

THEME 01 · DATA, MEASUREMENT & EVIDENCE

Table T1 · Data Quality & Integration

The plumbing under everything else.

WHY THIS TOPIC, WHY NOW

OEHS data infrastructure at most large organizations is built up from two decades of vendor decisions - incident systems bought one way, exposure monitoring another, occupational health a third, HRIS a fourth. The result: signals that don't talk to each other, manual reconciliation, and slow loss of confidence in the underlying numbers.

The pressure has shifted. ESG and SEC human-capital disclosure rules put OEHS data onto the official record. AI deployment needs training-data quality the legacy systems were never built to produce. Insurance carriers and auditors are asking integration and data-lineage questions they rarely asked five years ago. The cost of poor data quality is moving from an internal headache to an external risk. The question is no longer what a perfect architecture looks like - it's what "good enough to act on" looks like.

THE FRAMING QUESTION

"What is the minimum data-quality threshold your organization can actually defend for a real OEHS decision - and what happens, in practice, when the data falls below it?"

FOUR QUESTIONS

Q1 · Strengths (0:10–0:20)

Tell us about a data-integration or data-quality build in your organization that actually held up. What made it work?

Listening for: *Listen for specific design decisions, not vendor names. Strong answers name a particular integration pattern, a governance practice around master data, or a discipline about tracking where data came from. "We consolidated platforms" is less useful than "we made one system the source of truth for X and made the others read-only against it."*

Q2 · Gaps (0:20–0:30)

Where are your biggest integration failures or data-quality gaps right now? What's the day-to-day cost your organization is absorbing because of them?

Listening for: *The moment a participant names a specific recurring "reconciliation tax." Honest answers are usually about person-hours of manual work the official architecture doesn't acknowledge. The volume of that hidden work is often the most useful signal.*

Q3 · Needs (0:30–0:40)

If you had a measurement-infrastructure budget tomorrow, where would the first dollar go? What's the reason that isn't already happening?

Listening for: *The reason behind it. Sometimes it's budget. More often it's unclear ownership between IT, OEHS, and operations, or vendor lock-in, or the lack of a defensible business case. Which one it is matters.*

Q4 · Asks (0:40–0:50)

What does the field need to share or standardize to move the integration question from "every organization for itself" to a shared baseline we can all defend?

Listening for: *Asks here often cluster around shared definitions and data schemas. The ISO 45001 data appendix work and the various ESG-disclosure data schemas are reference points; which one the room is reaching for is itself the signal.*

PATTERNS TO REDIRECT FROM

The platform-consolidation reflex. “We need one system” is a real and partly correct answer, and it can also be a way to dodge the harder conversation about governance and data ownership. Single-vendor environments still produce integration problems if the governance underneath is missing. The redirect: surface where the room mixes up platform consolidation with data-quality governance.

The vendor-pitch drift. Solutions partners arrive at Rotation III. This table pulls toward tool talk more than any other in the Tech lane. If the conversation slides into comparing platform features, lightly: “Let’s park the tool question. What’s the data-governance question underneath?”

The data-as-byproduct epistemology. Most OEHS data was originally captured as a byproduct of compliance reporting, not as evidence to support operating decisions. The data-quality gap is partly a downstream effect of that origin. When a participant catches themselves - “the data isn’t good because no one needed it to be” - that observation is the substantive ground.

WHERE THIS TABLE SITS IN THE LARGER ROOM

Threads from this table connect to others through live theme mapping happening around the room. By the Rotation III handoff, what each table brings in is richer than any single one could produce on its own.

CONNECTED TABLE

WHERE IT REJOINS YOURS

Tech · Performance Metrics & Analytics (Theme 1 sibling)

Same theme, the analytics half of the same problem. If your table reaches the question of what the data is actually for, that one is the natural companion. Rotation III will thread the questions together.

Human · Defining the Healthy Worker (Theme 1 partner)

The measurement-system question your table looks at is the plumbing under their human-side measurement conversation. If the Human side reaches “what could we measure if the systems talked to each other,” this is the mirror.

Human · Measuring Human-Centric Risk (Theme 1 partner)

Validation and integration are the same problem from two sides. Risk measurement sits downstream of integration.

Tech · AI in EHS Maturity & Frameworks (Theme 2)

AI deployment can’t outrun the data layer underneath it. If your conversation reaches the AI-readiness question at the “needs” stage, the governance table is the closest companion.

WHAT THE TABLE IS AND ISN’T TRYING TO PRODUCE

What this table is holding for: Working familiarity with the integration patterns the field actually uses - ESB-style, point-to-point, data lake, API-mediated. Not as expertise; just enough to recognize them. Be comfortable asking the manual-reconciliation question: the honest data on hidden work lives in “how many hours per month are people spending stitching this together by hand?”

What this table isn’t trying to produce: A reference architecture - the room is past that, and forcing one would produce thinner captures than the working-trade-offs conversation will. Vendor recommendations; Rotation III handles the tool conversation. A clean resolution: the data-quality problem is structural across the field and won’t be solved in fifty minutes.

THEME 01 · DATA, MEASUREMENT & EVIDENCE

Table T2 · Performance Metrics & Analytics

What counts as proof at the operating layer.

WHY THIS TOPIC, WHY NOW

Lagging indicators still dominate executive OEHS reporting. TRIR, DART, lost-time frequency, recordable counts - the metrics that became standard under OSHA in the 1970s are still what moves at the board level. Leading indicators have multiplied but haven't replaced them; they sit alongside, often more for political reasons than analytical ones.

The analytics layer has matured faster than the executive layer can absorb - predictive analytics on incident data, leading-indicator dashboards, near-miss density measures, exposure-trend modeling. Microsoft Research's 2025 work on cognitive offloading and the broader literature on algorithmic dashboards suggest that moving from descriptive to predictive analytics changes how decisions get made, often in ways the organization didn't expect. This table opens up: what counts as proof at the operating layer now - and what does the field still have to defend by hand?

THE FRAMING QUESTION

"Which OEHS metrics actually drive executive decisions in your organization - and which are reported, accepted, or ignored? What does the gap say about what counts as proof?"

