ChatGPT

State of the Union: Project Management, Consulting and Generative AI in June 2026

Executive summary

Generative AI has already changed project work in 2026, but not in the simplistic “AI replaces project managers” way. The biggest confirmed shift is that AI has become normal infrastructure for knowledge work: most organizations now use AI somewhere, major collaboration suites have embedded copilots or agents, and project platforms increasingly treat AI as a native layer for drafting, synthesis, search, and workflow assistance. At the same time, most firms still have not deeply redesigned project workflows around AI, and most remain in experimentation or partial scale rather than enterprise-wide transformation. McKinsey’s 2025 global survey found 88% of respondents reporting regular AI use in at least one business function, but only about one-third said their companies had begun to scale AI; Deloitte’s 2026 enterprise survey says worker access to AI rose by 50% in 2025, yet only 34% are truly reimagining the business. [1]

For project management specifically, the most meaningful current change is not autonomous delivery. It is the compression of administrative and content-heavy work: first-draft scoping, meeting capture, documentation, status reporting, backlog shaping, risk extraction, and knowledge retrieval are now materially faster when tools are connected to real project context. PMI’s June 2026 AI standard explicitly frames AI as changing how project-based work is planned, governed, and delivered, while Microsoft, Atlassian, Google, OpenAI, Asana, Notion, Smartsheet, and ClickUp all now market AI as embedded in day-to-day planning, search, summarization, and task execution workflows. [2]

What remains mostly hype is the idea that current tools can replace senior judgment in ambiguous, political, cross-functional project environments. Gartner’s 2025 warning that more than 40% of agentic AI projects would be scrapped by 2027 due to cost and unclear value is consistent with a broader pattern: vendors are shipping “agent” features faster than many clients can govern, trust, or operationalize them. Even optimistic BCG work argues that value comes when agents are embedded into redesigned workflows with clear human oversight and controls, not when they are used as free-floating autonomy. [3]

For independent senior consultants in a 7N-style context, the role is shifting away from being paid mainly for producing artifacts and coordination overhead, and toward being paid for judgment, orchestration, adoption, and accountability. The consultant who only writes status decks or meeting minutes is more exposed to commoditization. The consultant who can redesign workflows, define decision rights, set safe AI operating patterns, supervise AI-generated outputs, and help clients move from pilots to scaled practice becomes more valuable. This is not just a market opinion: McKinsey links value capture to workflow redesign and defined human validation; Deloitte reports that productivity gains are outpacing business reimagination; BCG reports rising enterprise demand for service providers that can design, deploy, and operate agentic systems. [4]

The practical implication for webinar participants is straightforward. After basic prompting, the differentiator is no longer “can you use ChatGPT?” It is whether you can combine business context,

project data, permissions, governance, facilitation, stakeholder judgment, and human accountability into AI-enabled project leadership. That is the next professional threshold. [5]

Current state of project management and generative AI in June 2026

Confirmed current developments. In June 2026, AI is no longer a side experiment in project environments. It is embedded across the dominant ecosystems consultants meet inside clients: Microsoft is converging tasks, plans, portfolios, and projects in Planner while adding Planner Agent, Teams Channel Agents, SharePoint Agents, and Copilot Studio governance; Atlassian has made Rovo central to Jira, Confluence, Loom, and Teamwork Graph; Google Workspace includes Gemini across the suite and is exposing MCP servers for agent connections; OpenAI Business and Enterprise have apps, company knowledge, deep research, and agent capabilities; and work-management platforms such as Asana, monday.com, Notion, Smartsheet, ClickUp, and others now frame themselves as AI work platforms rather than just task boards. [6]

What has actually changed in project management. The work has shifted fastest in three areas. First, administrative compression: meetings, notes, summaries, first-draft reports, task extraction, and document synthesis are quicker and more standardized. Second, contextual search and retrieval: project teams can increasingly query project knowledge across documents, chats, and work systems rather than manually hunt through folders and tools. Third, workflow assistance: several platforms can now generate task structures, suggest plans, write updates, and in bounded cases execute actions under approvals. PMI's framing of AI in project management, together with vendor capabilities in Planner, Jira/Rovo, Gemini, and work-management suites, supports this as the practical 2026 baseline. [7]

