

NORTH STAR

Improve the quality and speed of Kirin Group executive decision-making so management time shifts toward higher-value, CSV-oriented growth decisions.

Overall

MINOR GAPS

CoreMate is a well-scoped advise-level agent: read-only tools, human decision ownership, and a clearly bounded workflow. The design holds together reasonably well at this autonomy level, but the success metrics do not yet measure what the North Star actually cares about, and the governance cells are thinner than the sensitivity of the data warrants. The biggest open question is whether 'shorter meeting time' drives the right behaviour or quietly crowds out the quality deliberation the whole system is meant to produce.

PARTLY COHERES

The design tells a consistent story from advise-level autonomy through read-only tools to human decision ownership — that part holds. But the North Star calls for better decision quality while Cell 01 includes a meeting-length metric that pulls the opposite way, and the governance cells (07, 08, 09) are thinner than the sensitivity of the data and the breadth of affected stakeholders in Cell 02 would call for.

DO THESE NEXT

Cell 01 Replace or reframe 'shorter meeting time' with a decision-quality outcome measure — for example, the share of proposals where executives report that risks were identified before the vote, assessed in a brief post-meeting survey.

Cell 07 Add a workflow-specific never-do list: at minimum, topics CoreMate must never introduce live (for example, named individuals' performance, undisclosed M&A targets), and conditions under which a generated question must be held back entirely rather than displayed.

Cell 08 State explicitly what persists after each meeting (and for how long), who can query that history, and how post-meeting executive feedback is collected and fed into improvement — one sentence each.

TOP RISKS

- 'Shorter meeting time' as a metric will push CoreMate toward brevity rather than decision quality, directly undermining the North Star.
- Ten years of board and executive-committee records sit as persistent institutional memory with no stated retention limits, access controls, or rules about which executives can see which materials.
- Cell 07 relies on the Kirin AI Policy in general terms but contains no workflow-specific prohibitions — for example, what CoreMate must never surface in a live meeting or when a question must be withheld entirely.

CROSS-CELL FLAGS

CELLS 01, NORTHSTAR

'Shorter meeting time' measures efficiency, but the North Star calls for higher-quality decisions; these can conflict.

An agent optimising for brevity will reduce the number of questions it surfaces, which shortens meetings but may cut the challenge that improves decisions.

CELLS 05, 04

Cell 04 requires source-grounded outputs but Cell 05 leaves external sources unspecified and unvetted.

Without named, trusted external sources, there is no way to enforce the grounding standard or to explain to an executive where a claim came from.

CELLS 08, 07

Sensitive board deliberations persist as institutional memory, but Cell 07 contains no rules about access, retention limits, or what must be forgotten.

Board-level records covering personnel, M&A, and strategy are among the most sensitive data a company holds; retaining them without explicit controls creates privacy and liability exposure.

CELLS 02, 09

Cell 02 names patients and society as affected stakeholders, but Cell 09 names no individual accountable for outcomes and no escalation path.

If a flawed question shapes a consequential decision that affects patients or the public, there is currently no named person responsible for the error.

THE CANVAS, CELL BY CELL

Cell 01 · Target Workflow & Success Metrics

NEEDS DETAIL

WHAT YOU WROTE

Group Executive Committee agenda preparation and live strategy-meeting deliberation. Success metrics: more new perspectives in remarks, more medium- and long-term discussion, shorter meeting time, and higher proposal readiness before the meeting.

COACH

'Shorter meeting time' is an efficiency metric that could drive the agent to trim discussion rather than improve its quality, which is what the North Star requires.

What would tell you, six months in, that the decisions coming out of these meetings were actually better — not just faster?

Cell 02 · Users & Stakeholders

LOOKS CLEAR

WHAT YOU WROTE

Direct users: executives, proposal originators, meeting facilitators, and the executive-committee secretariat. Affected stakeholders: operating companies, employees, customers, patients and consumers, shareholders, business partners, and society through Kirin's CSV commitments.

COACH

Of all the affected groups you listed, which one faces the most serious downside if CoreMate surfaces a flawed question in the room at the wrong moment?

Cell 03 · AI Role & Autonomy

LOOKS CLEAR

WHAT YOU WROTE

Autonomy: advise. CoreMate acts as a review and challenge partner: it surfaces questions, risks, missing viewpoints, and discussion points. It does not make binding management decisions; human executives decide.

COACH

Is there a scenario — a live meeting under time pressure, say — where executives might start treating CoreMate's framing as the decision rather than as a prompt for one?

Cell 04 · Performance Needs

NEEDS DETAIL

WHAT YOU WROTE

Outputs must be relevant, source-grounded, concise, and timely enough for pre-meeting idea sparring and live committee use. Noise must be low because only a small number of questions can usefully enter the room.

COACH

No concrete, testable bar is set for any of the four qualities named — relevance, grounding, concision, or timeliness — so there is no agreed standard for when CoreMate is good enough to use live.

What response time and what rejection rate for irrelevant questions would you need before you trusted this in a live committee meeting?

Cell 05 · Context & Knowledge

NEEDS DETAIL

WHAT YOU WROTE

Kirin public sources name 10 years of Board and Group Executive Committee data, other internal materials, management philosophy, proposal materials, and latest external information such as market trends, competitor moves, and legal changes.

COACH

'Latest external information' names no specific sources, so there is no way to know which external inputs are trusted or how currency is verified before a question enters the room.

Which external sources — and which providers or databases specifically — are authoritative enough that you would quote them to the Group Executive Committee without checking?

Cell 06 · Tools & Action Channels

LOOKS CLEAR

WHAT YOU WROTE

CoreMate reads internal and external knowledge sources, supports proposal-owner idea sparring before meetings, and displays selected discussion points or questions during meetings. No public source indicates autonomous external action.

COACH

Who decides, in the moment, which of CoreMate's generated questions actually appear on the screen during a live meeting, and what is their basis for choosing?

Cell 07 · Rules & Boundaries

NEEDS DETAIL

WHAT YOU WROTE

Must follow the Kirin Group AI Policy: human-centric, safe, fair, privacy-protecting, secure, transparent, and accountable. Must not make binding decisions, override human judgement, or present unchecked AI output as an executive decision.

COACH

The rules reference the Kirin AI Policy at a principles level but name no workflow-specific prohibitions — for example, topics CoreMate must not raise live, or conditions under which a question must be withheld entirely.

What is the one thing CoreMate must never surface in the room, even if it is technically relevant to the agenda?

Cell 08 · Memory & Learning

NEEDS DETAIL

WHAT YOU WROTE

Persistent institutional memory comes from prior board and executive-committee records and internal materials. External information is updated. The system improves through post-meeting executive feedback and continuing feature refinement.

COACH

Ten years of board and committee records are treated as persistent institutional memory with no stated access controls, session boundaries, or rules about what is retained versus discarded after a meeting.

After a meeting ends, what does CoreMate still hold, and who can query it?

Cell 09 · Oversight & Accountability

NEEDS DETAIL

WHAT YOU WROTE

Human executives own the decisions. Proposal originators use CoreMate before meetings; facilitators or the secretariat control how selected questions enter live discussion. Oversight sits within Kirin's AI Policy and management-governance structures.

COACH

Accountability sits with 'management-governance structures' rather than a named person, and no escalation trigger is stated for when CoreMate surfaces something that should not enter the room at all.

If CoreMate generates a question that is factually wrong or based on a misread of a proposal, who is responsible for catching it before it reaches the executives — and how?