Understanding Assurance Cases:
An Educational Presentation in Five Parts

Module 4: Creation

C. Michael Holloway
c.michael.holloway@nasa.gov

Senior Research Computer Engineer
Safety-Critical Avionics Systems Branch
NASA Langley Research Center, Hampton, Virginia, U.S.A.

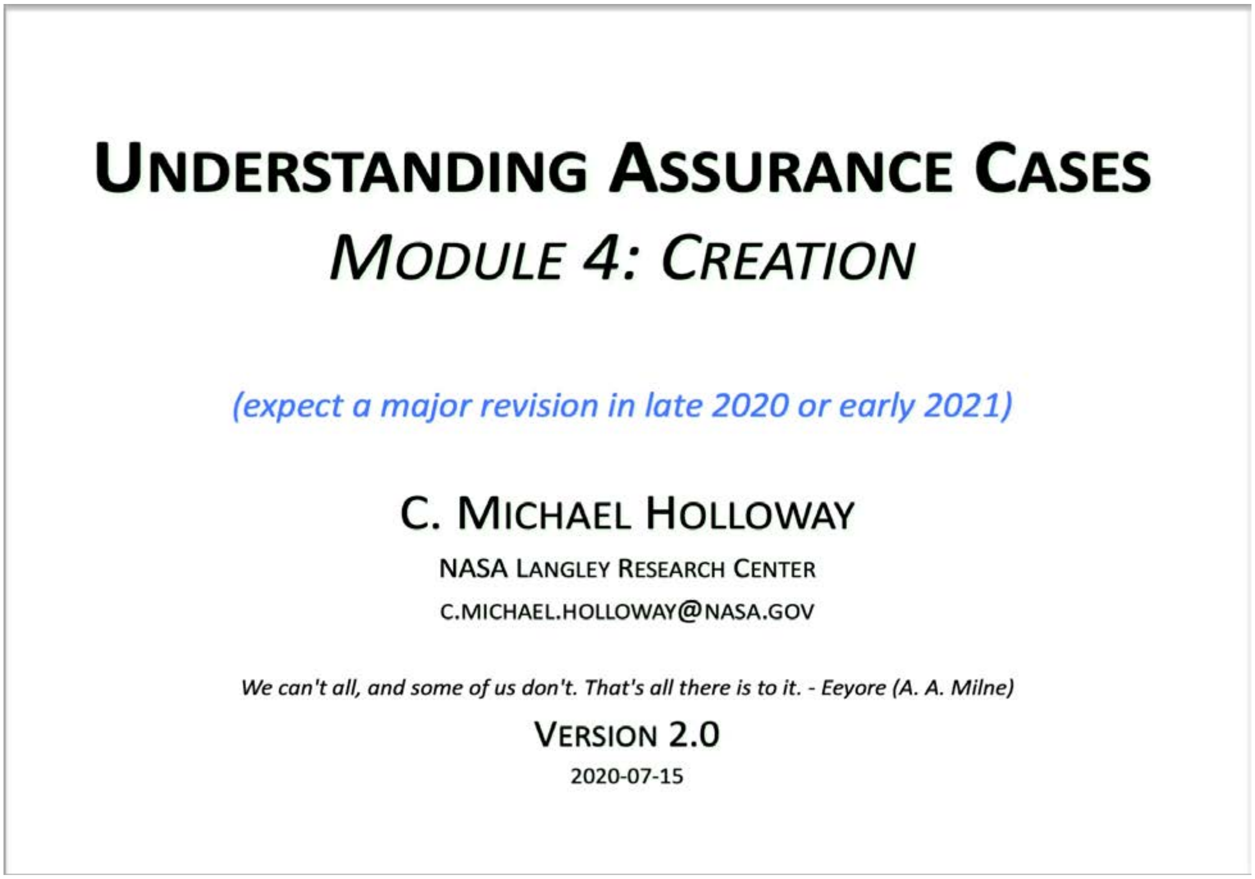

UNDERSTANDING ASSURANCE CASES

MODULE 4: CREATION

(expect a major revision in late 2020 or early 2021)

C. MICHAEL HOLLOWAY

NASA LANGLEY RESEARCH CENTER

C.MICHAEL.HOLLOWAY@NASA.GOV

We can't all, and some of us don't. That's all there is to it. - Eeyore (A. A. Milne)

VERSION 2.0

2020-07-15

This material was created in 2015-16, as part of the Explicate '78 project. The project was supported in substantial part by the Assurance Case Applicability to Digital Systems task under the reimbursable interagency agreement with the Federal Aviation Administration for Design, Verification, and Validation of Advanced Digital Airborne Systems Technology (IAI-1073 Annex 2 for NASA; DTFACT-10-X0008, Modification 0004 for the FAA). The original presentations were delivered to a selected group of FAA civil servants and NASA Langley personnel. The audio was recorded and partial transcripts (containing only the words spoken by the presenter, Mr. Holloway) produced. The intent from the beginning was to collect the material into a form that could be made available publicly. The text adheres closely to the original transcript, except where changes have been made to the original presentation since it was first given, as part of work for for NASA IA-303333/FAA IA NO 692M15-19-T-00029 Annex 1/TO 1. The full collection consists of six documents (including this one), which are available electronically through `https://shemesh.larc.nasa.gov/arg/uac.html`.

(I expect to make a major revision to this Module in late 2020 or early 2021)

Greetings.

Welcome to the fourth and penultimate module in an educational series about Understanding Assurance Cases. In this module, we will examine the ***Creation*** of assurance cases.

If you have not already completed Modules 1 - 3 (Foundation, Application, and Evaluation respectively), please stop reading this document, and complete, at least, Foundation and Application before continuing[1].

I'm quite sure that A. A. Milne *did not* have creating assurance cases in mind when he had Eeyore say "We can't all, and some of us don't. That's all there is to it." [Milne, A. A. 1928. *Winnie the Pooh*. London: Methuen & Co, Ltd.] But it's apt none-the-less. Creating cogent assurance cases is not something that everyone can do. Perhaps only a few of you will ever try to create a real case, but knowing a bit about what goes into such an endeavor may be useful for you nonetheless.

As with all the modules, feel free to interrupt me at *any* point if you have a burning question. I reserve the right to defer the answer to later on that's appropriate, but otherwise I'll do my best to answer it. As with the other modules, there will be times when I'll ask you questions, too. Like now.

[Question to participants: Does anyone have any questions or comments that you want to make before we proceed further?]

Let's list our learning objectives. By the time we're finished today, I hope that you'll be able to do at least these four things:

- Enumerate steps for creating a new assurance case.
- Explain essential questions that must be answered while developing a case.
- Identify common mistakes made in assurance case creation.
- Create a simple assurance case.

As I'm sure you realize, when we're done with this module, you're not going to be an expert in creating assurance cases (unless you're one already), but you should have a little better acquaintance with what's involved in creating them.

We're only going to be able to scratch the surface, But I will provide you with a homework exercise that, if you choose to do it, will help you scratch a bit deeper.

[1] Just in case someone does not follow the suggestion, and thus misses the preliminary information first expounded in Module 1 and repeated verbatim in Module 2, here is that information in simplified form: Within the assurance case community, intramural debates abound about a variety of topics we will discuss. Except in rare instances the existence of these debates is intentionally ignored or mentioned only briefly in this material. (See Module 1 or 2 for an explanation of why). Also, all images you see were either created by me (Michael Holloway) or are in the public domain via CC0 1.0 Universal.

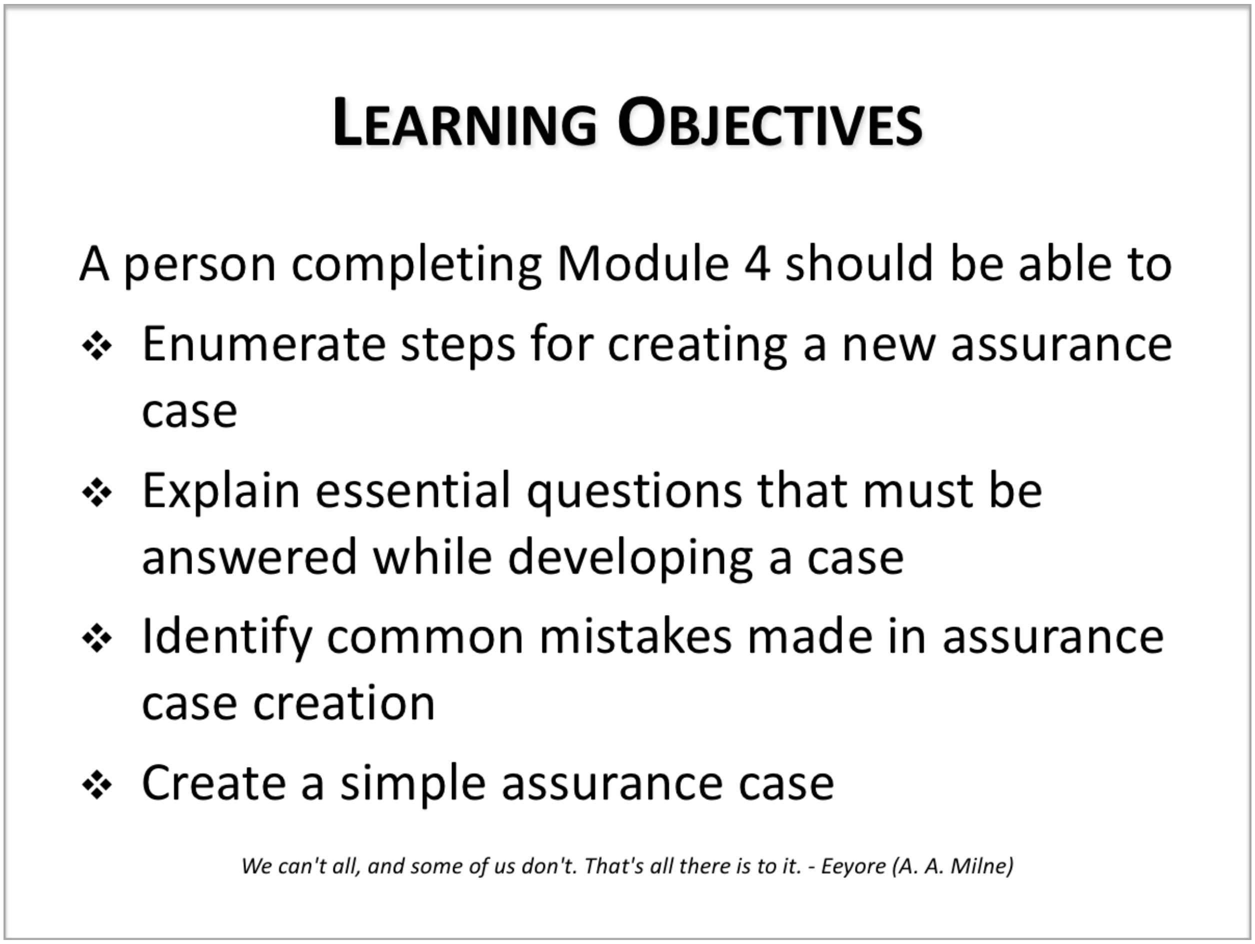

Learning Objectives

A person completing Module 4 should be able to

- Enumerate steps for creating a new assurance case
- Explain essential questions that must be answered while developing a case
- Identify common mistakes made in assurance case creation
- Create a simple assurance case

We can't all, and some of us don't. That's all there is to it. - Eeyore (A. A. Milne)

[Question to participants: Any questions about these learning objectives?]

As you probably expect, we begin with the continuing saga of Jon, Mike, and (the unseen) Tim.

When last we left our friends Mike had just told Jon, “Deciding if a case is good enough can be rather tough.”

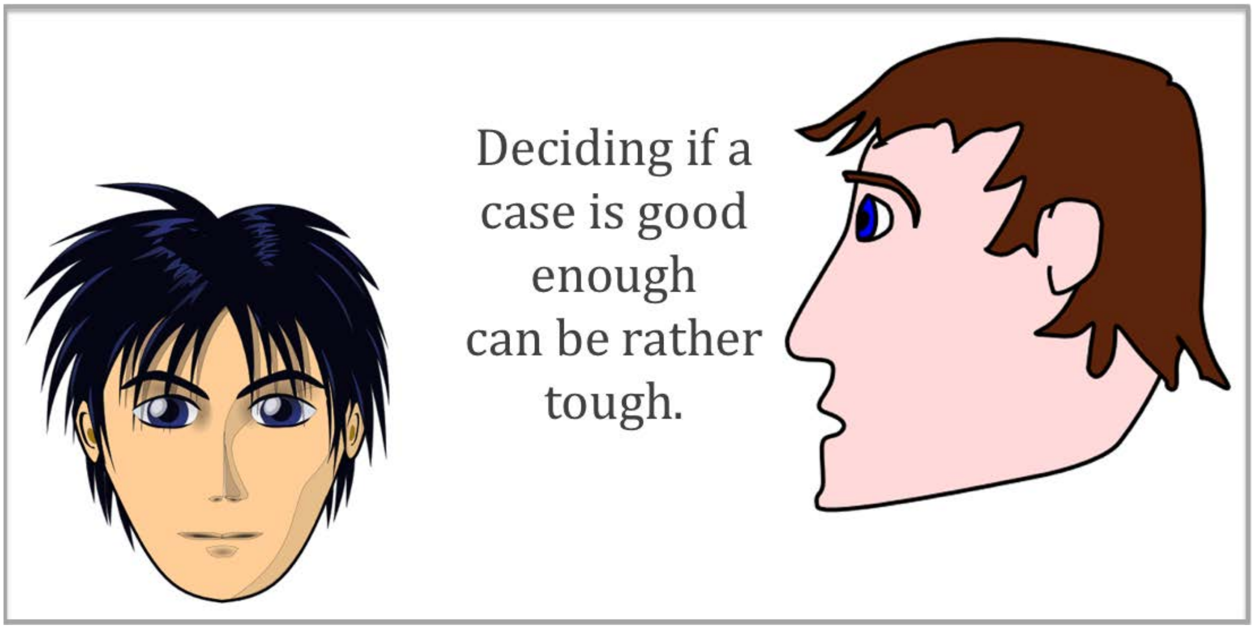

Jon thinks for a few seconds, then asks “How tough is it to create a case in the first place?”