What has changed for consultants embedded inside client organizations. External specialists now enter clients where AI use is often already present at the individual level, but unevenly governed and weakly integrated into delivery. That changes the consultant's job. Clients are less impressed by generic prompting demos and more interested in whether AI can reduce reporting friction, accelerate delivery decisions, keep governance intact, and work inside existing security and permission boundaries. Microsoft's 2025 Work Trend Index shows leaders explicitly thinking in terms of digital labor and hybrid human-agent teams; McKinsey shows high performers are redesigning workflows and defining human validation; Deloitte shows AI access and efficiency gains are rising faster than business redesign. [8]

Where clients are ahead. Many clients are ahead on personal productivity and embedded-suite adoption. They already use AI to summarize meetings, tighten prose, generate slide drafts, accelerate spreadsheets, or get first-pass project content. This is especially true in Microsoft and Google environments where AI is bundled into daily tools, and in Atlassian or work-management platforms that have integrated AI natively. [9]

Where clients are still immature. Most organizations are still immature on workflow redesign, decision-rights redesign, project-data readiness, auditability, permissions architecture, and evidence discipline. McKinsey reports that most organizations still have not embedded AI deeply enough in workflows to realize material enterprise-level benefit, and only a minority have scaled; Deloitte says only 34% are truly reimagining the business; PMI says many organizations are adopting AI faster than they can align on responsible and effective use. That maturity gap is exactly where senior consultants can add value. [10]

The common gap between experimentation and business value. The load-bearing gap is not model quality alone. It is operating-model quality. Organizations often have tools before they have clean project data, stable taxonomies, role-based access, approval logic, output validation, or redesigned workflows. BCG's 2025 and 2026 work is blunt on this point: the gains come from zero-based redesign, data foundations, and calibrated autonomy. Deloitte's Nordic 2025 report adds a useful European nuance: respondents cite regulation, data quality, and real-world mistakes as major barriers, and the report quotes executives saying poor data quality is a primary blocker. [11]

What is real now versus what is hype and how the consultant role is changing

What is real now. Real in 2026 means AI is already reliable enough to draft, summarize, cluster, compare, translate, retrieve, classify, and generate structured first versions of common project artifacts. It is also real that enterprise assistants can act inside bounded environments: Microsoft Planner Agent can execute assigned tasks and write status reports, Teams Channel Agents can answer project questions and draft AI status reports inside a channel, Google Workspace MCP servers can let agents read and act within Gmail, Drive, Calendar, and Chat under user permissions, and OpenAI apps can search, sync, and act across connected tools with admin controls. Those are not future concepts; they are documented platform capabilities. [12]

What is still hype. It is still hype to imply that current AI can independently run the social, political, and fiduciary parts of project leadership. It is also hype to assume that an "AI-enabled" platform delivers value without process cleanup and data governance. Gartner's warning on agentic project failure, together with NIST's emphasis on prompt injection, human oversight, testing, incident disclosure, and provenance, shows that "autonomous" does not mean "safe," "trusted," or "worth paying for." In practice, many "agent" products are still preview-stage, tightly bounded, or heavily dependent on permissions and connector setup. [13]

How the role is changing. The role is shifting away from coordination as a scarce service and toward judgment as a scarce service. Meeting notes, routine progress summaries, basic risk logs, draft plans, and first-pass stakeholder updates are increasingly automatable. The role that rises is the person who can ask better questions, pressure-test AI outputs, detect missing assumptions, shape trade-offs, broker alignment, and decide when not to automate. PMI itself now frames the profession as both "using AI to manage work" and "managing AI-driven initiatives," which is a strong signal that project leadership is becoming more strategic and governance-heavy. [14]

The consultant role shift in plain language. The old center of gravity was task tracking, project administration, coordination, workshop documentation, and status reporting. The new center of gravity is decision support, workflow redesign, AI-enabled facilitation, orchestration of people and tools, supervisory control of AI outputs, governance, and adoption advice. Microsoft's "agent boss" language is vendor-coined, but it captures a real trend: experienced knowledge workers are increasingly expected to delegate to AI, manage it, and remain accountable for its outputs. McKinsey's evidence that high performers define human validation and redesign workflows supports the same conclusion from a less vendor-centric angle. [15]

