AdvancedBITS | Business IT Solutions

AI Readiness for M&A Firms

Building the governance, data, security, and operational foundation for responsible AI across sourcing, diligence, execution, and integration

Web publishing notes and keyword plan (remove before posting)

Title tag (60 characters): AI Readiness for M&A Firms: A Practical Guide | AdvancedBITS

Meta description (144 characters): A practical AI readiness guide for M&A firms: governance, NDAs, FINRA and NIST context, AI agents, Microsoft 365, vendors, and a 90-day roadmap.

URL: /resources/ai-readiness-for-ma-firms

Primary keywords: AI readiness for M&A firms, AI governance for M&A, AI policy for investment banks and advisory firms, generative AI in M&A.

Secondary keywords: AI acceptable use policy for deal teams, shadow AI in M&A, AI agents in dealmaking, FINRA generative AI guidance, NIST AI RMF for financial firms, Microsoft 365 Copilot for deal teams, AI vendor due diligence, AI and NDAs.

Long-tail questions covered in the FAQ: what is AI readiness for M&A firms, do FINRA rules apply to AI tools, can deal teams use ChatGPT or Copilot on deal documents, what should an AI policy for an M&A firm include, what controls do AI agents need, do we need ISO 42001.

Page roles (avoid keyword overlap): this guide owns the firm-wide readiness framework, maturity model, regulatory map, policy outline, self-check, and the adoption and governance statistics, which are not repeated in the case study. The AI in M&A Diligence case study owns the diligence scenario and NDA workflow. The M&A technology due diligence landing page owns commercial intent and the independent controls benchmark table. The Ownership Transition Technology Operating Model owns succession readiness. Link all four with descriptive anchor text.

Before publishing: confirm the SOC 2 Type II wording, confirm whether your firm is a FINRA member before keeping the FINRA row prominent, and have counsel review the regulatory map. **Schema:** Article + FAQPage. **Last updated:** October 22, 2026.

About this guide

This is a framework guide, not a client case study. It describes no client outcomes, and every statistic comes from independent third-party research. AdvancedBITS does not replace your legal, compliance, or deal advisors. It strengthens the technology foundation beneath your AI program.

> **Executive takeaway**
>
> **AI readiness for an M&A firm is the ability to use AI across sourcing, diligence, execution, and integration while keeping deal confidentiality, NDA commitments, source accuracy, access control, vendor accountability, and professional judgment intact.** It depends less on which AI tool you pick and more on governance, data, Microsoft 365 permissions, identity, vendor terms, human review, and recovery. This guide gives you a seven-domain readiness model, a four-level maturity scale, a regulatory and standards map, a two-page policy outline, controls for AI agents, a 90-day roadmap, and a ten-question self-check. AdvancedBITS helps you build the foundation first, then scale AI with confidence.

M&A work runs on speed, evidence, confidentiality, and judgment. Advisors, buyers, sellers, lenders, attorneys, accountants, and operators all need controlled access to sensitive information while timelines compress. AI can help your teams find risk faster, summarize materials, compare documents, draft first-pass memos, and prepare integration workstreams. It also exposes every weakness in how your firm already handles information.

AI is not safe because a tool is popular, or because a vendor says it is enterprise-grade. You still have to decide which tools are allowed, what data they may touch, what each NDA permits, who checks the output, and how you would prove all of it to a client or regulator. AdvancedBITS treats AI readiness as a security and operating-model challenge first and a software decision second.

Why AI readiness matters now for M&A firms

Dealmakers have moved from AI pilots to AI in production, and governance is trailing adoption. Independent research shows widespread use, real value, and consistent gaps in policy, access control, and verification.

Table 1. AI adoption and governance in M&A, by the numbers

Finding	Source
90% of surveyed corporate and private equity organizations are using generative AI in M&A, and **37%** apply it across multiple deal stages. **49%** use it in deal closing and **52%** in post-close integration, but only about **one-third** use it for judgment-intensive negotiating steps.	Deloitte, 2026 Generative AI in M&A Pulse Study
Human review is the top requirement for high-stakes AI use. More than half of organizations have only **partial** integration of AI tools across the deal lifecycle, and fewer than **one-third** have full integration. Integration with approved deal data sources is the most important capability they cite.	Deloitte, 2026 Generative AI in M&A Pulse Study
Adoption of AI tools in M&A more than doubled in 2025, with **45%** of surveyed M&A executives now relying on the technology. About **one-third** deploy AI systematically or are redesigning M&A processes for it.	Bain & Company, M&A Report 2026 (January 2026; 300+ M&A executives)

