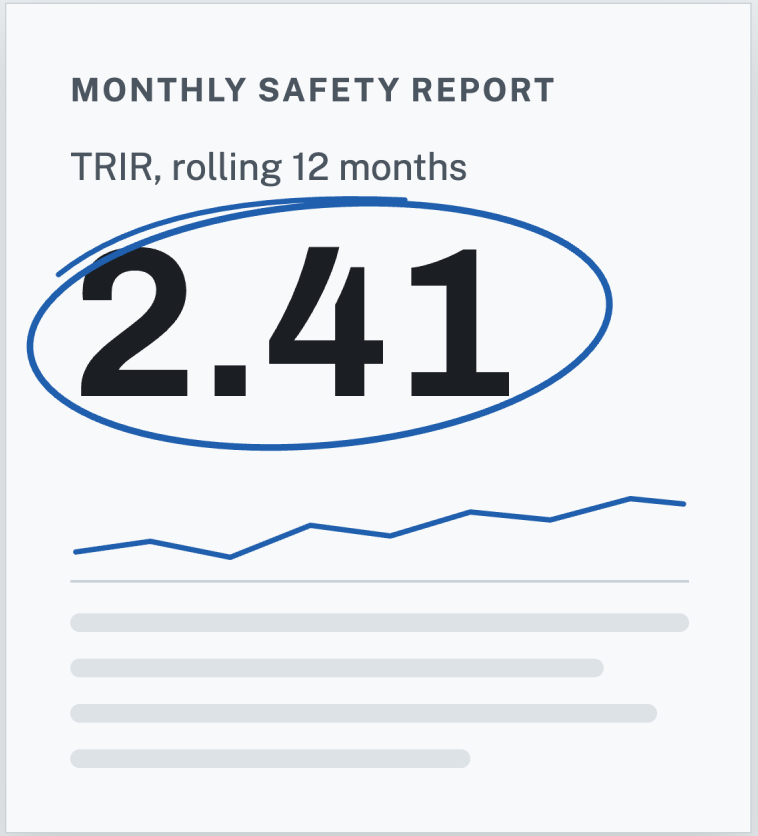
ASSP PACIFIC RIM

Ask better questions

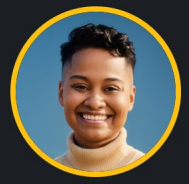
Using AI to surface what EHS data is hiding



Arianna Howard
Syncra Group

Heather Chapman
Paradigm Safety

Two lenses on the same problem



Arianna Howard

EHS Technology Consultant & Managing Partner, Syncra Group

15+ years across safety, risk and operations

EHS technology consulting: system transformation, AI enablement and product-market advisory

Founding member, ASSP AI Task Force

SME, ASSP AI Standards-Based User Group

US Army officer, 16+ years



Heather Chapman

Founder, Paradigm Safety

25 years across safety, risk and emerging technologies

General industry safety consulting, work comp and return-to-work

Insurance experience and technology product advisory

Member, ASSP SBUG

NW Kellogg: AI Applications for Growth

What you'll leave with

by the end of the hour

1. ☐ **Spot the limits**
   Why TRIR, incident counts and trend percentages rarely drive preventive action.
2. ☐ **Build better queries**
   A structured way to surface leading indicators, exposure trends and precursor signals.
3. ☐ **Match the tool to your maturity**
   Which approaches fit your data quality, volume and organizational readiness.
4. ☐ **Leave with a question**
   One new question for your data next week, and a way to start answering it.

PART ONE

01

Why the usual questions aren’t enough

Learning objective 1: the limits of lagging-indicator reporting.

THE USUAL QUESTIONS

Three questions every safety review asks

TRIR

2.41

What's our TRIR?

a rate, after the fact

INCIDENTS, Q3

17

How many incidents this quarter?

a count, with no cause

TREND VS. LAST QUARTER

+8%

Are we trending up or down?

a direction, with no location

AI makes these three faster… not necessarily better.

