Chapter 20. Group Visibility Mirror for problem solving

Introduction

The Group Visibility Mirror is designed to make two things visible at the same time: the problem a group is trying to solve, and the way the group is going about solving it. Poor group behaviour can damage the solution, while a difficult problem can expose weaknesses in the way the group works together.

AI can help make human judgement more visible, more evidence-based, more disciplined and more correctable by making evidence, assumptions, level changes, disagreement, bias, participation and drift visible while work is happening. The Group Visibility Mirror provides a second control-and-correction loop around the ordinary problem-solving process: while the group corrects the problem, the visibility system helps the group correct the way it thinks, interacts and solves problems.

The Group Visibility Mirror is not merely watching the group. It helps the group solve the substantive problem while making the human process visible enough in real time to improve it. Over time, that visibility should increase both the quality of decisions and the future capability of the people and the group.

How this chapter unfolds

1. **Start with the finished state -** Describes the mature form of the Group Visibility Mirror so the destination is clear before implementation is discussed.
2. **Individual visibility preconditions -** Sets out what each participant and their Personal Visibility Mirror need before group-level visibility can work safely and usefully.
3. **Group visibility preconditions -** Tests whether the group has the level, authority, knowledge, perspective and relational capability required for the problem.
4. **Define and locate the problem -** Uses CQQRT, From What to What, domain, level and time horizon to make sure the group is solving the right problem.
5. **Run the Group Visibility Mirror -** Shows how AI supports the problem and the human process at the same time, using private and group visibility with selective intervention.
6. **Use one disciplined problem-solving process -** Builds the Old World/New World comparison and the full ten-step process directly into the Group Visibility Mirror.
7. **What this means for humans -** Explains why visibility matters especially for people working on AI and how the personal, group, national and transnational pilots fit together.

8. **What this means for AI** — Shows how AI can provide the operating engine of the Group Visibility Mirror by holding the problem-solving process, group dynamics, permissions, memory and recursive learning together under human-set aims and authority.

1. Start with the finished state

It is useful to describe the mature version first. Actual implementation would be progressive because permissions, trust, Personal Visibility Mirrors and AI capability would develop over time. The finished state gives us a clear direction without pretending that every part can be switched on at once.

In the mature state, every participant and the group itself can see, in real time, both the problem being solved and the human process by which it is being solved. The group therefore works in two visible domains at once: problem performance and relational performance.

The finished-state Group Visibility Mirror allows a group to see the problem, the process being used to solve it, and the behaviour of the people participating in that process. AI provides permissioned feedback to individuals and to the group so that both the solution and the group's future capability can improve.

2. Individual visibility preconditions

The Group Visibility Mirror depends on the people inside it. Before a participant enters the mature group process, that person should have an operating Personal Visibility Mirror and a clear permissions arrangement. The purpose is not to expose the person to the group. It is to give the individual and their AI enough authorised visibility to support better judgement and better participation.

The Personal Visibility Mirror should hold the information that is relevant to the work, subject to explicit permission. That includes the person's aims and responsibilities; level, skills, knowledge and experience; behavioural and problem-solving tendencies; strengths and frailties; communication patterns; previous commitments and performance; and development objectives. Physiological or emotional information belongs in the system only where the individual has explicitly permitted it and where it is relevant to the work.

This visibility matters because people do not enter a group as neutral reasoning machines. They bring experience, expertise, preferences, confidence, fear, status concerns, habits, biases and different ways of responding under pressure. Some of those qualities strengthen the group. Others can distort judgement. The Personal Visibility Mirror gives the individual a better chance of seeing those tendencies while they are happening rather than only after the event.

Permissions are therefore fundamental. At minimum, the architecture must distinguish three layers: information that remains private to the person; information available to that person's AI; and information that the person has permitted to be used within the Group Visibility Mirror. These are not the same thing.

The Group Visibility Mirror does not mean that everybody can see everybody else's Personal Visibility Mirror. An individual's AI may know something privately and use it to coach the individual without disclosing it to the group. That separation is essential if visibility is to increase trust rather than destroy it.

