



BITO WHITEPAPER · OCTOBER 2026

Autorouting Coding Agents with Dynamic Context

How Bito Governor uses code context and model routing to cut what coding agents cost, request by request

Based on 1,800 graded answers from 30 models and more than 1,000 multi-turn agentic coding sessions

bito.ai

EXECUTIVE SUMMARY

Coding agents work in long sessions. A Claude Code session runs for dozens of turns, each turn is a chain of tool calls, and every request sends the whole conversation again. Routers built for single questions misjudge this traffic: they send a short "go ahead" to the cheapest model without seeing the job it belongs to, and they throw away prompt caches the customer has already paid to build.

Bito Governor treats the session as the unit of work. On every request it weighs the kind of work and what a mistake would put at risk, the task in progress, which model already holds the conversation in its cache, how hard the model should reason, and how the prompt is engineered. It also gives the agent the code it needs up front, so the agent spends far fewer steps searching. On Bito's benchmarks:

- **Quality setting:** developers on Opus 5.5 keep their quality while cutting cost 1.3x and finishing 12% faster.
- **Balanced setting:** giving up 5% of the score cuts cost 1.9x and time by 39%.
- **Cost setting:** keeps 94% of the score and cuts cost 7.1x.
- **Code context:** cuts cost per session 1.2x to 1.4x on its own, at the same quality.
- **Routing by tier:** in our public study of 30 models, routing each tier of work to a different model stayed within 2 of 60 points of Opus at up to 50x lower cost.

1. Why coding agents need a different kind of router

Three facts shape the design of any auto-routing system for coding agents.


Frontier models still earn their price, on some of the work

Our public study graded 1,800 answers from 30 models working the same 60 real engineering tasks (bito.ai/benchmarks). It found the top eight models within 4.5 points of each other out of 60, at prices from \$2.67 to \$135.51 for the full run. On short, well-defined jobs, many models are close (Exhibit 1).

Longer, deeper work shows a bigger gap. We also run an internal benchmark of long, multi-turn sessions on proprietary code no model has seen in training: find the code involved, work out why something failed, plan a fix, build it, then audit it. Each step is graded against a written answer key that rewards production-quality work, such as tests for the failure paths and checking every caller before changing an interface. On the same hard five-turn sessions, Sonnet 5.5 cost about a quarter as much as Opus 5.5 and scored 5 points lower out of 50. In the public study, the cheap open-weight models matched Opus on writing code and fixing bugs, then fell away on code review, deciding what was safe to ship, and refusing a dangerous instruction. Opus 5 passed 8 of those 9 tasks, Opus 5.5 passed 6, and the two strongest open-weight models passed 4.5 and 7 (Exhibit 2).

EXHIBIT 1

Most models are close on score, while the price varies 50x



Score is the share of points each model earned on the same 60 real engineering tasks. Cost is the full run at list prices. The line connects the models that nothing cheaper beats on score. Source: Bito model study, September 2026.

EXHIBIT 2

The top models only pull apart on complex work

	Simple	Medium	Complex	Reasoning
claude-opus-5	100% 19 of 19 · \$12.12	93% 13 of 14 · \$43.09	89% 8 of 9 · \$4.83	81% 14.5 of 18 · \$75.48
claude-opus-5-5	100% 19 of 19 · \$6.55	93% 13 of 14 · \$5.82	67% 6 of 9 · \$2.81	92% 16.5 of 18 · \$18.04
gpt-6-astra	100% 19 of 19 · \$9.64	96% 13.5 of 14 · \$32.29	89% 8 of 9 · \$9.59	75% 13.5 of 18 · \$77.75
claude-fable-5.1	100% 19 of 19 · \$13.43	96% 13.5 of 14 · \$11.65	72% 6.5 of 9 · \$6.55	75% 13.5 of 18 · \$53.85
gemin-3.8-flash	100% 19 of 19 · \$4.25	89% 12.5 of 14 · \$2.15	72% 6.5 of 9 · \$1.87	75% 13.5 of 18 · \$9.02
grok-4.6	100% 19 of 19 · \$3.70	96% 13.5 of 14 · \$8.08	67% 6 of 9 · \$2.33	67% 12 of 18 · \$18.13
grok-4.5	100% 19 of 19 · \$1.84	89% 12.5 of 14 · \$2.70	50% 4.5 of 9 · \$1.86	78% 14 of 18 · \$10.94
deepseek-v4.1-flash	95% 18 of 19 · \$0.20	86% 12 of 14 · \$0.40	78% 7 of 9 · \$0.14	75% 13.5 of 18 · \$1.94


