




Commercialization Workbench

A free, private self-assessment for researchers, founders and innovators taking a technology from the lab to the market.



Version 0.5 · commercializationworkbench.com

Contents

1. What the Workbench is
2. Your privacy
3. Getting started
4. The six steps
5. Market sizing: TAM, SAM and SOM
6. How the ARL score works
7. Discovery log and interview guide
8. Compare versions
9. Exports
10. Cohort View for mentors and program managers
11. Tips for an honest assessment
12. Questions and answers
13. About

1 What the Workbench is

The Commercialization Workbench helps you judge how ready a new technology is to be adopted, and what to do next. It combines four approaches that are usually taught separately:

Approach	Source	The question it answers
Adoption Readiness Levels (ARL)	U.S. Department of Energy	What could stop adoption? 17 risks, from cost and demand to regulation, permitting and community acceptance.
Business Model Canvas	Osterwalder and Pigneur	How will the business work: customers, revenue, partners, costs?
Customer discovery	U.S. National Science Foundation I-Corps	How do you know? Every claim is tested in real conversations.
TRIZ	Genrich Altshuller	How do you resolve a trade-off instead of settling for a compromise?

ARL is the backbone. The other three fill its gaps. Your ratings and business model statements are treated as **claims**, your interviews as **evidence**, and the tool flags every claim that has no evidence behind it.

Who it is for: university researchers and faculty, students, lab scientists, independent developers, and founders of hardware and software startups. You can complete it alone and take the exported report to a mentor, advisor, program manager or investor.

The tool is in English. This manual keeps button and screen names in English so you can find them, for example **Save file** or **Rate 17 risks**.

2 Your privacy

The Workbench runs entirely in your web browser. **Nothing you type is uploaded:** there is no account, no server that stores your work, and no tracking of what you enter.

- Your work is saved automatically in this browser on this computer.
- The site counts visits with Cloudflare Web Analytics, so we see only totals such as page views, countries and browsers. It never sees anything you enter in the tool.
- To keep a copy, move to another computer or share with a mentor, use **Save file**. It downloads a .json file.
- To continue later from a file, use **Open file**.
- Clearing your browser data, or using a private window, deletes the copy kept in the browser. Save a file regularly.
- To remove your work from a shared computer, use **Erase this assessment from this browser** in step 6.

3 Getting started

1. Open the Workbench. It starts with a finished example, **GridKeep**, a fictional battery company. Click around it to see what a complete assessment looks like.
2. When you are ready, select **Start your own assessment** in the blue banner.
3. Work through the six steps in the left menu, in order. Your first pass takes one to two hours. Expect to come back as you learn more.
4. Save a file at the end of every session.

The left menu also has three tools you can open at any time: the **Discovery log** (your interviews), the **Glossary** (69 terms in plain language) and **Compare versions**. Underlined terms anywhere in the tool show a short definition when you point at them.

4 The six steps

Step	What you do	Tips
1. Scope	Name and describe the technology, give its Technology Readiness Level (TRL), say what is in and out of scope, and set your commercialization window (6 months to 15 years).	If you have taken an NSF I-Corps or Lean LaunchPad course, choose that track. It shows how your canvas maps onto ARL.
2. Business model	Write down the parts of the business model that ARL does not cover: beachhead customer, market size (TAM, SAM, SOM), the customer's job, revenue model, relationships, partners, alternatives and competitors, and unit economics.	Each card shows who to ask and which ARL risks it feeds. Start the competitor table with "doing nothing": it is usually your strongest rival.
3. Rate 17 risks	For each risk, pick the description that fits your technology today (Low, Medium or High risk, or N/A), choose an evidence tier and explain your reasoning.	Rate what is true now, not what you plan. An N/A needs a written reason.
4. Score and evidence	Read your two scores, your risk profile, the evidence audit and the what-if scenarios.	Fix the flags marked Contradicted first.
5. Resolve trade-offs	Pick up to three focus risks and work through the TRIZ worksheet: resources, ideal result, the obvious fix, what it makes worse, and principles as questions.	The 17 × 17 contradiction matrix suggests principles for any pair of risks.
6. Plan and export	Turn each move into a test with an owner, a date and a target rating. Add questions for your mentor. Export.	A specific question gets a better answer than "any feedback?"

Evidence tiers describe how strong the evidence behind each rating is:

Tier	Meaning	Examples
0	None yet	An assumption the team has not tested
1	Our own analysis	Internal models, desk research, team expertise

Tier	Meaning	Examples
2	Published third-party data	Market reports, government data, peer-reviewed studies
3	Direct outside signal	Conversations with customers, suppliers, regulators or investors; pilot results
4	Commitment or independent validation	Signed letter of intent or contract, certification, independent test, term sheet

5 Market sizing: TAM, SAM and SOM

Highly recommended. Mentors, investors, grant reviewers and I-Corps instructors ask for market size in almost every first meeting. This section also gives your Market Size rating (B2) real evidence instead of a guess.

