



X1® Introduces X1 Search MCP Connector for Claude, Bringing Enterprise AI Search In-Place at a Fraction of the Cost

New Capability Enables Legal, Compliance, and Government Teams to Securely Search Emails, Chats, Files, and More in Natural Language Across Their Data Sources, Directly Through Their Own Claude Instance—Significantly Reducing AI Token Costs

LOS ANGELES, CA – July 29, 2026 – X1®, the global leader in index-in-place data discovery software for legal, compliance, governance, and investigation use cases, today extended its AI In-Place architecture with a new capability that establishes X1 Search as the infrastructure layer beneath an organization's Claude deployment, bringing secure, natural-language enterprise search to AI.

This new capability enables legal, compliance, and government teams to search their organization's email, chats, files, documents, and more across their data sources using natural language through their own Claude instance. Organizations receive enriched, contextual results without their data ever leaving their environment and at a fraction of the AI token cost of traditional AI retrieval approaches.

The Model Context Protocol (MCP) makes it possible for AI models like Claude to access enterprise data. But when an AI assistant is asked to search that data, its default approach is to retrieve and stream large volumes of raw emails, files, and attachments into its own context window, then read through them to find what matters. Because every item pulled into context is billed as tokens—whether or not it turns out to be relevant—that approach can sharply increase AI token costs, slow response times, and expose more sensitive information than necessary. The **X1 Search MCP connector** takes a fundamentally different approach.

X1 performs enterprise search first using its existing enriched local index to identify the most relevant information before Claude ever processes the request. Instead of transmitting entire emails, files, and attachments, X1 returns only ranked results, targeted snippets, and file locations—dramatically reducing AI token consumption while improving speed, security, and contextual relevance. The capability integrates seamlessly with an organization's existing Claude deployment through X1 Search-specific licensing, giving IT and security teams full implementation control.

The economics are compelling. At current large-model rates, a single broad search across one year's worth of emails or a file share can cost roughly \$20 in AI tokens when the raw content is streamed into the model—versus a fraction of a cent when only the relevant results are passed to it. Across an organization with thousands of users searching throughout the day, that difference can translate into millions of dollars in avoided AI token costs each year, while also improving speed and reducing data exposure.

“Large language models are incredibly good at reasoning, but they aren't enterprise search engines,” said Larry Gill, CEO at X1. “X1's patented index-in-place architecture has already solved how to find the right information across an organization's data. Our MCP Connector combines those strengths, giving

Claude precisely the information it needs without moving data or overwhelming the model with unnecessary content.”

Search Everywhere, Instantly, Without Moving Data

Search Across All Data Sources in One Query — Search email, chats, files, Microsoft 365, Google Workspace, and more across the entire organization through a single natural-language search.

Reduce AI Token Costs at Any Scale — Large attachments and lengthy document threads return as a single compact result, with only the specific document needed ever opened—keeping AI token costs low, often just a fraction of a cent per search, regardless of data volume.

Keep Data Where It Lives — Search runs against X1's local index on the user's machine or behind the corporate firewall. Content never leaves the environment unless explicitly requested.

Get Instant, Reliable Results — Local index-in-place search delivers fast, consistent performance whether searching thousands or millions of items.

“Without an enterprise search layer, connecting AI to corporate data often means sending entire repositories for the model to sift through,” said Chas Meier, Chief Product Officer at X1. “X1 searches the data locally first, delivering only the most relevant information to Claude—not entire mailboxes or file shares. This allows AI to reason over enterprise data with far greater efficiency, precision, and cost control.”

The X1 Search MCP Connector for Claude is the latest extension of X1's AI In-Place strategy, combining X1's patented index-in-place architecture with the Model Context Protocol standard to securely bring natural-language AI search to enterprise data. It represents the first phase of X1's broader MCP connector roadmap, with support for additional leading AI platforms planned.

To learn more about the X1 Search MCP Connector for Claude or discuss how it can be integrated within your organization, contact info@x1.com.

About X1

X1 is the pioneer of micro-indexing software and global leader in index-in-place data discovery solutions. With more than 600,000 users in 20,000 different organizations worldwide, X1 solutions empower enterprises to meet mission-critical needs in legal, compliance, governance, and investigation data discovery. For more information, visit x1.com.

Contacts

X1 Media Contact
Karen Haworth
E: khaworth@x1.com