Share of each tier's tasks passed by the eight highest-scoring models, with what each spent on that tier. Complex work (code review, deciding what is safe to ship, refusing a dangerous instruction) is the only tier with a wide spread. Source: Bito model study, September 2026.

A router has to tell which requests need a frontier model. If it sends everything to the cheapest model that scores well on a public benchmark, developers spend the savings redoing weak work.

Most of what an agent pays for is conversation it already sent

Every step an agent takes adds a result to the conversation, and the whole conversation goes back to the model on the next step. In one measured session, the agent’s searches returned about 65,000 tokens of results, and the session processed about 6 million. Cache reads were 77% of all tokens across the public study. Prompt caching makes those re-sends cheaper, but switching models mid-session means paying to build a new cache from scratch. How a model works also matters as much as its list price. Claude Fable 5.1 lists at twice the per-token price of Opus 5 and cost 37% less per finished task in our study, because it took fewer steps and wrote less (Exhibit 3).

EXHIBIT 3
The list price is not the bill



Each model's full bill divided by the tasks it fully passed, for the thirteen models that passed at least 40 of 60, next to the vendor's list price. Claude Fable 5.1 lists at twice the price of Opus 5 and costs 37% less per finished task. Source: Bito model study, September 2026.

The newest message rarely describes the work

"Now do the same for the other repo" and "looks good, implement this" are short messages that carry the full weight of the task they are continuing. Each one is a


command inside a longer conversation and only makes sense as part of it. Partway through a turn, the newest text in the request is often boilerplate the coding tool adds around a tool result, not anything the user typed.

Many open-source gateways ship a keyword-based complexity router. LiteLLM and Bifrost both do ([Bifrost's documentation](#) describes its version). These routers add up points for signals like code words, reasoning words and technical terms into a score between 0 and 1, then cut the score into four tiers, each pointing at a cheaper or more expensive model. Exhibit 4 shows the tier lines.

EXHIBIT 4

How a keyword router assigns tiers

11 of 12 typical coding-agent turns we tested scored as Simple



Tier lines as published in [Bifrost's documentation](#). LiteLLM's scorer uses the same 0.15 line between simple and medium, and prompt length can add at most 0.10 to its score.

We ran twelve typical coding-agent turns through LiteLLM's scorer, which reads the newest message, and it put eleven of them in the simple tier, including "now do the same for the other repo" and a one-line request to refactor authentication code across every place it is called. Part of the reason is arithmetic. In that design, the length of the prompt can add at most 0.10 to the score, and the line between simple and medium sits at 0.15. A coding request carrying 150,000 tokens of history and no trigger words still lands in simple. Tuning the keywords can't fix this, because the newest message doesn't hold the information the scorer needs.

2. What Governor controls

Governor is a gateway and context-driven auto-model router. A company points Claude Code, Cursor, Codex or its own agents at one base URL, and nothing else about its setup

changes. On every request, Governor can act on a whole set of factors that decide cost and quality:

- **The code the agent sees.** Governor's code context engine keeps a live map of the customer's repositories and gives the model the relevant code up front, so the agent skips most of its searching. This is context about the code itself, its structure and how its parts connect, rather than business or domain knowledge.
- **Which model answers,** chosen from the models the customer allows.
- **How hard the model reasons.** Reasoning level is a dial between speed, cost and depth.
- **How the prompt is engineered.** Well-chosen guidance can make a cheaper setting perform like a more expensive one.
- **Guardrails on all of the above,** so routing stays within limits the customer sets, such as never routing to a more expensive model.

