



TIO

The Intelligent Organization

Executive AI & Automation Vocabulary Quick Reference

A plain-English reference for senior leaders.

01

The Essential Terms Every Senior Leader Should Know

Sharper definitions for faster executive use

Term	Executive meaning	Leader question
Artificial Intelligence	Computer systems, methods, and applications that perform work normally requiring human-like perception, language, pattern recognition, prediction, generation, planning, or decision support.	What task is the system being asked to perform, and where does human judgment remain required?
Machine Learning	A way of building AI systems by training them on examples so they can detect patterns, classify information, make predictions, or recommend actions without being explicitly programmed for every possible case.	What data was used, what pattern is being learned, and how do we know it works in our context?
Generative AI	AI that produces new content or structured output in response to instructions, examples, data, or prompts. The output may include text, code, images, audio, video, summaries, plans, analysis, or synthetic records.	Where must the output be reviewed, verified, cited, or grounded before use?
Large Language Model	A model trained on very large collections of text, code, and related material so it can predict, generate, summarize, translate, classify, reason over, or transform language-based inputs and outputs.	Which model is being used, for which task, with what data, controls, and evaluation?
Foundation Model	A large, general model trained to support many downstream uses and then adapted, prompted, connected, or embedded into applications for specific tasks.	What application, data, workflow, and decision context sits on top of the foundation model?
Frontier Model	One of the most capable current AI models at the leading edge of general performance, often strong in reasoning-like tasks, coding, multimodal work, long context, tool use, and complex instruction following.	Do we need frontier capability, or would a smaller, cheaper, more controllable model be sufficient?
Open-Weight Model	A model whose learned weights are made available for others to download, run, adapt, or deploy under a license. Open weights do not necessarily mean open training data, open code, open safety methods, or unlimited commercial rights.	What exactly is open, what does the license permit, and can we operate it safely?
Proprietary Model	A model controlled by a vendor or provider, usually	What are the data, contract, cost,

Term	Executive meaning	Leader question
	accessed through an application, cloud service, or API rather than released as downloadable model weights.	resilience, and exit implications?
Multimodal Model	A model that can process or generate more than one kind of input or output, such as text, images, documents, audio, video, code, charts, screenshots, or sensor-like data.	What input types are allowed, and which ones may contain sensitive or regulated information?
Parameter	A learned numerical value inside a model that helps determine how the model transforms input into output. Parameters are part of the model's learned machinery. They are not individual facts, files, or memories.	How does this model perform on our use case, not just how large is it?
Token	A unit of text or content that a model uses to measure input, output, context, and often cost. A token may be a word, part of a word, punctuation mark, code fragment, or other model-specific unit.	How many tokens will typical, peak, and worst-case use consume?
Context Window	The maximum amount of instructions, conversation history, retrieved material, input, and generated output that a model can process within one active interaction. It is the model's working space, not its permanent memory.	How much information can the model use reliably, not merely accept technically?
Training	The process of adjusting model parameters using data so the model develops patterns, capabilities, behaviors, and response tendencies. Training may happen at large scale before release or at smaller scale for adaptation.	Are we training a model, fine-tuning one, or simply using one?
Inference	The act of using a trained model to produce an output from a prompt, input, document, image, tool result, or other request. Inference is what happens when the model is called in operation.	Where does inference occur, and what data is sent there?
Prediction	A model's estimate of a likely, useful, or probable output based on its training, the input, and the surrounding system design. In AI, many outputs that look like answers are still forms of prediction.	How will we verify outputs where correctness matters?
Reasoning	The apparent ability of a model to work through multi-step problems, compare information, follow constraints, draw inferences, and produce a coherent answer. In TIO language, AI reasoning is not the same as human judgment or accountable decision-making.	What evidence shows the model can reason reliably in our use case?
Hallucination	A plausible, fluent, or confident AI output that is incorrect, fabricated, unsupported, misleading, or not grounded in the sources or facts required for	What must be grounded, cited, checked, or reviewed before use?

