



pdfAssistant Launches MCP Server, Bringing Enterprise PDF Processing Directly into Claude, ChatGPT, and Codex

pdfAssistant launches official MCP server, bringing 55+ professional PDF tools to Claude, ChatGPT, and Codex for secure, high-fidelity AI document workflows.

CHICAGO, IL, UNITED STATES, September 8, 2026 /EINPresswire.com/ -- pdfAssistant, by Datalogics, announced the launch of its official MCP (Model Context Protocol) server, bringing more than 55 professional PDF operations directly into Claude, [ChatGPT](#), and [Codex](#). The server is available now on paid pdfAssistant plans and authenticates via OAuth 2.0.

An AI assistant can read a contract and tell you exactly which clauses need to go. Making that change in the file, without breaking the structure, the fonts, the tag tree, or the color profile, is a different kind of work.

The gap shows up most clearly in redaction. Without purpose-built tooling, removing sensitive content from a PDF often means drawing a black rectangle over it, leaving the original text sitting underneath, still selectable and still extractable. The document looks correct. The file is not. That failure has produced real breach notifications at real organizations, and it is a fidelity problem rather than an intelligence problem. Knowing what to redact is not the same as redacting it.

The pdfAssistant MCP server solves this problem by dividing the work according to where each system excels. While the AI assistant interprets user intent, pdfAssistant executes the operation on a duplicate file with full document fidelity. The original document remains untouched.

Connected through MCP, complex document tasks are handled in seconds through simple natural language prompts, right inside the active conversation or terminal:

Legal and Compliance: An attorney reviewing a settlement agreement can prompt: "Redact the account numbers and SSNs in this agreement, then apply my digital signature." pdfAssistant generates a visual preview of proposed redactions, permanently removes the underlying text upon approval, and applies a tamper-evident digital signature.

Records and Archival: A team receiving scanned board minutes can request: "Apply OCR to these scans and convert them to PDF/A for our retention archive." Scanned paper becomes searchable, standards-compliant archival files ready for long-term storage.

Business Operations: A team finalizing a proposal can say: "Convert this document to PDF, add a CONFIDENTIAL watermark, and restrict printing and editing." The deliverable is converted and secured in the same workspace where it was created.

Developer Workflows: A developer working in the OpenAI Codex CLI can issue a single command: "Compress these PDFs and convert them to PDF/A." Operations execute using the same enterprise engine, delivering consistent, reliable results without requiring custom scripts or tool switching.

Built on What the Assistant Cannot Do Well, Not What It Already Does

The server deliberately contains no summarization, translation, question-answering, or content generation tools. Those capabilities already exist in the assistant, and duplicating them inside it would add nothing.

What the assistant gains instead is the deterministic half: operations that produce byte-accurate files, identical results on every run, and structural changes that hold up when the document is opened somewhere else. Compression. Permanent redaction. Digital signatures. Archival conversion. Form data import and export. Color profile conversion. The work that requires an engine rather than a model.

"Most AI platforms can tell you what a document says," said Eric Shore, Chief Innovation Officer at Datalogics. "Knowing what to change is the easy half. Executing that change without breaking the structure that makes the document work is the half that requires an engine. We deliberately did not ship summarization or translation tools, because the assistant already does that well. The model handles intent. We handle fidelity. That division is the whole point."

Nothing You Cannot Undo

When an AI assistant is choosing which operation to run, the question that matters is what happens when it chooses wrong.

Every pdfAssistant operation works on a copy. Files are processed into new documents and the original is never modified, so an unintended conversion, compression, or watermark costs nothing but the credit it consumed. The source document is still there.

Redaction, the one operation that genuinely destroys content, runs in two stages. The first produces a visual preview of exactly what will be removed. The second applies it permanently. Nothing is destroyed without a chance to look at it first.

And everything, originals and outputs alike, is permanently deleted from our servers within 24 hours.

