



AI Presentation Coach

Using AI to make you a better public speaker

COACHING REPORT


AUGUST 28, 2026

Speaker: Michael Geiser (mgeiser@mgeiser.net)

File: 1_Introducing AgentCore.mp3

Presentation: logo fix test



Description: One more last run?



Your presentation is competent with a solid foundation. Prioritizing improvements in audience engagement will yield the greatest score gains.

TALK TIMELINE

Presentation duration: 5 min 06 sec



Effective starts at 6.5 ↑

High: 9

Medium: 21

Low: 5

Strengths: 21

MEASURED, NOT JUDGED

Speaking Rate	234 WPM	Target Range	130–170 WPM
Filler Words	2	"So" Openers	1
Pauses > 1s	1	Longest Run	177.8s
Close Share	15.0%	Enunciation	100%

You scored 6.0 out of 10, placing you in the Competent band. Your three highest-leverage improvements are: Audience Interaction, Storytelling, and Rhetorical Questions. Addressing these will create measurable progress toward the next performance band.

Executive Summary

Your presentation scores a 6.0 out of 10, placing you in the Competent band — and that designation is accurate in the truest sense: you are technically capable, clearly knowledgeable, and able to hold an audience's attention without losing them. But Competent also means you are leaving a significant amount of impact on the table. The gap between where you are and where you could be is not a gap in knowledge or preparation — it is a gap in performance craft. You know this material deeply, and that expertise is not yet translating into the kind of commanding, memorable delivery that moves an audience from interested to convinced.

What you do well is substantial and worth naming clearly. Your technical command is genuine and immediately credible — you speak with authority about production-scale challenges, framework compatibility, and infrastructure abstraction without stumbling or hedging. Your use of second-person language creates a naturally conversational, inclusive tone that makes complex content feel approachable, and your instinct to frame every feature against a real developer pain point is exactly the right strategic move. The prototype-to-production narrative arc you use as a unifying thread is smart structural thinking, and your opening ten seconds — establishing who you are, what the product does, and why it matters — is one of the strongest moments in the entire presentation. Your Delivery and Technical Communication scores, both at 6.8, confirm that your foundational communication skills are already working.

The two or three changes most likely to move your overall score are storytelling, pause discipline, and eliminating your over-reliance on filler transitions. Storytelling is the single highest-leverage investment because its absence is felt across your Audience Engagement, Persuasion, and Executive Presence scores simultaneously. Right now your presentation is almost entirely descriptive — a feature walkthrough dressed up with good language. There is not one customer scenario, one before-and-after moment, or one concrete human story that shows what actually happens when a team uses Agent Core in production. Without that narrative anchor, claims like "generating real business value quickly" float without weight. Equally damaging is your near-total absence of pauses — your rhetorical questions at around 223 and 226 seconds land with zero silence after them, which functionally cancels their persuasive purpose. And the word "So" as a sentence opener appears more than 25 times across a five-minute presentation, which signals stream-of-consciousness delivery and chips away at the executive gravitas your content deserves.

These weak dimensions are not isolated problems — they are compressing your ceiling across every other area. The lack of storytelling suppresses your Audience Engagement score, but it is also the primary reason your Persuasion dimension cannot climb despite your solid pain-point framing and your strong live-demo CTA, because without evidence or narrative, the emotional arc never builds. The absence of pauses and the monotonous "So"-driven cadence are simultaneously dragging your Pacing, Delivery, and Executive Presence scores, because they create a flat, single-register delivery that makes your genuine enthusiasm — phrases like "super, super easy" and "pretty cool if you ask me" — land as informal noise rather than earned energy. Fix the pauses and the transitions, and your Delivery, Pacing, and Executive Presence scores all rise together without changing a single word of your actual content.

For your next recording, make these specific changes in the first 30 seconds: replace your opening with a one-sentence micro-story before your value proposition — something like "Last year, a team I worked with had a working AI prototype that collapsed the moment they tried to move it to production — that's exactly the problem Agent Core was built to solve." Then state your thesis. Remove the word "So" as a sentence opener entirely — replace it with either silence or a purposeful connector like "Here's why that matters" or "The second capability builds on this." When you ask your two rhetorical questions about memory and conversation history, stop talking for a full two seconds after each one. Let the question sit. That silence is not dead air — it is the moment your audience actually absorbs the problem you are about to solve.

Overall Coaching Assessment

You are solidly in competent territory across the board, which tells a clear story: you have genuine foundations to build on, but you have not yet broken through to the next level of presenter authority. Your delivery and technical communication are already operating at an effective level, and that is not a small thing — your vocal clarity, your command of complex material, and your ability to make agentic AI accessible to a broad audience are real assets that many experienced presenters lack. Protect those strengths fiercely as you develop. They are the reason people trust what you say.

Your highest-return focus right now is executive presence, and specifically the gap between how much you know and how much authority you project. The casual opening, the absence of a verbal roadmap, and the double-coverage of capabilities all send the same unintentional signal: that the presentation is being discovered in real time rather than led with intention. Fixing this — combined with deliberate use of strategic silence and pace variation at key moments — will do more for your overall score than any other investment you can make right now.

