guide

Grow Commercial Transformation with AI You Can Trust

Considerations for SME and mid-market leaders to plan AI, automation and orchestration intelligence with confidence

Commercial Transformation

Contents

Introduction

AI, automation and orchestration intelligence present a practical opportunity for SME and mid-market organisations to improve commercial performance. Used well, they can help leadership grow revenue, reduce operating cost, remove manual drag, improve employee experience and create more confident management visibility. The pressure to adopt is real, but the strategic question is not simply "how quickly can we use AI?" The sharper NSG question is "where can Microsoft-aligned AI, automation and orchestration improve the commercial operating model safely and measurably?"

For all the interest around AI, business leaders are also right to be cautious. Poorly governed AI can expose sensitive data, create unreliable outputs, introduce compliance risk and generate more operational noise. In the following sections, we outline the considerations leaders should plan for when using Microsoft AI, Power Platform, Copilot, Azure AI, Fabric and security tooling as part of a trusted Commercial Transformation strategy.

Establishing responsible and secure AI practices helps your organisation implement intelligent tools without creating avoidable risk. For SME and mid-market businesses, this is not about building an enterprise bureaucracy. It is about setting the right operating controls: who owns the use case, what data is used, what Microsoft environment it sits in, who reviews outputs, how risks are monitored and how value is measured. As AI regulation increases, responsible AI foundations will help leaders adopt innovation with more confidence and less operational exposure.

NSG POSITIONS

Trustworthy AI as a commercial operating requirement, not a compliance afterthought. With Microsoft capabilities that improve security, safety, privacy, governance and observability, SME and mid-market firms can use and build AI solutions that support growth, control and operational intelligence.

42%

of business leaders said they were **'equally concerned and excited'** about generative AI.

Practice responsible AI

Consider responsible AI principles

Public policy and industry best practice are still evolving alongside AI technology. That leaves SME and mid-market leaders needing practical guidance on how to adopt AI systems in a way that improves business outcomes without creating unnecessary risk for customers, employees, data or commercial decision-making.

Taking time to define a responsible AI approach helps leadership move forward with confidence. For NSG, this means connecting responsible AI to the commercial operating model: how work is done, how data moves, how decisions are made, how Microsoft tools are governed and how teams are supported. The six responsible AI principles below provide a strong foundation for any SME or mid-market business planning AI, automation and orchestration intelligence.

These principles include:

- Privacy and security
- Reliability and safety
- Accountability
- Inclusiveness
- Transparency
- Fairness

01.
Privacy and security

Ensure your AI projects serve strategic business objectives

AI systems should follow the same privacy and security standards that your business applies to its most sensitive customer, employee, financial and operational data.

Know where data is located, who can access it, how it is used and whether it is protected at rest and in transit.

Check that AI tools align with your company's policies for privacy, security, customer confidentiality and commercial sensitivity.

When you implement strong permissions across Microsoft 365, SharePoint, Teams, OneDrive, Dynamics and other connected systems, only authorised individuals can access sensitive information. This reduces the risk of internal oversharing and helps prevent AI tools from surfacing content to the wrong people.

You can also maintain security and meet compliance requirements with governance solutions that retain, log and monitor interactions with AI applications. This helps detect organisational policy violations, investigate incidents and understand how AI is being used across the business.

AI implementation must also comply with local laws and regulations covering data use, privacy and security. For SME and mid-market leaders, the safer commercial position is to be overcautious early, especially where customer data, financial information, HR records, contracts, board material or regulated content is involved.

An example of privacy and security in practice:

Use of an AI tool to analyse customer communication and resolve a support ticket could involve access to sensitive or identifiable customer data. Understanding local laws and regulations, applying strong Microsoft 365 permissions, and enforcing adequate controls can help keep that sensitive data private while still improving service speed.

COMPLYING WITH AI REGULATIONS

NSG is committed to building products and solutions that comply with regulations like the EU AI Act to help customers use AI compliantly. NSG helps translate that product commitment into practical SME and mid-market operating controls.

