AI that Acts:
The Agentic Revolution in Business

Dr. Andrew Schwarz

LSU | E. J. Ourso College of Business
Stephenson Department of
Entrepreneurship & Information Systems



Agentic artificial intelligence is an AI capability where a technological system can independently set goals, plan multi-step actions, and execute tasks with minimal human oversight by acting as a proactive “agent” rather than a reactive “tool” by understanding context, reasoning, and adapting to new information to achieve objectives

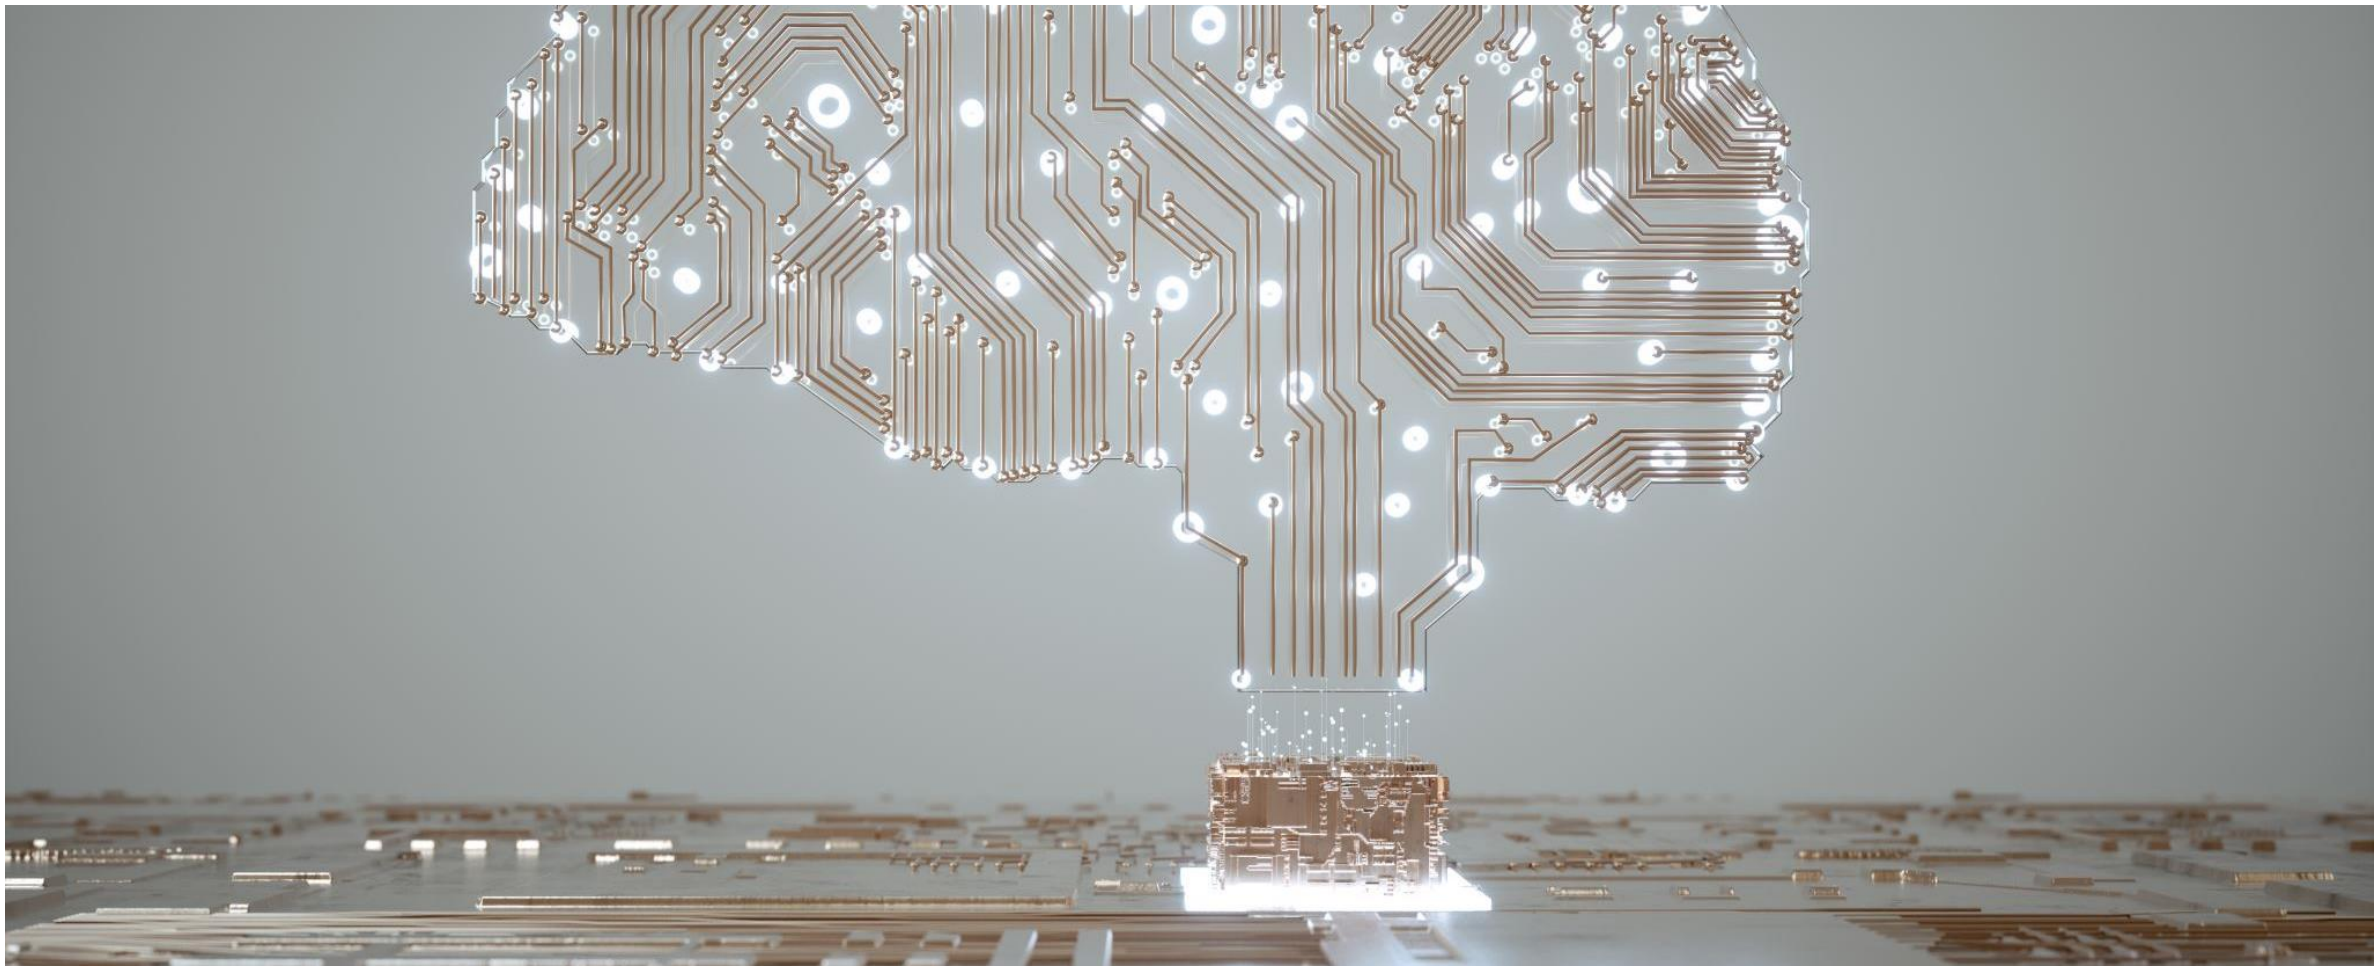

The Long Arc:
Positioning Agentic AI Historically

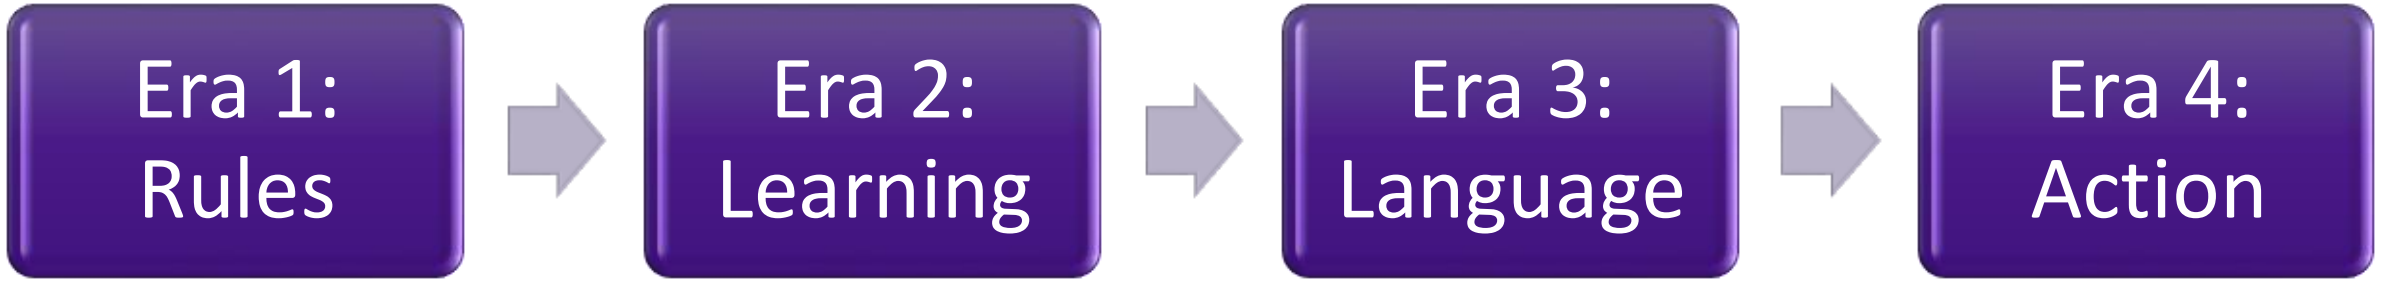

Era 1:
Rules-Based Systems (1950s to 1990s)

- Approach:
  - “If X, then Y”
- Early AI succeeded in structured, predictable domains using explicit logic:
  - Expert systems, decision trees, rule engines
- Strengths:
  - Transparent and auditable (“we can explain the rule”)
  - Reliable in narrow contexts
- Limitations:
  - Hard to scale: rule maintenance becomes expensive
  - Brittle when reality doesn’t match assumptions
  - Poor with ambiguity, language, and changing environments
- Historical lesson:
  - Organizations learned that automation must be **governed**, not just deployed.

Era 2.1:
Statistical Machine Learning (1990s to 2015)

- Approach:
  - “Learn patterns from structured data”
- The rise of ML reframed AI as prediction from examples
  - Spam detection, churn prediction, recommendation systems
- Strengths
  - Handles complexity better than rules
  - Works well when the goal is prediction/classification
- Limitations:
  - Most ML systems didn’t plan
  - Outputs often required humans or software pipelines to act on them
  - ML was “inside the product,” not “across workflows”
- Historical lesson:
  - ML helped us predict, but it didn’t autonomously execute.

Era 2.2:
Deep Learning (2012 to 2020)

- Approach
  - Extra insights from unstructured inputs
- Strengths:
  - High accuracy in perception-like tasks
  - Automation improved dramatically in specific functions
- Limitations:
  - Still not naturally interactive or goal-seeking
  - Required specialized teams and training pipelines
  - Not accessible to most knowledge workers directly
- Historical lesson:
  - Deep learning gave AI strong insights, but not conversational or agentic behavior.

