Morgan Stanley



Building Better Agents: Designing for Enterprise Work

Microsoft Ignite NYC Community Summit

Kristen Holmes | November 16, 2026

About Me



Kristen Holmes

Morgan Stanley · Application Infrastructure

Power Platform · Copilot Studio · AI adoption

kristen.holmes@morganstanley.com

Should These Two Get the Same Approval Process?

Agent A

Built in Copilot Studio

About 5 minutes, no code

Built by a business user

Published to Teams

Agent B

Built in Copilot Studio

About 5 minutes, no code

Built by a business user

Published to Teams

Should These Two Get the Same Approval Process?

Agent A:

Lunch Menu Helper

Who uses it	my team
Data	the cafeteria menu
Actions	answers questions
If it's wrong	someone gets the wrong soup

Agent B:

Customer Follow-Up Agent

Who uses it	the whole sales org
Data	customer records
Actions	updates records, sends emails as you
If it's wrong	a customer gets the wrong email

THE PROMISE

Guardrails, not gates

What's Changed

Agents vs. traditional development

Why Gates Don't Work

Where the risk shows up

Finding the Right Path

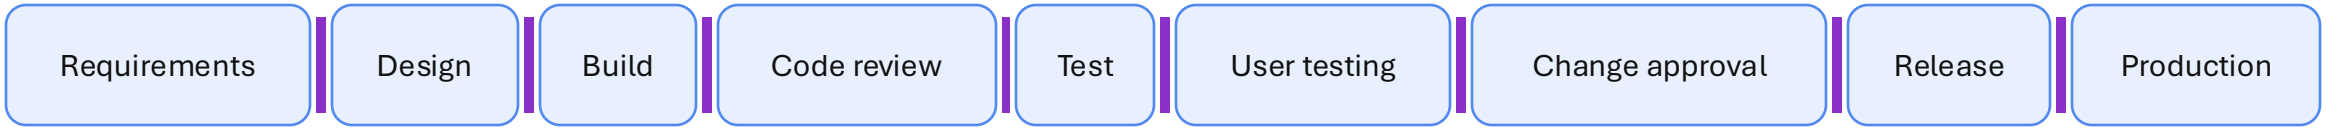
Four questions before you build

Designing the Agent

What goes inside, and how much access it gets

Traditional Development vs. Agents

Traditional development · weeks to months



Agent development · minutes



46% of leaders say their companies use agents to fully automate workflows or processes

<1% → 33% of enterprise software apps forecast to include agentic AI, 2024 to 2028

Sources: Microsoft (2025), Work Trend Index; Gartner (2025), agentic AI forecast.

When the Code Is Natural Language

	Traditional code	Natural-language agents
What you review	Lines of code, shown as a diff	Instructions, knowledge sources, tool permissions
What a small change does	Changes a specific piece of logic	A reworded sentence can change behavior everywhere
Testing	Same input, same output: pass or fail	Probable answers, scored across many test runs
What can change it	A developer deploying a change	An edited instruction, an updated document or a new model version
Who can review it	A developer	The business expert (is the answer right?) and a technologist (is the access right?)
Adaptability	Weeks to change	Changes in seconds, and can drift on its own

Adaptability is both a superpower, and it's also the risk.

Sources: AWS, Prompt, agent and model lifecycle management; Microsoft (2026), Manage the AI model lifecycle.

How People Use AI at Work Today

48%

uploaded company information into public AI tools

66%

relied on AI output without checking it

56%

made mistakes at work because of AI

Only 47% had any AI-related training.

The training paradox

Correlation, not causation

Trained employees reported more complacent or inappropriate use: 63% vs. 46%.

Source: Gillespie et al. (2025), Trust, Attitudes and Use of AI: A Global Study, University of Melbourne & KPMG. Survey of 48 ,340 adults in 47 countries.

Lock It Down or Open It Up?

LOCK IT DOWN

Gates on everything

→ People work around the blocks

OPEN IT UP

No gates at all

→ No owners and no inventory

45% of employees in one survey said they find workarounds to access blocked AI applications.

Guardrails let you drive at 65. A toll gate every mile just sends people looking for another road.