02.
Reliability and safety

Reliability and safety mean AI systems perform as expected, without avoidable errors, interruptions or harmful outputs. Microsoft and other AI tool providers are responsible for product-level testing and documentation, but the business still needs oversight to verify that each tool is working safely for its intended commercial use.

Systems should undergo regular monitoring, maintenance, feedback and evaluation to identify new use patterns, troubleshoot issues, improve outputs and ensure the AI system remains fit for purpose.

Regularly conduct stress testing.

Stress testing prepares an AI system, workflow or agent to handle the type and volume of use it is expected to support without producing avoidable errors or creating unmanaged risk.

Red teaming is a type of stress testing that simulates real-world attacks and uses techniques commonly used to gain access to secure systems. In 2018, Microsoft established its dedicated AI Red Team and has expanded the team's mission to map risks beyond traditional security, including risks from non-adversarial users compromising responsible AI standards. For example, red teaming a generative AI tool may involve testing whether a user can generate content that stereotypes a marginalised group using the tool.

An AI model can also be red teamed to identify potential misuse, scope capabilities and understand limitations. The insights can then be applied to future versions of the model so it operates more reliably and safely.

For SME and mid-market businesses, due diligence should happen before purchase, before rollout and throughout usage. This is especially important when AI outputs support financial reporting, customer communication, compliance review, operational decisions or commercial planning.

COMPLYING WITH AI REGULATIONS

NSG is committed to building products and solutions that comply with regulations like the EU AI Act to help customers use AI compliantly. NSG helps translate that product commitment into practical SME and mid-market operating controls.

An example of reliability and safety in practice:

An AI tool is used to model financial outcomes and report on performance. Testing is conducted regularly to ensure the AI tool reliably produces accurate results, avoiding adverse impact on the organisation's financial health.

03.
Accountability

In many instances, responsible AI is human-centred. Clear oversight helps people remain in control of the AI, automation and orchestration tools the business implements, while staying accountable for the outcomes those tools produce.

Establish a system of oversight that clearly defines roles and responsibilities at every stage of the AI journey.

Implementing oversight, impact testing and response processes keeps people at the centre. This helps protect against adverse impacts and ensures adequate controls are built into each stage of adoption.

Ensure AI tools are fit for purpose.

Regularly assess whether AI tools still provide the right solution for the problem they were intended to solve. Determine how your organisation will respond if a tool fails, produces poor outputs, creates user confusion or no longer delivers measurable commercial value.

An example of accountability in practice:

Use of an AI tool to review legal contracts involves oversight by an individual with adequate context and expertise to verify compliance with applicable laws and regulations, and to sign off on the final outcome of the AI-supported review.

04.
Inclusiveness

At its core, inclusivity requires that AI tools are accessible to people of all abilities. The tools that SME and mid-market leaders create, configure or procure should follow accessible design principles and comply with relevant standards such as the European accessibility standard, EN 301 549; Section 508 of the U.S. Rehabilitation Act; and the Web Content Accessibility Guidelines (WCAG).

Identify opportunities to build inclusivity into your organisation with AI.

For NSG, inclusiveness is also an adoption question. If AI and automation are only usable by technically confident employees, value will remain narrow. Commercial Transformation requires tools, training and workflows that help people across functions use Microsoft capabilities responsibly and confidently.

An example of Inclusiveness in practice:

Recipients of Microsoft funding are creating a hiring platform for neurodiverse applicants, building better and more affordable Braille displays for students with visual impairment, and creating a web app to help individuals with speech disabilities communicate more effectively.

05.
Transparency

Transparency is a building block of trust. To achieve and maintain transparency, be clear about how and when AI is being used, as well as its capabilities and limitations.

Be open about how AI is implemented and used across your organisation.

Employees and stakeholders are more likely to trust AI-supported work when they understand how tools are being used, where outputs come from and where human judgement is still required. This transparency builds capability and helps teams know when to supplement AI outputs with information outside the tool's scope.

Customers may also need to know when they are interacting with an AI tool, when AI has contributed to decision-making or when an audio, visual or written asset has been generated or manipulated with AI. Business leaders should consider how and when that information will be disclosed.

