OPENAI ARCHITECTURE PRIMER

Codex Non-Interactive Mode

A governed execution contract
for scripts and pipelines

A visual guide to predictable input, session state, output, permissions, and automation identity with codex exec.

FIRST-PARTY DOCUMENTATION / INDEPENDENT LUMINITY SYNTHESIS



Luminity Digital
INTELLIGENCE, ENGINEERED. DELIVERED.

Codex Non-Interactive Mode

Source verified Nov 5, 2026

OpenAI Select Partner

Codex Exec Turns an Agent Run Into an Execution Contract

The command is scriptable only when its inputs, session state, authority, outputs, and identities are explicit.

FIVE BOUNDARIES

01
Input
Prompt + piped context

02
State
Ephemeral or recorded

03
Authority
Sandbox + rules

04
Output
Text, events, or schema

Identity boundary

05

Keep the model credential distinct from the workflow identity that reviews or applies a consequence.

Interpretation boundary

RULE

A completed command proves execution. It does not prove that the inputs, authority, result, or downstream action were appropriate.

IN THIS PRIMER

Key Points

01 Compose the run

Separate the instruction, piped context, working directory, and persistence choice.

02 Bound authority + output

Choose the downstream interface independently from sandbox, rules, tools, and credentials.

03 Preserve evidence

Produce a report, structured result, or patch before a separately authorized consequence.

The five-boundary execution contract is Luminity synthesis. The underlying controls are documented OpenAI product behavior.

© 2026 Luminity Digital, Inc. | Independent synthesis of official OpenAI documentation.

Source: https://learn.chatgpt.com/docs/non-interactive-mode · verified Nov 5, 2026

OpenAI Select Partner

Compose the Run

Prompt delivery, working context, and persistence determine what the run knows and what evidence remains.

01

What enters the run?

- ✓ Positional instruction — PROMPT
- ✓ Piped stdin with a prompt — CONTEXT BLOCK
- ✓ codex exec - — STDIN IS PROMPT
- ✓ Working directory — REPOSITORY CONTEXT
- ! Unbounded ambient context — DO NOT ASSUME

Instruction → Context → codex exec

02

Which state should survive?

Input · prompt + context | Scope · working directory | State · deliberate choice

1 **Ephemeral** — No rollout file

2 **Recorded** — New session history

3 **Resumed** — Prior context continues

4 **Run** — One effective request

RETENTION
--ephemeral removes the continuation path

CONTINUITY
resume uses an identifier or --last

EVIDENCE
State choice changes recovery and review

03

Four decisions to record

Instruction source — INPUT

Record the prompt file, inline instruction, or generated stdin source.

Context boundary — SCOPE

Record piped content, repository revision, and working directory.

Persistence — STATE

Choose ephemeral, newly recorded, or resumed execution deliberately.

Continuation key — RESUME

When resuming, preserve the exact session identifier or the rule behind --last.

DO NOT ASSUME

- stdin always replaces the prompt
- ephemeral can be resumed
- --last identifies one global run
- cwd is incidental

© 2026 Luminity Digital, Inc. | Independent synthesis of official OpenAI documentation.

Source: https://learn.chatgpt.com/docs/non-interactive-mode · verified Nov 5, 2026

OpenAI Select Partner

Separate Output From Authority

The downstream interface and the run's permissions solve different problems and must be chosen independently.

01

Which output is consumed?

✓	Final message on stdout	DEFAULT
✓	Progress and diagnostics	STDERR
✓	JSONL event stream	--JSON
✓	Schema-constrained result	--OUTPUT-SCHEMA
!	Free-form text parsed as data	AVOID

Human · Event consumer · Machine contract

02

How authority is assembled

Runner · host boundary · Identity · API credential · Tools · required services

1 **Sandbox** — read-only by default

2 **Approvals** — review or never

3 **Rules** — config + execpolicy

4 **Tools** — required MCP fails closed

DEFAULT
Read-only unless explicitly widened

AUTOMATION
--full-auto retired; use explicit controls

CREDENTIAL
Sandboxing does not isolate an exposed key

03

Four controls to retain

Explicit sandbox — FILES

Use read-only, workspace-write, or controlled full access; do not rely on a retired shortcut.

Approval route — REVIEW

Use --approve-for-me only as a deliberate automation control with bounded write scope.

Configuration inputs — CONFIG

Record whether user config and execpolicy rules were loaded or ignored.

Credential isolation — IDENTITY

Keep the model key out of repository-controlled steps and separate it from applying authority.

CONTROL RULE

- Output shape is not authorization
- Sandbox layers do not substitute
- Required tools should fail closed
- Credentials remain separate

Luminity Digital
INTELLIGENCE, ENGINEERED. DELIVERED.