FOUR QUESTIONS

Q1 · Strengths (0:10–0:20)

Tell us about a metric or analytical product in your organization that actually changed an executive decision. What made it land?

Listening for: *Listen for the conditions, not the metric itself. Substantive answers name a particular relationship with finance or operations, a particular review cadence, or a particular framing that made the metric defensible.*

Q2 · Gaps (0:20–0:30)

Where have leading indicators worked well in your organization, and where have they failed? What's the reason behind the failure?

Listening for: *Look for the gaming pattern. Behavioral observation scores, near-miss reporting volumes, and audit findings are especially vulnerable to being optimized for the score rather than the underlying signal.*

Q3 · Needs (0:30–0:40)

What's the analytical capability gap that's most holding your team back? What would it take to close it?

Listening for: *Watch for the tendency to reach for tool answers. The more useful answer is usually about translation - the relationship between OEHS analytics and the finance, operations, or HR functions that have to act on them.*

Q4 · Asks (0:40–0:50)

If you could swap one OEHS metric in your executive report for another tomorrow, what's the swap? What does the field need to standardize, validate, or defend together to make it possible?

Listening for: *The strongest asks usually involve peer benchmarking. "If I could show my CEO that ten of our peers report X, the swap is easy." The specific swap and the specific benchmark that would unlock it are what matter.*

PATTERNS TO REDIRECT FROM

The TRIR-defense reflex. Some participants will defend lagging metrics on the grounds that they're still required for compliance and external reporting. The defense is correct but beside the point; the question is what justifies the decision, not what gets reported. If the conversation collapses into that debate, redirect: "What's the metric that actually changed how you allocated capital last year?"

The predictive-analytics overclaim. Several vendors and consulting practices now market predictive OEHS analytics with confidence levels their validation can't actually support. If a participant cites a predictive model, ask what it was validated against and over what time horizon. The honest answers are usually narrower than the headline.

The dashboard-as-substitute trap. Visualization tools have improved faster than the underlying analytics. Some organizations have spent heavily on dashboards that present old metrics more attractively without changing what gets decided. Surface where the room recognizes the pattern and what it took to move past it.

WHERE THIS TABLE SITS IN THE LARGER ROOM

Threads from this table connect to others through live theme mapping happening around the room. By the Rotation III handoff, what each table brings in is richer than any single one could produce on its own.

CONNECTED TABLE

WHERE IT REJOINS YOURS

Tech · Data Quality & Integration (Theme 1 sibling)

The data layer underneath the analytics layer. Analytics sits downstream of integration; if your table reaches the question of what could be measured if the systems were better, that's the conversation.

Human · Defining the Healthy Worker (Theme 1 partner)

The "what should we measure" question runs through both tables. Your conversation about what counts as proof and theirs about what we're willing to measure meet at the executive translation layer.

Human · SIF/HOP, Systems, and Strategy (Theme 2)

SIF precursor metrics are a specific case of the leading-indicator question. If your room reaches the precursor-measurement question, that table is the closest companion.

Tech · Building the Business Case (Theme 4)

The analytics layer feeds the business case. Your conversation about what counts as proof in the operating review and theirs about what counts in the funding review are the same conversation at different layers.

WHAT THE TABLE IS AND ISN'T TRYING TO PRODUCE

What this table is holding for: Working familiarity with the leading and lagging indicator landscape - TRIR, DART, near-miss density, audit closure rates, behavioral observation scores. Be comfortable with the gaming question; most leading-indicator programs have a gaming story, and the honest capture is how the organization caught it and what they did about it.

What this table isn't trying to produce: A defense of leading or lagging indicators as a category (the room is past that). A single recommended set of metrics; the field doesn't have one and the room won't produce one in fifty minutes. Vendor pitches; the work here is to surface the analytical substance underneath the platform layer.

THEME 02 · GOVERNANCE, RISK & DECISION RIGHTS

Table T3 · AI in EHS Maturity & Frameworks

Reading the field's position honestly.

WHY THIS TOPIC, WHY NOW

AI in OEHS has moved past proof-of-concept at most leading organizations. Computer vision on the production line, natural-language processing on incident narratives, predictive risk scoring, exposure modeling - the deployments are real. The field hasn't yet settled on shared language for what "mature" looks like, and most of the available maturity frameworks were built by vendors with skin in the game.

The 2024–2025 NIST AI Risk Management Framework, ISO/IEC 42001, and the EU AI Act coming into force have given the broader field reference points OEHS is still translating. The profession has started working out its own position on AI in occupational health practice. This table opens up: what does "ready" mean in operating terms - and what would an honest readout be if someone outside the field audited the deployments tomorrow?

THE FRAMING QUESTION

"If you had to brief your CEO honestly tomorrow on where you actually sit on AI in EHS - deployed, governed, learning, struggling, or still posturing - what would you say?"

FOUR QUESTIONS

Q1 · Strengths (0:10–0:20)

Tell us about an AI deployment in your OEHS function that's actually running at production scale - not a pilot. What conditions made it mature?

Listening for: *Listen for the difference between a pilot that grew and a deployment built for scale from the start. Strong answers name a specific governance practice, a specific data-pipeline discipline, or a specific human-review setup.*

Q2 · Gaps (0:20–0:30)

Where in your AI work is your organization still posturing rather than producing real results? Why?

Listening for: *The posturing question is the hardest to surface and the most valuable. Many organizations have inherited AI maturity claims from RFPs, board updates, or ESG disclosures the operating layer can't actually back up. The silence itself is part of the signal.*

Q3 · Needs (0:30–0:40)

What would have to be true to move your AI work from where it is to where the organization needs it to be? What's blocking it?

Listening for: *The blocker is rarely the technology. It's usually some combination of unclear governance, executive risk appetite, and the lack of a defensible business case for the next investment.*

Q4 · Asks (0:40–0:50)

What would the field need for a defensible shared maturity language and a defensible readout standard? What's the deliverable the room is asking for?

Listening for: *Asks here often cluster around an OEHS-specific overlay on NIST AI RMF or ISO/IEC 42001 - not a new framework, an adaptation of an existing one.*

PATTERNS TO REDIRECT FROM

The maturity-model marketing drift. Several vendors and consultancies have published OEHS-specific AI maturity models with their own interests baked into the levels. If the conversation reaches for a specific framework, ask who built it and what level the framework rewards. The honest captures often come from participants who've had to translate a vendor framework into something they could defend internally.