An example of transparency in practice:

Generative AI is used to create content for a marketing campaign, and the organisation identifies which elements are created by AI. A specialist reviews the AI-created content to verify its accuracy so it doesn't mislead customers about the features or capabilities of the advertised product.

06.
Fairness

Fair AI implementation allocates opportunities, resources and information equitably among the people who use and are impacted by AI.

Ensure that your AI systems provide a similar quality of service and delivery of resources and opportunities to all who use them or are affected by them, across demographic groups.

When using AI tools that describe, depict or otherwise represent people, minimise the potential for stereotyping or demeaning people, especially people from marginalised groups.

Include people with different backgrounds, roles, experience levels, education levels and perspectives in the team that manages AI implementation. Identify statistical bias in datasets and include human review by subject matter experts in decisions that use AI. This helps protect against biased outcomes and gives the business a more balanced view of how tools behave in practice.

An example of fairness in practice:

An AI tool is used to review applications and identify priority candidates in the hiring process with oversight from a human resources representative. This person confirms that the tool assesses information accurately, without statistical bias, and has the final review in decision-making.

Make informed decisions about AI and safety

Implementing AI responsibly minimises risks while allowing your business to benefit from practical uses across commercial growth, revenue control, customer service, operations, finance and management intelligence. Use the following questions as conversation starters with your team as you plan Microsoft AI, automation and orchestration initiatives.

Privacy and security

- Is the data accessed by AI systems secured according to your organisation's policies for handling sensitive data?
- Are you in control of your data, including where it's stored and how it's used?
- Is your data secured at all times, including when it is in transit from one system to another?
- Do you have quality security tools in place to defend against third-party access or cyberattacks?
- Do you have threat identification and response tools in place in case of cyberattacks?

Reliability and safety

- Have the tools you intend to use been adequately tested to minimise errors?
- Do you have a plan in place to remediate any failures that occur?
- Will the tools be regularly monitored for reliability issues?
- Are you prepared to comply with all requirements to operate the tools safely?

Accountability

- Have you assessed the impact this implementation would have on your employees, organisation and customers?
- Have you established a system of oversight and response in case of potential negative impacts?
- Have you implemented data governance and management best practices?
- Have you determined who will have oversight of AI tools and ensured they have adequate training and control?
- Have you made sure this solution is fit for its intended purpose?

Inclusiveness

- Have you confirmed that the tools you intend to use meet accessible design principles?
- Do the tools comply with the European accessibility standard, EN 301 549?
- Do the tools comply with Section 508 of the U.S. Rehabilitation Act?
- Do the tools comply with the Web Content Accessibility Guidelines (WCAG)?

Transparency

- Have the tools you intend to use been adequately tested to minimise errors?
- Do you have a plan in place to remediate any failures that occur?
- Will the tools be regularly monitored for reliability issues?
- Are you prepared to comply with all requirements to operate the tools safely?

Fairness

- Have you ensured this implementation will provide the same quality of service to all affected by it?
- Have you tested the system's outputs to make sure it will fairly allocate resources and opportunities across demographic groups?
- Are outputs of the system free of stereotypes and negative portrayals of marginalised groups?

Protect your data and AI systems

Secure your AI tools now and in the future

Implemented incorrectly, any technology that has access to sensitive data can create risk. Because AI systems may use large volumes of proprietary business data, security must be prioritised from the start when assessing Microsoft AI, Power Platform automation, Copilot deployment, Azure AI services or third-party AI solutions.

Microsoft launched its Secure Future Initiative to address and prepare for the increasing scale and heightening risk of cyberattacks in the age of AI. The Secure Future Initiative identifies three principles that Microsoft upholds to help secure the wider digital ecosystem:

Secure Future Initiative principles

SME and mid-market organisations that accept and manage the learning curve are more likely to build internal confidence and realise value sooner. Progress comes from engaging with the system, improving the foundations and shaping responsible use over time.

1 Secure by design
Security comes first when designing any product or service.

2 Secure by default
Security protections are enabled and enforced by default, require no extra effort and aren't optional.

