CloverDX

WEBINAR SERIES: AI in CLOVERDX

Building Data Pipelines with the AI Assistant

- Quick intro to authoring CloverDX jobs with AI
- Integrated AI Assistant vs External IDEs
- Document existing or inherited project
- Build a project from scratch from a spec

Pavel Najvar
Petr Drabek
September 30, 2026

CloverDX Data Integration Platform

CloverDX

Role of AI for development

A force-multiplier, not a substitute



- Drag boxes onto a canvas
- Code, use documentation

- Hand over a spec, approve proposed architecture
- Review the results in the editor
- Troubleshoot
- Collaborate

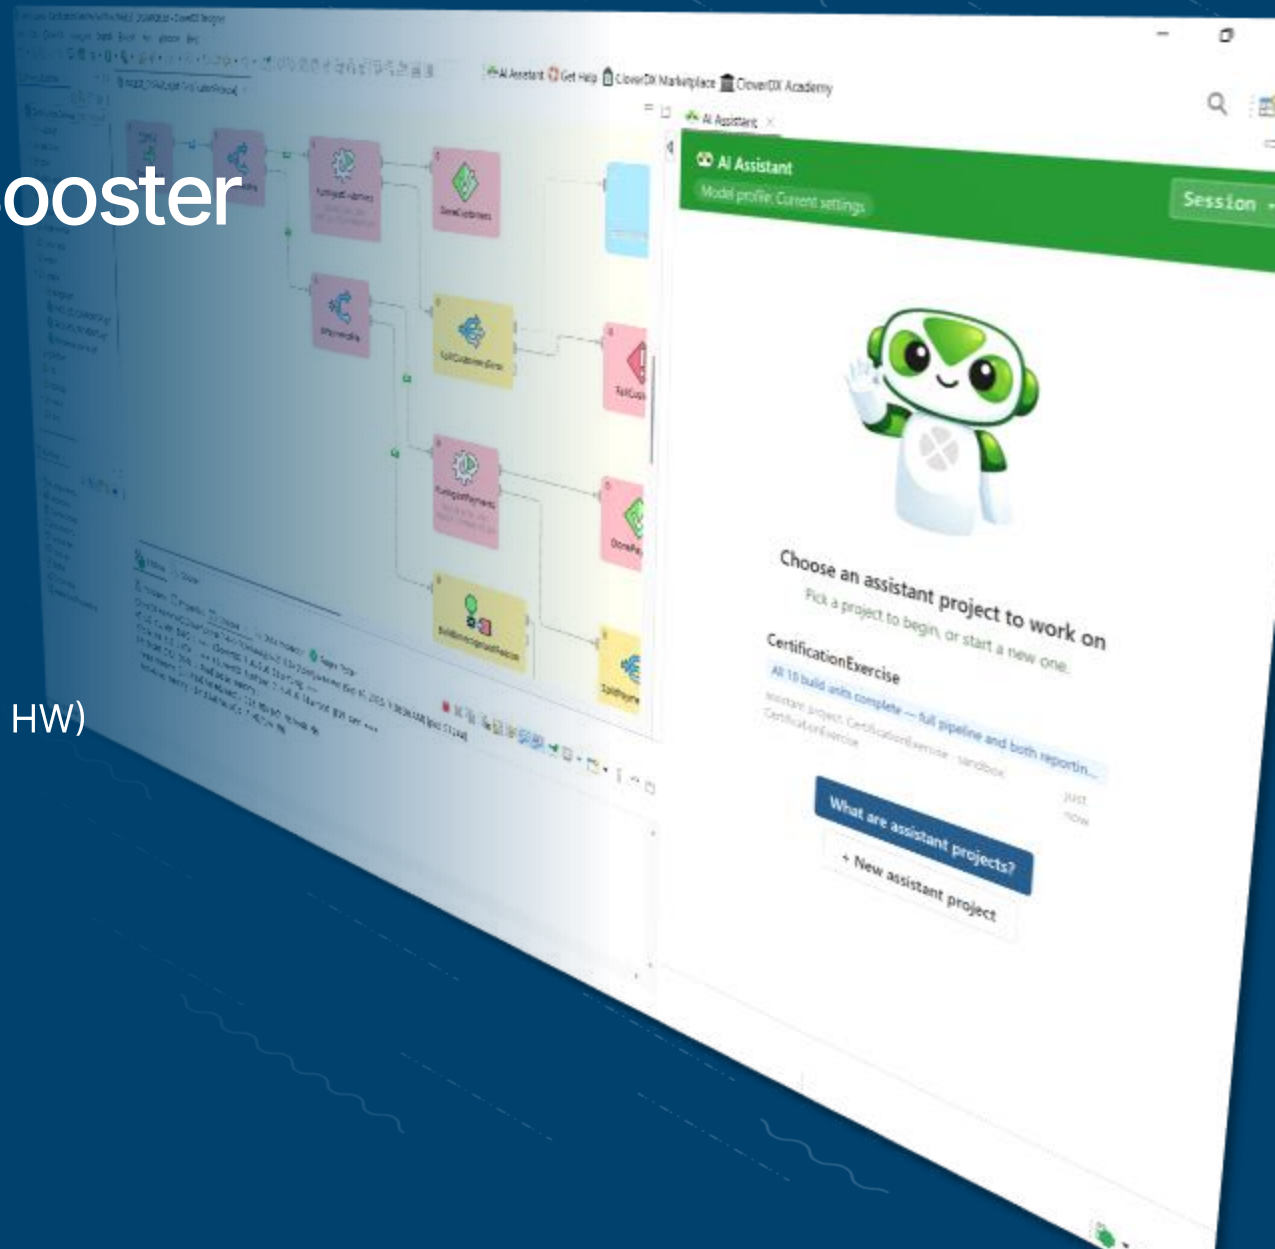
CloverDX

AI Assistant
Developer Productivity Booster

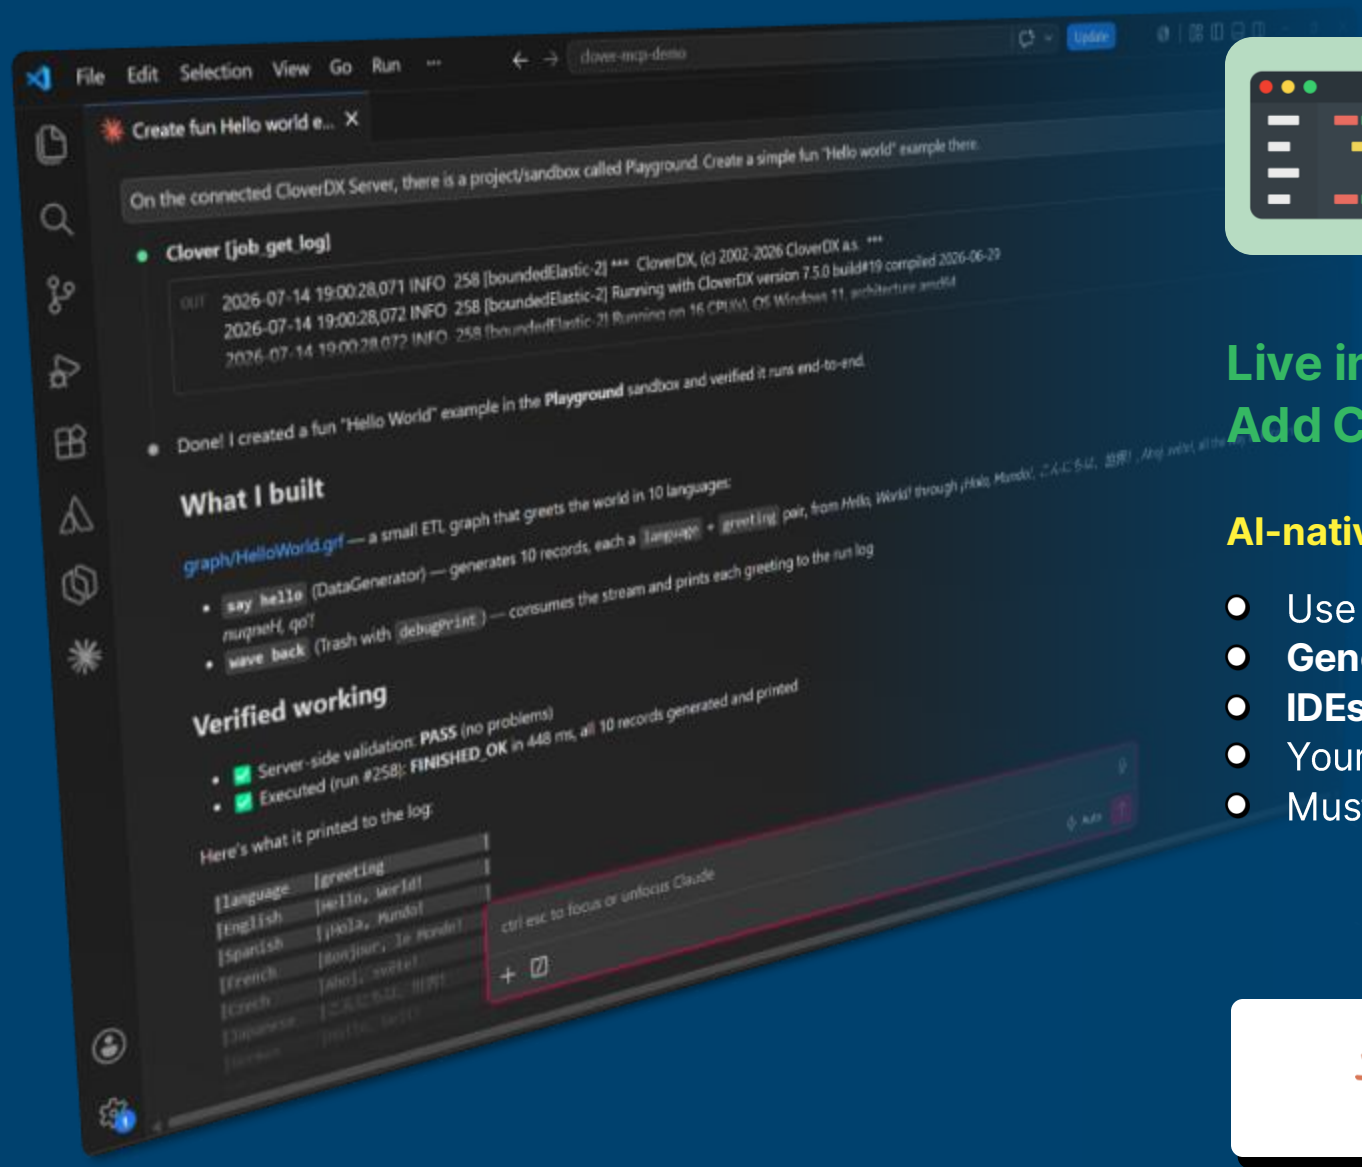
Think bigger. Prototype quickly.
Let AI do the mechanical work.

AI Assistant for developers in Designer:

- Use it as a **companion**, **architect** and **troubleshooter**
- **Rich integrated data-centric agentic environment**
- BYOK – You control the AI – Use your own AI API keys
- BYOM – Use your local LLM (100B+ models, enterprise GPU HW)
- No CloverDX service, no credits
- Designer connects to LLM, No Server AI connectivity