3. Group visibility preconditions

Individual readiness is not enough. The group itself has to be capable of solving the problem. A highly visible process will not compensate for a team that lacks the necessary level, authority, knowledge or domain coverage. The first group-level visibility test is therefore whether the right people are present for the work that has to be done.

The group needs a shared understanding of the aim, enough decision authority to act, and a level of capability that matches the level of the problem. It also needs the required domain knowledge, technical skills, experience and resources. Where the problem crosses functions or systems, the group needs enough perspective to see those boundaries rather than treating the problem as though it belongs to only one part of the organisation.

Relational capability matters as well. A team can contain excellent technical people and still perform poorly if expertise is ignored, disagreement becomes personal, hierarchy blocks challenge, or people withdraw. Group visibility should therefore show not only who is in the room but what each person is there to contribute, where authority sits, what expertise is present, what is missing, and whether the group can work constructively enough to use what it knows.

AI should test this before the problem-solving process begins. It may show, for example, that the problem is operating at Level V while most of the team's experience and authority sits at Levels III-IV, or that an essential functional domain is missing. That is valuable visibility. Without it, the system could become very good at observing a group that was structurally unable to solve the problem.

Group composition and problem definition are also linked. The initial team defines the problem as well as it can; the emerging definition then tests whether the right people are present; and the group is changed or the problem reframed where necessary. The visibility process therefore begins before Step 1 and continues as the group learns.

4. Define and locate the problem

The problem needs an explicit starting definition. CQQRT and ‘From What to What’ provide a simple discipline: Context, Quantity, Quality, Resources and Time. Together they describe the current state, the required future state and the important constraints between them.

Instead of “Improve customer service”, the group might define the task as: “Reduce average complaint-resolution time from 72 hours to 24 hours within six months, while maintaining customer satisfaction above 85% within the existing operating budget.” The second statement gives the group something it can see, measure and later correct against.

The Group Visibility Mirror should also locate the problem before the group tries to solve it. The three tests are domain, level and time horizon. A problem that appears to sit in Sales may actually be a cross-functional P&L problem, a competitive-position problem or an industry-structure problem. A symptom at one level may be caused by a problem one level higher.

Problem definition remains provisional. Investigation may show that the original framing was wrong. The visibility system should preserve the path from initial problem, through evidence and learning, to reframing and a revised problem. That history makes the change in understanding visible rather than allowing the group to rewrite the story after the event.

5. Run the group visibility mirror

The Group Visibility Mirror operates on two tracks at the same time. The first track helps solve the problem. AI can support research, retrieval, calculation, hypothesis generation, scenario modelling, comparison, evidence testing, identification of missing information, contradiction detection, assumption testing, recording decisions and tracking the agreed process.

The second track makes the way the group is solving the problem visible. AI can identify withdrawal, ignored expertise, premature consensus, domination, repeated interruption, defended positions, confirmation bias, movement from evidence toward felt truth, unresolved disagreement, and interpersonal conflict contaminating substantive reasoning.

Level and domain errors are especially important. Groups can appear to be discussing one problem while participants have silently moved to different levels or different domains. The evidence then stops answering the same question, proposed solutions address different systems, and genuine analytical disagreement can be mistaken for interpersonal conflict. Group visibility should make those shifts explicit.

Private and group visibility

The visibility system needs two channels. The private channel allows AI to coach an individual without unnecessarily exposing that person to the group. It might point out repeated interruption, defended reasoning, unused expertise or a mismatch between confidence and experience. The group channel is reserved for matters that affect collective performance: unclear aims, competing problem definitions, unresolved hypotheses, weak evidence, missing expertise or drift away from the agreed process.

The distinction between private visibility and collective visibility is critical. Not every useful observation belongs in the room. Permissions determine what can be used, and good judgement determines where it should be used.

Selective intervention

A Group Visibility Mirror that comments continuously would become intolerable. AI should observe continuously but intervene selectively according to materiality, confidence and timing. A minor interruption can be ignored. A repeated pattern that damages the quality of the decision may justify private feedback. A structural failure affecting the whole group's reasoning may justify group feedback.

Two outputs and a learning loop

The process should produce two outputs. The problem output records what was decided, why, what evidence supports it, what remains uncertain, who does what next and how success will be measured. The visibility output records how well the group solved the problem together, where the process worked, where it broke down, what each person contributed, what recurring patterns appeared and what should change next time.

