


Side Panels

Inspector / Details

The Inspector (opened via the **Details** side panel) shows everything about the currently selected node. It has a **Pretty / JSON** toggle and displays the node's type and status (e.g. `LLM_CALL · COMPLETED`), name, an ownership line ("*in {parent} · Iter N · {duration}*"), a *Final results* block, and collapsible sections.

User can close it with the **x**.



The panel is read top-to-bottom. There are several sections:

1. Format toggle + identity header

At the very top a **Pretty / JSON** toggle switches the whole panel between a human-readable layout (**Pretty**) and the raw payload (**JSON**). Under it is the node's identity:

- **Type badge + status** — e.g. `LLM_CALL` with a green **COMPLETED** pill (other statuses: `RUNNING`, `PENDING`, `FAILED`).
- **Name** — the step's name (e.g. *Thinking*).
- **Ownership line** — The parent name is a link

Inspector

Pretty JSON

• LLM_CALL COMPLETED

Thinking

in Template Analyzer · Iter 1 · 6.04s

Final results:

Open LLM call → Parsed messages

PROMPT 2 B Pretty JSON

RESPONSE 1.6 KB Pretty JSON

Timing 6.04s

Started 20:14:38 · 36 d ago

2. Final results

The **Final results** block is the node's output.

- **LLM node** — a quick **Open LLM call → Parsed messages** button that takes you to the full panel, plus short
- **PROMPT** and **RESPONSE** previews. Each preview shows its size and has its own **Pretty / JSON** toggle.

Inspector

Pretty

{ } JSON

• LLM_CALL

COMPLETED

Thinking

in Template Analyzer · 🔁 Iter 1 · ⌚ 6.04s

Final results:

🔗 Open LLM call
Parsed messages

⬇️ PROMPT 2 B
Pretty JSON

⬆️ RESPONSE 1.6 KB
Pretty JSON

Timing

6.04s

▶ Started

20:14:38 · 36 d ago

3. Timing

Timing shows exactly when the step was ran:

- **Started** — absolute clock time plus a relative age
- **Completed** — when it finished.
- **Duration** — how long it took.
- **Progress** — completion percentage (100% for a finished step).

The screenshot shows the 'Inspector' panel with a close button (X) in the top right. Below the title bar, there are tabs for 'Pretty' (selected) and 'JSON'. The main content area is divided into two sections: 'PROMPT 2 B' and 'RESPONSE 1.6 KB', each with 'Pretty' and 'JSON' view options and a right arrow. Below these is a 'Timing' section, which is highlighted with a red border. It shows a total duration of '6.04s' with a dropdown arrow. The timing details are as follows:

Event	Time
Started	20:14:38 · 36 d ago
Completed	20:14:44
Duration	6.04s
Progress	100%

Below the timing section is a 'References' section, also with a dropdown arrow showing '3' items. It contains two entries:

- # JOB ID
llm-call-TemplateAnalyzer-b442-1d6490f7c080
- ← PARENT JOB
TemplateAnalysisThinker-bb67-04b5221fbd56

4. References

References displays the identifiers that tie this node into the run:

- **JOB ID** — this node's own id (the value used in `sel=`).
- **PARENT JOB** — the id of the step that owns it.
- **WORKFLOW** — the trace id it belongs to.

Inspector

Pretty

{ } JSON

Started

20:14:38 · 36 d ago

Completed

20:14:44

Duration

6.04s

Progress

100%

References

3

JOB ID

llm-call-TemplateAnalyzer-b442-1d6490f7c080

← PARENT JOB

TemplateAnalysisThinker-bb67-04b5221fbd56

>

WORKFLOW

tga-sicherheits--un-8188-e6925bc244fb

Tokens

6 100 tok

↓ Input tokens

5 963

↑ Output tokens

137

Total tokens

6 100

5. Tokens

Tokens displays token breakdown that drives cost:

- **Input tokens / Output tokens** — the split between prompt and completion.
- **Total tokens** — their sum.
- **LLM requests** — how many model calls this node made.

