CPNI

Centre for the Protection of National Infrastructure

CAE:

BLOCKS AND CONNECTION RULES

CONTENTS

FIGURES

TABLES

01. INTRODUCTION

This document provides guidance on a set of building blocks and the Claims, Argument and Evidence (CAE) ‘normal form’, discussing the important connection rules that place constraints on the way that claims, argument and evidence are combined in a CAE structure.

Building blocks are fragments that are useful for expressing the safety justification. These can be used to decide which type of argument to apply for a specific type of claim, and guide the user through the process of elaborating that fragment in a careful manner, aiming at creating a complete and clear argument.

Connection rules place constraints on the manner in which the components of CAE are linked. These rules are important as they help to achieve consistency in the presentation of a CAE structure, and more importantly help to avoid some of the risks that may arise from a free form approach.

02.
SIGNPOSTING

This is the third CAE guide in the stack of resources for security-informed safety assurance. Figure 1 below shows its location in the set of guides (highlighted in red).

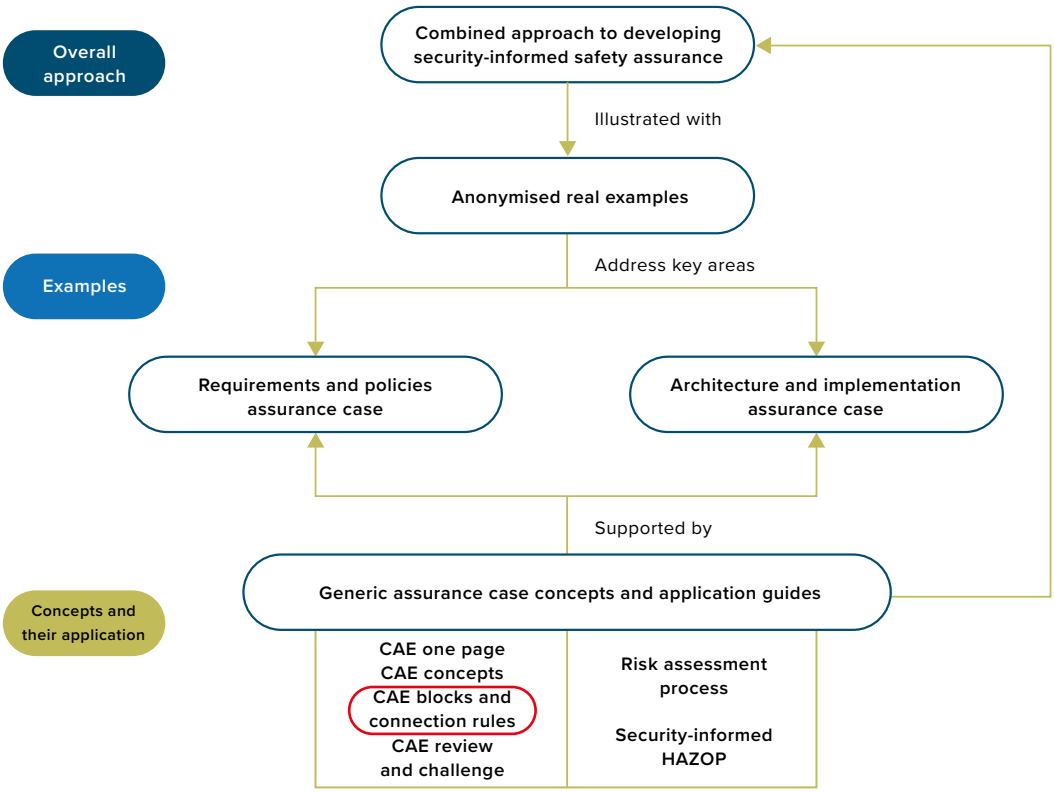

Figure 1: Location of this guide in the set of resources

03.
CAE BLOCKS

3.1 WHAT ARE CAE BUILDING BLOCKS?

CAE building blocks are a series of archetypal CAE fragments that were derived from an empirical analysis of real cases in various domains, where cases were analysed to determine what they were trying to express. They enhance the classical CAE approach [1] with a standardised structure and an approach to how arguments are addressed. The five basic CAE building blocks are:

- Decomposition – this partitions some aspect of the claim in a “divide and conquer” approach
- Substitution – refines a claim about an object into another claim about an equivalent object
- Concretion – gives a more precise definition to some aspect of the claim
- Calculation or proof – used when some value of the claim can be computed or proved
- Evidence incorporation – incorporates evidence that directly supports the claim.

The summary and the structure of these basic blocks are provided in Table 1. Additional information and guidance has been published [2], with recent developments described in Assurance 2.0 reports [3]. CAE building blocks are based on the CAE normal form, described in Section 4, with further simplification and enhancements. The block structure contains enhancements in how arguments are addressed. Specific rules of the argument called ‘side-claims’ explain why the top-level claim can be deduced from the subclaims, and under what circumstances the argument is valid.

The side-claim is in fact a type of claim and there may be a need to challenge and demonstrate this for the specific case. This can be done either by justifying the side-claim directly or by supporting the side-claims with further subclaims and argument. The graphical scheme of a generic CAE block structure is shown in Figure 2 below. It shows subclaims supporting an argument that justifies a top-level claim, with some of the key properties of the argument expressed as the side-claim and supported by the system information and external backing.