Era 3:
Generative AI + LLM's (2020 to 2023)

- Approach:
  - "Language interfaces for intelligence"
- Strengths:
  - Universal interface: language becomes the API
  - Works across domains without retraining in many cases
- Limitations:
  - LLMs primarily generate responses, not completed workflows
  - They can be confident-but-wrong (hallucinations)
  - They don't inherently verify actions or outcomes
  - Practical gap that emerged:•"We can produce text faster… but the work still isn't done."
- Historical lesson:
  - GenAI scaled content creation but didn't reliably close the loop on execution.

Era 4:
Agentic AI (2023 to present)

- Approach:
  - “Goal-seeking systems that act”
- The critical change:
  - AI systems now plan steps and use external tools
  - The model becomes part of an “execution loop,” not just a “response engine”
- What agentic AI adds beyond generative AI:
  - Planning: break a goal into tasks
  - Tool use: call APIs, search databases, trigger workflows
  - Statefulness: track progress across steps
  - Verification: check whether results are correct
  - Escalation: hand off to humans when uncertainty is high

Agentic AI Unlocks

- Workflow completion
  - Not just content generation
- Consistency at scale
  - SOPs, checklists, routing
- Human oversight where it matters
  - Approvals + audits

We've moved from AI that answers → AI that acts

Enablers of Agentic AI

- Smarter AI
  - Models can now reason through complex, multi-step tasks
- Connected Enterprise Systems
  - AI can securely interact with enterprise applications and data
- Demand for Productivity
  - AI can coordinate work across multiple systems
- Organizations expect business outcomes
  - Shift in focus from creating to completing work

Key Characteristics of Agentic AI

Agentic artificial intelligence is an AI capability where a technological system can independently set goals, plan multi-step actions, and execute tasks with minimal human oversight by acting as a proactive "agent" rather than a reactive "tool" by understanding context, reasoning, and adapting to new information to achieve objectives

- Can act independently
- Goal directed
- Planning and decomposition
- Action-taking
- Awareness
- Adaptability

Agentic AI = Goals + Actions + Autonomy (with guardrails)