Source: UpGuard (2025), The State of Shadow AI. Vendor survey of 1,000 employees and 500 security leaders.

If You Give an Executive an Agent

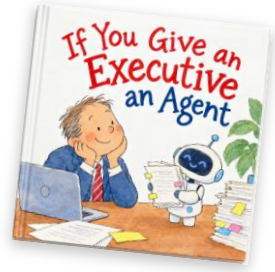
Front-door review: Approved ✓

Reads his email and drafts replies he checks.

After launch: no review

> *If you give an executive an agent, he'll want it to read his email.*
>
> *If it reads his email, he'll want it to draft his replies.*
>
> *If it drafts his replies, he'll want it to send them for him.*
>
> *If it sends his emails, his whole team will want one too.*
>
> *If his whole team has one, they'll want it to approve their expenses.*
>
> *And one quiet Tuesday, the model gets updated…*

…and that's how Agent A becomes Agent B.



- Trigger created
- Instructions edited
- Tool or action added
- Audience expanded
- Tool or action added
- Model updated

Where Should the Review Happen?

STEP 1	STEP 2	STEP 3	STEP 4	STEP 5	STEP 6
Reads email	Drafts replies	Sends as him	Whole team	Approves expenses	Model updated
Guardrails	Guardrails	Review	Depends	Review	Review

Guardrails are enough

Depends on reach

Should trigger a review

Finding the Right Path

Maker

Who's building it?
Do they have the skills yet?





Impact

Who does it reach?
What can it do?
What if it's wrong?

Connectors

What data and tools does it need?
Whose identity does it use?





Support

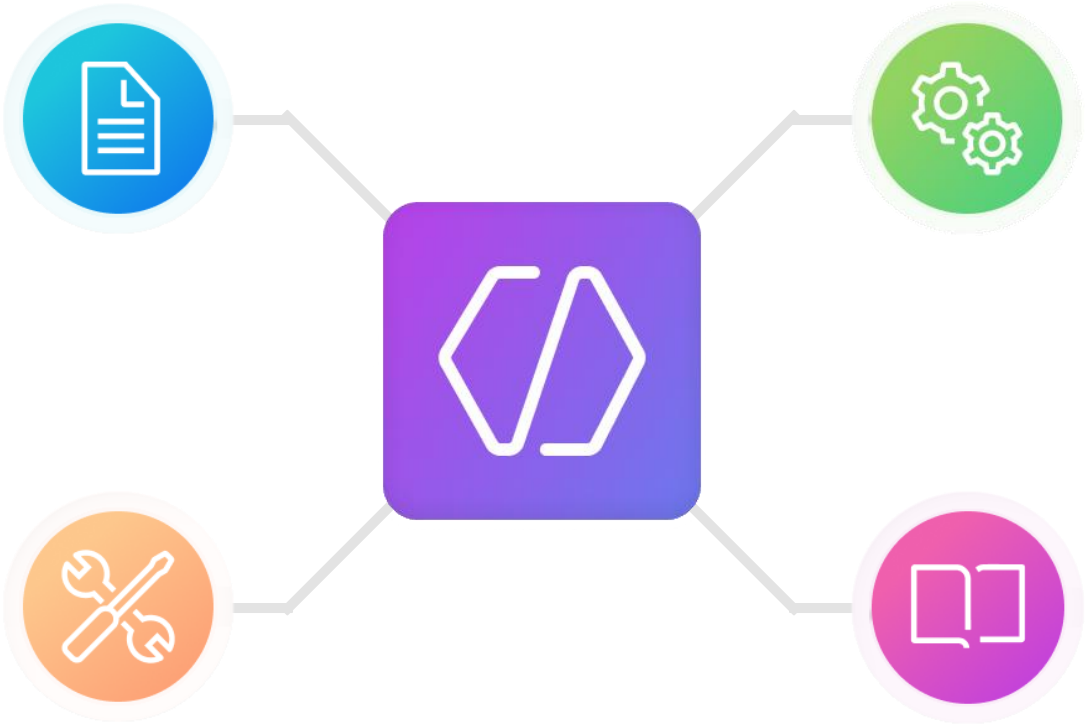
Who owns and tests it?
Who retires it?
Who needs called at 2am?