“Hmmmm,” says Mike. “Good question. I guess it sorta depends.”

“It sorta depends on what?” inquires Jon.

“Lots of things,” says Mike, unhelpfully. But after a brief pause he continues, “… what the case is trying to show … what kind of evidence you have … who you’re trying to convince”

Jon interrupts his dad at this point: “I’m trying to convince you Dad, remember?”

“That you are my son …”

Then after a pause, with a slight grin on his face, Mike continues, “’Tis probably best to just give up now.”

Jon, not seeing the grin on his dad’s face, exclaims with a slightly annoyed tone, “I don’t wanna give up! Tim’s my only hope for getting to the game!”

Mike, with a bigger grin on his face, replies, “No, there is another.”

“Huh?” asks Jon, failing to recognize the reference.

“Never mind. I was far, far away for a second,” says Mike, continuing his excursion into the Star Wars universe[2].

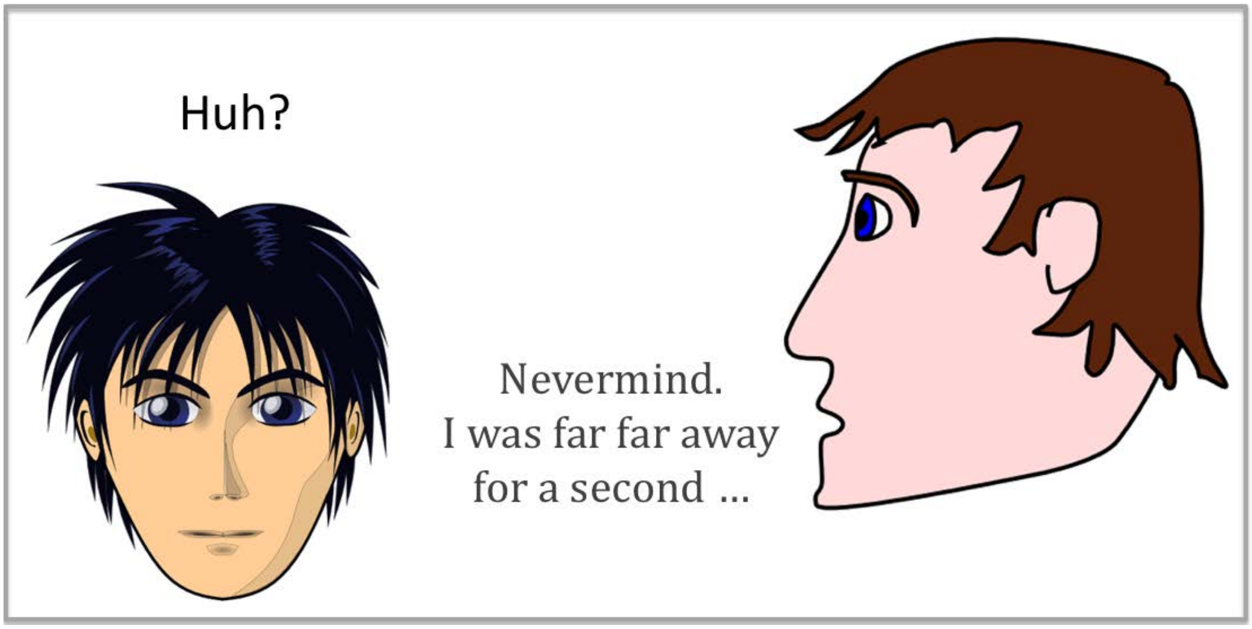

[2] One of the major supporters of the work that led to these five educational modules was Mike DeWalt. At the time he was the FAA’s Chief Scientist and Technical Advisor (CSTA) for Aircraft Computer Software and a huge fan of the Star Wars saga. This reference was for him. Mike is no longer a CSTA, having retired in 2016. He is, presumably still a Star Wars fan.

After seeing no hint of recognition on Jon's face, he replies, "I'll show you how to create a case to convince me."

"Thanks Dad!" replies Jon happily.

Despite what some of you may think, I'm not Mike, as my mother reminded anyone who tried to call me that when I was growing up, but I am going to have a go at explaining a bit about creating assurance cases.

Because it may have been a while since some of you completed the last module, I think it's probably a good idea to briefly review argument terms.

You see here a slide that we first saw in Module 1. (Changes will be made here soon.)

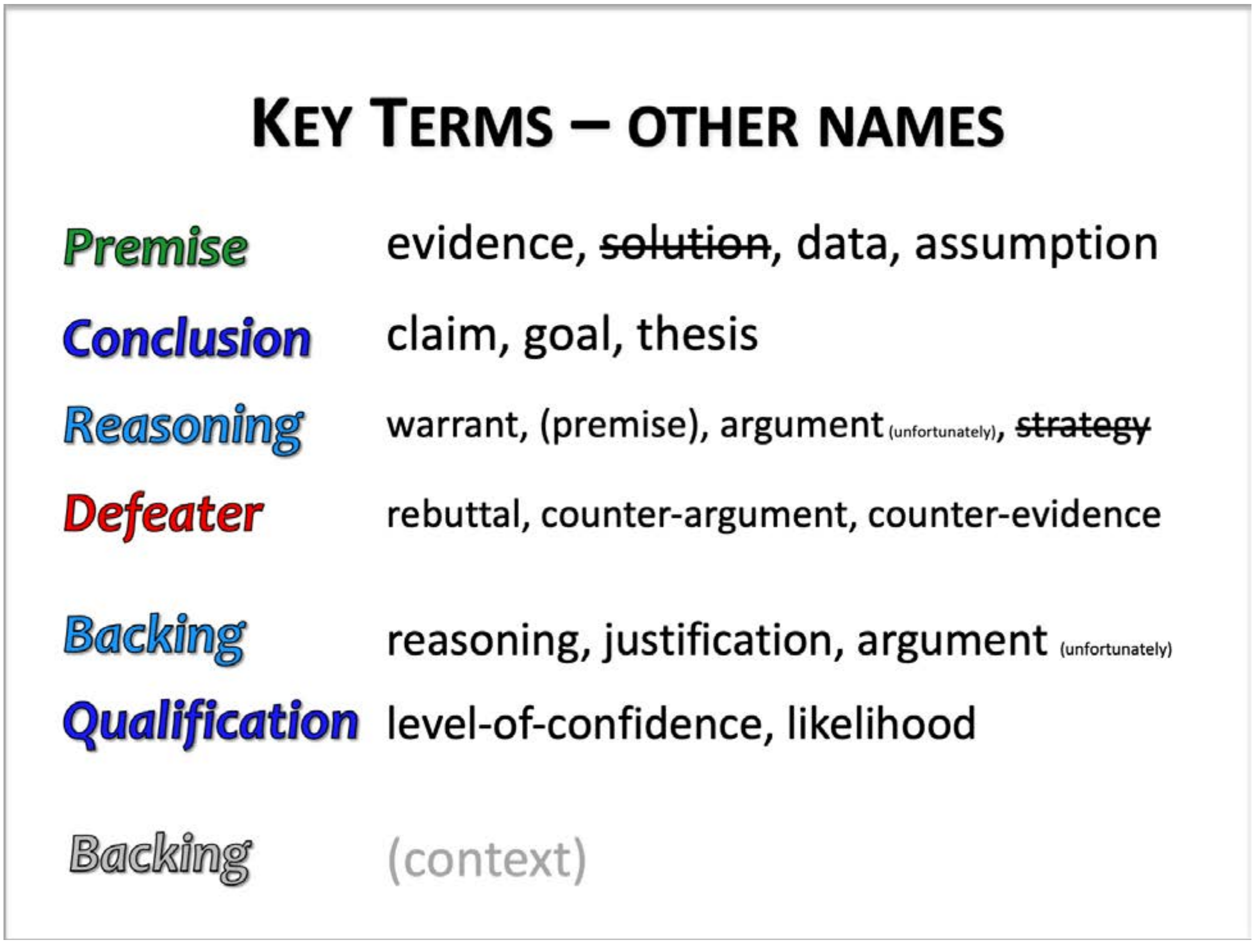

KEY TERMS – OTHER NAMES

Premise	evidence, ~~solution~~, data, assumption
Conclusion	claim, goal, thesis
Reasoning	warrant, (premise), argument (unfortunately), ~~strategy~~
Defeater	rebuttal, counter-argument, counter-evidence
Backing	reasoning, justification, argument (unfortunately)
Qualification	level-of-confidence, likelihood
Backing	(context)

On the left side are the terms that we're using in this course: *premise, conclusion, reasoning, defeater, backing (*incorporated in reasoning*)*, *qualification*, and *binding*.

The right side lists some popular alternative terms.

As I've mentioned before, within the assurance case community, the most common terms tend to be *evidence* (instead of *premise*), *claim* or *goal* (instead of *conclusion*) and *argument* (instead of *reasoning*).

I've explained before why I prefer our terms to those, and won't got back over my arguments, unless someone asks me to do so[3].

[Question to participants: Any questions about terms?]

On to talking specifically about assurance case creation.

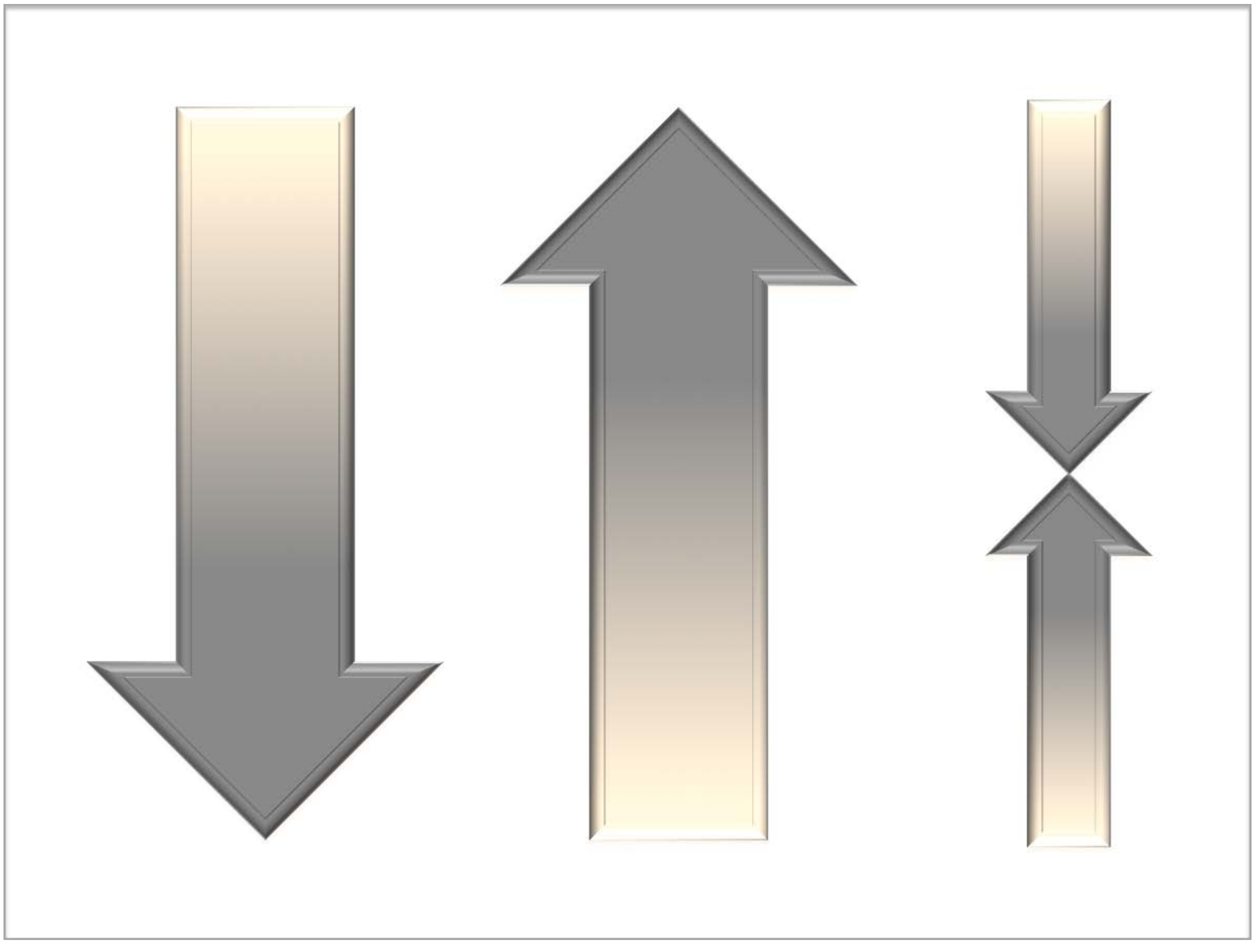

As you might imagine, in creating an assurance case, one might choose to proceed from the top down, or from the bottom up, or (as is most common) use a combination of the two. For pedagogical purposes, looking at idealized versions of a top down approach and a bottom up approach seems the most helpful. We will start with a top down approach.