Generative AI versus Agentic AI

Dimension	Generative AI	Agentic AI
Primary purpose	Generate content (text, images, code)	Achieve goals by taking actions
Typical output	Answers, drafts, summaries	Completed tasks + outcomes
Core capability	Language generation	Planning + tool use + execution
Interaction pattern	Single-turn or short back-and-forth	Multi-step workflows over time
Tool use	Optional (often none)	Central (APIs, apps, databases, workflows)
Autonomy	Low—responds to prompts	Moderate to high—chooses next steps
Verification	May not verify; relies on user review	Built-in checks, validation, escalation
Risk profile	Wrong info → low direct impact	Wrong action → operational impact

The Long Arc: Positioning Agentic AI Historically

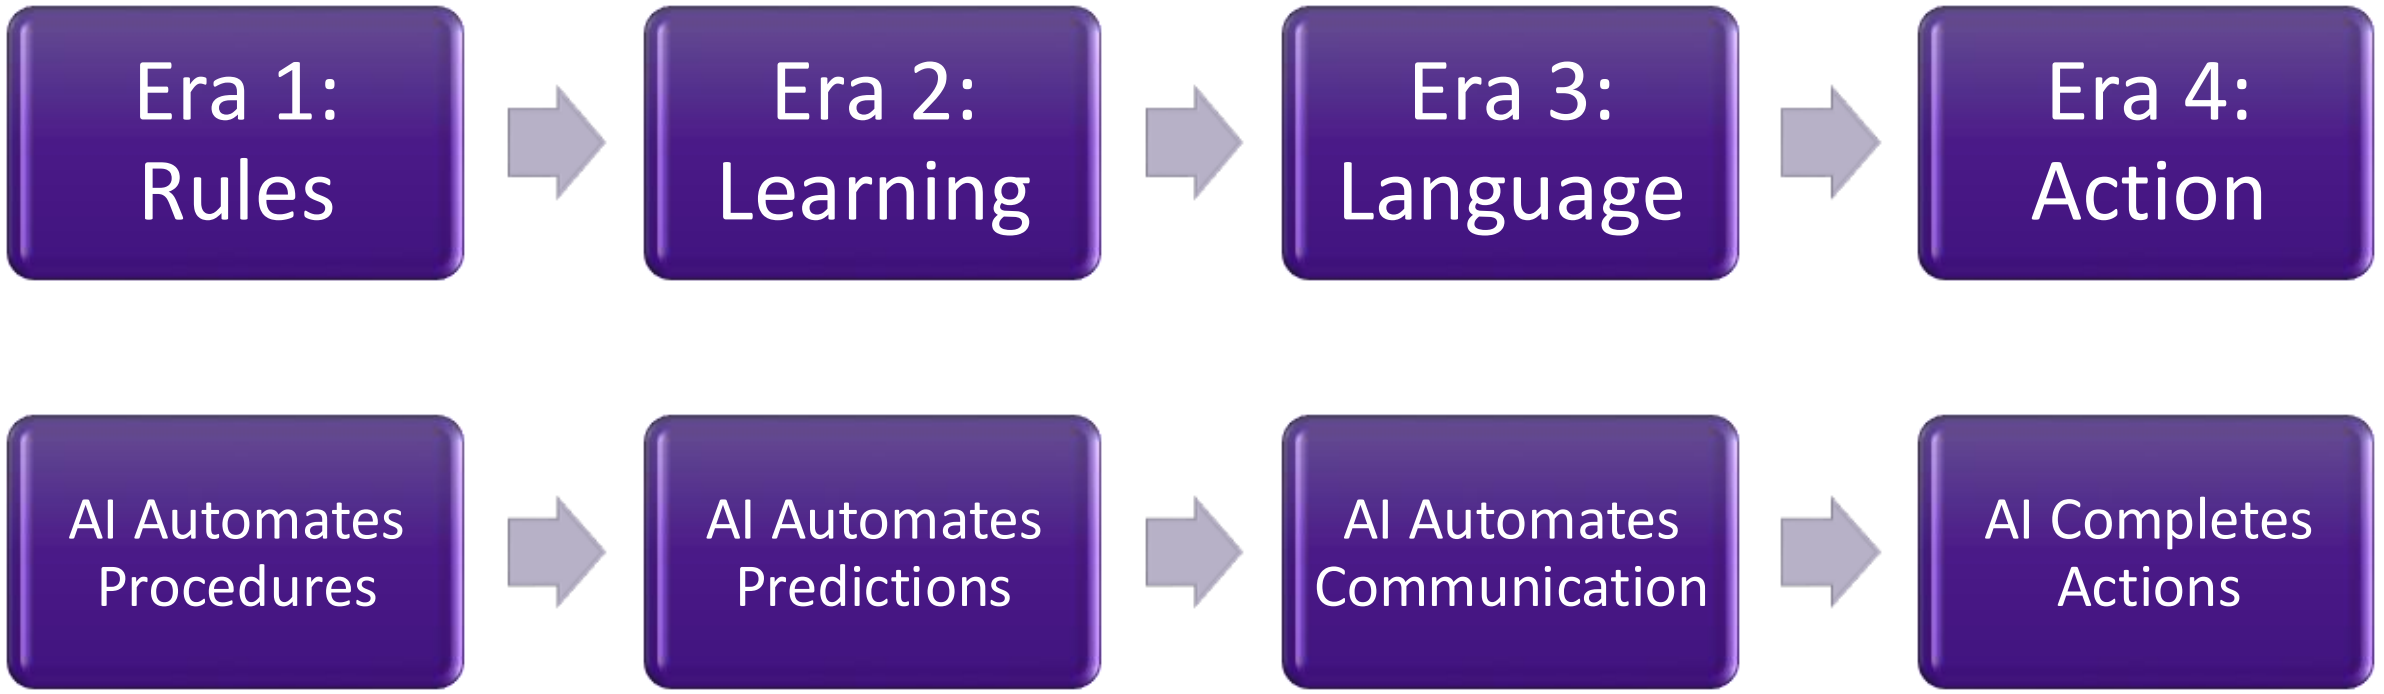

What is an Agent?

- An AI agent is a system that can decide what to do next to achieve a goal, and can use tools to complete steps on your behalf
- Core Capabilities:
  - Goal-directed planning (breaks tasks into steps)
  - Tool use (APIs, databases, workflow tools)
  - Verification / self-checks (did it work?)
  - Escalation (asks a human when uncertain)
  - Memory/state (optional; retains context over time)

Agentic AI is not….