The pilot-to-production conflation. “We’re running AI in EHS” often means “we have a vendor pilot in one region.” The honest distinction is whether the deployment is producing decisions the organization actually defends, with the governance, audit trail, and human review to back them up. If a participant claims production scale, ask what would happen under audit.

The regulatory-anticipation trap. The EU AI Act, the U.S. AI executive orders, and various state-level frameworks are creating real regulatory exposure - and they're also producing a pattern of organizations claiming compliance with frameworks they haven't actually built against. If a participant cites compliance, ask how it's being evidenced.

WHERE THIS TABLE SITS IN THE LARGER ROOM

Threads from this table connect to others through live theme mapping happening around the room. By the Rotation III handoff, what each table brings in is richer than any single one could produce on its own.

CONNECTED TABLE**WHERE IT REJOINS YOURS****Tech · AI Ethics, Governance & Strategy
(Theme 2 sibling)**

The ethics-and-strategy companion to your maturity-and-frameworks conversation. Same theme, viewed from the design-decisions side. If your conversation reaches the question of which deployments shouldn't mature, that's the natural companion.

**Human · SIF/HOP, Systems, and Strategy
(Theme 2 partner)**

The decision-rights side of the governance question. Your maturity question runs into their decision-rights question whenever an AI signal enters an operating decision.

**Human · Leadership and Legal Risks
(Theme 2 partner)**

Maturity claims, defensibility, and legal exposure are tightly linked. If your conversation reaches what would survive an audit, that table is the mirror.

**Tech · Data Quality & Integration
(Theme 1)**

AI maturity can't outrun the data layer. The Theme 1 Tech table is the plumbing under your maturity question; if your “needs” stage reaches the data-readiness layer, that's the companion.

WHAT THE TABLE IS AND ISN'T TRYING TO PRODUCE

What this table is holding for: Working familiarity with NIST AI RMF, ISO/IEC 42001, and the EU AI Act risk categories - not as expertise, just enough to recognize them. Be comfortable with the posturing question; the honest captures come out when the table makes it safe for participants to name what they've been claiming but can't actually defend.

What this table isn't trying to produce: A decision on which maturity model to adopt. A defense of any specific framework. Vendor pitches on AI governance tools - return to the practice underneath. A clean resolution; the deliverable is three lines on what the field needs in shared maturity language.

THEME 02 · GOVERNANCE, RISK & DECISION RIGHTS

Table T4 · AI Ethics, Governance & Strategy

What the OEHS function should and shouldn't carry.

WHY THIS TOPIC, WHY NOW

AI deployment isn't an OEHS-adjacent question anymore. Worker monitoring, predictive risk scoring, algorithmic discipline, ergonomic tools that look a lot like surveillance, AI-driven coaching - these are deployments the OEHS function is being asked to put its name on, without a shared field practice to lean on.

The institutional pressure is real. ESG disclosure rules are asking more and more specific questions about algorithmic worker management; works councils in European and Latin American operations have formal consultation rights over algorithmic surveillance. The OEHS function is being asked to take a position on AI ethics that the field hasn't collectively articulated yet. This conversation is about what the function should and shouldn't let AI do under its name - and what we need to share to support that work.

THE FRAMING QUESTION

"What is an AI deployment in your organization or your peers' that the OEHS function should not be lending its name to - and what governance would let you change that?"

FOUR QUESTIONS

Q1 · Strengths (0:10–0:20)

Tell us about an AI ethics or governance decision in your organization that actually held - where the OEHS function took a position that shaped what got deployed or how. What made it defensible?

Listening for: *Listen for the defensibility infrastructure. Strong answers name a specific governance body, a specific review practice, a specific principle the function stuck to, or a specific relationship with legal, HR, or executive leadership.*

Q2 · Gaps (0:20–0:30)

Where has the OEHS function been carried along by AI deployments it should have shaped or pushed back on, because the governance authority wasn't there? What was the reason behind it?

Listening for: *The "carried-along" pattern is common and rarely named in public. The honest answers tend to come in the second or third response, not the first, once the silence feels safe.*

Q3 · Needs (0:30–0:40)

What governance practice or institutional position would your organization need to take a defensible stand on an AI deployment the OEHS function shouldn't carry? What's stopping it from getting built?

Listening for: *The blocker is often executive risk appetite or a relationship with HR or operations that the OEHS function hasn't had room to renegotiate. If the room reaches for it, the relationship layer is where the substance is.*

Q4 · Asks (0:40–0:50)

What does the field need to take a position on, build together, or share so OEHS executives have collective cover for the harder governance calls in 2026 and beyond?

Listening for: *The strongest asks tend to be about field-level position statements or principle documents that give individual OEHS leaders cover when they're in front of their executive teams.*

PATTERNS TO REDIRECT FROM

The ethics-as-marketing collapse. “Ethical AI” has become a marketing posture as much as a governance practice. If the room reaches for ethics language without operational specifics, lightly: “What does that look like in the deployment we just discussed?” Operating practice, not principles.

The worker-monitoring euphemism drift. Several deployments currently sold to OEHS functions under names like “wellness monitoring,” “fatigue analytics,” or “predictive ergonomics” are operationally indistinguishable from worker surveillance. The honest conversation about which is which is rare and valuable. If a participant names one of these categories, ask what it actually does.

The let-legal-decide deferral. AI ethics in OEHS deployments often gets handed off to legal counsel, which is reactive and limited by privilege. OEHS leadership is often the function that should be carrying the ethics conversation upstream of legal review. Surface where the handoff has become structural avoidance rather than a real handoff.

WHERE THIS TABLE SITS IN THE LARGER ROOM

Threads from this table connect to others through live theme mapping happening around the room. By the Rotation III handoff, what each table brings in is richer than any single one could produce on its own.

CONNECTED TABLE

WHERE IT REJOINS YOURS

Tech · AI in EHS Maturity & Frameworks (Theme 2 sibling)

The maturity-and-frameworks companion. Same theme, design side versus operations side. The two together produce what the field needs at Rotation III – a coherent answer on which deployments to mature and which to refuse.

