

AI Vocabulary Reference • Tier 2: Practitioner

The language of doing • Aurorae Group, LLC

This reference builds on foundational AI vocabulary (Tier 1). It is designed for people who are actively using AI tools and want to work with them more effectively, intentionally, and safely. If you haven't reviewed Tier 1, start there.

Working with GenAI Tools

Terms that help you understand how these systems actually function — and why they behave the way they do.

Context Window	<p>The amount of text a GenAI model can process and hold in memory at one time — including your prompt, any documents you upload, and its own response. When content exceeds this limit, earlier material gets dropped.</p> <p>Why it matters in practice: If you paste a long policy document and then ask a question, the model may have 'forgotten' the beginning by the time it reaches the end. Knowing this helps you structure inputs strategically — most important information first, or in chunks.</p>
Token	<p>The basic unit of text that AI models process — roughly equivalent to three-quarters of a word. Models have token limits that govern both inputs and outputs.</p> <p>Why it matters in practice: Token limits explain why a model might cut off a long response mid-sentence, or why uploading a very large document may not work as expected. Paid plans typically offer higher token limits.</p>
System Prompt	<p>A hidden set of instructions given to an AI model before the conversation begins, used to shape its persona, constraints, and behavior. Users typically don't see the system prompt in consumer tools.</p> <p>Why it matters in practice: Understanding that system prompts exist explains why the same underlying model can behave very differently across different tools or platforms — the model is the same, but the instructions shaping it are different.</p>
Temperature	<p>A setting that controls how creative or consistent an AI model's outputs are. High temperature produces more varied, creative responses. Low temperature produces more focused, predictable ones.</p> <p>Why it matters in practice: For HR tasks requiring precision — policy language, compliance summaries, factual documentation — lower temperature settings (where available) reduce variability and risk. For brainstorming or drafting, higher temperature may be useful.</p>
Iteration	<p>The process of refining AI outputs through successive prompts — asking the model to revise, expand, shorten, reframe, or adjust its response until it meets your needs.</p> <p>Why it matters in practice: First outputs from GenAI are rarely final. Building iteration into your workflow — treating AI as a drafting partner rather than a one-shot tool — consistently produces better results and builds your own prompting fluency.</p>

Multimodal AI	<p>AI systems capable of processing and generating multiple types of content — text, images, audio, video, and documents — within a single interaction.</p> <p>Why it matters in practice: Multimodal capability means you can upload a scanned policy document, an org chart image, or a meeting recording and ask the AI to work with it directly. This expands what's possible well beyond text-only interactions.</p>
----------------------	--

Prompting & Output Quality

The vocabulary of getting better results — faster, with less rework, and with more confidence in what you're producing.

Prompt Engineering	<p>The practice of crafting inputs to a GenAI system with deliberate structure and specificity to improve the quality, accuracy, and usefulness of outputs.</p> <p>Why it matters in practice: This is a learnable skill, not a technical one. HR professionals who invest in prompting fluency consistently get more useful outputs with less back-and-forth — and are better positioned to shape how their teams use AI.</p>
Role Prompting	<p>A technique in which you instruct the AI to respond from a specific perspective or expertise — for example, 'Act as an experienced HR business partner' or 'Respond as a compliance attorney reviewing this policy.'</p> <p>Why it matters in practice: Role prompting is one of the most immediately useful techniques for HR. Assigning a role shapes tone, depth, and framing in ways that generic prompts don't — and produces outputs that are more directly applicable to your context.</p>
Chain-of-Thought Prompting	<p>A technique that asks the AI to reason through a problem step by step before providing an answer, rather than jumping directly to a conclusion.</p> <p>Why it matters in practice: For complex HR scenarios — performance cases, policy decisions, or org design questions — asking the model to 'think through this step by step' surfaces reasoning you can evaluate and challenge, rather than just an output you have to accept or reject.</p>
Prompt Injection	<p>A security vulnerability in which malicious instructions are embedded in content fed to an AI system — for example, in a document or email the AI is asked to summarize — causing the AI to behave in unintended ways.</p> <p>Why it matters in practice: As HR teams use AI to process documents, resumes, or survey responses, prompt injection is a real risk. Knowing it exists prompts the right question: what are we feeding this tool, and who created that content?</p>
Output Types	<p>AI outputs can be structured (tables, JSON, formatted lists) or unstructured (prose, summaries, narratives). The format you request significantly affects how useful the output is for your intended purpose.</p> <p>Why it matters in practice: Specifying output format in your prompt — 'Give me this as a table with three columns' or 'Summarize this in three bullet points for a leadership audience' — dramatically reduces editing time and improves usability.</p>

Grounding	<p>The practice of anchoring an AI's responses to specific, provided source material — documents, data, or references — rather than relying solely on its training data.</p> <p>Why it matters in practice: Grounded prompts reduce hallucination risk significantly. Instead of asking 'What does FMLA require?' paste in your organization's current policy and ask the model to work from that. The output is then traceable to a source you control.</p>
------------------	---

Workflow Integration

Terms that matter once AI moves from experimentation into deliberate, repeatable practice within a team or organization.

AI Workflow Integration	<p>The deliberate embedding of AI tools into existing work processes — identifying where AI adds value, designing the human review steps, and building the habit of use into team routines.</p> <p>Why it matters in practice: Integration is different from experimentation. Experimentation is trying a tool. Integration is deciding where it lives in your workflow, who uses it, how outputs get reviewed, and what gets documented. HR is well-positioned to model this distinction for the broader organization.</p>
Human-in-the-Loop (applied)	<p>In practice, this means defining — in advance — which tasks require human review before an AI output is acted upon, who conducts that review, and what they are checking for.</p> <p>Why it matters in practice: At the practitioner level, human-in-the-loop isn't a principle — it's a checklist habit. The Human Review & AI Disclosure Checklist you received as a session resource is a practical implementation of this concept.</p>
Custom GPT / AI Assistant	<p>A pre-configured version of a GenAI tool built with a specific set of instructions, persona, and knowledge base — designed to perform a defined set of tasks consistently without requiring users to re-prompt from scratch each time.</p> <p>Why it matters in practice: Custom assistants reduce the prompting burden on individual users and help standardize outputs across a team. An HR team might build a custom assistant for drafting job descriptions, summarizing exit interview themes, or generating performance review language.</p>
Retrieval-Augmented Generation (RAG)	<p>A technique that connects a GenAI model to an external knowledge base or document repository, allowing it to retrieve and reference specific organizational content when generating responses.</p> <p>Why it matters in practice: RAG is what makes it possible to build an AI assistant that 'knows' your organization's specific policies, handbooks, and procedures — without those documents ever being used to train the underlying model. This is an important distinction for data privacy.</p>
AI Policy vs. AI Guidelines	<p>An AI policy is a formal, binding document governing how AI may be used in the organization. AI guidelines are advisory — recommended practices that inform use without mandating it. Both serve a role; they are not interchangeable.</p> <p>Why it matters in practice: Many organizations start with guidelines and formalize into policy as usage matures. HR typically owns the rollout of both. Understanding the distinction helps you advise leadership on what level of structure is appropriate for your organization's current stage.</p>