3 Secure in operations
Security controls and monitoring will be continuously improved to meet current and future threats.

Effective practices for security hygiene include

- Enabling multifactor authentication to protect against compromised user passwords and help provide extra resilience for identities.
- Applying Zero Trust principles, which involves explicit verification, use of least privileged access and assumption of breach, to limit the impact of an attack.
- Using extended detection and response and antimalware to automatically block attacks and gain insight into the security operations software for faster response.
- Ensuring systems are up to date with the latest versions of firmware, operating systems and applications.
- Implementing the right protections for critical data, which requires knowledge of which data is most important and where it's located.

Security is the essential underpinning of any AI implementation. When you ensure basic security hygiene practices, you protect your data, your people and your devices from more than 98% of cyberattacks.

This matters because organisations often want the productivity gains of AI before their Microsoft environment is ready. The right sequence is security hygiene first, then AI readiness, then controlled rollout.

Manage your AI systems with governance

A strong governance model helps build the foundation for responsible AI implementation. Governments and regulatory bodies maintain baseline requirements to minimise adverse effects of AI use on society. Commercial organisations also have a responsibility to manage their own development, configuration or use of AI systems according to company policy, customer expectations, local laws, regulatory obligations and responsible AI principles.

For organations, governance should be practical and proportionate. It should protect the business while still allowing useful innovation. The goal is not to slow down AI adoption. The goal is to make adoption trusted, observable and commercially useful.

Establishing your own governance

When creating your organisation's governance system, remember that the purpose is to align AI solutions with company policy and responsible AI principles through clear policies and procedures. This includes policies for assessing third-party AI tools, approving Microsoft Copilot or agent use cases, coordinating stakeholder involvement, documenting acceptable use, educating employees and communicating with customers where required.

Governance should also define how the business handles AI-enabled automations, Power Platform workflows, Copilot Studio agents, data connections and reporting outputs. In NSG language, this is orchestration control: understanding what has been built, who owns it, what data it touches, what decisions it supports and how risk is reviewed.

Risks

Map

Mapping risks is the first stage of governing AI and should inform decisions about a tool's safety, reliability and fitness for purpose. Mapping risks involves running AI impact assessments and privacy and security reviews, including red teaming and stress testing where appropriate.

Measure

Measuring risks involves developing metrics by which to assess identified risks and testing planned mitigation methods to determine how effective they will be. Metrics may include output accuracy, user adoption, incident volume, access exceptions, review failure rates, cost, cycle time and commercial impact.

Manage

Managing risks requires organisations to consistently monitor performance. At this stage, you should identify opportunities for user agency and educate stakeholders about responsible use. Human review and oversight should be included in the management process, alongside transparency practices aligned to responsible AI principles.

Build observability at every layer

When creating your organisation's governance system, remember that the purpose is to align AI solutions with company policy and responsible AI principles through clear policies and procedures. This includes policies for assessing third-party AI tools, approving Microsoft Copilot or agent use cases, coordinating stakeholder involvement, documenting acceptable use, educating employees and communicating with customers where required.

Although organisations want agentic AI at the heart of the business, governance pressures, regulatory scrutiny and internal risk controls keep them grounded. Without full visibility into agents, interactions, automations and data flows, AI remains difficult to trust.

To deploy agentic AI, Commercial Transformation leaders need understanding

- What agents exist?
- Who is using them?
- What systems and data do they access?
- What workloads do they drive and what outcomes do they produce?

According to **Cyber Pulse (2026**), an AI security report, 29% of employees have turned to unsanctioned AI agents.

Without clear governance and answers to the questions above, organisations risk the rise of shadow AI and double agents operating outside established controls.

These unsanctioned tools can introduce security vulnerabilities, create entry points for malicious actors, disrupt operations and increase the likelihood of sensitive data leakage.

For organisations, shadow AI is particularly dangerous because control models are often informal. A single unsanctioned agent or workflow can touch customer records, commercial documents, HR data or finance processes without leadership visibility.