Human · Trust & Psychological Safety (Theme 4)

The deployment side of the trust-infrastructure conversation. Your table asks what the OEHS function should refuse to deploy; theirs asks what trust conditions any deployment needs to land.

Human · Leadership and Legal Risks (Theme 2 partner)

Ethics calls are often legal-exposure calls. If your conversation reaches the question of where the governance position needs legal weight behind it, that table is the closest companion.

Human · SIF/HOP, Systems, and Strategy (Theme 2 partner)

Decision-rights infrastructure is what the ethics conversation runs on. If your conversation reaches the question of who has the authority to refuse an AI deployment, that's the conversation.

WHAT THE TABLE IS AND ISN'T TRYING TO PRODUCE

What this table is holding for: Working familiarity with the post-2023 AI ethics literature relevant to workplace deployment – the NIST AI RMF, the OECD AI Principles, the EU AI Act's prohibited and high-risk categories, the algorithmic-management literature. Be comfortable holding the operational-specifics frame against the pull of principle-statement talk.

What this table isn't trying to produce: A defense of AI ethics as a category (the room is past that). A consensus position; the field doesn't have one and the room won't produce one in fifty minutes. Vendor pitches on AI governance tools – when the conversation slides toward governance-platform features, return to the practices.

THEME 03 · ADOPTION, SCALE & CHANGE

Table T5 · Use Cases & Lessons Learned

What the field is actually learning from what it deployed.

WHY THIS TOPIC, WHY NOW

The OEHS-technology deployment landscape has matured enough to produce a real body of operating lessons - wearables programs that worked and ones that didn't, vision-system rollouts that landed and ones that quietly faded, predictive analytics that earned trust and ones that lost it. Most of the available case material is vendor-curated, which skews heavily toward success stories and away from the more useful failure stories.

Professional events and similar venues run case-and-lesson sessions, but the format rewards the polished retrospective over the candid one. This table is built for the candid version: executives comparing what their organizations actually learned, including the deployments that had to be quietly reset. The companion pilot-to-scale table picks up the question of what to do with the lessons.

THE FRAMING QUESTION

“Walk us through an OEHS-technology deployment that taught you something the polished case study wouldn't - the lesson the conference circuit hasn't seen.”

FOUR QUESTIONS

Q1 · Strengths (0:10–0:20)

Tell us about a deployment in your organization that produced an honest lesson the field could use. What was the deployment, and what was the lesson?

Listening for: *Listen for the unexpected finding. Strong answers are moments where the participant says “we thought it would do X and it actually did Y.” Those Y findings are the field-level signal the case-study circuit hasn't produced yet.*

Q2 · Gaps (0:20–0:30)

Where has your organization deployed something and then quietly reset it? What did you learn from that experience that the field doesn't talk about?

Listening for: *The quiet-reset pattern is the most undercaptured category in OEHS technology learning. Vendor-driven case material rarely surfaces these. The silence itself is part of the signal; the honest answers often come on the second response, not the first.*

Q3 · Needs (0:30–0:40)

What does the field need to capture and share so the lessons in this room don't stay in this room?

Listening for: *Asks here often cluster around field-level case-study infrastructure. Participants who've tried to share lessons publicly will have the sharpest sense of what blocks it - legal review, vendor relationships, competitive concerns.*

Q4 · Asks (0:40–0:50)

If peer practice or this convening could produce one shared output from the deployment lessons in this room, what would be most useful? Who builds it, and what makes it stick?

Listening for: *The strongest asks tend to name both the deliverable and the convening practice that would keep it alive.*

PATTERNS TO REDIRECT FROM

The vendor-curated success story drift. Several participants will reach for case material they've presented at vendor events. Those stories are real and partial; they were filtered through a positional lens. If a participant cites a polished case, ask what the unpublished part was. The honest captures live in the unpublished part.

The IP-and-disclosure caution. Some organizations will hesitate to share deployment lessons because of competitive or vendor-relationship concerns. Those concerns are real. The room's working norm is that captures are anonymized for the field-level readout; if a participant hesitates, name the norm and offer the protection.

The recency-bias trap. AI deployments from the last 18 months will dominate the conversation. Some of the most useful lessons come from older deployments - wearables programs from 2018–2022, ergonomic-monitoring deployments from before LLMs, behavioral-analytics rollouts that taught the field about gaming. Pull the room toward the longer arc when it gets stuck on the most recent technology.

WHERE THIS TABLE SITS IN THE LARGER ROOM

Threads from this table connect to others through live theme mapping happening around the room. By the Rotation III handoff, what each table brings in is richer than any single one could produce on its own.

CONNECTED TABLE**WHERE IT REJOINS YOURS****Tech · Moving from Pilot-to-Scale
(Theme 3 sibling)**

Same theme, the next-step conversation. The lessons your table captures feed into the pilot-to-scale conversation. Expect direct cross-pollination during Rotation III; some participants will move between the two tables.

**Human · Creating Actionable Learning Loops
(Theme 3 partner)**

The human-side companion. Their table asks how organizations turn learning into change; yours asks what the field is actually learning in the first place. The two together produce the field-level loop question.

**Human · Engagement & Risk Tolerance
(Theme 3 partner)**

Many of the most useful deployment lessons are about workforce engagement. If your conversation reaches the question of why a deployment landed or didn't with workers, that table is the mirror.

**Tech · New Tools & Approaches
(Theme 4)**

The trust-and-business-case Tech table picks up the question of what the field should invest in next. Your captures feed theirs directly; if a participant bridges, expect it at the "asks" stage.

WHAT THE TABLE IS AND ISN'T TRYING TO PRODUCE

What this table is holding for: Be comfortable holding the candid-lesson layer instead of letting the conversation drift to the conference-polish layer. "What's the unpublished part?" is the question this conversation returns to most. Working familiarity with the OEHS-technology deployment landscape over the last five to seven years - wearables, vision systems, predictive analytics, conversational interfaces. The convening norm: captures from this table are anonymized in the field-level readout.

What this table isn't trying to produce: A polished case-study collection (the room is past that). A vendor recommendation; Rotation III handles tool conversations. A clean resolution on case-format standards; the deliverable is what the field needs to capture honestly, not the format itself.

