ETSI TS 103 523-1 V1.1.1 (2020-12)

CYBER;
Middlebox Security Protocol;
Part 1: MSP Framework and Template Requirements

Reference
DTS/CYBER-0027-1

Keywords
cyber security

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - NAF 742 C
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° 7803/88

Contents

Intellectual Property Rights

Essential patents

IPRs essential or potentially essential to normative deliverables may have been declared to ETSI. The information pertaining to these essential IPRs, if any, is publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (https://ipr.etsi.org/).

Pursuant to the ETSI IPR Policy, no investigation, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Trademarks

The present document may include trademarks and/or tradenames which are asserted and/or registered by their owners. ETSI claims no ownership of these except for any which are indicated as being the property of ETSI, and conveys no right to use or reproduce any trademark and/or tradename. Mention of those trademarks in the present document does not constitute an endorsement by ETSI of products, services or organizations associated with those trademarks.

Foreword

This Technical Specification (TS) has been produced by ETSI Technical Committee Cyber Security (CYBER).

The present document is part 1 of a multi-part deliverable covering Middlebox Security Protocols (MSP), defining a generic security blueprint for a family of profiles of MSP, as identified below:

Part 1: "MSP Framework and Template Requirements";

Part 2: "Transport layer MSP, profile for fine grained access control";

Part 3: "Enterprise Transport Security";

Part 5: "Enterprise Network Security".

Later parts of this multi-part deliverable will define new protocols, or provide profiles for the use of existing protocols, that satisfy the security characteristics in this abstract protocol specification.

Modal verbs terminology

In the present document "**shall**", "**shall not**", "**should**", "**should not**", "**may**", "**need not**", "**will**", "**will not**", "**can**" and "**cannot**" are to be interpreted as described in clause 3.2 of the ETSI Drafting Rules (Verbal forms for the expression of provisions).

"**must**" and "**must not**" are **NOT** allowed in ETSI deliverables except when used in direct citation.

Executive summary

Middleboxes are vital in modern networks - from new 5G deployments, with ever-faster networks that need performance management, to resisting new cyber attacks with evolved threat defence that copes with encrypted traffic, to VPN provision. Industry needs middlebox technology to keep pace with these and other evolving and diverse use cases. However, middlebox deployments often raise complex and multi-layered questions around the security, privacy and trust of using middleboxes.

The MSP series [i.22] addresses this gap by specifying a new generation of middlebox protocols that allow middleboxes to perform vital functions securely whilst keeping up with the rapid pace of technical development. The Middlebox Security Protocol (MSP) series contains a set of specifications for protocols that enable secure and functional operation of next generation middleboxes.

Outside of the MSP series, "middlebox" refers to a variety of devices, use cases and functions (see Clause A.7). However, in the MSP series, middlebox refers to devices that perform functions higher up the stack, at the security layer or for security-specific functions (see Clause A.1). Such security use cases are many and varied, described non-exhaustively in Annex A and selectively listed here:

- to provide security services in NFV and SDN environments, as described in Clause A.2;
- system and user security, including cyber defence and protection of user data, as described in Clause A.3;
- operational use cases, as described in Clause A.4, including in Content Delivery Networks [i.16];
- compliance by network operators with obligations and service agreements, and discharge of transparency and audit obligations in regulated industries: see Clause A.5, CIS® Controls [i.3] and ETSI TR 103 305-1 [i.4];
- maintaining enterprise network and data centre visibility [i.1], as described in Clause A.6.

Creating a new generation of protocols brings an opportunity to set a new standard of security and privacy; the present document, as the first part of the MSP series, does this through the MSP Framework. The MSP Framework allows the MSP series to be extensible and allows MSP profiles to be consistently analysed for security, allowing flexibility for individual profiles, whilst maintaining a high security standard across the set of Middlebox Security Protocols.

The present document, Part 1 of the MSP series, defines the security properties of a Middlebox Security Protocol. Each subsequent part of the MSP series defines an MSP profile - which is a protocol that meets the security defined in the MSP Framework and is demonstrated through a Conformance Analysis. This allows multiple MSP profiles with common underlying security properties.

The MSP series [i.22] is driven by four important principles, defined in the present document, that are vital for MSP middlebox deployments to perform their functions securely. These principles underpin the MSP Framework and therefore the MSP series. These are:

1) Data Protection (DP): protecting data from network attackers and malicious actors.
2) Transparency (T): having knowledge of which parties have what access to the data.
3) Access Control (AC): allowing endpoints meaningfully to grant access to parties with this knowledge.
4) Good Citizen (GC): preventing complexity that adds DDoS attack vectors to the network.

Defined in the present document, the MSP Framework defines provisions in the area of each of these principles, called MSP Template Requirements. Using the MSP Framework gives both flexibility and consistency across different MSP profiles to MSP profile developers, MSP profile implementors and MSP specification writers. This methodology permits an array of use cases, as well as thorough security analysis, for the next generation of middlebox protocols: MSP.

1 Scope

The present document is the first part of the Middlebox Security Protocol (MSP) series [i.22]. It is intended to be used by MSP profile developers, MSP profile implementors and MSP specification writers to create MSP profiles and analyse their security. The present document does not specify an MSP profile itself.

The present document defines a security baseline that MSP profiles fulfil to be included in the MSP series. This baseline (defined via the MSP Framework and MSP Template Requirements) facilitates creation of MSP profiles for a wide array of implementations and applications, by simplifying the security analysis required for each profile.

The present document is intended to be a human-readable guide to the security methodology and principles applied to create the MSP Framework and resulting MSP Template Requirements (see Clauses 6.2 to 6.5). Clause 5 and Clause 6, together with profile-specific analysis, form the security analysis for the MSP series. The present document is not and does not attempt to be a security proof; security proofs are only as strong as the assumptions made [i.20] and can lead to a false sense of security [i.25].

The present document describes the motivations behind MSP's creation, how MSP differs to previous middlebox protocols and some of MSP's architectures. The present document introduces the MSP Framework: a common set of security provisions that underpins all MSP specifications (MSP Template Requirements). The present document describes the motivation for having such a framework, the issues addressed by it and the four principles that guided its creation. The remainder of the present document defines the usage and applicability of the MSP Framework to subsequent parts of the MSP series: how the MSP Framework is to be used by MSP profile developers, MSP profile implementors and MSP specification writers.

The present document includes informative annexes to aid readers in its usage. Annex A contains a non-exhaustive list of use cases for MSP. Annex B contains an exemplar MSP Conformance Analysis, performed against the Enterprise Transport Security (ETS) profile, ETSI TS 103 523-3 [i.6]. Annex C describes how a traditional TLS split proxy does not meet the MSP standard set out in the present document.

Comprehensive mitigations for all potential attacks are out of scope. Security requirements of generic and well-known cryptographic algorithms, and assessment of security properties of cryptographic primitives, are out of scope. Attacks that are not attacks on the MSP specification itself, such as implementation vulnerabilities, are out of scope.

2 References

2.1 Normative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

Referenced documents which are not found to be publicly available in the expected location might be found at https://docbox.etsi.org/Reference.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long-term validity.

The following referenced documents are necessary for the application of the present document.

Not applicable.

2.2 Informative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

[i.1] draft-fenter-tls-decryption-00: "Why Enterprises Need Out-of-Band TLS Decryption", S. Fenter, IETF, 2018.

[i.2] ETSI TR 103 421: "CYBER; Network Gateway Cyber Defence".

[i.3] CIS®: "CIS Controls", Version 7.0, 2018.

NOTE: Available at https://www.cisecurity.org/controls/.

[i.4] ETSI TR 103 305-1: "CYBER; Critical Security Controls for Effective Cyber Defence; Part 1: The Critical Security Controls".

[i.5] draft-dolson-transport-middlebox-03: "An Inventory of Transport-centric Functions Provided by Middleboxes", D. Dolson, IETF, 2018.

[i.6] ETSI TS 103 523-3: "CYBER; Middlebox Security Protocol; Part 3: Enterprise Transport Security".

[i.7] "An Untold Story of Middleboxes in Cellular Networks".

NOTE: Available at http://web.eecs.umich.edu/~zmao/Papers/netpiculet.pdf.

[i.8] cloudbric®: "What Are Application Level Attacks?".

NOTE: Available at https://www.cloudbric.com/blog/2015/03/application-level-attacks/.

[i.9] "Network Security: Current Status and Future Directions", Chapter 7: "Intrusion Detection Versus Intrusion Protection", Luis Sousa Cardoso, pp.102.

NOTE: Available at https://books.google.co.uk/books?id=dHys9OXMFMIC&pg=PA102.

[i.10] "Advances in Network Security and Applications: 4th International Conference", David C. Wyld et al. pp. 409, "Difference between IDS and IPS Systems".

NOTE: Available at https://books.google.co.uk/books?id=V7eqCAAAQBAJ&pg=PA499.

[i.11] NETSCOUT®: "Application Layer Attacks".

NOTE: Available at https://www.netscout.com/what-is-ddos/application-layer-attacks.

[i.12] Andreas Müller: "Analysis and Control of Middleboxes in the Internet".

NOTE: Available https://pdfs.semanticscholar.org/d6b2/a37a32af58768e7440d2fc52e48cc483baec.pdf.

[i.13] Ashley Wainwright: "WAN Optimization: What is it and what are the benefits?".

NOTE: Available at https://www.securedgenetworks.com/blog/WAN-Optimization-What-is-it-and-what-are-the-benefits.

[i.14] P. Julian Benadit and F. Sagayaraj Francis: "Improving the Performance of a Proxy Cache Using Very Fast Decision Tree Classifier".

NOTE: Available at https://www.sciencedirect.com/science/article/pii/S187705091500695X.

[i.15] A. Feldmann, et al.: "Performance of Web proxy caching in heterogeneous bandwidth environments", 1999.

NOTE: Available at https://www.researchgate.net/publication/3789953_Performance_of_Web_proxy_caching_in_heterogeneous_bandwidth_environments.

[i.16] CLOUDFLARE®: "The Monsters In The Middleboxes".

NOTE: Available at https://blog.cloudflare.com/monsters-in-the-middleboxes/.

[i.17] Z. Durumeric et al.: "The Security Impact of HTTPS Interception", 2017.

NOTE: Available at https://jhalderm.com/pub/papers/interception-ndss17.pdf.

[i.18] K. Edeline et al.: "mmb: Flexible High-Speed Userspace Middleboxes", 2019.

NOTE: Available at https://arxiv.org/abs/1904.11277.

[i.19] draft-camwinget-tls-ns-impact-00: "Impact of TLS 1.3 to Operational Network Security Practices", N. Camwinget, IETF, 2020.

NOTE: Available at https://datatracker.ietf.org/doc/draft-camwinget-tls-ns-impact/.

[i.20] N. Drucker and S. Gueron: "Selfie: reflections on TLS 1.3 with PSK", Cryptology ePrint Archive, Report 2019/347, 2019.

NOTE: Available at https://eprint.iacr.org/2019/347.

[i.21] K. Bhargavan et al.: "A Formal Treatment of Accountable Proxying over TLS", IEEE 2018 Symposium on Security and Privacy (SP), DOI 10.1109/SP.2018.00021, 2018.

NOTE: Available at https://rud.is/dl/ieee-sp-2018/435301a339.pdf.

[i.22] ETSI TS 103 523 (all parts): "CYBER; Middlebox Security Protocol".

[i.23] "Threat Modeling: Designing for Security", A. Shostack. Chapter 3: STRIDE Methodology, ISBN: 9781118809990.

[i.24] draft-ietf-tls-esni: "TLS Encrypted Client Hello", E. Rescorla et al., IETF, 2020.

NOTE: Available at https://datatracker.ietf.org/doc/draft-ietf-tls-esni/.

[i.25] D. Diemert and T. Jager: "On the Tight Security of TLS 1.3: Theoretically-Sound Cryptographic Parameters for Real-World Deployments", Cryptology ePrint Archive, Report 2020/726, 2020.

NOTE: Available https://eprint.iacr.org/2020/726.pdf.

[i.26] draft-wang-opsec-tls-proxy-bp: "TLS Proxy Best Practice", E. Wang, IETF, 2020.

NOTE: Available https://datatracker.ietf.org/doc/draft-wang-opsec-tls-proxy-bp.

3 Definition of terms, symbols and abbreviations

3.1 Terms

For the purposes of the present document, the following terms apply:

access: legitimate addition into a connection

NOTE: A middlebox can have legitimately-granted access to a connection, but still have insufficient middlebox permissions to modify data. (See also: authorized, permissions).

application: software designed to be used by people, which often has a user interface

application data: information content created or managed by an application, often relating to a user

authorized: having legitimately-granted access

conformance analysis: description of how an MSP profile satisfies its Profile Requirements

context: portion of a datastream or datagram to which middlebox permissions can be granted independently of other datastream or datagram portions

datagram: packet of data

datastream: ordered sequence of datagrams

identity: value or collection of values that identifies and distinguishes an entity sufficiently for the use case

EXAMPLE: A designer of an MSP profile can, depending on use case and threat model, determine which values or collections of values are considered as identities for purposes of the profile.

in-band: occurring in the network traffic stream in real time (at the same protocol layer and connection)

NOTE: This is commonly referred to as "having on-path access", and is the opposite of out-of-band.

initiator: party that begins a connection

Middlebox: physical or virtual device, system or process, in a connection between network endpoints, that requires access to the traffic with specified permissions at the security layer or for security-specific functions

NOTE: Whether in-band or out-of-band, a middlebox is in the connection between network endpoints.

modification: deletion, change, duplication, reordering and/or insertion

NOTE: A party who can delete, change and insert data can also reorder data and duplicate data.

MSP framework: abstract architecture of requirements and process according to which an MSP profile can be defined

MSP profile: single protocol specified in the MSP series [i.22]

MSP Template Requirements: boilerplate provisions, defined in Clause 6 of the present document, from which a subset is chosen for a given MSP profile

out-of-band: occurring separately from the network traffic stream, not necessarily in real time

NOTE: This is commonly referred to as "off-path", and is the opposite of in-band.

permissions: permitted privileges (including any of read, modify, insert and/or delete) that an entity has been granted to use on data carried by an MSP profile

EXAMPLE: A non-exhaustive list of permissions are: none (cannot observe data and cannot edit data); read only (can observe data only); read and write (can observe data and can edit data); read and delete (can observe data and can delete data, but cannot edit).

profile requirements: subset of the MSP Template Requirements that are specific to an MSP profile and threat model

resource attack: attack designed to overwhelm the CPU, memory or network resources of a victim, typically by sending a large number of inauthentic requests

responder: party that replies to the initiator in an MSP connection

selection process: procedure that transforms MSP Template Requirements to Profile Requirements

sensitive data: data requiring heightened confidentiality protection

suitable knowledge: knowledge that is appropriate taking account of use case and threat model of an MSP profile

unauthorized: exercising a greater privilege level than authorized

verify: cryptographically assure the fact of

3.2 Symbols

Void.

3.3 Abbreviations

For the purposes of the present document, the following abbreviations apply:

5G	5th Generation
AC	Access Control (Principle)
ASIC	Application-Specific Integrated Circuit
CDN	Content Delivery Network
CIS®	Center for Internet Security
DDoS	Distributed Denial of Service
DP	Data Protection (Principle)
E	Endpoint (Objective)
ETS	Enterprise Transport Security
GC	Good Citizen (Principle)
HTTP	Hyper Text Transfer Protocol
IDS	Intrusion Detection System
IoT	Internet of Things
IP	Internet Protocol
IPS	Intrusion Prevention System
M	Middlebox (Objective)
MAC	Message Authentication Code
MITM	Man-In-The-Middle
MM	MSP-Mandatory
MSP	Middlebox Security Protocol
NAT	Network Address Translation
NFV	Network Functions Virtualisation
NIS	Network and Information Security
PM	Profile-Mandatory
PNA	Profile-Not-Applicable
PO	Profile-Optional
PR	Profile-Rejected
PSK	Pre-Shared Key
RSA	Rivest-Shamir-Adleman
RTT	Round Trip Time
SDN	Software Defined Network
STRIDE	Spoofing, Tampering, Repudiation, Information disclosure, Denial of service, Elevation of privilege

NOTE: See [i.23].

T	Transparency (Principle)
TCP	Transmission Control Protocol
TLS	Transport Layer Security
TM	Threat Model
UDP	User Datagram Protocol
VNF	Virtualised Network Functions
VPN	Virtual Private Network
WAN	Wide Area Network

4 The Middlebox Security Protocol (MSP) Series

4.1 Introduction

Middleboxes, often referred to as proxies, are vital in modern networks - from new 5G deployments, with ever-faster networks that need performance management, to resisting new cyber attacks with evolved threat defence that copes with encrypted traffic, to VPN provision. Annex A of the present document describes the variety of modern use-cases for middleboxes.

Industry needs middlebox and proxy technology to keep pace with these and other evolving and diverse use cases. However, middlebox deployments often raise complex and multi-layered questions around the security, privacy and trust of using middleboxes. These issues vary: from low-level but crucial questions, such as applying suitable data confidentiality and integrity protection between authorized parties, to high-level issues, such as what knowledge authorized parties have of other parties' identities and functions, which actions parties trust each other to perform independently and which require verification, and the extent to which one or both endpoints can audit the activity of middleboxes. Further, a middlebox deployment could assist other, unrelated network attackers, such as by inadvertently concealing the origin of a Denial of Service attack.

The MSP series [i.22] addresses this gap by specifying a new generation of middlebox protocols that allow middleboxes to perform vital functions [i.19] securely whilst keeping up with the rapid pace of technical development. Further, whilst these common issues can arise in middlebox deployments, mechanisms are being created to solve them and respond as the underlying technologies evolve. These mechanisms depend on use case and threat model, but there has been no industry standard or general framework into which the varied use cases fit, and within which the complex issues can be analysed and techniques measured against. ETSI TR 103 421 [i.2] points out and discusses this deficiency, which led to establishing the Middlebox Security Protocol (MSP) series.

4.2 Introduction to the Middlebox Security Protocol (MSP)

Outside of the MSP series, the term "middlebox" refers to a variety of devices, use cases and functions: see Annex A. However, in the MSP series, a middlebox means a device that perform functions higher up the stack, at the security layer or for security-specific functions. (The term "proxy" is also commonly used outside the MSP series, to cover either this or the wider usage). Such security use cases are many and varied, described non-exhaustively in Annex A, and selectively listed here:

- to provide security services in NFV and SDN environments, as described in Clause A.2;
- system and user security, including cyber defence and protection of user data, as described in Annex A.3;
- operational use cases, as described in Clause A.4, including in Content Delivery Networks [i.16] and [i.19];
- compliance by network operators with obligations and service agreements, and discharge of transparency and audit obligations in regulated industries: see Clause A.5, CIS® Controls [i.3] and ETSI TR 103 305-1 [i.4];
- maintaining enterprise network and data centre visibility [i.1], as described in Clause A.6.