The coaching philosophy here is about closing the credibility gap. Your content credibility is already high. What is lagging is structural and presence credibility — the signals that tell an audience they are in the hands of someone who has thought carefully about their experience. Once those two things align, you will feel like a fundamentally different presenter.

Your Three Moves

Your three highest-leverage improvements are interconnected. Addressing Audience Interaction will create a foundation for the other changes. Working on Storytelling reinforces that foundation and amplifies your impact. Finally, improving Rhetorical Questions ties everything together, helping you deliver a more cohesive and compelling presentation.

+5.2 Audience Interaction

Insert 1–2 interactive checkpoints, such as 'Raise your hand if you've already built an agent prototype' or a quick poll asking which framework (LangChain, CrewAI, etc.) attendees currently use. This directly leverages the topic of framework flexibility. The presenter uses inclusive 'you' language extensively throughout ('get you up and running,' 'your agentic AI prototypes,' 'your organization'), which creates a conversational tone. However, there are zero direct interactive elements — no polls, no check-ins, no invitations to respond or participate actively.

Lifts: Delivery

+5.2 Storytelling

Open with a brief story: 'Imagine your team just built a brilliant AI agent in a weekend hackathon — it works great on your laptop, but the moment you try to scale it for 10,000 users, everything breaks. That's exactly the problem Agent Core was built to solve.' This creates an immediate emotional hook. The presentation is almost entirely descriptive and feature-listing in nature. There are no personal anecdotes, no customer stories, no 'before-and-after' scenarios, and no concrete narrative examples of what a real team experienced going from prototype to production. The closest attempt is a generic mention of 'generating real business value for your organization.'

Lifts: Delivery

+5.2 Rhetorical Questions

Increase the use of rhetorical questions to provoke thought before each major capability section. For example, before introducing Gateways: 'Have you ever spent days trying to connect your agent to a legacy API that has no MCP support? What if that took five minutes instead?' This pattern of question-then-answer is a powerful attention retention device. There are two genuine rhetorical questions in the presentation: 'How do I actually keep track of these conversations over time?' and 'How do I have a conversation with a user that can actually reference back to earlier conversations?' and one mild engagement check ('right?'). These are functional but limited in frequency and impact.

Lifts: Delivery

Do not lose these while you fix the rest

- Highly conversational and approachable delivery style that makes complex technical content (serverless deployment, MCP servers, agentic AI) accessible to a broad audience, using relatable analogies like 'nitty gritty,' 'fuss around,'.
- Excellent articulation and projection throughout the presentation — transcription confidence scores are consistently near 1.0 for the vast majority of words, including difficult technical acronyms like ECR, ECS, SDK, and.
- Effective use of repetition and signposting ('So,' 'Big one is memory,' 'Right after that') to guide listeners through a dense technical overview, helping maintain audience orientation across multiple capability categories.
- Natural enthusiasm and genuine passion for the subject matter is evident through expressive language choices and the informal, engaging tone, which helps sustain listener interest during a technically dense introduction.

Projected impact scores are estimates derived from rubric weightings and do not guarantee exact score increases.

4.8

/10

Competent

Audience Engagement

Your audience engagement is competent but has 2 high-priority areas to strengthen.

Weakest dimension · 2 high-severity findings

What is working

- Consistent and natural use of second-person 'you' language creates an inclusive, conversational tone that makes the audience feel directly addressed and personally relevant throughout the presentation.
- The presenter effectively frames the value proposition around audience pain points (e.g., production deployment challenges, memory management at scale, authentication complexity), demonstrating an understanding of what the audience cares about.
- Informal, relatable vocabulary ('mishmash,' 'nitty gritty,' 'nasty problem,' 'fuss around') lowers the perceived barrier to a complex technical topic and keeps the tone accessible to a broad developer audience.

HIGH Storytelling

The presentation is almost entirely descriptive and feature-listing in nature. There are no personal anecdotes, no customer stories, no 'before-and-after' scenarios, and no concrete narrative examples of what a real team experienced going from prototype to production. The closest attempt is a generic mention of 'generating real business value for your organization.'

Try: Open with a brief story: 'Imagine your team just built a brilliant AI agent in a weekend hackathon — it works great on your laptop, but the moment you try to scale it for 10,000 users, everything breaks. That's exactly the problem Agent Core was built to solve.' This creates an immediate emotional hook.

Effort: moderate · Impact: high

HIGH Q&A Handling

There is no Q&A segment, no invitation for questions, and no moment where the presenter anticipates or addresses likely audience questions. The presentation ends abruptly with a transition to a demo ('let's get started') without any opportunity for audience input or clarification.

Try: Even in pre-recorded or structured course formats, explicitly acknowledge anticipated questions: 'You might be wondering — does this work with my existing Terraform setup? Great question, and we'll address that in Module 3.' Alternatively, close the intro segment with: 'Before we dive into the demo, what questions do you have so far? Drop them in the chat and I'll address them as we go.'

Effort: moderate · Impact: high

MEDIUM Rhetorical Questions

There are two genuine rhetorical questions in the presentation: 'How do I actually keep track of these conversations over time?' and 'How do I have a conversation with a user that can actually reference back to earlier conversations?' and one mild engagement check ('right?'). These are functional but limited in frequency and impact.