Finding	Source
52% of finance leaders use AI for diligence analysis, and **60%** name technology and AI-driven transformation as their top value-creation priority for the next 12 to 24 months.	Grant Thornton, 2026 Q2 CFO survey (n=232 corporate finance leaders)
40% of respondents using generative AI in M&A said it shortened deal cycles by an average of **30 to 50 percent**.	McKinsey, AI meets M&A (March 2026)
GPs report the highest returns from generative AI in deal sourcing and due diligence, while about **40%** do not expect material financial impact from AI in portfolio companies in 2026.	Bain & Company and StepStone, 2026 GP Outlook (March 2026; 100+ respondents)
63% of breached organizations lacked an AI governance policy. **20%** reported a breach tied to shadow AI, adding up to **$670,000** to breach cost. Of the **13%** with breaches involving AI tools, **97%** lacked proper AI access controls.	IBM Cost of a Data Breach Report 2025 (600 organizations, 17 industries)
93% of IT leaders were confident their Microsoft 365 governance could support AI, yet **29%** said AI had already surfaced something it should not have had access to. **53%** lack full visibility into what their Copilot agents can access.	ShareGate, State of Microsoft 365 2026 (851 IT leaders)

Survey figures are self-reported and use different samples. Deloitte and Bain cover corporate and private equity dealmakers, Grant Thornton covers corporate finance leaders, and IBM and ShareGate cover all industries, not M&A advisory firms alone.

The pattern is consistent: adoption is high, controls are thin, and confidence runs ahead of evidence. Your clients, counterparties, and regulators will increasingly ask you to show your controls, and that evidence lives in your systems.

The AdvancedBITS M&A AI readiness model

AI readiness has seven domains, and a gap in any one can undermine the rest. The model replaces tool-by-tool decisions with one view that leadership can review. It can also be cross-walked to the four functions of the NIST AI Risk Management Framework: govern, map, measure, and manage.

Table 2. Seven readiness domains

Readiness domain	Question to ask	Risk if ignored	Evidence you should have
1. Governance and accountability	Who owns AI decisions, and how are tools approved, monitored, and retired?	Tools spread without an owner, a policy, or a decision record.	AI policy, executive owner, approved tool list, review cadence.
2. Deal data and NDA readiness	What may each AI tool touch, and what does each NDA and client agreement allow?	Use exceeds the permitted purpose, or confidential material reaches a non-permitted party.	Data classes, prohibited-data list, per-deal AI tool register, counsel review.
3. Microsoft 365 and Copilot readiness	Are deal sites, Teams, OneDrive, and mailboxes ready for AI search?	AI surfaces old deals, stale links, or over-shared files.	Permission review, site owners, labels, audit logs, retention.

Readiness domain	Question to ask	Risk if ignored	Evidence you should have
4. Identity, security, and resilience	Can we prevent account compromise and recover from ransomware or deletion?	AI workflows depend on accounts and systems that are not protected or recoverable.	MFA, Conditional Access, endpoint protection, tested restores, incident response plan.
5. Vendor and VDR oversight	What do data room, CRM, AI, and diligence vendors do with prompts, files, and outputs?	Retention, training use, and connections to other AI tools go unexamined.	Vendor inventory, terms review, audit-log and export evidence.
6. Workflow, accuracy, and human review	Which workflows are mature enough for AI, and who checks the output?	Fluent but wrong summaries reach clients, memos, or models.	Source-linked findings, named reviewers, exception log.
7. Workforce, records, and measurement	Do staff know the rules, and can we retain and show the results?	Policy exists on paper only, AI-assisted work is not retained properly, and value cannot be proven.	Training records, retention rules, pilot metrics reported to leadership.

Four levels of AI readiness

Most M&A firms sit at level 1 or 2 today. The scale shows what each level looks like, what it costs you, and the next move.