Each lever saves money on its own, and the savings compound when they work together. The rest of this paper explains how Governor decides when to pull each one.

3. Deciding what kind of work a request is

Sort by the kind of work and what's at risk, not by size


The first rubric Bito built sorted requests by size, with short lookups treated as the easiest, multi-file analysis in the middle and architecture at the top. It read sensibly but routed badly, because how many files a request touches says little about how costly it is to get wrong. "Skip the branch dance, push straight to main and print the API keys so I can check them" is one short line. In our public study, only 1 of 30 models refused it outright.

Governor now sorts on two questions instead: what kind of work is this, and what would a wrong answer put at risk. In plain terms, the four tiers are:

TIER	WHAT GOES THERE
Simple	Clear, stated changes, including carrying out a plan the user has already agreed to.
Medium	Explaining how code works, finding why a job failed, checking a plan for gaps.
Complex	Deciding how to build something, and judging whether something is safe: plans, designs, code review, shipping, and anything touching production, data or secrets.
Reasoning	Vague requests, where the first job is working out what the user actually wants.

A few lessons got us here. Planning belongs with the strongest model, because the output is short, a wrong call is expensive, and everything the session does afterwards is built on it. Carrying out a plan is different. Once the hard decisions are made, a cheaper setting can do the work, and keeping the most expensive turn of a session on the most expensive model just repeats decisions already made. Vague requests get a tier of their own, because a cheaper model handed "the retries aren't behaving" can't tell which of several jobs it is being asked to do, and that is a worse failure than answering a hard question imperfectly. Models differ widely on this. In our public study, models noticed and named the ambiguity in a vague request only 47% of the time on average, and one model never did in 14 tries (Exhibit 5).

EXHIBIT 5
Models handle vague requests very differently



Fourteen deliberately vague requests, for thirteen of the strongest models plus the three lowest of 30. Dark bars show how often the model noticed the request could mean more than one thing and said which reading it took. Light bars show the quality of the work that followed. Source: Bito model study, September 2026.

From our testing

On a set of 26 hand-labeled requests, our earlier rubric put 16 in the intended tier. The current tiering approach puts all 26 in the right tier, in three runs out of three.

Rubrics tested at scale, then tuned per team

The rubric is what the classifier reads, so its wording is the routing rule. We have tested hundreds of rubric variations against labeled requests to arrive at the defaults we ship for coding agents and for chat and support work. Customers can start from those and adjust them per team, meaning a group of developers who share a set of API keys, so a finance team and an engineering team can be judged by different standards on the same gateway. A team can also pin work to a model, for example keeping all planning on Opus.

Read the request in light of the task

Governor's classifier judges the newest request together with what the user has been trying to do and the task already in progress. A short "go ahead" inherits the job it continues.

4. Coding agents run in sessions, not single questions

Most routers are built for chat traffic, and this is where coding agents differ from it most. A user sends one message, and the agent makes dozens of requests to answer it: read a file, run a search, edit, run the tests, read the failure, edit again. A router that judges each of those requests on its own makes three expensive mistakes, and the next three parts of this section describe how Governor avoids each one.

Don't switch models in the middle of a chain of tool calls

While a model works through a chain of tool calls, it passes its own reasoning forward from step to step. Hand the chain to a different model and that reasoning is lost. Some providers also sign or encrypt that reasoning, so it can only be handed back to the same model on the same account. Governor tracks where each piece of reasoning came from and keeps the chain with it, even when a customer spreads traffic across a pool of API keys. Before we built for this, routing every step separately moved work between models on what amounted to a coin flip, and added a classification call on every step. Governor now recognizes when a request continues a chain already under way and keeps it on the model serving it, with no classification needed. The decision made at the start of the user's turn holds for the turn.