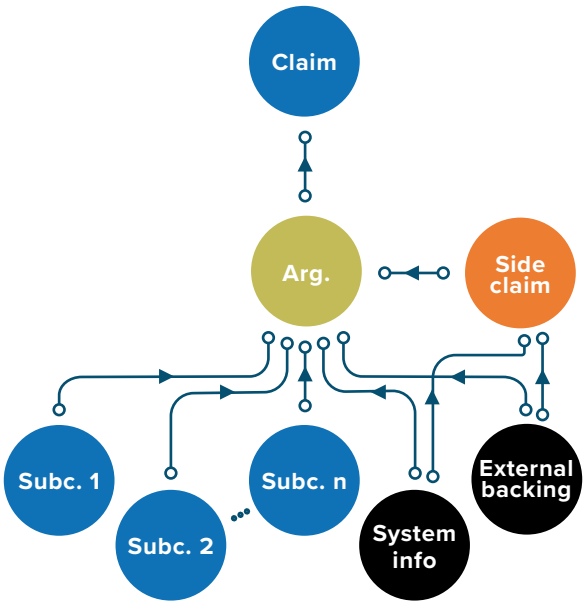

Figure 2: Generic CAE Block Structure

The overall justification for the block and its application can be included in the argument node narrative or accompanying text for the block. As shown in Figure 2, both the argument node and the side-claim can be supported by additional data: system information and external backing. The former includes any system-related information that drives the justification: models of system objects and properties, information from the product specification or user documentation, etc. The latter includes facts, guidance, theorems and theories that are appealed to as true statements of facts external to the claim. The side-claim might rely on external backing when demonstrating the claim, so if the backing itself is questionable, it must also be justified.

The side-claim serves to remind of the reasoning that is claimed to be true of the block. This may prompt further detailing of the CAE structure, as breaking it into more steps may identity how the justification can be made valid. Or it may be that the block is a simple one (perhaps the top-level claim is just the conjunction of the subclaims). Alternatively it may be that some general result or authority can be used to justify the side-claim. The five CAE Building Blocks as shown in Table 1.

Structure	Description
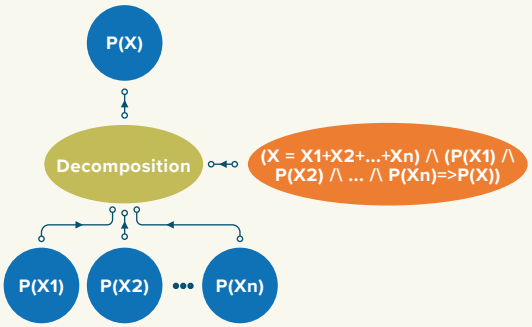	Decomposition block This block is used to claim that a conclusion about the whole object or property can be deduced from the claims or facts about constituent parts. Decomposition blocks can also be used to incorporate defeaters into the case.
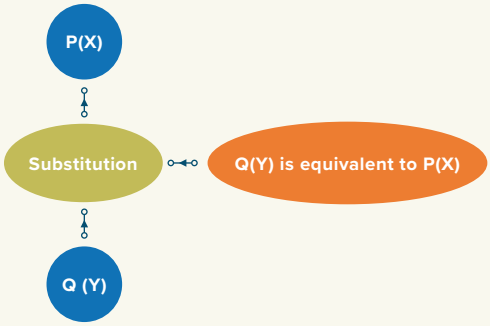	Substitution block This block is used to claim that if a property holds for one object, then it holds for an equivalent object. Similarly, if a property holds for some object, then an equivalent property will also hold for the same object. The nature of the 'equivalence' will vary with the object and property and will need to be defined.
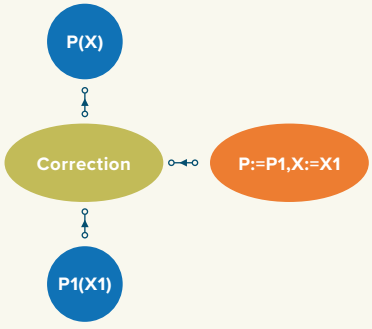	Concretion block This block is used when a claim needs to be given a more precise definition or interpretation.

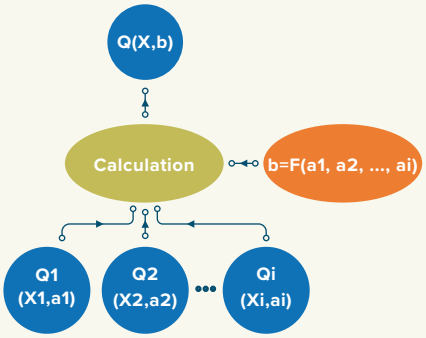	Calculation block This block is used to claim that the value of a property of a system can be computed from the values of related properties of other objects. Show that the value *b* of property Q (X, b, E, C) of system X in environment E and confidence C can be calculated or proved from values Q_1 (X_1, a_1, E, C), Q_2, (X_2, a_2, E, C),..., Q_n (X_n, a_n, E, C).
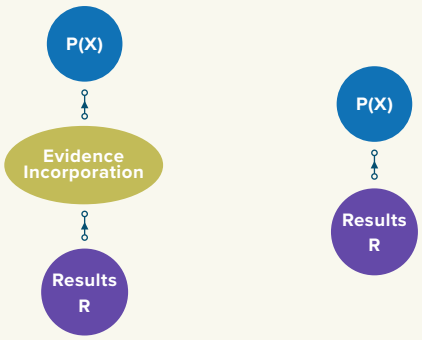	Evidence incorporation block This block is used to incorporate evidence elements into the case. A typical application of this block is at the edge of a case tree where a claim is shown to be directly satisfied by its supporting evidence.

Table 1: Basic Building Blocks for Assurance Cases

3.2 HOW CAN THEY BE USED?

The CAE blocks are meant to support the creative process of constructing a case. They do not themselves show how to architect cases, but provide a series of standardised ways of proceeding, either when a case is being developed top down or bottom up. The question that the case developer has to address is “which block could I use now?”.