What this means in a 7N-style consulting context. For senior independents, trust now depends less on being the fastest producer of deliverables and more on being the safest and sharpest interpreter of AI-assisted work. A trusted external consultant must be able to say: which tool was used, what data it touched, what assumptions were embedded, what was machine-generated, what was human-reviewed, and where the decision still sits. That posture increases credibility rather than reducing value. It also

protects against the obvious client fear: paying senior rates for junior-level AI output. This is a qualified assessment, but it is strongly supported by the enterprise trend toward role-scoped controls, audit logs, human validation processes, and AI-specific governance in the major platforms. [16]

Project lifecycle impact

Project initiation. AI can now help turn rough goals into draft charters, milestone structures, stakeholder questions, and dependency maps faster than manual blank-page work. Microsoft Planner and Copilot can generate work breakdown structures; work-management platforms can scaffold projects from templates and prompts. Human judgment is still decisive in defining why the project exists, what problem is worth solving, and which constraints are politically or commercially real. Consultants should use AI to shorten setup time, then spend more time on objective clarity, sponsor alignment, and kill-criteria. [17]

Business case and scoping. AI is useful for scenario drafting, assumption harvesting, benchmark summarization, benefit hypothesis generation, and first-pass option comparison. It is weak at validating whether a proposed business case rests on reliable internal numbers, realistic change effort, or credible adoption assumptions. Consultants should move from “writing the business case” to “stress-testing the business case,” with explicit assumption logs and source checks. McKinsey’s high-performer data and NIST’s accuracy and oversight guidance both point that way. [18]

Discovery and current-state analysis. This is one of the strongest immediate use cases. AI can rapidly summarize interviews, policies, meeting transcripts, ticket patterns, and fragmented documentation, especially when grounded in enterprise search or approved connectors. But discovery still needs human framing: what counts as signal, what contradictions matter, and what the organization is not saying. Consultants should learn to run AI-assisted evidence synthesis while protecting confidentiality and documenting uncertainty. [19]

Stakeholder analysis. AI can help draft stakeholder maps, message variants, engagement plans, and influence hypotheses from notes and prior communications. What it cannot safely infer on its own is political salience, informal power, trust history, or the risk of saying the wrong thing to the wrong audience. Consultants should use AI to improve completeness and message preparation, not to delegate stakeholder judgment. [20]

Planning and scheduling. AI is now good at creating draft plans, dependencies, goals, backlogs, sprint structures, and resource prompts inside platforms such as Planner, Jira, and work-management suites. It remains weak at negotiating resource contention, cross-team trade-offs, and uncertain external dependencies without strong local context. Consultants should get faster at generating planning options, then focus their added value on sequencing judgment, resource politics, and escalation design. [21]

Workshop design and facilitation. AI already helps assemble agendas, synthesize pre-readings, produce icebreakers, cluster inputs, and draft outputs in real time. Yet live facilitation still depends on reading the room, handling resistance, redirecting conflict, and surfacing what participants avoid saying. Consultants should redesign workshops so AI handles capture and synthesis while humans handle energy, challenge, and decision moments. [22]

Requirements and documentation. This is among the most mature augmentation areas. AI can convert conversations into requirements drafts, user stories, acceptance criteria, process notes, RAID entries, and decision logs. Atlassian, ClickUp, Asana, and Microsoft all push this direction. Human

judgment is still essential for ambiguity removal, testability, prioritization, and regulatory interpretation. Consultants should stop treating first-draft documentation as premium work and instead sell precision, traceability, and decision quality. [23]

Backlog and delivery management. AI can summarize discussions into tasks, detect duplicates, suggest subtasks, cluster themes, and surface blockers. It can also improve backlog hygiene when the work graph is clean. But prioritization remains deeply contextual: urgency, value, risk, architecture, stakeholder pressure, and delivery capacity rarely live in one place. Consultants should use AI to keep the backlog coherent and current, then apply human judgment to priority and sequencing. [24]

Risk management. AI is genuinely useful for extracting risks from notes, comparing them to historical patterns, suggesting mitigations, and drafting escalation language. PMI's research and vendor tools support this. But current systems still produce false confidence and generic mitigations if not grounded in project context. Consultants should shift from manually listing risks to curating risk signal, validating materiality, and defining thresholds, owners, and responses. [25]