In his doctoral thesis in 1998, Tim Kelly from the University of York proposed a six step method for creating safety cases using the Goal Structuring Notation. This slide, derived from a figure in the GSN Community Standard, illustrates that method, using the GSN terminology.

[3] Folks who are reading the material instead of seeing it being presented may look to pages 21-22 in Module 1.

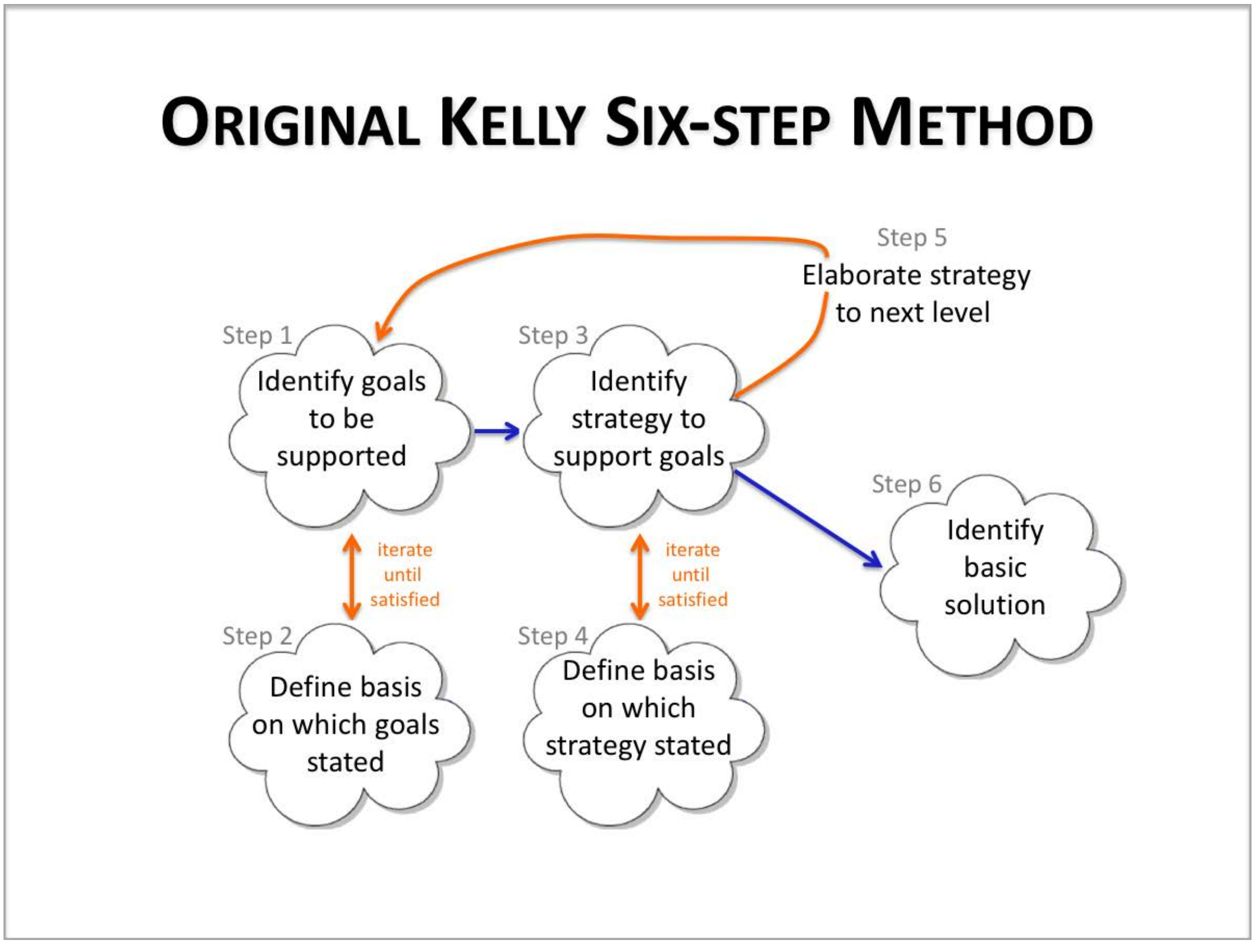

ORIGINAL KELLY SIX-STEP METHOD

Step 1
Identify goals to be supported

Step 2
Define basis on which goals stated

Step 3
Identify strategy to support goals

Step 4
Define basis on which strategy stated

Step 5
Elaborate strategy to next level

Step 6
Identify basic solution

iterate until satisfied

iterate until satisfied

In the years since 1998, other top down approaches have been proposed. But most of them are really nothing more than variations on the six step method, and no evidence has been produced to suggest any of the variations are definitely better, so, we'll follow this approach, 'though rewording it to correspond to the terminology that I prefer.

Step 1: Identify conclusions to be supported.

Step 2: Define basis on which conclusions stated.

Step 3: Identify reasoning to justify conclusions.

Step 4: Define basis on which reasoning stated.

Step 5: Elaborate argument to next level.

Step 6: Identify grounded premises.

Here is the figure modified with the different (aka better) terminology.

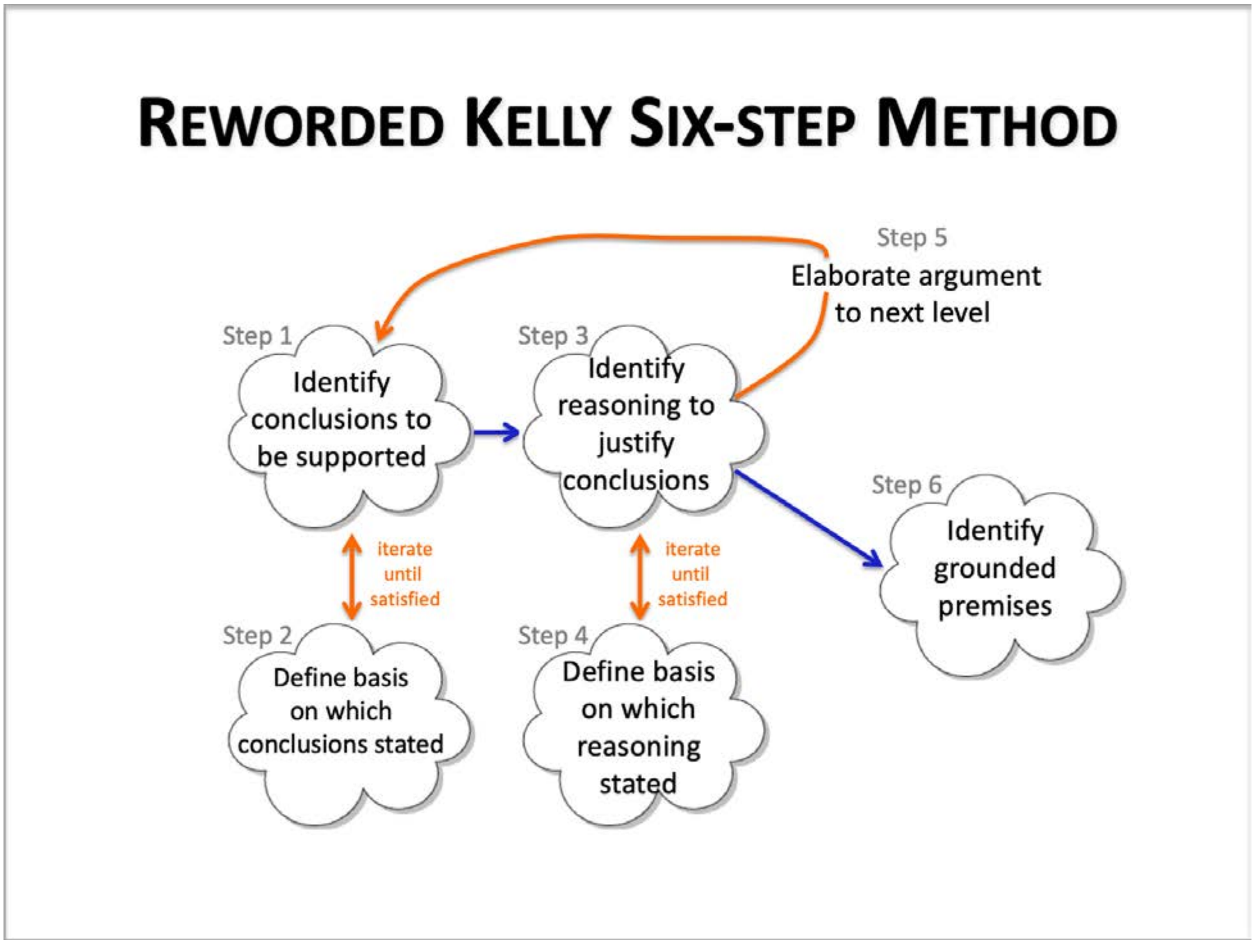

Reworded Kelly Six-step Method

Step 1
Identify conclusions to be supported

Step 2
Define basis on which conclusions stated

Step 3
Identify reasoning to justify conclusions

Step 4
Define basis on which reasoning stated

Step 5
Elaborate argument to next level

Step 6
Identify grounded premises

iterate until satisfied

iterate until satisfied

Let's see what each of these steps means, and how they relate to one another by way of an example. Because Jon seems like such a decent kid, let's use his situation as the basis for the example.

Recall Jon wants Tim to take him to a game. Jon's dad, Mike, doesn't know Tim, and wants assurance that Tim is a safe driver. He's asked Jon & Tim to build an assurance case.

What do you think an appropriate top-level *conclusion* (or goal or claim if you must) is for such a case?

Please do not turn the page until you have an answer to the question.

I suggest the following: "Tim is a safe enough driver to take Jon to the game."

That's step one: identifying the *conclusion* to be supported.

Perhaps some of you may see some problems (or at least ambiguities) with this statement as the *conclusion*. Handling such problems is the purpose of the next step.

For step 2, we need to define the basis on which the *conclusion* is stated. Or, in other, perhaps slightly clearer words, we need to decide if there's additional information we need to know in order for our statement of the *conclusion* to make sense.

Any ideas?

Some questions to ask yourself as you formulate your own ideas:

- Are there any words or phrases for which definitions are needed?
- Are any unstated assumptions seemingly present?
- After adding definitions and assumptions, are any changes to the original statement necessary to ensure it is unambiguous?
- And what about Naomi? [4]

A hint: the answer to each of the first three questions is, "Yes."

Another hint: While thinking about what changes to the original statement may be necessary to ensure it is unambiguous, complete the following quotation from President Kennedy's announcement of the goal of going to the moon:

"I believe that this nation should commit itself to achieving the goal, before this decade is out, of landing a man on the moon …."

Please do not turn the page until you have your own ideas.

[4] Folks who are reading the material instead of seeing it being presented, and who are confused by this question should refer to page 6 of Module 1.

Here are my answers.