Try: Increase the use of rhetorical questions to provoke thought before each major capability section. For example, before introducing Gateways: 'Have you ever spent days trying to connect your agent to a legacy API that has no MCP support? What if that took five minutes instead?' This pattern of question-then-answer is a powerful attention retention device.

Effort: moderate · Impact: medium

MEDIUM Eye Contact References

As this is a recorded/transcribed audio presentation, there is no evidence of direct visual engagement cues such as scanning the room, pausing for visual acknowledgment, or referencing audience reactions. The delivery is uniform and unidirectional throughout the entire ~5-minute segment.

Try: Even in recorded or online formats, simulate eye contact engagement by periodically addressing the viewer directly: 'I want you to look at this diagram specifically' or 'Keep this screen in mind as we move forward.' If presenting live, practice deliberate pausing to scan different parts of the room after key statements.

Effort: moderate · Impact: medium

MEDIUM Attention-Holding Strategies

The presenter uses informal language effectively ('mishmash of stuff,' 'fuss around,' 'nitty gritty,' 'nasty problem,' 'that's pretty cool if you ask me') which maintains a conversational, accessible tone. However, the delivery lacks structural variety — it is a continuous information dump without surprises, humor, dramatic pauses, or vivid analogies beyond surface-level ones.

Try: Introduce at least one memorable analogy or surprising contrast to anchor a key concept. For example: 'Think of Agent Core like a Swiss Army knife for deployment — your agent is the blade you choose, and Agent Core is everything else that makes it actually usable in the field.' Also consider using a deliberate dramatic pause after 'It's really cool.' to let that land before diving into details.

Effort: moderate · Impact: medium

Practice Drill 5 min

Practice drill for Audience Engagement: Record yourself for the allotted time focusing on 'Storytelling'. Open with a brief story: 'Imagine your team just built a brilliant AI agent in a weekend hackathon — it works great on your laptop, but the moment you try to scale it for 10,000 users, everything breaks. That's exactly the problem Agent Core was built to solve.' This creates an immediate emotional hook. Play back the recording and note improvements.

5.2

/10

Competent

Executive Presence

Your executive presence is competent but has 3 high-priority areas to strengthen.

Rank 2 of 7 · 3 high-severity findings

What is working

- Deep and accurate technical command of the subject matter, including correct identification of production-scale challenges (memory persistence, container management, authentication), competing frameworks, and AWS service integrations — this expertise is the foundation of genuine credibility.
- Clear value proposition articulation: the core message of 'taking agentic AI prototypes to production' is communicated consistently and accessibly, and the business outcome ('generating real business value for your organization quickly') is explicitly named.
- Effective use of contrast to explain the product: consistently framing Agent Core against the problems it solves (e.g., local SQL database vs. production-scale memory, manual Docker/ECR vs. serverless abstraction) demonstrates strong explanatory logic.

HIGH Gravitas

Informal and casual language significantly reduces gravitas. The presenter uses phrases like 'super, super easy,' 'fuss around,' 'nitty gritty,' 'nasty problem,' 'that's not gonna cut it,' 'pretty cool if you ask me,' and 'really cool' multiple times. While these may foster relatability in a teaching context, they erode the weight and seriousness expected in an executive-level presentation.

Try: Calibrate language to the audience. For executive audiences, replace 'super, super easy' with 'dramatically reduces operational complexity' and 'nasty problem' with 'a well-documented production challenge.' Reserve casual language intentionally and sparingly for emphasis only.

Effort: moderate · Impact: high

HIGH Leadership Demeanor

The presenter frames the session as instructional rather than visionary. Phrases like 'we're just going to start with,' 'let's get started,' and 'we'll see how that works' position the speaker as a guide rather than a thought leader. There is no articulation of a strategic vision, no call to action framed in business terms, and no expression of a point of view on where agentic AI is heading.

Try: Open with a bold strategic statement that establishes a point of view. For example: 'The gap between an AI prototype and a production-ready agent has killed more initiatives than poor models ever have. Agent Core closes that gap.' This immediately positions the speaker as a leader, not just an instructor.

Effort: moderate · Impact: high

HIGH Credibility Signals

The presenter offers no data points, no customer case studies, no quantified business outcomes, and no personal experience references to anchor credibility. The statement 'I'm sure there'll be more as time goes on' regarding built-in tools actually introduces speculative uncertainty where a confident observation would serve better.

Try: Introduce at least one credibility anchor early: a quantified claim (e.g., 'Teams report reducing deployment complexity by over 60%'), a reference to a real-world implementation pattern, or a personal observation from hands-on work with the platform. Replace speculation ('I'm sure there'll be more') with framed confidence ('Amazon has a pattern of rapidly expanding these toolsets — expect the libr

Effort: moderate · Impact: high

MEDIUM Composure

Two filler words 'uh' are detected in the transcript (e.g., 'It also has built-in, uh, tools for handling memory' and 'The, uh, agent SDKs'). While infrequent, they occur at moments of topic transition, suggesting slight hesitation when moving between points.