Creating a new generation of protocols brings an opportunity to set a new standard of security and privacy; the present document, as the first part of the MSP series, does this through the MSP Framework (Clause 5).

Any single protocol in the MSP series is an "MSP profile". Each MSP profile is independent of others, but they are all linked by a common baseline established by the MSP Framework defined in the present document (Clause 5.4.2). The MSP Framework allows the MSP series to be extensible and allows MSP profiles to be consistently analysed for security, allowing flexibility for individual profiles, whilst maintaining a high security standard across the set of Middlebox Security Protocols.

MSP is driven by two simple functionality objectives (for endpoints (E) and middleboxes (M) respectively) that are vital for MSP middlebox deployments:

E: Endpoints are able to access middlebox services securely

M: Middleboxes are able to provide services to endpoints securely

Four principles underpin these two functionality objectives to form the foundation of the MSP Framework. Further described in Clause 5.3, these four principles are:

1) Data Protection (DP): protecting data from network attackers and malicious actors.

2) Transparency (T): having knowledge of which parties have what access to the data.

3) Access Control (AC): allowing endpoints meaningfully to grant access to parties with this knowledge.

4) Good Citizen (GC): preventing complexity that adds DDoS attack vectors to the network.

Current middlebox solutions typically fail one or more of these four principles, showcasing the need for MSP. Annex C contains such an analysis for a MITM split proxy and describes its various failures against these four principles.

NOTE: Other efforts are ongoing to document best practice for proxies [i.26].

The MSP series, supported by the present document, allows the creation and analysis of standardized, secure middlebox protocols to support vital use cases and retain important network functionality.

4.3 MSP architecture

4.3.1 Flexibility and extensibility

The range of architectures permitted by MSP is broad and open-ended: the entire MSP series is deliberately not limited to any architect. This is important to span a variety of use cases, to provide for the range of stakeholders and technologies that exist. Clauses 4.3.2 to 4.3.7 describe a non-exhaustive list of architectures that MSP can support, to show the breadth and flexibility of MSP deployment. The architectures in these clauses can also be combined.

An MSP profile specification should describe the architecture or range of architectures for which that MSP profile is designed. As architectural properties can introduce constraints and considerations to the security of the protocol, such architectural properties should be included in any threat modelling assumptions and analysis.

4.3.2 Generic in-band MSP middlebox

The architecture depicted in Figure 4.1 shows a generic in-band middlebox setup, where the MSP middlebox shown is in an arbitrary in-band network position.

NOTE: This is commonly called "on-path".

Many use cases, such as performance functionality described in Clause A.4 and the architectures in Clauses 4.3.4 to 4.3.7, depend on a middlebox being in-band to perform its functions.

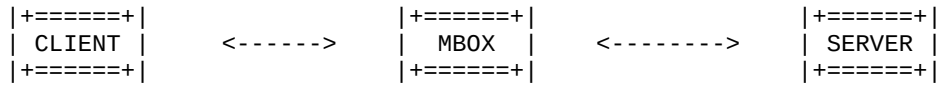

Figure 4.1: In-band MSP architecture

4.3.3 Generic out-of-band MSP middlebox

The architecture depicted in Figure 4.2 shows a generic out-of-band middlebox setup, where the MSP middlebox shown is in an arbitrary out-of-band position.

NOTE: This is commonly called "off-path".

This setup often uses a selective middlebox that forwards (on a separate traffic path) or stores data for further, slower analysis. This is done where action does not need to be taken on live traffic, and can be for the use cases described in Clause A.3 (system security and user security) and Clause A.5 (compliance obligations), analytics, creation of logs, troubleshooting or other purposes.

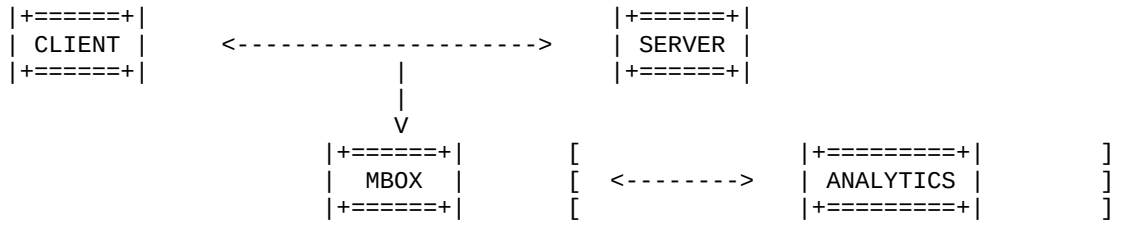

Figure 4.2: Out-of-band MSP architecture

4.3.4 Multiple clients per MSP middlebox

The architecture depicted in Figure 4.3 shows multiple clients connecting to one in-band MSP middlebox. This is often used for VPN provision or anonymization, and in home or enterprise firewalls for defence and filtering purposes.

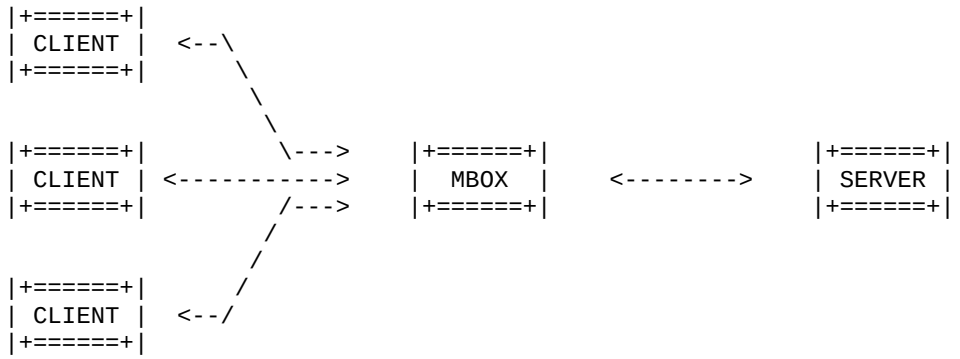

Figure 4.3: MSP architecture with multiple clients per MSP middlebox

4.3.5 Multiple servers per MSP middlebox

The architecture depicted in Figure 4.4 shows multiple servers connecting to one in-band MSP middlebox. This is a common deployment model for CDNs. These middleboxes often go by a different name, such as TLS-terminating "reverse proxies" or "client-facing server" [i.24].

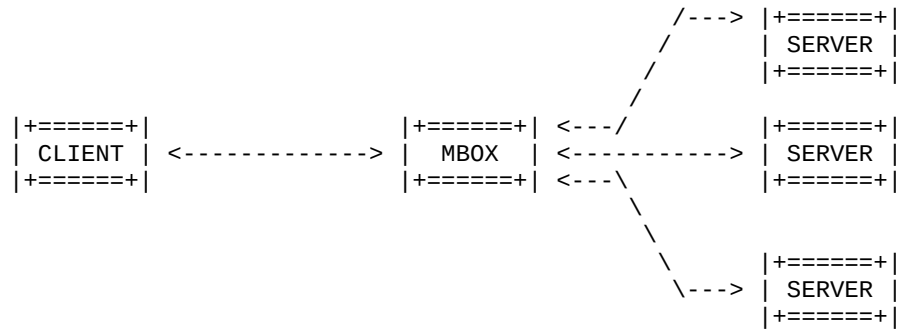

Figure 4.4: MSP architecture with multiple servers per MSP middlebox

4.3.6 Multiple MSP middleboxes per connection

The architecture depicted in Figure 4.5 shows multiple in-band middleboxes in a connection, as well as one example out-of-band middlebox for logging. Such architectures often occur when both the client and server have a middlebox deployment; some middleboxes are within a client's network and some are within the server's.

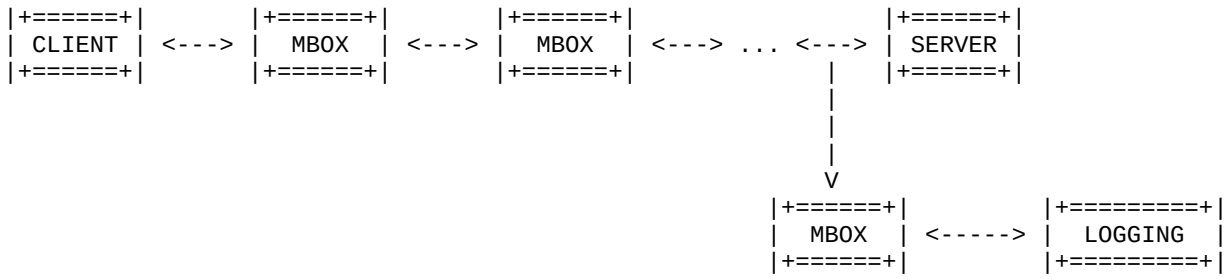

Figure 4.5: MSP architecture with multiple MSP middleboxes

4.3.7 MSP middlebox locations

The architectures depicted in Figure 4.6 and Figure 4.7 describe the location of MSP middleboxes in relation to the endpoints. This not does necessarily refer to geographical location but can hint to ownership of the MSP middlebox, its expected capabilities and capacities, and its common use cases.

In Figure 4.6, the MSP middlebox is closer to the client; this is commonly used for advertisement blockers, for personal data protection and for anonymization purposes.

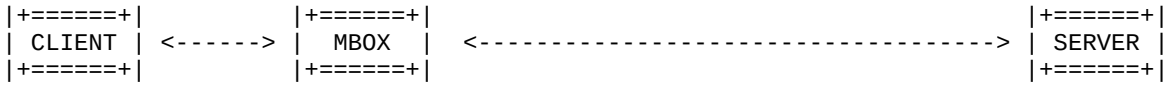

Figure 4.6: MSP middlebox closer to the client

In Figure 4.7, the MSP middlebox is closer to the server; this is commonly the case for data centres, CDNs and other hosting services.

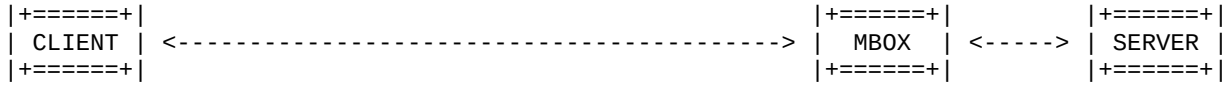

Figure 4.7: MSP middlebox closer to the server

5 MSP Framework

5.1 Introduction

Creating a new generation of middlebox security protocols brings an opportunity to set a new high standard of security and privacy for middlebox deployments. This new high standard is set through the MSP Framework described in Clause 5.4.2. The MSP Framework is formed from four principles, described fully in Clause 5.3. The MSP Framework underpins all MSP specifications, allowing flexibility for individual profiles, whilst maintaining a high security standard across the set of Middlebox Security Protocols.

Clause 5 is intended to be a human-readable guide to the security methodology and principles applied to create the MSP Framework and resulting Template Requirements (see Clauses 6.2 to 6.5). Clause 5 and Clause 6, along with profile-specific analysis, form together the security analysis for MSP. This is not and does not attempt to be a security proof; security proofs are only as strong as the assumptions made [i.20] and can lead to a false sense of security [i.25].

5.2 Motivation

The MSP Framework formed of four principles, which are driven by two top level objectives (for endpoints (E) and middleboxes (M) respectively) that are vital for middlebox deployments:

E: Endpoints are able to access middlebox services securely;

M: Middleboxes are able to provide services to endpoints securely.

These two objectives were driven by real industry use cases, some of which are listed in Annex A, and fulfilling these objectives is commonly agreed to be vital for the management of modern networks. However, analysis showed that perceived gaps in middlebox security centre around four main concerns:

1) data not being protected from network attackers;

2) not fully knowing which parties have what access to the data;

3) the difficulty for endpoints meaningfully to grant access to parties without this knowledge;

4) complexity that adds DDoS attack vectors to the network.

As a result, four principles to address these concerns form the foundation of the MSP Framework, for both middleboxes and endpoints in an MSP connection. Further described in Clause 5.3, these four principles are respectively:

1) Data Protection (DP): protecting data from network attackers and malicious actors;

2) Transparency (T): having knowledge of which parties have what access to the data;

3) Access Control (AC): allowing endpoints meaningfully to grant access to parties with this knowledge;

4) Good Citizen (GC): preventing complexity that adds DDoS attack vectors to the network.

Starting from these four principles, analysis was done to find specific provisions that would cover and meet each principle to a high standard, using the STRIDE methodology [i.23], common for Information Security analysis. These specific provisions are defined in the present document as "Template Requirement". Template Requirements are comprised of:

- Mandatory Template Requirements: requirements that every MSP profile shall meet. These are based on a realistic threat model, with the assumption that network adversaries are always in scope.

- Optional Template Requirements: additional Template Requirements that shall be selected or not selected according to the Selection Process defined in Clause 6.1.3. In the Selection Process an Optional Template Requirement can remain an Optional Template Requirement or become mandatory for a profile, among other possible outcomes. Optional Template Requirements are useful to consider but are not applicable or possible in all use cases. They are based on a threat model where particular adversaries are sometimes in scope, such as malicious endpoints or malicious middleboxes.

The Template Requirements are defined in Clauses 6.2 to 6.5.

NOTE: Traditional MITM split proxies fall short of meeting all the Mandatory Template Requirements derived from the four founding principles of Data Protection, Transparency, Access Control and Good Citizen. Therefore MITM split proxies fall short of the new standard set by the MSP Framework, as described in detail in Annex C.

<table>
<tr><td colspan="4">Template Requirements (Mandatory and Optional)</td></tr>
<tr><td>Data Protection (DP)</td><td>Transparency (T)</td><td>Access Control (AC)</td><td>Good Citizen (GC)</td></tr>
<tr><td colspan="4">M: Middleboxes are able to provide services to endpoints securely</td></tr>
<tr><td colspan="4">E: Endpoints are able to access middlebox services securely</td></tr>
</table>

Figure 5.1: Two key objectives underpin the four principles, which result in Template Requirements

5.3 Principles

5.3.1 Data Protection

The Data Protection principle refers to the well-understood concept of applying communications security to traffic carried by an MSP profile.

Data Protection: Sensitive data is protected by endpoints and middleboxes.

The Data Protection principle relates to two threats in the STRIDE methodology [i.23]: Information Disclosure and Tampering. These are mitigated by the MSP Framework Data Protection principle, which respectively provides both Confidentiality and Integrity.

5.3.2 Transparency

The Transparency principle refers to participants having appropriate information about middleboxes and endpoints in an MSP connection.

Transparency: Endpoints and middleboxes have suitable knowledge of other endpoints and middleboxes.

The Transparency principle relates to one aspect of the STRIDE methodology [i.23]: Spoofing. This is mitigated by the MSP Framework Transparency principle, which provides Authentication.

This principle is not applied uniformly to endpoints and middleboxes; the transparency given to endpoints is mandatory, whereas all transparency given to middleboxes is optional. This is to ensure that middlebox knowledge of parties in the connection will not exceed both endpoints' knowledge of parties in the connection.

5.3.3 Access Control

The Access Control principle refers to participants' ability to opt-out but also to opt-in to an MSP connection, without abuse of the participants.

Access Control: Endpoints are in control of middlebox service provision.

The Access Control principle relates to one aspect of the STRIDE methodology [i.23]: Elevation of Privilege. This is mitigated by the MSP Framework Access Control principle, which provides Authorization.

5.3.4 Good Citizen

The Good Citizen principle is perhaps the least obvious principle; it means an MSP protocol cannot be used to attack other network entities, which typically manifests as traffic amplification in a DDoS attack [i.21]. This principle means MSP does not create significant amplification attacks on processing resources or network resources and that any attempt to perform resource starvation using MSP can be attributed to a source.

Good Citizen: Undetectable resource attacks on any entity are not made easier or more effective.

The Good Citizen principle relates to two aspects of the STRIDE methodology [i.23]: Repudiation and Denial of Service. These are mitigated by the MSP Framework Good Citizen principle, which respectively provides both Non-Repudiation and Availability.

5.4 Structure

5.4.1 Labels and hierarchy

The MSP Framework is hierarchical; each of the four principles is subdivided and forms a 'tree'. The level of detail increases through the 'branches' and 'sub-branches', until reaching detail where it is reasonably easy to judge whether each property is met ('leaves'). This results in the MSP Framework having different levels of depth, as can be seen in Figure 5.2.

Labels are applied to every level of the tree in the MSP Framework. These labels relate to the parent branches from where the requirement is derived, to show the hierarchical nature of the requirement derivation. This can be seen in Figure 5.2.

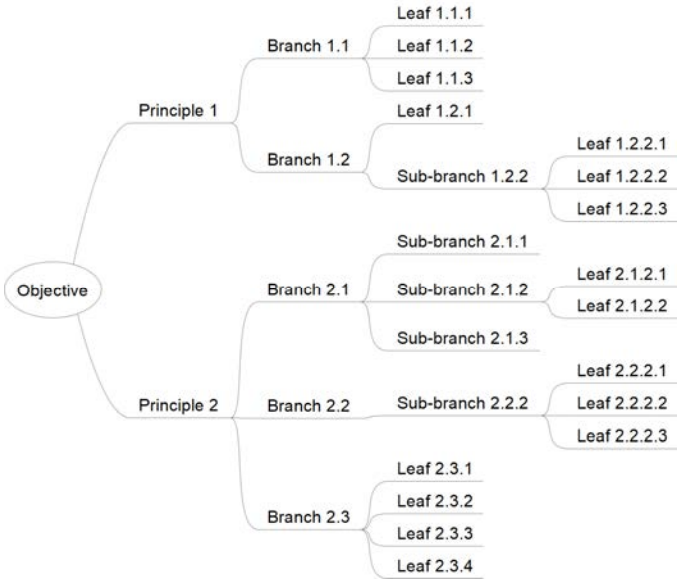

Figure 5.2: Example tree

EXAMPLE 1: The Template Requirement "E.AC.1.2: Middlebox permissions shall be granted by at least one endpoint" is derived from Endpoint Requirements (E), then the Access Control principle (AC) and its parent is E.AC.1 ("Only endpoints shall grant or deny middlebox access and middlebox permissions"). For a particular use case, E.AC.1 can be too generic to be understood without further expansion.

It is possible to satisfy a branch of the tree by meeting all mandatory requirements contained in its sub-branches or leaves.

EXAMPLE 2: The Template Requirement E.AC.1 ("Only endpoints shall grant or deny middlebox access and middlebox permissions") can be satisfied by meeting all of its 'leaves': E.AC.1.1, E.AC.1.2, E.AC.1.3, E.AC.1.4 and E.AC.1.5.

