StellisAI

LEADER LAB | BRIEFING

GPT-5.6 and ChatGPT Work

What Business Leaders Need to Know

Prepared by: Stellis AI Team | July 2026

A Leader Lab briefing, prepared by Stellis AI to keep leaders ahead of what matters in applied AI.

> **The model improvements are significant. The larger business change is that ChatGPT is evolving from a tool that primarily answers questions into an application environment that can act across your systems and carry work forward over time.**

OpenAI released GPT-5.6 on July 9, together with ChatGPT Work and a redesigned desktop experience that brings Chat, Work, and Codex into one application. For business leaders, this creates both an opportunity and a governance obligation. The same shift is underway across the major platforms: Microsoft Copilot and Anthropic's Claude Cowork are moving in the same direction, so the guidance in this note applies well beyond ChatGPT.

What You Need to Know

> **Bottom line:** This is the step we have seen coming for a while: capable, agentic AI that completes real multi-step work, now packaged for everyone, not just developers. GPT-5.6 is a model upgrade you can adopt today with little disruption. ChatGPT Work is the larger change, because it means you are now delegating a degree of authority, not just access to information. The competitive question is no longer whether the technology is ready. It is who inside your organization is going to use it to enable the business, and how many people you have who can. Adopt the model now, find those people, and put the tools in their hands, while governing Work deliberately before turning it on broadly.

What we recommend

- **Adopt GPT-5.6, but hold on the new desktop app.** Employees get the model improvements today through the browser or existing ChatGPT Classic. Do not broadly deploy the redesigned desktop application until feature consistency and sync issues settle.
- **Assign an owner and check your defaults before Work turns on.** Name an administrative and business owner (your AI Champion). Enterprise and Edu preview with Work off; ChatGPT Business is more open by default, so do not assume a managed workspace is tightly restricted.
- **Separate what AI can see from what it can do.** Classify every integration as read-only, selected-actions, pilot-only, or prohibited, and require explicit human verification for high-impact actions (email, CRM, finance, HR, customer records, deletion, external publishing).
- **Update your AI use policy and require review of all output.** Cover connected apps, local files, managed vs. personal accounts, cross-system aggregation, and scheduled work. Treat "finished"-looking deliverables as unverified until a person checks them.
- **Start with a controlled pilot, not a launch.** This is how you move quickly without moving recklessly: give your strongest people a small set of approved apps and real workflows, start read-only where

the stakes are high, and expand as each use case proves out. A pilot is the accelerator, not the brake.

The sections below explain each announcement and expand on these recommendations for teams that want the full detail.

What OpenAI Announced

1. GPT-5.6: Three models for different kinds of work

Sol is the newest flagship model for difficult work that requires deeper reasoning, technical ability, judgment, or coordination across multiple tools. Likely uses include complex research, strategic analysis, software development, cybersecurity, scientific work, and sophisticated business problems.

Terra balances capability, speed, and cost. **It is likely to be the practical default for everyday professional work such as research**, planning, data analysis, document creation, and general business support.

Luna is the fastest and least expensive option. It is intended for routine or high-volume tasks such as summarization, classification, extraction, content transformation, first drafts, and standardized workflows.

Model	Best suited for	Input / 1M tokens	Output / 1M tokens
Sol	Complex, high-value work	$5.00	$30.00
Terra	Everyday business work	$2.50	$15.00
Luna	Fast, high-volume tasks	$1.00	$6.00

Terra is 50% less expensive than Sol, while Luna is 80% less expensive. Sol is priced at the same level as OpenAI's previous flagship model (GPT 5.5). The practical benefit is that organizations can reserve Sol for difficult, high-value work and route routine activity to Terra or Luna.

> **Point of confusion worth flagging:** For most employees using ChatGPT through a subscription, these token prices do not directly affect the monthly license fee. They matter most when a company builds API-based workflows or begins paying for usage beyond included allowances.

GPT-5.6 is rolling out across ChatGPT, Codex, and the OpenAI API. Availability may vary temporarily by plan, workspace configuration, and platform while the rollout is completed.