- General intelligence
  - It is still narrow and can be brittle
- “Autopilot for everything”
  - Good agents operate within constraints
- Automatically reliable
  - It can be confidently wrong
- A replacement for governance, security, and accountability

The Agentic AI Loop

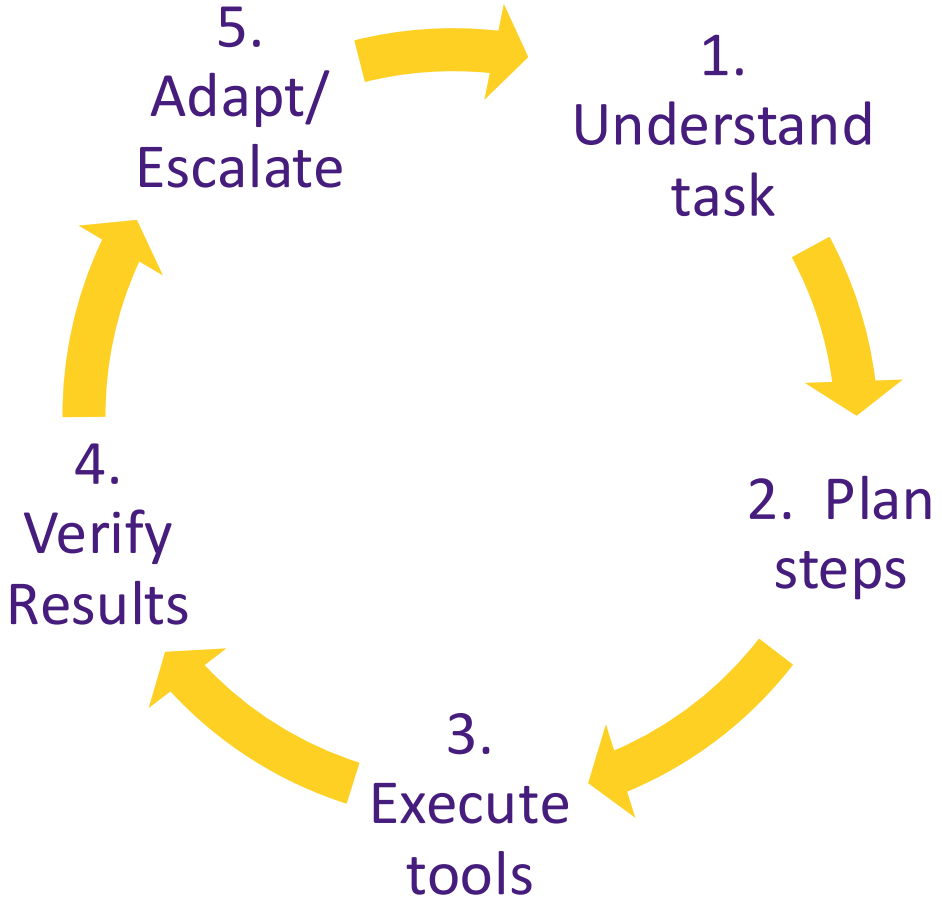

Step 1
Understand the Task

- Clarify the goal and desired outcome (what ‘done’ looks like).
- Identify constraints: time, policy, permissions, budget, risk.
- Extract key entities and context (people, systems, deadlines, data sources).
- Detect ambiguity and ask targeted questions only when necessary.
- Set the success criteria and format of deliverables.

Agent ingests the request, frames the objective, and confirms boundaries before acting.

Step 2
Plan the Steps

- Break the goal into smaller, ordered sub-tasks.
- Select which tools/systems are needed for each step (email, calendar, CRM, database).
- Choose an autonomy level: suggest-only vs execute-with-approval vs execute-then-notify.
- Anticipate failure modes (missing data, access denied, conflicting rules).
- Create a lightweight execution plan with checkpoints.

Agent produces an actionable plan—what to do, in what order, using which tools.

Step 3
Execute via Tools

- Call the appropriate tools/APIs to perform actions (query, draft, schedule, update).
- Follow least-privilege permissions and policy constraints.
- Record intermediate results and decisions (for traceability).
- Handle errors with retries, fallbacks, or alternative tools.
- Keep actions atomic and reversible when possible.

Agent performs work in external systems to move from intent to real-world changes.

Step 4
Verify Results

- Validate outputs against success criteria (correctness, completeness, formatting).
- Run sanity checks (dates, names, totals, links, file locations).
- Cross-check sources when accuracy matters (retrieval/citations, system-of-record).
- Detect inconsistencies or low-confidence outcomes.
- Prepare a concise status report of what was done and what remains.

Agent checks whether the actions actually achieved the intended outcome.

Step 5
Adapt or Escalate

- If results are insufficient, revise the plan and try the next best step.
- Escalate to a human when uncertainty is high or the action is high-stakes.
- Request approvals for irreversible actions (payments, external commitments).
- Log learnings and update memory/state for continuity (where allowed).
- Close the loop with next steps and handoff artifacts (emails, tickets, notes).

Agent adjusts to reality, involves humans when needed, and completes the workflow responsibly.

Tools are the Game Changer

- Without tools: “trapped in text”
- With tools: end-to-end workflow completion
- More capability → more risk → need stronger permissions
- Practical mantra: “capability is access; safety is constraint”
- Tools examples
  - Email + calendar
  - CRM (Salesforce), ticketing (ServiceNow/Jira)
  - Databases / SQL / analytics
  - Document retrieval / knowledge bases
  - Workflow automation (Zapier/Power Automate)

AGENTIC AI IN ACTION

Month-End Close Assistant

One agent coordinates the close—humans retain judgment and approval.