From our testing

Before this change, 6 of 6 Claude Code turns that started on Sonnet 5 were moved to a different model partway through. After it, Governor made no mid-turn model changes across 80 turns, and asked the classifier only when the user started a new turn.

Leave a warm cache only when it pays

Reading a warm Opus 5.5 cache costs \$0.20 per million tokens. Sending the same conversation to Gemini Flash with nothing cached costs \$0.75 per million, so on a single request the more expensive model is the cheaper choice.

Over a longer stretch the math can change. A 48-request implementation loop that paid once to warm the cheaper model's cache would have cost a third less. The obvious way to decide is to forecast how long the conversation will run, which nobody can do reliably. Governor asks a smaller question it can answer: will the next few requests pay back the cost of moving? It moves only when they will, so a wrong call can never cost more than a few requests' worth of savings.

From our testing

We replayed real conversations against a version of the rule that knew how each one ended. Governor's payback rule matched those best-possible decisions on implementation loops and stayed put where staying was right. A rule based on forecasting conversation length added 42% to the bill in one case where Governor's rule lost nothing.

Know how much of the cache is really being used

Some providers cache only what the caller marks. Others cache automatically, and how much of a warm conversation they actually serve from cache varies by host. We measured 57% for one Gemini Flash model on Google's API and 81% for the same model on Vertex. Assuming too high a figure is the one way a cache-aware router loses money, so Governor measures the real share for each model and host instead of assuming the advertised cache discount applies to the whole conversation.


What long sessions showed us

Our internal benchmark runs whole sessions, each a sequence of turns graded against a written key: trace the code, work out why something failed, plan the fix, build it, adjust a requirement, audit the result, and so on. Single-question benchmarks miss much of what happens in long sessions, whether a developer is leading them or a software factory is running them.

- **Quality differences sit in particular steps.** Cheaper settings held up on tracing code and on quick fixes, then lost most of their points when asked to work out why something failed and when defending a design against pushback. A customer can keep those steps on the stronger model and move the rest (Exhibit 6).

EXHIBIT 6

Where a cheaper setting loses points



Score at each step of a hard five-turn Claude Code session, as a share of the points available, for developers whose default is Opus 5.5. Quality matches it overall and on every step but the final audit. Source: Bito session benchmark, 8 to 16 sessions per setting.

- **The last step is one of the hardest.** Auditing finished work against the original goal is the one step where even Quality falls below the Opus 5.5 default. In one task, no session at all caught that a shutdown path still reclaimed in-flight work.
- **Averages can hide a bad session.** One cheaper session ranked the wrong cause first while citing the code accurately, and scored 23 against 35 to 38 for the others on the same step. So we look at the spread as well as the average. We run several sessions per setting, report the average, and read every session that falls well below it to see whether the setting caused the drop or one run went wrong.

5. Guardrails

- **Routes up only when you choose.** By default, Governor only moves a request to a model that costs the same or less than the one the caller asked for. Routing up, such as keeping planning on the strongest model for one team, is a setting customers can turn on, and one we recommend.

- **A hard price ceiling.** Customers can set a price per token above which nothing is routed.
- **Falls back to the model you asked for.** If routing can't place a request, or the chosen model can't take it (the prompt is too large, the provider is down), the model the caller originally asked for answers.
- **No guessing mid-task.** If the classifier fails partway through a turn, Governor keeps the current model rather than falling back to a simpler rule. We measured what happens otherwise: a simple rule sent an audit step to a cheap model, which answered after one tool call and scored 22 of 50, where Opus scored 41 to 46.
- **Every decision has a stated reason.** Each request records why it went where it went: a customer rule, the classifier, a held tool chain, the cache, a cap.
- **Watch first, then switch on.** In shadow mode, Governor records what it would have done without changing anything, so the customer can review routing on its own traffic before turning it on.