The **Market size** card in the Business model step sizes your market from the bottom up. Instead of taking a share of a large industry figure, you count real customers and multiply by what one of them spends a year.

Ring	What it counts
TAM: total addressable market	Every customer who could use a solution like yours, anywhere, if nothing stood in the way.
SAM: serviceable available market	The customers your product, channels, regulations and geography let you serve today.
SOM: serviceable obtainable market	The customers you can realistically win within your commercialization window.

1. Decide **what counts as one customer**: a company, a site, a hospital, a household. Use the same unit in all three rings.
2. Enter how many units one customer buys a year. The tool multiplies it by your price per unit from Unit economics. You can also type the revenue per customer directly.
3. Enter the number of customers in each ring, with a source and an evidence tier for each.
4. If you like, enter the market value from a published report as a cross-check.


The checks. The tool tells you when the rings are out of order; when your SOM is an ambitious (10% or more) or aggressive (25% or more) share of your SAM within your window; when your bottom-up TAM and the published figure differ by more than three times; and when even your whole SOM would not sell enough units to cover your overhead. The share thresholds are rules of thumb, not hard limits.

Example. GridKeep counts tower sites. A site replaces a battery cabinet about every four years (0.25 a year) at USD 18,000, so one customer is worth USD 4,500 a year. 400,000 sites give a TAM of USD 1.8 billion, 60,000 a SAM of USD 270 million and 3,000 a SOM of USD 13.5 million, which is 5% of the SAM. But 3,000 sites buy only about 750 cabinets a year, below the 1,715 GridKeep needs to cover its overhead. The tool flags this: a sign to widen the SAM, raise the price or cut overhead.

Your sizing appears on the B2 rating page, in the PDF report, on its own Excel sheet and in the Cohort View.

6 How the ARL score works

The Adoption Readiness Level is a number from 1 to 9. Most confusion about it disappears once you know these five points.



The score page for the GridKeep example: ARL as rated, evidence-adjusted ARL and the rating counts.

1. **Only Medium and High risks count.** The score depends on just two numbers: how many of your 17 risks are Medium and how many are High. Low, N/A and unrated risks add nothing.
2. **A table turns those two numbers into the score.** It is a lookup, not an average or a sum. With no Medium or High risks the score is 9.
3. **High risks cost much more than Medium ones.** One High risk takes you from 9 to 8, two to 7, three to 5, and five or more to 1, whatever else you rated.
4. **Levels fall into three bands:** 1 to 3 is low readiness, 4 to 6 is medium readiness, 7 to 9 is high readiness.
5. **You get two scores.** ARL as rated uses your ratings as given. The evidence-adjusted ARL uses the same table, but counts any Low or N/A rating that your evidence does not support as Medium. The gap between the two shows how much of your result rests on assumptions.

Medium risks ↓ / High risks →	0	1	2	3	4	5	6	7	8+
0	9	8	7	5	3	1	1	1	1
1	8	7	6	4	2	1	1	1	1
2	8	7	6	4	2	1	1	1	1
3	7	6	5	3	1	1	1	1	1
4	7	6	5	3	1	1	1	1	1
5	6	5	4	2	1	1	1	1	1
6	5	4	3	1	1	1	1	1	1
7	3	2	1	1	1	1	1	1	1
8+	1	1	1	1	1	1	1	1	1

DOE's default look-up table. Find your number of Medium risks (rows) and High risks (columns); the cell is your ARL.

Example. GridKeep has 5 Medium and 2 High risks. Row 5, column 2 gives ARL 4. One of its Low ratings has no outside evidence, so the evidence-adjusted count is 6 Medium and 2 High, which gives ARL 3.

When is a Low rating supported? When it has tier 2 evidence or higher, or an interview that supports it. It is not supported when interviews only contradict it. An N/A rating is supported only when you write down why the risk does not apply.

Why didn't my score change? The table moves in uneven steps. Going from 2 Medium risks to 1, or from 4 to 3, gives the same score. With five or more High risks the score stays at 1 until some come down. The **What could move the score** panel on the score page tests these cases for you.

Why did my score fall when a risk improved? Each High risk you bring down to Medium adds one to your Medium count. DOE's table treats seven or more Medium risks as a warning sign on its own, so a change that takes you to seven or more can lower the score. For example, GridKeep's 5 Medium and 2 High risks give ARL 4, but if both Highs become Medium, 7 Medium and no High give ARL 3. Your risk is still lower; the score falls only because of how the table is built. This is a known quirk of DOE's table, not an error. To see a gain, bring some Medium risks down to Low as well.