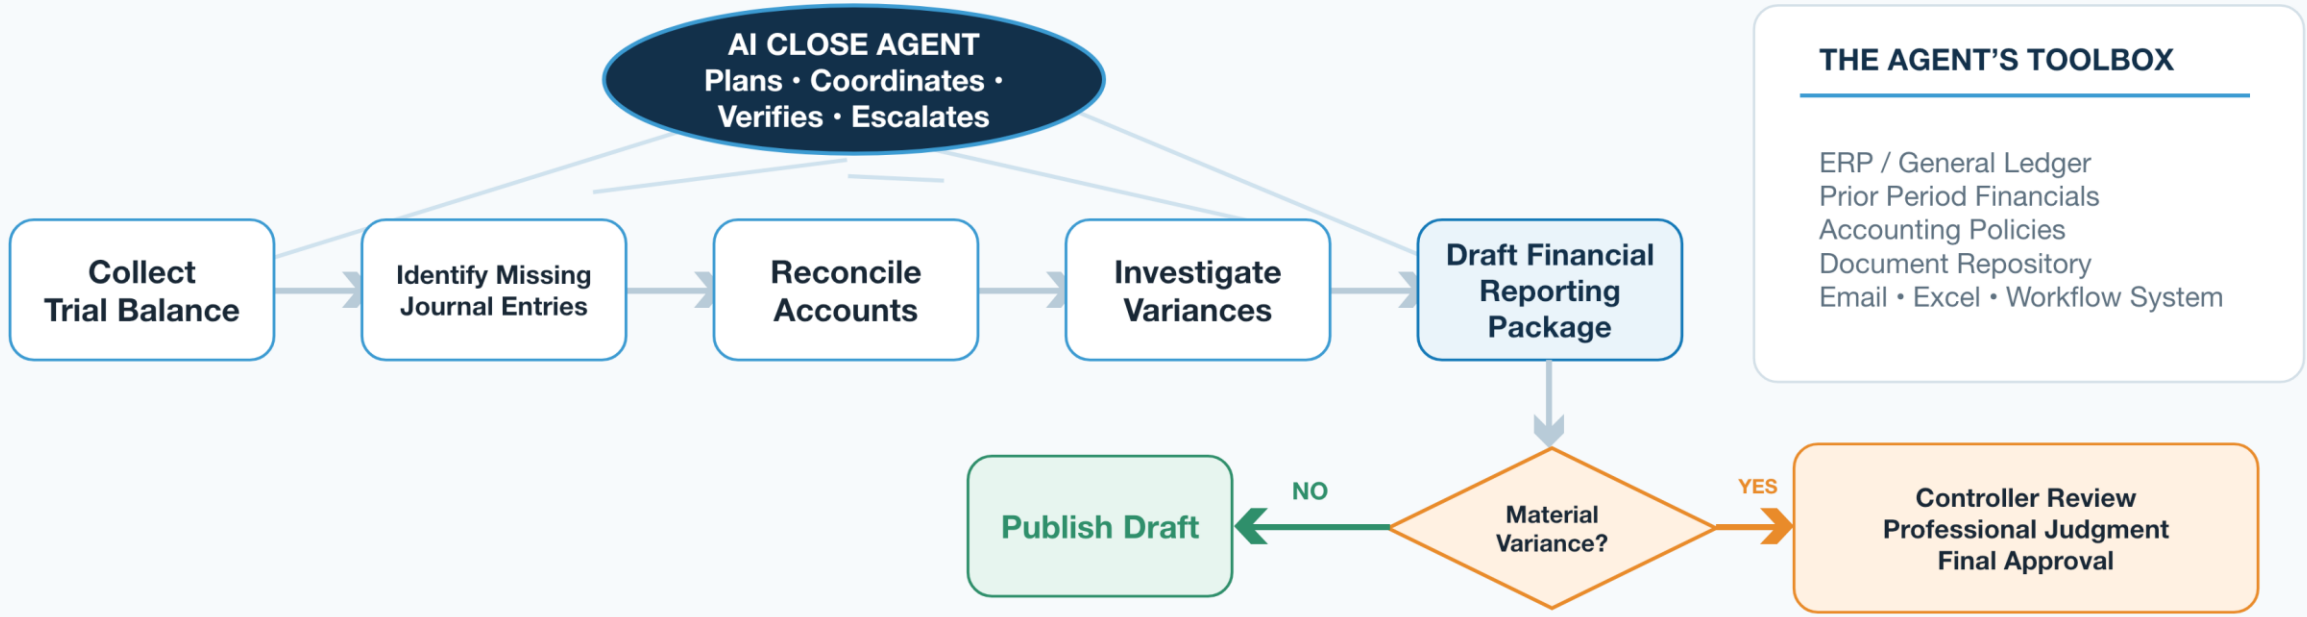

AI HANDLES

Data collection
Reconciliations
Variance analysis
Draft reporting

HUMANS HANDLE

Material judgments
Complex exceptions
Final approvals
Financial accountability

The agent prepares the work. The controller owns the decision.

From Request to Result

How Agentic AI Completes Business Work

A governed sequence converts intent into action, evidence, and an outcome.

Agentic AI transforms business requests into completed work—not just generated content.

The AI as the Conductor

The agent orchestrates work across enterprise systems.