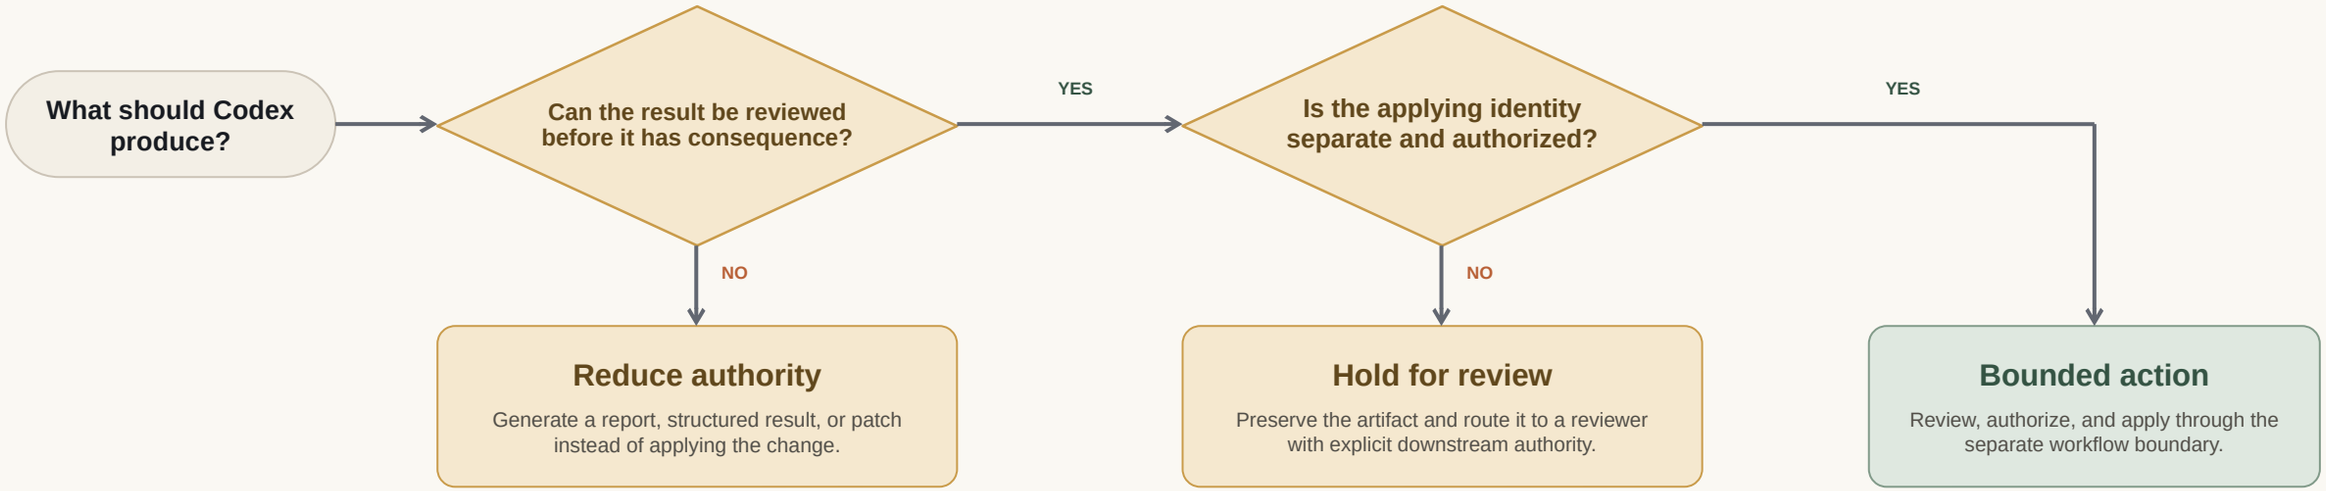
© 2026 Luminity Digital, Inc. | Independent synthesis of official OpenAI documentation.

Source: https://learn.chatgpt.com/docs/non-interactive-mode · verified Nov 5, 2026

OpenAI Select Partner

Evidence Before Consequence

Use a restricted Codex run to create a reviewable artifact before a separate identity applies a change.



RETAIN THIS RUN RECORD WITH EVERY ARTIFACT

Input + revision — SOURCE

Prompt source, supplied context, repository revision, and working directory.

State + authority — RUN

Persistence mode, sandbox, approvals, rules, tools, and network boundary.

Output + diagnostics — EVIDENCE

Chosen interface, final result, error stream, and artifact hash.

Identity + decision — CONTROL

Model credential boundary, reviewer, applying identity, and authorization record.

ARCHITECTURE RULE

- Generation does not grant consequence
- A patch is evidence, not authorization
- Apply through a separate identity

Luminity Digital
INTELLIGENCE, ENGINEERED. DELIVERED.

© 2026 Luminity Digital, Inc. | Independent synthesis of official OpenAI documentation.

Source: https://learn.chatgpt.com/docs/non-interactive-mode · verified Nov 5, 2026

OpenAI Select Partner