Table 3. AI readiness maturity model

Level	What it looks like	What it costs you	Next step
1. Ad hoc	Bankers and analysts use whatever AI tools they find. There is no inventory, policy, owner, or NDA check.	Unknown deal-data exposure, no evidence for a client or regulator, and no safe basis for approving Copilot.	Inventory AI use and name an executive owner.
2. Aware	Leadership knows AI is in use. A draft policy exists, but permissions, vendor terms, and NDA checks are incomplete.	Policy without enforcement. Staff cannot tell which tools are approved, and old deal content is still searchable.	Publish approved and prohibited tool lists and clean up permissions.
3. Governed	Approved tools, prohibited data, NDA checks, source-linked review, labels, DLP, and vendor reviews are in place for pilot workflows.	Manual monitoring effort, and scaling slows if reviews are not routine.	Run controlled pilots with measured outcomes and audit-log review.
4. Managed	Continuous monitoring, periodic access reviews, tested restores, vendor reassessment, agent controls, and regular reporting to leadership.	Ongoing discipline, which is the price of scale.	Expand to more teams and review against external frameworks.

AI regulation and standards: what applies to M&A firms

No single law governs AI in M&A advisory work, so firms combine securities rules, contract terms, professional duties, and voluntary frameworks. The table summarizes what applies. It is general information, not legal advice, and which rows apply depends on how your firm is registered and who your clients are.

Table 4. Regulatory and standards map

Framework	Status and scope	What it means for AI readiness
FINRA rules and guidance	Applies if your firm is a FINRA member broker-dealer. FINRA's technology-neutral rules apply to generative AI, including tools from third parties and AI features embedded in products you already run (Regulatory Notice 24-09, June 2024, which creates no new requirements). FINRA's 2026 Annual Regulatory Oversight Report adds a standalone generative AI section.	AI use can implicate supervision, communications, recordkeeping, and fair dealing. FINRA suggests enterprise-level supervisory processes, controls for accuracy and bias, testing and monitoring, and cybersecurity programs that cover AI risks including vendors. Its report also discusses AI agents, and summaries of it point to human-in-the-loop oversight.
SEC Regulation S-P (as amended 2024)	Covers broker-dealers, SEC-registered investment advisers, and other covered institutions.	Requires written incident response programs, customer notification, and service-provider oversight. AI vendors that hold customer information are service providers to review.
NDAs and client engagement terms	Contract terms for each deal. Law-firm commentary reports that NDAs increasingly address AI, that some tools retain inputs or take license rights, and that an AI provider may not qualify as a permitted recipient.	Check the NDA before the tool, keep a per-deal AI tool register, and propose targeted wording on training and approved tools instead of blanket bans.
ABA Formal Opinion 512	Ethics guidance for lawyers (July 2024) applying competence, confidentiality, communication, and fee rules to generative AI.	If counsel on your deals use AI, expect questions about confidentiality and supervision. It also informs your own verification standards.
NIST AI RMF 1.0 and Generative AI Profile (AI 600-1)	Voluntary frameworks. The AI RMF Core has four functions: govern, map, measure, and manage. NIST published the Generative AI Profile in July 2024, and says the AI RMF is being revised.	A neutral structure for your policy and risk register that clients and regulators recognize.
ISO/IEC 42001:2023	A certifiable international standard for AI management systems, published December 2023. Not a legal requirement.	Worth considering only if clients or lenders ask for independent assurance. Many firms will not need it.
Other jurisdictions	Cross-border deals and staff can bring non-U.S. AI and data rules into scope.	Ask counsel which jurisdictions apply and add them to your AI inventory.

General information, not legal advice. FINRA obligations apply to member firms, and not every M&A advisory firm is a member. The FINRA and Regulation S-P rows are summaries from FINRA, the SEC, and law-firm analyses.

Readiness is how you meet all of these at once. A firm with an AI inventory, a written policy, NDA checks, human review points, vendor reviews, retention rules, and an incident plan is positioned for FINRA examinations, client questionnaires, and counterparty scrutiny, without treating each as a separate project.

Deal data and NDA readiness

The first AI question on any deal is what the contract allows. Counsel decides what each NDA permits. Your technology team keeps the facts: which tools are in use, what each does with the data, and proof of what touched what. In practice, check the NDA before any document enters an AI tool, keep a per-deal AI tool register, define prohibited data, prefer approved enterprise tools with written no-training terms, and plan for return or destruction. Our AI in M&A Diligence case study walks through this workflow in detail, including the technology control for each common NDA concern.

