INTREPID™
CONNECT. PRACTICE. PERFORM.

AI Agent Onboarding:

A Buyer's Guide for Healthcare Payer Operations Leaders

intrepidlearning.com

Table of Contents

INTRODUCTION

Onboarding Programs Have a Performance Problem

If you run member services operations and want to significantly improve ramp time and attrition, this guide is for you.

Your agents are leaving before they're ready. Across contact centers, annual turnover commonly falls in the 30–40% range, and for health plans, every new member services hire who exits early takes critical knowledge, coaching time, and member experience with them.

The ones who stay still face a long ramp. In complex customer-facing roles, time-to-proficiency can stretch for months. Intrepid reported one onboarding program that reduced proficiency time from nine months to six weeks, a reminder that readiness is not automatic just because onboarding is complete.

That's the problem with most onboarding platforms. They deliver content. They track completions. They generate dashboards full of modules watched and quizzes passed. They can tell you what agents finished.

They can't tell you whether those agents are ready to handle a Medicare appeals call, navigate a complex Commercial benefits dispute, or de-escalate a frustrated member whose coverage question has already bounced through three channels.

That gap between completion and readiness is where health plans are losing money, momentum, and member trust.

This guide exists because we've watched healthcare payer operations leaders make significant investments in onboarding technology that looks impressive in a demo and produces no measurable change in ramp time, attrition, or quality scores six months after launch. The goal here isn't to help you find the most modern-looking platform. It's to help you find the one that actually gets agents ready faster and keeps them longer.

THE REAL COST:

Agent turnover and slow ramp time aren't training problems. They're EBITDA problems. The right onboarding platform doesn't just improve learning metrics, it shows up in operational results.

CHAPTER 1

Why Most Agent Onboarding is Still Broken

Let's start with an uncomfortable truth: most agent onboarding programs weren't working before AI showed up. Adding AI to a broken program doesn't fix it. It makes it more expensive and harder to unwind.

The fundamental problem with agent onboarding isn't content. It's practice.

We've spent years building onboarding programs around what agents need to know: benefit structures, compliance requirements, appeals processes, system navigation, de-escalation scripts. We've created detailed training modules. We've run classroom sessions and virtual ILT. We've deployed eLearning libraries with branching scenarios and knowledge checks.

And yet when an agent takes their first live Medicare appeals call, a confused member disputing a denied prior authorization, too many of them freeze, default to the script, or escalate immediately because they've never actually practiced handling that interaction. They've watched videos about it, but they haven't done it.

Knowledge doesn't create performance. Practice does.

The Gap That's Costing You

There's always been a gap between knowing and doing. The research on how people actually build skills is unambiguous: deliberate practice, spaced repetition, real feedback tied to real behavior—these are the mechanisms. But they're expensive to operationalize at scale.

A human trainer working with one agent at a time can't reach your 500 new hires onboarding in waves across multiple locations. A two-week classroom program can't replicate the dozens of practice interactions it takes to build call-ready confidence. So lots of content got built and we measured whether people watched.

That's not an onboarding strategy. That's a compliance strategy. And compliance doesn't build agents who can handle the call.

30-40%

of new call center agents leave within the first year (SQM Group)

4-6 mos

average time to get contact center new hires proficient (McKinsey)

6 weeks

UnitedHealthcare's time-to-proficiency after move to applied, cohort learning

The Real Cost of Content-First Onboarding

Low transfer	Content doesn't automatically become behavior. Agents learn procedures in training and revert to improvisation under pressure on live calls. The transfer gap is a readiness gap.
Invisible readiness	Completion rates tell you who finished the modules. They tell you nothing about whether anyone is actually ready to handle a real member interaction.
High attrition cost	Every new agent who leaves represents recruiting cost, onboarding and training cost, and lost productivity before they reach full performance. Gallup estimates that replacing frontline employees costs about 40% of salary, while broader turnover research puts replacement costs anywhere from one-half to two times annual salary, depending on role complexity and seniority.
Slow ramp drag	A nine-month ramp means nine months of quality risk, supervisor burden, and operational drag per new hire cohort. Multiply that across hundreds of new hires per year and it becomes a material EBITDA problem.

What AI Makes Possible, and What It Doesn't Automatically Fix

AI creates the possibility of scalable practice. For the first time, a new Medicare agent in a remote location can engage in a realistic, challenging member interaction, getting real-time feedback, trying again, adjusting their approach, without requiring a human trainer on the other end. That's a genuine shift in what's possible.

But AI doesn't automatically create this. You have to design for it.

Most AI tools being sold into agent onboarding programs right now are doing something far simpler: they let agents ask questions about procedures, generate summaries of policy documents, or run basic check-your-understanding prompts after a video. These are useful features. They are not practice.

Practice requires a challenge. It requires feedback tied to specific behavior. It requires repetition with variation. It requires a loop between doing and learning from doing.

When you're evaluating AI-powered onboarding platforms, the first question isn't what AI model does it use. It's: where does the practice happen, and how does the feedback loop work?

CHAPTER 2

What AI Onboarding Actually Means

The term ‘AI-powered’ is doing a lot of heavy lifting right now. Vendors use it to describe wildly different capabilities. Before you evaluate anything, you need a shared vocabulary. Here are the four modes of AI in agent onboarding and how to think about each mode duing platform demos.

MODE 1: AI as Content Delivery

What it looks like: Personalized content recommendations, AI-generated summaries of policy documents, adaptive learning paths based on role or prior knowledge.

What it actually is: A smarter content distribution system. The AI is improving access to information, not actively helping agents build the skills they need for live calls.

When it’s useful: Initial orientation, compliance knowledge transfer, just-in-time procedure lookup.

Where it gets oversold: Personalized content recommendations are not the same as personalized practice. Delivering the right policy document faster doesn’t prepare an agent for an irate member.

MODE 2: AI as a Reflection Partner

What it looks like: Post-call reflection prompts, guided self-assessment, AI-facilitated sense-making after training activities.

What it actually is: A conversational reflection tool. Useful for building self-awareness, but still working from the agent’s own interpretation of events.

When it’s useful: Post-interaction debriefs, processing feedback from supervisors, preparing mentally for difficult call types.

Where it gets oversold: Reflection can deepen awareness. It doesn’t replace the experience of actually practicing the interaction before it’s live.