Try: Practice transitions between capability topics using deliberate pauses rather than filler words. A brief, confident silence is far more authoritative than an 'uh.' Rehearse the segue from deployment features to built-in tools as a distinct transition point.

Effort: moderate · Impact: medium

MEDIUM Composure

The structure of the presentation is repetitive in its use of 'So' as a transitional connector. The word 'So' appears as a sentence opener dozens of times throughout the transcript, creating a monotonous rhythm and suggesting stream-of-consciousness delivery rather than a prepared, structured narrative.

Try: Replace repeated 'So' transitions with purposeful signposting language such as 'More critically,' 'Building on that,' 'This is where it gets strategically important,' or 'The third pillar is.' This creates the impression of a deliberate, well-architected narrative rather than improvised flow.

Effort: moderate · Impact: medium

Practice Drill 5 min

Practice drill for Executive Presence: Record yourself for the allotted time focusing on 'Gravitas'. Calibrate language to the audience. For executive audiences, replace 'super, super easy' with 'dramatically reduces operational complexity' and 'nasty problem' with 'a well-documented production challenge.' Reserve casual language intentionally and sparingly for emphasis only. Play back the recording and note improvements.

5.8

/10

Competent

Persuasion

Your persuasion is competent but has 1 high-priority area to strengthen.

Rank 3 of 7 · 1 high-severity finding

What is working

- Clear identification of audience pain points: the presenter effectively names specific developer headaches (Docker, ECR, ECS, memory management at scale, authentication) that resonate with a technical audience, making the solution feel relevant and timely.
- Strong framework-agnostic positioning: emphasizing compatibility with OpenAI Agents SDK, LangChain, LangGraph, CrewAI, and AWS Strands is a differentiated and persuasive selling point that directly addresses a common objection about vendor lock-in, and it is articulated with concrete examples.
- Practical hands-on CTA within the session: the immediate transition to a live build-and-deploy demo is an effective persuasion technique, allowing the audience to experience the product's ease firsthand rather than simply being told about it.

HIGH

Evidence Usage

The presentation contains zero quantitative evidence — no performance metrics, deployment time comparisons, cost savings, user adoption statistics, or case studies. Claims like 'start generating some real business value quickly' and 'super easy' are entirely unsupported by data. The presenter acknowledges a limitation ('as of the time of this recording, there's only a couple') which inadvertently weakens the built-in tools argument.

Try: Introduce at least 2-3 concrete data points: e.g., 'Deploying with Agent Core takes 3 steps vs. 15+ with manual container setup,' or a customer testimonial. Show a before/after deployment complexity comparison to make the ease-of-use claim tangible.

Effort: moderate · Impact: high

MEDIUM

Call-to-Action Clarity

The immediate call-to-action is clear and actionable: 'let's get started' followed by a specific hands-on sequence (build an OpenAI Agents SDK prototype, then deploy it with Agent Core). This practical, step-by-step framing gives the audience a concrete next action. However, there is no broader CTA for the audience to adopt Agent Core in their own organizations beyond the course context.

Try: Add an organizational CTA beyond the course activity: for example, direct the audience to try Agent Core on one of their existing agent prototypes within the next week, and link to a getting-started resource or free tier offer to reduce friction.

Effort: moderate · Impact: medium

MEDIUM Emotional Appeal

The presenter uses some conversational enthusiasm ('It's really cool,' 'pretty cool if you ask me,' 'visceral feeling') and light humor ('mishmash of stuff') to create an approachable tone. The appeal to developer frustration ('you don't have to fuss around with Docker or ECR or ECS') resonates with the target audience's known pain points. However, emotional engagement remains shallow and does not build a compelling narrative arc.

Try: Create a mini-narrative around a relatable developer pain point: e.g., describe the experience of spending weeks configuring infrastructure only to have a prototype never reach production. Use this as the emotional hook before presenting Agent Core as the solution, making the audience feel understood before persuading them.

Effort: moderate · Impact: medium

MEDIUM Logical Structure

The presentation follows a broadly logical flow: introduction of the problem (prototype-to-production gap) → what Agent Core is → capability breakdown → hands-on demo preview. However, the middle section becomes repetitive and circular, with the framework flexibility point reiterated multiple times ('you're not tied to... you're not even tied to...') and the same benefit re-explained from slightly different angles, diluting momentum.

Try: Restructure the capability overview using a clear framework such as 'Problem → Agent Core Solution → Benefit' for each feature. Eliminate redundant repetition of the framework-agnostic point, which is stated at least four times, and instead use that time to deepen coverage of more complex features like memory and gateways.

Effort: moderate · Impact: medium

MEDIUM Credibility

The presenter self-identifies as Frank Kane but provides no credentials, professional background, or indicators of expertise with AWS or AI systems. There are no citations, third-party validations, or references to documentation. A self-deprecating phrase ('mishmash of stuff') and hedging language ('I'm sure there'll be more as time goes on,' 'that does seem to get a little bit more support') slightly undermine authority. The casual, conversational tone fits the educational context but does not

Try: Open with a brief credibility statement (e.g., 'As an AWS instructor who has worked with agentic AI deployments across X organizations...') and reference official AWS documentation or blog posts at key points. Replace hedging phrases with more confident language where the facts support it.