Inside an Agent

SET THE DIRECTION

Instructions

What the agent is for, and what's out of scope.



DEFINE THE APPROACH

Skills

Shared steps, written once and reused across agents.

PROVIDE CAPABILITIES

Tools

Only the tools it needs, with the least access. Fixed steps for anything you can't undo.

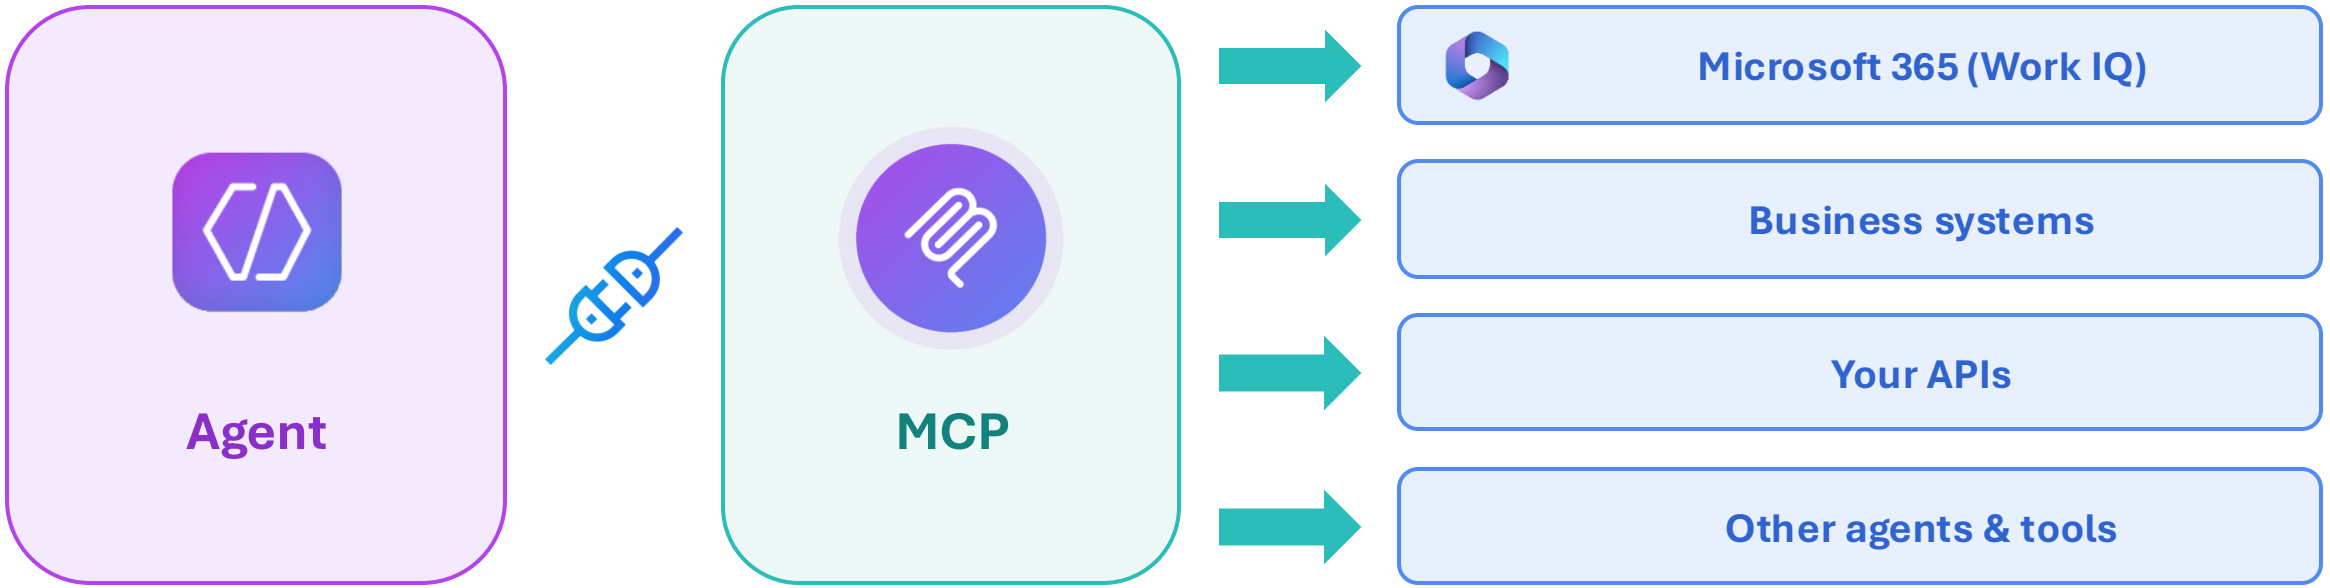
GROUND YOUR AGENT

Knowledge

Only the sources the job actually needs.

Sources: Anthropic (2024), Building effective agents; OpenAI (2026), A practical guide to building agents.

Connecting Tools with MCP



Sources: Model Context Protocol, Introduction (modelcontextprotocol.io); Microsoft Learn (2026), Work IQ MCP overview.

Work IQ



Inference

Memory

Data

WORKLOADS

Mail	Calendar	Teams	SharePoint	OneDrive	Word	People

Sources: Microsoft Learn (2026), Work IQ MCP overview; Work IQ in Microsoft Copilot Studio.