MODE 3: AI as a Practice Environment

What it looks like: Roleplay simulations where the agent plays themselves and the AI plays the member. The agent has to respond, not choose from a list. The AI reacts to what the agent actually says.

What it actually is: This is the real thing. A new Medicare agent practices a denial appeal call. A Commercial agent practices navigating a benefits dispute with a frustrated member. The AI evaluates the response. Feedback is tied to what they said and how they handled it.

When it's useful: High-stakes call preparation for appeals and grievances, complex benefits questions, CMS compliance scenarios, and de-escalation practice.

When it's oversold: When the practice is just branching scenarios with canned feedback. If agents can't give a free-form response and receive specific feedback on what they actually said, it's not a practice environment, it's a quiz with better UX.

MODE 4: AI as an Evaluation Engine

What it looks like: AI scoring of written or verbal responses, automated quality assessment, supervisor-augmentation tools that flag patterns across entire new-hire cohorts.

What it actually is: Scale without loss of quality. Supervisors can't manually review every practice interaction across hundreds of new hires. AI makes consistent, high-volume evaluation possible.

When it's useful: Any program where you need consistent feedback at volume, such as call simulations, Medicare compliance scenarios, Commercial onboarding cohorts.

Where it gets oversold: As a full replacement for human oversight. AI evaluation is most powerful when it surfaces patterns for human supervisors to act on, not when it operates in isolation.

THE BOTTOM LINE ON AI MODES:

You want a platform where Mode 3 (Practice) and Mode 4 (Evaluation) are first-class capabilities, not afterthoughts. If a vendor spends the entire demo on Mode 1 and Mode 2 and can't articulate what Mode 3 looks like, keep looking.

CHAPTER 3

The Buyer's Framework: 5 Questions That Separate Practice from Content

Most platform evaluations are designed to compare features on a spreadsheet. That's the wrong lens. Features are table stakes. What you need to evaluate is whether this platform will actually get agents ready faster, reduce early attrition, and produce measurable operational impact.

Here are five questions that tell you whether you're looking at a practice platform or content platform.

Question 1

Where Does Practice Actually Happen?

This is the most important question in your entire evaluation. Ask the vendor to show you, not tell you, where an agent practices a skill. Not watches. Not reads. Not reflects. Practices a real member interaction.

What you're looking for: A scenario where the agent has to produce a response (written or verbal), receives specific feedback tied to that response, and has the ability to try again with a different approach.

RED FLAGS: Walk Away	GREEN FLAGS: Good Signs
The vendor shows content with a knowledge check at the end	The vendor shows a real member interaction simulation right now, unedited
Practice is positioned as optional enrichment, not core program	Practice is embedded in the program sequence, not a side feature
Feedback is generic: 'Great response! Here are some tips...'	Feedback references what the agent specifically said and why it matters
The AI plays a generic 'caller' with no healthcare context	Scenarios reflect actual call types: Medicare appeals, benefits disputes, de-escalation

Question 2

How is the AI Configured for Healthcare Payer Context?

Generic AI onboarding is almost always weaker than context-specific AI onboarding. The difference between an AI that says 'listen actively to the member' and one that says 'in a Medicare Part D coverage gap dispute, here's how an experienced agent navigates the first 60 seconds of that call' is enormous.

You need to understand how the vendor allows you to configure the AI for your environment. That includes:

- The specific call types your agents handle: Medicare Advantage, Commercial, Medicaid, dual eligible
- CMS compliance requirements and regulatory context embedded in the scenarios
- Your quality standards, call handling frameworks, and member experience expectations
- The right difficulty calibration for a new hire in week two versus week six

Ask the vendor directly: how does your platform know anything specific about health insurance member services? If the answer is it doesn't, that's a problem.

Question 3

What Does the Feedback Loop Look Like?

Feedback is the mechanism of skill development. Without it, practice just reinforces existing habits—good or bad. Evaluate the feedback loop on three dimensions.

Specificity	Does feedback reference what the agent actually said, or is it templated? Generic feedback is marginally better than no feedback, but not much. You want the AI to say 'you opened with the member's name and validated their concern before explaining the denial... that's the right sequence' not 'good job showing empathy.'
Timeliness	Does the agent get feedback immediately, while the practice interaction is fresh? Delayed feedback breaks the learning loop. In a fast-paced onboarding program, same-session feedback is the standard.
Actionability	Does the feedback tell the agent what to do differently and give them the chance to try it? A coaching observation with no retry opportunity is just criticism. The agent needs to practice the correction, not just read about it.

Question 4

How Does it Handle High-Volume Cohort Onboarding?

Healthcare payers don't onboard one agent at a time. You're managing cohorts of 20, 50, or 100 new hires starting in waves and often in multiple locations, across Medicare and Commercial lines, with different role-specific training paths.

AI onboarding tools that exist outside a structured program architecture get skipped. Busy agents and their supervisors won't navigate to a separate tool to do optional practice activities. The practice needs to be in the program, sequenced into the workflow, and non-negotiable.

- Can the platform manage hundreds of concurrent cohorts with different start dates, roles, and program tracks?
- Can supervisors see real-time readiness data across their cohort before agents go live on the floor?
- Can Medicare and Commercial onboarding run on the same platform with distinct content and compliance requirements?
- What's the administrative overhead to manage a cohort of 50 new hires versus 500?

Question 5

How Will You Know if It's Working?

This is the question most operations leaders never ask, and every vendor hopes you won't. Completion rates will go up when you launch something new with manager support and a push notification. That tells you nothing about readiness.

What you need is leading indicators that agents are actually becoming more capable before they're live. That might look like:

- Quality scores on practice simulations improving across a cohort over time
- Time-to-first-unsupported-call decreasing vs prior cohorts
- 90-day attrition rates falling, the number that most directly ties onboarding quality to EBITDA
- Supervisor-rated readiness scores pre- and post-program for specific call types
- QA scores in the first 60 days on the floor compared to prior onboarding cohorts

ASK EVERY VENDOR:

What measurement framework do you recommend, and what data does your platform generate to support it? If they show you a completion dashboard, keep asking.

PROOF POINT

What This Looks Like in Practice

> "Our new employees consistently challenge us to deliver quality, hands-on learning that is relevant and impactful to their real-world experiences. Intrepid allows us to deliver collaborative and relevant learning programs while onboarding hundreds of employees across the globe in a virtual setting.
>
> - Christy Pletcher, VP Learning Solutions, UnitedHealthcare

UnitedHealthcare faced a challenge that every large health plan recognizes: how do you onboard hundreds of new member services agents in a fully remote environment, prepare them for complex, compliance-sensitive interactions, and get them to proficiency fast without a classroom and without a trainer for every new hire?