Your data already tells a story

WHAT THE DASHBOARD SAYS

BTL-2 near-miss reports

+30%

this month

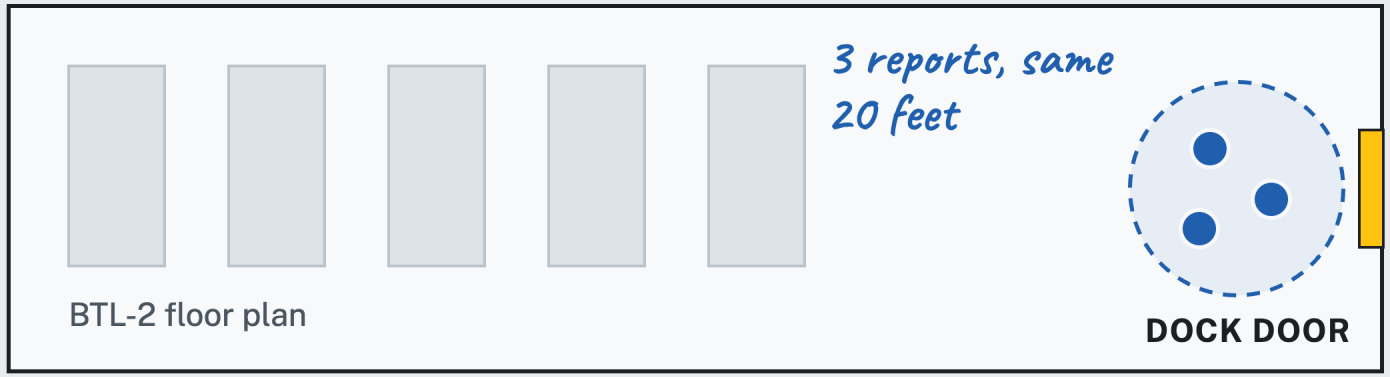
more risk, or more reporting?

Did **exposure** change?

Did **reporting behavior** change?

Did the **form** change? BTL-2's did, in March.

WHAT THE RECORDS ACTUALLY SAY

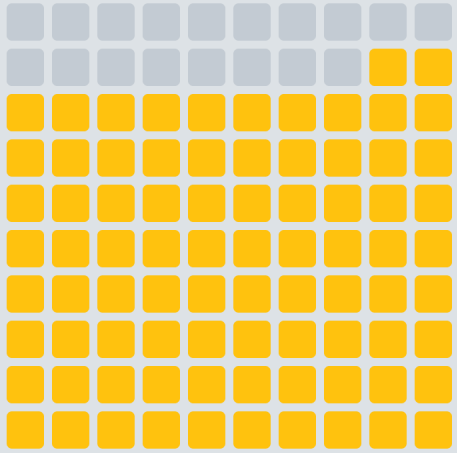


> *"slipped by the dock, night shift"*

> *"wet floor at the loading bay again, nights"*

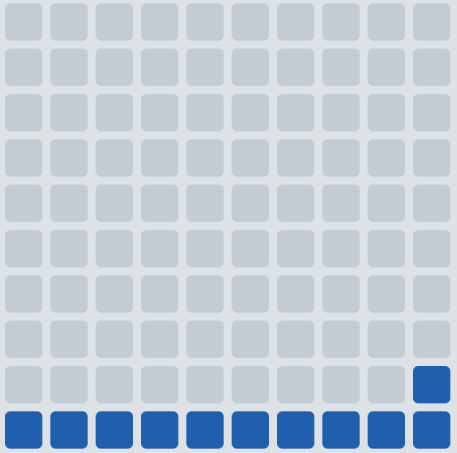
> *"nearly went down near D-2, same twenty feet, 11pm"*

Most teams use AI. Few have the data underneath it.



82%

report extensive, moderate or limited AI use



11%

have fully digital, integrated EHS systems

62% cite **data quality** as a barrier to AI adoption

82% / 11%: NSC + Enablon, The Safety Shift 2026 (n=1,053). 62%: ASSP / Kalorama, Trust, Ethics and Governance (n=992).

