RAIL Self-Grading Scorecard

For: C-suite / VP and function leads · **Type:** Scorecard

SYSTEM BEING SCORED _____ SCORED BY / DATE _____

Score one AI tool or system against the six RAIL components, reach a shared verdict, and identify the gaps to close.

How to use it

1. Choose one system and bring together the relevant business, IT, and legal or compliance owners.
2. Score the system from 1–3 on each component using the rubric below and evidence you can verify. Record the evidence alongside each score.
3. Total the score, apply the floor rule, and use the weakest components to set your next actions.

Your score

Component	Evidence	Score (1–3)
Human Centricity		
Value Creation		
Orchestration		
Governance ⚑		
Security ⚑		
Testing & Certification		
Total		___ / 18

Floor check: Governance = 3? ☐ Yes ☐ No · Security = 3? ☐ Yes ☐ No

Tier: ☐ High ☐ Medium ☐ Low

The rubric

Component	1 Absent / ad hoc	2 Partial / inconsistent	3 Established / enforced
Human Centricity role architecture · context grounding · human ownership of consequential action	Generic, role-agnostic assistant; no human checkpoint on consequential action	Some role tailoring and context; human-in-the-loop exists but is informal or inconsistent	Built role-by-role from real workflows; rich job context; human ownership of consequential action explicit and enforced
Value Creation value defined upfront · before/after baseline · ROI mapped to a number	Value asserted, not defined or measured; no baseline, no metric	Value defined and a metric or two tracked, but no clean before/after comparison	Target value defined upfront, measured against the human baseline, ROI mapped to a real number leadership can compare
Orchestration connected systems · agent-to-agent hand-offs · context carried across data sets, systems, and workflows	System-bound AI working in isolation; agents cannot work across systems or with other agents, hand-offs to other tools are manual	Connects to some systems or agents, but hand-offs are brittle or context is lost between steps	Agents coordinate across systems, other agents, and approved data sets; context and guardrails are upheld at every hand-off
Governance ⚑ named owner & escalation · in-system SOPs, guardrails & disclaimers · enforced at agent and sub-agent level	No named owner; compliance lives in human review only; no in-system guardrails	Named owner; guardrails and disclaimers exist at the top level, but the sub-agent / skill level is uneven	Accountable named owner and escalation path; SOPs, guardrails, and disclaimers enforced in-system at both agent and sub-agent (skill) level
Security ⚑ data protection · IP containment · access controls · deploy and revoke	No clear data boundary; IP could leak to a model or downstream system; no one can say who can switch it off	Standard protections in place; gaps in IP containment or access control; revocation possible but manual and untested	Enforced data protection, IP contained, least-privilege access controls, tested deploy-and-revoke
Testing & Certification observability · audit trail · performance graded & re-certified	A black box: outputs are not logged or reviewed, and performance is never graded	Basic logging and spot checks; performance reviewed informally, with no re-certification schedule or trigger	A glass box: full observability and audit trail; performance graded against defined standards and re-certified on a schedule or based on a trigger, just like a human

⚑ **= floor component.** Security and Governance are safety-critical, so a strong total can't paper over a weak score on either. See the floor rule in the tier table below.

The tiers

Tier	RAIL Score	What it means
High	15–18	Trustworthy enough for use in Life Sciences
Medium	10–14	Specific gaps to close before use in Life Sciences
Low	Below 10	Not fit for use in Life Sciences
Floor rule	Security or Governance below 3	Caps the system at Low, regardless of the total

Recommended next steps

High

- Continue monitoring value, safety, and adoption.
- Re-certify the system on a defined schedule or trigger.
- Expand to additional workflows using the same governance model.
- Capture lessons learned to accelerate future deployments.

Medium

(both floor components met)

- Start with the lowest-scoring component, and agree an owner and date to close each gap.
- **Value Creation:** define one before/after measure against the human baseline, and take the baseline reading.
- **Human Centricity:** make human checkpoints on consequential actions explicit and enforced.
- **Orchestration:** close manual hand-offs so context carries between systems without re-entry.
- **Testing & Certification:** grade a sample of outputs against a defined standard, and set a re-certification schedule or trigger.
- Re-score within 30 days of closing the gaps.

Low: floor breach

(Governance or Security below 3)

- Fix the floor component first. The total doesn't count until both score 3.
- **Governance:** name an accountable owner and escalation path; enforce SOPs, guardrails, and disclaimers in-system at agent and sub-agent level.
- **Security:** set a clear data boundary, contain IP, apply least-privilege access, and test deploy-and-revoke.
- Re-score once both floor components reach 3.

Low: total below 10

- Define measurable business value before further investment.
- Address the lowest-scoring components first.
- Reassess the system only after critical deficiencies have been addressed.

NEXT For a system you're piloting, scope it with the **Agent JD Template**. For a system you're buying, put the vendor through the **Vendor-Grading Checklist**.