Their Medicare Appeals and Grievances program is one of the highest-stakes onboarding environments in health insurance. Agents handle formal disputes from Medicare members about denied benefits, coverage decisions, and quality of care. Getting those interactions wrong has regulatory and financial consequences. Getting agents ready fast has direct EBITDA implications.

UnitedHealthcare rebuilt the program on Intrepid using applied, cohort-based learning. Instead of modules and knowledge checks, new hires were given real case numbers to research and resolve, just as they would with live member cases. They documented their findings, submitted them for review, and received structured feedback. By the end of each module, every agent had worked through eight complete cases end-to-end.

The results were significant.

Supervisors and subject matter experts noted that new hires completing the Intrepid program arrived better prepared than some long-tenured staff. That's not a learning metric. That's an operational outcome.

The Medicare Appeals and Grievances context matters. This wasn't a generic onboarding program. It was designed for one of the most complex, regulated, judgment-intensive member services roles in health insurance. The fact that applied learning produced a 9-month-to-6-week ramp reduction in that environment is the proof point worth holding vendors to.

9 months to 6 weeks

time-to-proficiecy reduction after implementing Intrepid

35 to 72

Net Promoter Score improvement among new hire cohorts

Hundreds

of learners onboarded through Intrepid at UnitedHealthcare

CHAPTER 4

The Vendor Evaluation Checklist

Use this as your working tool through vendor evaluations. These aren't features to check off, they're capabilities to pressure-test in the context of your specific onboarding environment.

Program Architecture

- Can the platform manage concurrent cohorts with different start dates, role tracks, and program configurations without high administrative overhead?
- Are AI practice activities embedded into the program sequence as non-negotiable steps in the journey rather than optional add-ons?
- Can Medicare and Commercial onboarding run on the same platform with distinct content, compliance requirements, and call type scenarios?
- Does the platform support combinations of self-paced learning, live virtual sessions, AI practice, and cohort collaboration in one experience?
- When CMS guidance changes or Commercial compliance requirements shift, how quickly can program content and AI scenarios be updated?
- Walk me through exactly what happens to our costs as we scale from 300 to 3,000 agents per onboarding cycle.
- What governance controls do I have if I need to adjust AI activity volume to manage costs?
- What's the implementation timeline from contract to first live cohort? What are the dependencies on your side versus ours?

AI Coaching Depth

- Does the platform support open-ended agent responses in simulations, not just multiple choice or scripted branching?
- Can AI member characters be configured to simulate actual call types: Medicare appeals, benefits denials, CMS compliance inquiries, de-escalation situations?
- Is feedback tied to what the agent specifically said, not generic encouragement or templated coaching tips?
- Can agents retry simulations with adjusted approaches and receive different, responsive feedback based on what they do differently?
- Can the AI assess agent responses against your specific quality standards and call handling framework?

- Show me a new Medicare Advantage agent completing a full denial appeal simulation from the start of the call to feedback—right now, unedited.
- What happens when the agent gives a completely wrong response? Show me that.
- How does this simulation connect to what the agent does before and after it in the program?
- Who designed the call handling rubric for this scenario, and what was their process?

Supervisor and Operations Visibility

- Can supervisors see readiness scores across their cohort before agents go live so they can intervene before performance problems start?
- Does the platform surface patterns across cohorts such as where agents are consistently struggling and which call types need more practice time, not just individual scores?
- Does the vendor have a framework for connecting platform readiness data to actual QA scores in the first 60 days on the floor?
- Can the platform flag early disengagement patterns that correlate with 90-day attrition risk?

Governance and Compliance

- Who controls AI scenario configuration? Your team, the vendor, or a services partner? What's the turnaround when regulatory requirements change?
- How is agent performance data handled, stored, and governed? What are the data residency commitments relevant to your compliance requirements?
- Is the vendor locked to a single AI model, or model-agnostic? As the AI landscape evolves, your onboarding program should stay current without a platform migration.
- What happens to your AI costs as you scale from 200 to 2,000 agents per onboarding cycle? Get the numbers explicitly.
- CMS updates its Medicare guidance regularly. Walk me through exactly what happens to our AI scenarios when that occurs. Who updates them, how fast, and who controls the change?
- How do you handle the difference between Medicare Advantage and Commercial onboarding on the same platform?
- Can your AI scenarios be configured to reflect our specific QA framework and call handling standards, not a generic call center rubric?

Proven Results in Healthcare Payer Environments

- Can the vendor point to health plan clients with programs at your agent volume who have achieved measurable reductions in ramp time or 90-day attrition?
- What does implementation look like specifically? Do you get a dedicated team or a ticketing queue?
- What is the typical time from signed contract to first agent cohort live on the platform?
- Can they connect you with a reference from a health insurance operation (not a generic enterprise L&D reference) who will have a real conversation with you?
- Can you connect me with a health insurance client in Medicare or Commercial who is 18 months into the platform?
- What's the one thing clients who underperform with your platform have in common?

CHAPTER 5

What to Watch for in the Demo

Here are the patterns that can quickly reveal whether you're looking at something that will actually move the needle on ramp time and attrition, or something that will move the needle on your invoice.

RED FLAGS: Walk Away

- The demo never shows actual practice. It's all content, dashboards, and admin interfaces. Practice gets five minutes at the end.
- 'AI-powered' means a chatbot that answers questions about your policy documents. Useful for lookups. Not useful for building call-ready agents.
- The feedback loop is closed. Agents complete a simulation and feedback goes nowhere. No retry. No supervisor visibility. No next step.
- Measurement equals a completion dashboard. When you ask about outcomes, they show you logins, time-on-platform, and module completions.
- Scenarios aren't configurable. The vendor's library of generic call center scenarios is as specific as it gets. Your Medicare A&G environment doesn't exist in their platform.
- No healthcare payer references. They have enterprise clients. None in health insurance. None with your agent volumes.