Microsoft 365 and Copilot readiness for deal teams

Copilot is powerful because it works across what a user can already access, which is also why permissions come first. Old deal sites, stale sharing links, broad Teams channels, and forgotten OneDrive folders become far easier to find once AI search is on. ShareGate found that 93% of IT leaders felt ready, 29% had already seen AI surface something it should not have, and full Copilot deployment rose from 29% to 56% of organizations in a year. The tools moved faster than the environment underneath them.

Table 5. Microsoft 365 readiness checklist for AI

Readiness item	Why M&A firms should care	AdvancedBITS action
Oversharing review	Copilot can surface old deals, seller materials, buyer lists, or internal valuation work a user technically can open.	Review SharePoint, Teams, OneDrive, guest access, and sharing links before deployment.
Deal-site ownership and lifecycle	Closed, lost, or paused deals often stay active and searchable.	Assign owners and define create, review, archive, close, and revoke-access rules.
Sensitivity labels	Confidential transaction data should not rely on folder names alone.	Define labels for active deals, clean-team, client-only, restricted, and internal analysis content.
Clean-team access	Some users may need limited access to specific workstreams only.	Review access groups, Conditional Access, and audit trails around restricted material.
Audit logging and DLP	You need evidence of how AI is used and a way to catch risky movement of data.	Enable audit logs and DLP policies, with a review cadence for high-risk work.

Readiness item	Why M&A firms should care	AdvancedBITS action
Controlled pilot	Firm-wide rollout before cleanup increases exposure.	Start with a pilot group, validate permissions, monitor use, and adjust before scaling.

An interim control while you clean up. Microsoft documents Restricted Content Discovery, a temporary control that limits discovery of selected SharePoint sites in organization-wide search and Copilot responses while permissions are reviewed. It does not change existing permissions, and Microsoft advises using it selectively because heavy use reduces what Copilot can draw on. It buys time for old deal sites and does not replace cleanup.

AI agents: the next readiness question

FINRA's 2026 report discusses AI agents for the first time, and you should treat agents as a distinct readiness item. An agent is a system that can take actions on its own, such as searching a data room, drafting and sending messages, or updating records. According to a consulting-firm summary of the report, the risks include autonomy without human validation, actions beyond intended scope or authority, difficulty with auditability and transparency, data sensitivity and potential leaks, and lack of domain knowledge for complex tasks. ShareGate also found that 53% of IT leaders lack full visibility into what their Copilot agents can access.

Table 6. Controls for AI agents in deal environments

Control	What it means in practice
Approve each agent	Register every agent, its owner, its purpose, and the data it may reach, before it runs.
Least privilege	Give each agent its own identity with the minimum access needed, and never reuse a person's credentials.
No unattended external actions	Require human approval before an agent sends external messages, changes records, or touches payment or wire information.
Respect NDAs and clean teams	Keep agents away from deal material until the NDA check and access boundaries are confirmed.
Log and review	Keep activity logs that show what the agent accessed and did, and review them on a schedule.
Ability to stop	Document how to suspend an agent and revoke its access quickly.
Vendor review	Review third-party agents and embedded agent features the same way you review any vendor with access to deal data.

AI governance in practice: a two-page policy staff can follow

A short policy that people read beats a long one that they do not. The outline below fits on two pages and answers six questions: which tools, which data, what the contract says, who reviews, what is off limits, and what is retained. Pair the policy with approved alternatives so

staff are not pushed toward unapproved tools. IBM linked one in five studied organizations to a breach involving shadow AI.

Table 7. AI acceptable use policy outline

Section	What it should say
Purpose and scope	Who and what the policy covers: employees, contractors, advisors, all AI tools including features embedded in products, and agents.
Approved, restricted, and prohibited tools	A maintained list by workflow and data type, with an owner and a request process for new tools.
Prohibited data	What may never enter public or unapproved tools, including NDA-covered material, CIMs, models, buyer lists, board materials, and personal data.
NDA and client terms check	A required step before AI use on a deal, with counsel confirmation and an entry in the per-deal register.
Human review and sources	Named reviewers, a source reference on every finding, and the rule that reviewers open the source before relying on it.
Decision boundary	AI does not set valuation, give legal advice, make buyer recommendations, or clear red flags without a person deciding.
Agents	Approval, least privilege, human approval for external actions, and logging.
Records and retention	How AI-assisted work, prompts, and outputs are retained or deleted in line with your recordkeeping duties and deal terms.
Incidents and exceptions	How to report a suspected leak or error, and how exceptions are approved and recorded.
Training and enforcement	Required training, periodic attestation, and what happens when the policy is broken.