The second output feeds back into the Personal Visibility Mirrors and the Group Visibility Mirror, subject to permissions. The architecture therefore becomes recursive: Personal Visibility Mirrors -> Group Visibility Mirror -> problem solving -> outcome -> measurement -> reflection -> updated visibility -> next problem. The system is not mainly a meeting technology. It is a learning and behavioural-change architecture wrapped around disciplined problem solving.

6. Use one disciplined problem-solving process

The Group Visibility Mirror works best when the group agrees in advance to use one common problem-solving process. People remain free to disagree strongly about the substance, but they agree on how the disagreement will be worked through. That shared process becomes part of the group's institutional architecture.

Chapter 6 established the ten-step process for solvable, complex, levelled problems. Its first three steps are deliberately different from conventional problem solving. Before defining or analysing the problem, the group first recognises the condition that has caused it to act, locates the level at which the problem actually sits, and asks whether the people trying to solve it have the capability and authority to work at that level. Only then does the group define the problem and move through analysis, possible solutions, evaluation, choice, implementation and learning.

The Group Visibility Mirror makes the process itself visible while the group is using it. It locates the group in the ten steps, shows what evidence it is using, what assumptions have entered, makes visible if problem has silently shifted level or domain, and reflects whether the way people are working together is helping or damaging the quality of the solution.

Old World and New World

The ten steps do not fundamentally change. What changes is the information environment and the visibility and feedback architecture around them.

In the Old World, the group depends largely on what participants know, remember, disclose and notice. Research is episodic. The chair or team leader tries to maintain process discipline. Biases and interpersonal dynamics may be noticed, but often imperfectly and after the event. Much of the reasoning disappears when the meeting ends.

In the New World, AI is present throughout the process. It can draw on authorised Personal Visibility Mirrors, the Group Visibility Mirror, external evidence and the accumulating record of the problem. It can support individuals privately while supporting the group collectively. It can make both the substantive problem-solving process and the relational process visible, while intervening selectively according to materiality, confidence and timing.

The distinction is not Old World human problem solving versus New World AI problem solving. Humans remain responsible for aims, judgement, values, trade-offs and decisions. The change is from a largely episodic, partly visible and memory-dependent process to a continuously supported, evidence-rich, permissioned and correctable human problem-solving process.

Old World and New World problem solving. Exhibit 20.1

Ten-step process	Old World	New World — AI + Group Visibility Mirror
1. Recognise the Condition	The group responds to a failure, deterioration, risk or opportunity using what participants can immediately see. The condition may quickly become mixed with an assumed explanation.	Separates observed change from assumed cause, helping prevent the group from locking onto an explanation before the condition is clear.
2. Locate the Problem by Level	The problem is often accepted at the level where the symptom first appears. A functional issue may be treated as a functional problem even when its cause sits above or across functions.	Tests domain, level and time horizon, reducing the risk of solving the problem only where the symptom first appears.
3. Locate Yourself	Teams may begin work without testing whether they have the capability, knowledge or authority required at the level where the problem sits. Missing expertise or authority may emerge only after substantial work has been done.	Tests capability, knowledge and authority before work begins, exposing missing expertise or the need to escalate rather than discovering it late.
4. Define the Problem	The group frames the problem from the information already available. Important facts may sit in separate systems or in the memories of different participants. Assumptions can become embedded in the definition without being visible.	Defines the gap precisely and maintains the Fact Book, reducing ambiguity and preventing assumptions or felt truths from becoming accepted facts.
5. Analyse the Problem	Causal explanations can be shaped by hierarchy, experience, functional interests, selective evidence and early certainty. Groups may stop at the first plausible explanation.	Tests competing explanations and contradictory evidence, reducing the risk that hierarchy, bias or the first plausible explanation becomes the answer.
6. Generate Possible Solutions	Alternatives depend heavily on the experience and imagination of the people present. Groups can converge early around a familiar solution or the preferred idea of the most influential participant.	Broadens the option set and draws on unused expertise, reducing early convergence around familiar or influential ideas.
7. Evaluate the Alternatives	Options may be compared unevenly, with assumptions, second-order effects, risks and trade-offs receiving different levels of attention.	Applies the same criteria across options and makes assumptions, trade-offs and wider effects visible, reducing uneven or selective evaluation.
8. Choose the Solution	Judgement, evidence, negotiation, hierarchy and institutional preference all influence the decision. Dissent and uncertainty can disappear once the choice has been made.	Humans decide, while the Mirror preserves evidence, uncertainty and dissent so hierarchy or convenience cannot quietly erase the reasoning.
9. Implement the Solution	Decisions are translated into actions through normal management systems. Responsibilities, dependencies and slippage may become visible only during formal review.	Makes responsibilities, dependencies, milestones and slippage visible early, allowing correction before implementation problems become entrenched.
10. Review and Recalibrate	Reviews are often periodic and depend heavily on reporting quality, memory and willingness to acknowledge that the original diagnosis or solution may have been wrong.	Compares outcomes with the original aim and assumptions, reducing hindsight rewriting and turning success or failure into explicit learning.

