

TRAINING MODULE

Improvement Opportunity Summary

The bird's-eye view, built from what a Plant Head actually asked for

A working module for engineers with 1–2 years on the floor. By the end you will be able to tell “biggest theme” apart from “biggest lever right now” — and know exactly what this tool commits to NOT telling an executive.

Built directly from stated executive feedback — a Voice-of-Customer discipline applied to report design itself

What You Will Be Able to Do After This Module

1

Explain why an executive needs a different first page than an engineer, without losing any underlying rigor.

2

Explain the difference between “biggest theme” and “biggest lever right now” — and why conflating them misleads.

3

Describe all four layers this tool adds on top of the existing candidate pool.

4

Explain why Cost Tier deliberately stops short of an ROI or rupee estimate.

5

Recognize the two most common mistakes junior IEs make when they first present this summary to leadership.

“Where Should I Focus?” — Before Page 2

Every other tool in this series produces real, detailed findings. This tool exists because of one specific, honest piece of feedback:

“A Plant Head or CEO opening this report doesn't want to read every MTM deviation, C2C variation, bottleneck, hidden-factory, line-balance, buffer, ECRS, or Kaizen finding one by one.”

EVERY OTHER TOOL

Built for the engineer doing the work — full detail, one finding at a time.

THIS TOOL

Built for the person deciding where to spend budget and attention — one page, one clear answer, full detail still one click away.

“Biggest Theme” and “Biggest Lever” Are Different Questions

Watch how these two honest, correct answers can point in completely different directions:

BIGGEST THEME

Material Handling — 12 findings across the plant. The most COMMON pattern, plant-wide.

BIGGEST LEVER RIGHT NOW

A single Fixture/Jig finding — but it sits AT the plant's current TOC constraint station. Fixing it moves output today.

An executive who only reads “Material Handling has the most findings” and pours resources there misses the one fix that actually raises throughput today — the same Goldratt lesson from the Bottleneck and Priority Recommendation modules, now aimed at an executive summary instead of an engineering worksheet.

Even a Report Has a Customer

Voice of the Customer isn't just for product design — it applies to how you present findings, too. This tool's whole structure comes from actually listening to what its reader said they needed, not from an engineer's assumption about what would be impressive. Every other module in this series teaches you how to find real findings. This one teaches something different: real findings still fail if the person who needs to act on them can't find the one that matters first.

Rigor and readability are not opposites — this tool is built to prove that.

VOICE OF THE CUSTOMER

*A Lean / Six Sigma
design discipline*

Four Layers, Added on Top of Findings You Already Trust



It does NOT re-investigate anything — it reads the same candidate pool the OPL tool already builds.

1

LAYER 1

Theme Grouping

Built on `opportunity_taxonomy.py` — which itself reuses the already-validated Lean/TPM/ILO/MTM evidence dimensions.

No new keyword bank invented just for this tool — the same vocabulary you've already trained on.

FOUR RAW FINDINGS → ONE THEME

Stn 2: operator reaches into deep bulk bin (MTM EXCEED)

Stn 4: connector missing/double-loaded (PFMEA finding)

Stn 6: parts staged too far from work envelope (ECSR observation)

Stn 7: kitting mismatch causes visual search (Kaizen root cause)

↓ `opportunity_taxonomy.py` groups all four

THEME: MATERIAL HANDLING

2

LAYER 2

Cost Tier, Honestly Bounded

Same discipline as the Control Plan module: commit to what your data can actually support, nothing more.

An ROI number invented from outside data would put a confident-looking guess in front of the one audience most likely to act on it directly.

FOUR TIERS, READ OFF THE FIX TEXT ITSELF

NO COST

“Relocate the parts bin within the Golden Zone.”

LOW COST

“Add a 3D-printed locating fixture.”

MEDIUM COST

“Install a dedicated kitting cart with size-separated bins.”

CAPITAL INVESTMENT

“Install an automated pick-to-light system.”

DELIBERATELY NOT an ROI or rupee estimate — that decision stays with the customer, per standing instruction. A tier tells a manager how big a decision this is; a fabricated ROI number would pretend to know their cost structure.

3

LAYER 3

Fixture/Chute Signal

Explicit and inspectable — you can always click through to see exactly which evidence produced the call.

Never a black-box score with no way to check its reasoning.

THREE LEVELS, EACH WITH A CLEAR BAR

MANDATORY

Repeated misorientation or missing-part evidence with real downstream consequence — a locating device is clearly warranted.

RECOMMENDED

Some reorientation observed, moderate frequency — worth doing, not yet an emergency.

CONSIDER

Isolated or low-frequency evidence — flagged for awareness, not urgency.

Our Stn 2 bulk-bin example: 3 of 8 cycles show reorientation — calls **RECOMMENDED**, with the underlying evidence card one click away for anyone who wants to check the call themselves.


Same Tie-Break Rule as the Bottleneck Module

Every finding is checked against ONE question: does this sit at the line's CURRENT single constraint station?

Finding	Theme	At Constraint?
Stn 2 bulk-bin reorientation	Material Handling	no
Stn 3 fixture missing detent	Fixture / Jig	✓ YES — Stn 3 is today's constraint
Stn 6 layout distance	Layout	no
Stn 7 kitting mismatch	Material Handling	no

Material Handling still wins on raw count (2 of 4 findings) — but the flagged Stn 3 Fixture/Jig finding is the one worth pointing at FIRST for immediate output, exactly the Slide 4 distinction made concrete.

One Bar Chart, Then a Drill Path to Full Evidence



Nothing is lost by simplifying the first page — every finding is still one click from full traceability.

Five Mistakes That Give Junior IEs Away

-  Presenting “biggest theme by count” as if it were automatically the top priority — check the TOC Constraint flag first.
-  Converting Cost Tier into a rupee ROI figure to sound more decisive — that decision was deliberately left with the customer.
-  Treating the Fixture/Chute Signal as a final verdict instead of an inspectable call — always be ready to show the evidence card behind it.
-  Re-investigating findings for this summary instead of trusting the same candidate pool the OPL tool already validated.
-  Forgetting that a LOW theme count can still contain the single most urgent finding on the whole line.

What generate_opportunity_summary.py Does For You

Two outputs, built automatically from the same trusted candidate pool — no re-investigation, no black boxes.

- 1

Improvement_Opportunity_Summary.htm
1

Theme bar chart → station list → evidence card, exactly the Slide 11 drill path, live and clickable.
- 2

Executive Summary sheet (xlsx)

Theme × count × cost-tier table — built to be literally the first page of the printed report.
- 3

Opportunity Detail sheet (xlsx)

Every candidate, one row each, full traceability back to its exact station and source finding.
- 4

Fully dynamic

No hardcoded industry, station count, or station type — inherits the OPL tool's own auto-discovery.

KEY TAKEAWAYS

Recap — Say This Back In Your Own Words

- This tool exists because a real executive asked a specific question — “where should I focus” — before reading any detail.
- “Biggest theme” (most common pattern) and “biggest lever right now” (at the actual constraint) are different questions — never conflate them.
- Cost Tier stops at No/Low/Medium/Capital on purpose — an ROI number was never this tool's decision to make.
- Every summarized number stays one click from full, inspectable evidence — simplicity never means hiding the reasoning.

TALK IT THROUGH — 3 QUESTIONS FOR YOUR NEXT LINE WALK

1. If you handed today's findings to a Plant Head as a raw list, how long would it take them to find the one that matters most?
2. Pick a theme with a high count on your line — does it actually contain the current TOC constraint, or just a lot of noise?
3. Have you ever seen a summary quietly convert a cost tier into a specific rupee figure? Who actually validated that number?