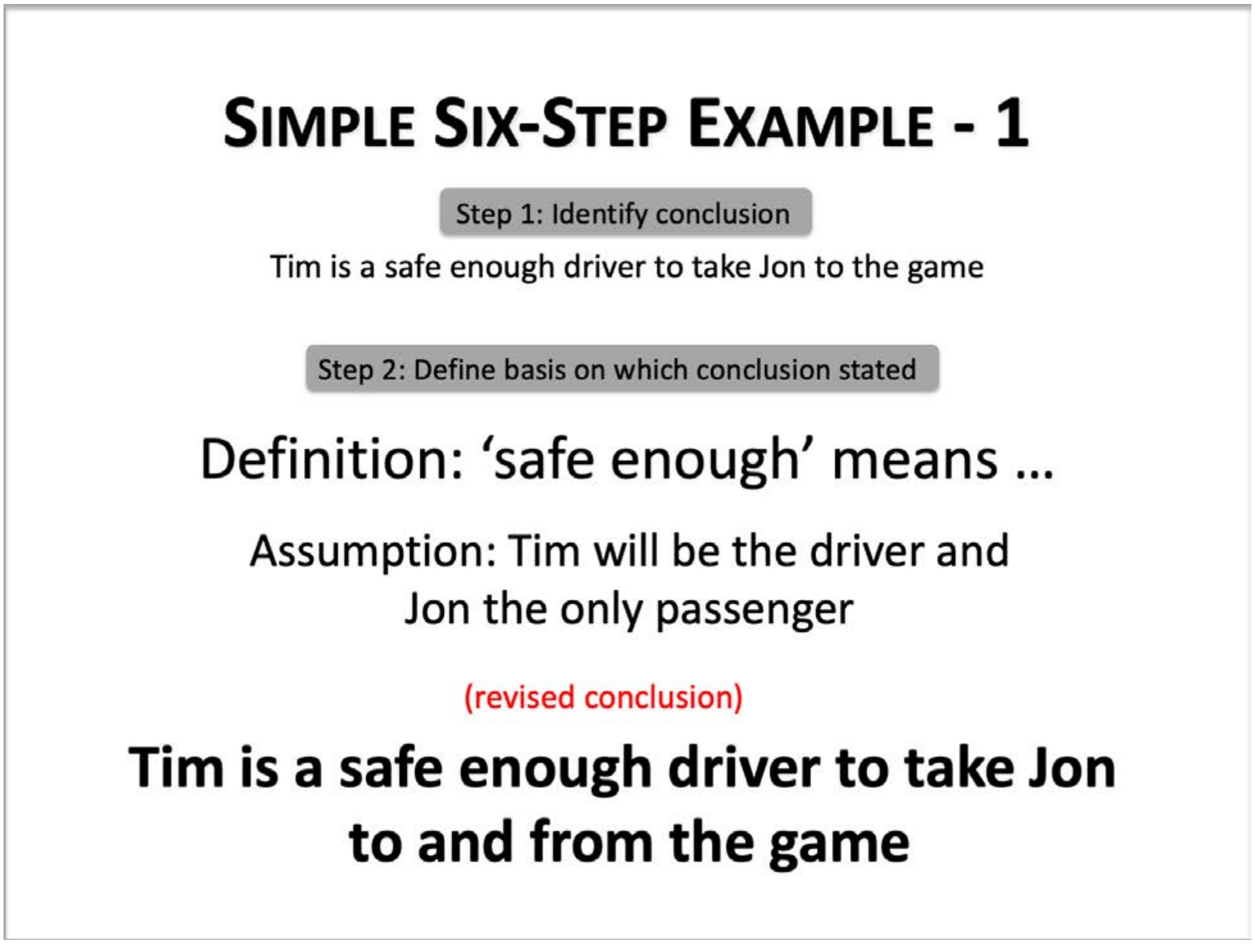

Simple Six-Step Example - 1

Step 1: Identify conclusion

Tim is a safe enough driver to take Jon to the game

Step 2: Define basis on which conclusion stated

Definition: ‘safe enough’ means …

Assumption: Tim will be the driver and Jon the only passenger

(revised conclusion)

Tim is a safe enough driver to take Jon to and from the game

One thing we certainly need is to know is the meaning of the phrase ‘safe enough’.

[Question to participants: What do you think might be an appropriate definition?]

There are a variety of options, but perhaps “posing no greater risk to Jon than Mike would …” would be a good one.

It seems to me that we might need to make at least one assumption, something along the lines of “Tim will be the driver and Jon the only passenger.”

In writing the original conclusion, I was thinking of ‘to the game’ as being equivalent to ‘to and from the game’; meaning it isn’t okay for Tim to just get Jon safely to the game, but he also needs to get Jon back home afterwards. To avoid possible ambiguity, perhaps the conclusion ought to be as “Tim is a safe enough driver to take Jon to and from the game.”[5]

[5] The rest of JFK’s statement was, “… and returning him safely to the Earth.” Landing wasn’t enough; returning safely to Earth was equally as important.

We have now completed steps 1 and 2 for our simple example.

[Question to participants: What questions or comments do you have at this point?]

Returning to the graphical illustration of the method, you see that steps 3 & 4 are similar to steps 1 & 2 but applied to the reasoning instead of to the conclusion. Step 5 involves elaborating the argument to identify premises for the top level conclusion, which will likely be conclusions that need to be supported themselves.

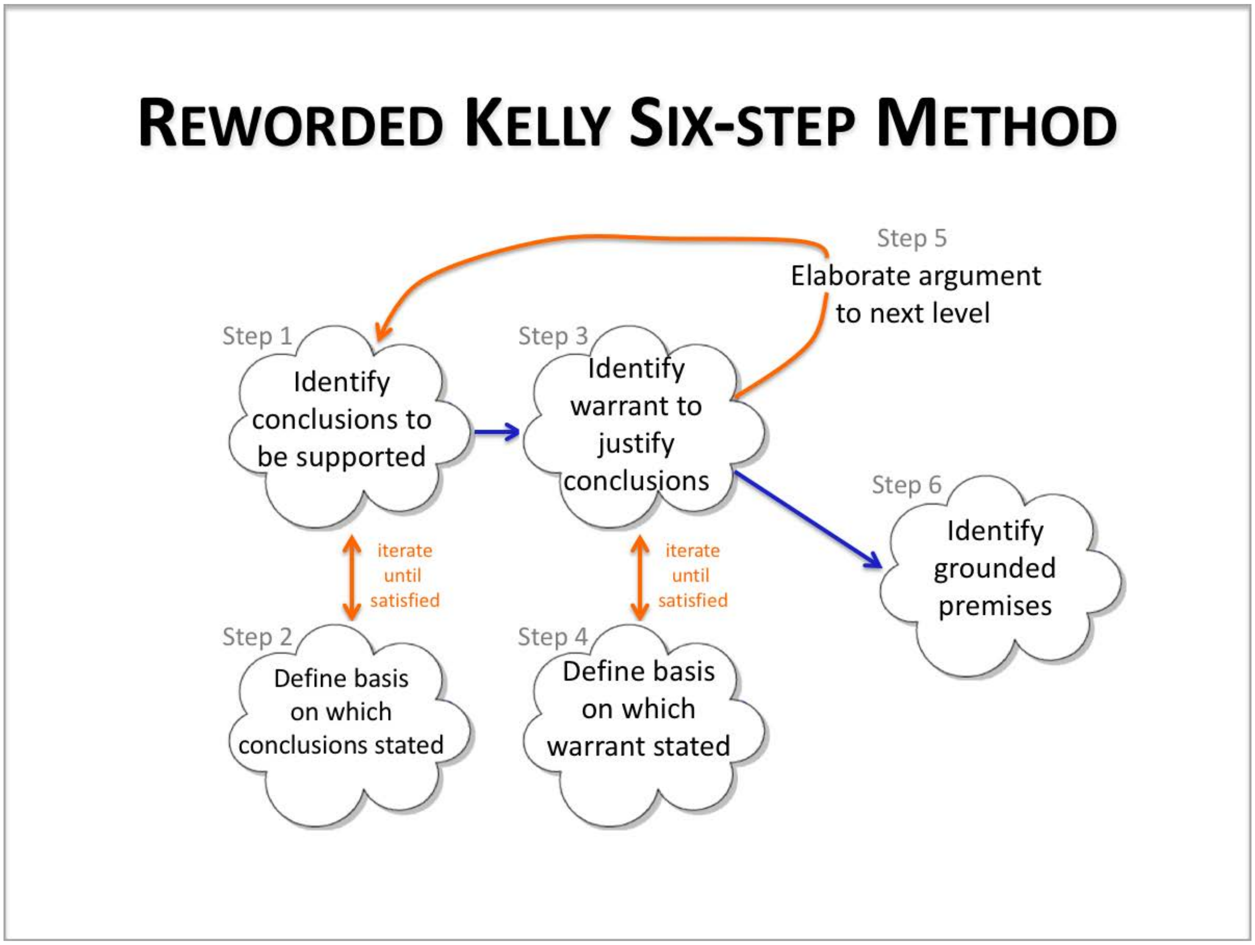

REWORDED KELLY SIX-STEP METHOD

Step 1
Identify conclusions to be supported

Step 2
Define basis on which conclusions stated

iterate until satisfied

Step 3
Identify warrant to justify conclusions

Step 4
Define basis on which warrant stated

iterate until satisfied

Step 5
Elaborate argument to next level

Step 6
Identify grounded premises

So, what we want to do next is think about the sort of reasoning that we'd want to use to establish the conclusion that "Tim is a safe enough driver to take Jon to and from the game."

[Question to participants: Does anyone want to suggest possible reasoning?]

If you're having trouble thinking of the reasoning, try instead to think about the sorts of premises that you think you'd want to see for the conclusion (skipping mentally to Step 5). Then think about the reason those premises would give you confidence in the conclusion.

The 6-step method isn't intended to be a straightjacket that restricts your thinking into a strictly sequential order. It is really just a guideline to help prompt your thinking. Often considering Steps 3, 4, & 5 together may be the most useful approach to creating a case.

There plenty of different possibilities for plausible and sufficient reasoning. For the purposes of continuing the example, I will suggest something mundane.

Reasoning: “Four independent indicators of driver safety suffice.”

[Question to participants: What do we need to know for this reasoning to make sense?]

Well, at the very least we’d need to have a common understanding of what constitutes an ‘independent indicator’. For the purposes of the example, let us assume that we have completed Steps 3 and 4.

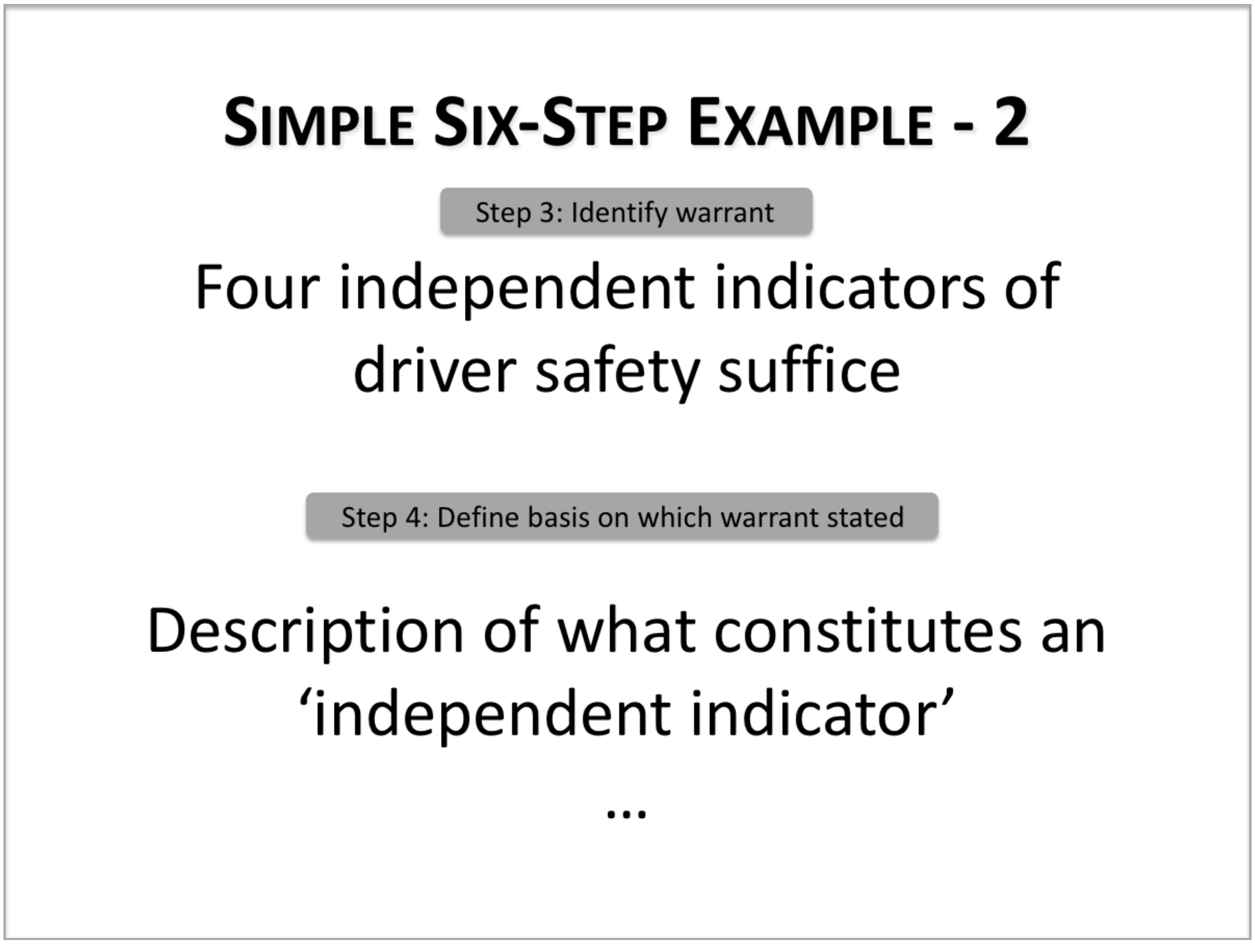

Simple Six-Step Example - 2

Step 3: Identify warrant

Four independent indicators of driver safety suffice

Step 4: Define basis on which warrant stated

Description of what constitutes an ‘independent indicator’

...

