

# Agentic AI Identity 101

**Agentic AI** represents a turning point in computing. This cheat sheet traces how AI evolved and shows why identity defines the boundary between innovation and exposure.

## The Evolution of *Autonomy*

Generation	Rule-Based Chatbots & Widgets (“Press 1 for Billing”)	Conversational AI (Amazon Alexa, etc.)	Generative AI (ChatGPT, etc.)	Agentic AI
Model	Predefined logic and if/then menus	Natural-language recognition and intent mapping	Content and code generation from large models	Planning, reasoning, and autonomous execution
Value	Cost reduction through scripted responses	Improved containment and user experience	Acceleration of knowledge and creative work	Full workflow automation and independent action
Time	1990s - 2010s	2010s - 2020	2020 - 2023	2023 - Present
Limitation	Fragile, narrow domain; required exact input	Reactive; limited to single-domain context	Suggestive, not executable	Demands oversight of autonomous access to enterprise systems and data
Identity Implication	No distinct identity – actions mapped directly to human operators or systems.	Shared credentials across instances; limited attribution or contextual control.	Human-linked identity; outputs tied to users, not models.	Requires distinct <b>non-human identity-based</b> access with real-time control and auditable scope.

## 10 Core Concepts in *Agentic AI Identity*

Term	Definition
1 <b>Agentic Identity</b>	Defines an agent’s verified origin, capabilities, and access scope.
2 <b>Blended Identity</b>	Combines an agent’s own credentials with the entitlements of the human or service activating it.
3 <b>Persistent Identity</b>	Maintains continuity across sessions so the agent’s actions can be attributed and audited over time.
4 <b>Role-Based Capabilities</b>	Grants the agent privileges aligned with its defined function and operational scope.
5 <b>Autonomous Authority</b>	Recognizes the agent’s ability to make and execute decisions independently, necessitating oversight.
6 <b>Authentication and Authorization</b>	Establishes how the agent proves what it is and what it may access within enterprise systems.
7 <b>Contextual Identity Adaptation</b>	Adjusts the agent’s access and behavior dynamically based on environment and posture.
8 <b>Behavioral Consistency and Alignment</b>	Ensures the agent’s decisions remain predictable and aligned with human and organizational objectives.
9 <b>Accountability and Audit Trails</b>	Captures every action and decision under the agent’s identity for traceability and compliance.
10 <b>Human-Agent Boundary</b>	Distinguishes between delegated human actions and autonomous machine execution.

## Enterprise Use Cases

### Where *Agentic AI* Is Already Active

- CI/CD and DevOps Agents**  
 Review code, test builds, deploy services, and monitor systems across environments.
- Data and Analytics Agents**  
 Query data, create reports, and update dashboards within enterprise data flows.
- Customer and Workflow Agents**  
 Resolve support issues across CRM, ticketing, and internal applications.
- Cross System Orchestration Agents**  
 Connect actions between SaaS, APIs, and partner platforms.
- Security and Operations Agents**  
 Detect anomalies, rotate keys, and perform remediation tasks.
- Research and Synthesis Agents**  
 Collect data, summarize insights, and initiate follow-up actions.
- Governance and Compliance Agents**  
 Track policy, create audit logs, and manage data retention.

## Key Risks Most Teams Miss

- Persistent Credentials**  
 Agents operate under static or shared credentials that, once exposed, give lasting access.
- Privilege Creep**  
 Access grows as agents connect to more systems without periodic review.
- Cross-Boundary Access**  
 Agents move between clouds, SaaS, and on-prem, extending trust beyond intended limits.
- Context Gaps**  
 Access decisions ignore real-time posture or environmental state.
- Limited Traceability**  
 Agentic actions are difficult to map back to a specific identity or decision chain.

## Practical Controls That Work

- Inventory Every Agent**  
 Treat each as a managed identity with defined privileges and lifecycle.
- Adopt Identity-First Design**  
 Use short-lived, identity-bound credentials instead of stored secrets.
- Authorize by Context**  
 Grant access only when identity, purpose, and posture align.
- Keep the Chain Auditable**  
 Link every action to the agent, credential, and task.
- Enforce Policies Everywhere**  
 Apply consistent controls across all environments.

## Why the Identity Layer Matters Now

**Agentic AI** operates with the same privileges as human users, but with speed, scale, and autonomy that traditional IAM cannot govern. **Agent identities** often share credentials, persist without oversight, and cross trust boundaries unchecked – creating risk and operational drag.