Inside an AI Agent

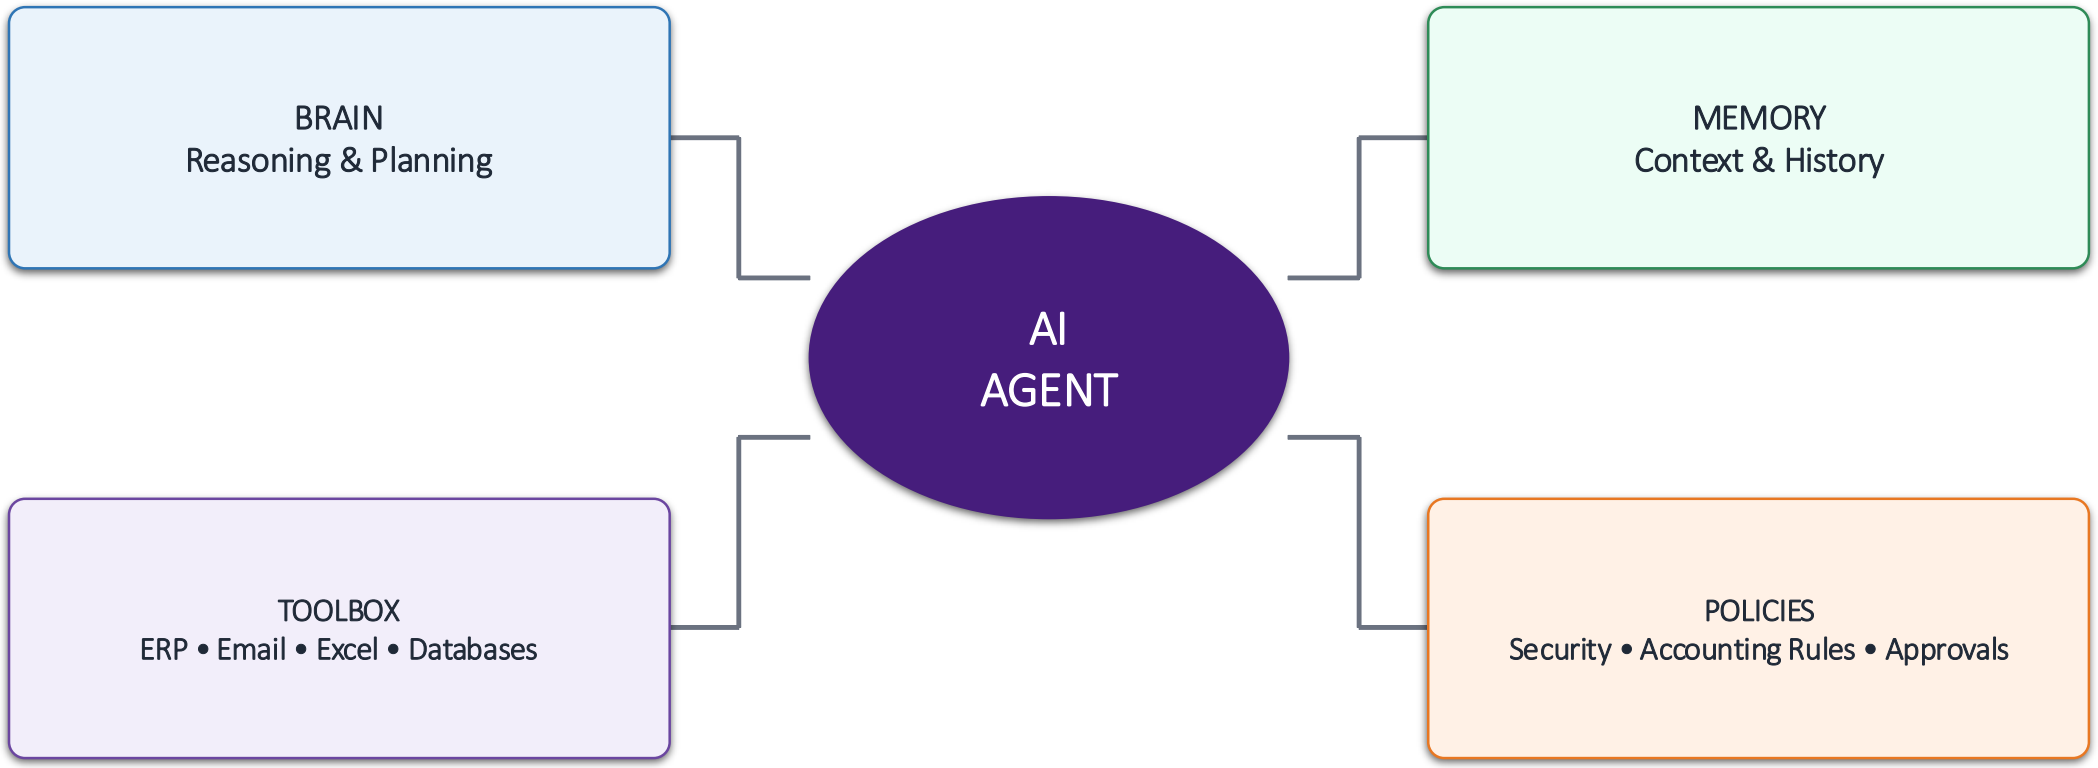

An agent combines reasoning, context, tools, and governance.

Example: AI-Assisted Month-End Close

Routine work advances automatically; judgment and approval remain with the controller.

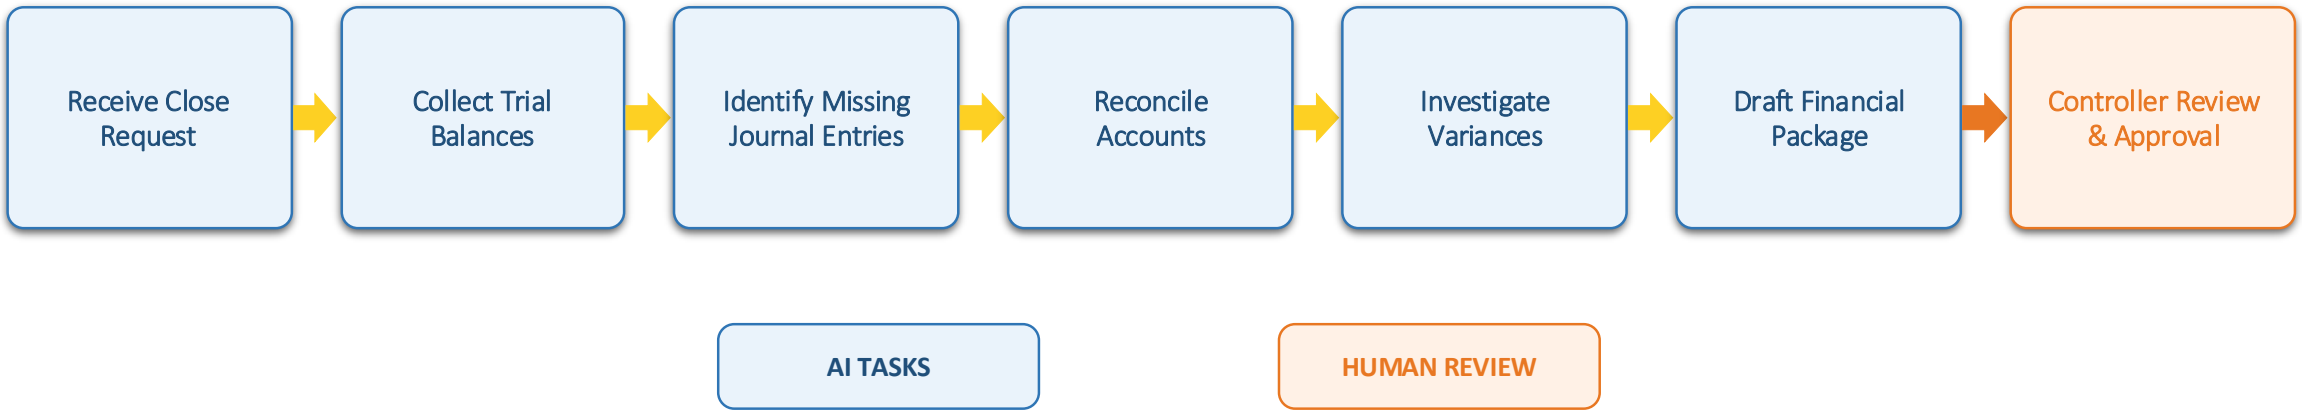

The agent handles repeatable close work; the controller owns judgment and approval.

Where Humans Stay in Control

AI HANDLES

- Data Collection
- Reconciliations
- Variance Detection
- Draft Reporting

HUMANS HANDLE

- Professional Judgment
- Material Exceptions
- Financial Accountability
- Final Approval

AI augments accountants—it does not replace professional judgment.

Building Trust in Agentic AI

1

TRANSPARENCY

"Show me what happened"

Actions and decisions are visible.

2

VERIFICATION

"Show me why"

Results are checked against rules and evidence.

3

AUDITABILITY

"Provide evidence"

Records support review, controls, and assurance.

Trust is built through visibility, validation, and governance.

Levels of Autonomy within Agentic AI

Level 0	Level 1	Level 2	Level 3
Suggest only	Execute with approval	Execute then notify	Execute autonomously (with constraints)

Recommendation: start at Level 1 for most business workflows, then increase autonomy only after you have strong logging + evaluation.

Keeping the Human-in-the-loop

- Human checkpoints
  - Money-moving actions (payments, refunds, pricing exceptions)
  - Legal commitments and contracts
  - Customer-facing statements in sensitive contexts
  - Safety-critical decisions (medical, physical safety)
  - Exceptions and novel edge cases
- Best practices
  - Show the agent's plan and evidence
  - Make approvals one-click where possible
  - Escalate only when needed (confidence thresholds)
  - Record decisions to improve future policy and training

Humans are not a bottleneck — they are the governance layer.

Business Value and Agentic AI

- Value drivers
  - Time saved per case
  - Cycle-time reduction (faster throughput)
  - Lower error and rework rates
  - Higher capacity (do more with same team)
  - Improved customer experience / satisfaction
- Advice on defining business value
  - Start with high-volume workflows
  - Prefer “bounded” tasks with clear definitions of done
  - Measure before/after on a small cohort
  - Increase autonomy only when metrics are stable

KPIs for Agentic AI

Operational metrics

- Time-to-resolution / cycle time
- Throughput (cases per week)
- Escalation rate
- Cost per case

Quality metrics

- Accuracy audits / sampling
- Rework rate
- Policy compliance rate
- Customer satisfaction (if applicable)

Adoption metrics

- Active users
- Repeat usage
- Opt-out frequency
- Time saved per user

Tip: pick 3–5 primary KPIs per pilot. Too many metrics becomes noise.