Communication and steering-group reporting. AI can now draft executive summaries, status packs, milestone updates, and audience-specific messages very well. Microsoft, Planner Agent, Teams agents, Loop, and enterprise assistants all make this easier. What still requires a consultant is calibration: what to emphasize, what to omit, what is politically safe, and what the governance body actually needs to decide. Consultants should use AI to eliminate format labor and spend more time on decision framing. [26]

Execution and follow-up. AI helps monitor activity, summarize progress, propose follow-up actions, and in some tools execute bounded tasks. But execution quality still depends on context changes, commitment, vendor dependencies, and exception handling. Consultants should supervise AI-generated follow-up, not assume completion equals progress. This is especially true where agents can act but still need approvals and well-scoped permissions. [27]

Benefits tracking. AI can help structure benefit trees, compare plan versus actual narratives, detect recurring blockers, and automate some reporting logic. What it cannot do on its own is establish causality, isolate project impact from external factors, or arbitrate benefit ownership between functions. Consultants should become better at benefit data design and evidence discipline, not just dashboard production. [28]

Lessons learned and retrospectives. AI is valuable for clustering feedback, extracting themes, comparing retrospective notes across waves, and identifying repeated failure modes. Human facilitation remains critical for candor, psychological safety, and deciding what action is real rather than rhetorically attractive. Consultants should use AI to improve pattern recognition while keeping retrospectives human-led. [29]

Handover and capability building. AI can accelerate playbooks, SOP drafts, onboarding material, FAQs, and searchable knowledge bases. But sustainable handover still requires role clarity, process ownership, governance, and actual capability transfer inside the client. Consultants should increasingly treat handover as "operating model installation," not just documentation delivery. That means training teams in both process and AI-enabled practice. [30]

AI agents in project and consulting work

Are AI agents actually useful in 2026? Yes, but only in bounded ways. The strongest evidence says agents are useful where work is repetitive, multistep, tool-connected, and governed: plan generation, task routing, status drafting, workflow execution, retrieval across systems, and narrow operational monitoring. Microsoft Planner Agent, Teams Channel Agents, Google Workspace MCP servers, OpenAI apps and agent features, Rovo Agents, and several work-management platforms all support this practical picture. BCG's work suggests significant value when agents are embedded into process-heavy functions. [31]

Where agents are still immature. Agents remain immature in ambiguous project environments with partial data, fuzzy objectives, unsafe permissions, or politically sensitive trade-offs. Gartner's 2025 warning about project failure and "agent washing" matters because many vendors market autonomy before governance is mature. Preview status is another clue: Microsoft's Copilot Studio workflows were still public preview in May 2026, and several higher-autonomy features across suites remain gated by admin setup, licensing, or early-access programs. [32]

What "human in the loop" means in real project environments. It does not mean a human glances at outputs occasionally. In project work, human-in-the-loop means a named person remains accountable for goal framing, permission scope, exception handling, sign-off thresholds, and downstream decisions. NIST explicitly emphasizes human oversight roles, testing, incident disclosure, and documentation retention; the EU AI Act similarly ties deployer obligations for high-risk systems to human oversight and logging. In practice, that means consultants should design agent workflows with checkpoints, not just prompts. [33]

What governance is needed when agents act on behalf of a team. At minimum: agent inventory, named owner, approved data sources, scoped permissions, audit logs, retention policy, escalation rules, output-review thresholds, and a clear rule for when agents may recommend versus act. Microsoft Copilot Studio now exposes data policy controls, audit visibility, agent identities, and DLP governance; OpenAI Business and Enterprise let admins disable apps by default and assign app access by role; Google emphasizes inherited permissions and enterprise protections; Atlassian recommends reviewing source data types and permissions before enabling Teamwork Graph connectors. [34]

How to explain agentic AI to clients without overselling it. The sober explanation is: agents are best treated as workflow participants, not replacements for accountable professionals. A good test is whether the task has a clear goal, tool access, acceptable failure cost, and a human escalation path. If not, call it assistance, not autonomy. That framing aligns with both the optimistic and skeptical evidence: BCG sees value in process-heavy functions with controls; Gartner warns many projects will fail from unclear value and immaturity. [35]