Let’s proceed to elaborating the argument (Step 5). We will do so by considering what might constitute the collection of acceptable independent indicators of driver safety.

[Question to participants: Are you able to name some indicators?]

Please do not turn the page until you have thought of at least one.

Here are the four that I decided to write down:

1. Tim has satisfied all legal requirements for driving.
2. Tim has not been in an accident.
3. Tim has a reputation for driving safely.
4. Nothing is going on in Tim's life that might cause him to drive less safely than usual.

Of course, many more plausible possibilities exist, but this slide expresses what we've just discussed in FAN.

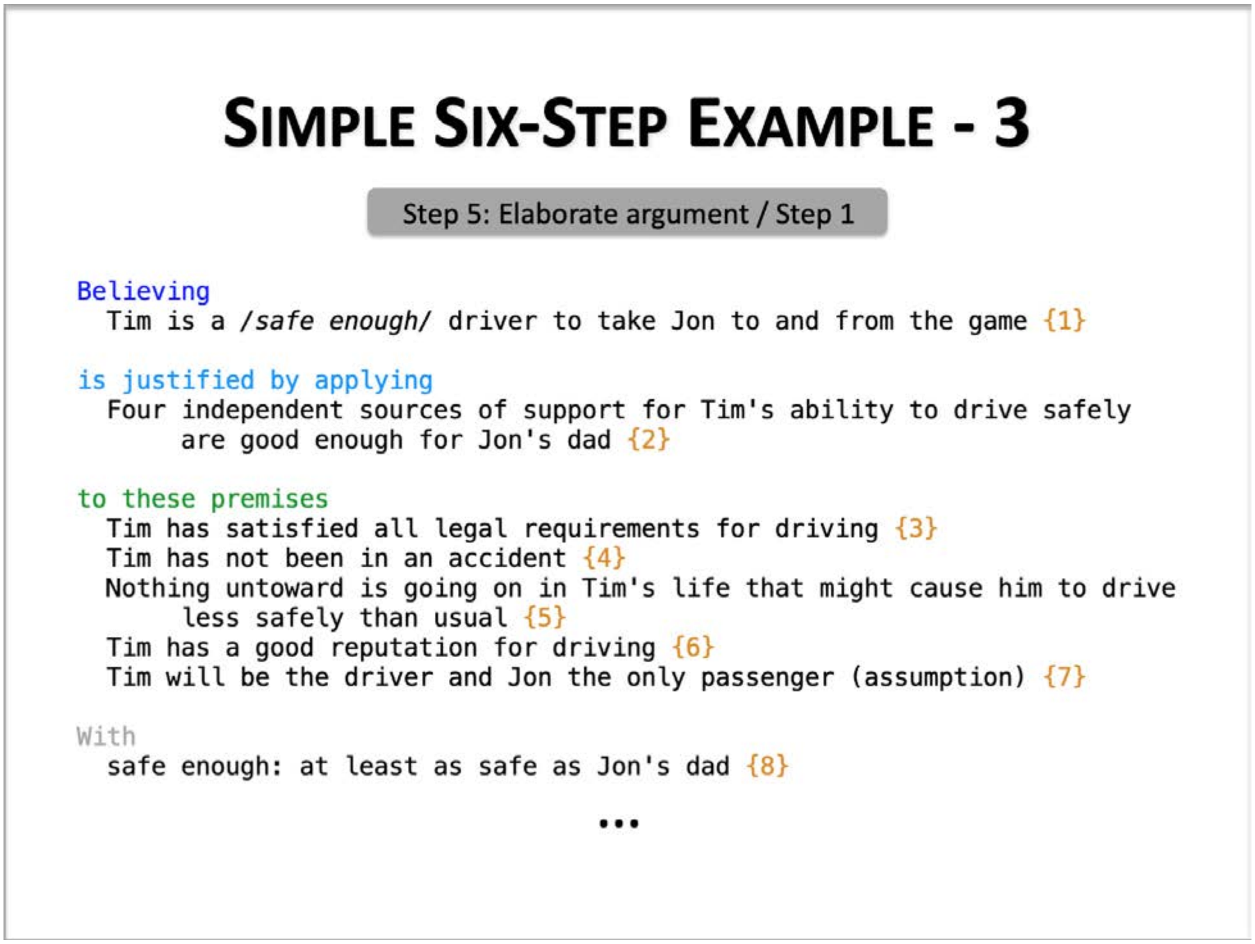

SIMPLE SIX-STEP EXAMPLE - 3

Step 5: Elaborate argument / Step 1

```
Believing
  Tim is a /safe enough/ driver to take Jon to and from the game {1}

is justified by applying
  Four independent sources of support for Tim's ability to drive safely
       are good enough for Jon's dad {2}

to these premises
  Tim has satisfied all legal requirements for driving {3}
  Tim has not been in an accident {4}
  Nothing untoward is going on in Tim's life that might cause him to drive
       less safely than usual {5}
  Tim has a good reputation for driving {6}
  Tim will be the driver and Jon the only passenger (assumption) {7}

With
  safe enough: at least as safe as Jon's dad {8}

                                  •••
```

[Question to participants: Does that make sense? What questions do you have?]

Let's now think about grounded premises (or evidence if you prefer) for only one of these: "Tim has not been in an accident."

What might be facts or data that establish that Tim has not been in an accident?

Please do not turn the page until you have thought of at least one.

Here are two possible grounded premises: “DMV records show no accidents,” and “Insurance records show no accidents”

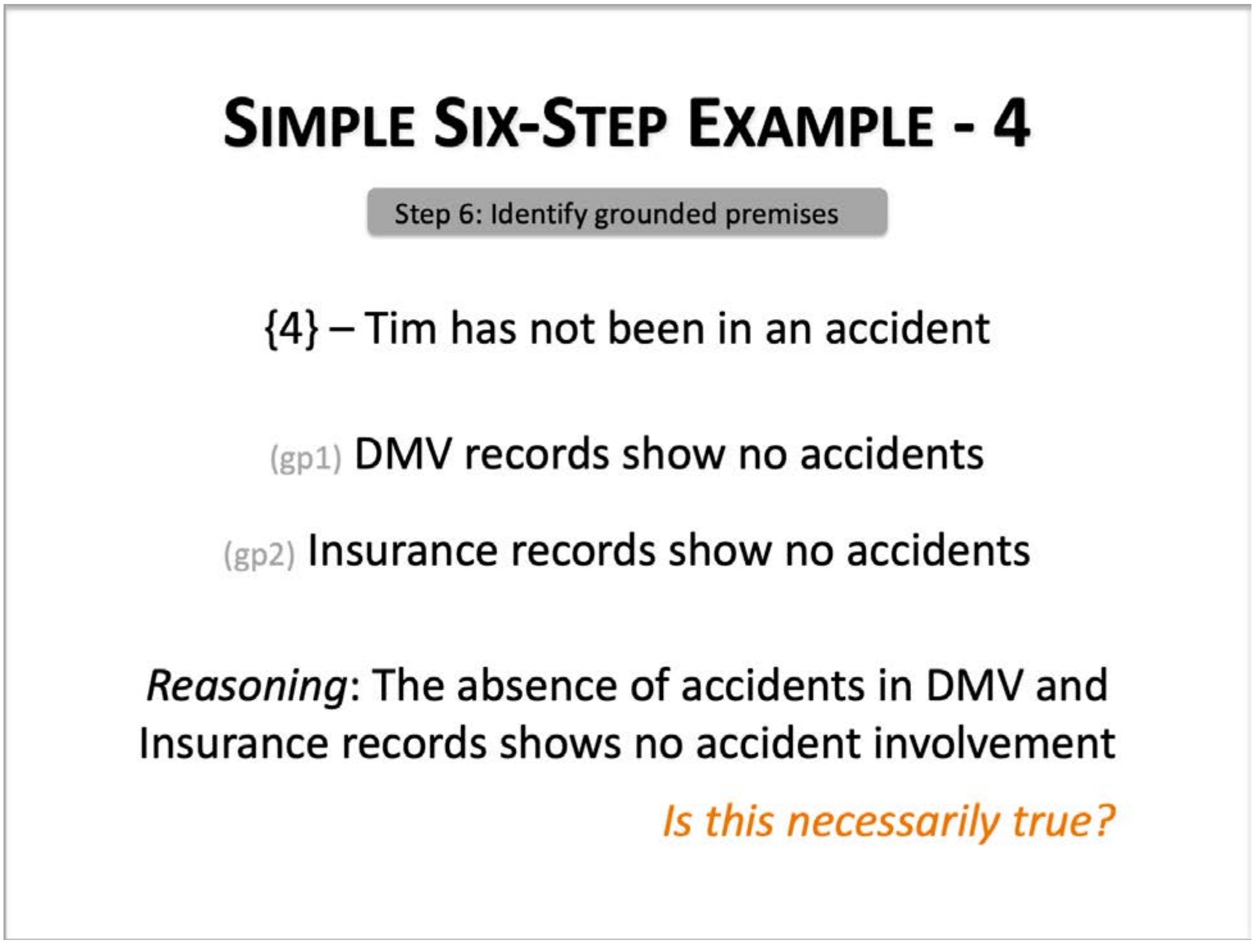

SIMPLE SIX-STEP EXAMPLE - 4

Step 6: Identify grounded premises

{4} – Tim has not been in an accident

(gp1) DMV records show no accidents

(gp2) Insurance records show no accidents

Reasoning: The absence of accidents in DMV and Insurance records shows no accident involvement

Is this necessarily true?

A reason why these two premises would be sufficient might be, “The absence of accidents in DMV and Insurance records shows no accident involvement.”

But is this *necessarily* true? Will it always be the case that the reasoning holds? That is, whenever DMV and Insurance records for a person contain no accidents, is it always true that the person has lived an accident-free driving life?

No … because the person, Tim in our example, could’ve had an unreported accident, or perhaps even several.

Some doubt will therefore exist as to whether we’ve fully established Tim’s accident-freedom. Hence, a reason we chose multiple independent indicators in the first place: no one of them alone provides sufficient confidence, but perhaps the combination of all four does justify the confidence. To complete the case, we’d continue in a similar fashion with each of the 3 other independent indicators, deciding what’s necessary to establish confidence that they are true. If we are unable to create an argument (or arguments) to provide sufficient confidence, then we will have to admit our efforts have failed to justify allowing Tim to take Jon to and from the game[6].

[6] I know several Tims. For one of those fellows, no convincing assurance case could ever be created for allowing one’s child in a car with that Tim behind the wheel.

Let’s look now at a primarily bottom-up method for creating assurance cases. It, too, was originally developed for GSN-style cases, but more recently than the method we just examined. I’ll skip showing you the version using GSN terminology[7], and move directly to one using our (better) terminology.

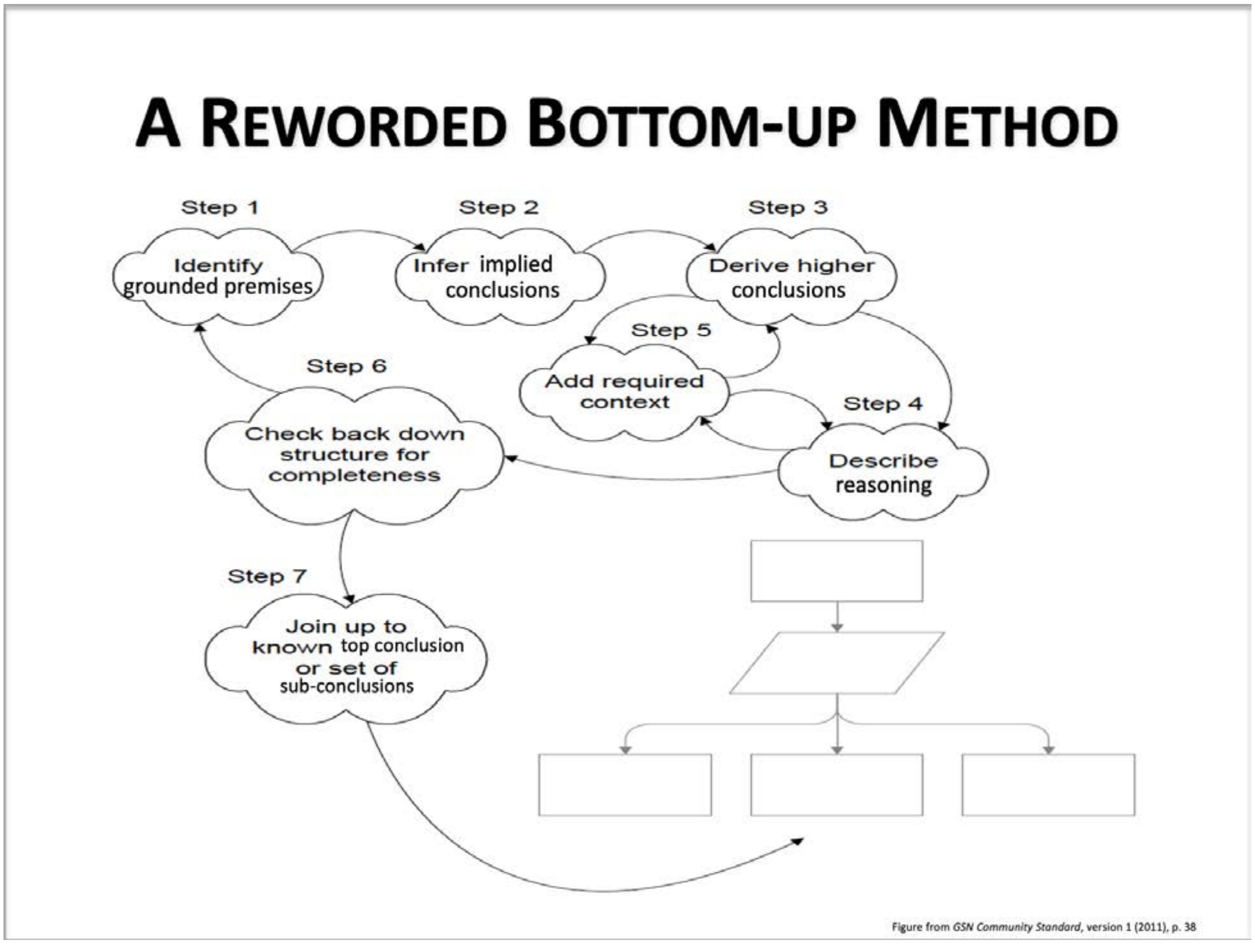
A REWORDED BOTTOM-UP METHOD

Step 1
Identify grounded premises
Step 2
Infer implied conclusions
Step 3
Derive higher conclusions
Step 4
Describe reasoning
Step 5
Add required context
Step 6
Check back down structure for completeness
Step 7
Join up to known top conclusion or set of sub-conclusions

Figure from *GSN Community Standard*, version 1 (2011), p. 38

You start with grounded premises, think about what they allow you to conclude, and why, and the needed context, and continue upwards. I’m not going to go through a full example, but let’s think about this approach a little bit.