Not every AI use carries the same risk. The tiers below help you decide where to start. FINRA's 2026 report found summarization and information extraction is the top generative AI use case among member firms, which fits the lower tiers.

Table 8. AI use-case risk tiers for M&A

Tier	Examples	Controls required
Lower risk: organizing and drafting	Data-room indexing, question trackers, internal meeting notes, first-draft issue lists, policy search.	Approved tools, NDA check, clean permissions, human review before sharing.
Moderate risk: extraction and comparison	Contract term extraction, change-of-control flags, financial schedule comparisons, first-pass memo summaries, buyer research.	Source link on every finding, reviewer opens the source, negative sampling, exception log.
Restricted: conclusions and autonomy	Valuation conclusions, legal advice, buyer recommendations, clearing red flags, anything sent outside the firm, and agents that act on their own.	Human decision-making only, senior or counsel approval, and agent controls from Table 6.

Vendor and data room oversight

Vendor review should not stop at whether a vendor says it is secure. Data rooms, CRM platforms, diligence tools, and AI vendors all touch confidential deal data, and each new AI feature or connection to an external assistant is another path to review. FINRA's notice says its rules apply whether a firm builds a tool, uses a third party's, or uses AI embedded in a product it already runs.

Table 9. Vendor review questions

Topic	Question to ask	Why it matters
Retention	How long are uploads, prompts, outputs, and extracted data kept?	NDA return-or-destroy obligations and breach exposure.
Training use	Will deal data be used to train or improve models?	Confidentiality and competitive sensitivity.
Audit logs	Can we see who viewed, downloaded, searched, uploaded, or generated outputs?	Defensibility and client confidence.
Access controls	Can access be limited by workstream, clean team, guest, advisor, buyer, or seller role?	Role-specific boundaries on sensitive work.
Connections to other AI tools	Does the platform connect to external AI assistants, and who controls what they reach?	A new data path that must fit the NDA.
Incident commitments	How fast will the vendor notify us of a security incident?	Incident response duties and Regulation S-P service-provider oversight.
Export and archive	Can we preserve final evidence and logs after close?	Post-close integration and dispute readiness.

Cyber resilience is an AI prerequisite

AI workflows depend on protected accounts and recoverable systems, and attackers are using AI too. IBM found that 16% of breaches studied involved attackers using AI, most often for AI-generated phishing (37%) and deepfake impersonation (35%). On a deal, a convincing voice or email can redirect a wire or extract a data room link. IBM also reported that among organizations with breaches involving AI tools, 97% lacked proper AI access controls.

The controls are familiar: enforce MFA and Conditional Access, protect email, train staff to verify urgent requests by a second channel, monitor identity activity, and test restores. They protect your AI program as much as the rest of the firm. For a deal-phase security view and independent evidence on how cyber issues affect transactions, see our M&A technology due diligence and cybersecurity page.

Ways to get support

Most firms do not build AI readiness alone. Deloitte found that 35% of organizations realize the most value from AI in M&A with advisory-led use and 20% prefer fully managed solutions. That survey is about AI solutions in M&A, not IT support, but it shows many teams want help. The right model depends on your internal capacity.

Table 10. Support models for AI readiness

Model	Where it works	Where it falls short
Self-managed	Firms with a staffed IT, security, and compliance team that can own policy, permissions, vendor review, and monitoring.	Stretched teams let reviews lapse, and no one tests the controls independently.
Advisory-led	Firms that want a one-time assessment, policy, and roadmap, then run it themselves.	The plan can stall without owners and recurring reviews.
Co-managed	Firms that keep an internal lead and add monitoring, security, Microsoft 365 governance, and recurring reviews.	Responsibilities must be written down before the engagement starts.
Fully managed	Firms without internal IT that want one accountable partner for operations, security, recovery, and AI governance.	Requires clear communication and documentation you own.

A 90-day AI readiness roadmap for M&A firms

A focused 90-day plan can move most firms from level 1 or 2 to level 3. Each phase ends with evidence leadership can review.