Term	Executive meaning	Leader question
	the task.	
Grounding	Providing the model with relevant, approved, and current source material so its response is tied to known information rather than only to general training patterns.	Which sources is the model grounded in, and are they authoritative, current, and permission-appropriate?
Retrieval-Augmented Generation	A system pattern in which relevant information is retrieved from approved documents, databases, or knowledge stores and then supplied to a model to help generate a response.	Can the system show which sources it used and when retrieval failed?
Evaluation	The structured testing of whether an AI system performs well enough for its intended use, under realistic conditions, with appropriate attention to accuracy, reliability, risk, cost, user behavior, and failure modes.	What must be true before this use case passes its Evidence Gate?
Accuracy	How often the system produces a correct output for a defined task, data set, condition, and measurement method.	What errors matter most, and how are they measured?
Reliability	The degree to which an AI or automation system performs dependably across users, data, conditions, edge cases, time, workload, and real operating situations.	How will performance be monitored after launch?
Explainability	The ability to explain, in terms useful to the relevant human audience, why a system produced a particular output, recommendation, action, or decision support result.	What level of explanation is required for this decision or process?
AI Assistant	An AI system that helps a person draft, summarize, search, analyze, plan, translate, classify, or perform other tasks through interaction. The assistant supports the user; the user normally remains responsible for the work.	What tasks is the assistant approved to support, and what data can it use?
AI Agent	An AI system that can pursue an objective by planning steps, using tools, taking actions, checking results, and adapting its path within defined boundaries.	What actions can the agent take, and what approvals or stop conditions apply?
Agentic AI	AI designed to work in goal-directed sequences, often using planning, tools, memory, retrieval, feedback, and monitoring to move from answering questions toward completing work.	Where does the system require human approval before acting?
Workflow Automation	Technology that moves work through predefined steps, rules, notifications, handoffs, approvals, checks, or system actions so a known process runs	Has the work been redesigned before automation?

Term	Executive meaning	Leader question
	more consistently.	
Robotic Process Automation	Software automation that imitates user actions across existing systems, often to move data, complete repetitive steps, or bridge systems that are not well integrated.	Is RPA a bridge, or are we avoiding a deeper system or process fix?
Tool Use	A model's ability to call external tools, functions, or systems such as search, databases, calculators, email, calendars, code execution, workflow applications, or APIs.	Which tools can the model use, with what permissions, logs, and limits?
Function Calling	A controlled pattern where a model requests a predefined function or tool to return information or perform an action, often using structured inputs and outputs.	What functions exist, who approved them, and what can they read, change, or trigger?
Autonomous System	A system that can operate with limited or delayed human intervention within defined objectives, boundaries, controls, and monitoring.	What decisions or actions are autonomous, and what are the limits?
Human-in-the-Loop	A design in which a human must review, approve, correct, or decide before the system completes a material action or outcome.	Can the human realistically catch and correct errors before harm occurs?
Stop Condition	A predefined condition that requires pausing, escalating, disabling, redesigning, rolling back, or retiring an AI or automation use case.	What evidence would make us stop, roll back, or redesign this use case?
Data Quality	The degree to which data is accurate, complete, current, consistent, accessible, traceable, and fit for the specific purpose, decision, workflow, automation, or AI Use Case in question.	Is this data fit for this decision, workflow, or AI Use Case?
Data Lineage	A view of where data comes from, how it moves, how it changes, who touches it, and where it is used in reports, workflows, automations, and AI Use Cases.	Can we trace the output back to source data and transformations?
Data Classification	A practical way to label data by sensitivity, permitted use, audience, access rules, protection needs, and AI-use restrictions.	Which data classes are allowed in which tools and AI Use Cases?
Embedding	A numerical representation of text, images, or other content that captures meaning in a form software can compare, search, cluster, or retrieve.	What content is embedded, where is it stored, and who can search it?
Vector Database	A database designed to store and search embeddings so systems can retrieve content with similar meaning, related context, or semantic relevance.	What information is indexed, how is access controlled, and how is retrieval evaluated?