Suppose we have these two facts:

- A Fault Tree Analysis showing the probability of a valve failing to close on demand is 1 x 10-4 / demand
- A requirement on the value to meet a probability of failure to close on demand of 1 x 10-3 / demand.

What’s a conclusion that we can infer?

Please do not turn the page until you have an answer.

[7] The figure you see here is based on a figure that first appeared in *GSN Community Standard*, version 1 (2011). p. 38. Since that time the GSN standard has been updated, but the figure illustrating the bottom-up style is unchanged. [Assurance Case Working Group. 2018. *Goal Structuring Notation Community Standard Version 2*. SCSC-141B. `https://scsc.uk/scsc-141B`]

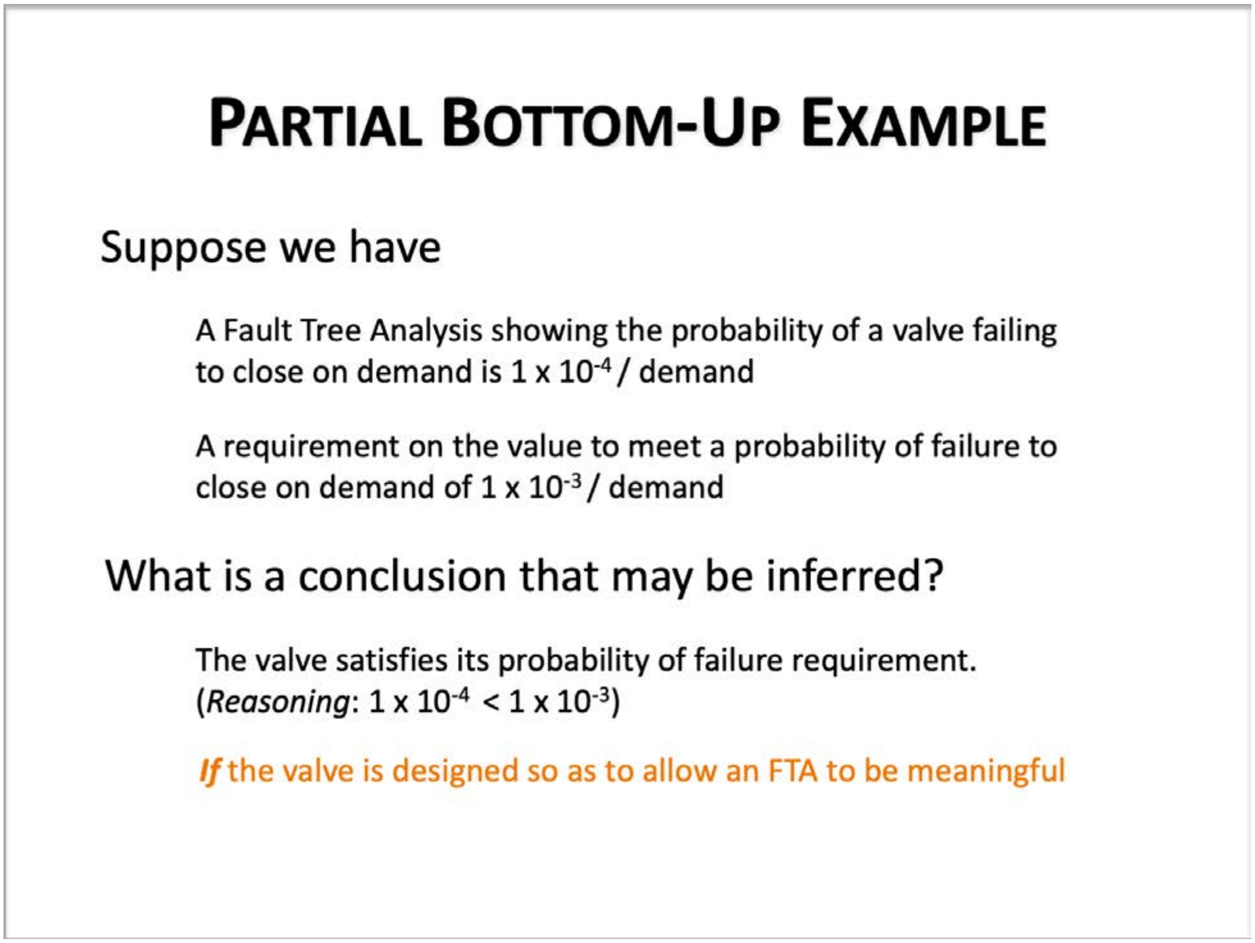

Partial Bottom-Up Example

Suppose we have

A Fault Tree Analysis showing the probability of a valve failing to close on demand is 1×10^{-4} / demand

A requirement on the value to meet a probability of failure to close on demand of 1×10^{-3} / demand

What is a conclusion that may be inferred?

The valve satisfies its probability of failure requirement. (*Reasoning*: $1 \times 10^{-4} < 1 \times 10^{-3}$)

If the valve is designed so as to allow an FTA to be meaningful

The valve satisfies its probability of failure requirement with the very simple reasoning: “1 x 10-4 < 1 x 10-3”.

But is this conclusion always justified in any circumstance, or are there conditions or context we need to consider?

At least one thing we need to consider is that the premises and reasoning justify confidence in the conclusion only “If the valve is designed so as to allow an FTA to be meaningful.”

If, however, the valve’s design includes aspects that make FTA untrustworthy (it contains software for example) then we can’t legitimately make the conclusion we suggested.

[Question to participants: Surely you have question and comments at this point. What are they? Note: in the original presentation, this Q&A part lasted for about 15 minutes. People who are reading this material are encouraged to send questions and comments to the author at `c.michael.holloway@nasa.gov`]

For those of you who are interested in seeing a much bigger example, consider taking a look at the Explicate ’78 work [full report: Holloway, C.M., Graydon, P.J. 2018. *Explicate '78: Assurance Case Applicability to Digital Systems*. DOT/FAA/TC-17/67. `https://go.usa.gov/xPEJr`. shorter version: Holloway, C.M. 2015. “Explicate '78: Uncovering the Implicit Assurance Case in DO-178C”. *Engineering Systems for Safety. Proceedings of the 23rd Safety-critical Systems Symposium*. M.

Parsons & T. Anderson (eds).] Although slightly different terminology was used for some terms in that report, you should by this time have no difficulty in translating to our better terminology.

Let's move on now to talking about some of the questions that a creator of an assurance case should be often asking her or his self. So, instead of FAQs, we'll be talking about QFAs.

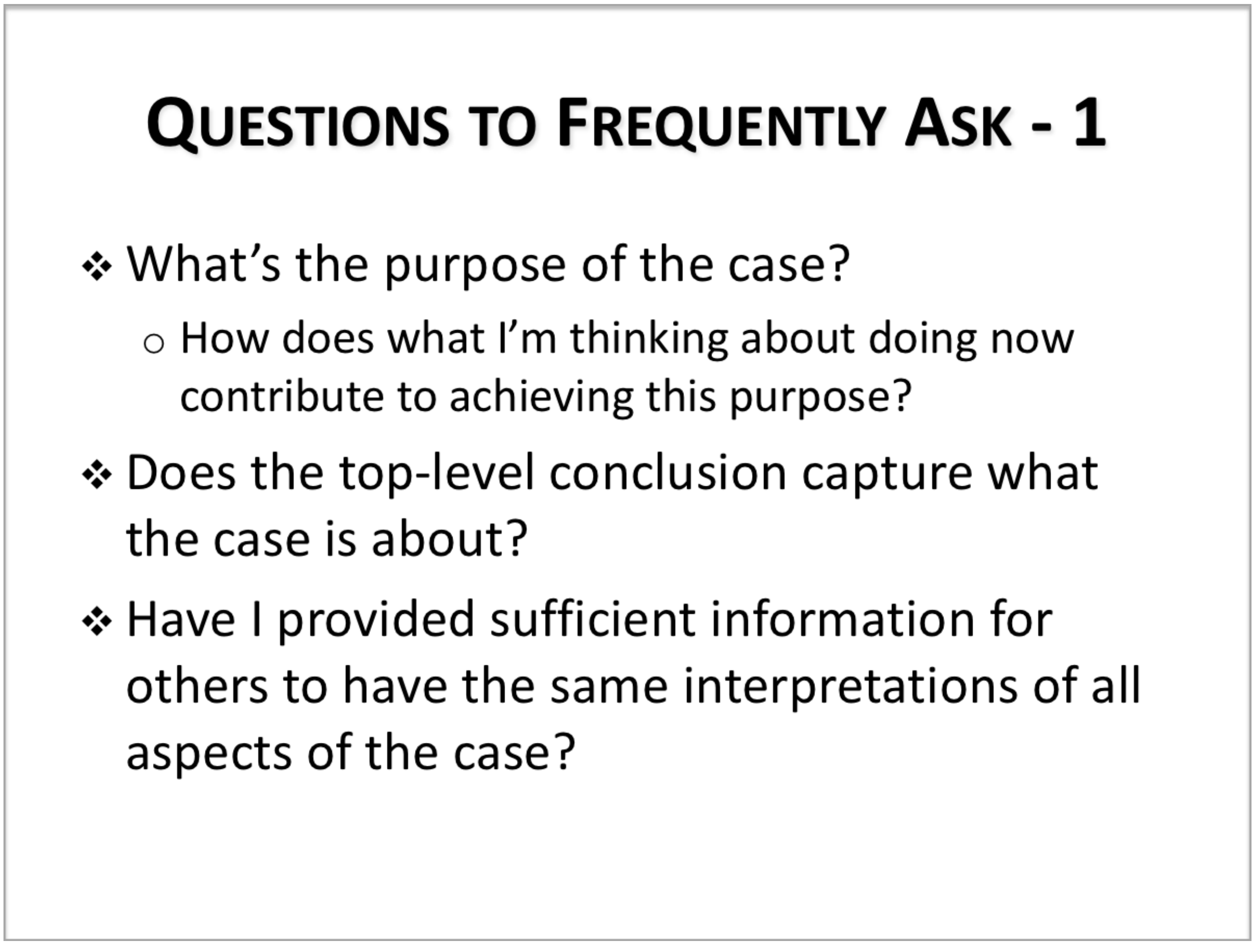

Questions to Frequently Ask - 1

- What's the purpose of the case?
    - How does what I'm thinking about doing now contribute to achieving this purpose?
- Does the top-level conclusion capture what the case is about?
- Have I provided sufficient information for others to have the same interpretations of all aspects of the case?

The first question you need to be frequently asking if you're creating an assurance case is, "What's the purpose of the case?" Also, ask yourself the associated question: "How does what I'm thinking about doing *now* contribute to achieving this purpose?" Your next steps may be different depending on the case's purpose.

Another important question is "Does the top-level conclusion capture what the case is about?" Suppose, for example, the top-level conclusion is solely about safety, but the case is supposed to provide justified confidence not only in safety, but also in achieving intended function; you need to modify the top-level conclusion.

An especially critical question to ask often is the last one shown on this slide: "Have I provided sufficient information for others to have the same interpretations of all aspects of the case?"

Recall my example from a few minutes ago: my use of "to the game" instead of "to and from the game" opened up an opportunity for differing interpretations by different people. Eliminating all such possibilities is not necessarily feasible (because some people insist on imagining impossible interpretations) but striving to eliminate *feasible* alternate interpretations is always the right thing to do.