Table 11. 90-day roadmap

Timeframe	Priority	Evidence at the end of the phase
Days 1 to 30: See	Inventory AI tools and AI features across Microsoft 365, the data room, CRM, and email. Review deal workspaces, guest access, and sharing. Collect NDA and engagement templates and active-deal NDAs for AI terms. Name an executive owner.	An AI inventory, a permissions findings report, a vendor list, and an NDA review list for counsel.
Days 31 to 60: Govern	Publish the AI acceptable use policy with approved and prohibited tools matched to NDA categories. Set source-validation rules, agent controls, labels, DLP priorities, vendor criteria, and retention guidance. Close the highest-risk permission gaps and apply interim discovery limits on old deal sites.	A signed policy, a per-deal AI tool register, cleaner deal workspaces, and an agent register.

Timeframe	Priority	Evidence at the end of the phase
Days 61 to 90: Pilot and prove	Pilot two or three lower-risk workflows with human review and audit-log monitoring. Train staff. Test a Microsoft 365 restore. Report results to leadership and decide what to scale.	Pilot metrics, audit-log findings, a restore test result, and a scale or stop decision with named owners.

Measure the pilot so leadership can decide with evidence. Useful indicators include:

- Time saved per task compared with the manual baseline
- Share of AI findings confirmed by the reviewer, and issues found by negative sampling that the AI missed
- Share of findings that carry a correct source reference
- Unapproved AI tools detected, then approved or retired
- DLP matches and audit-log findings in AI-assisted workflows
- Time to restore a sample of Microsoft 365 content during a test

M&A AI readiness self-check

Answer yes or no to each question and count your yeses. This is an AdvancedBITS planning tool, not a validated or certified assessment.

Table 12. Ten-question readiness self-check

#	Readiness check	Domain
1	We have a written AI acceptable use policy that staff have been trained on.	Governance
2	We know which AI tools and AI features staff and vendors are using today, approved or not.	Governance
3	An executive owner is accountable for AI decisions, with compliance, IT, and deal-team input.	Governance
4	We check each deal's NDA for AI terms before using AI, and keep a per-deal tool register.	NDA and data
5	We have reviewed Microsoft 365 deal sites, guest access, and sharing links in the last 12 months.	Microsoft 365
6	Every AI finding that could affect price, terms, or advice carries a source reference and a named reviewer.	Accuracy and review
7	MFA is enforced on email, remote access, and administrator accounts.	Security
8	We have test-restored Microsoft 365 data and know how long it took.	Resilience
9	Our data room, CRM, and AI vendors have been reviewed for retention, training use, and AI connections.	Vendors
10	Any AI agents are registered, limited to least privilege, and barred from unattended external actions.	Agents

Table 13. Scoring

Yes answers	Where you stand	Recommended move
0 to 3	Level 1 to 2: foundation first	Pause new AI rollouts. Inventory use, publish a policy, check NDA terms, and fix permissions and MFA.
4 to 7	Level 2 to 3: partly ready	Close the gaps you answered no to, then pilot lower-risk workflows.
8 to 10	Level 3 and above: ready to pilot	Run measured pilots, review audit logs, and plan scale-up with leadership.

Questions M&A leaders should ask

- Which AI tools are bankers, analysts, advisors, operating partners, and marketing teams already using?
- Has each live deal's NDA been checked for AI terms, and who confirmed it?
- What deal data is prohibited from entering public or unapproved AI tools?
- Can AI outputs link back to exact source documents, clauses, or schedules?
- Who validates AI-assisted findings before they influence valuation, terms, or client advice?
- Are our data room, CRM, and AI vendors clear on retention, training use, and connections to other AI tools?
- Are Microsoft 365 permissions ready for Copilot or AI-powered deal search?
- Do we have any AI agents, and what can they access and do?
- If a client, counterparty, or regulator asked how we control AI tomorrow, what evidence could we produce?

Where AdvancedBITS fits

AdvancedBITS brings more than 11 years of managed services experience and a SOC 2 Type II examined security practice to M&A firms building a secure, governed foundation for AI. We do not replace legal review, compliance advice, deal judgment, or advisor accountability. We strengthen the technology environment around your deal work, whether you outsource IT entirely or co-manage with an internal lead.

Table 14. How AdvancedBITS supports AI readiness for M&A firms

AdvancedBITS capability	How it supports AI readiness	Result for your firm
AI readiness assessment	Scores the seven domains, builds the AI and agent inventory, and defines approved tools, prohibited data, review owners, and pilot criteria.	A maturity level, a gap list, and a prioritized plan.