THEME 03 · ADOPTION, SCALE & CHANGE

Table T6 · Moving from Pilot-to-Scale

What survives the second region.

WHY THIS TOPIC, WHY NOW

Most OEHS technology deployments stall between pilot and production. The pilot succeeds, the case study gets written, the next region or business unit doesn't adopt it, and the deployment quietly turns into a one-region program with a vendor relationship nobody's maintaining. The pattern is widely recognized inside organizations and rarely named in public.

The structural reasons cluster: pilots usually run with conditions that don't carry over - dedicated people, executive attention, white-glove vendor support, a hand-picked site. The pilot proves the technology works under conditions the rest of the organization doesn't have. Scaling up means reproducing the conditions, not just the technology. Most organizations underinvest in the conditions and then read the scale-up failure as a technology problem.

THE FRAMING QUESTION

"Walk us through a deployment in your organization that successfully crossed from pilot to scale - or one that didn't. What were the operating conditions that made a difference?"

FOUR QUESTIONS

Q1 · Strengths (0:10–0:20)

Tell us about a deployment that actually scaled past the pilot region and held. What were the operating conditions that carried it?

Listening for: *Listen for the conditions, not the technology. Strong answers name a specific governance practice, a specific role setup, or a specific budget structure that let the scale-up survive without dedicated attention.*

Q2 · Gaps (0:20–0:30)

Where have you watched a successful pilot fail to scale, or your own deployment fail to cross into the second region? What was the reason behind it, and what would you have done differently?

Listening for: *The reason is rarely "the technology didn't work in the second region." It's usually about the conditions: dedicated people that didn't carry over, executive attention that moved on, vendor support that thinned out.*

Q3 · Needs (0:30–0:40)

What infrastructure would your organization need to make pilot-to-scale the default rather than the exception? What's stopping it from getting built?

Listening for: *Asks here often cluster around a different pilot-design discipline at the front end - designing pilots that are honest tests of scale conditions, not dedicated demonstrations.*

Q4 · Asks (0:40–0:50)

What does the field need to share about pilot-to-scale infrastructure so the pattern moves from individual-organization learning to a shared baseline?

Listening for: *The strongest asks will be specific: a pilot-design standard, a peer scale-up benchmark, a vendor-evaluation framework that includes scale-readiness as a criterion.*

PATTERNS TO REDIRECT FROM

The “champion left” explanation. Many scale-up failures get explained with “the regional champion left.” That’s partly true and partly a way to avoid the harder question: why was the deployment dependent on a single person in the first place? Press past the champion-departure framing toward the design question - what would the deployment have needed to survive any single person’s departure?

The vendor-support thinning problem. Pilots get white-glove vendor support that production deployments don’t. The thinning happens whether or not the contract changes. When a participant catches this pattern, the honest captures are about how organizations have built internal capability to replace that white-glove support.

The funded-pilot, unfunded-scale trap. Capital and project funding cover pilots; operating budgets have to absorb the scale-up. The transition is often invisible until it breaks. If a participant cites a scale failure, ask what the funding model looked like in year two. The honest answer is sometimes “nobody owned it.”

WHERE THIS TABLE SITS IN THE LARGER ROOM

Threads from this table connect to others through live theme mapping happening around the room. By the Rotation III handoff, what each table brings in is richer than any single one could produce on its own.

CONNECTED TABLE

WHERE IT REJOINS YOURS

Tech · Use Cases & Lessons Learned (Theme 3 sibling)

Same theme, the input conversation. The lessons captured at the sibling table feed directly into your scale-up conversation. Expect direct cross-pollination; some captures may need to move between tables.

Human · Engagement & Risk Tolerance (Theme 3 partner)

Workforce engagement is the most under-priced scale-up condition. If your conversation reaches the question of why a deployment landed in the pilot region but didn’t in the next one, that table is the mirror.

Human · Creating Actionable Learning Loops (Theme 3 partner)

Pilot-to-scale failures are partly a learning-loop failure - the organization knows the pilot didn’t generalize and doesn’t turn that into change. The sibling Human table is the companion at the “needs” stage.

Tech · Building the Business Case (Theme 4)

Scale-up funding is often the binding constraint. The business-case table is asking what makes the next investment defensible; your table is asking what makes the existing investment survive. The two meet at the operating-budget transition.

WHAT THE TABLE IS AND ISN'T TRYING TO PRODUCE

What this table is holding for: Be comfortable holding the conditions-not-technology frame against the pull of the technology-comparison frame. “What’s the condition underneath the technology?” is the question this conversation returns to most.

What this table isn’t trying to produce: A pilot-design framework; the room will reach for several, and the deliverable is what the field needs, not which framework to adopt. A clean resolution; pilot-to-scale is a structural problem in the field and the room won’t solve it. Vendor pitches; the move is to return to the conditions underneath.

THEME 04 · TRUST, CULTURE & THE BUSINESS CASE

Table T7 · New Tools & Approaches

What's genuinely new, what's rebranded, what's worth funding.

WHY THIS TOPIC, WHY NOW

The OEHS-technology market produces a new wave of “new tools” every eighteen months - tied to venture funding cycles, not operational need. Most of what gets marketed as new is older capability with a generative-AI wrapper or a different go-to-market angle. Some of it is genuinely new. Telling the difference at the moment of purchase is one of the hardest jobs an OEHS leader carries.

The post-2023 generative-AI wave has accelerated the marketing churn. Vendor decks describe nearly every product as AI-enabled; the substance ranges from genuinely transformative to UI-level. Buyer fatigue is real, and so is the cost of missing the transformative tools because they look like more noise. This table comes before the business-case table: what's actually new, and how does the field tell the difference?

THE FRAMING QUESTION

“What is a new OEHS tool or approach that you’ve evaluated in the last year that is worth the field’s attention - and what is one that’s being marketed as new but isn’t?”

FOUR QUESTIONS

Q1 · Strengths (0:10–0:20)

Tell us about a new tool or approach your organization has evaluated or adopted recently that's genuinely worth the field's attention. What makes it actually new, and what does it need to land?