The Ten-Step Process

1. Recognise the condition — Establish what has caused the group to act. Something may be deteriorating, failing or becoming risky, or there may be an opportunity to move to a much better condition. The Group Visibility Mirror helps separate the observed condition from early explanations of why it exists. At this stage the discipline is to establish what is happening and why attention is required without prematurely deciding what the problem is.

2. Locate the problem by level — Determine where the problem actually sits. The place where the symptom becomes visible may not be the place where its cause resides. The Group Visibility Mirror helps the group test domain, level and time horizon and move up and down the surrounding system before analysis becomes locked onto the wrong problem.

3. Locate yourself — Test whether the people attempting to solve the problem have both the capability to understand it and the authority to act at the level where it resides. The Group Visibility Mirror can make gaps in experience, knowledge, level, domain coverage and authority visible before the group proceeds. Where the fit is poor, the answer may be to change the group, bring in additional expertise, escalate the problem or redefine which part can legitimately be addressed locally.

4. Define the problem — State the problem precisely enough that it can be investigated, without assuming its cause. The definition should establish the relevant outcome, scope, level and time frame and make clear the gap between the present condition and the required one. CQQRT and From What to What provide useful disciplines for doing this.

This is also where the Group Visibility Mirror creates the **Fact Book**. The Fact Book is the structured shared record on which the problem-solving process will operate. It holds the relevant evidence and sources, separates facts from assumptions, interpretations and felt truths, records what remains uncertain and shows where important information is still missing. As the evidence develops, both the Fact Book and the problem definition can change. The important point is that the change in understanding remains visible rather than being rewritten after the event.

5. Analyse the problem — Determine what is causing the defined condition. Break the problem into its constituent parts, examine relationships, gather and test evidence and distinguish causes from symptoms. The Group Visibility Mirror helps keep competing explanations alive long enough to be tested. It can also make visible when hierarchy, confirmation bias, defended positions or changes of level and domain are beginning to distort the analysis.

6. Generate possible solutions — Develop a range of credible ways of acting on the causes that have been identified. The purpose is not to produce ideas for their own sake but to find different mechanisms that could move the system toward the required condition. The Mirror can help widen the range of alternatives, draw in unused expertise and expose premature convergence on a familiar or favoured answer.

7. Evaluate the alternatives — Compare the alternatives against the aim, evidence, constraints and likely consequences. Effectiveness, feasibility, cost, risk, speed, resources, reversibility and effects on neighbouring parts of the system all matter. The Group Visibility Mirror helps ensure that the same questions are applied across the alternatives and that assumptions, dependencies and trade-offs remain visible rather than disappearing inside advocacy for a preferred solution.

8. Choose the solution — Decide which option or combination of options best addresses the problem given the evidence, aims, constraints, authority and acceptable risk. The decision remains human. The Group Visibility Mirror records the reasoning, assumptions, trade-offs, uncertainty and disagreement that led to it. It also records what success should look like and what evidence would justify reopening the decision.