AdvancedBITS capability	How it supports AI readiness	Result for your firm
Microsoft 365 governance	Reviews SharePoint, Teams, OneDrive, permissions, labels, DLP, retention, audit logs, backup, and Copilot readiness.	Less oversharing and a safer AI search foundation.
Deal workspace design	Creates templates, access groups, closeout steps, and archive rules.	Repeatable collaboration and cleaner records.
Vendor and data room review support	Assesses data rooms, AI platforms, CRM, eSignature, diligence tools, and outsourced providers.	Third-party risk made visible before confidential data is exposed.
Cybersecurity and identity	Strengthens MFA, Conditional Access, endpoint and email protection, monitoring, and account recovery.	Lower risk of account compromise and deal confidentiality loss.
Backup and resilience	Reviews Microsoft 365 backup, restore testing, continuity, and incident response.	Tested recoverability for deal-critical records.
Policy and training support	Drafts AI acceptable use rules, review procedures, and staff training that fit your registration and clients.	Rules staff can follow and evidence leadership can show.

For a worked diligence example, see our AI in M&A Diligence case study. For deal-phase technology and cybersecurity diligence, see our M&A technology due diligence and cybersecurity page. For succession and ownership transition readiness, see our Ownership Transition Technology Operating Model.

Closing perspective

AI can help M&A firms move faster, but speed cannot come at the expense of confidentiality, source accuracy, client trust, or deal judgment. The firms that use AI well will be the firms that first strengthen the foundations AI depends on.

AdvancedBITS helps build that foundation. By aligning governance, NDA checks, Microsoft 365, vendor oversight, cybersecurity, backup, and agent controls, M&A firms can move from experimentation to responsible adoption with greater confidence.

Frequently asked questions

What is AI readiness for M&A firms?

It is the ability to use AI across sourcing, diligence, execution, and integration while protecting deal confidentiality, honoring NDAs, keeping outputs accurate and traceable, controlling access, overseeing vendors, and keeping humans accountable for conclusions. It covers governance, data, Microsoft 365, security, vendors, workflows, and workforce training.

Do FINRA rules apply to AI tools?

If your firm is a FINRA member, yes. FINRA says its technology-neutral rules apply to generative AI, including third-party tools and AI embedded in products you already use, and that

AI can implicate supervision, communications, recordkeeping, and fair dealing. Notice 24-09 created no new requirements. Not every M&A advisory firm is a member, so confirm your status with counsel.

Can deal teams use ChatGPT or Copilot on deal documents?

Only approved tools, and only after checking the deal's NDA with counsel. For Copilot, review Microsoft 365 permissions first, because it works across what a user can already access. Our AI in M&A Diligence case study shows the NDA workflow.

What should an AI policy for an M&A firm include?

Approved and prohibited tools, prohibited data, an NDA check before use, source-linked outputs, named reviewers, a decision boundary around valuation and legal advice, agent controls, retention rules, incident reporting, and training. A two-page policy is enough to start. IBM found 63% of breached organizations lacked any AI governance policy.

What controls do AI agents need?

Register each agent, give it least-privilege access under its own identity, require human approval for external actions and payments, keep it away from deal material until NDA and clean-team checks are done, log its activity, and document how to stop it. Summaries of FINRA's 2026 report highlight autonomy without human validation, out-of-scope actions, and auditability as key agent risks.

Do we need ISO 42001 certification?

No. It is a voluntary, certifiable standard for AI management systems and is not a legal requirement. Consider it only if clients or lenders ask for independent assurance. The NIST AI RMF is a lighter, voluntary starting framework.

Who should own AI governance in an M&A firm?

An executive owner should be accountable, supported by compliance, IT and cybersecurity, and deal-team leaders. If you are a broker-dealer, FINRA suggests enterprise-level supervisory processes for generative AI, which fits this structure.

How long does it take to become AI ready?

A focused 90-day plan can inventory AI use, check NDA and vendor terms, clean up the highest-risk Microsoft 365 permissions, publish governance, set agent controls, and pilot lower-risk workflows. Scaling then follows from what the pilot shows.

> **Take the next step**
>
> **Find out where your firm stands.** AdvancedBITS offers a no-obligation AI readiness review covering AI and agent inventory, Microsoft 365 deal workspaces, NDA and vendor terms, and backup readiness, with a plain-language summary leadership can act on.
>
> **Call 561-300-3015** | sales@advancedbits.com | www.advancedbits.com
>
> AdvancedBITS | Business IT Solutions. Secure. Compliant. Outcome-Focused.