LIVE POLL

What's getting in the way of using your safety data today?

A

I can't get the data out of our systems

B

I don't trust the data we have

C

I don't know where to start

D

I know what to do and have no time

PART TWO

02

Build context-driven queries

Learning objective 2: work backwards from a decision to surface leading indicators, exposure trends and precursor signals.

Work your data pipeline backwards

HOW MOST OF US WORK

The data we have → A report we always run → Hope something stands out

WORK IT BACKWARDS

START HERE

The decision we need

Change the dock surface, the cleaning schedule, or neither?

→

THEN

The question we need answered

What is actually driving slips at the BTL-2 dock?

→

THEN

Evidence that could support or challenge it

Near-miss text, cleaning logs, shift records

We're testing possibilities, not asking AI to justify the answer we prefer.

Three examples, each starting from an open decision

start on the left, and let the evidence push back →

1 · Decision we need	2 · Question we need answered	3 · Evidence that could support or challenge it
Change the dock surface, the cleaning schedule, or neither?	What is actually driving slips at the BTL-2 dock?	Near-miss text, cleaning logs, shift records
Fund the lift assist, retrain, or neither?	Which tasks drive strain claims, and how often are they done?	Incident notes, comp claims, task frequency
Change the summer work-rest cycle, or keep it?	When does heat illness happen, and to whom?	Incidents, schedules, temperature logs

CRAFT: five parts of a prompt that works

One prompt, built in five parts

C	**CONTEXT**	I run EHS for a bottling plant, four lines, about 400 people, and our slip reports keep climbing.
R	**ROLE**	Act as an EHS data analyst who has spent years cleaning up messy observation data.
A	**ACTION**	Find the three strongest patterns, and use code for every count rather than estimating.
F	**FORMAT**	A table with the pattern, the record count, and the row numbers so I can verify.
T	**TONE**	Plain enough that a line supervisor can act on it without calling me first.

LIVE POLL

How much of your safety data lands in “Other”?

A

Hardly any

B

Some, maybe a tenth

C

A lot, more than a quarter

D

No idea, and now I want to check

Ask your data whether it can answer you

PROMPT EXCERPT — 12 months of observations, ~1,400 rows

I want to answer one question:

Where are our near-misses actually clustering?

Audit first, analyse nothing.

What percent is blank or “Other”, by field and by site. Entries repeated word for word. Location values that might mean the same place, with the source rows for each. Don’t merge anything: keep the original text and tell me how sure you are.

Finish with what this data cannot tell me, and the three cheapest things to fix.

1 **One useful question**
Give the task a specific purpose.

2 **Audit before analysis**
Check missing fields, duplicates and inconsistent names.

3 **Suggest, don’t correct**
AI flags possible aliases. A person confirms them.

4 **Make limits visible**
Ask what the data cannot support and what to fix first.

Goal: find out whether it can answer where near-misses cluster.

Cleaning up location names

sample rows

Row	As entered (kept)	AI suggests	How sure	Person confirms
412	“dock”	Dock door, BTL-2	High	Confirmed
413	“loading bay”	Dock door, BTL-2	Medium: could mean the north bay	Confirmed: same place
418	“D-2”	Dock door, BTL-2	Medium: a door code, not a place name	Confirmed
421	“Other”	Maybe dock door: text says “same twenty feet”	Low	Not changed. Flagged to follow up

Suggested cleanup: AI’s guess, with its source rows and how sure it is

Verified correction: a person who knows the site confirmed it

Build the context once, then reuse it

PASTE BEFORE EVERY PROMPT ON THIS DATASET

Data dictionary

Near miss	Unplanned event with no injury
Dock door	Also entered as dock, loading bay, D-2
Site codes	BTL-1 through BTL-4
BTL-2	Reporting form changed in March

> *If the data cannot support a pattern, say so rather than reaching.*

For each pattern: what it is, the record count, and the row numbers so I can verify it myself.

20 minutes talking it out loud buys a data dictionary you keep.

Let AI interview you first

For when you know something is wrong but can’t name the question.