CloverDX



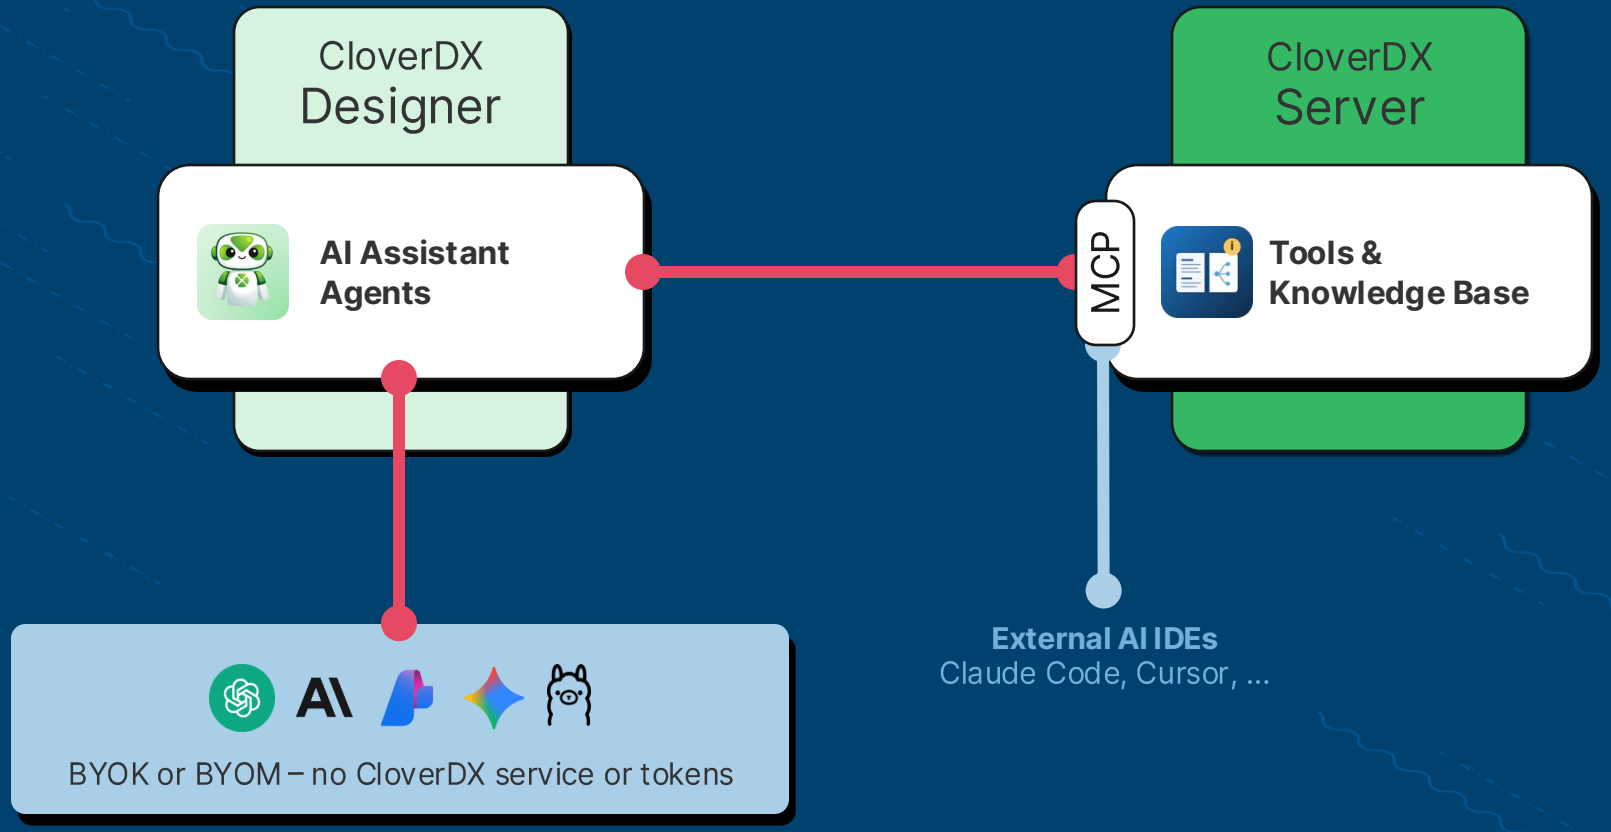
AI-native IDEs
3rd party MCP tools

Live in your favorite development tools. Add CloverDX graph authoring to

AI-native IDEs (Claude Code, Codex, ...):

- Use your own skills and agent orchestration
- **Generic capability but plain-text only**
- **IDEs are built primarily for coding, not data**
- Your own subscription and LLM of choice (MCP-enabled)
- Must connect to CloverDX Server MCP

How does Designer Assistant work?



AI Assistant
Developer Productivity Booster



Companion

Turn to AI with questions.
Offload mechanical tasks to AI.

*"**Help me understand this data?**"*
*"**Build me a graph that does this.**"*
"Add this logic to this piece of CTL code."
"How do I connect to this API?"



Architect

Provide a spec or chat through it.
Let AI document, plan and execute.