Excessive Agency

Too much functionality
Tools it doesn't need

Too-broad permissions
Access wider than the job

Too much autonomy
Acts without anyone checking

With MCP, the model picks the tool, so:

- Least privilege on every tool
- A person approves anything you can't undo
- Only trusted servers, reviewed when they change

Sources: OWASP (2025), LLM06:2025 Excessive Agency; OWASP Cheat Sheet Series, MCP Security Cheat Sheet.

Defaults vs. What the Job Needs

On by default

- Added to every new agent
- All workloads selected
- Write operations on

Narrow it

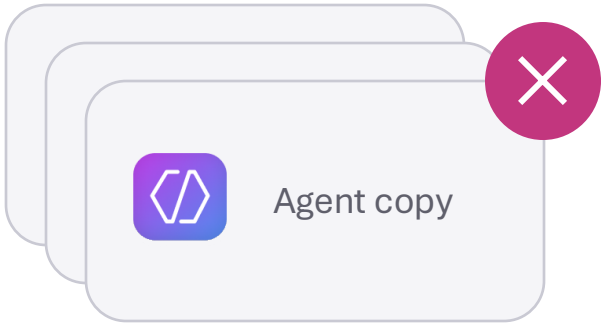
What the job needs

- Only the workloads it uses
- Read-only until trusted
- Runs as the user

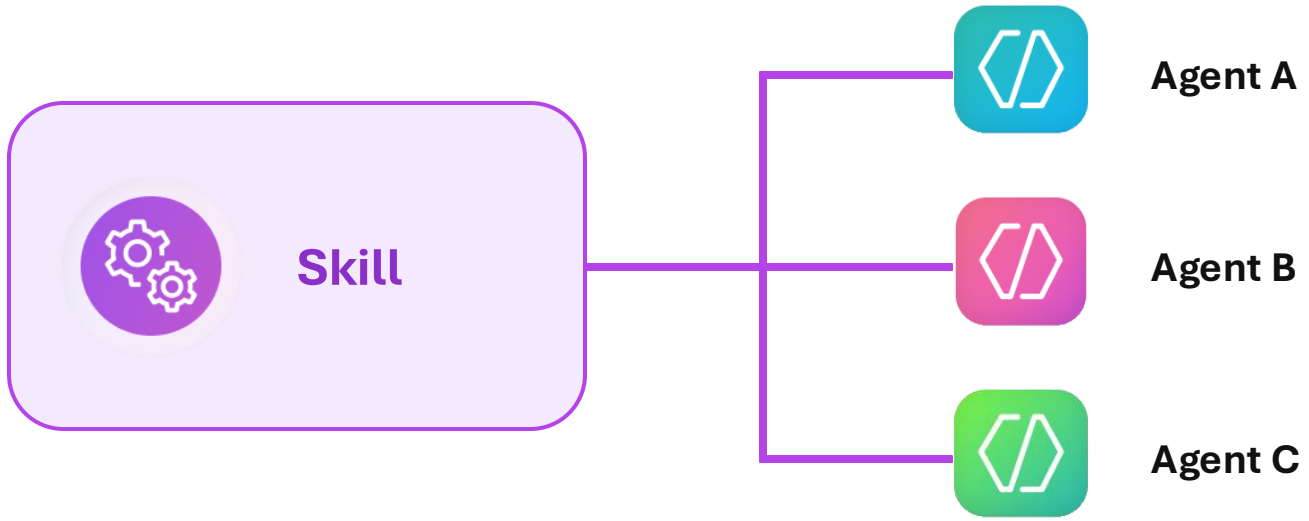
Source: Microsoft Learn (2026), Work IQ in Microsoft Copilot Studio.