Sources

Statistics and legal summaries were checked against the sources below as of October 22, 2026. This guide is general information and not legal advice.

1. Deloitte. 2026 Generative AI in M&A Pulse Study. https://www.deloitte.com/us/en/what-we-do/capabilities/mergers-acquisitions/articles/m-and-a-generative-ai-study.html
2. Bain & Company. M&A Capability for a New Era: Five Ways AI Is Creating More Value in M&A Right Now (January 27, 2026). https://www.bain.com/insights/capability-for-a-new-era-m-and-a-report-2026/
3. Grant Thornton. AI is changing the M&A playbook (June 17, 2026). https://www.grantthornton.com/insights/articles/advisory/2026/ai-changing-ma-playbook
4. McKinsey. AI meets M&A (March 12, 2026). https://www.mckinsey.com/featured-insights/charts/ai-meets-m-and-a
5. StepStone Group. Bain & Company and StepStone Group Release 2026 Private Equity GP Outlook (March 3, 2026). https://www.stepstonegroup.com/news-insights/bain-company-and-stepstone-group-release-2026-private-equity-gp-outlook/
6. IBM. What data leaders need to know from the Cost of a Data Breach Report 2025 (November 12, 2025). https://www.ibm.com/think/insights/data-matters/cost-of-a-data-breach
7. Cybersecurity Dive. ‘Shadow AI’ increases cost of data breaches, report finds (July 30, 2025). https://www.cybersecuritydive.com/news/artificial-intelligence-security-shadow-ai-ibm-report/754009/
8. ShareGate. State of Microsoft 365 2026, Chapter 1: AI confidence is not control. https://sharegate.com/state-of-microsoft-365/chap-1
9. Microsoft Learn. Restrict discovery of SharePoint sites and content. https://learn.microsoft.com/en-us/sharepoint/restricted-content-discovery
10. FINRA. 2026 Annual Regulatory Oversight Report: GenAI: Continuing and Emerging Trends (December 9, 2025). https://www.finra.org/rules-guidance/guidance/reports/2026-finra-annual-regulatory-oversight-report/gen-ai
11. FINRA. Regulatory Notice 24-09: FINRA Reminds Members of Regulatory Obligations When Using Generative Artificial Intelligence and Large Language Models (June 27, 2024). https://www.finra.org/rules-guidance/notices/24-09
12. Debevoise & Plimpton. FINRA’s 2026 Regulatory Oversight Report: Continued Focus on Generative AI and Emerging Agent-Based Risks (December 12, 2025). https://www.debevoise.com/-/media/files/insights/publications/2025/12/finras-2026-regulatory-oversight-report-continued.pdf
13. Sia Partners. 2026 FINRA Regulatory Oversight Report: Key Insights for GenAI & Compliance (April 8, 2026), summary of the AI agent risks. https://www.sia-partners.com/en/insights/publications/2026-finra-regulatory-oversight-report-key-insights-genai-compliance
14. DLA Piper. FINRA flags generative AI risks and governance expectations (December 19, 2025). https://www.dlapiper.com/en-us/insights/publications/2025/12/finra-flags-generative-ai-risks-and-governance-expectations
15. SEC. Regulation S-P: Privacy of Consumer Financial Information and Safeguarding Customer Information, Final Rule 34-100155 (May 2024). https://www.sec.gov/files/rules/final/2024/34-100155.pdf
16. Kohrman Jackson Krantz. AI and M&A NDAs: Managing Artificial Intelligence Risks in Confidentiality Agreements (March 12, 2026). https://kjk.com/2026/03/12/ai-and-ma-ndas-managing-artificial-intelligence-risks-in-confidentiality-agreements/
17. Carta. AI clauses in NDAs: Protecting confidentiality without killing collaboration (July 6, 2026). https://carta.com/blog/ai-clauses-nda/
18. American Bar Association. ABA issues first ethics guidance on a lawyer’s use of AI tools (July 29, 2024). https://www.americanbar.org/news/abanews/aba-news-archives/2024/07/aba-issues-first-ethics-guidance-ai-tools/

19. NIST. AI Risk Management Framework. https://www.nist.gov/itl/ai-risk-management-framework
20. NIST. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, AI 600-1 (July 26, 2024). https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence
21. ISO. ISO/IEC 42001:2023, AI management systems. https://www.iso.org/standard/42001