Conversely, if it is possible to satisfy a branch outright, it is not necessary to show compliance with every leaf below that branch.

EXAMPLE 3: The Template Requirement E.AC.1 ("Only endpoints shall grant or deny middlebox access and middlebox permissions") is satisfied by cryptographic mechanisms used. This is enough to meet the Template Requirement; it is not necessary to describe compliance to each of E.AC.1.1, E.AC.1.2, E.AC.1.3, E.AC.1.4 and E.AC.1.5.

The MSP Framework is designed to be flexible. Within the MSP framework, profiles can be defined to address various use cases, whilst still ensuring that each profile satisfies the four important principles defined in Clause 5.3.

5.4.2 MSP Framework contents

The MSP Framework is outlined in the present clause in Figure 5.3. The Template Requirements that fall within it are listed in full in Clauses 6.2, 6.3, 6.4 and 6.5; their usage is described in Clause 6.1.

E: Endpoints are able to access middlebox services securely

- E.DP: Data protection:
  - E.DP.1: Endpoints shall protect confidentiality of sensitive data.
  - E.DP.2: Endpoints may add protection to externally visible characteristics of application data to protect confidentiality of sensitive information about application activity. (This is commonly referred to as Traffic Analysis Protection).
  - E.DP.3: Endpoints shall protect integrity of application data:
    - Sub-nodes E.DP.3.1 - E.DP.3.4: see Clause 6.2.
  - E.DP.4: Endpoints shall protect sensitive information about session state from Unauthorized disclosure, discovery, manipulation and creation.
    - Sub-nodes E.DP.4.1 - E.DP.4.2: see Clause 6.2.
- E.T: Transparency:
  - E.T.1: Endpoints shall receive suitable knowledge of all middlebox identities:
    - Sub-nodes E.T.1.1 - E.T.1.3: see Clause 6.3.
  - E.T.2: Endpoints shall receive knowledge of all middlebox permissions and knowledge of all security mechanisms for data protection.
  - E.T.3: Each endpoint shall be able to verify or otherwise confirm that they have the same knowledge (of middlebox permissions and security mechanisms for data protection) as the other endpoint.
  - E.T.4: Endpoints may receive knowledge of the peer endpoint identity:
    - Sub-nodes E.T.4.1 - E.T.4.2: see Clause 6.3.
  - E.T.5: Endpoints may verifiably audit activity of middleboxes:
    - Sub-nodes E.T.5.1 - E.T.5.2, and sub-sub-nodes: see Clause 6.3.
  - E.T.6: Endpoints may verify or otherwise confirm that middlebox access and middlebox permissions have been granted or denied.
- E.AC: Access Control:
  - E.AC.1: Only endpoints shall grant or deny middlebox access and middlebox permissions:
    - Sub-nodes E.AC.1.1 - E.AC.1.5: see Clause 6.4.
  - E.AC.2: The endpoint(s) that grant(s) access to a middlebox shall authenticate or otherwise confirm its identity before granting access.
  - E.AC.3: At least one endpoint shall choose all security mechanisms for data protection.
  - E.AC.4: Endpoints may grant middlebox access and middlebox permissions only through mutual agreement with the peer endpoint.
  - E.AC.5: Endpoints may authenticate or otherwise verify the identity of all middleboxes whose access is granted by the other endpoint.
- E.GC: Good Citizen:
  - E.GC.1: Resource attacks that use an endpoint action or request shall have some attribution to the attacker:
    - Sub-node E.GC.1.1: see Clause 6.5.

- E.GC.2: An MSP profile shall not provide a significant amplification factor for a resource attack that uses an endpoint action or request:
  - Sub-node E.GC.2.1: see Clause 6.5.

M: Middleboxes are able to provide services to endpoints securely

- M.DP: Data Protection:
  - M.DP.1: Middleboxes shall protect confidentiality of sensitive data that they send.
  - M.DP.2: Middleboxes may add protection to externally visible characteristics of application data to protect confidentiality of sensitive information about application activity. (This is commonly referred to as Traffic Analysis Protection).
  - M.DP.3: Middleboxes shall protect integrity of application data:
    - Sub-nodes M.DP.3.1 - M.DP.3.4: see Clause 6.2.
  - M.DP.4: Middleboxes shall protect sensitive information about session state from Unauthorized disclosure, discovery, manipulation and creation:
    - Sub-nodes M.DP.4.1 - M.DP.4.2: see Clause 6.2.
- M.T: Transparency:
  - M.T.1: Middleboxes may receive knowledge of all middlebox identities:
    - Sub-nodes M.T.1.1 - M.T.1.3: see Clause 6.3.
  - M.T.2: Middleboxes may receive knowledge of all middlebox permissions and knowledge of all security mechanisms for data protection.
  - M.T.3: Middleboxes may be able to verify or otherwise confirm that they have the same knowledge (of middlebox permissions and security mechanisms for data protection) as the other participants.
  - M.T.4: Middleboxes may receive knowledge of either or both endpoint identities:
    - Sub-nodes M.T.4.1 - M.T.4.2: see Clause 6.3.
  - M.T.5: Middleboxes may verifiably audit activity of other middleboxes:
    - Sub-nodes M.T.5.1 - M.T.5.2, and sub-sub-nodes: see Clause 6.3.
- M.AC: Access Control:
  - M.AC.1: Middleboxes shall authenticate or otherwise confirm any participant identity they use for an identity-dependent action. This action is not granting or denying access to an MSP connection, which shall fall within endpoint remit only (E.AC.1). (This stops a middlebox unlocking access to data or services for an identity that has not been checked by the middlebox).
  - M.AC.2: Middleboxes may authenticate or otherwise confirm participant identities:
    - Sub-nodes M.AC.2.1 - M.AC.2.3: see Clause 6.4.
  - M.AC.3: A middlebox may know that its access has been withheld. (Meeting this requirement implies it is not possible to deceive a middlebox into believing it has access).
- M.GC: Good Citizen:
  - M.GC.1: Resource attacks that use a middlebox action or request shall have some attribution to the attacker:
    - Sub-node M.GC.1.1: see Clause 6.5.

- M.GC.2: An MSP profile shall not provide a significant amplification factor for a resource attack that uses a middlebox action or request:
  - Sub-node M.GC.2.1: see Clause 6.5.
- M.GC.3: Middleboxes may be able to drop out of a connection, without breaking or degrading the connection for other participants, to counter an attempted resource attack.

Figure 5.3: MSP Framework

6 MSP Template Requirements

6.1 Usage

6.1.1 Overview

MSP Template Requirements, taken from the MSP Framework defined in Clause 5.4.2, are exhaustively listed in Clauses 6.2, 6.3, 6.4 and 6.5. The MSP Template Requirements should be used by MSP profile developers, MSP profile implementors and MSP specification writers to write an MSP profile specification, analyse the security of an MSP profile, and consider the threat model for an MSP profile.

The Mandatory MSP Template Requirements form a minimum set of requirements that guarantee a security baseline and ensure a large part of the security analysis for an MSP profile is done by satisfying them.

All Mandatory MSP Template Requirements shall be met by an MSP profile. If any Mandatory MSP Template Requirements are not met by a profile, that profile shall not be part of the Middlebox Security Protocol series [i.22].

The Selection Process, defined in Clause 6.1.2, transforms MSP Template Requirements into Profile Requirements that are particular to an MSP use case and threat model. These Profile Requirements shall be replicated in the MSP profile specification, as defined in Clause 6.1.2.

For each profile specification, the Selection Process shall be applied to all MSP Template Requirements (exhaustively listed in Clauses 6.2 to 6.5) to create Profile Requirements. Each Profile Requirement shall be assigned a Profile Requirement type, defined in Table 6.2.

For each profile specification, a Conformance Analysis shall be completed as defined in Clause 6.1.3:

- A Conformance Analysis shall be completed for every MSP Template Requirement to which is assigned any of the following Profile Requirement labels: MSP-Mandatory, Profile-Mandatory, Profile-Optional.
- A Conformance Analysis may be completed for every MSP Template Requirement to which is assigned any of the following Profile Requirement labels: Profile-Not-Applicable, Profile-Rejected.

Annex B of the present document provides an Exemplar Conformance Analysis for ETSI TS 103 523-3 [i.6] that can be used as a guideline when creating a Conformance Analysis for a given MSP profile specification.

The MSP Template Requirements are not, and are not intended to be, exhaustive; collectively, the MSP Template Requirements mitigate the most important threats and provide the functionality of two key objectives, as described in Clause 5.2.

All MSP profile specifications shall contain all the following:

1) A set of Profile Requirements for the specification, derived using the Selection Process defined in Clause 6.1.2, and replicated in the specification as defined in Clause 6.1.2.3.
2) A Conformance Analysis that describes how each MSP-Mandatory, Profile-Mandatory and Profile-Optional Profile Requirement is satisfied, as defined in Clause 6.1.3.

An MSP profile specification may also contain any of the following:

3) A Conformance Analysis that describes why each Profile-Not-Applicable and Profile-Rejected Profile Requirement is labelled as such, as defined in Clause 6.1.3.

4) Further security analysis as needed for the profile, as described in Clause 6.1.4 of the present document.

5) A description, definition or explanation of the following key terms as they relate to the profile: "sensitive data", "suitable knowledge", "middlebox permissions". Definitions may be inherited from the present document or defined in each profile specification.

6) Any dependencies or inherited properties from protocols underpinning the MSP profile that are relevant to the security analysis.

6.1.2 Stage 1: Selection Process

6.1.2.1 Purpose

The MSP Template Requirements provide a template from which MSP profile developers, MSP profile implementors and MSP specification writers can select Profile Requirements that are particular to their use case and threat model.

The procedure to transform MSP Template Requirements into Profile Requirements is the Selection Process, defined in the present clause and shown in Figure 6.1.

Exemplar applications of the Selection Process are described in Clause B.2 and Clause C.2.

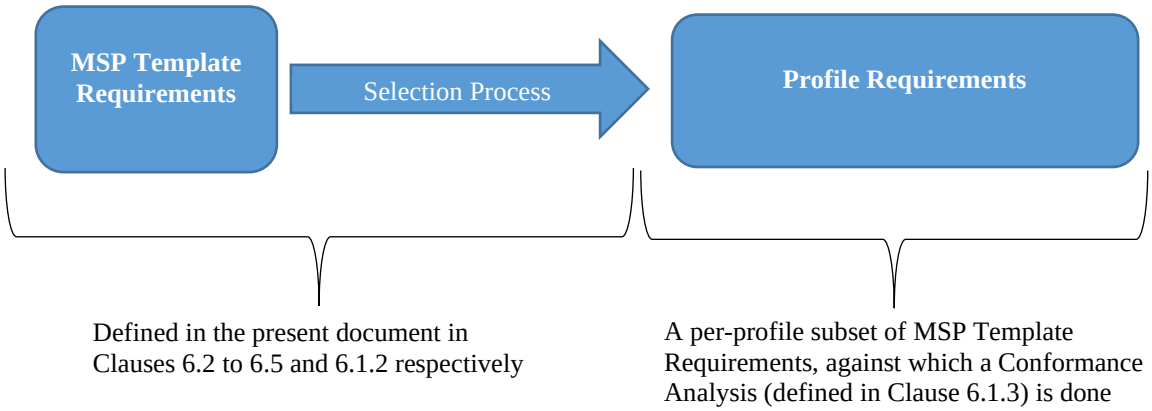

Figure 6.1: The relationship between MSP Template Requirements and Profile Requirements

6.1.2.2 Input: MSP Template Requirements

The MSP Framework, defined in Clause 5.4.2, was designed for flexibility and to respect a wide variety of use cases. This results in a breadth of optional MSP Template Requirements available for MSP profile developers, MSP profile implementors and MSP specification writers to select or not select during the Selection Process.

Table 6.1 describes the differences between Mandatory and Optional MSP Template Requirements.

Table 6.1: Notation for MSP Template Requirement tables

MSP Template Requirement Type	Description of MSP Template Requirement Type
Mandatory	All mandatory MSP Template Requirements shall be selected in every Selection Process. Equivalently, all MSP Template Requirements shall be present in Profile Requirements. If any Mandatory MSP Template Requirements are not met by a profile, that profile shall not be part of the Middlebox Security Protocol series [i.22].
Optional	The MSP Template Requirement may be selected to become part of the Profile Requirements. Satisfying a requirement of this type is specific to the profile and use case, and does not affect a profile's inclusion in the Middlebox Security Protocol series [i.22].

6.1.2.3 Output: Creating Profile Requirements

The Selection Process transforms MSP Template Requirements into Profile Requirements. The Selection Process assigns to each MSP Template Requirement one of five labels:

1) MSP-Mandatory (MM)

2) Profile-Mandatory (PM)

3) Profile-Optional (PO)

4) Profile-Not-Applicable (PNA)

5) Profile-Rejected (PR)

The present clause contains three representations to define how MSP Template Requirements are assigned one of these five labels: Table 6.2, Figure 6.2 and bullet points in the present clause.

Table 6.2: Notation for MSP Profile Requirements after the Selection Process

MSP Template Requirement Type	Selected	Reason for selection/not selection	Profile Requirement Type	Label
Mandatory	YES	Is always selected (see note).	MSP-Mandatory	MM
Optional	YES	Important to the threat model and is always met by the profile.	Profile-Mandatory	PM
	YES	To provide feature(s) important for the use case and is included optionally in the profile.	Profile-Optional	PO
	NO	Selection was precluded by external factors, such as the underlying protocol and its features or the possibility of different permissions within the profile. EXAMPLE 1: A protocol running over UDP does not support the optional functionality of confirming receipt of packets, as this would be an extra feature to incorporate.	Profile-Not-Applicable	PNA
	NO	Not important for the threat model or use case, which can be informed by profile-specific threat analysis. EXAMPLE 2: Audit functionality is not supported as the threat model assumes trustworthy middleboxes.	Profile-Rejected	PR
NOTE: If any Mandatory MSP Template Requirements are not met by a profile, that profile will not be part of the Middlebox Security Protocol series [i.22].				

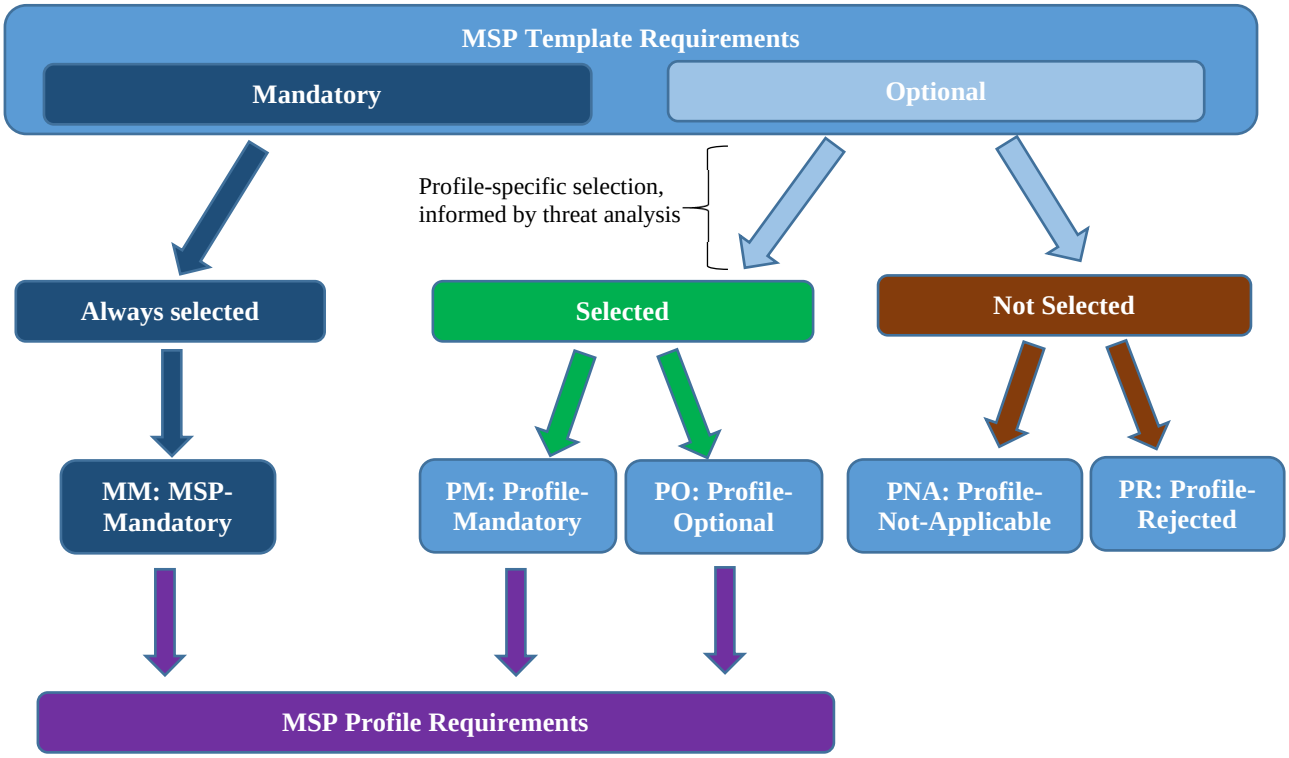

Figure 6.2: The Selection Process

The Selection Process in bullet point form:

- Mandatory MSP Template Requirements (always selected) become MSP-Mandatory (MM) Profile Requirements.
- Optional MSP Template Requirements (can be selected or not selected):
  - If selected, they become one of:
    - Profile-Mandatory (PM) - important to the threat model and is always met by the profile; or
    - Profile-Optional (PO) - to provide feature(s) important for the use case and is included optionally in the profile.
  - If not selected, they become one of:
    - Profile-Not-Applicable (PNA) - selection was precluded by external factors, such as the underlying protocol and its features or the possibility of different permissions within the profile.
    - Profile-Rejected (PR) - not important for the threat model or use case, which can be informed by threat analysis.

EXAMPLE: A Profile-Optional Profile Requirement is included as an extension in an MSP Profile, but only invoked when desired by the profile setup.

The final stage of the Selection Process is based on the Profile Requirement type that has been assigned:

- MSP Template Requirements designated as "**MSP-Mandatory**" shall be copied without changes into the profile specification.
- MSP Template Requirements designated as "**Profile-Mandatory**" shall be copied into the MSP profile specification. The only change performed on Profile-Mandatory Profile Requirements shall be changing the "may" into a "shall".

- MSP Template Requirements designated as "**Profile-Optional**" shall be copied into the profile specification. This shall be done one of two ways:

  1) without changes. In this case, Profile-Optional Profile Requirements contain a "may".