2. ChatGPT Work: From answering questions to completing work

ChatGPT Work is a new agent-style interface for complex, multi-step tasks. A conventional chat helps a user think, draft, analyze, or answer questions. ChatGPT Work is intended to take a broader objective, gather information from permitted sources, break the assignment into steps, use tools, and produce a more complete deliverable across a series of steps.

For example, an employee could ask Work to:

- Review customer interviews and produce a presentation summarizing the main findings.
- Analyze operating data and create a management report explaining significant variances.
- Prepare for a sales meeting using CRM records, internal documents, company messages, and external research.

- Research a new market and produce an editable report, spreadsheet, presentation, or website.
- Monitor an approved source and update a document when meaningful changes occur.

ChatGPT Work can use connected applications, websites, uploaded files, local desktop files, and supported desktop applications. It can also create documents, spreadsheets, presentations, reports, websites, and other work products.

> **For teams using other platforms:** ChatGPT Work is closely analogous to Microsoft Copilot and to Anthropic's Claude Cowork, both agentic workspaces for non-developers. The governance questions in this note, what the AI may see, what it may do, who approves, and who owns the result, apply equally across all of them. A companion guide for Claude is planned.

How Work differs from Agent Mode

ChatGPT's prior Agent Mode introduced the ability to delegate an individual task within a conversation. Work is broader, and replaces Agent Mode. It provides an environment for longer-running assignments, connected applications, generated deliverables, scheduled work, plugins, and reusable workflows.

- Chat helps an employee think and create.
- Work provides a broader environment for managing agent-assisted work.
- Workspace agents package repeatable workflows that an organization can share and govern.

Work should therefore not be treated as simply another chat feature. It materially increases what users may allow ChatGPT to see, combine, generate, and, in some cases, change.

3. A new desktop application combines Chat, Work, and Codex

OpenAI's GPT-5.6 release included two separate changes.

The first was the model upgrade itself. GPT-5.6 improves the underlying AI capabilities available across ChatGPT, Codex, and the API. Organizations can benefit from those improvements without changing how employees currently access ChatGPT.

The second was a redesign of the desktop application. OpenAI began combining ChatGPT, ChatGPT Work, and Codex into a new desktop application (still called ChatGPT). The long-term goal is a unified environment for conversation, agent-assisted work, and software development.

The initial rollout, however, has created confusion.

The new interface hides familiar features such as recent conversations, Projects, and Custom GPTs, confusing users. The transition between the new application and ChatGPT Classic has also been inconsistent, with some users retaining both applications and others reporting that the older version was replaced or moved during installation.

The desktop and cloud experiences are not yet fully consistent. Some Work activity and local files may remain tied to the desktop computer, while conversations and context may not appear uniformly across the browser, desktop app, and other devices.

For organizations with trained IT staff, these issues can be managed through testing, controlled deployment, endpoint review, and user guidance. Where that support is not available, a broad rollout may create unnecessary confusion around installation, local data, account access, and which interface employees should use.

This does not affect access to GPT-5.6

The desktop rollout should be treated separately from the GPT-5.6 model release. Employees can continue to benefit from GPT-5.6 through:

- ChatGPT in a web browser
- The existing ChatGPT Classic application, where it remains installed
- Their current managed ChatGPT workspace and plan

An organization can therefore adopt GPT-5.6 without immediately migrating employees to the new combined desktop application.

> **Stellis recommendation**
> Continue using ChatGPT in the browser or the existing ChatGPT Classic application for normal work. Do not proactively install or broadly deploy the new desktop application until OpenAI improves feature consistency, synchronization, installation behavior, and user guidance.
> Organizations that specifically need ChatGPT Work or Codex should test the new application with a small group of users first. IT should confirm how conversations, Projects, Custom GPTs, local files, and desktop Work threads behave before recommending a wider rollout.
> This approach preserves access to GPT-5.6 and its benefits while avoiding unnecessary disruption from a desktop migration that is still settling.

What Business Leaders Need to Understand

The constraint is no longer the technology. It is people.

Agentic AI that can carry out real, multi-step work is something we have watched approach for a while. What changed with this release is availability: these capabilities are now packaged for ordinary knowledge workers, not just developers and technical specialists. The gating question for most organizations is no longer whether the tools are capable enough. They are. The question is who is going to put them to work enabling the business, and how many people you have who can actually do that well.