New consulting opportunities around agents. The most credible opportunities are not "build me a magical PM agent." They are agent readiness assessment, workflow decomposition, connector and permission design, approval architecture, operating-model design for human-agent teams, and supervision protocols. BCG's 2026 tech-services report is particularly relevant here: enterprise buyers increasingly expect providers to design, deploy, and operate autonomous systems that deliver outcomes. [36]

Tools and platform landscape

Microsoft ecosystem. Microsoft is currently the most complete enterprise stack for AI-enabled project work, especially where clients already live in Teams, SharePoint, Planner, Office, and Power Platform. The genuinely useful parts for consultants are unified work management in Planner; premium features such as dependencies, backlogs, sprints, and portfolios; Planner Agent for plan and task assistance; Teams Channel Agents for project Q&A and AI status reporting; SharePoint Agents for site-grounded retrieval; Copilot Pages and Loop for collaborative drafting; and Copilot Studio for building governed automations and agents. The main dependency is data hygiene and permissions in Microsoft 365. The main governance need is access control, DLP, audit, agent identity, and environment routing. The most likely maturity curve over the next 12–24 months is from "copilot inside documents and plans" to "role-scoped agents working across Teams, SharePoint, Planner, and Power Platform with stronger admin controls." [37]

Google Workspace Gemini. Google is strong wherever the client's collaboration backbone is Gmail, Docs, Sheets, Meet, Drive, and Chat. The practical wins are summarization, drafting, chart and spreadsheet assistance, document analysis, and enterprise search across Workspace content, initially for productivity and increasingly for agentic action through Google Workspace MCP servers. The major dependency is still permissions and the fact that Gemini can only work on content the user already has access to; client-side encryption remains a relevant limitation for some use cases. Governance is strongest when Gemini is used under Workspace licenses, where submissions are not used to train models and enterprise protections and DLP apply. Over the next 12–24 months, Google's MCP direction suggests stronger third-party and cross-tool agent workflows, but today much of the value is still inside the suite. [38]

Atlassian Jira, Confluence, Rovo, and Jira Product Discovery. Atlassian is particularly strong for consultants in delivery-heavy, product-heavy, and software-adjacent environments. The genuine value is in contextual search, knowledge retrieval, idea capture, summarization, and AI grounded in the Teamwork Graph across Jira, Confluence, Loom, and connected systems. It is one of the more explicit examples of AI working from a work graph rather than a standalone prompt box. But it also illustrates the maturity constraint clearly: usefulness depends heavily on admin setup, connector strategy, permissions sync, and the quality of the underlying work graph. Over the next 12–24 months, expect increasing value from Rovo agents and connectors, but only where the client treats its Atlassian environment as a maintained knowledge system rather than a debris field of half-finished tickets and stale pages. [39]

Asana, monday.com, Smartsheet, Notion, ClickUp, and similar platforms. These platforms have moved fastest in packaging AI around practical work-management workflows. Asana AI Studio is compelling for no-code workflow logic; monday.com is pushing an AI work-platform model with AI credits and sidekick-style assistance; Notion is becoming an AI workspace with agents, meeting notes, and connectors; ClickUp Brain is increasingly a context-aware project assistant; Smartsheet is positioning Smart Agents for guided action under approval. Their biggest strength is speed to value for teams willing to standardize workflows. Their biggest weakness is that they often magnify existing process inconsistency: if naming conventions, project structures, statuses, or ownership are messy, AI simply automates the mess faster. Governance needs vary, but role-based access, connector review, and AI feature controls are now central. In maturity terms, Asana, Notion, monday, and ClickUp are already useful for mid-market workflow augmentation; Smartsheet's more agentic layer still looks earlier-stage. [40]

General-purpose assistants such as ChatGPT, Claude, and Gemini app. For senior consultants, these remain extremely useful for analysis, synthesis, red-teaming, draft generation, scenario design, meeting preparation, and document transformation. Their value increases sharply when connected to company knowledge or approved apps. OpenAI's Business and Enterprise offerings now support apps, company knowledge, deep research, and agent capabilities with admin controls and no training on connected business data by default; Anthropic's Enterprise plan emphasizes audit logs, retention controls, and compliance features; Google Workspace with Gemini brings enterprise controls when used under licensed Workspace accounts. The main governance question is simple: are you using enterprise-approved, role-scoped, audited access, or are you copying client-sensitive material into a personal or consumer-tier tool? Consultants should treat that distinction as non-negotiable. [41]