  2) changing only the word "may" to "should". In this case, Profile-Optional Profile Requirements contain a "should".

- MSP Template Requirements designated as "**Profile-Not-Applicable**" may be copied without changes into the profile specification.

- MSP Template Requirements designated as "**Profile-Rejected**" may be copied without changes into the profile specification.

For each MSP profile specification, the Selection Process shall be applied to all MSP Template Requirements (exhaustively listed in Clauses 6.2 to 6.5) to create Profile Requirements as defined in the present clause and a Conformance Analysis shall be completed as defined in Clause 6.1.3.

6.1.3 Stage 2: Conformance Analysis

Conformance Analysis shall verify that MSP-Mandatory and Profile-Mandatory Profile Requirements are always met by any specification-compliant implementation. Conformance Analysis shall verify that Profile-Optional Profile Requirements are provided by the profile to be used when desired.

A Conformance Analysis shall be performed against Profile Requirements, which are created by the Selection Process as defined in Clause 6.1.2, and it shall be included in the MSP profile specification:

- A Conformance Analysis shall be completed for every MSP Template Requirement assigned any of the Profile Requirement type: **MSP-Mandatory**, **Profile-Mandatory**, **Profile-Optional**.

- A Conformance Analysis may be completed for every MSP Template Requirement assigned with Profile Requirement type: **Profile-Not-Applicable**, **Profile-Rejected**.

A Conformance Analysis shall describe how Profile Requirements are satisfied, by describing the mechanism or mechanisms that a profile provides to achieve the functionality defined by the Profile Requirement. Mechanisms can be either technical (including cryptographic or architectural) or procedural (policy-driven or human-driven). For each Profile Requirement, a technical mitigation should be used.

Multiple mechanisms may combine to satisfy one Profile Requirement. Conversely, many Profile Requirements may be satisfied by one mechanism.

The description of how a Profile Requirement is satisfied (or not) should be clear, readable and have enough detail to allow an MSP profile implementor to understand the security model of the MSP profile from the MSP profile specification.

A Conformance Analysis shall describe how each Profile Requirement is satisfied by the MSP profile, according to its type (MM, PM, PO, PNA or PR), as follows:

1) Profile Requirements that are "**MSP-Mandatory**": each profile specification shall describe how each MSP-Mandatory Profile Requirement is satisfied by the profile.

2) Profile Requirements that are "**Profile-Mandatory**": the profile specification shall describe how each Profile-Mandatory Profile Requirement is satisfied by the profile.

3) Profile Requirements that are "**Profile-Optional**": the profile specification shall describe how each Profile-Optional Profile Requirement is optionally satisfied by the profile.

4) Profile Requirements that are "**Profile-Not-Applicable**": if included in the Conformance Analysis, the profile specification shall describe why each Profile-Not-Applicable Profile Requirement is not applicable.

5) Profile Requirements that are "**Profile-Rejected**": if included in the Conformance Analysis, the profile specification shall describe why each Profile-Rejected Profile Requirement is rejected.

NOTE: For some use cases, rejecting Profile Requirements is done to prevent negative effects on security or availability.

When a Profile Requirement is satisfied, a profile specification need not contain a description of how all its leaf Profile Requirements are satisfied.

EXAMPLE 1: A mechanism satisfies E.DP.3 ("Endpoints shall protect integrity of application data") using a bespoke integrity check at the application layer. This data is carried by the MSP profile which operates at a lower layer. The profile specification describes the cryptographic mechanism that satisfies E.DP.3 and does not describe how E.DP.3.1, E.DP.3.2, E.DP.3.3 or E.DP.3.4 are satisfied.

Similarly, when all 'leaf' Profile Requirements are satisfied, a profile specification need not contain a description of how all its 'parent' Profile Requirement is satisfied.

EXAMPLE 2: A mechanism satisfies all of E.DP.3.1, E.DP.3.2, E.DP.3.3 and E.DP.3.4. The profile specification describes the mechanism(s) that satisfy these Profile Requirements and so does not describe how E.DP.3 is satisfied.

Annex B of the present document provides an Exemplar Conformance Analysis for Enterprise Transport Security (ETS) [i.6] that can be used as a guide when creating a Conformance Analysis for an MSP profile specification.

Annex C of the present document provides a Conformance Analysis for a Man-In-The-Middle TLS split proxy, which shows that the split proxy does not meet the requirements to be part of the Middlebox Security Protocol series.

6.1.4 Stage 3: Further analysis (optional)

An individual profile specification may include more security analysis specific to the profile that depends on its features and use case.

The process of creating an MSP profile and writing an MSP profile specification is in practice iterative, but ultimately is presented in the profile specification as linear: analysing threats and the use case, stating threat model assumptions, using these to select MSP Template Requirements, and finally performing a Conformance Analysis for the profile to its Profile Requirements. MSP profile developers and specification writers can decide how to structure the profile security analysis so as to describe and analyse threats that arise, in the context of the threat model and use case, by not meeting specific MSP Template Requirements. This analysis can inform the assignment of Profile Requirement types in the Selection Process (defined in Table 6.2 in Clause 6.1.2.3) as well as the profile's Conformance Analysis.

6.1.5 Attacks external to MSP profiles (optional)

The attacks in the present clause are relate to factors outside of MSP profiles, so are out of scope for the security analysis in the present document; however, these factors and attacks still can be considered by MSP profile developers, MSP profile implementors and MSP specification writers when writing, using or creating MSP profiles.

- Exploitation of software that implements the MSP profile or its cryptographic algorithms; this is a software assurance challenge.
- Sending or receiving malicious data using an MSP profile by an authorized party. This is a cyber defence consideration; however, a cyber defence solution using an MSP-compliant middlebox would provide an extra layer of defence against this threat.
- Insecure random bit generator seeding; sourcing entropy is vital for the security of many protocols.
- Divulging of keys or data outside of the protocol to Unauthorized parties; this is a trust and deployment challenge.
- Bypassing a middlebox using a mechanism outside of the protocol. This is a deployment consideration.
- Side channel and timing attacks against the cryptographic implementation, hardware failures and vulnerabilities, and tampering with devices. This is a hardware-specific challenge.
- Misconfiguration of the system, deliberately or accidentally; this is a deployment challenge.

6.2 MSP Template Requirements - Data Protection

The present clause defines the Data Protection Template Requirements, based on the principle defined in Clause 5.3.2.

Table 6.2: Data Protection Template Requirements

Ref	Data Protection Template Requirement	Mandatory / Optional
E.DP.1	Endpoints shall protect confidentiality of sensitive data that they send.	Mandatory
E.DP.2	Endpoints may add protection to externally visible characteristics of application data to protect confidentiality of sensitive information about application activity. (This is commonly referred to as Traffic Analysis Protection).	Optional
E.DP.3	Endpoints shall protect integrity of application data.	Mandatory
E.DP.3.1	Endpoints shall protect application datagrams from modification in transit between authorized participants.	Mandatory
E.DP.3.2	Endpoints may protect application datagrams from Unauthorized modification by a middlebox.	Optional
E.DP.3.3	Endpoints may protect the datastream from modification in transit between authorized participants.	Optional
E.DP.3.4	Endpoints may protect the datastream from Unauthorized modification by a middlebox.	Optional
E.DP.4	Endpoints shall protect sensitive information about session state from Unauthorized disclosure, discovery, manipulation and creation.	Mandatory
E.DP.4.1	Endpoints shall protect the sensitive cryptographic state from Unauthorized disclosure, discovery, manipulation and creation.	Mandatory
E.DP.4.2	Endpoints may protect the application state from replay and pre-play of data.	Optional
M.DP.1	Middleboxes shall protect confidentiality of sensitive data that they send.	Mandatory
M.DP.2	Middleboxes may add protection to externally visible characteristics of application data to protect confidentiality of sensitive information about application activity. (This is commonly referred to as Traffic Analysis Protection).	Optional
M.DP.3	Middleboxes shall protect integrity of application data.	Mandatory
M.DP.3.1	Middleboxes shall protect application datagrams from modification in transit between authorized participants.	Mandatory
M.DP.3.2	Middleboxes may protect application datagrams from Unauthorized modification by a middlebox.	Optional
M.DP.3.3	Middleboxes may protect the datastream from modification in transit between authorized participants.	Optional
M.DP.3.4	Middleboxes may protect the datastream from Unauthorized modification by a middlebox.	Optional
M.DP.4	Middleboxes shall protect sensitive information about session state from Unauthorized disclosure, discovery, manipulation and creation.	Mandatory
M.DP.4.1	Middleboxes shall protect the sensitive cryptographic state from Unauthorized disclosure, discovery, manipulation and creation.	Mandatory
M.DP.4.2	Middleboxes may protect the application state from replay and pre-play of data.	Optional
M.DP.5	Middleboxes may protect against protocol data fields being used as covert channels by validating the contents or otherwise. (This does not eliminate covert channels from externally visible characteristics such as timings and sizes).	Optional

6.3 MSP Template Requirements - Transparency

The present clause defines the Transparency Template Requirements, based on the principle defined in Clause 5.3.2.

Table 6.3: Transparency Template Requirements

Ref	Transparency Template Requirement	Mandatory / Optional
E.T.1	Endpoints shall receive suitable knowledge of all middlebox identities.	Mandatory
E.T.1.1	Both endpoints shall receive suitable knowledge about the identity of all middleboxes authorized.	Mandatory
E.T.1.2	Endpoints may receive knowledge about the identity of all refused middleboxes.	Optional
E.T.1.3	Endpoints shall be able to verify or otherwise confirm that they have the same knowledge as the peer endpoint of all middleboxes' identities that are authorized.	Mandatory
E.T.2	Endpoints shall receive knowledge of all middlebox permissions and knowledge of all security mechanisms for data protection.	Mandatory
E.T.3	Each endpoint shall be able to verify or otherwise confirm that they have the same knowledge (of middlebox permissions and security mechanisms for data protection) as the other endpoint.	Mandatory
E.T.4	Endpoints may receive knowledge of the peer endpoint identity.	Optional
E.T.4.1	The initiator endpoint may authenticate or otherwise verify the identity of the responder endpoint.	Optional
E.T.4.2	The responder endpoint may authenticate or otherwise verify the identity of the initiator endpoint.	Optional
E.T.5	Endpoints may verifiably audit activity of middleboxes.	Optional
E.T.5.1	The destination endpoint may verifiably audit the activity of middleboxes.	Optional
E.T.5.1.1	The destination endpoint may verify that data has transited and not bypassed each middlebox.	Optional
E.T.5.1.2	The destination endpoint may verify whether a middlebox has modified data.	Optional
E.T.5.1.3	The destination endpoint may verify the full change history of received data.	Optional
E.T.5.2	The sending endpoint may verifiably audit the activity of middleboxes.	Optional
E.T.5.2.1	The sending endpoint may verify that data has transited and not bypassed each middlebox.	Optional
E.T.5.2.2	The sending endpoint may verify whether a middlebox has modified data.	Optional
E.T.5.2.3	The sending endpoint may verify the full change history of received data.	Optional
E.T.6	Endpoints may verify or otherwise confirm that middlebox access and middlebox permissions have been granted or denied.	Optional
M.T.1	Middleboxes may receive knowledge of all middlebox identities.	Optional
M.T.1.1	Middleboxes may receive knowledge about the identity of all middleboxes authorized.	Optional
M.T.1.2	Middleboxes may receive knowledge about the identity of all refused middleboxes.	Optional
M.T.1.3	Middleboxes may be able to verify or otherwise confirm that they have the same knowledge as other participants of all middleboxes' identities that are authorized.	Optional
M.T.2	Middleboxes may receive knowledge of all middlebox permissions and knowledge of all security mechanisms for data protection.	Optional
M.T.3	Middleboxes may be able to verify or otherwise confirm that they have the same knowledge (of middlebox permissions and security mechanisms for data protection) as the other participants.	Optional
M.T.4	Middleboxes may receive knowledge of either or both endpoint identities.	Optional
M.T.4.1	Middleboxes may receive knowledge about the identity of the responder endpoint.	Optional
M.T.4.2	Middleboxes may receive knowledge about the identity of the initiator endpoint.	Optional
M.T.5	Middleboxes may verifiably audit activity of other middleboxes.	Optional
M.T.5.1	Middleboxes may verifiably audit activity of other participants on received data.	Optional
M.T.5.1.1	Middleboxes may verify that received data has transited and not bypassed each middlebox.	Optional
M.T.5.1.2	Middleboxes may verify whether another middlebox has modified received data	Optional
M.T.5.1.3	Middleboxes may verify the full change history of received data.	Optional
M.T.5.2	Middleboxes may verifiably audit activity of other participants on sent data.	Optional
M.T.5.2.1	Middleboxes may verify that sent data has transited and not bypassed each middlebox.	Optional
M.T.5.2.2	Middleboxes may verify whether another middlebox has modified sent data.	Optional
M.T.5.2.3	Middleboxes may verify the full change history of sent data.	Optional

6.4 MSP Template Requirements - Access Control

The present clause defines the Access Control Template Requirements, based on the principle defined in Clause 5.3.3.

Table 6.4: Access Control Template Requirements

Ref	Access Control Template Requirement	Mandatory / Optional
E.AC.1	Only endpoints shall grant or deny middlebox access and middlebox permissions.	Mandatory
E.AC.1.1	Middlebox access shall be granted by at least one endpoint.	Mandatory
E.AC.1.2	Middlebox permissions shall be granted by the same endpoint or endpoints that granted access.	Mandatory
E.AC.1.3	The profile may support multiple levels for middlebox permissions.	Optional
E.AC.1.4	Endpoints may authorize middlebox permissions per context.	Optional
E.AC.1.5	Only endpoints shall deny middlebox access or middlebox permissions. (A middlebox can still block the entire connection between suspected malicious endpoints).	Mandatory
E.AC.2	The endpoint(s) that grant(s) access to a middlebox shall authenticate or otherwise confirm its identity before granting access.	Mandatory
E.AC.3	At least one endpoint shall choose all security mechanisms for data protection.	Mandatory
E.AC.4	Endpoints may grant middlebox access and middlebox permissions only through mutual agreement with the peer endpoint.	Optional
E.AC.5	Endpoints may authenticate or otherwise verify the identity of all middleboxes whose access is granted by the other endpoint.	Optional
M.AC.1	Middleboxes shall authenticate or otherwise confirm any participant identity they use for an identity-dependent action. This action is not granting or denying access to an MSP connection, which shall fall within endpoint remit only (E.AC.1). (This stops a middlebox unlocking access to data or services for an identity that has not been checked by the middlebox).	Mandatory
M.AC.2	Middleboxes may authenticate or otherwise confirm participant identities.	Optional
M.AC.2.1	Middleboxes may authenticate or otherwise confirm the initiator endpoint identity.	Optional
M.AC.2.2	Middleboxes may authenticate or otherwise confirm the responder endpoint identity.	Optional
M.AC.2.3	Middleboxes may authenticate or otherwise confirm all middlebox identities.	Optional
M.AC.3	A middlebox may know that its access has been withheld. (Meeting this requirement implies it is not possible to deceive a middlebox into believing it has access).	Optional

6.5 MSP Template Requirements - Good Citizen

The present clause defines the Good Citizen Template Requirements, based on the principle defined in Clause 5.3.4.

Table 6.5: Good Citizen Template Requirements

Ref	Good Citizen Template Requirement	Mandatory / Optional
E.GC.1	Resource attacks that use an endpoint action or request shall have some attribution to the attacker.	Mandatory
E.GC.1.1	Any party being asked to expend significant resource due to an endpoint request, shall have some attribution of the request to the endpoint.	Mandatory
E.GC.2	An MSP profile shall not provide a significant amplification factor for a resource attack that uses an endpoint action or request.	Mandatory
E.GC.2.1	Where an endpoint sends MSP protocol messages that request a significant amplification factor on resource expenditure, then one of the following two things shall happen: either the recipient is not forced to accept the request or the requesting endpoint expends commensurately amplified resource as a consumer of the result.	Mandatory
M.GC.1	Resource attacks that use a middlebox action or request shall have some attribution to the attacker.	Mandatory
M.GC.1.1	Any party being asked to expend significant resource due to a middlebox request, shall have some attribution of the request to the middlebox.	Mandatory
M.GC.2	An MSP profile shall not provide a significant amplification factor for a resource attack that uses a middlebox action or request.	Mandatory

Ref	Good Citizen Template Requirement	Mandatory / Optional
M.GC.2.1	Where a middlebox sends MSP protocol messages that request a significant amplification factor on resource expenditure, then one of the following two things shall happen: either the recipient is not forced to accept the request or the requesting middlebox expends commensurately amplified resource as a consumer of the result.	Mandatory
M.GC.3	Middleboxes may be able to drop out of a connection, without breaking or degrading the connection for other participants, to counter an attempted resource attack.	Optional

Annex A (informative): Use Cases of MSP

A.1 Introduction

Use cases for middleboxes are broad and varied; the present Annex describes some of the more common deployment scenarios for middleboxes and outlines some MSP-applicable use cases. The use cases in the present Annex describe the functionality and benefits enabled by physical and virtual instantiations of middleboxes in networks and services.

There is some tension regarding the technical and operational need for and role of middleboxes. Some users seek unfettered transmission bandwidth directly between endpoints, free from middleboxes [i.12], Section 2.2. However, operators of networks, service providers, and regulatory communities use middleboxes to meet real-world physical, economic, performance, regulatory, and societal needs.

Over the past several years, this tension has become more obvious. End-to-end encryption impaired the ability of existing middleboxes to perform their essential operational functions [i.18], and also caused issues with deployment as networks became more poorly performing as middleboxes could not properly handle such encrypted traffic [i.12], Section 2.3.3 and [i.17]. If the traffic is encrypted, a middlebox needs to have some manner of controlled awareness/exposure to that traffic to remain functional.

The MSP series is needed because such functionality is not currently provided by the best security methods available (see Annex C). MSP profiles, when implemented, will provision high security across four principles (see Clause 5.3) for users across the use cases described below.

A.2 New infrastructure, services, and innovation

One of the most significant sets of middlebox use cases fundamentally enable new infrastructure, services, and innovations.

Middleboxes provide tailored capabilities within transport paths; as a result, middleboxes are core components of new infrastructure such as Network Functions Virtualisation (NFV) and Software Defined Networks (SDNs) and their implementations as 5G. Middleboxes are part of basic NFV specifications, in the form of VNFs (Virtualised Network Functions).

Innovative middlebox techniques have been applied to satellite, mobile, IoT, industrial control systems, automobile communications, WiFi-clustered installations, and new reductions in power consumption.