Effort: moderate · Impact: medium

Practice Drill 5 min

Practice drill for Persuasion: Record yourself for the allotted time focusing on 'Evidence Usage'. Introduce at least 2-3 concrete data points: e.g., 'Deploying with Agent Core takes 3 steps vs. 15+ with manual container setup,' or a customer testimonial. Show a before/after deployment complexity comparison to make the ease-of-use claim tangible. Play back the recording and note improvements.

6.2

/10

Competent

Pacing

Your pacing is competent but has 2 high-priority areas to strengthen.

Rank 4 of 7 · 2 high-severity findings

What is working

- Clear and energetic opening that immediately establishes the topic, the presenter's identity, and the value proposition within the first 10 seconds, setting a strong hook for the audience.
- Consistent use of forward-referencing phrases ('we'll see how that works in a moment,' 'we'll talk about there,' 'right after that') creates narrative continuity and maintains listener engagement across the presentation.
- Natural section transition at ~178-180s (identifiable 1.4-second gap) between the conceptual explanation and the capabilities overview demonstrates an instinctive sense for structural pacing at major topic shifts.

HIGH Appropriate Pauses

Measurable pauses are rare and brief. The longest detectable gap is approximately 1.4 seconds between segment 7 ending (~178.69s) and segment 8 beginning (~180.02s), coinciding with a natural section transition. Most sentence-level gaps are under 0.5 seconds. Critical rhetorical questions at ~223s ('How do I actually keep track of these conversations over time?') and ~226s ('How do I have a conversation with a user...?') receive no pause after them, undermining their intended impact.

Try: Insert deliberate 1-2 second pauses after rhetorical questions to let the audience reflect. Add a pause before and after the 'Big one is memory' statement (~220s) for emphasis. Pause at major section transitions (e.g., between the capabilities overview and the 'let's get started' conclusion) to signal structural shifts.

Effort: moderate · Impact: high

HIGH Content Density

Segment 6 (~119s to ~149s, 30 seconds) introduces the starter toolkit, testing shortcuts, AND additional tools/capabilities in rapid succession — roughly one new concept every 6-7 seconds. Similarly, segments 8-11 (~180s to ~268s) cover identity, authentication, built-in tools, memory (short-term and long-term), and gateways in approximately 88 seconds. This clustering of 6+ distinct features without pauses risks cognitive overload.

Try: Break the capabilities section into clearly demarcated sub-sections, each with a brief pause and verbal signpost (e.g., 'The first capability is identity management... [pause]. Second, let's talk about authentication...'). Consider splitting the capabilities overview into a dedicated segment rather than rushing through all features in a single continuous block.

Effort: moderate · Impact: high

MEDIUM Rhythm

The presentation relies heavily on repetitive transitional phrases — 'So,' appears as a sentence starter over 25 times throughout the ~307-second presentation (approximately every 12 seconds on average). This creates a monotonous cadence that dulls the natural flow. Additionally, filler words like 'uh' appear at timestamps ~159s and ~240s, disrupting rhythm at key points.

Try: Vary transitional language by using alternatives such as 'Additionally,' 'Building on that,' 'Here's why that matters,' or 'Let's now turn to...' Reduce filler words ('uh') through rehearsal or by inserting deliberate silent pauses instead.

Effort: moderate · Impact: medium

MEDIUM Speed Variation

Analysis of word density across audio segments reveals a largely consistent speaking pace throughout the entire presentation — roughly 140-155 words per minute across all segments. There is minimal slowing for high-importance moments (e.g., 'Big one is memory' at ~220s is stated briefly without significant deceleration) and no noticeable acceleration to convey excitement about key differentiators like framework agnosticism.

Try: Deliberately slow down when introducing key concepts (e.g., 'memory,' 'gateways,' 'agent runtime environment') to signal their importance. Consider speeding up slightly during list enumerations (Docker/ECR/ECS, framework names) to maintain energy while saving time for deeper concepts.

Effort: moderate · Impact: medium

MEDIUM Segment Balance

The presentation has an identifiable three-part structure: Introduction/hook (0-35s, ~11%), Core explanation/What is Agent Core (35-180s, ~47%), Capability deep-dive and wrap-up (180-307s, ~42%). The introduction is proportionally short for a conceptual overview course, and the transition to a conclusion ('So let's get started') at ~294s is abrupt — the closing section covering CloudWatch observability and the upcoming demo plan is compressed into roughly 26 seconds (~8.5% of total time).

Try: Expand the introduction slightly (15-20% of total time) to better orient learners. Give the concluding 'what comes next' section more breathing room — at least 45-60 seconds — to properly set up the hands-on activity and leave learners with a clear mental model of what they are about to do.

Effort: moderate · Impact: medium

Practice Drill 5 min

Practice drill for Pacing: Record yourself for the allotted time focusing on 'Rhythm'. Vary transitional language by using alternatives such as 'Additionally,' 'Building on that,' 'Here's why that matters,' or 'Let's now turn to...' Reduce filler words ('uh') through rehearsal or by inserting deliberate silent pauses instead. Play back the recording and note improvements.