GREEN FLAGS: Good Signs

- They show you something hard. The simulation is genuinely challenging: an irate member, an ambiguous benefits question, a CMS compliance edge case. If the practice feels easy, it won't build the skill.
- They lead with your problem. Before showing the product, they ask about your ramp time, your attrition rate, your call types. That's a vendor thinking about outcomes.
- They can articulate the behavior change theory. Ask why their approach builds skill, not just knowledge. If they can explain deliberate practice, feedback loops, and spaced repetition clearly, they have real learning science behind the product.
- Supervisors can see readiness before agents go live. The best platforms give operations leaders a view into cohort readiness that lets them intervene before performance problems hit the floor.
- They have healthcare payer proof. Not a generic case study. Actual ramp time or attrition numbers from a health plan that looks like yours.
- They measure what you care about. The vendor leads with time-to-proficiency, 90-day attrition, and QA correlation, not satisfaction scores.

CHAPTER 6

The Future of AI Agent Onboarding

The platforms that exist today are significantly ahead of where they were three years ago. They're also primitive compared to what's coming. Here's where it's going and what it means for the decisions you're making now.

From Static Scenarios to Dynamic Conversations

Today's AI practice activities are largely scenario-based: a defined call type, a defined member situation, a constrained interaction. The next generation will be fully dynamic: AI member characters that remember prior interactions with this agent, adapt their behavior based on how the conversation develops, and escalate or de-escalate based on what the agent does. That's not a scripted simulation. That's a genuine practice environment.

Implication for buyers: Look for platforms investing in conversational depth and adaptability, not just adding more scenario templates to a static library.

From Individual Practice to Cohort Intelligence

Right now, AI onboarding data mostly tells you how an individual agent performed on a specific simulation. The next wave will aggregate this into cohort-level intelligence: where are agents in this cohort consistently struggling on de-escalation scenarios? Which call types need more practice time before this group is ready for the floor? This turns AI onboarding from a development tool into a readiness diagnostic tool for operations leaders.

Implication for buyers: Ask vendors about their roadmap for aggregate analytics across cohorts. Platforms building this now will create significant differentiation over the next 18 months.

From Human or AI to Human *and* AI

The most effective onboarding programs don't replace human trainers and supervisors, they make them more effective. AI handles the volume, the consistency, and the 24/7 availability of practice at scale. Human trainers get richer data, better-prepared agents, and the ability to focus on the judgment calls that only humans can make: the contextual edge case, the agent who needs a real conversation, the cohort dynamic that the data can't see.

Platforms that understand this build for trainer and supervisor augmentation. Platforms that don't are positioning AI as a headcount reduction play, and the results in those programs reflect it.

From Generic Compliance to Applied Call Readiness

The future of AI-powered onboarding for healthcare payers isn't 'practice being a good call center agent in general.' It's 'practice the exact interaction you'll have tomorrow with a Medicare member disputing a coverage decision.' The closer AI practice gets to the real work, the actual call type, the actual regulatory context, the actual member population, the faster agents reach proficiency and the longer they stay.

Configurability isn't a nice-to-have feature. It's the investment protection for everything you build on the platform. Generic AI onboarding has a short shelf life.

CHAPTER 7

From Evaluation To Decision

You've done the demos. You've asked the hard questions. Multiple vendors have told you they're the best. Here's how to make the call.

Step 1: Get Clear on What You're Solving For

Before you score vendors, score yourself on clarity. Can you answer these questions without hedging?

- What is our current 90-day attrition rate for new member services agents, and what would a 10-point reduction mean for our operating costs?
- What is our current time-to-proficiency for Medicare and Commercial agents?
- Which specific call types produce the most QA failures, supervisor escalations, and agent anxiety in the first 60 days on the floor?
- What does success look like at 12 months and what data would prove it?

If you can't answer these cleanly, you're not ready to buy yet. Any platform will disappoint you if you don't know what you're measuring against.

Step 2: Run a Real Proof of Concept

Most vendor pilots are designed to generate adoption numbers, not prove outcomes. They give you 30 agents and 90 days and measure completion rates.

A real proof of concept has a defined hypothesis: 'We believe that agents who complete three AI simulation activities on Medicare denial appeals will score higher on QA reviews in their first 30 days than agents who don't.' Then you test it.

If the hypothesis holds, you scale. If it doesn't, you learn something real.

Step 3: Evaluate the Vendor Relationship, Not Just the Product

Healthcare payer onboarding is complex. It's regulated. It changes when CMS guidance changes. The product is table stakes. What matters is whether this vendor understands your operational environment well enough to be a real partner, not just a platform provider waiting for support tickets.

Ask for an introduction to your prospective customer success contact before you sign. Have a real conversation with them about your Medicare and Commercial onboarding context. Your instinct about that relationship is worth more than any feature comparison spreadsheet.

Step 4: Build for Proof, Not Perfection

The best onboarding programs we've seen weren't the most elaborate ones. They were the ones that launched with something specific, proved it worked, and expanded from there. Start with one high-stakes call type in one new-hire cohort. Measure ramp time and early QA scores. Build the proof.

Scope creep kills programs. An AI onboarding program that runs for three cohorts and produces a measurable reduction in 90-day attrition is worth more than a platform-wide rollout that achieves 70 percent completion and a good satisfaction score.

FINAL THOUGHT:

The question isn't whether AI has a role in agent onboarding. It does... and it's significant. The question is whether you're going to invest in AI that makes your onboarding program look modern, or AI that makes your agents perform better.

Those aren't the same thing. Stop measuring completions. Start measuring readiness. Build something that shows up in your operational results.

CHAPTER 8

How Intrepid Puts it Together

Everything in this guide—the practice gap, the feedback loop, the need for healthcare-specific configurability, the demand for measurable operational outcomes—shaped how Intrepid built its platform for member services onboarding. This isn't a generic learning platform with healthcare case studies. It's an applied learning platform designed around the specific challenges of getting agents call-ready at health plan scale.

AI Roleplay: Practice Before the Live Call

AI Roleplay puts new hires in realistic member interactions before they're on the floor. The agent plays themselves. The AI plays the member: a frustrated beneficiary disputing a denied prior authorization, a member confused about their out-of-pocket costs, an escalating caller who won't accept an answer.

The AI character responds to what the agent actually says. If they escalate prematurely, the scenario reflects that. If they validate the member's concern before explaining the denial, the interaction shifts accordingly. Agents get specific feedback tied to what they said and how they handled the interaction, mapped to your call handling standards and quality framework.

Agents can retry. They can try a different approach. They can practice the same scenario until the behavior is natural, before a single live member interaction depends on it.