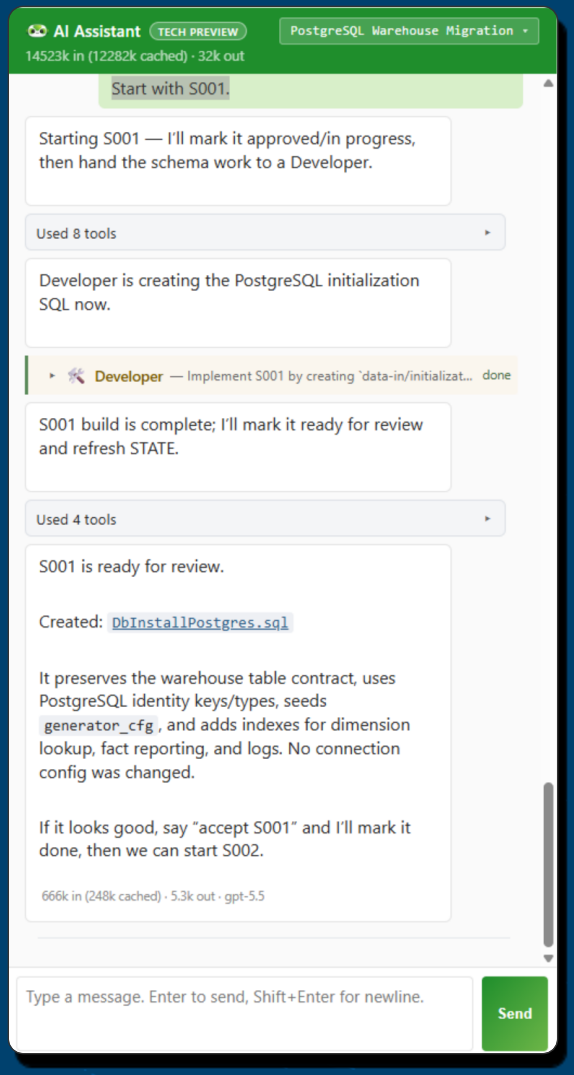
*"**Read specs in this folder and let's plan**"*
*"**Go ahead with sprint 002**"*
*"**Orient yourself in this old project**"*
"Where are we with implementation?"



Troubleshooter

Save time sifting through logs.
Identify and resolve issues.

*"**Why is this graph failing?**"*
*"**Go ahead with the proposed fix.**"*
"Ops reported failed jobs last night, investigate."
"Users have reported incorrect data from yesterday's run..."



Building Data Pipelines
with the AI Assistant

Key Takeaways

Full agentic development workflow in Designer
Architect → Develop → Review → Troubleshoot
AI Assistant ↔ AI-native IDEs via MCP

Next Up

Oct 22: Using your own AI IDE for CloverDX authoring

Q & A

hello@cloverdx.com

CloverDX

Engineering-led data pipelines,
collaborative business tools,
supercharged with AI.

Best practices

Start fresh sessions often

- New session per new task (use the session selector in upper-right corner)
- Short, focused sessions work best

Bring a spec

- Write the spec down first and hand it over (markdown, PDF, ...)
- Use flat-rate subscription chat tool to craft a spec doc

ARCHITECTURE and DECISIONS is your “plan mode”

- Refine until you're happy, then let it plan reviewable sprints and green-light one at a time
- Edit by hand when faster, or ask to walk you through decisions

Experiment with models

- We recommend frontier models for reasoning (Master, Architect).
- No need to use top-tier models for mechanical work (e.g. Sonnet for Developer is great).

CloverDX

Undo Assistant's change

- Ask the Assistant **revert its last change** if you don't like what it did
- It keeps an internal git versioning just for the Assistant

Useful chat commands

- **/read-only**
  - ✓ Puts the session TEMPORARILY into read-only mode: Useful when you want to make absolutely sure the Assistant won't make any changes.
- **/read-write**
  - ✓ Restores full read-write access.
- **/compact**
  - ✓ Summarizes the older part of the conversation and replaces it with that summary, which frees up room in the context window. Use it when a long session is filling up and you would rather continue in it than start a new one.

CloverDX

RELEASES TIMELINE

TECH PREVIEW

GA - AI AUTHROING

7.5.0 → 7.5.1 → 7.6.0

- **Idea & core**

- **AI capabilities**
- **Optimization**

- Moves **out of Tech Preview**
- Commercial “**AI Authoring Seat**” on Server → unlocks Assistant & MCP for external IDEs
- **Global knowledge base** shared across all projects (managed on Server)
- **OAuth2** authorization for MCP and APIs (with “Built-in” for hassle-free MCP setup)
- **File access control** via allow/deny wildcards

How you're billed

BYOK/BYOM. No CloverDX credits/tokens.