In order to support the teaching and deployment of CAE Building Blocks, a visual guidance shown in Figure 3 has been created. The ‘helping hand’ is designed to help people structure assurance cases in an easier and more intuitive way by providing a cheatsheet with some hints and questions to answer. Instead of wondering what to do next and how to better expand the case, this approach shifts the question to an easier one: “which block is best to use?” and helps to find the answer by following the provided guidance.

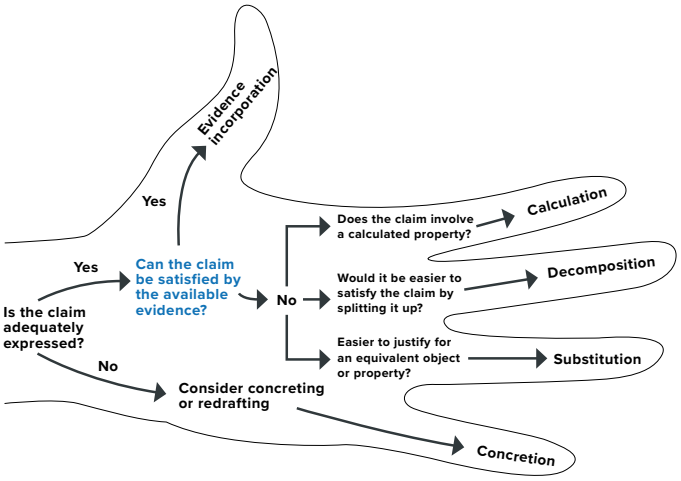

Figure 3: ‘Helping hand’ – high-level guidelines for selecting the CAE building block

The CAE blocks can be applied in three steps:

1. The first step involves selection of the block using the helping hand and its instantiation to the claims being considered. A rationale for selecting this block is given in narrative.
2. The next step involves adding the side-claim that defines the argument rule being used.
3. The third step is to support the argument rule with narrative. Usually this requires further support either with a single evidence incorporation block or with a more detailed CAE structure.

The three steps are summarised in Figure 4 below.

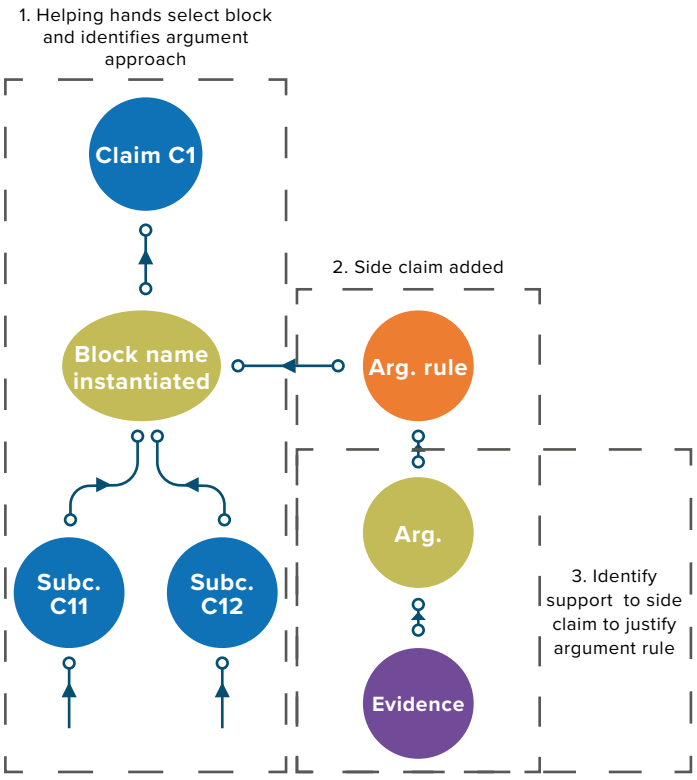

Figure 4: Summary of three steps of using CAE Block

Part of the skill in architecting a CAE fragment is in identifying the key properties that should be justified separately in the side-claim. Taking the simple example of Figure 5, it might be found that the only side-claim identified is C11 /\ C12 => C1, which makes the verification trivial (just modus ponens) [4] but pushes the justification into that for the claim. Or it may be found that if a property referenced in the subclaims distributes then C1 can be inferred from the subclaims.

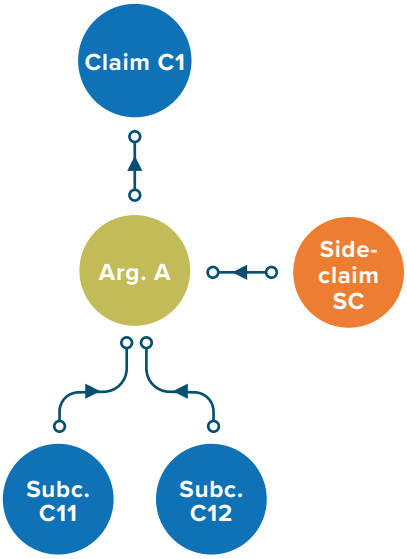

Figure 5: Simple example of CAE fragment with side-claim

The side-claim serves to remind of the reasoning that is claimed to be true of the block. This may prompt to detail the CAE structure further, as breaking it into more steps may identity how the justification can be made valid. Or it may be that the block is a simple one (perhaps C1 is just the conjunction of C11 and C12). Alternatively it may be that some general result or authority can be appealed to in order to justify the side-claim.