PM-specific AI and low-code workflow builders. PMI Infinity is important because it signals the professionalization of AI in project work: PMI now has both a June 2026 AI standard and a PM-specific AI assistant grounded in PMI standards and community knowledge. That makes it useful as a coaching and method-support layer. But compared with enterprise-native suites, its value is more in guidance than in being a live execution surface connected to client work systems. On the low-code side, Microsoft Copilot Studio currently stands out for enterprise governance depth, while Google MCP, OpenAI apps/custom connectors, Asana AI Studio, and Atlassian Rovo connectors show that the broader market is moving toward connector-driven orchestration. Over the next 12–24 months, this low-code orchestration layer is likely to matter more than standalone prompting. [42]

Governance, risks, professional responsibility, and consulting opportunities in a 7N-style context

The non-negotiable governance issues. For consultants working inside client organizations, the core risks are now well understood: confidentiality leakage, use of client-sensitive material in the wrong tier of tool, hallucinations presented with false certainty, weak audit trails, broken decision accountability, prompt injection and connector abuse, and unclear ownership of AI-assisted deliverables. NIST's GenAI profile treats governance, provenance, testing, and incident disclosure as first-class concerns, and specifically calls out prompt injection, data poisoning, human oversight, and retention. The EDPB's 2024 opinion adds a crucial GDPR point: AI models trained on personal data cannot automatically be treated as anonymous; that assessment is case-by-case. [43]

European compliance in June 2026. The EU AI Act is now a live planning constraint, not a distant concept. The European Commission states that the Act entered into force on 1 August 2024; prohibited-practice and AI-literacy obligations applied from 2 February 2025; GPAI obligations applied from 2 August 2025; and, following the 2026 political agreement on simplification, many high-risk system rules move to December 2027 while product-integrated high-risk rules move to August 2028. For project consultants, the practical point is this: even where a project team is not building a "high-risk AI system," AI literacy, transparency, documentation, governance, and role clarity are already becoming expected. [44]

Practical rules for consultants handling client data. Do not put client-sensitive data, personal data, contracts, steering materials, or internal roadmaps into public or personal AI accounts. Use enterprise-approved tools whenever prompts could expose confidential project context, regulated data, or identifiable personal information. Prefer tools where admins can control apps, roles, logs, retention, and training settings. Document assumptions and source boundaries in every important AI-assisted output. Label whether content is AI-assisted and whether it has been human-reviewed. Keep humans explicitly accountable for recommendations, decisions, approvals, and external communications. Those are not "nice to have" habits anymore; they are the minimum credible professional standard. [45]

How to be transparent about AI use without undermining value. The right posture is factual and disciplined: say which enterprise tool was used, what it was used for, what sources grounded the output, what checks were performed, and where human judgment overrode or refined the machine. Transparency increases trust when it signals control. Hidden AI use creates client suspicion because it invites the question of whether the consultant is billing senior rates for unreviewed machine output. This is a qualified assessment, but it aligns closely with the market trend toward auditability, role-based access, and human validation in the major enterprise platforms. [46]

What becomes commoditized. Low-complexity drafting, routine meeting documentation, first-pass status decks, generic RAID updates, straightforward stakeholder email drafting, and basic backlog hygiene are all more exposed to commoditization now. Clients can do more themselves, and they increasingly expect consultants to do these things faster. The risk is not that the work disappears; it is that it stops justifying premium pricing on its own. [47]

What becomes more valuable. More valuable are problem framing, sequencing judgment, executive translation, governance design, adoption planning, facilitating hard trade-offs, risk curation, evidence discipline, operating-model redesign, and human-agent workflow design. BCG's 2026 services analysis supports this strongly: enterprise buyers are looking for help designing, deploying, and operating outcome-oriented autonomous systems, not just for standalone tooling. [48]

Credible offerings for independent senior consultants. The most credible offerings in a 7N-style market are pragmatic and close to delivery work: AI readiness assessments for project organizations; AI-enabled PMO redesign; reporting and steering-pack automation; AI-assisted risk and issue management; requirements and documentation acceleration with review controls; stakeholder communication systems; workflow and agent readiness redesign; governance design for project teams; and training in AI-enabled project leadership. These offerings fit what clients actually lack now: integration, controls, redesign, and safe scaling. [49]