Reusable Skills

INSTEAD OF



Copying whole agents



Check for an approved skill before you build a new one.

Sources: Microsoft Learn (2026), Skills overview for agents; Build or reuse skills.

Key Takeaways



Find the right path

Four questions
before you build.



Scope what you’re given

Only the data, tools and
access the job needs.



Design for change

Review when it changes,
not just when it launches.

Guardrails, not gates.

Thank you · kristen.holmes@morganstanley.com

References (1 of 3)

Al-Otaibi, M. (2026, April 29). Aramco's citizen developers accelerate digital transformation. MEED. https://www.meed.com/aramcos-citizen-developers-accelerate-digital-transformation

Amazon Web Services. (n.d.). Prompt, agent, and model lifecycle management. AWS Prescriptive Guidance. https://docs.aws.amazon.com/prescriptive-guidance/latest/agentic-ai-serverless/prompt-agent-and-model.html

Anthropic. (2024, December 19). Building effective agents. https://www.anthropic.com/engineering/building-effective-agents

Cloud Security Alliance. (2026, April 21). Autonomous but not controlled: AI agent incidents now common in enterprises [Press release]. https://cloudsecurityalliance.org/press-releases/2026/04/21/new-cloud-security-alliance-survey-reveals-82-of-enterprises-have-unknown-ai-agents-in-their-environments

Gartner. (2025, June 25). Gartner predicts over 40% of agentic AI projects will be canceled by end of 2027 [Press release]. https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027

Gartner. (2026, May 26). Gartner says applying uniform governance across AI agents will lead to enterprise AI agent failure [Press release]. https://www.gartner.com/en/newsroom/press-releases/2026-05-26-gartner-says-applying-uniform-governance-across-ai-agents-will-lead-to-enterprise-ai-agent-failure

Gillespie, N., Lockey, S., Ward, T., Macdade, A., & Hassed, G. (2025, April 28). Trust, attitudes and use of artificial intelligence: A global study 2025. University of Melbourne & KPMG. https://doi.org/10.26188/28822919

Microsoft. (2025, April 23). 2025: The year the Frontier Firm is born. Work Trend Index. https://www.microsoft.com/en-us/worklab/work-trend-index/2025-the-year-the-frontier-firm-is-born

Microsoft. (2025, July 28). Configure user authentication for Microsoft Copilot Studio. Microsoft Learn. https://learn.microsoft.com/en-us/microsoft-copilot-studio/configuration-end-user-authentication

References (2 of 3)

Microsoft. (2026). Configure data policies for agents. Microsoft Learn. https://learn.microsoft.com/en-us/microsoft-copilot-studio/admin-data-loss-prevention

Microsoft. (2026). Establish an ALM strategy. Microsoft Copilot Studio guidance. https://learn.microsoft.com/en-us/microsoft-copilot-studio/guidance/alm

Microsoft. (2026). Manage the AI model lifecycle. Microsoft Copilot Studio guidance. https://learn.microsoft.com/en-us/microsoft-copilot-studio/guidance/plan-agent-model-lifecycle

Microsoft. (2026). Work IQ in Microsoft Copilot Studio. Microsoft Learn. https://learn.microsoft.com/en-us/microsoft-copilot-studio/agents-experience/add-work-iq