04.
CAE CONNECTION RULES

A number of rules (referred to in the rest of the document as the 'CAE normal form') make a CAE structure more consistent and easier to read. These rules place constraints on the way that claims, arguments and evidence may be linked in a CAE structure. These rules are to help avoid some issues arising from a free-form style of construction, yet recognise that different styles are appropriate. For example, in the initial stages of case exploration, a more brainstorming and free-form approach is helpful.

CAE normal form has the following connection rules:

1. Claim nodes may only be connected to argument nodes, i.e. evidence cannot support a claim without an intervening argument. Claims cannot be split into subclaims without an argument.
2. Argument nodes may only be connected to claim and evidence nodes, i.e., argument nodes are not connected to other argument nodes.
3. Each argument node may only have one outbound link to a claim node, i.e. it can only support one claim.
4. Evidence nodes may only be connected to argument nodes.
5. Each claim is to be supported by one and only one argument. If two arguments appear to be reinforcing the same claim, consider why this is so and explain the increase in confidence or reduction in assumption doubt that might be brought about. This will involve making the claims more precise and adding an additional argument or merging two arguments into one.
6. Argument nodes must be supported by at least one subclaim or evidence node.
7. Evidence nodes represent the bottom of the safety argument and are not supported; they represent agreed facts.
8. A claim or a subclaim may support more than one argument and similarly, one evidence node may be used by more than one argument.

These connection rules do not apply to context nodes, which can be connected to any type of node.

The table below (Table 2) summarises what is allowed and what is not when linking the various components of the CAE, assuming the direction of links is flowing upwards and towards the top-level claim.

Allowed	Not allowed
Claim to Argument or several Arguments	Claim to Evidence Claim to Claim Unsupported Claims
Argument to single Claim	Argument to Evidence Argument to Argument (needs claim between) Argument to multiple Claims Unsupported Arguments (but might occur in case development)
Evidence to Argument	Evidence to Claim Evidence to Evidence (sometimes used to show structure of evidence)

Table 2: CAE linking rules

The use of the CAE normal form has the effect of encouraging the safety case author to be more precise about the claims being made and more explicit about the supporting arguments for those claims. Furthermore, following the CAE normal form will help to achieve consistency in how CAE is used (developed and read) within an organisation.

An example of the application of the rules is shown in Figure 6, where the parts not in CAE normal form are highlighted in red.