Skills consultants should build next, implications for training, and the next 24 months

The next capability curve after prompting is **AI-enabled consulting and project leadership**. That includes context engineering, workflow decomposition, project-data literacy, source checking, structured evaluation of outputs, supervision of AI actions, decision-rights design, facilitation in AI-assisted environments, and clear communication with clients about what AI can and cannot safely do. PMI's current education stack, Microsoft's "agent boss" framing, and NIST's emphasis on governance and oversight all point in the same direction: the valuable professional is the one who can make AI useful inside real organizational constraints. [50]

For training programs aimed at experienced consultants, the next step should therefore move decisively beyond prompt tactics. The curriculum should focus on redesigning project processes rather than merely accelerating old ones; calibrating human review thresholds by task criticality; using enterprise connectors and approved knowledge sources; creating audit-friendly deliverables; and learning how to explain AI use to sponsors, PMOs, works councils, compliance teams, and steering groups. In other words: from "prompt engineering" to "operating model engineering." This is a qualified assessment, but it is strongly grounded in the evidence that workflow redesign, governance, data, and adoption are the main determinants of value. [51]

Confirmed current direction for the next 12–24 months. Project management is likely to become more assistant-rich, more connector-dependent, and more governance-sensitive. PMOs are likely to

evolve from document-and-template custodians into workflow, control, and enablement hubs for AI-assisted delivery. Human teams are likely to get smaller on low-value coordination work but heavier on supervision, design, and escalation. Client expectations of consultants are likely to shift toward “use AI where it is appropriate, but keep accountability and quality visibly human.” Those points are the most plausible extrapolation from current Microsoft, McKinsey, Deloitte, PMI, and BCG evidence. [52]

Forecasts and opinions. My best forecast is that the next 12–24 months will not eliminate project managers or independent consultants, but will punish undifferentiated delivery. The likely losers are roles defined mainly by coordination overhead, format production, and generic documentation. The likely winners are consultants who can combine domain expertise, organizational sensing, process redesign, AI governance, and commercial clarity. Agent use will grow, but most organizations will still run mixed models where humans own decisions and agents handle bounded execution. That forecast is consistent with current platform direction, but it remains a forecast, not a confirmed outcome. [53]

Slide-worthy headline ideas for a webinar

- The admin work is shrinking; judgment work is growing.
- AI did not kill project management; it raised the bar for it.
- The real bottleneck is no longer drafting. It is deciding.
- Most clients have AI access. Far fewer have AI operating discipline.
- The next premium skill is not prompting. It is orchestration.
- Agents are useful when the workflow is clear and the guardrails are real.
- PMOs are moving from template police to workflow governors.
- External consultants will be judged less by output volume and more by control quality.
- If the data, permissions, and ownership are messy, AI scales the mess.
- The consultant advantage now is trusted human accountability on top of machine speed.

State of the Union: Project Management, Consulting and Generative AI, June 2026

- AI is now normal in project environments, but scaled business value still lags behind experimentation. [1]
- The work most changed so far is administrative, content-heavy, and retrieval-heavy work. [54]
- The work least replaceable is political judgment, prioritization, sponsorship alignment, and accountable decision-making. [18]
- Independent senior consultants are moving from “producer of deliverables” to “designer of trustworthy AI-enabled work.” [55]
- Clients are usually ahead on AI experimentation and behind on governance, workflow redesign, and data readiness. [56]
- Agents are real, but mostly for bounded workflows with strong permissions, tooling, and supervision. [57]
- The consultant who cannot explain how AI was used, checked, and governed will lose credibility. [58]
- Routine documentation is becoming cheaper; workflow redesign and governance are becoming more valuable. [59]
- The next learning curve is context engineering, supervision, evidence discipline, and AI-enabled project leadership. [60]
- The winning posture for 2026 is sober ambition: use AI aggressively where it is reliable, and stay visibly human where judgment and accountability matter most. [61]

---

1 4 10 18 46 49 51 56 61 The State of AI: Global Survey 2025 | McKinsey
https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