Microsoft. (2026, March 3). Environment routing. Microsoft Learn. https://learn.microsoft.com/en-us/power-platform/admin/default-environment-routing

Microsoft. (2026, March 6). Share an agent with other users. Microsoft Learn. https://learn.microsoft.com/en-us/microsoft-copilot-studio/admin-share-bots

Microsoft. (2026, August 18). Secure and govern Microsoft Copilot: Foundational deployment guidance. Microsoft Learn. https://learn.microsoft.com/en-us/microsoft-365/copilot/secure-govern-copilot-foundational-deployment-guidance

Microsoft. (2026, September 30). Build or reuse skills. Microsoft Learn. https://learn.microsoft.com/en-us/microsoft-365/copilot/extensibility/build-reuse-skills

Microsoft. (2026, October 5). Microsoft Purview DLP for Microsoft 365 Copilot and Cowork. Microsoft Learn. https://learn.microsoft.com/en-us/purview/dlp-microsoft365-copilot-location-learn-about

References (3 of 3)

Microsoft. (2026, October 11). Restrict discovery of SharePoint sites and content. Microsoft Learn. https://learn.microsoft.com/en-us/sharepoint/restricted-content-discovery

Microsoft. (2026, October 13). Work IQ MCP overview. Microsoft Learn. https://learn.microsoft.com/en-us/microsoft-365/copilot/extensibility/work-iq/mcp/overview

Microsoft. (2026, October 21). Skills overview for agents. Microsoft Learn. https://learn.microsoft.com/en-us/microsoft-copilot-studio/agents-experience/skills-overview

Model Context Protocol. (n.d.). Introduction. https://modelcontextprotocol.io/docs/getting-started/intro

OpenAI. (2026, January 15). A practical guide to building agents. https://openai.com/business/guides-and-resources/a-practical-guide-to-building-ai-agents/

OWASP Foundation. (n.d.). MCP security cheat sheet. OWASP Cheat Sheet Series. https://cheatsheetseries.owasp.org/cheatsheets/MCP_Security_Cheat_Sheet.html

OWASP Foundation. (2025). LLM06:2025 Excessive agency. OWASP Gen AI Security Project. https://genai.owasp.org/llmrisk/llm062025-excessive-agency/

UpGuard. (2025). The state of shadow AI. https://www.upguard.com/resources/the-state-of-shadow-ai

Xia, B., Lu, Q., Zhu, L., Xing, Z., Zhao, D., & Zhang, H. (2025). Evaluation-driven development of LLM agents [Preprint]. arXiv. https://arxiv.org/html/2411.13768v2

Slide 10 is an affectionate parody of Laura Numeroff's If You Give a Mouse a Cookie (1985).

Disclosures

This material is provided for informational purposes only and does not provide individually tailored investment or financial advice. It has been prepared without regard to the individual financial circumstances and objectives of persons who receive it. Information and data contained herein is from multiple sources considered to be reliable and Morgan Stanley makes no representation as to the accuracy or completeness of the information or data from sources outside of Morgan Stanley. Past performance is not indicative or a guarantee of future results.

Environmental, Social and Governance ("ESG") investments in a portfolio may experience performance that is lower or higher than a portfolio not employing such practices. Portfolios with ESG restrictions and strategies as well as ESG investments may not be able to take advantage of the same opportunities or market trends as portfolios where ESG criteria is not applied. There are inconsistent ESG definitions and criteria within the industry, as well as multiple ESG ratings providers that provide ESG ratings of the same subject companies and/or securities that vary among the providers. Certain issuers of investments may have differing and inconsistent views concerning ESG criteria where the ESG claims made in offering documents or other literature may overstate ESG impact. ESG designations are as of the date of this material, and no assurance is provided that the underlying assets have maintained or will maintain and such designation or any stated ESG compliance. As a result, it is difficult to compare ESG investment products or to evaluate an ESG investment product in comparison to one that does not focus on ESG. Investors should also independently consider whether the ESG investment product meets their own ESG objectives or criteria.