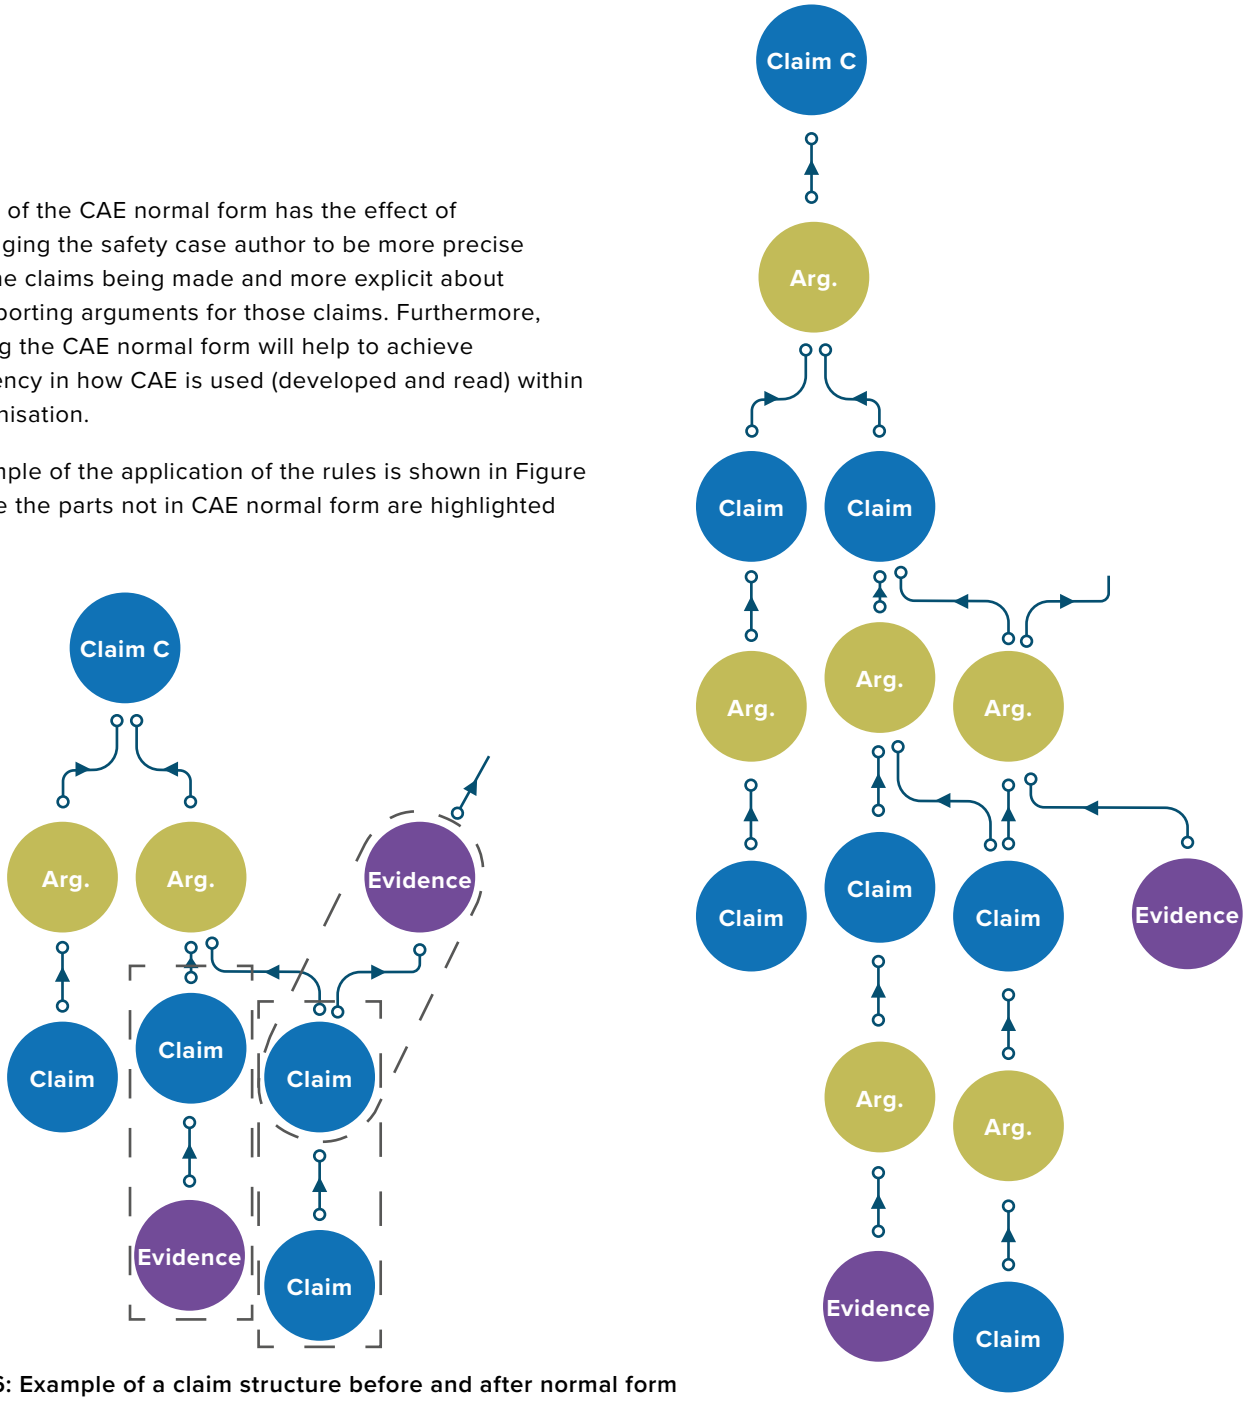

Figure 6: Example of a claim structure before and after normal form

In this example, it is possible to see an evidence node being directly connected to a claim. This goes against the rules suggested by the CAE normal form. An argument is needed between them, explaining how and why the evidence supports the claim, potentially including statements regarding the quality and trustworthiness of the evidence.

A lot of these practices are common in the CAE community, typically because people instinctively wish to 'tell the story' as it flows in their minds and naturally in conversation. It can leave room for assumptions to go unnoticed, positive bias to occur, misunderstanding to take place, or just reduce efficiency by requiring the audience to ask more questions until the argument is eventually rephrased. CAE normal form helps avoid risks associated with this, and also helps achieve consistency in how an organisation uses CAE.

In addition to the restrictions posed by CAE normal form, arguments to be conjunctions of the subclaims and evidence are required. Note that when satisfying a subclaim, it must not be forgotten that the other subclaims are also true: the graphical format may hide significant dependencies. A disjunctive combination of claims (logical OR) is not normally appropriate: even when there is diversity; both subclaims are usually required to ensure that the parent claim holds with sufficient confidence [5].

4.1 EVOLVING THE TOPOLOGY OF THE CASE

This section shows how a CAE structure may evolve as it is developed in light of the above rules.

The first stage of developing a case might be a brainstorm that identifies a top-level claim along with the five supporting arguments as shown in Figure 7.

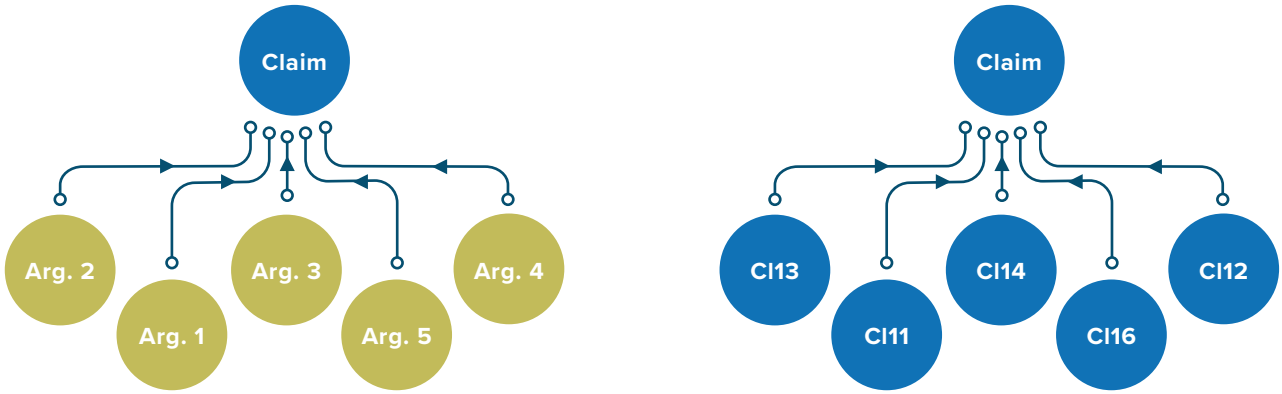

Figure 7: Initial structure

When reviewed it is likely that these are not arguments but actually supporting subclaims. So the first stage of evolving the structure is to redraw this as shown in the right-hand side of the figure. If in fact there really were five arguments then more analysis of the claim is probably needed.