2 42 Artificial Intelligence | Project Management Institute
https://www.pmi.org/standards/artificial-intelligence

3 13 32 53 Over 40% of agentic AI projects will be scrapped by 2027, Gartner says
https://www.reuters.com/business/over-40-agentic-ai-projects-will-be-scrapped-by-2027-gartner-says-2025-06-25/?utm_source=chatgpt.com

5 14 50 60 Artificial Intelligence in Project Management | PMI
https://www.pmi.org/learning/ai-in-project-management

6 37 Task and Project Management Software | Microsoft Planner
https://www.microsoft.com/en-us/microsoft-365/planner/microsoft-planner?utm_source=chatgpt.com

7 20 25 28 29 54 Benefits of Adopting Generative AI for Project Management | PMI
https://www.pmi.org/learning/thought-leadership/benefits-of-ai-for-project-management

8 15 52 2025: The year the Frontier Firm is born
https://www.microsoft.com/en-us/worklab/work-trend-index/2025-the-year-the-frontier-firm-is-born

9 Gemini AI features now included in Google Workspace ...
https://knowledge.workspace.google.com/admin/generative-ai/workspace-with-gemini/gemini-ai-features-now-included-in-google-workspace-subscriptions?utm_source=chatgpt.com

11 Agents Accelerate the Next Wave of AI Value Creation | BCG
https://www.bcg.com/publications/2025/agents-accelerate-next-wave-of-ai-value-creation

12 26 31 Frequently asked questions about Planner agent | Microsoft Support
https://support.microsoft.com/en-us/planner/copilot/frequently-asked-questions-about-planner-agent

16 41 45 58 Admin Controls, Security, and Compliance in apps (Enterprise, Edu, and Business) | OpenAI Help Center
https://help.openai.com/en/articles/11509118-admin-controls-security-and-compliance-in-apps-enterprise-edu-and-business

17 21 Frequently asked questions about Copilot in Planner ...
https://support.microsoft.com/en-us/planner/copilot/frequently-asked-questions-about-copilot-in-planner-preview?utm_source=chatgpt.com

19 30 Company knowledge in ChatGPT (Business, Enterprise, and Edu) | OpenAI Help Center
https://help.openai.com/en/articles/12628342-company-knowledge-in-chatgpt-business-enterprise-and-edu

22 Get started with Microsoft 365 Copilot Pages
https://support.microsoft.com/en-us/microsoft-365-copilot/get-started-with-microsoft-365-copilot-pages?utm_source=chatgpt.com

23 Rovo in Jira: AI features
https://www.atlassian.com/software/jira/ai?utm_source=chatgpt.com

24 Manage tasks with Brain AI
https://help.clickup.com/hc/en-us/articles/24998833529751-Manage-tasks-with-Brain-AI?utm_source=chatgpt.com

27 Execute tasks with Planner agent
https://support.microsoft.com/en-us/planner/copilot/execute-tasks-with-planner-agent?utm_source=chatgpt.com

33 43 Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile
https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf

34 Security and governance - Microsoft Copilot Studio | Microsoft Learn
https://learn.microsoft.com/en-us/microsoft-copilot-studio/security-and-governance

35 How Agentic AI is Transforming Enterprise Platforms | BCG
https://www.bcg.com/publications/2025/how-agentic-ai-is-transforming-enterprise-platforms

36 48 55 The $200 Billion AI Opportunity in Tech Services | BCG
https://www.bcg.com/publications/2026/the-200-billion-dollar-ai-opportunity-in-tech-services

38 Google Workspace with Gemini | Generative AI
https://knowledge.workspace.google.com/admin/generative-ai/workspace-with-gemini/google-workspace-with-gemini?utm_source=chatgpt.com

39 Atlassian Teamwork Collection
https://www.atlassian.com/collections/teamwork?utm_source=chatgpt.com

40 Asana AI & Agentic Work Management
https://asana.com/product/ai?utm_source=chatgpt.com

44 AI Act | Shaping Europe’s digital future
https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

47 59 The State of AI in the Enterprise - 2026 AI report | Deloitte US
https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/content/state-of-ai-in-the-enterprise.html

57 Configure the Google Workspace MCP servers | Google for Developers
https://developers.google.com/workspace/guides/configure-mcp-servers