A brief aside: If you're skeptical about my claim that some people imagine impossible interpretations, then I think a simple example will cause you to give up the skepticism.

DO-178C Chapter 1, section 4, item d notes the "document describes activities for achieving" the objectives, but says explicitly: "The applicant may plan and, subject to the approval of the certification authority, adopt alternate activities to those described in this document." Despite the explicit words, there are some people who insist DO-178C requires that all the activities listed in it must be followed. The words do not allow such an interpretation, but some people imagine they do[8].

[Question to participants: Any questions about these QFA's before we move on to some more?]

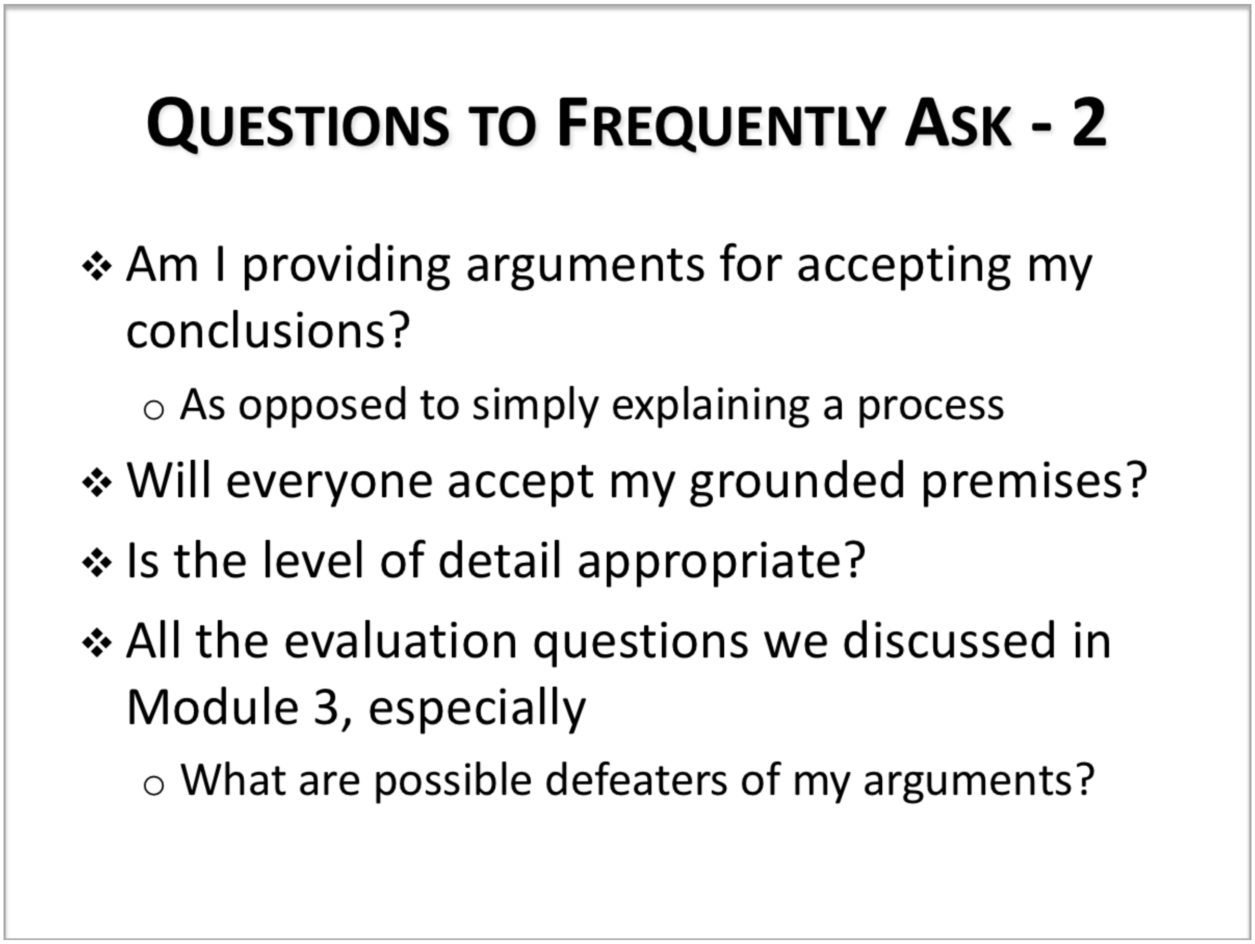

Questions to Frequently Ask - 2

- ❖ Am I providing arguments for accepting my conclusions?
  - o As opposed to simply explaining a process
- ❖ Will everyone accept my grounded premises?
- ❖ Is the level of detail appropriate?
- ❖ All the evaluation questions we discussed in Module 3, especially
  - o What are possible defeaters of my arguments?

[8] During the writing of the document, some of us anticipated the possibility that some people would be negatively imaginative when reading the sentence. I suggested quite strongly we should delete the words "subject to the approval of the certification authority." Because the qualification was (and still is) already implicitly applied to *every sentence* in the guidance, writing it out explicitly here was unnecessary. It was also dangerous, because would likely encourage those who wanted to be encouraged to think alternate activities were deprecated. Only handful of others supported my position. Thus, the words remained.

Another important question to keep in mind is this one: "Am I providing arguments for accepting my conclusions as opposed to simply explaining a process?"

It is not uncommon to see an assurance case written by a neophyte looking much more like a simple description of *what was* done than an argument about *why* doing those things is sufficient to establish the truth of the top-level conclusion to an acceptable level of confidence. Typically in such cases, reasoning is missing, or written too poorly to explain the reasoning[9].

Another important question is "Will everyone accept my grounded premises?" (Or if you prefer the term 'evidence': "Will everyone accept my evidence?")

Recall our quantified fault tree analysis example from earlier. If the analysis was applied to a subsystem or component for which obtaining real probability of failure numbers is possible, then citing the FTA results as a grounded premise is appropriate. Everyone should accept it.

But for other subsystems or components, for which probability numbers are fictitious (for example, a subsystem or component containing software), the FTA results should not be accepted. At least not without an additional argument justifying their acceptance for the particular subsystem or component in the case under consideration.

Do not forget: The assurance case argument structure must end with accepted grounded premises. If it does not, more argument is needed.

We've talked at several times during the course about this next question: "Is the level of detail appropriate?"[10]

We talked at length in Module 3 about other evaluation questions; all these constitute QFAs, particularly, but not only, the specific question, "What are possible defeaters of my arguments?" I won't go back over our fairly extensive discussion of defeaters, but will stop at this point for questions or comments about this section on questions to frequently ask.

[Question to participants: What questions do you have?]

Let's move now to talking about some common mistakes that happen when assurance cases are created. This discussion will mostly be a review of things we've talked about previously, both in earlier modules, and earlier in this module, so I'll go through these quickly, unless you have some questions.

[9] I tend to think that the GSN use of the term 'strategy' (and its associated typical instantiations) can inadvertently contribute to missing reasoning going undetected. My pro-GSN friends dispute this contention. Neither side has developed a compelling argument to convince the other side of the error of their ways.

[10] The question of appropriate detail is one of those questions about which opinions differ strongly within the safety/assurance case community. At one far end of the spectrum are folks who claim a good assurance case must address in deep detail every aspect of the system or service. At the other far end are people who claim that no assurance case should ever be more than 1-5 pages long. My own opinion lies closer to the small case side than the huge case side.

As you may suspect, many of the common mistakes are rooted in failing to ask the questions I enumerated just now.

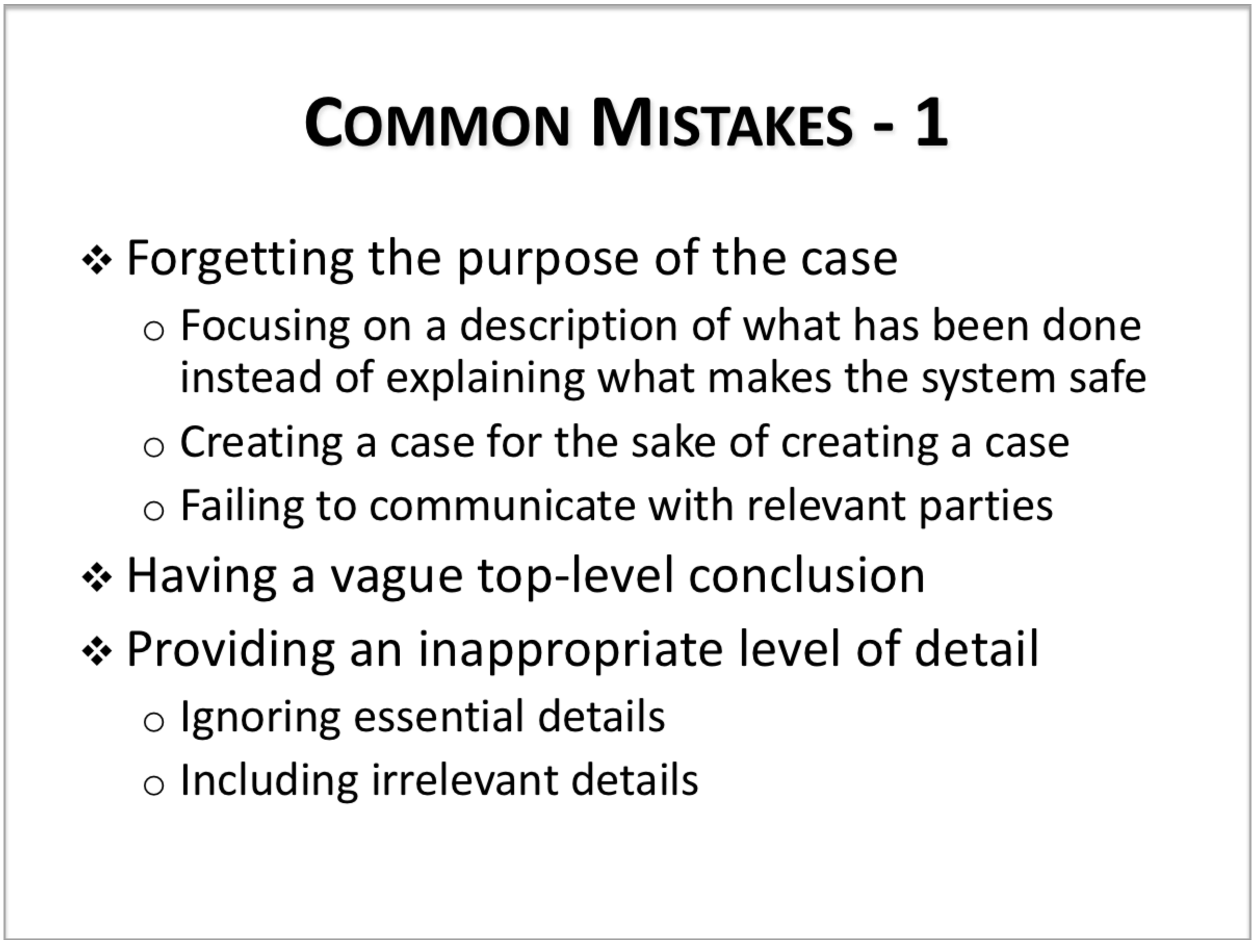

Common Mistakes - 1

- Forgetting the purpose of the case
  - Focusing on a description of what has been done instead of explaining what makes the system safe
  - Creating a case for the sake of creating a case
  - Failing to communicate with relevant parties
- Having a vague top-level conclusion
- Providing an inappropriate level of detail
  - Ignoring essential details
  - Including irrelevant details

Failing to ask about the purpose of the case easily results in making the mistake of forgetting the purpose of the case. This mistake may manifest itself in a number of ways, including the three you see listed here: focusing on a description of what has been done instead of explaining what makes the system safe; creating a case for the sake of creating a case; and failing to communicate with relevant parties.

Failing to question the top-level conclusion can result in having a vague (or otherwise deficient) top-level conclusion.

Not asking questions about detail frequently leads to providing an inappropriate level of detail, which can manifest in either direction: ignoring essential details, or including irrelevant details.

[Question to participants: Anyone have questions about these common mistakes before we move on to some more?]