In reclassifying the arguments as claims, there will probably be a need to rework and update these new claims.

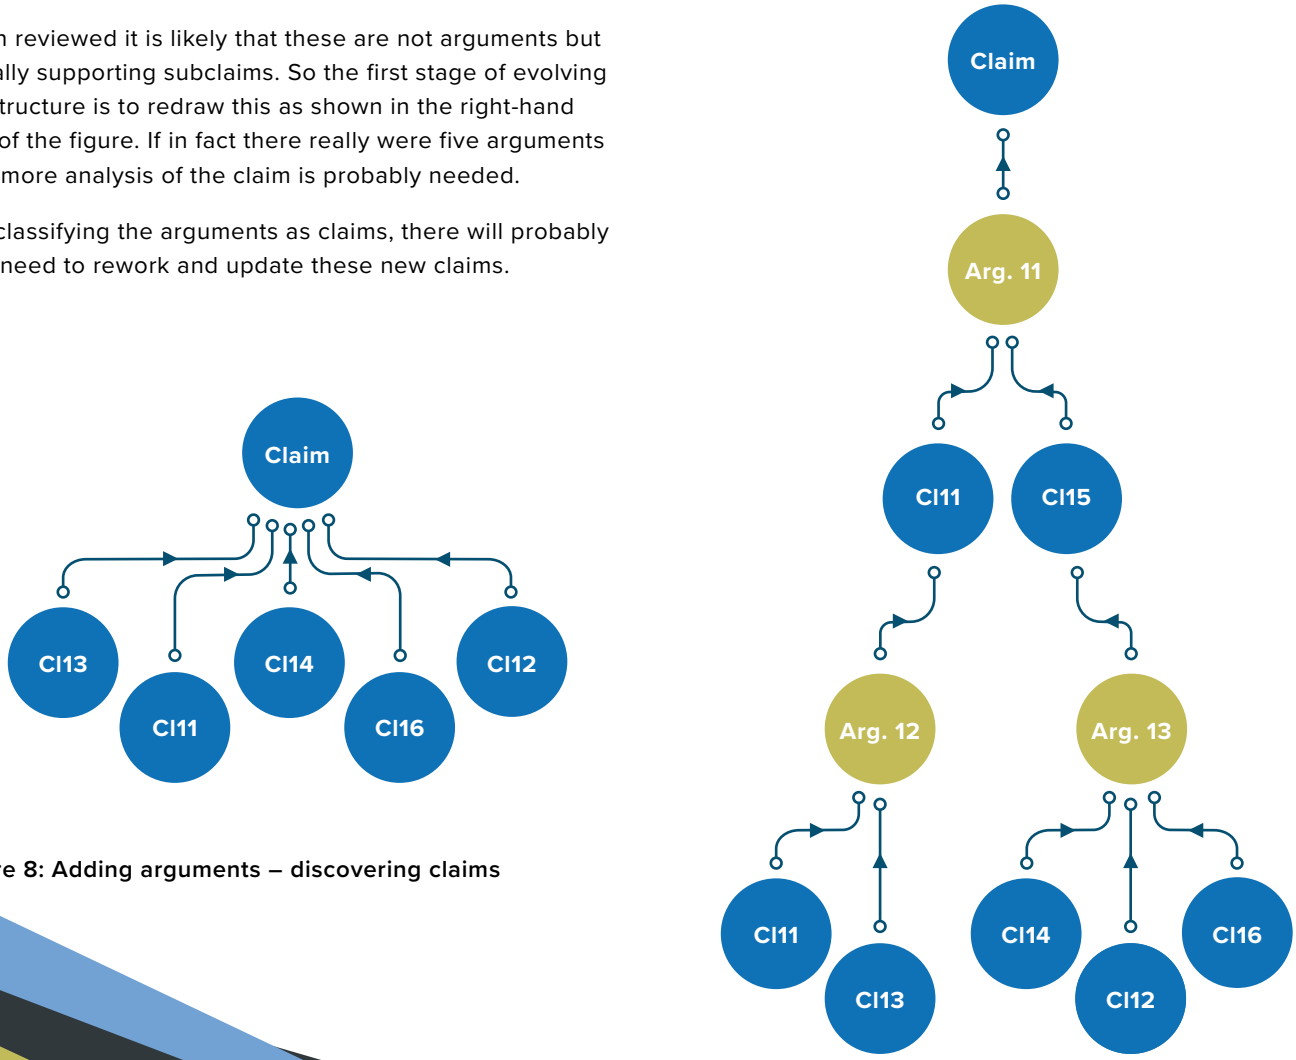

Figure 8: Adding arguments – discovering claims

Having established the top-level claim and some of the subclaims that support it, it is necessary to identify the arguments that give the reasons why these subclaims support the top-level claim. This is where the CAE blocks are helpful, as they can be used either bottom up or top down to arrive at a structure as in the right-hand side of Figure 8. In doing so, intermediary subclaims have been identified. Also it is found that a claim was missing in Figure 6. Now it can be seen how the top claim is supported by two main “legs”.

The next stage of the CAE evolution is to identify and map the evidence to the subclaims. Here it is found that one subclaim is unsupported by evidence and that some evidence supports several claims. This may warrant a further decomposition of the structure to understand in what way the evidence contributes: more precise claims may be necessary to see the role of the evidence and to assess if it is redundant, as there may be savings from not using it. The resulting structure is shown Figure 9 below.