Use Cases Beyond Accounting: Where Agentic AI Creates Value

Operations

Triage, routing, SOPs • Workflow automation • Reporting

Customer & Revenue

Lead qualification • Proposals & onboarding • Renewals

Finance & Risk

Exceptions & approvals • Policy checks • Audit prep

IT & Analytics

Incident response • Knowledge base • Ask-your-data

HR & Learning

Training plans • Recruiting support • Policy Q&A

Agentic AI Ecosystem

Category	Examples
Consumer AI Agents	ChatGPT Agent Mode, Gemini, Claude
Productivity Agents	Microsoft Copilot
Enterprise Workflow Agents	Salesforce Agentforce, ServiceNow AI Agents
Custom Business Agents	Organization-specific agents built on enterprise platforms

ChatGPT: Conversation vs. Agent Mode

Conversation Mode	Agent Mode
Responds to individual prompts	Works toward a defined goal
User directs every step	AI plans and executes multiple steps
Primarily generates content	Produces completed deliverables
Limited tool use	Uses available tools and connected systems
User manages the workflow	AI manages the workflow
Best for brainstorming and drafting	Best for complex, multi-step business tasks

ChatGPT @ Work:
Latest Release Changes

- ChatGPT is becoming an Agent
  - Earlier versions produced text – latest version can complete larger tasks
- Improved planning
- Better Tool use available through plugins
- Richer deliverables
- More persistent workflows

- But….this is still not autonomous

Best Practices of Agent Mode

- Add verification checkpoints (“confirm these 3 assumptions”)
- Require sources for claims (when researching)
- Use human approval for sending/publishing decisions
- Store workflows in a Custom GPT for repeatability

Value of Agentic AI in ChatGPT

Autonomy

Risk

Higher Risk / Low Autonomy

Research + recommendations

- Use sources & citations
- Score options & tradeoffs
- Human validates facts

Higher Risk / Higher Autonomy

External actions (guardrails)

- Sending/posting needs integrations
- Approvals required
- Security + permissions matter

Low Risk / Low Autonomy

Draft + structure (safe default)

- Outlines, checklists, rubrics
- Rewrite emails & memos
- Summaries of your notes

Low Risk / Higher Autonomy

Workflow orchestration

- Meeting → action plan
- Create deliverables in one run
- “Approve & send” model

Rule of thumb: automate low-risk work fully; keep humans in the loop as risk increases.

Choosing your First Pilot

- Target:
  - Clear inputs and outputs
  - Known playbook / SOP
  - Measurable success (time, quality, throughput)
  - Low risk if wrong (or easy to undo)
  - High frequency and repetition
  - Available tool integrations (APIs) or safe proxies
- Avoid:
  - High-stakes, irreversible actions (payments, terminations, legal commitments)
  - Ambiguous success criteria (“make it better”)
  - No data access and no tools (agent can’t actually act)
  - Highly political or compliance-sensitive domains without oversight
  - Unstable requirements and shifting ownership

Choosing your First Pilot

- Target:
  - Clear inputs and outputs
  - Known playbook / SOP
  - Measurable success (time, quality, throughput)
  - Low risk if wrong (or easy to undo)
  - High frequency and repetition
  - Available tool integrations (APIs) or safe proxies
- Avoid:
  - High-stakes, irreversible actions (payments, terminations, legal commitments)
  - Ambiguous success criteria (“make it better”)
  - No data access and no tools (agent can’t actually act)
  - Highly political or compliance-sensitive domains without oversight
  - Unstable requirements and shifting ownership

Risks of Agentic AI

- Reasons for failure
  - Hallucinated assumptions (invented facts)
  - Tool misuse (wrong record, wrong environment)
  - Overstepping permissions
  - Automation bias (humans stop checking)
  - Silent failure (agent looks done but isn’t)
- Why this is new
  - Actions can change systems of record
  - Errors can propagate across multiple tools
  - Higher need for logging, approvals, and rollback
  - Security becomes central (not optional)

Essential Guardrails

- Access controls
  - Role-based access + least privilege
  - Tool allow-lists and safe defaults
  - Approval gates for high-impact steps
  - Rate limits + spend limits
- Quality and safety controls
  - Verification steps (checks, constraints, sanity tests)
  - Confidence thresholds → ask a human
  - Logging + traceability (audit trail)
  - Evaluation: test suites, red-teaming, regression checks