Most organizations have a small number of people who instinctively see how to apply these tools to real work, redesign a process, wire up the right sources, and get a genuinely useful result. That capacity, not the model, is the scarce resource. The organizations that pull ahead will be the ones that identify those people early, give them room to work, and let what they learn spread to everyone else.

> **Practical move:** Identify the handful of people in your organization who are naturally good at turning these tools into business results. Get the tools into their hands first, inside the pilot structure described below, and treat them as the engine for wider adoption. Finding and enabling those people is the differentiator, not being first to switch the feature on.

The larger shift is from intelligence to delegated authority

The central issue is no longer only which AI model is smartest. Organizations must now consider:

- What information the AI may access
- Which systems it may connect to
- What actions it may perform
- Whether it may continue working unattended
- Who approves consequential actions
- Who remains accountable for the result

Privacy remains important, but it is only one part of the risk. Once AI can gather information from multiple systems, produce operational artifacts, schedule future work, and act through connected applications, the organization is delegating a degree of authority.

Leaders should be able to answer four questions:

- What may the AI see?
- What may it do?
- Who approves consequential actions?
- Who owns the result?

What Organizations Should Do Now

1. Assign an administrative and business owner

Every organization should assign responsibility for the rollout, ideally your AI Champion. That owner should manage reviewing:

- Work enablement for staff
- User roles and permissions are still appropriate
- Plugins and connected applications are acceptable
- Read and action permissions are configured for your level of safety
- Application deployment
- Usage, cost, support, and policy issues

For an SMB, this role may be held by an operations leader, managed-service provider, security consultant, or technically capable administrator rather than a full-time IT employee.

2. Review Work enablement and plan-specific defaults

Enterprise and Edu organizations receive a preview period during which Work is disabled. It is then scheduled to turn on unless the organization opts out. Administrators should review the settings before employees begin using it.

ChatGPT Business presents a different concern. Plugins and connected applications are generally more open by default, while Enterprise and Edu begin with more restrictive controls. Some connected capabilities, including Google Workspace actions, may also be available by default in Business even though separate authorization from the external provider may still be required.

This is especially important for smaller companies that may have purchased ChatGPT Business primarily for stronger privacy and centralized billing. A managed workspace should not be assumed to be tightly restricted by default.

3. Govern plugins and their underlying applications separately

Also in this release OpenAI has renamed the previous Apps Directory the Plugin Directory. A plugin may contain:

- Skills or instructions
- One or more applications

- Application templates
- Preconfigured workflow components

The plugin and its underlying applications are separate permission layers. Installing a plugin does not necessarily authorize every application it uses. Disabling an application may block access to that system without removing the plugin's other skills or instructions.

Administrators must review both:

- Which plugins employees may install
- Which underlying applications, information sources, and actions those plugins may use

Existing connections should also be reassessed. An application that was previously approved mainly for document search may now support broader workflows, generated deliverables, or actions through Work. Earlier approvals may therefore grant more practical capability than the organization originally intended.

4. Separate information access from action authority

An application that searches company files presents a different risk from one that can send email, update CRM records, modify files, publish content, or change financial or customer information.

Organizations should classify integrations into four groups:

- Approved for read-only use
- Approved to perform selected actions
- Pilot-only
- Prohibited

Where supported, administrators should restrict applications to read-only access, permit only selected actions, and determine when ChatGPT must ask the employee for approval. Some applications do not support granular action controls. In those cases, the organization may need to permit or block the entire application.

Approval prompts reduce risk but do not eliminate it. Employees may approve an action without understanding its full scope, especially when moving quickly.

> **High-impact actions involving email, CRM systems, finance, HR, customer records, file deletion, external publishing, or contractual commitments should require an explicit human-verification step.**

5. Update the organization's AI use policy

Work can access and combine information across applications, websites, cloud repositories, and local files. This creates risks beyond ordinary chat use. Employees may:

- Expose confidential files without understanding the scope of access
- Combine information from multiple systems into a more sensitive aggregate
- Include restricted data in generated reports or presentations
- Create local-data and endpoint-security risks through desktop access
- Connect company systems from a personal ChatGPT account

- Use personal information inside a company-managed workspace

The AI use policy should explicitly address:

- Corporate, customer, employee, and regulated data
- Connected repositories and business applications
- Local files and desktop applications
- Cross-system data aggregation
- Personal versus company-managed accounts
- Read and action permissions
- Scheduled and monitoring tasks
- Human review of generated work
- Responsibility for errors and downstream decisions

Employees should be trained to confirm which account and workspace is active before connecting an application, opening a corporate file, or beginning sensitive work.