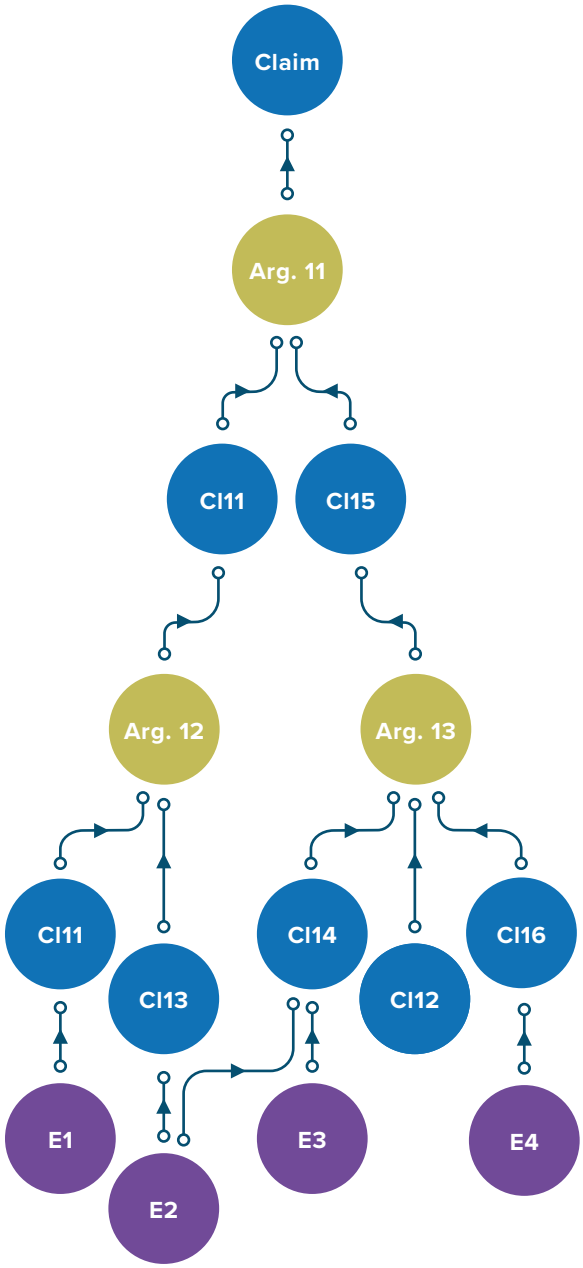

Figure 9: Identifying the role of evidence and gaps

Although the CAE structure in Figure 9 is small enough to be readily assimilated, in more complicated CAE structures there may be a need to provide a summary. Figure 10 illustrates two options for summarising. One takes the top of the case, showing the two legs, and provides the main claims and main evidence sources. The second approach is to suppress the arguments and provide a structure that shows all the claims. Of course, summaries by their nature omit things; in the first summary it is not shown that there is an unsupported claim.