YOU

“Interview me first, one question at a time.”

Push back if my answer is vague. Then propose three questions we should ask this data.

Don’t analyse anything yet.

AI ASKS BACK

- What decision are you trying to make?
- What do you already suspect?
- What would have to be true?

an analysis partner, not an analyst

PART THREE

03

Match AI to your data maturity

Learning objective 3: an honest read on where AI adds value, and where it just adds noise.

What AI is good (and bad) at

Strong

Reading thousands of narratives at once

Pulling structure out of free text

Grouping by what happened, not the box someone ticked

Drafting a first version fast

Weak

Math it does in its head

Rates per exposure hours, including TRIR

Anything where field codes are undefined

Knowing which question mattered

Still yours

Deciding what is worth asking

Checking the numbers

The context the data never captured

Signing your name to it

Right-size AI to where your data is

PAPER AND SPREADSHEETS → FULLY INTEGRATED

Starting out

YOUR DATA

Paper or spreadsheets, inconsistent fields, heavy free text

RIGHT-SIZED APPROACH

Chat-based exploration: audit prompts, one export at a time

WHERE NOISE CREEPS IN

Reporting a rate without checking its inputs; treating one pattern as proof

Building

YOUR DATA

One system of record, partially digital, some field discipline

RIGHT-SIZED APPROACH

Saved prompts and light agents; a data dictionary you reuse

WHERE NOISE CREEPS IN

Skipping verification because the tool “usually” gets it right

Mature

YOUR DATA

Fully digital and integrated, exposure and schedule data linked

RIGHT-SIZED APPROACH

Persistent agents on live data; automated first-pass drafts

WHERE NOISE CREEPS IN

Over-automating judgment calls that still need a human sign-off

CHECK THE WORK

Explore with AI → **Verify against the records** → **Report using agreed definitions**

Before you report a rate like TRIR, verify all five inputs:

Numerator	Which cases were counted? Can you find each one in the records?
Denominator	Whose hours, from which timekeeping source?
Period	Do the cases and the hours cover the same dates?
Definitions	Recordable, near miss, site: the agreed versions?
Formula	Cases × 200,000 ÷ hours worked, written out and checked

AND FOR EVERY PATTERN

- ✓ Have it write code for counts, then check two numbers yourself.
- ✓ Ask for the row numbers, then open a few.
- ✓ Delete confident claims your data doesn't support.
- ✓ Watch for smoothing: keep the specific detail.

Correct code can still calculate the wrong thing from the wrong inputs.

PART FOUR

04 Leave with one new question

Learning objective 4: a question for your own data, and a practical way to start answering it.

One finding, two audiences

Verified finding

FOR THE PERSON WHO FUNDS THE FIX

A one-page brief

- What changed
- What it means in downtime and cost
- The single decision I need

Under 250 words. End with one clear ask. Confident, no safety jargon.

FOR THE CREW DOING THE WORK

A two-minute talk

- What we found
- Why it matters to them
- What changes Monday
- What to report if they see it again

Spoken language, no statistics. Respectful: they already know this area is a problem.

then close the loop

- Who owns the response?
- What happens next, and by when?
- How will the people who reported it hear back?

Where to start on Monday

1 Write your data dictionary

- What your fields and codes actually mean
- How key numbers are calculated, and what counts
- The same definitions at every site
- Build it by voice in 20 minutes, then reuse it

2 Fix data where it's created

- Better prompts on the form itself
- Fewer boxes people click straight through
- Voice capture where typing is the barrier

3 Then pick one report

- The one leadership already ignores
- Run the audit prompt on it
- Tell one finding to one person who can act

One useful question. Enough trustworthy evidence. One person who can act.

BEFORE YOU GO

If you could ask your data anything Monday morning, what would you ask?



Thank you!

Questions?

take this home



Ask Better Questions Takeaway Guide

Prompts, CRAFT and checklists

Arianna Howard

Managing Partner, Syncra Group

ahoward@syncragroup.com

Heather Chapman

Founder, Paradigm Safety

heather@paradigm-safety.com