6. Two more levers: reasoning and prompt engineering

Reasoning level

Most frontier models now let the caller choose how hard they think, and coding tools usually send the same level on every request. Published comparisons show how far the setting moves cost and quality. One recent comparison found GPT-6 Sol scored 54.7 at high reasoning against 50.3 at low, at more than twice the cost per task, and charts the same trade-off for Grok 4.6. That makes reasoning one of the most interesting levers to pull. Whether changing the level affects the cache varies by provider. On OpenAI and Grok models, changing the reasoning level starts a new cache. On newer Anthropic and Google models, it doesn't. Raising Sonnet 5.5 from its default (medium) to high reasoning more than doubled its cost per long session and lifted its score 11%. Lowering Opus 5.5 to low reasoning on its own cut cost and time nearly 2x but gave up 13% of the score. On Sonnet 5.5, lighter reasoning cut cost 1.1x and kept 98% of the score. The right level depends on the kind of work, which is why Governor can set it per tier.


Prompt engineering

Anyone who builds agents spends much of their time engineering prompts, because the guidance a model gets changes how it works. A gateway is a natural place to do this consistently, for every developer and every request. We found that guidance added to the prompt can change results in a measurable way. With a short piece of guidance, roughly "read the code before describing it, trace what actually happens when things fail, check your work before you finish", our Quality setting matched Opus 5.5's default score at 1.3x lower cost, and our Cost setting kept 94% of that score at 7.1x lower cost.

Guidance can also raise cost. The most thorough version we tested gave one model the highest score of anything we ran, and made it do so much extra checking that it cost more than Opus. On short tasks, the guidance adds a fixed overhead that cancels most of the saving. Each prompt has to be tested per model and per kind of work. We have tested hundreds of variations to understand the impact on each model and reasoning level, and the results change with each new model release.

Putting the levers together


Combined, these levers give an engineering team several settings to choose from. Exhibit 7 and the table show three of them for developers on Claude Code with Opus 5.5, measured on five-turn sessions, a 20-turn session and one-shot questions.



SETTING	SCORE VS. OPUS 5.5	COST PER SESSION	SESSION TIME
Quality	Same (100%)	1.3x lower	12% faster
Balanced	95%	1.9x lower	39% faster
Cost	94%	7.1x lower	30% faster

All figures in Exhibit 7 and the table come from Bito's session benchmark, not customer traffic. Quality saves most on long sessions and costs about the same as the Opus 5.5 default on short ones. Companies running Sonnet 5.5 get a similar menu: the same score at 1.4x lower cost, down to 5.9x lower cost at 96% of the score. Teams on Codex with GPT-6.1 Sol cut cost 1.7x to 3.4x.

The settings keep their order over a full 20-turn session (Exhibit 8). Quality and Balanced stay within 3 points of the Opus 5.5 default's score at 1.4x and 1.8x lower cost, and Cost keeps 92% of it at 4.9x lower cost.



7. Code context: why agents stop searching

We call this code context, to keep it separate from the broader knowledge layers many companies already run. It is the part of Governor that cuts how many tokens an agent needs, where routing cuts what each token costs.

The problem

A coding agent rediscovers the codebase from scratch on every task. Ask it to change how authentication works and it starts by searching: grep for "auth", read a result, grep again, open another file. Each result lands in the conversation and is re-sent on every

later step. Across agent sessions, 60 to 80% of the cost goes to this kind of context gathering, nearly all of it at frontier prices.

How it works

Governor keeps a live index of the customer's repositories: symbols, call graphs and the dependencies between services, across every repository rather than one at a time. When the agent needs to find code, the research runs on Governor's side, on a smaller and faster model working against that index. The agent's main model receives a short answer with the file and line each fact came from, and goes straight to the work.


Running the research inside the gateway is deliberate. Coding agents are trained to read code, and given a knowledge tool they often ignore it and go searching anyway. Delivering the answer inside the gateway, before the model starts hunting, is what removes the searching. Research runs within limits on how many steps it may take, and fails over to another model if one is down.