Listening for: *Listen for the specific capability that's new, not the marketing language. Strong answers name a particular operational capability. “It uses AI” is not a capability claim.*

Q2 · Gaps (0:20–0:30)

What's being marketed as new in the OEHS space that isn't - or that is, but not in the way the marketing claims? Where is the gap between the vendor story and operating reality wider than we say out loud?

Listening for: *Look for the rebranding pattern. Generative-AI wrapping on existing capability is the most common; predictive-analytics rebranding of basic trend reports is close behind.*

Q3 · Needs (0:30–0:40)

What evaluation infrastructure would help the field tell the difference at scale, not just at the individual buyer level? What's blocking it?

Listening for: *Asks here often cluster around independent evaluation. Participants who've run internal pilots with real validation discipline will have the sharpest sense of what infrastructure would help.*

Q4 · Asks (0:40–0:50)

If the field could produce one shared output that helps OEHS executives read the new-tools market more reliably, what would be most useful? Who builds it?

Listening for: *The strongest asks will be specific about both the artifact and the convening practice that supports it. Field-level market-evaluation infrastructure is a long-game effort, and this room is one of the few places it could actually be started.*

PATTERNS TO REDIRECT FROM

The AI-wrapped-everything drift. Almost every vendor pitch in 2025–2026 describes its product as AI-enabled. Some of that is real capability; much of it is UI-level or marketing-level. If a participant cites an AI capability claim, ask what the model is actually doing, what it's been validated against, and what the operating workflow looks like. The honest answers are usually narrower than the headline.

The vendor-pitch drift. This table pulls toward tool-comparison talk more than most. Solutions partners arrive at Rotation III; if a participant slides toward platform comparison, lightly: "Let's park the platform comparison. What's the capability question underneath?"

The novelty-fatigue trap. Some senior OEHS executives have developed a defensive skepticism that misses genuinely new capability when it appears. It's the opposite of the novelty-marketing problem and just as costly. Surface both patterns and ask the room to name the cases where their skepticism cost them something the field needed.

WHERE THIS TABLE SITS IN THE LARGER ROOM

Threads from this table connect to others through live theme mapping happening around the room. By the Rotation III handoff, what each table brings in is richer than any single one could produce on its own.

CONNECTED TABLE

WHERE IT REJOINS YOURS

Tech · Building the Business Case
(Theme 4 sibling)

Same theme, the funding-side companion. The capability judgments your table builds are the input to the funding-case conversation. Rotation III will thread the two together.

Tech · Use Cases & Lessons Learned
(Theme 3)

The deployment-experience side of the new-tools question. The lessons captured at that table sharpen the judgment your table is building. If your conversation reaches the "how do we evaluate" question, that table is the input.

Human · Trust & Psychological Safety
(Theme 4 partner)

New tools - especially monitoring and predictive ones - land or don't depending on the trust infrastructure underneath. If your conversation reaches the workforce-reception question, that table is the mirror.

Tech · AI in EHS Maturity & Frameworks
(Theme 2)

Maturity-honesty and new-tool-evaluation are the same skill in different domains. If your table reaches the question of how the field evaluates capability claims at scale, that table is the companion.

WHAT THE TABLE IS AND ISN'T TRYING TO PRODUCE

What this table is holding for: A working read on the current OEHS-technology market - wearables, vision systems, NLP on incident narratives, conversational interfaces, predictive risk scoring. Not exhaustive; just enough to recognize when participants reach for specific categories. Be comfortable with the validation question: "What is that validated against?" is the most useful redirect.

What this table isn't trying to produce: A defense of any specific tool category (the room is past category-level debate). Vendor recommendations - Rotation III handles the tool conversation. A clean resolution on which tools to buy; the deliverable is what the field needs to read the market reliably, not which products to recommend.

THEME 04 · TRUST, CULTURE & THE BUSINESS CASE

Table T8 · Building the Business Case

What gets funded after the wellness-program credibility hole.

WHY THIS TOPIC, WHY NOW

Building a defensible business case for OEHS-technology investment is one of the most consequential and least-supported skills in the field. Executives in this room can name multiple investments they can't get funded. The translation problem - from operating need to a CFO-grade case - is structural, and the field hasn't built the shared evidence to scale the practice past individual leaders' translation skills.

The credibility environment has tightened. Corporate wellness took heavy damage between 2019 and 2023 - the Goldman Sachs internal 2019 readout, the JAMA workplace wellness studies, and the broader literature on employer-sponsored health interventions raised the bar. ESG disclosure, SEC human-capital reporting, and insurance audits have moved OEHS cases closer to the official record.

THE FRAMING QUESTION

"If you walked into your CFO's office tomorrow with an OEHS-technology investment case, what would the case need to include - in 2026 terms - to clear the bar?"

FOUR QUESTIONS

Q1 · Strengths (0:10–0:20)

Tell us about an OEHS-technology investment case in your organization that got funded and produced a result the finance team will defend. What did the case look like, and what made it land?

Listening for: *Listen for the specific financial signal. Strong answers name TCOR movement, retention numbers, a shift in workers' comp loss ratio, an ESG disclosure consequence, or a CFO-level metric. "Our team liked it" is not a financial signal.*

Q2 · Gaps (0:20–0:30)

Where is the OEHS-technology business case still costing you arguments rather than winning them? What part of the case doesn't land at the executive layer?

Listening for: *The wellness-program credibility hole will surface. When it does, what matters is what specifically the organization is now skeptical of - and what the OEHS-technology case has to prove to clear the bar.*

Q3 · Needs (0:30–0:40)

What financial signal, peer benchmark, ESG disclosure standard, or insurance-industry framework would have to exist for the OEHS-technology investment case to be defensible at scale rather than depending on an individual leader's translation skill?

Listening for: *Asks here often cluster around comparable-company benchmarks, alignment with ESG and SEC human-capital reporting, and insurance-carrier audit framing.*

Q4 · Asks (0:40–0:50)

What does the field need to share, build together, or anchor through professional bodies so the OEHS-technology business case becomes defensible by default rather than by translation?

Listening for: *The strongest asks tend to be about shared evidence rather than shared templates. A peer-organization case library with real financial outcomes is more useful than a template.*

PATTERNS TO REDIRECT FROM

The wellness-program ghost. Most of the executives in this room lived through the wellness-program credibility collapse. The frame still carries baggage. If the conversation starts to feel like a wellness pitch in OEHS-technology clothing, name it: “What’s the part of this case that has to be different from the 2015 wellness program to actually get funded?”

