HMAC_SIGNATURE_SIZE]; // HMAC-SHA-256
    uint32_t extensions[];
} certificate_hmac_hmac_t;

typedef struct {
    token_header_t header;
    uint8_t extensions_hash[HMAC_HASH_SIZE]; // SHA-256 Hash
    uint8_t signature[HMAC_SIGNATURE_SIZE]; // HMAC-SHA-256
    uint32_t extensions[];
} token_hmac_t;
```

Glossary

AP

Access Port.

APB

Advanced Peripheral Bus – Low speed, low complexity bus for peripherals.

CMSIS

Cortex Microcontroller Software Interface Standard – a vendor independent hardware abstraction layer for Cortex-M processors.

DAP

Debug Access Port - A block that acts as a master on a system bus and provides access to the bus from an external debugger.

DCU

Debug Control Unit.

Debug Client

Software on the debug host that controls the debug link.

Debug Host

The master component that performs debug operations on the debug target.

Debug Link

The connection between debug host and debug target through over which the debug client performs debug operations.

Debug Target

The slave component which is controlled by the debug host.

Debugger Mailbox

Generic term for a communications channel between a debug host and a software agent running on the device being debugged. The actual hardware IP consists of an AP on the debugger (external) side connected to an APB peripheral on the internal side.

DP

Debug Port.

ICV

IC Vendor, aka Silicon Partner (SiP).

IFR

Indexed Flash Region, an NXP term for reserved flash regions with special purposes. Usually not directly programmable by the OEM.

JTAG

Joint Test Action Group - An IEEE group focussed on silicon chip testing methods. Many debug and programming tools use a JTAG interface port to communicate with processors

NSPE

Non-Secure Processing Environment.

OEM

Original Equipment Manufacturer, the device owner.

PKI

Public Key Infrastructure.

ROTPK

Root of Trust public key. A public key programmed into immutable memory of a device.

Secure Debug Authenticator

Component residing in the debug target that receives and verifies requests to unlock debug access.

Secure Debug Manager

Component residing in the debug host that asks the Secure Debug Authenticator for debug access.

SiP

Silicon Partner.

SPE

Secure Processing Environment.

UDE

Unprivileged Debug Extension.