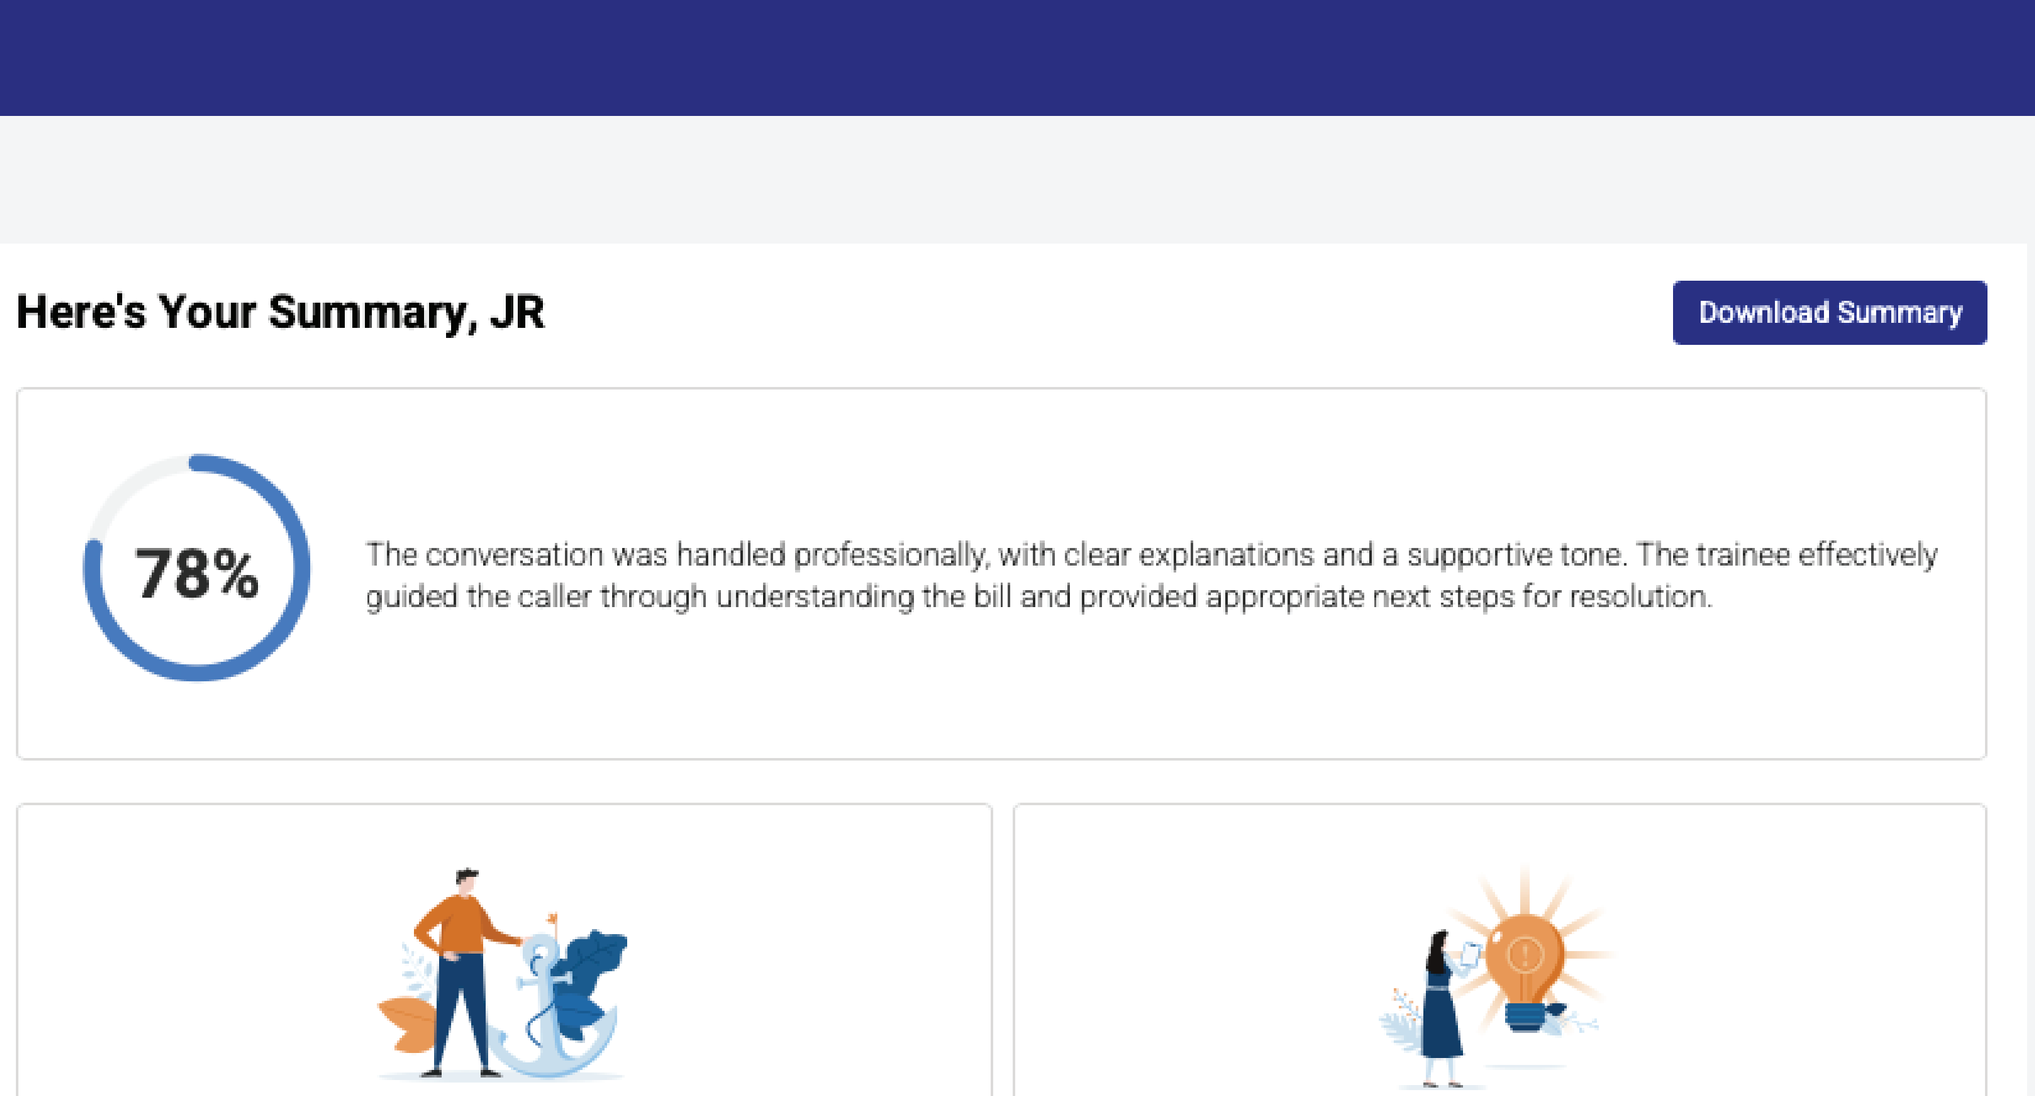

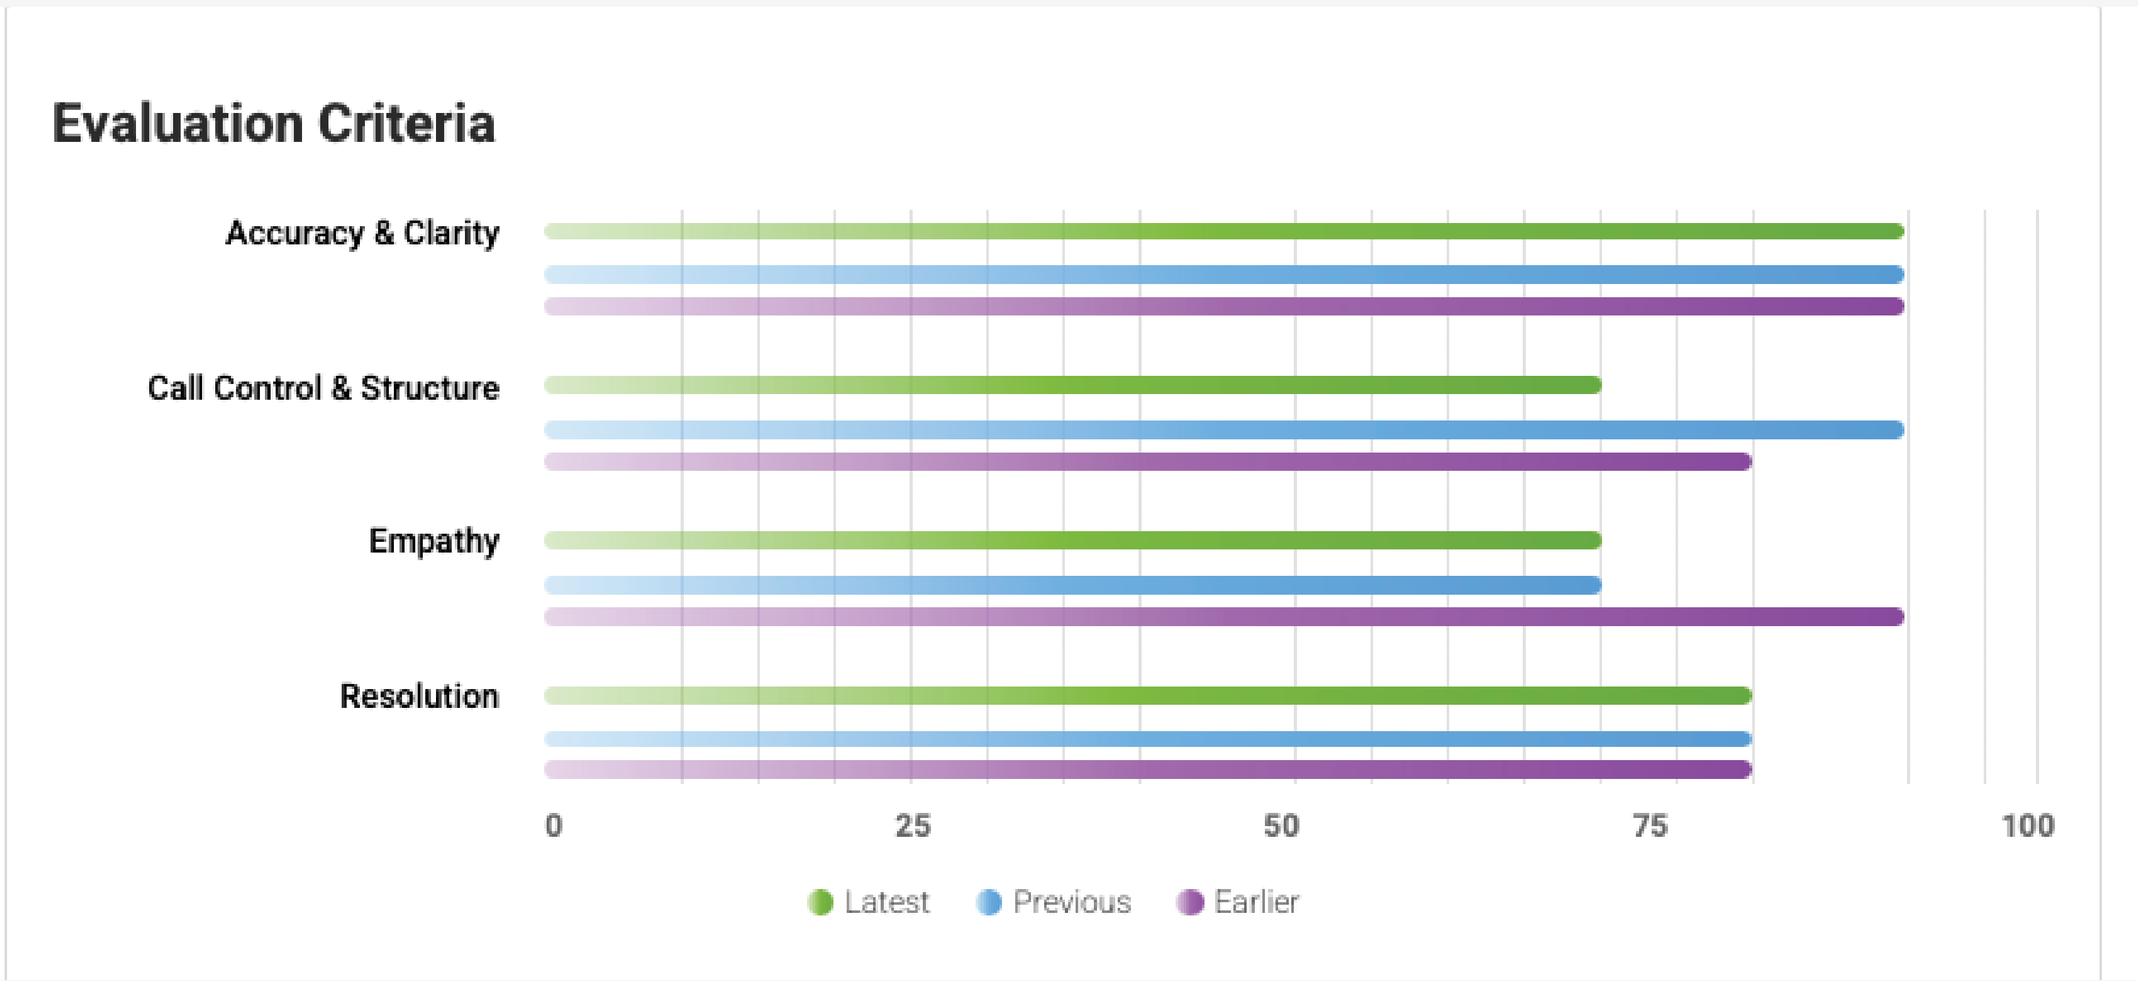

WHAT THIS SOLVES

The readiness gap. Agents arrive on the floor having practiced the actual call types they'll face, not just having watched videos about them. That's the difference between a 9-month ramp and a 6-week ramp.

AI Coaching: A Thinking Partner for Complex Situations

AI Coaching gives agents an on-demand resource that goes beyond answering procedural questions. Agents can bring a real challenge such as a call type they're struggling with, a compliance scenario they're unsure about, a member interaction they handled poorly and want to better understand, and work through it with an AI coach that pushes their thinking and surfaces what they missed.

Unlike a general-purpose AI tool, Intrepid's AI Coaching activities are purpose-built for your environment. Each one is configured with your call handling framework, your regulatory context, and the specific competencies you're developing. The AI responds to what this agent is working on right now, not a generic new hire.

And because AI Coaching is embedded in the program sequence, not a separate tool agents have to remember to use, it happens as part of the onboarding journey, not alongside it.

AI Evaluation: Readiness Visibility at Cohort Scale

AI Evaluation solves one of the oldest problems in high-volume onboarding: how do you assess whether hundreds of agents are actually ready, not just whether they completed the program, without overwhelming your supervisors with manual review?

With AI Evaluation, agents submit responses to simulated call scenarios, written case analyses, or compliance decision exercises, and receive structured feedback that's consistent across every agent in every cohort. Supervisors set the rubric. The AI applies it at scale and flags patterns for human follow-up.

Your supervisors spend their time on the agents and situations that need human judgment, not reviewing routine submissions across a cohort of 80 new hires.