The ROI-overclaim trap. Vendor case materials routinely overclaim ROI by cherry-picking time horizons, ignoring opportunity cost, or counting downstream benefits that don’t survive scrutiny. If a participant cites an ROI number, ask the source and the time horizon. The honest internal cases are usually narrower and more defensible than the vendor materials.

The CFO-translation gap. Many OEHS leaders can describe what they want to fund but can’t translate it into the language a CFO will defend to a board. The translation problem is structural and the room can do real work on it. When participants who’ve made the translation cleanly speak up, their exact language is what the rest of the room needs.

WHERE THIS TABLE SITS IN THE LARGER ROOM

Threads from this table connect to others through live theme mapping happening around the room. By the Rotation III handoff, what each table brings in is richer than any single one could produce on its own.

CONNECTED TABLE**WHERE IT REJOINS YOURS****Tech · New Tools & Approaches
(Theme 4 sibling)**

Same theme, the conversation that comes before this one. The capability judgments built at the sibling table feed directly into your funding-case conversation.

**Human · Safety Culture for 2026 & Beyond
(Theme 4 partner)**

The culture companion. Their table asks what culture has to become to hold up under AI-driven work and demographic change; yours asks what investment the organization is actually willing to fund to support it.

**Tech · Performance Metrics & Analytics
(Theme 1)**

Business cases sit downstream of analytics. The metrics conversation produces the evidence base your case rests on; if your table reaches the analytical-defensibility question, that table is the mirror.

**Tech · Moving from Pilot-to-Scale
(Theme 3)**

Scale-up funding runs through your table. If your conversation reaches the operating-budget transition question, pilot-to-scale is the closest companion.

WHAT THE TABLE IS AND ISN'T TRYING TO PRODUCE

What this table is holding for: Be comfortable with executive-financial language at the level a CFO uses: TCOR, retention cost, workers' comp loss ratio, SEC human-capital metrics, ESG disclosure frameworks, insurance-carrier audit categories. Be aware of the wellness-program credibility hole and what an OEHS-technology case has to do differently from a wellness business case to clear it.

What this table isn't trying to produce: A defense of any specific investment category (the room is past that). A single agreed-on ROI methodology - the field doesn't have one and the room won't produce one in fifty minutes. Vendor pitches; the work here is to surface the business-case substance underneath the platform layer.

THEME 05 · SKILLS & THE HUMAN-TECH INTERFACE

Table T9 · Human-in-the-Loop Design

Where the human stays, and on what terms.

WHY THIS TOPIC, WHY NOW

Human-in-the-loop has gone from being a design principle to being a regulatory expectation. The EU AI Act requires meaningful human oversight on high-risk deployments; NIST AI RMF treats it as something measurable; ISO/IEC 42001 builds it into management-system requirements. “A human reviews it” is no longer a defense unless the review is operationally meaningful, properly time-budgeted, and defensible.

The practical question is harder than the principle. Most human-in-the-loop deployments today are pro-forma - a human signs off on outputs they can't meaningfully evaluate at the cadence required. The human-factors literature on automation complacency and the 2025 Microsoft Research work on AI-driven cognitive offloading both predict this pattern. The speed hasn't slowed down.

THE FRAMING QUESTION

“In your live AI-assisted OEHS deployments, where is the human-in-the-loop genuinely doing meaningful oversight - and are there outputs the human can't evaluate at speed?”

FOUR QUESTIONS

Q1 · Strengths (0:10–0:20)

Tell us about a deployment in your organization where human-in-the-loop is operationally meaningful. What design choices made the oversight real?

Listening for: *Listen for the design discipline. Strong answers name a specific workflow choice, a specific time budget for review, a specific role structure, or a specific piece of friction the organization built in.*

Q2 · Gaps (0:20–0:30)

Where in your deployments is the human-in-the-loop currently pro-forma rather than operationally meaningful? What's the reason it hasn't been redesigned?

Listening for: *Look for the cadence-versus-capability gap. Many deployments fail human-in-the-loop because the volume of model output overwhelms the available review time - or because the reviewer doesn't have the foundational judgment to overrule the model.*

Q3 · Needs (0:30–0:40)

What would your organization need to make human-in-the-loop operationally meaningful at scale, not just in pilots? What's blocking it?

Listening for: *Asks here often cluster around upstream pilot-design choices and around the practitioner-training pipeline.*

Q4 · Asks (0:40–0:50)

What does the field need to share, take a position on, or build together so meaningful human-in-the-loop becomes a defensible operating practice rather than a regulatory aspiration?

Listening for: *Asks may include design standards, training and credentialing requirements, audit frameworks, and anchored field-level positions.*

PATTERNS TO REDIRECT FROM

The signature-as-oversight trap. A human signature on AI output is not the same as meaningful oversight. Some deployments treat the signature as the compliance artifact and back off the operating discipline entirely. If a participant cites human review as evidence of meaningful oversight, ask what the review actually does, on what time budget, with what authority to overrule.

The exception-flagging overload. Several common deployment patterns flag exceptions for human review at volumes that overwhelm the reviewers' capacity. The result is rubber-stamping under operational pressure. When participants name the volume their reviewers actually handle versus the volume the model is flagging, the gap is the substantive capture.

The reviewer-capability erosion. The 2025 cognitive-offloading research and the broader automation-complacency literature both predict that reviewers lose the judgment they need to overrule the model the longer they review at scale. If the table has people running reviewers in active deployments, ask about reviewer-skill preservation.

WHERE THIS TABLE SITS IN THE LARGER ROOM

Threads from this table connect to others through live theme mapping happening around the room. By the Rotation III handoff, what each table brings in is richer than any single one could produce on its own.

CONNECTED TABLE

WHERE IT REJOINS YOURS

Tech · Preparing the Profession (Theme 5 sibling)

Same theme, the workforce-pipeline companion. Meaningful human-in-the-loop requires reviewers with judgment that the next generation of practitioners may or may not be developing. Rotation III will thread your design conversation with theirs about the human in the loop.