What the score is not. It is not a probability of success, and ARL 4 does not mean four steps from market. It does not match TRL step for step. DOE itself warns that a single number carries false precision. The most useful output is the list of High risks: those are what could stop adoption.

7 Discovery log and interview guide

The **Discovery log** holds your interviews. For each one, record the date, the person's role (no names), the type of organization, their role in the buying decision and what you learned. Then link it to the claims it informs: click a claim once for **supports**, twice for **contradicts**, three times for **mixed**.

- Linked interviews feed the evidence audit and the evidence-adjusted score.
- Progress bars compare your count with I-Corps targets (20 for regional courses, 100 for the national program).
- **Import from a spreadsheet** accepts rows pasted from Excel or Google Sheets. A column of personal names is skipped.

The **Interview guide** turns your weakest claims into interview questions. Claims with an evidence problem are picked for you. Questions are grouped by who you would ask, and each one says what answer would point to lower or higher risk. Download it as a text file and take it to your next conversation.

Good interviews ask about past behavior, not opinions about the future. Don't pitch. Ask for specifics, listen more than you talk, and end by asking who else you should speak to.

8 Compare versions

Save a file before and after a round of work, then open **Compare versions** in the left menu. Load the earlier file and the later one, or compare a file with your current work. You will see what changed: scores, ratings, evidence tiers, interview links, evidence flags, business model status, new interviews and whether your planned targets were reached.

9 Exports

Format	What it contains	Use it for
PDF report	Both scores, your questions for a mentor, scope, business model, market size, competitors, risk profile, evidence audit, look-up table position, trade-offs, action plan and interview log	Sharing with a mentor, program manager, investor or grant reviewer
Excel workbook	A cover sheet with your scores, then eleven sheets: summary, business model, market size, competitors, ratings, interviews, evidence audit, trade-offs, resources, action plan and the look-up table	Your own analysis, or combining with other data
Save file (.json)	Your complete assessment	Continuing later, moving computers, comparing versions, and sending to a mentor who uses the Cohort View


File names include the product name, the tool version and the date, for example GridKeep_CommercializationWorkbench-v0.5_2026-10-29.pdf.

10 Cohort View for mentors and program managers

The **Cohort View** is a separate page for mentors, incubator and accelerator staff and program managers. It reads the save files of many teams at once and shows the whole cohort side by side. It is just as private: files are read in your browser and never uploaded.

1. Ask each team to complete the Workbench and send you its save file (from **Save file**).
2. Open the Cohort View and select **Add save files**, or drag the files onto the page. Adding a file for a team that is already loaded replaces it.

3. Read the results: a chart of each team's TRL against its ARL, a grid of all 17 risks by team, a team table and every team's questions for mentors.
4. Rows highlighted in orange are High risks shared by several teams. They are good topics for a cohort workshop, a guest speaker or an introduction.
5. Use **Download CSV** to take the data into a spreadsheet.



Cohort View: each team's TRL against its ARL. Filled dots are as rated, hollow dots evidence-adjusted.

Select **Load example cohort** to try it with four fictional teams.

11 Tips for an honest assessment

- Rate each risk for the technology as it is today, at the scale you will first sell it.
- Low ratings are the ones reviewers challenge. Back each one with outside evidence or an interview.
- Talk to people beyond customers: regulators, permitting staff, insurers, installers and neighbors can all slow adoption.
- Write your reasoning in each rating. A mentor learns more from your notes than from the color.
- Reassess after every major round of interviews, and keep the old save files so you can compare.

12 Questions and answers

Question	Answer
My work disappeared.	The browser copy is lost when browser data is cleared. Open your latest save file with Open file .

Question	Answer
Can I use it on a phone?	Yes, but a laptop or desktop is easier for rating and the TRIZ worksheet.
My mentor uses another version of the tool.	Save files from earlier versions open in later ones. Compare versions warns you when the versions differ.
Do I have to fill in everything?	No. Unrated risks count as Low in the score, so the score is marked provisional until all 17 are rated.
Is my score official?	No. It is a self-assessment using DOE's public rubric. It is not reviewed or endorsed by DOE.

13 About

Copyright 2026 Commercialization Workbench, developed by Fernand Gouveia.
commercializationworkbench.com

The code is licensed under the Apache License 2.0 and the content under Creative Commons Attribution-ShareAlike 4.0 (CC BY-SA 4.0). The name and logo are not covered by these licenses.

ARL rubric text and look-up table: U.S. Department of Energy, Adoption Readiness Assessment, Version: April 2025. Business Model Canvas by Strategyzer.com, used under a Creative Commons license. The evidence rules, question bank, evidence audit and TRIZ contradiction map are additions by this tool. They are a draft method that has not yet been validated against commercialization outcomes.

This tool is not affiliated with or endorsed by the U.S. Department of Energy, the U.S. National Science Foundation or Strategyzer.