6.2

/10

Competent

Structure

Your structure is competent but has 1 high-priority area to strengthen.

Rank 5 of 7 · 1 high-severity finding

What is working

- The opening hook is effective — the presenter immediately establishes who they are, what the technology does, and why it matters ('take agentic AI prototypes and turn them into production-level solutions'), giving the audience a clear reason to pay attention.
- Forward-referencing signposts ('we'll see how that works in a moment,' 'we'll talk about there,' 'let's get started') are used consistently throughout, creating a sense of continuity and momentum that keeps the audience engaged.
- The concrete value framing — positioning Agent Core as solving real production problems (scale, authentication, memory, infrastructure management) rather than just describing features — gives the content practical relevance and audience-centered focus.

HIGH Conclusion Impact

The presentation ends abruptly with 'Right after that, we're going to go ahead and deploy that using Agent Core and see how easy that is to do.' This is a preview of what comes next in the course, not a conclusion for this segment. There is no summary of key takeaways, no reinforcement of the core message, and no memorable close for this introductory segment.

Try: Add a brief closing summary before transitioning to the hands-on section: 'So to summarize — Agent Core solves the production deployment problem for AI agents. It works with any framework, abstracts away infrastructure, and provides built-in capabilities for memory, identity, tools, and observability. Now let's see all of that in action.' This gives the conceptual section a satisfying close.

Effort: moderate · Impact: high

MEDIUM Logical Flow

The presentation broadly follows a sensible sequence: problem statement → what Agent Core is → deployment features → framework flexibility → starter toolkit → capabilities deep-dive → hands-on plan. However, the deep-dive on capabilities (identity, authentication, built-in tools, memory, gateways, runtime, observability) is partially covered twice — once in passing during the 'What is Agent Core' section and again in the 'at a high level let's look at the capabilities' section — creating redundancy.

Try: Consolidate the capabilities discussion into a single, dedicated section. In the introductory 'What is Agent Core' section, name the capability categories briefly without detailing them, then give the full treatment in the dedicated capabilities walkthrough that follows.

Effort: moderate · Impact: medium

MEDIUM Organization

The content can be grouped into roughly four natural sections: (1) hook/value proposition, (2) what Agent Core is at a high level, (3) capability details, and (4) hands-on preview. While these sections exist implicitly, they are not formally demarcated. The transition from 'What is Agent Core' into framework flexibility, then into the starter toolkit, and back into capability details feels meandering rather than structured.

Try: Introduce explicit section markers — e.g., 'Now let's talk specifically about framework flexibility,' 'Moving on to the starter toolkit,' 'Now I want to walk through each capability in detail' — to signal clear section boundaries to the audience.

Effort: moderate · Impact: medium

MEDIUM Transitions

Transitions between sub-topics are predominantly additive ('It also...', 'And it also includes...', 'So...') rather than purposeful connectors that explain why the next point follows. The phrase 'So, at a high level, let's look at the capabilities' (~180s) is the only notable structural transition in the presentation.

Try: Replace generic connectors with signposted transitions that carry the audience forward, e.g., 'Now that you understand what Agent Core does at a high level, let me break down each individual capability' or 'With that foundation in place, let's see how the starter toolkit makes all of this even more accessible.'

Effort: moderate · Impact: medium

MEDIUM Signposting

The presenter uses forward-referencing phrases like 'we'll see how that works in a moment,' 'We'll see that in a moment,' and 'we'll talk about there,' which help maintain continuity. However, there are no backward-referencing summaries (e.g., 'So to recap what we've covered so far...') and no numbered list signposting (e.g., 'There are five capabilities — the first is...').

Try: When enumerating capabilities (identity, authentication, built-in tools, memory, gateways, runtime, observability), explicitly number or list them upfront: 'Agent Core has seven key capabilities — let me walk through each one.' This helps the audience track where they are in the enumeration.

Effort: moderate · Impact: medium

Practice Drill 5 min

Practice drill for Structure: Record yourself for the allotted time focusing on 'Logical Flow'. Consolidate the capabilities discussion into a single, dedicated section. In the introductory 'What is Agent Core' section, name the capability categories briefly without detailing them, then give the full treatment in the dedicated capabilities walkthrough that follows. Play back the recording and note improvements.

6.8

/10

Effective

Delivery

Your delivery is effective with minor refinements possible.

Rank 6 of 7 · 0 high-severity findings

What is working

- Highly conversational and approachable delivery style that makes complex technical content (serverless deployment, MCP servers, agentic AI) accessible to a broad audience, using relatable analogies like 'nitty gritty,' 'fuss around,' and 'mishmash of stuff.'
- Excellent articulation and projection throughout the presentation — transcription confidence scores are consistently near 1.0 for the vast majority of words, including difficult technical acronyms like ECR, ECS, SDK, and MCP, demonstrating strong vocal clarity.
- Effective use of repetition and signposting ('So,' 'Big one is memory,' 'Right after that') to guide listeners through a dense technical overview, helping maintain audience orientation across multiple capability categories.