Human · Critical Thinking Under AI (Theme 5 partner)

The human-side mirror to your design conversation. Their table asks what skills the practitioner still has to develop; yours asks what the design has to require to actually use those skills. The two together produce the field's position on the human-tech interface.

Tech · AI Ethics, Governance & Strategy (Theme 2)

Human-in-the-loop is one of the load-bearing pieces of defensible AI governance. The Theme 2 Tech table is the policy-side companion.

Human · Leadership and Legal Risks (Theme 2)

Pro-forma human-in-the-loop becomes a legal-exposure problem when an AI deployment produces a harm. The legal-risk table is the closest companion if your table reaches the defensibility-under-discovery question.

WHAT THE TABLE IS AND ISN'T TRYING TO PRODUCE

What this table is holding for: Working familiarity with the human-in-the-loop literature - EU AI Act requirements, how NIST AI RMF treats it, the automation-complacency research, and the Microsoft Research 2025 cognitive-offloading work.

What this table isn't trying to produce: A defense of human-in-the-loop as a design principle (the room is past that). A specific design framework, or vendor pitches on oversight tooling - the deliverable is shared practice, not choosing a framework.

THEME 05 · SKILLS & THE HUMAN-TECH INTERFACE

Table T10 · Preparing the Profession

What the 2030 practitioner has to know that the 2020 practitioner didn't.

WHY THIS TOPIC, WHY NOW

The OEHS profession is being asked to absorb the biggest skill transition in a generation while the pipeline has been shrinking for over a decade. ABET-accredited industrial hygiene programs have consolidated; BLS projects about five percent growth in OEHS roles through 2033, below the average for all occupations. The scope has gotten wider, and AI-fluency expectations are rising.

Credentialing and professional bodies have all been moving on this. AI-fluency components are showing up in credential update conversations; Total Worker Health® competencies are entering the credentialing landscape; recognition of non-traditional paths is on the table. The repositioning is real and partial; the executive view of what 2030 requires isn't consolidated yet, and the credentialing bodies need to hear from this room directly.

THE FRAMING QUESTION

"If you described the OEHS practitioner we will need in 2030 - what would they have to be able to do that the 2020 practitioner didn't, and how far are we from producing them?"

FOUR QUESTIONS

Q1 · Strengths (0:10–0:20)

Tell us about where your organization is genuinely developing the next-generation practitioner well. What operating elements make it work?

Listening for: *Listen for the operating elements, not the program names. Strong answers name a specific deployment-experience requirement, a specific cross-functional rotation, or a specific AI-fluency practice.*

Q2 · Gaps (0:20–0:30)

Where is the current pipeline most visibly failing to produce the practitioner the human-tech interface requires? What's the reason behind it?

Listening for: *Listen for the specific skill. The most useful answers are concrete: "we can't hire a senior IH with the AI-evaluation judgment to second-guess a vision-system deployment" is the kind of specificity the credentialing bodies need to hear.*

Q3 · Needs (0:30–0:40)

What does the credentialing ecosystem need to build to produce the 2030-ready practitioner?

Listening for: *Asks here often cluster around specific credential-modernization elements: an AI-fluency competency layer, a deployment-experience requirement, a recognition path for non-traditional entrants.*

Q4 · Asks (0:40–0:50)

If credentialing and professional bodies could hear one thing from this room about the human-tech-interface skills executives need in 2026 and beyond, what would it be? What's the message that has to land?

Listening for: *The field-level deliverable. What surfaces here goes directly into the leadership readout.*

PATTERNS TO REDIRECT FROM

The AI-fluency overclaim. “AI-fluent” is becoming a credential-marketing term faster than the field has defined what it actually means operationally. If a participant cites AI-fluency as a requirement, ask what it should let the practitioner do that they currently can’t. The honest answer is the substantive ground; the marketing claim is downstream.

The credentialism reflex. Senior practitioners often default to defending existing credentials as the answer to pipeline pressure. That’s partly right and partly path-dependent. Press the room: which credentials are actually predicting practitioner performance in human-tech-interface work, and which are just inherited? The honest answer is more useful than the defensive one.

The academia-blaming detour. It’s easy and partly fair to put the pipeline problem on the shrinking of academic OEHS programs. It’s also incomplete. The executives in this room have decision rights over hiring, development investment, deployment-experience design, and what gets funded at universities. If the conversation drifts toward pushing the problem outward, redirect: “What’s in this room’s decision rights?”

WHERE THIS TABLE SITS IN THE LARGER ROOM

Threads from this table connect to others through live theme mapping happening around the room. By the Rotation III handoff, what each table brings in is richer than any single one could produce on its own.

CONNECTED TABLE

WHERE IT REJOINS YOURS

Tech · Human-in-the-Loop Design (Theme 5 sibling)

Same theme, the deployment-side companion. The practitioner of 2030 needs to be the human in the loop that the design table is building for. Rotation III will thread your practitioner-development conversation with theirs on design discipline.

Human · Workforce Development & Skills (Theme 5 partner)

The human-side companion. Their table asks how the field develops the workforce broadly; yours asks what specifically the human-tech interface requires. The two together produce the field-level workforce position.

Human · Critical Thinking Under AI (Theme 5 partner)

The skill-preservation companion. The judgment your practitioner has to bring to a deployment is exactly the judgment the critical-thinking table is asking how to develop and preserve. The “asks” stage will converge.

Tech · AI in EHS Maturity & Frameworks (Theme 2)

Maturity claims require practitioners who can actually carry them. If your conversation reaches the question of whether the field can actually staff the deployments it claims to operate, that table is the mirror.

WHAT THE TABLE IS AND ISN'T TRYING TO PRODUCE

What this table is holding for: Working familiarity with the credentialing landscape, the BLS occupational outlook for OEHS roles, and the recent conversation on AI-fluency standards in adjacent professions. Be comfortable with the specifics question: “What should that practitioner be able to do that the 2020 practitioner couldn’t?”

What this table isn’t trying to produce: A defense of any specific credential (the room is past that). A single agreed-on theory of the OEHS practitioner of the future – the field doesn’t have one. Vendor pitches on training platforms; the work here is the field-level competency question underneath.