A.3 System security and user security

Middleboxes that perform functions to improve cyber defence posture typically fall into one of the following four categories:

- **Network firewalls:** Firewalls are security barriers in networks, frequently implemented with a middlebox in the path between endpoints - often at gateways between networks. Network firewalls use rules to prevent undesired or harmful traffic from reaching an endpoint; therefore, they can be the first line of defence against unwanted and malicious traffic entering the network and targeting users [i.7].

- **Application Firewalls:** Many attacks today are not at the network layer, but at the application layer, such as HTTP flooding attacks and Slowloris attacks [i.8]. Application firewalls are not usually specialized middleboxes but they are often included as part of defence-in-depth designs [i.11].

- **Intrusion Detection Systems (IDS):** These systems continuously monitor real-time network traffic to detect activity indicative of attempted or actual access of network endpoints by Unauthorized persons or computers [i.9].

- **Intrusion Prevention Systems (IPS):** These systems prevent attacks before they have achieved access and done damage. This can be done by traffic inspection to detect new types of attacks, by performing TCP segment reassembly, traffic analysis, application protocol validation, and signature matching [i.10].

Providing information about a provider's or enterprise's security, including where middleboxes are or are not used for defence purposes, helps to inform user choice. If a webserver stores personal information, a user would assume that the server is adequately protected with data loss prevention and intrusion detection mechanisms, in practice likely to be implemented using middleboxes. Currently, there is no standardized way for a user, host server, or another middlebox to determine middlebox cyber defence posture, aiding an informed decision on the entrusting of personal data.

A middlebox can also be used to store or forward traffic and act on it later; for cyber incident response, this allows identification of infected machines for remedial action.

A.4 Operations

Middleboxes for operational purposes can be broadly categorized as follows:

- **Content delivery networks:** Use of reverse proxies to cache static data to improve the speed of content delivery and offer security services such as DDoS mitigation [i.16].

- **Access control:** Access control middleboxes constrain the attachment of a device to a network, establishment of a communication endpoint and place restrictions on what that endpoint can do - spanning the execution of programs to its connectivity. This helps to prevent large-scale network attacks and breaches.

- **Billing and usage monitoring:** Support systems monitor customer use of networks for security, troubleshooting and billing purposes. Monitoring functions employ logging or other forms of audit, frequently implemented using middleboxes.

- **Asset tracking:** Discovery, identification, inventory, and management of physical and software assets at network endpoints is necessary for many purposes, including threat information exchange.

- **Name or tag resolution:** Telecommunication networks require high-speed lookups of addresses or tags, such as caller IDs or domain names in IP networks, which accompany the communications traffic. These lookups are accompanied by security verifications and are cached at different points in network paths and infrastructures - functions frequently performed by middleboxes.

- **Operations control:** Communication networks and services, including those at data centres, are complex. They require an array of management capabilities to monitor and allocate resources, including content delivery capabilities, load balancing, and collation of multiple sources - again, functions frequently performed by middleboxes.

A.5 Compliance obligations

Middleboxes implement requirements for legal and regulatory mechanisms, including Service Level Agreements. These fall into the following categories, some of which overlap with use cases in Clauses A.2 to A.4.

- **Availability/resilience:** Essential networks and services such as financial systems, public utilities and industrial control systems need a high level of availability. A combination of service level agreements, implied warranties and regulatory requirements can stipulate a desired level of availability. Middleboxes can provide functionality to prioritize access for users and to enable specific resilience and survivability design requirements such as for multiple redundant systems, backup, failure isolation capabilities, outage auditing and the elimination of single-point-of-failure components.

- **Emergency and public safety communication:** In emergencies, such as a tsunami warning, it is vital to be able to reach a wide variety of individuals through all available electronic communications (the so-called "authority-to-many" compliance obligation). Middleboxes are used to facilitate compliance capabilities including diverse structured information formats, authentication and delivery methods.

  The inverse of this "authority-to-many" compliance obligation is when an individual needs to reach emergency departments for assistance. During a serious emergency, telecommunication networks and designated public services can be limited to prioritize traffic. Such prioritization requires significant resilience and availability capabilities that are frequently provided using middleboxes.

- **Retained data:** Networks and services create significant amounts of information that is retained in temporary caches, log files, or auditing and accounting systems, often using middleboxes. Data can be retained to meet compliance obligations (especially in financial systems) or to meet contractual requirements, or business auditing requirements.

- **Identity management:** Identity management enables an entity, including a human user, to manifest an identifier at varying levels of trust and uniqueness in the context of the use or operation of a device, network, or service. Identity management capabilities can be implemented using middleboxes.

- **Cyber Security:** Along with Clause A.3, cyber security generally encompasses two sets of capabilities often provisioned by middleboxes - the instantiation of defensive measures and the ability to exchange structured cyber security threat information. Defensive measures include adaptive controls that are designed to be part of every lifecycle phase of a device, network, or service. These controls produce and leverage forensic information, and apply to the maintenance of that information's integrity. Defensive measures depend significantly on timely exchange of structured threat intelligence. Structured threat information exchange encompasses the ability of a device, network or service continuously to acquire, provide, and use current and relevant threat intelligence - all commonly implemented through middlebox platforms.

- **Content control, personal data and privacy:** Individuals and organizations frequently have requirements to protect individuals' personal data; in meeting these, personal privacy defensive middleboxes can play a substantial role. In addition, many companies and organizations have requirements to maintain and protect intellectual property rights. Implementing and maintaining network protections to meet these requirements can require an array of identity management, auditing, and filtering capabilities often provided by middleboxes. Organizations can also use such capabilities to filter network content based on organizational and societal norms; this can be to identify and filter content ranging from harassment, predatory behaviour, speech or depictions that are highly offensive or can cause substantial harm, to improper use of enterprise-provided equipment.

A.6 Enterprise networks and data centres

When enterprise networks encrypt at the transport layer in their internal networks, to protect against insider threat and/or for regulatory compliance reasons, these enterprises had the option to use RSA key exchanges and static RSA private keys for a small, privileged group to decrypt and inspect traffic out-of-band.

Out-of-band decryption provides ubiquitous packet payload visibility inside the enterprise that cannot be replaced by inline/MITM decryption solutions. Today enterprises have extensive packet broker networks to do out-of-band TLS decryption to help intrusion detection devices, fraud detection, malware detection, application performance monitoring tools, customer experience monitoring tools, and many other use cases. Other forms of troubleshooting and monitoring do not functionally replace the capability lost from removing these out-of-band decryption techniques [i.1].

The capability to do out-of-band decryption has been available for approximately twenty years and a large body of supporting tools exists that are dependent on performing efficient, scalable out-of-band decryption. These tools perform mission-critical functions for enterprises, and the loss of out-of-band decryption would create major operational problems for internal enterprise networks that are TLS-encrypted, ultimately preventing migration to later TLS versions. Scalable packet capture and decryption are required for enterprise troubleshooting; without this capability, there would be high severity outages that could not be solved in an acceptable time frame. Without an out-of-band decryption solution, enterprises would be left with the unattractive option of inline/MITM decryption at the data centre edge, or running traffic with legacy protocols or in the clear throughout the data centre if they need packet payload visibility. This would open enterprise networks up to regulatory and insider threat problems.

A.7 Non-MSP use cases

Middleboxes are used in several other contexts where security is not a focus or MSP is likely to be inapplicable.

Middleboxes are used for edge delivery of content, caching, transcoding, compression, forward-deployed insertion of cached content and functions to measure performance [i.5].

The application of performance-enhancing middleboxes in mobile networks, where many users are supported on a shared local radio access network bandwidth, has led to 5G initiatives such as Mobile Edge Computing, implemented through cloud-based middlebox arrays.

Other use cases, ranging from performance to transport reliability, include:

- **Proxy caches:** Proxy caching has been used since the early 1990s to speed up traffic flows significantly and reduce costs [i.14] and [i.15]. Proxy caching middleboxes detect repetitive requests for the same information and store it locally, and then transparently redirect the user to the local store.
- **WAN optimizers:** A WAN optimization device typically analyses network traffic from application clients and the edge gateway of a larger public network to attempt to predict data likely to be requested. Devices at these edge locations pre-fetch and store copies of static content near potential destinations to decrease transit time and latency [i.13].
- **Protocol accelerators:** Protocols used for network communication are highly layered; individual frames or packets and their headers contain considerable "wasted" capacity in structured bits that are not used. Middleboxes optimize the throughput and reduce latencies using these protocol characteristics, using high-performance hardware Application Specific Integrated Circuit (ASIC) chips and segregation of traffic types.
- **Protocol conversion:** Conversions between communication protocols are required for layering or encapsulation; these can be accomplished using specialized middleboxes optimized for the purpose.
- **Network Address and Port Translation:** Communication protocols use different traffic transport addresses and ports, so at network boundaries or gateways, translations can be required. Middleboxes can provide this translation functionality.

Annex B (informative): Exemplar MSP Conformance Analysis

B.1 Introduction

This exemplar is modelled to the MSP profile defined in ETSI TS 103 523-3 [i.6], known as "Enterprise Transport Security" (ETS). This exemplar is not modelled to the ETS variant described in Annex A of of the present document.

The following terms and abbreviations, defined in Clause 6.1.2, are used throughout:

- MSP-Mandatory (MM)
- Profile-Mandatory (PM)
- Profile-Optional (PO)
- Profile-Not-Applicable (PNA)
- Profile-Rejected (PR)

B.2 Exemplar Selection Process

B.2.1 Threat model and assumptions

Many of the assumptions for the ETS threat model, described in the present clause, are derived from its intended deployment: for private enterprise networks and not public networks (the Internet).

Assumptions are listed in the present clause and are labelled "**ETS.TM.X**", where X is a number. These assumptions are referenced throughout Annex B to give context for the MSP Conformance Analysis.

0) **ETS.TM.0:** An ETS middlebox is not malicious and only operates on traffic in a passive manner; specifically, an ETS middlebox reads traffic but does not write or modify traffic.

ETS.TM.0 can be further assured in implementation by the ETS profile deployment architecture.

1) **ETS.TM.1:** Entities with access to the TLS private key are trusted, as with TLS. Explicitly, entities with an ETS private key are trusted.

Some of the Profile Requirements are not guaranteed by technical mechanisms in the profile itself but because entities with access to the ETS private key are trusted. This guarantee holds for the underlying TLS protocol too. This trust in ETS entities - for all entities to share the ETS private key only according to that policy - assures some of the Profile Requirements, as part of the threat model for ETS.

2) **ETS.TM.2:** The ETS server provides accurate Visibility Information in its certificate. If a PSK is used for resumption, the ETS server provided accurate Visibility Information in the original connection.

This implies an assumption that the optional Annex A of ETSI TS 103 523-3 [i.6] is not implemented, and Clause 4 of [i.6] is fully implemented. Analogously to a TLS connection, a client trusts the server and so by extension trusts the server's infrastructure and its policies about handling client traffic. This trust in ETS entities - for the server to provide accurate Visibility Information - assures some of the Profile Requirements, as part of the threat model for ETS.

3) **ETS.TM.3:** In each TLS connection, the client for that connection presents either a certificate (initial connections) or a PSK (resumed connections) for authentication. When a PSK is presented, it is cryptographically tied to the certificate from the previous connection.

4) **ETS.TM.4:** An ETS middlebox can see and verify the handshake for the connection.

This does not violate ETS.TM.0 as a middlebox can still be passive while checking the handshake transcript.

B.2.2 Selection Process outcome

Based on ETS's use case (see Clause A.6 of the present document), these assumptions and threat modelling, the Selection Process assigned the label "Profile-Rejected" to many Optional MSP Template Requirements.

B.3 Conformance Analysis for ETS

B.3.1 Data Protection

B.3.1.1 Overview

Table B.1 contains an overview of the Profile Requirement type assigned to each Data Protection Profile Requirement for ETS, and whether each Profile Requirement is satisfied or not.

Table B.1: Data Protection Profile Requirements satisfied by ETS

Ref	Data Protection Profile Requirement	MM / PM / PO / PNA / PR	Met?
E.DP.1	Endpoints shall protect confidentiality of sensitive data that they send.	MM	Y
E.DP.2	Endpoints may add protection to externally visible characteristics of application data to protect confidentiality of sensitive information about application activity. (This is commonly referred to as Traffic Analysis Protection).	PR	N
E.DP.3	Endpoints shall protect integrity of application data.	MM	Y
E.DP.3.1	Endpoints shall protect application datagrams from modification in transit between authorized participants.	MM	Y
E.DP.3.2	Endpoints may protect application datagrams from Unauthorized modification by a middlebox.	PR	N
E.DP.3.3	Endpoints may protect the datastream from modification in transit between authorized participants.	PO	Y
E.DP.3.4	Endpoints may protect the datastream from Unauthorized modification by a middlebox.	PR	N
E.DP.4	Endpoints shall protect sensitive information about session state from Unauthorized disclosure, discovery, manipulation and creation.	MM	Y
E.DP.4.1	Endpoints shall protect the sensitive cryptographic state from Unauthorized disclosure, discovery, manipulation and creation.	MM	Y
E.DP.4.2	Endpoints may protect the application state from replay and pre-play of data.	PO	Y
M.DP.1	Middleboxes shall protect confidentiality of sensitive data that they send.	MM	Y
M.DP.2	Middleboxes may add protection to externally visible characteristics of application data to protect confidentiality of sensitive information about application activity. (This is commonly referred to as Traffic Analysis Protection).	PR	N
M.DP.3	Middleboxes shall protect integrity of application data.	MM	Y
M.DP.3.1	Middleboxes shall protect application datagrams from modification in transit between authorized participants.	MM	Y
M.DP.3.2	Middleboxes may protect application datagrams from Unauthorized modification by a middlebox.	PR	N
M.DP.3.3	Middleboxes may protect the datastream from modification in transit between authorized participants.	PNA	N
M.DP.3.4	Middleboxes may protect the datastream from Unauthorized modification by a middlebox.	PR	N
M.DP.4	Middleboxes shall protect sensitive information about session state from Unauthorized disclosure, discovery, manipulation and creation.	MM	Y

Ref	Data Protection Profile Requirement	MM / PM / PO / PNA / PR	Met?
M.DP.4.1	Middleboxes shall protect the sensitive cryptographic state from Unauthorized disclosure, discovery, manipulation and creation.	MM	Y
M.DP.4.2	Middleboxes may protect the application state from replay and pre-play of data.	PO	Y
M.DP.5	Middleboxes may protect against protocol data fields being used as covert channels by validating the contents or otherwise. (This does not eliminate covert channels from externally visible characteristics such as timings and sizes).	PR	N

B.3.1.2 MSP-Mandatory Profile Requirements

The MSP-Mandatory Profile Requirements for Data Protection are satisfied as described:

- E.DP.1 is satisfied because application data is encrypted under a session key, which is derived from an ephemeral-static Diffie-Hellman key exchange. This is protected from a replay attack, because the server random value that is used to derive the key will be different.
- E.DP.3 is satisfied because E.DP.3.1 is satisfied:
  - E.DP.3.1 is satisfied because the MAC that is added to application datagrams uses a key derived from the same source as the encryption key; therefore the MAC can only be created by the same participants that encrypted the datagram.
- E.DP.4 is satisfied as random nonces are present to prevent recreation of the session state. A session can be resumed, but this resumption is tied to keys known only to the participants of the session that is being resumed.
- E.DP.4.1 is satisfied as random nonces are present to prevent recreation of the session state. A session can be resumed, but this resumption is tied to keys known only to the participants of the session that is being resumed.
- M.DP.1 is satisfied because middleboxes are only permitted to read in-transit data, and not modify the data or change the encryption that was chosen by the sending endpoint (ETS.TM.0). The datagram is therefore unchanged before and after the middlebox operation.
- M.DP.3 is satisfied because M.DP.3.1 is satisfied:
  - M.DP.3.1 is satisfied because middleboxes will not modify the data or change the integrity protection that was chosen by the sending endpoint. The datagram is therefore unchanged before and after the middlebox operation.
- M.DP.4 is satisfied because M.DP.4.1 is satisfied:
  - The only cryptographic state a middlebox retains for future use is for resumption. M.DP.4.1 is satisfied because this retained state is bound to a key known only to the original participants of the connection.

NOTE: Colluding network adversaries, acting as client and server with captured data, can replay historic exchanges to the middlebox to re-use the same cryptographic state. However, this does not provide an oracle as the middlebox is read-only (ETS.TM.0).

B.3.1.3 Profile-Mandatory Profile Requirements

There are no ETS Profile Requirements of the type "Profile-Mandatory" for Data Protection.

B.3.1.4 Profile-Optional Profile Requirements

The Profile-Optional Profile Requirements for Data Protection are satisfied as described:

- E.DP.3.3 is satisfied optionally, as the TCP layer can correct for reordered, missing or duplicated datagrams, and TLS can be configured to provide cryptographic verification of TCP's reconstruction of the datagram sequence.
- E.DP.4.2 is satisfied optionally because a deployment of the ETS profile can choose to remove support for 0--RTT data, which is vulnerable to replay. The server can reject the use of 0-RTT.
- M.DP.4.2 is satisfied optionally because a deployment of the ETS profile can choose to remove support for 0-RTT data, which is vulnerable to replay. However, an ETS middlebox itself would not be able to reject support of 0-RTT or early data; such protection can be enforced only by the client or server.

B.3.1.5 Profile-Not-Applicable Profile Requirements

The Profile-Not-Applicable Profile Requirements for Data Protection are satisfied as described:

- M.DP.3.3 is not applicable, as an ETS middlebox is read-only and does not modify traffic as per ETS threat model (ETS.TM.0). An in-band ETS middlebox can reorder at the TCP layer to correct for reordered, missing or duplicated datagrams, but does not add protection to traffic.

B.3.1.6 Profile-Rejected Profile Requirements

The Profile-Rejected Profile Requirements for Data Protection are not satisfied and rejected as described:

- E.DP.2 is not satisfied because traffic analysis countermeasures are not provided by ETS. An application that requires traffic analysis protection will need to have its own scheme in the application layer.
- E.DP.3.2 is not satisfied because ETS's threat model assumes that middleboxes are trustworthy (ETS.TM.1) and comply with their permission restrictions (ETS.TM.0). A middlebox rewriting data would not be detectable by an endpoint based only on the features provided by ETS profile. Such detection would need additional mechanisms that are external to the profile.
- E.DP.3.4 is not satisfied because - although middleboxes can correct at the TCP layer for reordered, missing or duplicated datagrams - within the scope and threat model of ETS all middleboxes are assumed to be trustworthy (ETS.TM.1), so there is no requirement for protection against modification by a middlebox.
- M.DP.2 is not satisfied because traffic analysis is not considered an important threat to the ETS use case, so traffic analysis countermeasures are not provided by ETS. An application that requires traffic analysis protection will need to have its own scheme in the application layer.
- M.DP.3.2 is not satisfied as ETS's threat model assumes that middleboxes are trustworthy (ETS.TM.1) and comply with their permission restrictions (ETS.TM.0). A middlebox rewriting data would not be detectable by another middlebox based only on the features provided by the ETS profile. Such detection would need additional mechanisms that are external to the profile.
- M.DP.3.4 is not satisfied because - although middleboxes can correct at the TCP layer for reordered, missing or duplicated datagrams - ETS.TM.0 means there is no requirement for protection against modification by a middlebox.
- M.DP.5 is not satisfied because, in the ETS threat model and use cases, it is not an ETS-defined role of a middlebox to protect against protocol data fields being used as covert channels. This would need additional mechanisms, external to the profile, if it were required in deployments.