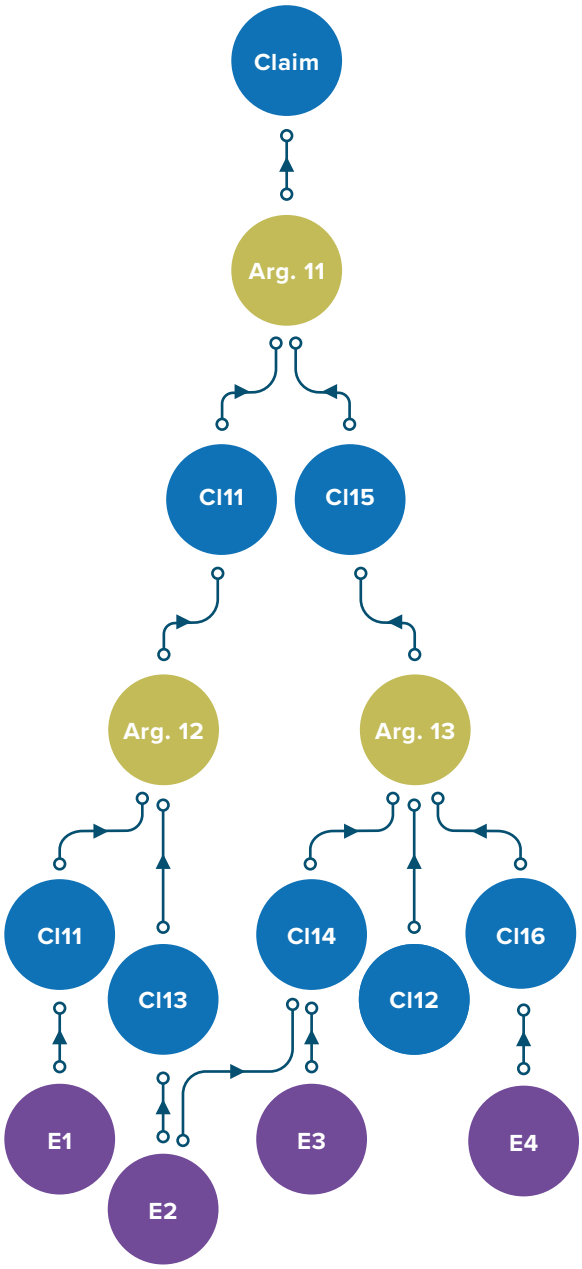

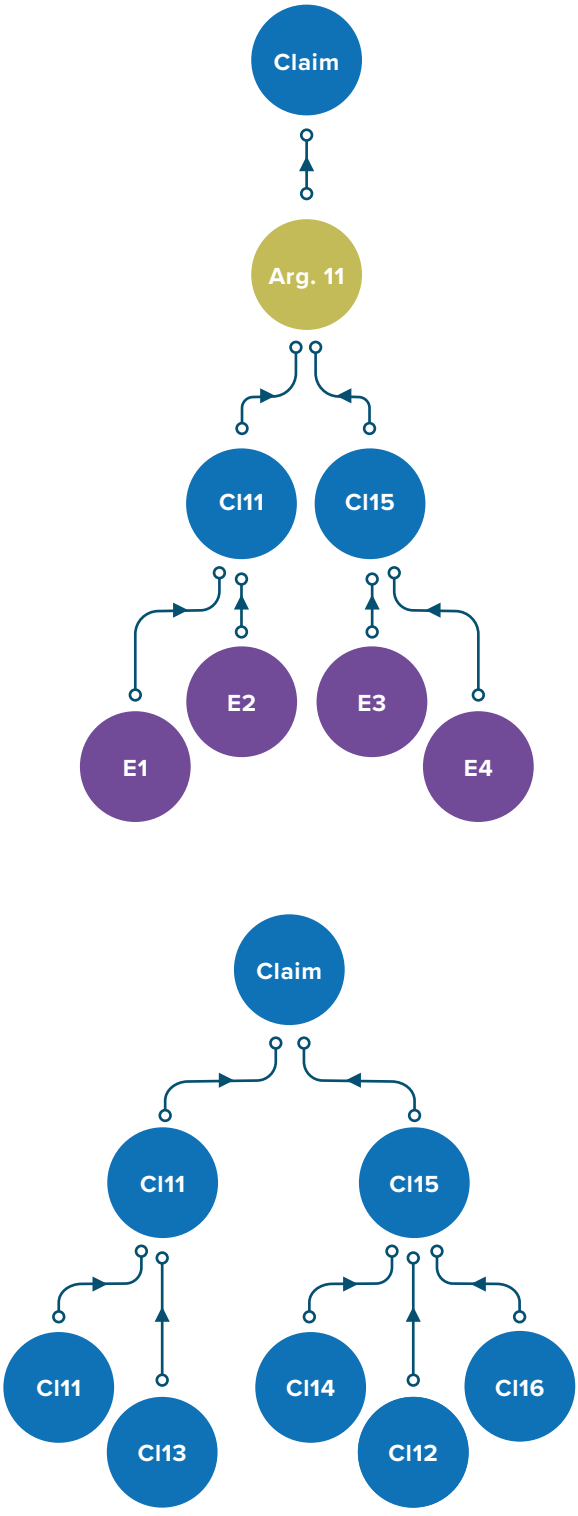

Figure 10: Options for summarising

Sometimes a mixed graphical and textual or tabular approach is useful. For example, the top set of claims could be defined in the CAE graphical notation as in the bottom of Figure 10, with the subclaims and evidence supporting the claims at the edges of the tree structure, and then expanded in tables within the supporting textual document.

05.
ACKNOWLEDGEMENTS

This document is based on material developed in earlier projects partially funded by the UK Control and Instrumentation Nuclear Industry Forum (CINIF) and guidance from previous CPNI projects and published research by Adelard.

06.
BIBLIOGRAPHY

[1] Bishop P.G., Bloomfield R.E. The SHIP Safety Case - A Combination of System and Software Methods, in SRSS95, Proc. 14th IFAC Conf. on Safety and Reliability of Software-based Systems, Brugge, Belgium, 12-15 September 1995; Bishop P.G, Bloomfield R.E., Van der Meulen M.J.P. “Public domain case study: An example application of the CEMSIS guidance”, v1.0, 26/3/2004, WP5 Deliverable, http://www.cemsis.org; and Bishop P.G., Bloomfield R.E. “A Methodology for Safety Case Development”, Safety-critical Systems Symposium 98, Birmingham, UK, Feb 1998, ISBN 3-540- 76189-6

[2] Bloomfield R. E., Netkachova K. “Building blocks for assurance cases”, 2014 IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW),3-6 Nov. 2014 doi: 10.1109/ISSREW.2014.72

[3] Robin Bloomfield and John Rushby. Assurance 2.0: A manifesto. In Mike Parsons and Mark Nicholson, editors, Systems and Covid-19: Proceedings of the 29th Safety-Critical Systems Symposium (SSS’21), pages 85–108, Safety- Critical Systems Club, York, UK, February 2021. Final draft available as arXiv:2004.10474; and John Rushby, Robin Bloomfield, Assessing Confidence with Assurance 2.0, arXiv:2205.04522, https://doi.org/10.48550/arXiv.2205.04522

[4] A simple logical semantic for the CAE notation comes from propositionalising the claims and the side-claim so: C11 /\ C12 /\ SC => C1

[5] In classical propositional logic one could move between a conjunctive and disjunctive normal form

Disclaimer

This guide has been prepared by CPNI and is intended to support the implementation of security-informed safety assurance and the Claims, Arguments and Evidence (CAE) methodology. This document is provided on an information basis only, and whilst CPNI has used all reasonable care in producing it, CPNI provides no warranty as to its accuracy or completeness.

To the fullest extent permitted by law, CPNI accepts no liability whatsoever for any expense, liability, loss, damage, claim or proceedings incurred or arising as a result of any error or omission in the report or arising from any person acting, refraining from acting, relying upon or otherwise using the [report]. You should make your own judgment with regard to the use of this document and seek independent professional advice on your particular circumstances.

No Endorsement

Reference to any specific commercial product, process or service by trade name, trademark, manufacturer or otherwise, does not constitute or imply its endorsement, recommendation or favour by CPNI. The views and opinions of authors expressed within this document shall not be used for advertising or product endorsement purposes.

CPNI
Centre for the Protection of National Infrastructure