MEDIUM Pace

The presentation covers approximately 307 seconds (~5.1 minutes) with a dense word count (~900 words), yielding roughly 176 words per minute. This is on the faster end for technical content. Segments like the capabilities overview (identity, authentication, memory, gateways, observability) are delivered in quick succession without sufficient breathing room between concepts.

Try: Slow down to approximately 140-150 WPM when introducing new technical components (e.g., gateways, MCP server integration). Consider adding a 1-second pause after each named capability before explaining it, signaling to listeners that a new concept is beginning.

Effort: moderate · Impact: medium

MEDIUM Pauses

Pauses are present but underutilized for emphasis. There is a noticeable gap (~1.4s) between 'Amazon Technologies' and the next segment (~178s to ~180s), which works well. However, key value statements like 'It's really cool' (9.67s) and 'Big one is memory' (220s) are followed immediately by elaboration without a beat of silence to let the statement land.

Try: After high-impact statements such as 'It's really cool,' 'Big one is memory,' and 'So let's get started,' insert a deliberate 1.5-2 second pause. This gives the audience time to absorb the point and creates natural emphasis without additional words.

Effort: moderate · Impact: medium

MEDIUM Vocal Variety

The delivery is conversational and generally pleasant, but the pitch and tonal range appear relatively flat across the ~5-minute segment. Enthusiastic phrases like 'super, super easy' and 'pretty cool if you ask me' suggest moments of higher energy, but these do not appear to be anchored by notable pitch shifts or dynamic contrast. The overall delivery maintains a single mid-range register throughout.

Try: Intentionally raise pitch and volume on key benefit statements (e.g., 'works with ANY agent framework,' 'you don't have to') and lower pitch when summarizing or transitioning. Practice vocal modulation by recording and listening back, targeting at least 3-4 notable pitch shifts per minute.

Effort: moderate · Impact: medium

LOW Energy

The speaker demonstrates genuine enthusiasm through word choice ('really cool,' 'nitty gritty,' 'super super easy,' 'mishmash') and a conversational, approachable tone. However, the energy level remains relatively consistent throughout, without notable peaks that signal to the audience which moments are most critical or exciting.

Try: Identify the 2-3 most important moments in the presentation (e.g., the introduction of Agent Core's framework flexibility, the memory solution, the 'let's get started' transition) and consciously amplify vocal energy — faster pace, louder volume, higher pitch — at those specific moments to create memorable highlights.

Effort: moderate · Impact: low

LOW Projection

The transcription confidence scores are consistently high (most words at 0.95-1.0 confidence), indicating clear articulation and good projection throughout. Technical terms like 'Amazon Bedrock Agent Core,' 'ECR,' 'ECS,' 'MCP server,' and 'CloudWatch' were all captured with reasonable confidence, demonstrating clear enunciation of domain-specific vocabulary.

Try: Maintain current projection standards. For terms with slightly lower transcription confidence (e.g., 'strands' at 0.132, 'Landgraf' at 0.088), consider slightly over-articulating these less common proper nouns to ensure audience clarity, especially in live settings with potential audio interference.

Effort: moderate · Impact: low

Practice Drill 10 min

Practice drill for Delivery: Record yourself for the allotted time focusing on 'Pace'. Slow down to approximately 140-150 WPM when introducing new technical components (e.g., gateways, MCP server integration). Consider adding a 1-second pause after each named capability before explaining it, signaling to listeners that a new concept is beginning. Play back the recording and note improvements.

6.8

/10

Effective

Technical Communication

Your technical communication is effective with minor refinements possible.

Rank 7 of 7 · 0 high-severity findings

What is working

- Strong audience-centric framing throughout: the presenter consistently grounds technical features in developer pain points ('you don't have to fuss around with Docker or ECR,' 'that's not gonna cut it at scale in production'), making abstract infrastructure concepts immediately relevant.
- Effective use of the prototype-to-production narrative arc as a unifying thread, which gives the entire introduction a clear purpose and helps the audience understand why each capability matters in context.
- Commendable transparency about platform limitations and honest caveats, such as noting that AWS Strands SDK gets better documentation support and acknowledging that built-in tools are currently limited to two, which builds credibility.

MEDIUM Appropriate Jargon Usage

The presenter references multiple technical terms and frameworks without definition: ECR (Elastic Container Registry), ECS (Elastic Container Service), MCP server, AWS Strands agent framework, LangGraph (transcribed as 'Landgraf'), and agentic AI. While some audience familiarity is assumed appropriately, the inconsistency is notable—Docker gets a pass as commonly known, but ECS/ECR are listed alongside it with no explanation. 'Agentic AI' is used in the opening sentence without definition.

Try: Add brief parenthetical definitions for AWS-specific acronyms on first use: 'Docker, ECR (Amazon's container registry), or ECS (Amazon's container orchestration service).' Also define 'agentic AI' briefly at the start: 'agentic AI—systems where an AI model autonomously takes actions and uses tools to complete tasks.'