B.3.2 Transparency

B.3.2.1 Overview

Table B.2 contains an overview of the Profile Requirement type assigned to each Transparency Profile Requirement for ETS, and whether each Profile Requirement is satisfied or not.

Table B.2: Transparency Profile Requirements met by ETS

Ref	Transparency Profile Requirement	MM / PM / PO / PNA / PR	Met?
E.T.1	Endpoints shall receive suitable knowledge of all middlebox identities.	MM	Y
E.T.1.1	Both endpoints shall receive suitable knowledge about the identity of all middleboxes authorized.	MM	Y
E.T.1.2	Endpoints may receive knowledge about the identity of all refused middleboxes.	PNA	N
E.T.1.3	Endpoints shall be able to verify or otherwise confirm that they have the same knowledge as the peer endpoint of all middleboxes' identities that are authorized.	MM	Y
E.T.2	Endpoints shall receive knowledge of all middlebox permissions and knowledge of all security mechanisms for data protection.	MM	Y
E.T.3	Each endpoint shall be able to verify or otherwise confirm that they have the same knowledge (of middlebox permissions and security mechanisms for data protection) as the other endpoint.	MM	Y
E.T.4	Endpoints may receive knowledge of the peer endpoint identity.	PO	Y
E.T.4.1	The initiator endpoint shall authenticate or otherwise verify the identity of the responder endpoint.	PM	Y
E.T.4.2	The responder endpoint may authenticate or otherwise verify the identity of the initiator endpoint.	PO	Y
E.T.5	Endpoints may verifiably audit activity of middleboxes.	PR	N
E.T.5.1	The destination endpoint may verifiably audit the activity of middleboxes.	PR	N
E.T.5.1.1	The destination endpoint may verify that data has transited and not bypassed each middlebox.	PR	N
E.T.5.1.2	The destination endpoint may verify whether a middlebox has modified data.	PR	N
E.T.5.1.3	The destination endpoint may verify the full change history of received data.	PR	N
E.T.5.2	The sending endpoint may verifiably audit the activity of middleboxes.	PR	N
E.T.5.2.1	The sending endpoint may verify that data has transited and not bypassed each middlebox.	PR	N
E.T.5.2.2	The sending endpoint may verify whether a middlebox has modified data.	PR	N
E.T.5.2.3	The sending endpoint may verify the full change history of received data.	PR	N
E.T.6	Endpoints may verify or otherwise confirm that middlebox access and middlebox permissions have been granted or denied.	PR	N
M.T.1	Middleboxes shall receive knowledge of all middlebox identities.	PM	Y
M.T.1.1	Middleboxes shall receive knowledge about the identity of all middleboxes authorized.	PM	Y
M.T.1.2	Middleboxes may receive knowledge about the identity of all refused middleboxes.	PNA	N
M.T.1.3	Middleboxes may be able to verify or otherwise confirm that they have the same knowledge as other participants of all middleboxes' identities that are authorized.	PO	Y
M.T.2	Middleboxes shall receive knowledge of all middlebox permissions and knowledge of all security mechanisms for data protection.	PM	Y
M.T.3	Middleboxes may be able to verify or otherwise confirm that they have the same knowledge (of middlebox permissions and security mechanisms for data protection) as the other participants.	PR	N
M.T.4	Middleboxes may receive knowledge of either or both endpoint identities.	PO	Y
M.T.4.1	Middleboxes may receive knowledge about the identity of the responder endpoint.	PO	Y
M.T.4.2	Middleboxes may receive knowledge about the identity of the initiator endpoint.	PO	Y
M.T.5	Middleboxes may verifiably audit activity of other middleboxes.	PR	N
M.T.5.1	Middleboxes may verifiably audit activity of other participants on received data.	PR	N
M.T.5.1.1	Middleboxes may verify that received data has transited and not bypassed each middlebox.	PR	N

Ref	Transparency Profile Requirement	MM / PM / PO / PNA / PR	Met?
M.T.5.1.2	Middleboxes may verify whether another middlebox has modified received data	PR	N
M.T.5.1.3	Middleboxes may verify the full change history of received data.	PR	N
M.T.5.2	Middleboxes may verifiably audit activity of other participants on sent data.	PR	N
M.T.5.2.1	Middleboxes may verify that sent data has transited and not bypassed each middlebox.	PR	N
M.T.5.2.2	Middleboxes may verify whether another middlebox has modified sent data.	PR	N
M.T.5.2.3	Middleboxes may verify the full change history of sent data.	PR	N

B.3.2.2 MSP-Mandatory Profile Requirements

The MSP-Mandatory Profile Requirements for Transparency are satisfied as described:

- E.T.1 is satisfied because E.T.1.1 and E.T.1.3 are satisfied, as described below.
- E.T.1.1 is satisfied because middleboxes are authorized by the server endpoint providing the private key via an out-of-band mechanism (ETS.TM.1). The server endpoint includes Visibility Information in its server certificate, which describes this key provision (ETS.TM.2). The server certificate is sent to the client. If the certificate was received correctly, the handshake transcript hash matches and only then does connection complete.
- E.T.1.3 is satisfied because knowledge of authorized middleboxes is contained in the Visibility Information in the server's certificate (ETS.TM.2). The transcript hash and certificate signature confirm that both endpoints have seen the same certificate.
- E.T.2 and E.T.3 is satisfied because the client and server negotiate the security mechanisms used for data protection, and this negotiation is integrity-protected by the ETS handshake transcript hash. Additionally, no choices are presented for middlebox permissions (ETS.TM.0) - so this satisfied by default.

B.3.2.3 Profile-Mandatory Profile Requirements

The Profile-Mandatory Profile Requirements for Transparency are satisfied as described:

- E.T.4.1 is satisfied because the server (responder endpoint) provides its certificate or a PSK as its identity (ETS.TM.3), and the client (initiator endpoint) is assured of its authenticity by a CertificateVerify message. The transcript hash and certificate signature are verified by the client.
- M.T.1 and M.T.1.1. are satisfied because the middlebox receives the server certificate, which contains Visibility Information that provides this information (ETS.TM.2).
- M.T.2 is satisfied because middleboxes are either excluded from the connection or given one level of permission (read). Middleboxes receive knowledge of the negotiated security mechanism for the connection from the ServerHello, which is integrity-protected by the ETS handshake transcript hash.

B.3.2.4 Profile-Optional Profile Requirements

The Profile-Optional Profile Requirements for Transparency are satisfied as described:

- E.T.4 is satisfied because the endpoint can verify the ETS handshake transcript hash, which integrity-protects client and server's key exchange negotiation.
- E.T.4.2 is satisfied optionally, as the client (initiator endpoint) does not have to send a certificate or PSK in the ETS handshake. If it sends a certificate, CertificateVerify messages provide the server (responder endpoint) with client authentication. If it sends a PSK (ETS.TM.3), this provides identity information that satisfies E.T.4.2.

- M.T.1.3 is satisfied optionally because knowledge of authorized middleboxes is contained in the Visibility Information in the server's certificate (ETS.TM.2). The transcript hash and certificate signature confirm that both endpoints have seen the same certificate. The middlebox can verify the hash and verify the certificate (ETS.TM.4).

- M.T.4 is satisfied because the server (responder) provides its certificate or a PSK as its identity (ETS.TM.3). The client (initiator) is assured of the certificate's authenticity by a CertificateVerify message or of the PSK's authenticity through ownership of that PSK. The transcript hash and certificate signature can be optionally verified by the middlebox (ETS.TM.4).

- M.T.4.1 is satisfied because the server (responder) provides its certificate or a PSK as its identity (ETS.TM.3). The client (initiator) is assured of the certificate's authenticity by a CertificateVerify message or of the PSK's authenticity through ownership of that PSK. The transcript hash and certificate signature can be optionally verified by the middlebox (ETS.TM.4).

- M.T.4.2 is satisfied if the client (initiator) provides a certificate or a PSK (ETS.TM.3), and the server (responder) is assured of the certificate's authenticity by a CertificateVerify message or of the PSK's authenticity through ownership of that PSK. The transcript hash and certificate signature can be optionally verified by the middlebox (ETS.TM.4).

B.3.2.5 Profile-Not-Applicable Profile Requirements

The Profile-Not-Applicable Profile Requirements for Transparency are not applicable as described:

- E.T.1.2 is not applicable the server can refuse middlebox access by using a different key, but there is no defined mechanism in the ETS profile to inform a client of the middleboxes that are no longer party to the connection.

- M.T.1.2 is not satisfied because the server can choose a private key that is known to some middleboxes but not known to all. This selection of middleboxes is not shared with other middleboxes. An ETS middlebox could compare Visibility Information in different certificates used (ETS.TM.2) for different connections (ETS.TM.4) to discern now-refused middleboxes, but there is no defined mechanism in the ETS profile to determine middleboxes that are no longer party to the connection.

B.3.2.6 Profile-Rejected Profile Requirements

The Profile-Rejected Profile Requirements for Transparency are not satisfied and rejected as described:

- E.T.5 is not satisfied because E.T.5.1 and E.T.5.2 are not satisfied:
  - E.T.5.1 is not satisfied because E.T.5.1.1, E.T.5.1.2 and E.T.5.1.3 are not satisfied:
    - E.T.5.1.1, E.T.5.1.2 and E.T.5.1.3 are not satisfied; middlebox auditing is not supported by the ETS profile, due to ETS.TM.0 and ETS.TM.1.
  - E.T.5.2 is not satisfied because E.T.5.2.1, E.T.5.2.2 and E.T.5.2.3 are not satisfied:
    - E.T.5.2.1, E.T.5.2.2 and E.T.5.2.3 are not satisfied; middlebox auditing is not supported by the ETS profile, due to ETS.TM.0 and ETS.TM.1.

- E.T.6 is not satisfied; due to ETS.TM.1, endpoints do not need assurance that middleboxes are present and behaving as expected.

- M.T.3 is rejected because there can be multiple middleboxes in a deployment (ETS.TM.5), and there is no mechanism for a middlebox to verify the knowledge of other middleboxes within ETS. Though a middlebox can receive and can verify the ETS handshake transcript hash (ETS.TM.4), which integrity-protects client and server's key exchange negotiation, this provides no knowledge about another middlebox's accesses.

- M.T.5 is not satisfied because M.T.5.1 and M.T.5.2 are not satisfied:
  - M.T.5.1 is not satisfied because M.T.5.1.1, M.T.5.1.2 and M.T.5.1.3 are not satisfied:
    - M.T.5.1.1, M.T.5.1.2 and M.T.5.1.3 are not satisfied; middlebox auditing is not supported by the ETS profile, due to ETS.TM.0 and ETS.TM.1.
  - M.T.5.2 is not satisfied because M.T.5.2.1, M.T.5.2.2 and M.T.5.2.3 are not satisfied:
    - M.T.5.2.1, M.T.5.2.2 and M.T.5.2.3 are not satisfied; middlebox auditing is not supported by the ETS profile, due to ETS.TM.0 and ETS.TM.1.

B.3.3 Access Control

B.3.3.1 Overview

Table B.3 contains an overview of the Profile Requirement type assigned to each Access Control Profile Requirement for ETS, and whether each Profile Requirement is satisfied or not.

Some of the Profile Requirements are not guaranteed by the profile itself (technical mechanisms) but are guaranteed to the extent that entities with access to the ETS private key are trusted. This guarantee holds for the underlying TLS protocol too. This trust in ETS entities - for the server to provide accurate Visibility Information, and for all entities to share the ETS private key only according to that policy - assures some of the Profile Requirements and is described in Clauses B.3.3.2 to B.3.3.5.

Table B.3: Access Control Profile Requirements met by ETS

Ref	Access Control Profile Requirement	MM / PM / PO / PNA / PR	Met?
E.AC.1	Only endpoints shall grant or deny middlebox access and middlebox permissions.	MM	Y
E.AC.1.1	Middlebox access shall be granted by at least one endpoint.	MM	Y
E.AC.1.2	Middlebox permissions shall be granted by the same endpoint or endpoints that granted access.	MM	Y
E.AC.1.3	The profile may support multiple levels for middlebox permissions.	PNA	N
E.AC.1.4	Endpoints may authorize middlebox permissions per context.	PNA	N
E.AC.1.5	Only endpoints shall deny middlebox access or middlebox permissions. (A middlebox can still block the entire connection between suspected malicious endpoints).	MM	Y
E.AC.2	The endpoint(s) that grant(s) access to a middlebox shall authenticate or otherwise confirm its identity before granting access.	MM	Y
E.AC.3	At least one endpoint shall choose all security mechanisms for data protection.	MM	Y
E.AC.4	Endpoints may grant middlebox access and middlebox permissions only through mutual agreement with the peer endpoint.	PR	N
E.AC.5	Endpoints may authenticate or otherwise verify the identity of all middleboxes whose access is granted by the other endpoint.	PR	N
M.AC.1	Middleboxes shall authenticate or otherwise confirm any participant identity they use for an identity-dependent action. This action is not granting or denying access to an MSP connection, which shall fall within endpoint remit only (E.AC.1). (This stops a middlebox unlocking access to data or services for an identity that has not been checked by the middlebox).	MM	N
M.AC.2	Middleboxes may authenticate or otherwise confirm participant identities.	PO	Y
M.AC.2.1	Middleboxes may authenticate or otherwise confirm the initiator endpoint identity.	PO	Y
M.AC.2.2	Middleboxes **shall** authenticate or otherwise confirm the responder endpoint identity.	PM	Y
M.AC.2.3	Middleboxes may authenticate or otherwise confirm all middlebox identities.	PR	N
M.AC.3	A middlebox may know that its access has been withheld. (Meeting this requirement implies it is not possible to deceive a middlebox into believing it has access).	PNA	Y

B.3.3.2 MSP-Mandatory Profile Requirements

The MSP-Mandatory Profile Requirements for Access Control are satisfied as described:

- E.AC.1 is satisfied as E.AC.1.1, E.AC.1.2 and E.AC.1.5 are all satisfied:
  - E.AC.1.1 is satisfied because middlebox access is granted by the server endpoint, as it provisions the Diffie-Hellman private key to the middleboxes.
  - E.AC.1.2 is satisfied by the server endpoint provisioning the Diffie-Hellman private key to the middleboxes, granting all middleboxes the same level of access.
  - E.AC.1.5 is satisfied because the mechanism to deny middlebox access is for the server endpoint to use a private Diffie-Hellman key that the middlebox does not have.
    Middleboxes cannot deny middlebox permissions, neither can the client endpoint.
- E.AC.2 is satisfied because the server confirms the middlebox identity out-of-band when provisioning the private key. Possession of this key is the only way for a middlebox to access data.
- E.AC.3 is satisfied because the client and server alone negotiate the cryptographic mechanisms used. This negotiation is integrity-protected by the ETS handshake transcript hash.
- M.AC.1 is satisfied as no participant identities are used for access control to data or services.

B.3.3.3 Profile-Mandatory Profile Requirements

The Profile-Mandatory Profile Requirements for Access Control are satisfied as described:

- M.AC.2.2 is satisfied because the server certificate is always provided. The transcript hash and certificate signature confirm that both endpoints have seen the same certificate. The middlebox can verify the hash and verify the certificate (ETS.TM.4).

B.3.3.4 Profile-Optional Profile Requirements

The Profile-Optional Profile Requirements for Access Control are satisfied as described:

- M.AC.2 is optionally satisfied as M.AC.2.1 is optionally satisfied and M.AC.2.2 is always satisfied. This means that the responder identity is always authenticated or otherwise confirmed, and the initiator identity is sometimes confirmed:
  - M.AC.2.1 is satisfied optionally because a client certificate can be provided optionally. When it is, the transcript hash and certificate signature confirm that both endpoints have seen the same certificate. The middlebox can verify the hash and verify the certificate (ETS.TM.4).

B.3.3.5 Profile-Not-Applicable Profile Requirements

The Profile-Not-Applicable Profile Requirements for Access Control are not applicable as described:

- E.AC.1.3 and E.AC.1.4 are not applicable because different permissions are not offered in ETS; when the server endpoint provisions the Diffie-Hellman private key to the middleboxes, all middleboxes are granted the same level of permission for a connection that uses that key.
- M.AC.3 is not applicable because, though a middlebox could check the key_share in the ServerHello message against the public counterpart to the static private key(s) it holds (and through this, know if its access has been withheld), this is not a mechanism defined in ETS.

B.3.3.6 Profile-Rejected Profile Requirements

The Profile-Rejected Profile Requirements for Access Control are not satisfied and rejected as described:

- M.AC.2.3 is not met because middleboxes have no means of authenticating other middleboxes or otherwise confirming the identities of all middleboxes. Though middleboxes are identified in the server certificate using Visibility Information (ETS.TM.2), ETS does not provide a means for middleboxes to authenticate each other. The middleboxes identified in the server certificate can be offline or not present to be authenticated and answer a challenge.

- E.AC.4 is not satisfied because middlebox access is granted by the server endpoint alone - when it provisions the Diffie-Hellman private key to the middleboxes. It also grants all middleboxes the same level of access. This accounted for through ETS.TM.1.

- E.AC.5 is not satisfied because only the server can have assurance of middlebox identities, through the fact the middlebox possesses the private key counterpart to the static Diffie-Hellman. This accounted for through ETS.TM.1. The client gains some knowledge from the Visibility Information in the certificate (ETS.TM.2), but it does not gain this same assurance level from the Visibility Information in the certificate.

B.3.4 Good Citizen

Table B.4 contains an overview of the Profile Requirement type assigned to each Good Citizen Profile Requirement for ETS, and whether each Profile Requirement is satisfied or not. All Good Citizen Profile Requirements are of the type "MSP-Mandatory".