There is no assurance that an ESG investing strategy or techniques employed will be successful. Past performance is not a guarantee or a dependable measure of future results.

Morgan Stanley Smith Barney LLC, its affiliates and Morgan Stanley Financial Advisors do not provide legal or tax advice. Each client should always consult his/her personal tax and/or legal advisor for information concerning his/her individual situation and to learn about any potential tax or other implications that may result from acting on a particular recommendation.

Bonds are subject to interest rate risk. When interest rates rise, bond prices fall; generally the longer a bond's maturity, the more sensitive it is to this risk. Bonds may also be subject to call risk, which is the risk that the issuer will redeem the debt at its option, fully or partially, before the scheduled maturity date. The market value of debt instruments may fluctuate, and proceeds from sales prior to maturity may be more or less than the amount originally invested or the maturity value due to changes in market conditions or changes in the credit quality of the issuer. Bonds are subject to the credit risk of the issuer. This is the risk that the issuer might be unable to make interest and/or principal payments on a timely basis. Bonds are also subject to reinvestment risk, which is the risk that principal and/or interest payments from a given investment may be reinvested at a lower interest rate.

Equity securities may fluctuate in response to news on companies, industries, market conditions and general economic environment. Companies paying dividends can reduce or stop payouts at any time.

Alternative investments often are speculative and include a high degree of risk. Investors could lose all or a substantial amount of their investment. Alternative investments are appropriate only for eligible, long-term investors who are willing to forgo liquidity and put capital at risk for an indefinite period of time. They may be highly illiquid and can engage in leverage and other speculative practices that may increase the volatility and risk of loss. Alternative Investments typically have higher fees than traditional investments. Investors should carefully review and consider potential risks before investing.

- Loss of all or a substantial portion of the investment due to leveraging, short-selling, or other speculative practices;
- Lack of liquidity in that there may be no secondary market for a fund;
- Volatility of returns;
- Restrictions on transferring interests in a fund;
- Potential lack of diversification and resulting higher risk due to concentration of trading authority when a single advisor is utilized;
- Absence of information regarding valuations and pricing;
- Complex tax structures and delays in tax reporting;
- Less regulation and higher fees than mutual funds;
- Risks associated with the operations, personnel, and processes of the manager; and
- Risks associated with cybersecurity.

Disclosures

The sole purpose of this material is to inform, and it in no way is intended to be an offer or solicitation to purchase or sell any security, other investment or service, or to attract any funds or deposits. Investments mentioned may not be appropriate for all clients. Any product discussed herein may be purchased only after a client has carefully reviewed the offering memorandum and executed the subscription documents. Morgan Stanley Wealth Management has not considered the actual or desired investment objectives, goals, strategies, guidelines, or factual circumstances of any investor in any fund(s). Before making any investment, each investor should carefully consider the risks associated with the investment, as discussed in the applicable offering memorandum, and make a determination based upon their own particular circumstances, that the investment is consistent with their investment objectives and risk tolerance.

As a diversified global financial services firm, Morgan Stanley engages in a broad spectrum of activities including financial advisory services, investment management activities, sponsoring and managing private investment funds, engaging in broker-dealer transactions and principal securities, commodities and foreign exchange transactions, research publication, and other activities. In the ordinary course of its business, Morgan Stanley therefore engages in activities where Morgan Stanley's interests may conflict with the interests of its clients, including the private investment funds it manages. Morgan Stanley can give no assurance that conflicts of interest will be resolved in favor of its clients or any such fund.

Past performance is no guarantee of future results. Actual results may vary. Diversification does not assure a profit or protect against loss in a declining market.

Alternative investments involve complex tax structures, tax inefficient investing, and delays in distributing important tax information. Individual funds have specific risks related to their investment programs that will vary from fund to fund. Clients should consult their own tax and legal advisors as Morgan Stanley Wealth Management does not provide tax or legal advice.

Interests in alternative investment products are only made available pursuant to the terms of the applicable offering memorandum, are distributed by Morgan Stanley Smith Barney LLC and certain of its affiliates, and (1) are not FDIC-insured, (2) are not deposits or other obligations of Morgan Stanley or any of its affiliates, (3) are not guaranteed by Morgan Stanley and its affiliates, and (4) involve investment risks, including possible loss of principal. Morgan Stanley Smith Barney LLC is a registered broker-dealer, not a bank.