Effort: moderate · Impact: medium

MEDIUM Complexity Management

The presentation follows a reasonable progression: what is Agent Core → high-level value → deployment abstraction → framework flexibility → starter toolkit → capability deep-dive → hands-on preview. This layered structure is effective. However, the capabilities section (identity, authentication, tools, memory, gateways, runtime, observability) is delivered in rapid succession as a list without clear separation between items, making it difficult to absorb the distinctions between, for instance, '

Try: When enumerating the seven capability areas, use explicit signposting like 'First capability...' or 'Moving on to the second component...' and provide a one-sentence use case for each before moving to the next. This pacing helps audiences mentally catalog each distinct capability rather than experiencing them as a single stream.

Effort: moderate · Impact: medium

LOW Visual Aid Effectiveness

The presenter makes several forward references to visuals or demonstrations ('we'll see how that works in a moment,' 'we'll see that in a moment,' 'let's look at the capabilities'), suggesting a diagram or slide accompanies this narration. However, based solely on the audio transcript, the architectural description of capabilities—identity, authentication, built-in tools, memory, gateways, runtime, observability—is delivered verbally without any confirmed visual anchor that can be evaluated. The

Try: Ensure that when describing the architectural capability layers (identity, authentication, tools, memory, gateways, runtime, observability), each layer is visually labeled on screen simultaneously as it is verbally introduced. Narration should explicitly reference what the audience sees: 'As you can see on the diagram, at the center is your agent prototype, and surrounding it are these capability

Effort: moderate · Impact: low

LOW Analogy and Metaphor

The presenter uses several effective informal analogies: 'containers under the hood' (abstraction metaphor), 'wrap anything you can dream up' (packaging metaphor), 'not gonna cut it at scale in production' (colloquial but clear), 'nitty gritty,' 'mishmash of stuff,' and 'toolbox.' These make the content approachable. The gateway explanation—'make those look like an MCP server...as a middle layer between the agent and that downstream service'—is a good functional analogy. However, the memory prob

Try: For the memory concept, add a relatable analogy such as: 'Think of it like a customer service rep who has no memory between calls—each time a user returns, they have to re-explain everything. Agent Core's memory component gives your agent the equivalent of a customer file it can pull up instantly.' This makes the scalability problem immediately visceral.

Effort: moderate · Impact: low

LOW Technical Accuracy

The technical claims are generally accurate: Agent Core does support multiple frameworks, provides serverless deployment, integrates with CloudWatch, and includes memory, identity, gateways, and observability components. The presenter correctly notes that AWS Strands SDK gets preferential documentation support. One minor accuracy concern: the presenter says 'deploy it to you in a serverless environment' which is slightly imprecise—Agent Core uses container-based runtime but abstracts server mana

Try: Clarify the deployment model by saying: 'It deploys to a managed container environment—you don't manage any servers, and you pay for what you use, similar to a serverless model.' This avoids potential confusion with pure serverless (Lambda) architectures that have different constraints like execution time limits.

Effort: moderate · Impact: low

Practice Drill 10 min

Practice drill for Technical Communication: Record yourself for the allotted time focusing on 'Appropriate Jargon Usage'. Add brief parenthetical definitions for AWS-specific acronyms on first use: 'Docker, ECR (Amazon's container registry), or ECS (Amazon's container orchestration service).' Also define 'agentic AI' briefly at the start: 'agentic AI—systems where an AI model autonomously takes actions and uses tools to complete tasks.' Play back the recording and note improvements.

Progress

⚠ Trend analysis is not yet implemented. This page will show your score history once three submissions are available.

This is your first submission. Trends need three data points before they mean anything.



PREVIEW: WHAT YOUR PROGRESS PAGE WILL SHOW



After three submissions, you'll see your overall score trend, per-dimension deltas, and which coaching moves produced measurable improvement.

⚠ Preview — This capability is not yet live. Submit two more presentations to activate trend tracking.

How This Was Scored

Score Bands

8.5–10.0	Exceptional	People quote it afterward and then act on it.
6.5–8.5	Effective	Lands with the room and holds attention throughout.
4.0–6.5	Competent	Clear and correct. Not yet memorable or persuasive.
0.0–4.0	Developing	Content or delivery is getting in the audience's way.

Methodology

This report is produced by an automated evaluation system consisting of seven specialist agents, each independently scoring one dimension of your presentation against a detailed rubric. A supervisor agent then merges their outputs, collapses duplicate observations into single findings, and ranks everything by projected impact.

Measured metrics (speaking rate, filler words, pauses, etc.) are computed directly from the transcript and word-level timings. They are reproducible — running the same audio through the system again will produce identical numbers.

Model judgment (dimension scores, findings, suggestions) involves language model reasoning and may vary slightly between evaluation runs on the same content.

Provenance

Report ID	599fc709-d82e-4d23-8e9a-a71a75c035da
Evaluator Release	2.0.0
Rubric Version	1.0.0
Prompt Set Version	1.0.0
Model	us.anthropic.claude-sonnet-4-6 (temp=0.0)
Transcription Service	aws-transcribe
Evaluation Window	PT30M
Run Completed	2026-08-28T17:07:43.145760+00:00

Disclaimer: Scores are guidance from an automated evaluator. They are most useful compared against your own earlier runs on the same rubric version. They are least useful compared against another speaker's scores. This tool measures; it does not judge worth.