

Deck review

SIH idea submission — what a juror will see, and what to change before you upload.

Sorbit Problem statement SIH1234 Generated 21 Aug 2026

SAMPLE REPORT — NOT A REAL TEAM

A full-length demonstration on a practice deck, shown so you can judge the depth before you pay. Your report has this structure, written about your deck.

1 Where you stand

4.6

/10

Strong engineering instinct, but the deck asserts impact it never measures and hides its whole stack inside an unreadable diagram.
Weighted across slides 2 to 5. Every parameter is scored out of 10.

Weights

Idea 30%

Technical Approach 30%

Feasibility and Viability 20%

Impact and Benefits 20%

BIGGEST GAP

Slide 4 — Feasibility and Viability

Slide 2 — Idea 30%		5.2/10
Slide 3 — Technical Approach 30%		4.6/10
Slide 4 — Feasibility and Viability 20%		4.0/10
Slide 5 — Impact and Benefits 20%		4.4/10

2 Slide 1 — Title page

Not scored. These are the template's own required fields, checked line by line.

NEEDS WORK the title page overall

✓ **Team members** 6 listed, which is what SIH requires

Missing from the slide

The line itself is gone. Add the field back first — there is nothing here to fill in.

- **Problem Statement ID** is not on the slide - add the line back
- **Theme** is not on the slide - add the line back
- **Team ID** is not on the slide - add the line back

On the slide, but blank

The label is there with nothing after it, or the template's own placeholder is still in place. Type the value in.

- **Problem Statement Title** is blank - fill it in

Filled in

- PS Category Software
- Team Name Sorbit

Also on this slide

- Leftover template text: the heading still reads TITLE PAGE.

3 Slide 2 — Idea

5.2/10

30% OF THE OVERALL

Solution fit

OKAY6/10

The problem statement asks which road segments become impassable and in what order. The deck promises depth maps but never sequencing.

FIX

Say on this slide that you rank segments by failure order.

Novelty strength

NEEDS WORK4/10

Novelty is claimed as a smart platform integrating terrain, rainfall and road data. Integration is plumbing, not novelty.

FIX

State what the simulation does that a rain-gauge threshold cannot.

Clarity

WORKS7/10

Three columns cleanly separate problem, solution and novelty. A juror gets the shape of it in twenty seconds.

✓

Works — nothing to change here.

Specificity

NEEDS WORK5/10

Every line is qualitative. No resolution, no update interval, no city and no accuracy figure anywhere on the slide.

FIX

Add the grid size and the prediction horizon to the solution column.

Differentiation

NEEDS WORK4/10

Nothing here separates this from any flood dashboard. The kinematic wave choice is the real differentiator and is not mentioned.

FIX

Move the low-compute modelling line up from the feasibility slide.

THIS SLIDE

Replace the smart platform line with one sentence naming your method and what it outputs, at what spacing and how often.

4.6/10

30% OF THE OVERALL

Architecture clarity

OKAY 6/10

Boxes and arrows exist, but the diagram never shows where rainfall enters or where a depth value leaves.

FIX Label the entry and exit arrows on the flow.

Stack appropriateness

NEEDS WORK 5/10

GDAL, PostGIS and NumPy fit the problem. Kubernetes for a hackathon prototype does not, and no caching layer appears anywhere.

FIX Drop Kubernetes from the diagram or justify it in one line.

Cost efficiency

NEEDS WORK 4/10

CesiumJS needs an Ion token and bills per tile. Nothing on the slide accounts for that recurring cost.

FIX Name the tile source and its cost, or change the renderer.

Buildability

NEEDS WORK 5/10

Eight technologies in thirty-six hours, with no split of who builds what. The GDAL pipeline alone is a day.

FIX Mark which two boxes are already working today.

Visual readability

NEEDS WORK 3/10

Logos carry the stack with no captions. Projected, the labels under twelve point will not read from the back row of the hall.

FIX Caption every logo at fourteen point or larger.

THIS SLIDE This slide extracts as four words because everything lives in the picture. Add a caption under the diagram naming the flow end to end.