Inspector

Pretty

{ } JSON

llm-call-TemplateAnalyzer-b442-1d6490f7c080

← PARENT JOB

TemplateAnalysisThinker-bb67-04b5221fbd56

→

WORKFLOW

tga-sicherheits--un-8188-e6925bc244fb

Tokens

6 100 tok ▾

↓ Input tokens

5 963

↑ Output tokens

137

Σ Total tokens

6 100

LLM requests

1

Metadata

7 ▾

State

COMPLETED

Type

LLM_CALL

Class

LLMCall

Package

ai.redouble.tools

6. Metadata

Metadata displays the low-level attributes of the step:

- **State**
- **Type**
- **Class**
- **Package**
- **Owner**
- **Iteration**
- **Status message**

Inspector

Pretty

{ } JSON

tga-sicherheits--un-8188-e6925bc244fb

Tokens

6 100 tok

↓ Input tokens

5 963

↑ Output tokens

137

Σ Total tokens

6 100

⚙ LLM requests

1

Metadata

7

✓ State

COMPLETED

🔑 Type

LLM_CALL

📄 Class

LLMCall

{ } Package

ai.redouble.tools

👤 Owner

Andrey Redouble.ai

Iteration

1

📄 Status message

Completed

Tool I/O

The **Tool I/O** panel shows the input and output of a tool call.

Traces

tga-sicherheits--un-8188-e6925bc244fb

91 nodes · 39 LLM · 1 116 853 tok

Graph

Waterfall

Tree

Deps

Tool I/O

Details

Tool I/O

Get Project Structure

TOOL · GETPROJECTSTRUCTURETOOL

1ms

↓ INPUT 24 B

Pretty JSON

PROJECT ID

201

↑ OUTPUT 1.8 KB

Pretty JSON

SECTIONS

4 items

Switch inventory and firmware

8 fields

ID

1 708

TITLE

Switch inventory and firmware

TOPIC

Device Inventory

DISPLAY ORDER

1

STATUS

InProgress

Inspector

Pretty

{ } JSON

• TOOL

COMPLETED

Get Project Structure

in Template Analyzer · Iter 1 · 1ms

Get all questions in this project and their answer status. Returns question IDs, titles, topics, and content snippets. Use this to see which questions already have answers and identify content that may overlap with yours.

Final results:

Open tool I/O

All fields · Raw JSON

↓ INPUT 24 B

Pretty JSON

↑ OUTPUT 1.8 KB

Pretty JSON

Header

The top of the panel names the tool and how the call went:

- Tool I/O

Get Project Structure

TOOL · GETPROJECTSTRUCTURETOOL

✓

🕒 1ms

↓ INPUT 24 B

Pretty JSON ▾

PROJECT ID

201

↑ OUTPUT 1.8 KB

Pretty JSON ▾

SECTIONS

4 items ▾

Switch inventory and firmware

8 fields ▾

ID

1 708

TITLE

Switch inventory and firmware

TOPIC

Device Inventory

DISPLAY ORDER

1

STATUS

InProgress

The **INPUT** block shows the parameters the tool was called with. Its header carries the payload **size** and a **Pretty / JSON** toggle, then each parameter is listed as a name/value pair.

Tool I/O

Get Project Structure

TOOL · GETPROJECTSTRUCTURETOOL

✓

🕒 1ms

↓ INPUT 24 B

PrettyJSON

PROJECT ID201

↑ OUTPUT 1.8 KB

PrettyJSON


SECTIONS4 items

Switch inventory and firmware8 fields

ID	1 708
TITLE	Switch inventory and firmware
TOPIC	Device Inventory
DISPLAY ORDER	1
STATUS	InProgress

Output

The **OUTPUT** block shows the result. Its header also carries the **size** and a **Pretty / JSON** toggle. When the result is structured, Retrax lays it out as **collapsible sections** — so a big response stays readable instead of dumping a wall of JSON.