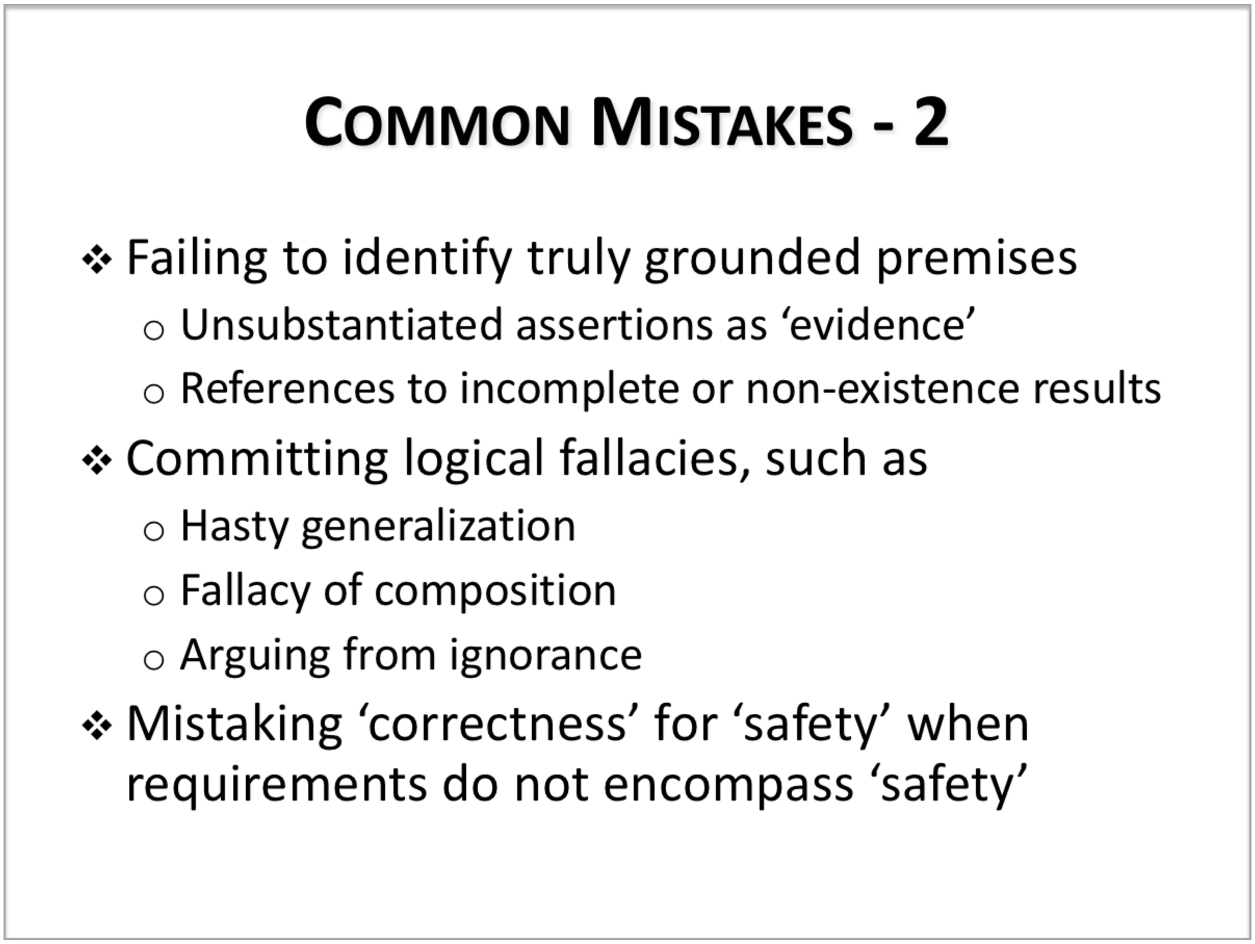

COMMON MISTAKES - 2

- ❖ Failing to identify truly grounded premises
  - o Unsubstantiated assertions as ‘evidence’
  - o References to incomplete or non-existence results
- ❖ Committing logical fallacies, such as
  - o Hasty generalization
  - o Fallacy of composition
  - o Arguing from ignorance
- ❖ Mistaking ‘correctness’ for ‘safety’ when requirements do not encompass ‘safety’

Another common mistake, well, really a category of mistakes, is (as you may have guessed) failing to identify truly grounded premises. This failure may manifest in several ways. Giving unsubstantiated assertions as ‘evidence’, which is what we just discussed a few minutes ago. There may also be references to incomplete or non-existence results. Perhaps the author of the assurance case expected certain tests to be conducted, and thus included the results of those tests as grounded premises in the argument, but in reality those tests were never conducted.

Another category of mistakes is committing logical fallacies in the argument. I’ve listed three such fallacies on the slide.

Hasty generalization refers (as its name suggests) to making a generalization from insufficient premises. One of the most common instantiations of it is generalizing from too few observations.

As an example, suppose you start looking at odd integers. You observe that 1 is a square number, 3 is a prime, 5 is a prime, 7 is a prime, 9 is a square, and 11 is prime. You conclude, “All odd numbers are either squares or primes.” If just looked at one more odd number, 13, you’d think your generalization still holds; but the next odd number, 15, refutes the generalization.

Fallacy of composition refers to inferring that a property that is true of a part is also true of the whole, without any other reasoning to establish the truth. This fallacy occurs in a safety case, for example, when the safety of individual subsystems is inferred to imply the safety of a whole system without also establishing the safety of interactions.

Arguing from ignorance is a name given to claims that something is true simply because it has not been proven false. “We ran lots of test cases and found no bugs; therefore, the software is necessarily bug-free” is a prototypical example.

A final mistake that may occur is to mistake ‘correctness’ for ‘safety’ when the requirements do not encompass ‘safety’. This mistake may be most likely to happen with software systems. If safety analysis is done in such a way that requirements are imposed on software to ensure safety (as is a fundamental assumption of DO-178C and its predecessors), then showing correctness does encompass ‘safety’. But in most other circumstances, ‘correctness’ and ‘safety’ are two different things. Conflating them is not a good thing.

That’s it for common mistakes. [Question to participants: Are there any questions or comments?]. [At this point in the original presentation I presented slide versions of a homework assignment. For this written version of Module 4, I will present the assignment at the end in straight text instead.]

At the beginning, I listed four things that I hoped you’d be able to do by the end of this module. Here are those four things recast in the form of questions. Think to yourself how you’d answer these questions.

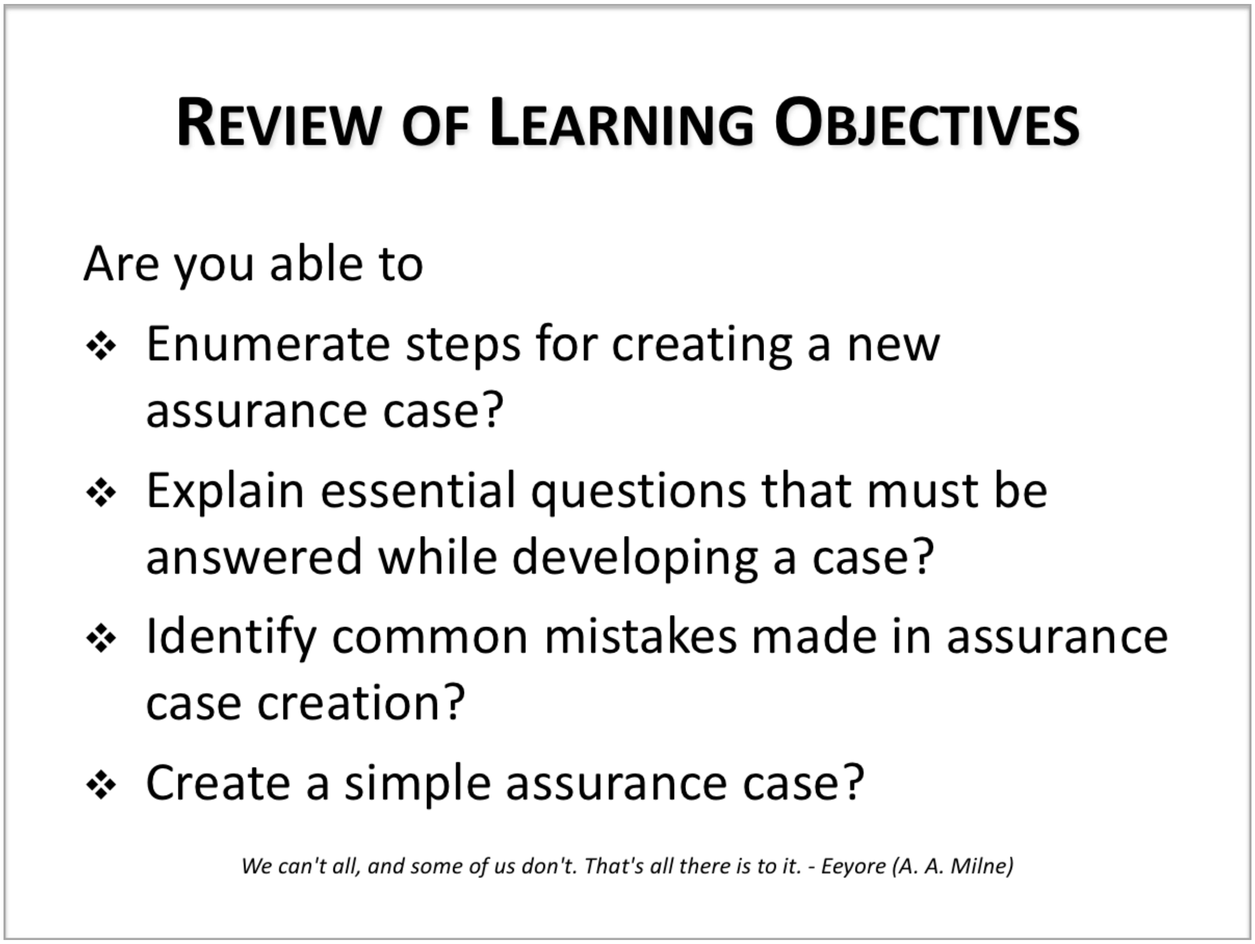

REVIEW OF LEARNING OBJECTIVES

Are you able to

- Enumerate steps for creating a new assurance case?
- Explain essential questions that must be answered while developing a case?
- Identify common mistakes made in assurance case creation?
- Create a simple assurance case?

We can't all, and some of us don't. That's all there is to it. - Eeyore (A. A. Milne)

After you’ve thought about the questions for a little bit, please ask me any questions that you still have.

For those of you who want to conduct a case study about how well you have learned the material in Module 4, here is an assignment developed by Mallory Graydon.

Jill Smyth wishes to operate her ultralight aircraft from a backyard aerodrome. Refueling this aircraft has hazards, including the potential for fire. Construct an operational safety argument illustrating why it is adequately safe for Jill to refuel her aircraft as planned.

You may either assume that an assessment of the hazards of the refueling operation has been completed, or do one yourself using whatever technique(s) you like. In either case, you will need to posit plausible assumptions about the following:

- The scope of the analysis (e.g., whether to include fuel storage)
- The environment where refueling will be done
- Persons who might be present, including bystanders
- Containers and equipment used to store, move, and dispense fuel
- The type of fuel used
- The design of the aircraft, including the placement of its fuel tank, engine, and
- other components

Construct an argument to support the claim that it is adequately safe to refuel the aircraft as planned.

- Use any argument notation you prefer (for example, prose, structured text, tables, Goal Structuring Notation).
- You may use any residual risk acceptance test you prefer. But it might suffice in this case to allow readers to judge mitigations without appealing to an explicit risk acceptance test.
- Make reasonable assumptions about the kind of grounded premises (evidence) that Jill might provide.
- Focus on how operational risks are mitigated. You may assume that a separate, complete safety case report will discuss remaining issues such as responsible parties and incident reporting.
- Elaborate the arguments regarding one or two hazards down to grounded premises. It is not necessary to elaborate the arguments for all hazards.

Here are answers to some questions that you may have about the assignment.

Q. How long should I spend on the exercise?

No more than 2-3 hours. You need not read about fire hazards or create a perfect argument to complete this exercise.

Q. How do I get started?

You might begin by defining an overall safety conclusion, elaborating what it means in the first argument step, and then arguing over hazard mitigations.

Q. What is the overall safety conclusion?

Specific overall conclusions are prescribed in some domains. For this exercise you might take a broad, intuitive claim such as this following for your conclusion: The refueling operation is adequately safe. Context for this conclusion might be written as, “Procedures for refueling are defined in the airstrip policies and procedures document.”

Q. How do I define ‘adequately safe’?

As you probably know, no uniformly accepted definition of adequate safety exists. In some domains (such as commercial aviation), developers access potential risk then follow a design and development process with commensurate rigor. In other domains, developers are operators perform a risk analysis to determine residual risk than apply a risk acceptance test such as As Low as Reasonably Practical (ALARP). But for the purposes of this exercise you might define ‘adequately safe’ and ‘adequately mitigated’ implicitly through the premises you supply.

Here is an example of using this implicit definition approach for the top-level conclusion, “My word burning stove is adequately safe to use.”

```
Conclusion: My word burning stove is adequately safe to use.
Premises:   The risk of carbon monoxide poisoning is adequately mitigated.
            The risk of a chimney fire is adequately mitigated.
            … …

Reasoning:    Establishing adequate mitigation of identified hazards is
            sufficient to show adequate safety.

Conclusion: The risk of carbon monoxide poisoning is adequately mitigated.
Premise:    My living room is fitted with a functioning carbon monoxide
            detector.
… … …
```

Q. How much detail do I need to include?

As much as you think appropriate to include while abiding to the 2-3 hour time limit. As noted in the module, level of detail is a subject of debate. It is usually possible to add more detail to any argument. But added detail might either illuminate important issues or clutter the argument. Case writers must balance explicitness and brevity. For this exercise, try to develop your argument (for at least one hazard) to a level that seems appropriate for both the matter at hand and the likely readers of the safety argument.

If you have questions or comments about this module, including the homework, contact its author at c.michael.holloway@nasa.gov.