4.0/10

20% OF THE OVERALL

Risk honesty

NEEDS WORK 3/10

The slide is headed Our Core Strengths. Feasibility asks what could fail, and no risk is named anywhere on it.

FIX Name the two things most likely to stop you.

Mitigation quality

NEEDS WORK 3/10

With no risks stated there is nothing to mitigate. Built with open elevation maps is a fact, not a mitigation.

FIX Pair each risk you name with one concrete fallback.

Resource realism

OKAY 6/10

Open elevation data and Python libraries are genuinely available, and the zero-install claim is reasonable for a web application.

FIX Say which elevation dataset, and at what spacing.

Adoption path

OKAY 6/10

Municipalities and emergency teams are named. Who signs it off, and what adopting it costs them, is not.

FIX Name the office that would run this day to day.

Evidence

NEEDS WORK 2/10

High ROI and Demonstrated Market Fit both appear with no number, no source and no pilot behind either phrase.

FIX Cite one flood-damage figure with the report it came from.

THIS SLIDE Retitle this from Our Core Strengths to the risks and what you do about them. A feasibility slide with no risk on it reads as unexamined.

6 Slide 5 — Impact and Benefits

4.4/10

20% OF THE OVERALL

Beneficiary clarity

WORKS 7/10

Emergency teams, municipalities and citizens each get a distinct benefit. That is better than most impact slides manage.

✓ Works — nothing to change here.

Impact measurability

NEEDS WORK 3/10

Eight benefits and not one number. Faster emergency response cannot be checked by anybody reading it.

FIX Put a target minutes-saved figure on the first benefit.

Proportionality

NEEDS WORK 4/10

Climate Resilience and Smarter Governance are both larger than a depth-prediction tool delivers. The claim outruns the build.

FIX Cut the two broadest headings.

Breadth vs depth

NEEDS WORK 3/10

Eight headings of two lines each. A juror remembers none of them, where three developed points would land.

FIX Keep three headings and develop each one.

Credibility

NEEDS WORK 5/10

A ministry official would accept the response and planning benefits, then question the economic loss line that has nothing behind it.

FIX Remove the economic loss heading or source it.

THIS SLIDE Cut this from eight headings to three, and give the first one a number a district officer could check afterwards.

7 Slide 6 — References

Not scored. A juror has to be able to follow a source from the slide.

NEEDS WORK

- References are a QR code only; no citation is readable on the slide.
- A juror reading a printout or a projection cannot follow a QR code.
- No source is named for the elevation or the rainfall data.

8 Technology review

Every tool on your stack, and where a lighter one would do the same job.

React.js

Standard choice for the dashboard, and the ecosystem covers mapping widgets.

✓ Keep it.

CesiumJS

A 3D globe engine doing city-scale overlay work, and it wants an Ion token.

SWAP **deck.gl** over a raster basemap

No Ion token and no per-tile billing.

FastAPI

Good fit. Async request handling, same language as your solver code.

✓ Keep it.

GDAL

Correct for reprojecting and clipping elevation tiles. Nothing lighter does it properly.

✓ Keep it.

NumPy / SciPy

Appropriate for a kinematic wave solver at city grid sizes.

✓ Keep it.

PostGIS

Right choice for the road-network spatial joins you need.

✓ Keep it.

Docker

Sensible. It makes the GDAL install reproducible, which is the usual time sink.

✓ Keep it.

Kubernetes

Orchestration for a single-city prototype nobody is scaling this weekend.

SWAP **one container on a managed runtime**

A day of setup and the cluster bill.

9 Cross-questions

The questions a juror will ask that the deck does not currently answer. Read these out loud before the round.

1 Open elevation data is thirty metre spacing. A city street is ten metres wide. How do you get street-level depth out of that?

Why they ask: Your central claim rests on a resolution you do not have.

SAY THIS

We do not claim absolute depth per street. We rank which segments flood first and how quickly, and that ranking holds at thirty metres because relative elevation is what drives the routing.

⚠ Only say this if you have actually built it.

If you have upsampled the terrain against municipal drain surveys, name the city and state the residual error.