User can perform several types of actions:


1. Select a **tool** node (e.g. "Get Project Structure"). The **Tool I/O** button becomes active.
2. Open **Tool I/O**.
3. Read the **input** — the parameters the tool was called with.
4. Read the **output** — the result the tool returned.
5. If a format toggle is available, switch between **Pretty** and **JSON**.

Dependencies (Deps)

The **Dependencies** panel reveals scheduling dependencies between nodes.

User can perform several types of actions:

1. Click **Deps** to open the dependencies panel / mode.
2. Find a node marked "Has scheduling dependencies" — the indicator confirms it has them.
3. Select a node that has dependencies — its related nodes / edges are highlighted.
4. Select an independent node — no dependencies are shown (empty), confirming it isn't waiting on anything.



LLM Call

When user selects an LLM node, the middle side-panel button becomes **LLM call**.

Opening it shows the details of that model call:

- **model** and provider (e.g. *claude-sonnet-4-6* · *BEDROCK*)
- a timestamp
- a status
- a token breakdown
- raw prompt/response.

LLM call

claude-sonnet-4-6
BEDROCK

20:14:38 SUCCESS

INPUT	4 773	OUTPUT	137
CACHE CREATE	4 763	BILLABLE IN	5 963

PROMPT - RAW

```
{}
```

RESPONSE - RAW

```
{
  "id": "msg_bdrk_012vozVcsVEDGfG9QojXW8YP",
  "content": [
    {
      "signature": "ErMECL0IDhABGAiQLHPLuL6CEJLtlFL2YW+Nl3igMKVJwS9VK0EjblwEGZmcJ5NbEhrYY3wcBXBN0KJIjhlCcwQ732PAu1BDiUHS298yEWnsYXVvZS1zb25uZXQTNc020AASDNQXA0FMaksQnvPnCBoM5Sc170xHw1rgeD111jBEz9Qb16BpIwVRhpaaljq047YbP+JciasznkSgoxocw29NommCFXxKiH1FD4dYqgwONtgcFyt0ny2AtCRXLEtGpsgamIwuiaorBacoJJva17wsUwahJEVhnJSRUjjRM0owbzs+vyQzQAmC+Z2Qnr/mjzh0gLMaCkvpak0i90DpEDE0yVvKAAjuytQA1FYhsd/b2xWnh9WX6s329qycmqdTG98b4M2FakJHHjuJ5RZftZDpfFMwv0QUw/OrFMsmognFfswenHHqtCVSUF9Xi4N8TAKg8DKHra8lbg78ZdG8Ldg5qU4yhLzq0HdCwIjLjBy2PCBwn7jjXF23w52dg4Vy/s1izQL2qKtNPDVP1t4GdCRLWTP0mLXL6yxZ52rda8GhIutJxkXQ/mL5z0nuZJuLdf/PA/jONhScz9rDCHHL+NDc58iutCm2o5d8hViWfAit93UHRNKxg+C0iQTqvtYaDZ9qXhvbRXGLPm0PXhdiErten1TWKE1mTSjRD0ZLXdHEiupHAUIxlcnR5ptCi+Q01YwkvXmcUz0BSGHakFzH9He1D9M6HuvXGUD6RIMX7haoYA0==",
      "thinking": "Let me analvze this structural"
    }
  ]
}
```

Inspector

Pretty {} JSON

LLM_CALL COMPLETED

Thinking

in Template Analyzer · Iter 1 · 6.04s

Final results:

Open LLM call
Parsed messages

PROMPT 2 B Pretty JSON

RESPONSE 1.6 KB Pretty JSON

Timing 6.04s

Started 20:14:38 - 36 d ago

Revision #5

Created 2026-07-02 14:54:30 UTC by Diana

Updated 2026-07-06 17:19:57 UTC by Diana