9. Implement the solution — Translate the chosen solution into coordinated action. Responsibilities, resources, timing, milestones, processes and measures have to become explicit. The Group Visibility Mirror helps keep commitments and dependencies visible and makes emerging deviations available early enough for correction. Implementation therefore remains connected to observation rather than becoming a period in which the group simply waits to see whether the original decision worked.

10. Review and recalibrate — Compare what happened with what was intended and determine what should now change. The group asks whether the intervention worked, whether unexpected consequences appeared, whether its original assumptions were correct and whether the environment changed during implementation. The Group Visibility Mirror preserves that history and feeds authorised learning back into the Personal and Group Visibility Mirrors. Review may result in a small correction, a different solution, a changed problem definition or a new cycle of problem solving at a different level of problem.

AI also acts as a custodian of the agreed process. It can make visible when the group begins solving before the problem has been located, moves from condition to assumed cause, defines the problem before checking its own capability and authority, skips an important hypothesis, ignores contradictory evidence or converges while major uncertainty remains.

The ten steps give the group enough common structure to know where it is, what it has established and what it has skipped. The Group Visibility Mirror then adds visibility around both the substantive work and the way the group is doing that work. Together they allow the group to solve the immediate problem while becoming better at solving the next one.

7. What this means for humans

The Group Visibility Mirror is ultimately about human beings. AI can increase visibility, but humans still bring the aims, values, judgement, relationships and institutional authority that determine what happens next. That makes self-understanding and group understanding part of the problem-solving architecture rather than an optional extra.

This matters especially for people working on AI. They are building systems with rapidly increasing capability while working inside laboratories shaped by competition, status, deadlines, investment, national pressure and strong technical cultures. They therefore need visibility of their own brains as well as visibility of the technology. Threat responses, felt truths, bias, defended identity, copying, overconfidence, deference to expertise and the pull of immediate advantage can all affect technical and safety decisions.

They also need visibility of institutional architecture. A technically excellent group can still produce difficult outcomes if authority is unclear, challenge is weak, incentives reward speed more strongly than correction, or responsibility sits below the level of the consequence. The Group Visibility Mirror should help make those institutional conditions visible while decisions are being made.

The pilots provide the practical path. The Personal Visibility Mirror helps the individual see their own patterns. The Group Visibility Mirror helps teams see how they are thinking and working together. National Visibility Mirrors help make larger institutional and societal tensions visible. Transnational Visibility Mirrors extend the same logic above the nation where problems, especially AI problems, cross borders. These are not separate ideas. They are the same visibility architecture operating recursively at different levels.

For AI laboratories, the implication is immediate. The visibility pilots should begin to be used inside the work: first in bounded teams and decisions, then across larger organisational questions as permissions, trust and capability mature. The purpose is not surveillance. It is to give the people building powerful systems a better chance of seeing themselves, their groups and their institutions clearly enough to make better decisions.

8. What this means for AI

The operating engine. The Group Visibility Mirror is something AI can help build and operate. Humans set the aims, permissions and boundaries. AI provides much of the engine that holds the information together, follows the problem-solving process, watches the group dynamics and keeps the whole system visible while the work is happening.

Scope of understanding. The AI engine has to understand the people and the group well enough to recognise patterns of threat, status, confidence, withdrawal, domination, ignored expertise, premature agreement, felt truth and changes of level or domain. With permission, it can also draw on Personal Visibility Mirrors so that individuals can receive private guidance while the group sees only what is relevant to collective performance.

The ten-step process. AI can also hold the ten-step problem-solving process together. It can help the group recognise the condition, locate the problem and the problem solver, build and maintain the Fact Book, test explanations, widen alternatives, preserve disagreement, track implementation and compare what happened with what was intended. The important change is that much of this can happen continuously rather than being reconstructed afterwards.

Recursive learning. The Group Visibility Mirror is recursive. What happens in one problem-solving cycle can improve the next one. With permission, learning can flow back to the individual, the group and the wider organisation. AI can remember what the group tried, where it went wrong, which patterns keep recurring and whether the group is actually getting better at solving problems together.

The larger possibility. AI can become the operating engine of a group that can increasingly **see itself while it works**. If that visibility is well designed, permissioned and kept under human aims and authority, the Group Visibility Mirror could become a continuing system for better judgement, better behaviour and better collective problem solving.