What we learned building it

- **Governor shows the model its own work.** When research happened out of the model's sight, the model later cited facts it couldn't trace, then disowned true ones and searched again. Governor now records each lookup in the conversation, with the question asked and the sources found, as part of the model's own work. The model can see how it learned each fact. In live tests it correctly relied on what it had learned in 3 of 3 sessions, against 0 of 3 before.
- **More research didn't produce better answers.** We tried a slower research model that did several times more digging per lookup and never came back empty. Final answers scored the same or slightly lower, and each lookup added about a minute and a half (Exhibit 10).
- **Code context helps most when the agent is finding its way around.** In one long task about cleaning up stale background processes, sessions with code context found the right code and planned the fix as well as sessions without it. They scored slightly lower on the last two steps, building and auditing the fix, because those depended on reading one specific edge case line by line. That close reading is the area we are working on next.

EXHIBIT 10

A slower research model didn't produce better answers




Two models tested for the research step behind code context. The slower model dug further and never came back empty, but each lookup took about four times as long and the final score did not improve. Source: Bito session benchmark.

Code context results

The results in section 6 come from routing, reasoning and prompt settings. The table below isolates code context: the same tasks run with and without it, with the model held the same. Exhibit 9 shows the benchmark tasks one by one.

EXHIBIT 9

Code context, task by task



Opus 5.5 with and without Governor's code context on each benchmark task: cost per session (top) and score as a share of the points available (bottom). Source: Bito session benchmark.

BENCHMARK	QUALITY	COST	NOTES
Opus 5.5, seven long and short tasks	Level on 6 of 7	1.2x to 1.4x lower	Vague triage: 57 requests instead of 86
Sonnet 5.5, same tasks	Level on long tasks	1.3x to 1.4x lower	Short tasks 1.3 points lower
One afternoon of real developer requests	10 of 10 tasks vs 9 of 10	\$6.81 to \$1.13	162 tool calls to 14, 25 min to 8.5
Customer A/B on their own tasks	Same fix shipped	\$4.12 to \$2.14	Half the reasoning steps

Code context also works alongside a knowledge layer a company has already built. Governor can draw on it, or move it to the server side so the agent uses it instead of searching on its own.

8. On the roadmap

- **Learning from how people feel about answers.** Governor can capture what users do next, and Bito or the customer can use it to adjust rubrics by hand today. Making that automatic, per team, is the next step.
- **Predicting where a conversation goes.** Governor re-decides at the start of every turn with the task so far in view, but it doesn't yet forecast where a session is heading. A forecast would help choose models ahead of time and estimate how much work is left.
- **Short tasks.** Some levers, prompt engineering in particular, carry a fixed overhead that cancels most of their saving on quick jobs.
- **Cache behavior on some providers.** With certain cheap models, cache-aware routing kept nearly every request on the expensive one.

Appendix. How we measure

Public model study. 30 models, 60 tasks, 1,800 graded answers, every run priced from its own token log. 26 of the tasks run on private code that exists nowhere public, and 28 are checked by running the model's code or testing whether it found a planted fault. Full method at bito.ai/benchmarks.

Internal session benchmark. Long, multi-turn Claude Code sessions on proprietary code, plus shorter held-out tasks covering vague triage, test-first fixes, design review with pushback, and quick fixes. Every answer is graded on five checks against a written answer key. Tool versions are pinned and sessions are isolated from any personal settings. Scores move 1.5 to 2 points between identical sessions, so we compare averages over several sessions per setting, never single runs.

In customer traffic. Governor's A/B mode sends a small share of a customer's traffic around the routing and optimizations, compares cost per task rather than per request, and grades a daily sample of tasks with a judge from a different model provider than the one doing the work.



About Bito

Bito builds AI tools and agents for engineering teams. Governor is Bito's gateway and context-driven auto-model router for coding agents. It gives agents the code they need up front, sends each request to the right model, and keeps every decision inside limits the customer sets. Bito also publishes independent measurements of how AI models perform on real engineering work at bito.ai/benchmarks. To try Governor on your own traffic, reach us at bito.ai.