An Implementation Roadmap for Agentic AI

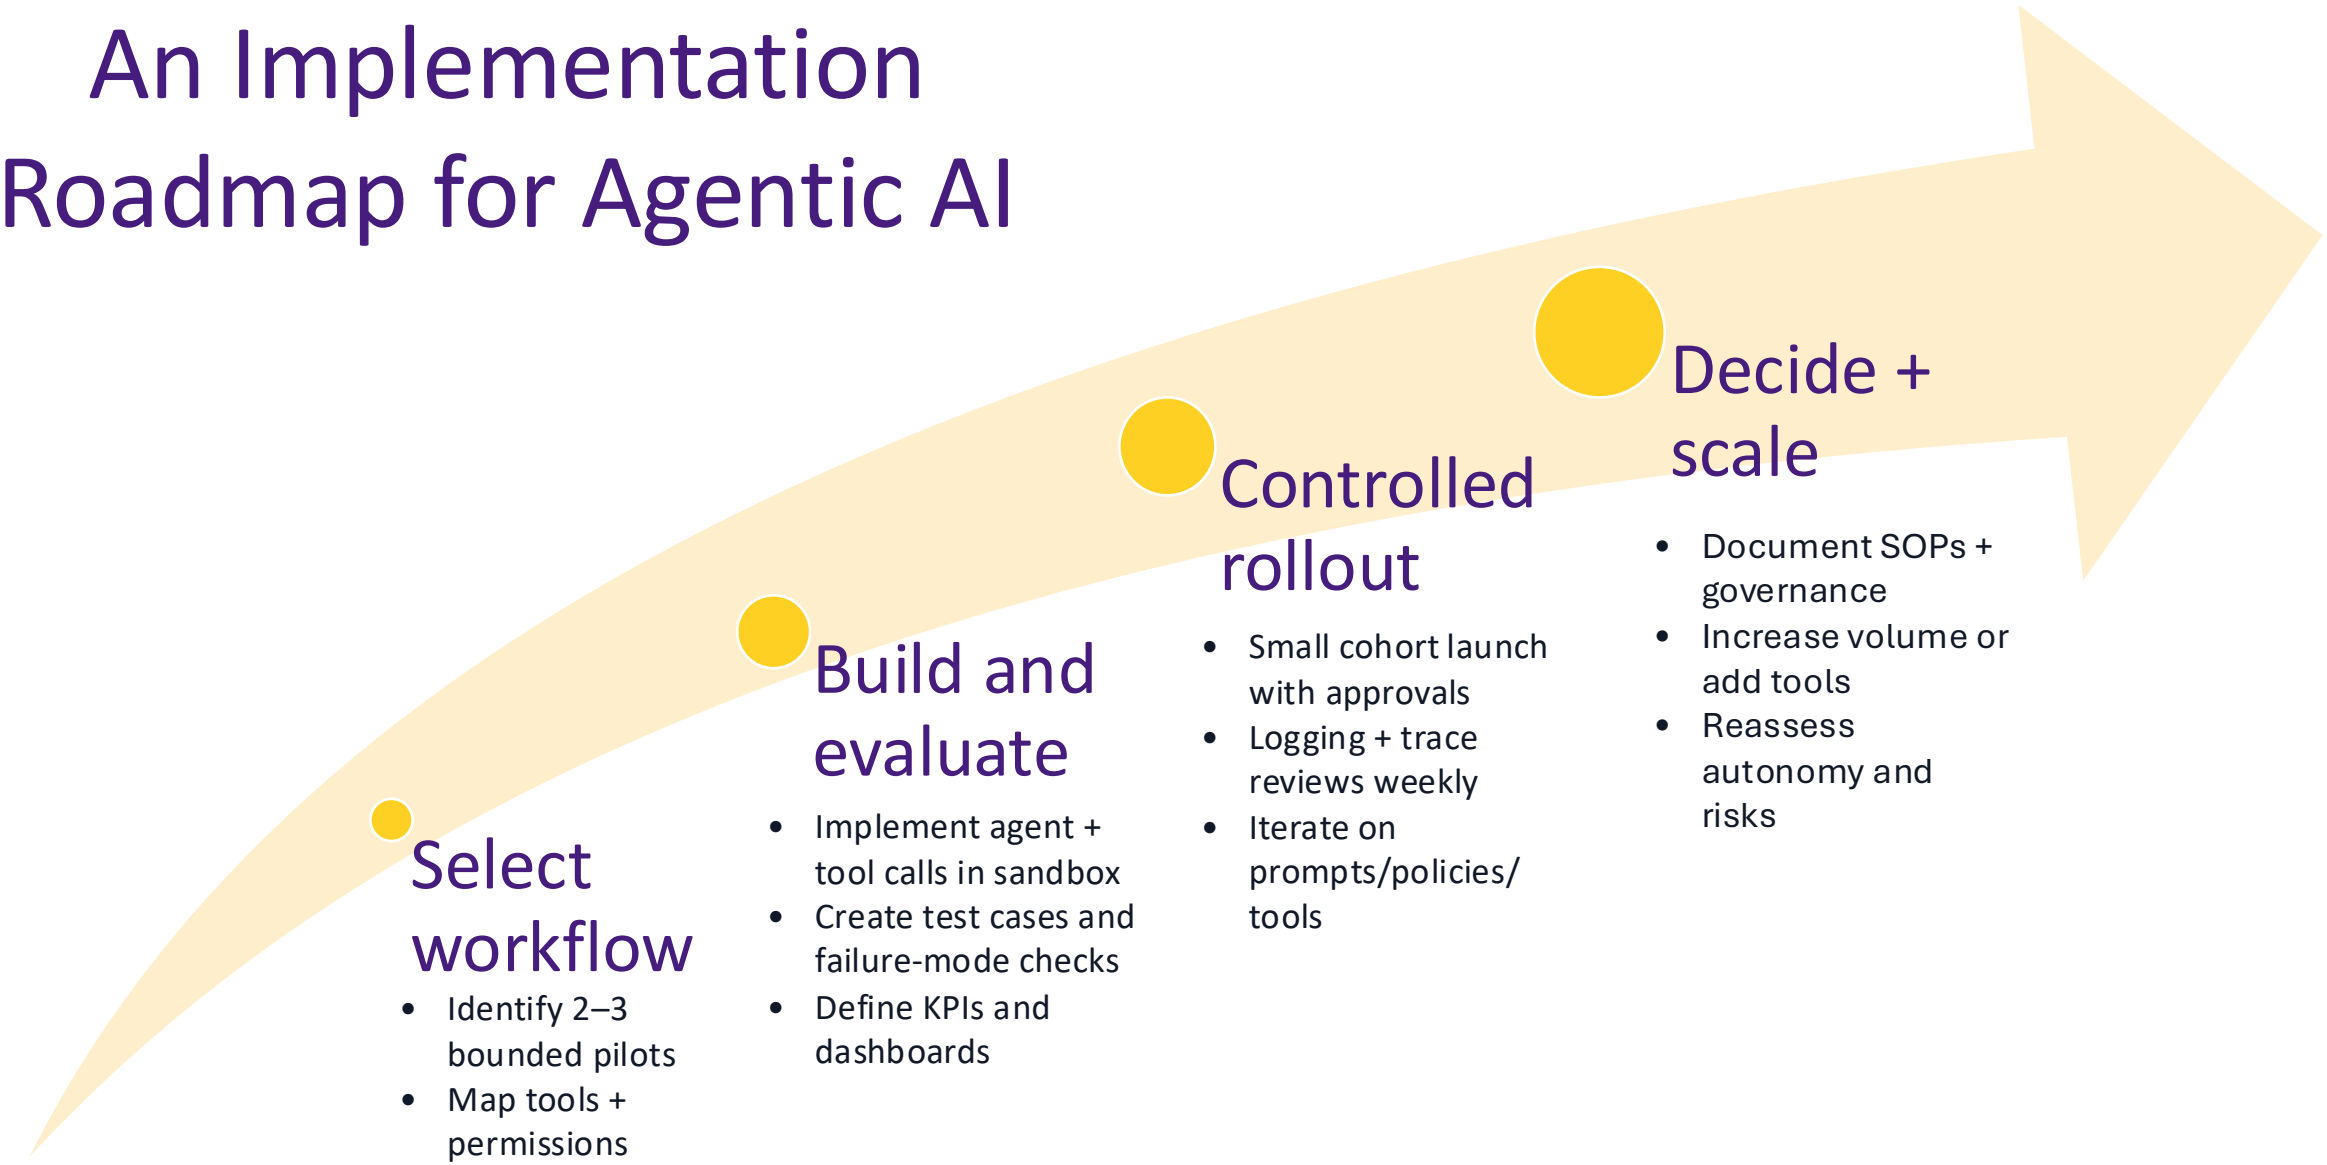

Select workflow

- Identify 2–3 bounded pilots
- Map tools + permissions
- Classify risk + autonomy level

Build and evaluate

- Implement agent + tool calls in sandbox
- Create test cases and failure-mode checks
- Define KPIs and dashboards

Controlled rollout

- Small cohort launch with approvals
- Logging + trace reviews weekly
- Iterate on prompts/policies/ tools

Decide + scale

- Document SOPs + governance
- Increase volume or add tools
- Reassess autonomy and risks

Key Takeaways:
5 Ideas to Remember

- Agentic AI = goal-driven, tool-using AI systems
- Most value comes from workflow integration, not clever prompts
- Governance requires autonomy levels + guardrails
- Start with bounded, measurable pilots and clear KPIs
- Scale responsibly: expand tools → stabilize metrics → increase autonomy

Thank you for attending!

Dr. Andrew Schwarz
aschwarz@lsu.edu

LSU | E. J. Ourso College of Business
Stephenson Department of
Entrepreneurship & Information Systems