Three First-Class AI Integrations

The pdfAssistant MCP server ships with dedicated setup guides and native integration support for the three leading AI platforms:

pdfAssistant with Claude: Full MCP integration with [Claude Desktop](#) via Streamable HTTP transport. Add the server with a single CLI command (`claude mcp add`) and authenticate via OAuth. Claude can then execute multi-step PDF pipelines -- such as OCR'ing a scanned contract, redacting sensitive fields, and archiving to PDF/A -- all from a single conversation prompt.

pdfAssistant with ChatGPT: Connect pdfAssistant to ChatGPT via Developer Mode using remote Streamable HTTP transport. Point to the pdfAssistant endpoint, complete OAuth login in seconds, and gain access to the complete document processing suite directly inside ChatGPT conversations.

pdfAssistant with Codex: Native CLI integration purpose-built for Codex developers and automation engineers. Map the server URL, authenticate via OAuth, and immediately integrate PDF conversion, compression, redaction, and OCR directly into build pipelines, CI/CD workflows, and scripted batch processes.

The Operations Your Assistant Gains

Once connected, the full pdfAssistant toolset becomes available to the AI assistant. Users can issue natural language prompts to orchestrate complex workflows across eight tool categories:

Convert into PDF. Microsoft Word, Excel, and PowerPoint files, HTML documents, raster images (JPEG, TIFF, BMP, PNG), .eml email files, and live webpages captured directly from a URL.

Convert out of PDF. Word, Excel, PowerPoint, Markdown, and five image formats (BMP, GIF, JPEG, PNG, TIFF), with control over page range, resolution, and color model including CMYK.

Redact, with review before it commits. Two stages. Preview what will be removed, then apply it permanently.

Secure and control access. Twelve distinct operations covering open passwords, permissions passwords, and restriction sets, addressing every combination of encrypted, unencrypted, restricted, and unrestricted starting states. Plus text and graphic watermarking with control over font, color, opacity, placement, and rotation.

Sign. Digital signatures using either a PFX or PKCS#12 credential or separate DER and PEM certificate and key files, with optional logo and full control over signature field appearance and placement.

Optimize and flatten. Compression profiles, linearization for fast web viewing, rasterization, color profile conversion, and flattening of forms, annotations, layers, and transparencies.

Extract and inspect. Text extraction to plain text or to JSON with per-word style and coordinate data, embedded image extraction, form data export to FDF, XFDF, XML, XDP, or XFD, and structural inspection covering tags, permissions, fonts, and page count.

Assemble and archive. Merge, split by page range, ZIP and unZIP, attach files inside a PDF, place images and styled text with precise positioning including right-to-left support, OCR for searchable text, XFA to AcroForm conversion, and conversion to PDF/A archival and PDF/X print standards.

Every one of these produces the same result each time it runs. That predictability is what makes them safe to put behind an AI prompt, and what makes them usable at volume.

Built for Enterprise Security Requirements

The pdfAssistant MCP server is engineered to meet the data protection standards expected in regulated industries. File payloads and tool calls are encrypted in transit using TLS 1.3. Uploaded files are processed in encrypted storage and permanently deleted within 24 hours of task completion. Datalogics maintains SOC 2 Type 2 audited security controls, and pdfAssistant operates a strict zero data training policy: documents, prompts, and extracted metadata are never used to train AI models. Authentication is handled exclusively via OAuth 2.0, so credentials are never exposed to the local AI environment.

Getting Started

The pdfAssistant MCP server is available now on paid pdfAssistant plans. Developers and AI power users can be connected in under 60 seconds:

Explore the MCP Integration Hub: Review the full overview and all supported platforms at pdfassistant.ai/integrations/ai-connections.

Connect Claude Desktop: Follow the step-by-step guide at pdfassistant.ai/integrations/ai-connections/claude.

Connect ChatGPT: Follow the Developer Mode setup guide at pdfassistant.ai/integrations/ai-connections/chatgpt.

Connect Codex: Follow the CLI integration guide at pdfassistant.ai/integrations/ai-connections/codex.

The server is also listed in the official MCP Registry.

About pdfAssistant

pdfAssistant, by Datalogics, is a comprehensive AI-powered PDF platform offering document processing, automation, and analysis tools for individuals and enterprise teams. With a lineage rooted in nearly 60 years of document technology expertise and the foundational development of PDF standards, pdfAssistant delivers professional-grade PDF capabilities through an intuitive interface, a robust REST API, and now a native MCP server for AI workspace integration. Learn more at pdfassistant.ai.

Eric Shore

Datalogics, Inc.

+1 480-221-1467

[email us here](#)

This press release can be viewed online at: <https://www.einpresswire.com/article/939744005>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.

© 1995-2026 Newsmatics Inc. All Right Reserved.