- Your provider bills you for tokens directly; CloverDX meters nothing.
- Assistant license and MCP Server: complimentary during the tech preview, per seat.
- Provider usage is the only variable cost under your control.

CloverDX

Control budget with provider

Set your budgets per API key with your provider

- Agentic work can use meaningful tokens.
- Set a spend limit or alert in your **provider dashboard**.

CloverDX

Experiment with models

Find most cost-effective fit for your needs

- We recommend frontier models for reasoning (Master, Architect).
- No need to use top-tier models for mechanical work (e.g. Sonnet for Developer is great).
- You could go lower and combine models depending on task

CloverDX

Start fresh sessions often

The single biggest cost saver

- Every turn re-sends the whole conversation.
- New session per new task (Use the session selector in upper-right corner).
- Short, focused sessions cost far less.

CloverDX

Bring a spec

Conversation is an expensive way to convey detail

- Write the spec down first and hand it over
- Use flat-rate subscription chat tool to craft a spec doc
- Reserve metered tokens for the build
- Markdown is ideal, but PDF, an email thread, a sample file, or rough notes all work; it normalizes messy input well.
- Short, focused sessions cost far less.

CloverDX

Architecture is your “plan mode”

Architect first, then build.

- Treat ARCHITECTURE.md as plan mode; refine until you're happy, then let it plan reviewable sprints and green-light one at a time
- Answer its questions – when it asks, the answer needs business judgment, your answers up front save rework and tokens.
- Stay the judge – open the graph editor, edit by hand when faster, rewrite a decision if you change your mind.
- Use the assistant project artifacts. Read and reference the architecture doc, decision log, journal, and plan in review.

CloverDX

Reuse work, don’t reinvent

Work with project files (ARCHITECTURE, JOURNAL, DECISIONS,...)

- For bigger tasks the Assistant crafts and stores session in the project's assistant folder in the sandbox – these files travel with the project and are visible to the team.
- Sessions are stored as well – there is nothing like a “local session” – all session history is stored in the sandbox

CloverDX

Should you trust it?

Human checkpoints are built in by design

- Load-bearing decisions are surfaced before action; the plan is proposed, not imposed.
- Every unit of work ends at a reviewable state you mark done.
- Interrupt anytime, do a quick course correction.
- Your hand-edits are the source of truth – it can re-read specs, project documentation.
- The audit trail is built-in: decision log (DECISIONS.md, choice + rationale + status), a journal, and recorded change history that captures why, not just what.

CloverDX

Can it hallucinate? Sure, it's AI.

You're the ultimate reviewer, but with help.

- Reduce hallucination at the source: give it real sample files, a real schema to inspect, an existing example to point at.
- It's not a black box – everything is a reviewable, modifiable, documented artifact.
- Ask it to revert a change if you don't like it
- Every unit of work ends at a reviewable state you mark done.
- Interrupt anytime, do a quick course correction.
- Your hand-edits are the source of truth – it can re-read specs, project documentation.
- The audit trail is built-in: decision log (DECISIONS.md, choice + rationale + status), a journal, and recorded change history that captures why, not just what.

CloverDX

Chat commands

- **/compact**
  - ✓ Summarizes the older part of the conversation and replaces it with that summary, which frees up room in the context window. Use it when a long session is filling up and you would rather continue in it than start a new one.
- **/read-only**
  - ✓ Puts the session into read-only mode: tools that only read stay available and everything that would write, run or commit is refused. Useful when you want the Assistant to investigate something with no chance of it changing the project.
- **/read-write**
  - ✓ Restores full access.
- **/profile**
  - ✓ Switches this chat view to a different set of models. Type /profile on its own to pick from the profiles you have configured. The switch applies to this view only and is forgotten when you close it

CloverDX

Finally, can I undo?

Yes!

- Ask the Assistant revert its last change if you don’t like what it did
- It keeps an internal git versioning just for the Assistant
- Don’t try to abuse it for your own version control (for now)