Table B.4: Good Citizen Profile Requirements met by ETS

Ref	Good Citizen Profile Requirement	MM / PM / PO / PNA / PR	Met?
E.GC.1	Resource attacks that use an endpoint action or request shall have some attribution to the attacker.	MM	Y
E.GC.1.1	Any party being asked to expend significant resource due to an endpoint request, shall have some attribution of the request to the endpoint.	MM	Y
E.GC.2	An MSP profile shall not provide a significant amplification factor for a resource attack that uses an endpoint action or request.	MM	Y
E.GC.2.1	Where an endpoint sends MSP protocol messages that request a significant amplification factor on resource expenditure, then one of the following two things shall happen: either the recipient is not forced to accept the request or the requesting endpoint expends commensurately amplified resource as a consumer of the result.	MM	Y
M.GC.1	Resource attacks that use a middlebox action or request shall have some attribution to the attacker.	MM	Y
M.GC.1.1	Any party being asked to expend significant resource due to a middlebox request, shall have some attribution of the request to the middlebox.	MM	Y
M.GC.2	An MSP profile shall not provide a significant amplification factor for a resource attack that uses a middlebox action or request.	MM	Y
M.GC.2.1	Where a middlebox sends MSP protocol messages that request a significant amplification factor on resource expenditure, then one of the following two things shall happen: either the recipient is not forced to accept the request or the requesting middlebox expends commensurately amplified resource as a consumer of the result.	MM	Y
M.GC.3	Middleboxes may be able to drop out of a connection, without breaking or degrading the connection for other participants, to counter an attempted resource attack.	PO	Y

The Conformance Analysis in the present clause describes how MSP-Mandatory Good Citizen Profile Requirements are met by the ETS profile.

- E.GC.1 is met by satisfying E.GC.1.1:

  - E.GC.1.1 is met as the endpoint establishing connection will provide their IP address and this will be correct, otherwise the TCP layer will not be established. This IP address provides some attribution.

- E.GC.2 is met as ETS has no messages that change the behaviour of a participant:
  - E.GC.2.1 is met by satisfying E.GC.2.
- M.GC.1 is met by satisfying M.GC.1.1:
  - M.GC.1.1 is met as ETS middleboxes are read-only (ETS.TM.0) and therefore there are no middlebox requests.
- M.GC.2 is met as ETS has no messages that change the behaviour of a participant:
  - M.GC.2.1 is met by satisfying M.GC.2.

The Conformance Analysis in the present clause describes in addition how the Profile-Optional Profile Requirement M.GC.3 is satisfied by the ETS profile:

- M.GC.3 is optionally satisfied by ETS because a conformant implementation can adhere to an architecture that ensures that other ETS participants continue to participate in a connection unaffected by a read-only middlebox (ETS.TM.0) dropping out at the ETS layer.

Annex C (informative): Insufficient MSP Conformance Analysis

C.1 TLS Man-In-The-Middle split proxy

A traditional TLS Man-In-The-Middle (MITM) split proxy is a middlebox with full read and write access to a TLS connection between a client and server (illustrated in Figure C.1). A TLS MITM split proxy does not meet MSP Template Requirements and so cannot and will not be included in the MSP series (see Clause C.3 for details).

The typical method of establishing the split proxy for a TLS connection is:

1) The client attempts to establish a connection to the server.
2) The middlebox interposes itself and becomes the terminating end of the client's attempted TLS connection.
3) The middlebox establishes a separate TLS connection to the server, acting as a client.
4) The middlebox decrypts data that it receives from one TLS connection, re-encrypts this data and sends it in the other TLS connection.

The data is decrypted ("broken out") at the endpoints and at the middlebox only. The middlebox can monitor and edit the data; neither client or server will be aware if any edits are made.

The client trusts the middlebox; this trust enables the middlebox to establish a TLS connection to the client (TLS #1 Connection in Figure C.1). Often, this is often done by configuring the client with a pre-installed public key, which allows the client to accept the middlebox as a certificate authority. The middlebox signs a certificate for itself with the server name that the client is expecting and uses this to authenticate the connection.

The client can also trust the middlebox because it believes the middlebox is the server. In this case, the middlebox the server has shared its private key and certificate with the middlebox. In this case, the server also clearly trusts the middlebox. This type of middlebox can be called a "reverse proxy".

For the middlebox TLS connection to the server (TLS #2 Connection in Figure C.1), the middlebox acts as the client. It ensures that the server certificate is valid and can perform other checks that the security properties of the connection are appropriate for the client. Mutual authentication between the middlebox and server is not guaranteed, as the middlebox does not always have a client certificate.

Multiple TLS MITM proxies can exist on the communications path from client to server (illustrated in Figure C.2).

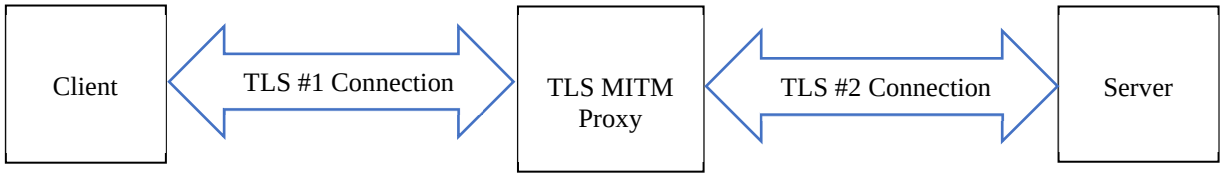

Figure C.1: Split proxy architecture

C.2 Template Requirement selection

C.2.1 Threat model and assumptions

The threat model for the TLS MITM split proxy described in the present clause considers only the mildest threats and the strongest assumptions of trust. This is to illustrate how strong the requirements in the MSP framework are: they cannot be met by a TLS MITM split proxy, even in this generous threat model. This guides the Selection Process of MSP Template Requirements, allowing selection and analysis for the use case based on the threat model.

Assumptions are listed in the present clause and are labelled "**Annex-C.TM.X**", where X is a number. Each one represents an assumption made in the threat model; these can be referenced throughout the rest of Annex C to give context for the MSP Conformance Analysis:

0) **Annex-C.TM.0**: Endpoints are not malicious and will only operate on traffic in an expected way.

1) **Annex-C.TM.1**: One or two TLS MITM split proxies can be present in a connection, so each split proxy has at least one adjacent endpoint.

2) **Annex-C.TM.2**: Each TLS MITM split proxy is strongly trusted by at least one of the endpoints.

NOTE 1: In Figure C.2, Annex-C.TM.2 means TLS MITM Proxy A is trusted by the Client, as the Client is configured to trust the certificate authority that signs the certificate used by the split proxy to set up the TLS connection. The Server trusts TLS MITM Proxy B; this is implied by the fact it sets up a connection. However, the Server cannot know that TLS MITM Proxy B is a proxy acting on behalf of the Client.

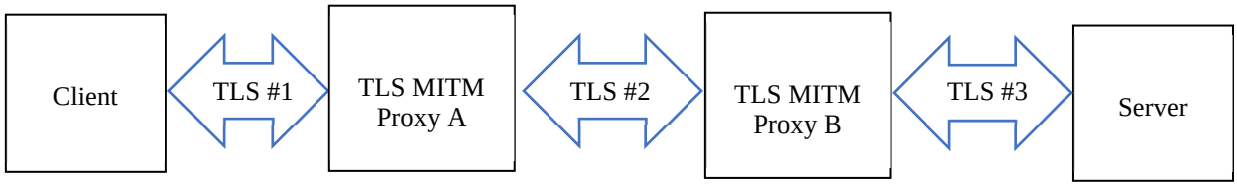

Figure C.2: Multiple split proxy architecture

3) **Annex-C.TM.3**: Endpoints trust that the TLS MITM split proxy will not be malicious and will only operate on traffic in an expected way. This does not extend to a split proxy being assumed always to make the same security-related choices that an endpoint would make.

4) **Annex-C.TM.4**: As with TLS, entities with access to the TLS private key of an endpoint, and cryptographic artefacts derived from it, are trusted.

5) **Annex-C.TM.5**: In each TLS connection, a cryptographic identity for each party is presented. Specifically, in each initial TLS connection, the client for that connection presents a certificate (or PSK). In each resumed TLS connection, a resumption ticket (or PSK) is sent, which satisfies this assumption too, in conjunction with Annex-C.TM.4.

NOTE 2: Annex-C.TM.5 is shown in Figure C.2, as TLS MITM Proxy A is aware of TLS MITM Proxy B and its presence as part of the connection.

C.2.2 Selection Process outcome

Based on this threat model and use case, these assumptions and threat modelling, the Selection Process assigned the label "Profile-Rejected" to many MSP Template Requirements. Additionally, the analysis of Clause C.3 shows that several MSP Mandatory (MM) Template Requirements are not met; therefore the TLS MITM split proxy profile does not conform to the MSP standard series.

C.3 Profile Requirements analysis

C.3.1 Data Protection

C.3.1.1 Overview

Table C.1 contains an overview of the Profile Requirement type assigned to each Data Protection Profile Requirement for a TLS MITM split proxy, and whether each Profile Requirement is satisfied or not.

With generous assumptions regarding trust (as detailed in Clause C.2.1), and an additional requirement which is not mandated by the split proxy, the TLS MITM split proxy satisfies the mandatory Data Protection MSP Template Requirements, defined in Clause 6.2.

Table C.1: Data Protection Profile Requirements met, and not met, by a TLS MITM split proxy

Ref	Data Protection Profile Requirement	MM / PM / PO / PNA / PR	Met?
E.DP.1	Endpoints shall protect confidentiality of sensitive data that they send.	MM	Y
E.DP.2	Endpoints may add protection to externally visible characteristics of application data to protect confidentiality of sensitive information about application activity. (This is commonly referred to as Traffic Analysis Protection).	PR	Y
E.DP.3	Endpoints shall protect integrity of application data.	MM	Y
E.DP.3.1	Endpoints shall protect application datagrams from modification in transit between authorized participants.	MM	Y
E.DP.3.2	Endpoints may protect application datagrams from Unauthorized modification by a middlebox.	PR	N
E.DP.3.3	Endpoints may protect the datastream from modification in transit between authorized participants.	PO	Y
E.DP.3.4	Endpoints may protect the datastream from Unauthorized modification by a middlebox.	PR	N
E.DP.4	Endpoints shall protect sensitive information about session state from Unauthorized disclosure, discovery, manipulation and creation.	MM	Y
E.DP.4.1	Endpoints shall protect the sensitive cryptographic state from Unauthorized disclosure, discovery, manipulation and creation.	MM	Y
E.DP.4.2	Endpoints **shall** protect the application state from replay and pre-play of data.	PM	Y
M.DP.1	Middleboxes shall protect confidentiality of sensitive data that they send.	MM	Y
M.DP.2	Middleboxes may add protection to externally visible characteristics of application data to protect confidentiality of sensitive information about application activity. (This is commonly referred to as Traffic Analysis Protection).	PR	N
M.DP.3	Middleboxes shall protect integrity of application data.	MM	Y
M.DP.3.1	Middleboxes shall protect application datagrams from modification in transit between authorized participants.	MM	Y
M.DP.3.2	Middleboxes may protect application datagrams from Unauthorized modification by a middlebox.	PR	N
M.DP.3.3	Middleboxes may protect the datastream from modification in transit between authorized participants.	PO	Y
M.DP.3.4	Middleboxes may protect the datastream from Unauthorized modification by a middlebox.	PR	N
M.DP.4	Middleboxes shall protect sensitive information about session state from Unauthorized disclosure, discovery, manipulation and creation.	MM	Y
M.DP.4.1	Middleboxes shall protect the sensitive cryptographic state from Unauthorized disclosure, discovery, manipulation and creation.	MM	Y
M.DP.4.2	Middleboxes **shall** protect the application state from replay and pre-play of data.	PM	Y
M.DP.5	Middleboxes may protect against protocol data fields being used as covert channels by validating the contents or otherwise. (This does not eliminate covert channels from externally visible characteristics such as timings and sizes).	PO	Y

C.3.1.2 MSP-Mandatory Profile Requirements

The MSP-Mandatory Profile Requirements for Data Protection are satisfied as described:

- E.DP.1 is satisfied because endpoints negotiate and agree a key and mechanism with the split proxy adjacent to it in the connection. This key and mechanism are then used to encrypt sensitive data. For the client endpoint, the split proxy chooses a mechanism from a list that the client presents; for the server endpoint, the split proxy presents a list of mechanisms and the server selects one that it supports.

- E.DP.3 is satisfied because E.DP.3.1 is satisfied:

  - E.DP.3.1 is satisfied because the datagrams in transit are integrity-protected by the endpoint using a keyed MAC, with a key that is only known by authorized parties.

- E.DP.4 is satisfied because E.DP.4.1 is satisfied:
  - E.DP.4.1 is satisfied because all parties that know the sensitive cryptographic state are trusted not to disclose it (Annex-C.TM.4) or to manipulate it (Annex-C.TM.0). Additionally, the TLS handshake transcript check, completed by both endpoints, prevents manipulation of the key exchange at the point of creation. In this key exchange, both parties in a connection contribute a nonce for that connection, ensuring a fresh session key every time. This prevents creation of an oracle and Unauthorized discovery.
- M.DP.1 is satisfied whenever one split proxy is present in the connection, as it negotiates the cryptography between itself and an endpoint. Where two split proxies are present in a connection, M.DP.1 can be met by placing an additional requirement on the split proxies, i.e. that the confidentiality mechanism used between them is acceptable to both endpoints. This mechanism is negotiated between the proxies without endpoint input and the endpoints will not know the outcome.

NOTE 1: In practice, this additional requirement can be ignored without endpoint knowledge, but it is consistent with our generous threat model and assumptions for Annex C to assume such a requirement can be imposed on split proxies.

- M.DP.3 is satisfied because M.DP.3.1 is satisfied.
- M.DP.3.1 is satisfied, as for M.DP.1, whenever one split proxy is present in the connection, as it negotiates the integrity mechanism between itself and an endpoint. Where two split proxies are present in a connection, M.DP.1 can be met by placing an additional requirement on the split proxies, i.e. that the integrity mechanism used between them is acceptable to both endpoints. This mechanism is negotiated between the proxies without endpoint input and the endpoints will not know the outcome.

NOTE 2: In practice, this additional requirement can be ignored without endpoint knowledge, but it is consistent with our generous threat model and assumptions for Annex C to assume such a requirement can be imposed on split proxies.

- M.DP.4 is satisfied because M.DP.4.1 is satisfied:
  - M.DP.4.1 is satisfied because all parties that know the sensitive cryptographic state are trusted not to disclose it (Annex-C.TM.4) or to manipulate it (Annex-C.TM.3). Additionally, the TLS handshake transcript check, completed by both endpoints, prevents manipulation of the key exchange at the point of creation. In this key exchange, both parties in a connection contribute a nonce for that connection, ensuring a fresh session key every time. This prevents creation of an oracle and Unauthorized discovery.

C.3.1.3 Profile-Mandatory Profile Requirements

There are no Profile Requirements of the type "Profile-Mandatory" for Data Protection for a TLS MITM split proxy.

C.3.1.4 Profile-Optional Profile Requirements

The Profile-Optional Profile Requirements for Data Protection are satisfied as described:

- E.DP.3.3 is satisfied optionally: the TCP layer can correct for reordered, missing or duplicated datagrams, and TLS can be configured to provide cryptographic verification of TCP's reconstruction of the datagram sequence.
- E.DP.4.2 is satisfied optionally when 0-RTT mode for TLS 1.3 is not supported, a new session key is created for every connection (including with resumption in non 0-RTT mode), which prevents replay. To create the session key, both parties in the connection contribute a nonce to the key exchange. TLS 1.3's 0-RTT is vulnerable to replay attacks, so this is subject to the implementation.
- M.DP.3.3 is satisfied optionally: the TCP layer can correct for reordered, missing or duplicated datagrams, and TLS can be configured to provide cryptographic verification of TCP's reconstruction of the datagram sequence.
- M.DP.4.2 is satisfied optionally when 0-RTT mode for TLS 1.3 is not supported, a new session key is created for every connection (including with resumption in non 0-RTT mode), which prevents replay. To create the session key, both parties in the connection contribute a nonce to the key exchange. TLS 1.3's 0-RTT is vulnerable to replay attacks, so this is subject to the implementation.

- M.DP.5 is satisfied optionally: since a MITM TLS proxy establishes two separate connections and has read and write access to the connections' content, the proxy can optionally implement examination and validation of protocol fields to protect against their possible use as covert channels.

C.3.1.5 Profile-Not-Applicable Profile Requirements

There are no Profile Requirements of the type "Profile-Not-Applicable" for Data Protection for a TLS MITM split proxy.

C.3.1.6 Profile-Rejected Profile Requirements

The Profile-Rejected Profile Requirements for Data Protection are not satisfied and rejected as described:

- E.DP.2 is rejected as the use case does not assume traffic analysis is a threat to the deployment; this feature is not supported and such protection can be implemented at the application layer as an additional measure.
- E.DP.3.2 is rejected as split proxies are strongly trusted (Annex-C.TM.2) and therefore this feature is not required.
- E.DP.3.4 is rejected as split proxies are strongly trusted (Annex-C.TM.2) and therefore this feature is not required.
- M.DP.2 is rejected as the use case does not assume traffic analysis is a threat to the deployment; this feature is not supported and such protection can be implemented at the application layer as an additional measure.
- M.DP.3.2 is rejected as split proxies are strongly trusted (Annex-C.TM.2) and therefore this feature is not required.
- M.DP.3.4 is rejected as split proxies are strongly trusted (Annex-C.TM.2) and therefore this feature is not required.

C.3.2 Transparency

C.3.2.1 Overview

The TLS MITM split proxy does not satisfy the mandatory Transparency MSP Template Requirements, defined in Clause 6.3.