Getting the most out of your AI agents

Effective observability is built on foundational capabilities that allow you to act quickly on performance, behaviour and risk signals before they affect the business:

Registry

Get a complete view of all agents in your organisation, including agents with agent ID, agents you register yourself and shadow agents.

Agent analytics

Track key agent metrics, including adoption, performance, speed, quality, cost and ROI to make more informed decisions.

Agent maps

Visualise agent usage, behaviour and trends in one comprehensive agent map to simplify monitoring and accelerate issue identification.

Role-specific oversight

Build extended visibility across your organisation through role-based reporting for security leaders to manage agent risks and business leaders to monitor business metrics and ROI.

Build observability at every layer

Observability is an organisation-wide responsibility

Observability is critical for deploying and managing trusted AI agents. This introduces responsibilities across every layer of the organisation so visibility and coordination are unified.

IT teams enable and govern every agent running in the environment. This includes operating the AI registry, managing identities and access controls and enforcing consistent policies. IT teams are also responsible for risk factors like over-permissioned agents.

Developers ensure that security, compliance and safety guardrails are implemented. They must also continuously evaluate models for vulnerabilities, monitor runtime performance and cost, observe production behaviour and apply guardrails consistently across all agents.

Security teams protect the entire AI ecosystem. They're responsible for detecting agent sprawl and inappropriate access, validating regulatory and compliance obligations, identifying data oversharing and leakage risks and evaluating AI-specific threats and vulnerabilities.

For NSG, the commercial layer also matters. Business owners should know which outcomes agents support, which workflows they affect, and whether they are improving speed, quality, control or customer experience.

Build observability at every layer

Observability is an organisation-wide responsibility

Observability is critical for deploying and managing trusted AI agents. This introduces responsibilities across every layer of the organisation so visibility and coordination are unified.

IT teams enable and govern every agent running in the environment. This includes operating the AI registry, managing identities and access controls and enforcing consistent policies. IT teams are also responsible for risk factors like over-permissioned agents.

Developers ensure that security, compliance and safety guardrails are implemented. They must also continuously evaluate models for vulnerabilities, monitor runtime performance and cost, observe production behaviour and apply guardrails consistently across all agents.

Security teams protect the entire AI ecosystem. They're responsible for detecting agent sprawl and inappropriate access, validating regulatory and compliance obligations, identifying data oversharing and leakage risks and evaluating AI-specific threats and vulnerabilities.

For NSG, the commercial layer also matters. Business owners should know which outcomes agents support, which workflows they affect, and whether they are improving speed, quality, control or customer experience.

Consulsion

This roadmap reframes AI readiness for the NSG Commercial Transformation strategy. The central argument is simple: SME and mid-market firms do not need AI theatre. They need practical operating improvement. They need to remove manual drag, improve data quality, automate the right workflows, strengthen governance, increase reporting confidence and create better management visibility.

Microsoft products can play a major role in that journey, but only when positioned properly. Copilot, Power Platform, Azure AI, Fabric, SharePoint, Teams and the wider Microsoft ecosystem should not be sold as disconnected tools. They should be orchestrated around the commercial operating model: how the business creates demand, converts revenue, serves customers, moves information, manages risk and makes decisions.

Take the next steps on your Commercial Transformation journey

01

MAP
where manual effort, document friction, weak reporting, poor handoffs and unclear ownership are limiting commercial performance.

02

PRIORITISE
the Microsoft AI, automation and orchestration use cases that can create measurable value in the next 90 days.

03

ENGAGE
with noodleSPARK Group to enable your AI, automation and Orchestration Intelligence Transformation.

noodleSPARK
a UMH company

Fix the friction between workflow drag, document friction, reporting delays, automation logic, AI prioritisation and governance.

Growing businesses rarely fail because effort is absent.

They fail because the operating logic behind growth is weak, inconsistent or too dependent on individual judgement.

NSG AI, Automation and Intelligence Advisory is designed to fix that. It helps UK B2B, B2B2C and Education organisations improve the specific operating conditions that are limiting performance in this part of the commercial system.

STRONGER TOGETHER

www.noodlespark.net

hello@noodlesparkgroup.co.uk