2 What does your system output when the rainfall feed drops mid-storm, which is exactly when the network is worst?

Why they ask: The problem statement requires it to stay usable when a source is unavailable.

SAY THIS

It keeps running on the last reading and labels the map as extrapolated, with the age of that reading shown. A stale prediction that says it is stale is still usable in a control room.

3 Kubernetes is on your architecture slide. What runs inside it, and who maintains that after the hackathon ends?

Why they ask: The stack gets handed to a municipality with no platform team.

SAY THIS

Kubernetes is on the diagram as a deployment target, not something we operate. For the pilot the whole system is one container, and we scoped orchestration out deliberately.

4 How is this different from a rain-gauge threshold wired to an SMS alert, which several cities already run today?

Why they ask: It is the first thing a juror who knows the domain reaches for.

SAY THIS

A gauge tells you it is raining hard where the gauge is. We route the water across terrain, so we can say which streets go under and in what order, before the water reaches them.

9 Cross-questions

5 Your impact slide claims reduced economic loss. Where does that come from?

Why they ask: It is the one claim on the deck with nothing behind it.

SAY THIS

We have not measured it and we are taking the line off the slide. What we can show is the response-time saving, which is what the tool actually changes.

⚠️ **Only say this if you have actually built it.**

If you have a district damage figure from a state disaster report, cite it as context rather than as your own result.

10 Pitch language

OPEN WITH

Heavy rain floods Indian cities street by street, and no city can say which street goes under first; we route rainfall across terrain to rank them hours ahead.

Terms to use

kinematic wave approximation

Names the method you already use and shows you read the hydrology. — slide 3 caption, and again in the Q&A

flow accumulation

The standard name for what your terrain step computes. — slide 3, on the diagram

Manning's roughness coefficient

The parameter that makes your routing physical rather than a heuristic. — the Q&A, if asked how water speed is set

impervious surface fraction

Explains in one term why urban flooding differs from river flooding. — slide 2, in the problem column

return period

How disaster-management officials size a rainfall event. — slide 5, when stating who benefits and when

nowcasting

The correct word for short-horizon rainfall, and it separates you from forecasting. — the opening line

Terms to drop

~~Smart platform integrating terrain, rainfall & road data~~ → **the method itself - terrain-routed flood simulation**

Says nothing, and invites a juror to ask what makes it smart.

~~Demonstrated Market Fit~~ → **the office that asked for this, or cut the line**

A startup phrase with no pilot behind it. Your jurors are ministry officials.

~~High ROI in Disaster Damage Prevention~~ → **one damage number from a named state report**

A return figure with no numerator and no source anywhere.

~~Advanced GIS Data Handling~~ → **the operation - elevation reprojection and road-network joins**

An adjective doing the work a noun should be doing.

11 Positioning

What to put first, what to leave alone, and how to answer honestly when the weak part comes up.

Lead with

You route water over terrain instead of thresholding a gauge.

Put it in the first sentence, before any problem framing.

The whole thing runs on open elevation data and commodity hardware.

Put it in slide 2, beside the solution column.

Three named beneficiaries with genuinely distinct uses.

Put it in slide 5, once it is cut down to three.

Do not volunteer

Runtime of the solver across a full-city grid

Nobody has timed it, and raising it invites a live performance question that costs more than the point is worth.

Comparison against a published hydrodynamic package

Accuracy against an established package is the one question you currently cannot answer at all.

If asked, say this

Thirty metre elevation resolution

Open elevation data is thirty metre spacing, so we report which segments go under first rather than a depth in centimetres. That ranking stays stable at that spacing, and ranking is what a control room acts on.

No pilot deployment and no municipal partner

We have not deployed with a city yet. We built on open data precisely so a municipality can run it without a procurement cycle, and that is the conversation we want after this round.

The economic loss claim on the impact slide

There is no figure behind that line and we are removing it. What we stand behind is the response-time saving, which follows directly from warning people earlier.

12 Spelling

Slide 4 **FEASIBILTY** → **FEASIBILITY**

Scores are computed from the parameter scores in code, not written by the model. Slide score is the mean of its five parameters; the overall is those slide scores at the weights shown above.