Table C.2: Transparency Profile Requirements met by a TLS MITM split proxy

Ref	Transparency Profile Requirement	MM / PM / PO / PNA / PR	Met?
E.T.1	Endpoints shall receive suitable knowledge of all middlebox identities.	MM	N
E.T.1.1	Both endpoints shall receive suitable knowledge about the identity of all middleboxes authorized.	MM	N
E.T.1.2	Endpoints may receive knowledge about the identity of all refused middleboxes.	PNA	N
E.T.1.3	Endpoints shall be able to verify or otherwise confirm that they have the same knowledge as the peer endpoint of all middleboxes' identities that are authorized.	MM	N
E.T.2	Endpoints shall receive knowledge of all middlebox permissions and knowledge of all security mechanisms for data protection.	MM	N
E.T.3	Each endpoint shall be able to verify or otherwise confirm that they have the same knowledge (of middlebox permissions and security mechanisms for data protection) as the other endpoint.	MM	N
E.T.4	Endpoints may receive knowledge of the peer endpoint identity.	PR	N
E.T.4.1	The initiator endpoint may authenticate or otherwise verify the identity of the responder endpoint.	PR	N
E.T.4.2	The responder endpoint may authenticate or otherwise verify the identity of the initiator endpoint.	PR	N
E.T.5	Endpoints may verifiably audit activity of middleboxes.	PR	N

Ref	Transparency Profile Requirement	MM / PM / PO / PNA / PR	Met?
E.T.5.1	The destination endpoint may verifiably audit the activity of middleboxes.	PR	N
E.T.5.1.1	The destination endpoint may verify that data has transited and not bypassed each middlebox.	PR	N
E.T.5.1.2	The destination endpoint may verify whether a middlebox has modified data.	PR	N
E.T.5.1.3	The destination endpoint may verify the full change history of received data.	PR	N
E.T.5.2	The sending endpoint may verifiably audit the activity of middleboxes.	PR	N
E.T.5.2.1	The sending endpoint may verify that data has transited and not bypassed each middlebox.	PR	N
E.T.5.2.2	The sending endpoint may verify whether a middlebox has modified data.	PR	N
E.T.5.2.3	The sending endpoint may verify the full change history of received data.	PR	N
E.T.6	Endpoints may verify or otherwise confirm that middlebox access and middlebox permissions have been granted or denied.	PR	N
M.T.1	Middleboxes may receive knowledge of all middlebox identities.	PR	N
M.T.1.1	Middleboxes **shall** receive knowledge about the identity of all middleboxes authorized.	PM	Y
M.T.1.2	Middleboxes may receive knowledge about the identity of all refused middleboxes.	PNA	N
M.T.1.3	Middleboxes may be able to verify or otherwise confirm that they have the same knowledge as other participants of all middleboxes' identities that are authorized.	PR	N
M.T.2	Middleboxes may receive knowledge of all middlebox permissions and knowledge of all security mechanisms for data protection.	PR	N
M.T.3	Middleboxes may be able to verify or otherwise confirm that they have the same knowledge (of middlebox permissions and security mechanisms for data protection) as the other participants.	PR	N
M.T.4	Middleboxes may receive knowledge of either or both endpoint identities.	PO	Y
M.T.4.1	Middleboxes may receive knowledge about the identity of the responder endpoint.	PO	Y
M.T.4.2	Middleboxes may receive knowledge about the identity of the initiator endpoint.	PO	Y
M.T.5	Middleboxes may verifiably audit activity of other middleboxes.	PR	N
M.T.5.1	Middleboxes may verifiably audit activity of other participants on received data.	PR	N
M.T.5.1.1	Middleboxes may verify that received data has transited and not bypassed each middlebox.	PR	N
M.T.5.1.2	Middleboxes may verify whether another middlebox has modified received data	PR	N
M.T.5.1.3	Middleboxes may verify the full change history of received data.	PR	N
M.T.5.2	Middleboxes may verifiably audit activity of other participants on sent data.	PR	N
M.T.5.2.1	Middleboxes may verify that sent data has transited and not bypassed each middlebox.	PR	N
M.T.5.2.2	Middleboxes may verify whether another middlebox has modified sent data.	PR	N
M.T.5.2.3	Middleboxes may verify the full change history of sent data.	PR	N

C.3.2.2 MSP-Mandatory Profile Requirements

All of the MSP-Mandatory Profile Requirements for Transparency are not satisfied by a TLS MITM split proxy, as described:

- E.T.1 is not satisfied because E.T.1.1 and E.T.1.3 are not satisfied, as described below:
  - E.T.1.1 is not satisfied because where two (or more) split proxies are on the connection, each of the client and server have no knowledge of at least one split proxy. In the case of one split proxy, the client endpoint can have no knowledge that its TLS connection terminates at a split proxy rather than the server it intends to connect to.
  - E.T.1.3 is not satisfied because, in the case of two split proxies, endpoints do not have the same knowledge of split proxies: they only learn about the one closest to them. With our generous assumption that client certificates are sent in each TLS connection, in the case of one split proxy endpoints do have the same knowledge of its identity, but E.T.1.3 is still not met because there is no mechanism available for the endpoints to confirm this.
- E.T.2 and E.T.3 are not satisfied because each endpoint has no knowledge of security mechanisms for data protection used by the other endpoint, even in a deployment with only one split proxy. In a more general case with *n* MITM split proxies, each endpoint has no knowledge of the data protection mechanisms for any of the TLS connection in which they are not a direct party, which is *n* connections.

C.3.2.3 Profile-Mandatory Profile Requirements

The Profile-Mandatory Requirements for Transparency are satisfied as described:

- M.T.1.1 is satisfied because split proxies will receive information about the identity of split proxies adjacent to it. As client certificates are always in use (Annex-C.TM.5) and there are at most two split proxies in the connection (Annex-C.TM.1), a split proxy will always receive a certificate from all split proxies in the connection (zero or one, in addition to itself).

C.3.2.4 Profile-Optional Profile Requirements

The Profile-Optional Profile Requirements for Transparency are satisfied as described:

- M.T.4 is satisfied optionally when M.T.4.1 and M.T.4.2 are both met:
  - M.T.4.1 is satisfied where one split proxy is deployed; where two split proxies are on the connection (Annex-C.TM.1), only the split proxy adjacent to the server will gain knowledge of the server endpoint identity from its certificate.
  - M.T.4.2 is satisfied where one split proxy is deployed; where two split proxies are on the connection (Annex-C.TM.1), only the split proxy adjacent to the client will gain knowledge of the client endpoint identity from its certificate.

C.3.2.5 Profile-Not-Applicable Profile Requirements

The Profile-Not-Applicable Profile Requirements for Transparency are not applicable as described:

- E.T.1.2 is not applicable as there is no reasonable definition of "refused" in the split proxy setup other than a party terminating a connection with a split proxy or deliberately routing around it. Where that occurs, there is no defined and general mechanism for informing an endpoint that was not party to the terminated connection or rerouting of the split proxy's identity.
- M.T.1.2 is not applicable as there is no reasonable definition of "refused" in the split proxy setup other than a party terminating a connection with a split proxy or deliberately routing around it. Where that occurs, there is no defined and general mechanism for a split proxy to determine if another split proxy on the connection rerouted the traffic around it or terminated the connection.

C.3.2.6 Profile-Rejected Profile Requirements

The Profile-Rejected Profile Requirements for Transparency are not satisfied and rejected as described:

- E.T.4 is not met as neither E.T.4.1 nor E.T.4.2 are met:
  - E.T.4.1 is not met as the client receives a certificate from the split proxy, acting as the server, not the certificate from the real server (the responder endpoint). Therefore, the client cannot verify the identity of the responder endpoint (the original server certificate).
  - E.T.4.2 is not met as the server receives a certificate from the split proxy, acting as the client, not the certificate from the real client (the initiator endpoint). Therefore, the server cannot verify the identity of the initiator endpoint (the original client certificate).
- E.T.5 is not satisfied because E.T.5.1 and E.T.5.2 are not satisfied:
  - E.T.5.1 and E.T.5.2 are not satisfied for the same reason; an endpoint cannot verify what activity a split proxy has performed. This is true, despite the strong level of trust in the split proxy (Annex-C.TM.3), because features relating to audit are not supported.
  - E.T.5.2 is not satisfied because the sending endpoint cannot verify what activity a split proxy has performed, even if by assumption it does trust the proxy to act according to its own expectations; nor can the sending endpoint attribute activity to a specific proxy when there are two (or more) in the connection.
- E.T.6 is not satisfied because, although one endpoint can always verify and confirm that the adjacent split proxy is present, in the case of two split proxies the other endpoint cannot. This is because, whilst one endpoint is configured to accept and recognize that split proxy (Annex-C.TM.2), the other is not necessarily. Beyond its adjacent split proxy, neither endpoint has visibility of other split proxies on the connection. Therefore, neither endpoint can always confirm split proxy permissions or access.
- M.T.1 is not satisfied because M.T.1.2 is not applicable (see Clause C.3.2.5) and M.T.1.3 is not satisfied:
  - M.T.1.3 is not satisfied because where there are two split proxies in the connection, the endpoints will only have knowledge of the identity of the split proxy adjacent to itself; in this case, either split proxy has greater knowledge of the split proxy identities than either endpoint. Where there is one split proxy in the connection, though a split proxy will always receive a certificate from all split proxies in the connection (zero or one, in addition to itself - Annex-C.TM1), it will not know if it is the only split proxy in the connection.
- M.T.2 and M.T.3 are not satisfied because each middlebox will only know the security mechanisms for data protections used to connect to adjacent participants. Therefore, not all participants have the same knowledge of security mechanisms for data protection.
- M.T.5 is not satisfied because M.T.5.1, M.T.5.2, M.T.5.1.1, M.T.5.1.2, M.T.5.1.3, M.T.5.2.1, M.T.5.2.2 and M.T.5.2.3 are not satisfied. They are not satisfied as one split proxy cannot verify what activity another split proxy has performed. This is true, despite the strong level of trust in the middlebox (Annex-C.TM.3), because features relating to middlebox audit are not supported.

C.3.3 Access Control

C.3.3.1 Overview

The TLS MITM split proxy does not satisfy the mandatory Access Control MSP Template Requirements, defined in Clause 6.4.

Table C.3: Access Control Profile Requirements met by a split proxy

Ref	Access Control Profile Requirement	MM / PM / PO / PNA / PR	Met?
E.AC.1	Only endpoints shall grant or deny middlebox access and middlebox permissions.	MM	Y
E.AC.1.1	Middlebox access shall be granted by at least one endpoint.	MM	Y
E.AC.1.2	Middlebox permissions shall be granted by the same endpoint or endpoints that granted access.	MM	Y
E.AC.1.3	The profile may support multiple levels for middlebox permissions.	PR	N
E.AC.1.4	Endpoints may authorize middlebox permissions per context.	PR	N
E.AC.1.5	Only endpoints shall deny middlebox access or middlebox permissions. (A middlebox can still block the entire connection between suspected malicious endpoints).	MM	Y
E.AC.2	The endpoint(s) that grant(s) access to a middlebox shall authenticate or otherwise confirm its identity before granting access.	MM	N
E.AC.3	At least one endpoint shall choose all security mechanisms for data protection.	MM	N
E.AC.4	Endpoints may grant middlebox access and middlebox permissions only through mutual agreement with the peer endpoint.	PR	N
E.AC.5	Endpoints may authenticate or otherwise verify the identity of all middleboxes whose access is granted by the other endpoint.	PR	N
M.AC.1	Middleboxes shall authenticate or otherwise confirm any participant identity they use for an identity-dependent action. This action is not granting or denying access to an MSP connection, which shall fall within endpoint remit only (E.AC.1). (This stops a middlebox unlocking access to data or services for an identity that has not been checked by the middlebox).	MM	Y
M.AC.2	Middleboxes may authenticate or otherwise confirm participant identities.	PO	Y
M.AC.2.1	Middleboxes may authenticate or otherwise confirm the initiator endpoint identity	PO	Y
M.AC.2.2	Middleboxes may authenticate or otherwise confirm the responder endpoint identity	PO	Y
M.AC.2.3	Middleboxes **shall** authenticate or otherwise confirm all middlebox identities	PM	Y
M.AC.3	A middlebox may know that its access has been withheld. (Meeting this requirement implies it is not possible to deceive a middlebox into believing it has access).	PR	N

C.3.3.2 MSP-Mandatory Profile Requirements

The MSP-Mandatory Profile Requirements for Access Control are not all satisfied as described:

- E.AC.1 is satisfied as E.AC.1.1, E.AC.1.2 and E.AC.1.5 are satisfied:
  - E.AC.1.1 and E.AC.1.2 are satisfied because the number of proxies in the connection is capped at two (Annex-C.TM.1) and other proxies will not add additional proxies to the connection (Annex-C.TM.3).
  - E.AC.1.5 is satisfied because there is no reasonable definition of "deny" in the split proxy setup, other than a party terminating a connection with a split proxy or deliberately routing around it, which would violate Annex-C.TM.3.
- E.AC.2 is not satisfied when two split proxies are present in the connection (Annex-C.TM.1). In this case, each split proxy is known only by one endpoint; each endpoint knows the identity of the middlebox adjacent to it (Annex-C.TM.5).
- E.AC.3 is not satisfied because the split proxy will select the mechanism for data protection between itself and the client.
- M.AC.1 is satisfied because the only mechanism to control access to the connection is the TLS resumption mechanism. This is tied to a key known only to the original participants; possession of the key for resumption therefore authenticates that it is the same party (Annex-C.TM.5).

C.3.3.3 Profile-Mandatory Profile Requirements

The Profile-Mandatory Profile Requirements for Access Control are satisfied as described:

- M.AC.2.3 is always satisfied because there are one or two split proxies in the connection (Annex-C.TM.1). Therefore, a split proxy will either be the only split proxy in the connection (and will have knowledge of its own identity), or it will receive information of the identity of the other split proxy in the connection (Annex-C.TM.5).

C.3.3.4 Profile-Optional Profile Requirements

The Profile-Optional Profile Requirements for Access Control are satisfied as described:

- E.AC.5 is optionally satisfied if only one split proxy is present in the connection (Annex-C.TM.1). In this case, the identity of the split proxy is verified by both endpoints, regardless of which endpoint granted the access. In the case of two split proxies, each endpoint can only verify the identity of the proxy closest to itself.
- M.AC.2 is optionally satisfied if only one split proxy is present in the connection (Annex-C.TM.1):
  - M.AC.2.1 is satisfied optionally where only one split proxy is present in the connection; this split proxy will receive a certificate from the client (Annex-C.TM.5). Where two split proxies are present in the connection, this is not satisfied as only one of them will receive the client certificate.
  - M.AC.2.2 is satisfied optionally where only one split proxy is present in the connection; this split proxy will receive a certificate from the server. Where two split proxies are present in the connection, this is not satisfied as only one of them will receive the server certificate.

C.3.3.5 Profile-Not-Applicable Profile Requirements

The Profile-Not-Applicable Profile Requirements for Access Control are not applicable as described:

- M.AC.3 is not applicable because there is no reasonable definition of "withheld" in the split proxy setup other than a party terminating a connection with a split proxy or deliberately routing around it. Where that occurs, there is no defined and general mechanism for informing an endpoint that was not party to the terminated connection or rerouting of the split proxy's identity.

C.3.3.6 Profile-Rejected Profile Requirements

The Profile-Rejected Profile Requirements for Access Control are not satisfied and rejected as described:

- E.AC.1.3 is rejected because there is no mechanism to support different levels of permission; once access is granted, a split proxy can read, modify and delete data.
- E.AC.1.4 is rejected because there is no mechanism to restrict access to certain types of data; permissions can only be granted to all data, and these permissions are not segmented within the stream.
- E.AC.4 is rejected because there is no mechanism to support mutual endpoint agreement for split proxy access. Access is granted by one endpoint, as stated in assumption Annex-C.TM.2.

C.3.4 Good Citizen

The TLS MITM split proxy does not satisfy the mandatory Good Citizen MSP Template Requirements (MSP-Mandatory Profile Requirements), as described below in the present clause.

Table C.4 contains an overview of the Profile Requirement type assigned to each Good Citizen Profile Requirement for a TLS MITM split proxy, and whether each Profile Requirement is satisfied or not. All Good Citizen Profile Requirements are of the type "MSP-Mandatory".

Table C.4: Good Citizen Profile Requirements met by a split proxy

Ref	Good Citizen Profile Requirement	MM / PM / PO / PNA / PR	Met?
E.GC.1	Resource attacks that use an endpoint action or request shall have some attribution to the attacker.	MM	N
E.GC.1.1	Any party being asked to expend significant resource due to an endpoint request, shall have some attribution of the request to the endpoint.	MM	N
E.GC.2	An MSP profile shall not provide a significant amplification factor for a resource attack that uses an endpoint action or request.	MM	Y
E.GC.2.1	Where an endpoint sends MSP protocol messages that request a significant amplification factor on resource expenditure, then one of the following two things shall happen: either the recipient is not forced to accept the request or the requesting endpoint expends commensurately amplified resource as a consumer of the result.	MM	Y
M.GC.1	Resource attacks that use a middlebox action or request shall have some attribution to the attacker.	MM	N
M.GC.1.1	Any party being asked to expend significant resource due to a middlebox request, shall have some attribution of the request to the middlebox.	MM	N
M.GC.2	An MSP profile shall not provide a significant amplification factor for a resource attack that uses a middlebox action or request.	MM	Y
M.GC.2.1	Where a middlebox sends MSP protocol messages that request a significant amplification factor on resource expenditure, then one of the following two things shall happen: either the recipient is not forced to accept the request or the requesting middlebox expends commensurately amplified resource as a consumer of the result.	MM	Y
M.GC.3	Middleboxes may be able to drop out of a connection, without breaking or degrading the connection for other participants, to counter an attempted resource attack.	PR	N

The Conformance Analysis in the present clause describes how MSP-Mandatory and Profile-Rejected Good Citizen Profile Requirements are not met by a TLS MITM split proxy:

- E.GC.1 and E.GC.1.1 are not met because, even though the split proxy has its own certificate to present to the server (Annex-C.TM.5), the split proxy removes attribution to the 'real' client and it is not possible for a server to know the 'real client'. This is worsened further if there are two split proxies in the connection (Annex-C.TM.1).
- E.GC.2 is met as E.GC.2.1 is met:
  - E.GC.2.1 is met as there are no messages that would cause this behaviour.
- M.GC.1 and M.GC.1.1 are not met because, although the split proxy has its own certificate to present to the server (Annex-C.TM.5), the split proxy removes attribution to the 'real' client and it is not possible for a middlebox to know the 'real client'. The same is true in reverse; the client does not know the 'real' server. This is worsened further if there are two split proxies in the connection (Annex-C.TM.1).
- M.GC.2 is met as M.GC.2.1 is met:
  - M.GC.2.1 is met as there are no messages that would cause this behaviour.
- M.GC.3 is not met because if the split proxy drops out of the connection, the connection has to restart.

C.4 Summary

Even with the generous threat model assumptions in Clause C.2.1, a TLS MITM split proxy does not meet many of the MSP-Mandatory Profile Requirements (equivalently, the Mandatory MSP Template Requirements). Of the four principles described in Clause 5.3 (Data Protection, Transparency, Access Control and Good Citizen), the TLS MITM split proxy does not meet Mandatory MSP Template Requirements in three of them, as described in Clauses C.3.2, C.3.3 and C.3.4. Further, it is only through generous assumptions (that can fail in practice) that the Mandatory MSP Template Requirements under the remaining principle, Data Protection, can be assessed as met, as described in Clause C.3.1.

This insufficient MSP Conformance Analysis indicates that the MSP series brings improvements in all four principles over existing MITM TLS split proxies, through use of the MSP Security Framework and subsequent MSP Conformance Analysis.

History

Document history		
V1.1.1	December 2020	Publication