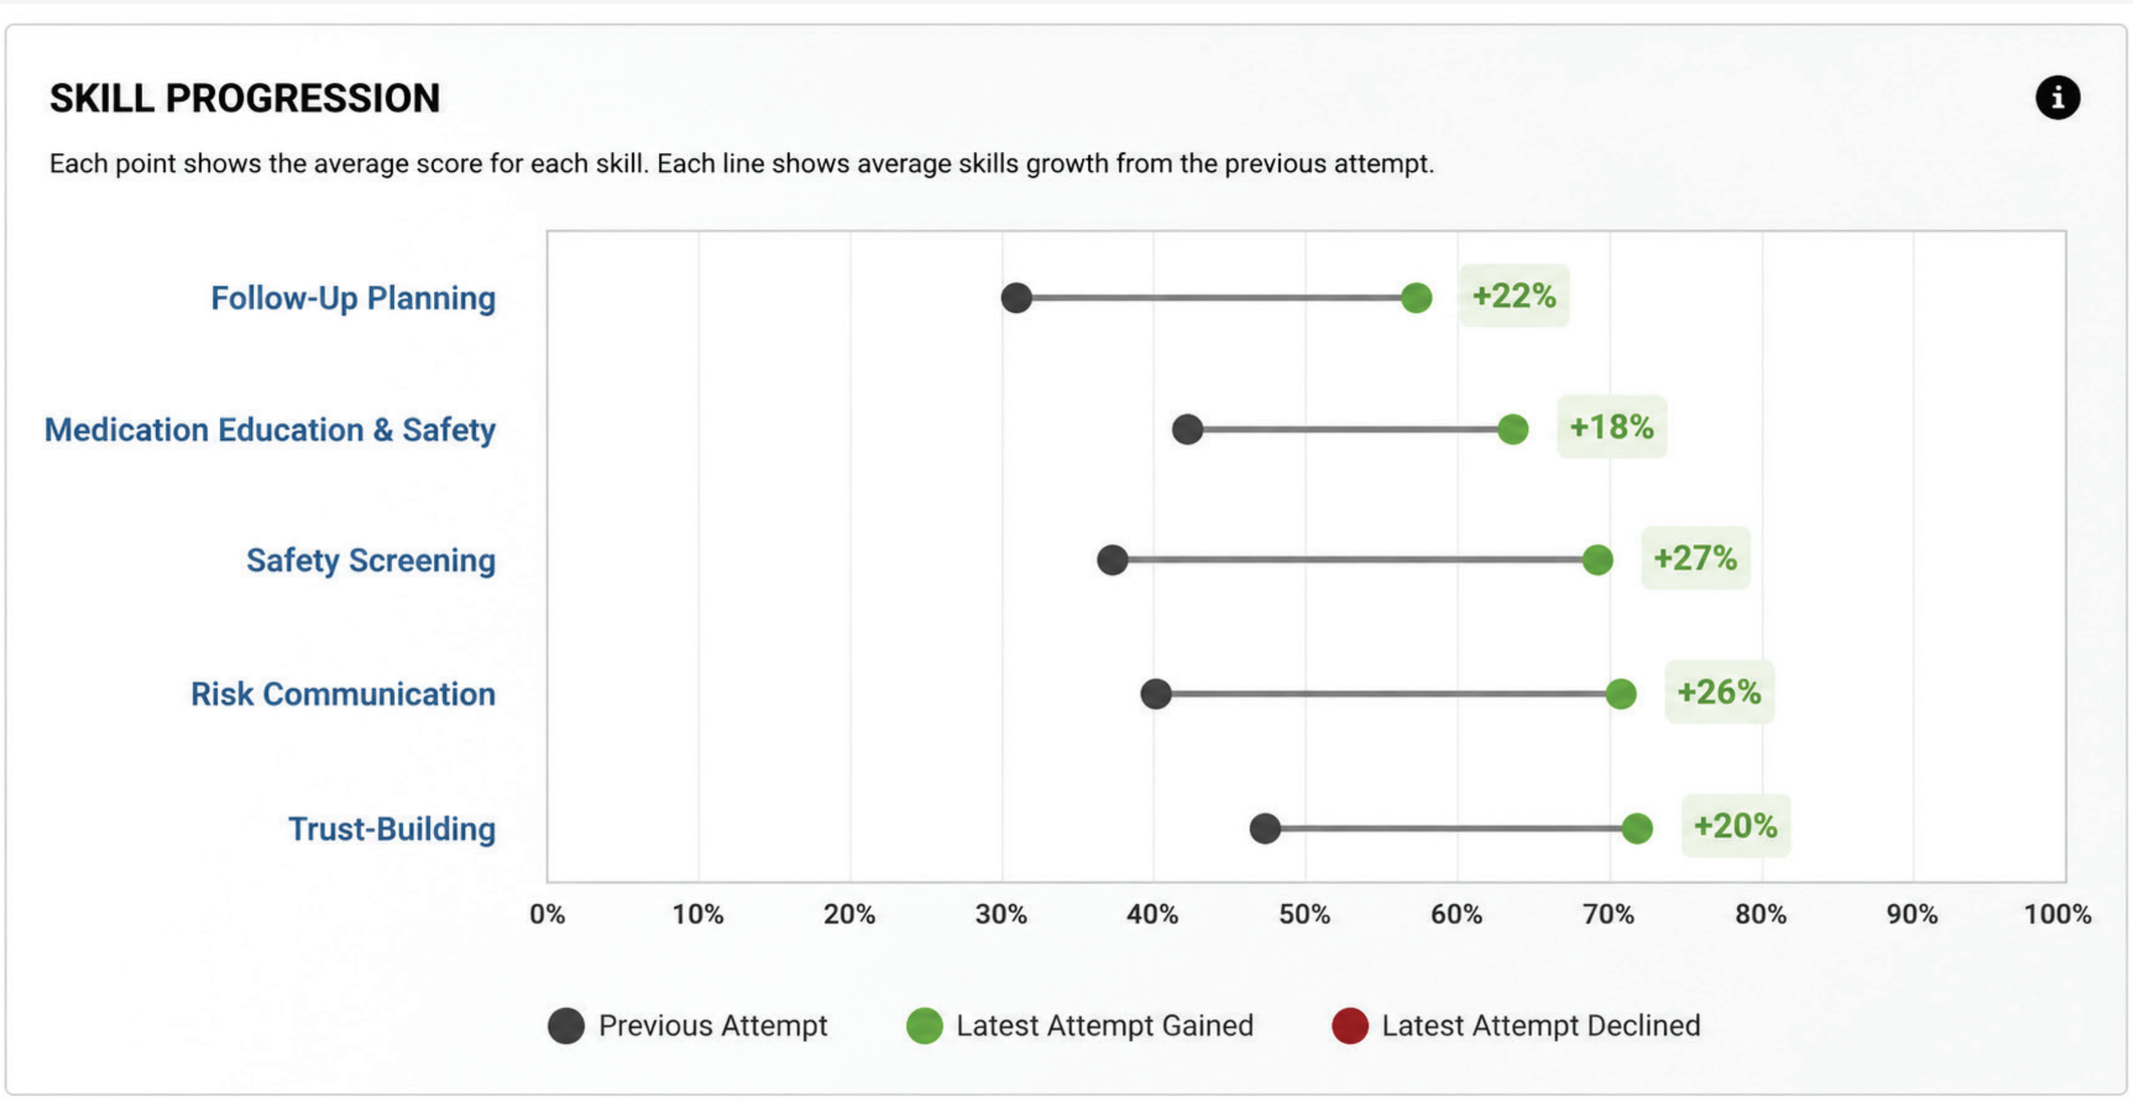

WHAT THIS SOLVES

The scale problem. You no longer have to choose between depth and reach. Supervisors can see who's ready before agents go live across the entire cohort, not just the agents who raised their hands.

All Three, Together, Inside a Structured Onboarding Program

The real capability of the Intrepid platform isn't any single AI feature. It's how AI practice, coaching, and evaluation work together inside a cohort-based onboarding architecture that makes readiness—not completion—the measure of success.

Stage	Activity	What it Builds
Learn	Content module + cohort discussion	Shared knowledge and compliance context
Practice	AI Roleplay: the call they'll face on the floor	Call-ready confidence through repetition
Apply	Real case simulation or on-the-job assignment	Application and accountability
Reflect	AI Coaching: debrief the interaction	Self-awareness and adjustment
Evaluate	AI Evaluation: assess their handling against rubric	Evidence of readiness, not just completion

Built for Health Plan Scale: Governance and Configurability

Intrepid is built for organizations where compliance isn't optional and regulatory context changes regularly. Every AI activity is configurable to your call handling framework, your CMS compliance requirements, and your specific member services environment, without requiring developer resources to manage. Your operations and L&D teams own the experience.

Intrepid is model-agnostic. We don't tie your onboarding program to a single AI vendor. As the AI landscape evolves, your program stays current.

And because cohort management at health plan scale is a real operational challenge, Intrepid's platform is built to handle hundreds of concurrent cohorts with different start dates, role tracks, and lines of business, all with the administrative efficiency a large member services operation requires.

THE BOTTOM LINE

AI Roleplay builds call-ready agents. AI Coaching deepens their judgment. AI Evaluation confirms their readiness. All three are embedded inside cohort programs designed for the complexity and volume of healthcare payer onboarding to improve operational outcomes.

Intrepid by VitalSource is an applied learning platform for healthcare payer onboarding. We help health plans build call-ready agents faster, reduce 90-day attrition, and turn onboarding investment into measurable operational results.

Ready to see what this would look like for your member services operation? Contact Intrepid today.

intrepidsales@vitalsource.com
www.intrepidlearning.com

INTREPID™
CONNECT. PRACTICE. PERFORM.