Work completed inside a managed account may be retained, administered, exported, audited, or deleted by the organization, depending on its configuration.

You may contact your Stellis AI Strategist for our latest templates.

6. Require review of generated deliverables

OpenAI describes Work as capable of producing finished spreadsheets, presentations, documents, reports, and websites. "Finished" refers to the format and polish of the output. It does not mean verified. Work may produce:

- Invented facts or citations
- Incorrect spreadsheet formulas
- Incomplete research
- Unsupported recommendations
- Misleading visualizations
- Documents that appear more authoritative than the underlying analysis warrants

The employee who submits, publishes, distributes, uploads, or relies upon an AI-generated deliverable remains accountable for its accuracy, sources, confidentiality, and downstream consequences.

7. Begin with a controlled pilot

Work is rolling out across plans and platforms, so employees may temporarily see different screens, settings, models, limits, and administrative controls. OpenAI's documentation and configuration options may also evolve during the launch period.

A pilot is not a delay tactic. It is the fastest safe way to get your strongest people working with these tools and to turn what they learn into repeatable practice for everyone else. Describe the initial deployment as a controlled pilot rather than a full organizational launch, and start with:

- A small group of users

- A limited set of approved applications
- Read-only permissions where possible
- Low-risk internal workflows
- Clear human-review requirements

Suitable first use cases include:

- Internal research using non-sensitive information
- Meeting preparation
- Drafting internal documents
- Summarizing approved materials
- Creating first drafts of reports or presentations
- Read-only analysis of selected repositories

Each pilot workflow should have:

- A named owner
- A defined purpose
- Approved information sources
- Clear action boundaries
- A review step
- Success and risk measures
- An expiration or reassessment date

8. Use Work for supervised acceleration, not production automation

ChatGPT Work is designed for flexible knowledge work. It should not initially replace:

- Audited business processes
- Deterministic workflows
- Regulated decision systems
- Financial controls
- Safety-critical operations
- Customer-facing production automation
- Processes requiring guaranteed completion

Work tasks may pause, usage limits may apply, connected services may change, and model outputs remain probabilistic.

> **Organizations should position Work as supervised work acceleration.** A workflow should move into formal production automation only after it has been tested and provided with appropriate controls, logging, monitoring, exception handling, and clearly assigned technical and business ownership.

Our Recommendation

The immediate response is not to disable these capabilities reflexively, nor to enable everything by default. The five recommendations on page one are the core of a deliberate response: confirm and assign ownership, check your defaults, separate access from action authority, update policy and require review of output, and pilot with your strongest people to move fast without moving recklessly. The preceding sections expand each one.

> **GPT-5.6 improves what the model can do. ChatGPT Work expands where and how that capability can be used.**
> The organizations that benefit most will not be those that simply turn it on first. They will be the ones that find the people who can turn these tools into results, put the capability in their hands, and decide deliberately what the system may see, what it may do, and who remains accountable.

At Stellis, we help organizations evaluate these changes, configure appropriate controls, select practical pilot workflows, and prepare employees to use AI responsibly and productively. As an OpenAI partner, we can also assist with licensing, including rebates on new ChatGPT licenses.

> **About Leader Lab.** Leader Lab is for those who intend to lead in the AI era. Members get briefings like this one, plus direct advice and a practical path to putting AI to work. Learn more or join at **stellis.ai**.

— The Stellis AI Team

hello@stellis.ai | stellis.ai

Selected OpenAI Sources

- GPT-5.6 announcement
- ChatGPT Work announcement
- ChatGPT release notes
- Workspace agents for Enterprise and Business
- Admin controls for apps and plugins
- Plugins in ChatGPT and Codex