Artificial intelligence (AI) is subject to limitations, and you should be aware that any output from an AI-supported tool or service made available by the Firm for your use is subject to such limitations, including but not limited to inaccuracy, incompleteness, or embedded bias. You should always verify the results of any AI-generated output.

Morgan Stanley at Work services are provided by Morgan Stanley Smith Barney LLC, member SIPC, and its affiliates, all wholly owned subsidiaries of Morgan Stanley.

© 2026 Morgan Stanley & Co. LLC and Morgan Stanley Smith Barney LLC. Members SIPC.

CRC # 5350120 4/8/2026

Morgan Stanley

Appendix

How to Review an Agent

Three smaller reviews instead of one code review

Intent

Is the job clear and narrow? Do the instructions say what we mean?

Who: Business owner

Access

What data can it reach, whose permissions does it use, what can it do?

Who: Technology / platform

Behavior

Does it answer correctly across a set of tests, and does it still?

Who: Both, mostly automated

Review again when any of these change

- Instructions edited
- Knowledge source added
- Tool or action added
- Model updated
- Audience expanded

93% of academic agent evaluations happen before deployment; only 4% are continuous.

Source: Xia et al. (2025), Evaluation-driven development of LLM agents (preprint).

Maker Access Levels

Training unlocks access to advanced tools.

Colors match the review levels: guardrails only, lightweight SDLC, full SDLC



STEP 1

Explorer

Approved tools
Non-sensitive data
No actions

STEP 2

Builder

Core training
Own environment
Standard connectors and reusable components

STEP 3

Publisher

Named owner
Audience and data review
Test evidence
Support and retirement plan

STEP 4

Privileged builder

Advanced tools or sensitive data
Least privilege
Security review
Renewal and monitoring

Example: Aramco reports more than 2,000 citizen developers, trained and certified before they build, with risky features switched off.

Example source: Al-Otaibi, M. (2026). Aramco's citizen developers accelerate digital transformation. MEED. Self-reported case.

Keeping Track of Agents

What agent debt looks like

- **Orphaned agents**
- **Duplicates**
- **Accidental production**
- **Stale knowledge**
- **Copied prompts**

82%

of IT and security professionals surveyed found an AI agent they didn't know about in the past year

21%

have a formal process to retire agents

Vendor-sponsored survey: directional

Every agent gets

- ✓ An owner
- ✓ An inventory record
- ✓ A defined audience
- ✓ A retirement plan

Reuse approved skills before building new ones.

Sources: Cloud Security Alliance (2026, April 21), survey of 418 IT and security professionals, commissioned by Token Security; Microsoft Learn (2026), Build or reuse skills.

Which Agents Need SDLC?

The highest-risk answer sets the level

	Agent A		Agent B
	Guardrails only	**Lightweight SDLC**	**Full SDLC**
Who uses it	Just me	My team	Across the firm, or customers
What data	My own, or public	Shared team data	Confidential or regulated
What it does	Reads and answers	Drafts for a person to approve	Writes to systems or takes actions
How it connects	Standard connectors	Shared connectors	Custom connectors, APIs, MCP
How it runs	User in the loop	Scheduled, with review	Autonomous triggers
If it's wrong	I notice	My team is inconvenienced	A customer or the business is affected

One answer in a right-hand column moves the whole agent there.

What Each Review Level Includes

	Guardrails only	Lightweight SDLC	Full SDLC
Always-on guardrails	✓	✓	✓
Named owner	–	✓	✓
Test set, re-run on every change	–	✓	✓
Version history	–	✓	✓
A way to report issues	–	✓	✓
Separate dev, test and production, with pipelines	–	–	✓
Change review and approval	–	–	✓
Security review	–	–	✓
Monitoring and a rollback plan	–	–	✓
	Learn and experiment	Team scale	Business critical

Sources: Microsoft Copilot Studio guidance, Establish an ALM strategy; Gartner (2026).