Term	Executive meaning	Leader question
Data Sovereignty	The organization's practical ability to control where data is stored, processed, accessed, moved, retained, deleted, backed up, exposed to law, and carried if providers change.	Can we control where data goes and what happens if we change providers?
Enterprise Memory	The organization's retained, governed, searchable, and usable knowledge across documents, systems, decisions, interactions, lessons learned, operating data, and approved AI memories.	What memory is retained, for whom, for how long, and under whose control?
Intelligence	The effective reasoning, creativity, synthesis, learning, and decision capacity available to people, teams, agents, and the enterprise.	Whose reasoning or decision capacity improves, using which approved sources, and how will the organization know?
Intelligence Expansion	A leader-first, role-based pattern for expanding approved AI use from direct leadership experience to personal, collective, and enterprise intelligence.	Which leadership or role group should be enabled next, for what work, with what support and data boundaries?
Enterprise Autonomy	The freedom, control, capability, and resilience the organization retains to protect what is uniquely its own and change direction.	Can we understand, govern, stop, recover, replace, and change direction without losing the business capability?
Delegated Autonomy	Authority granted to a system, agent, automation, or bounded unit of work without a person approving every instance. It is limited, evidence-gated, reviewable, revocable, and recoverable.	What exact unit of work is being delegated, within which limits, evidence, monitoring, stop, and recovery conditions?
Alpha	The TIO reference name for the organization's complete Sovereign Enterprise Intelligence System. Alpha is an integrated enterprise capability, not a single model, database, or product.	Which capabilities, knowledge, authority, controls, and operating responsibilities make up Alpha in this organization?
Intelligence Core	The replaceable models, agents, reasoning methods, orchestration, evaluations, tools, and human expertise used by Alpha to understand, plan, coordinate, and learn.	Which model, agent, tool, solver, or human specialist should handle this task, and can it be replaced?
Enterprise Knowledge and State	The organization's governed knowledge, definitions, intellectual property, authoritative records, events, evidence, and current operating condition.	What source establishes truth now, who owns it, how current is it, and what may be inferred or remembered?
Work and Decision Map	A shared view connecting the outcome, bounded work, decisions, people, AI and automation, information, authority, systems of record, controls, and evidence.	What outcome, work, decisions, people, information, authority, controls, and evidence must fit together?
Assurance and Operations	The lifecycle used after launch: Register, Approve, Operate, Observe, Review, Change, Recover, Retire.	Who owns the live system, what is observed, when is it reviewed, and how will it be changed, recovered, or retired?
End-State Architecture	The TIO enterprise reference architecture describing the future operating condition and the controls the organization must preserve.	What future operating condition are we building, and what control boundaries must remain explicit?
Transition Architecture	The designed route and temporary operating states used to move a bounded part of the organization from its current operation to its target operation.	Will this bounded unit evolve in place or be established beside the current operation, and how will work and authority move?
Establish the New Operation Beside It	A primary transition route in which a bounded target operation is created beside the current one, proven, given work and authority in controlled steps, and followed by an explicit legacy outcome.	Can the target operation be bounded, proven, supported, reconciled, and given authority without creating uncontrolled duplication?

02

Common Distinctions

Use these before approving a proposal

Do not confuse	With	Practical distinction
AI model	AI application	The model produces outputs. The application packages the model into usable work with controls and support.
Assistant	Agent	An assistant supports a user. An agent can plan and take actions toward an objective.
Automation	AI	Automation executes defined steps. AI interprets, generates, predicts, classifies, or recommends.
Open source	Open weight	Open source implies broader code rights. Open weight may only mean the model weights are available.
Context	Memory	Context is current working space. Memory is retained information available later.
Training	Inference	Training creates or modifies the model. Inference uses the model to produce outputs.
Fine-tuning	Grounding	Fine-tuning changes behavior. Grounding supplies reference information at use time.
Pilot	Production	A pilot tests a bounded idea. Production means the organization relies on it in real work.
Accuracy	Reliability	Accuracy is correctness on a test. Reliability is dependable performance in real conditions over time.
Tool	Capability	A tool is acquired. A capability is built across people, process, data, technology, governance, and management.