

AI Coding Exercise

Build this module's OpsMonitor sample with ChatGPT or Claude from a ready-made prompt that points the model at this module's specification, then review, run and extend what it gives you — including against our own reference build on GitHub.



The assistant is the fast pair of hands. You are still the architect.

An assistant can produce this module's sample in minutes. What it cannot do is decide whether the result is the right design for your application — that call stays with you. You supply the judgement: brief it properly, then check what comes back against what Module 7 just taught you.

What you are asking for

Deliver the WisejPerfLab capstone. Derive a capacity model from your own measurements: memory retained per idle session, CPU per common scenario, expected peak concurrency and the headroom you are keeping. Write a HealthCheck.json whose maxSessions, maxMemory and maxCPU come from those numbers and explain each threshold, returning 503 with a retry-after so a load balancer routes new users elsewhere while existing sessions keep running. Write the deployment notes for two instances behind a load balancer with session affinity and WebSocket support, and assemble the final report: scenarios, environment, baseline evidence, root causes, changes, after evidence, remaining risks and the production counters that will show if a regression returns.

The module asks for these deliverables, and the prompt below asks the assistant for exactly them:

- ✓ Capacity model deriving sessions per server from measured per-session memory and per-scenario CPU
- ✓ HealthCheck.json with maxSessions, maxMemory, maxCPU, return code and retry-after, each threshold justified
- ✓ Deployment notes covering session affinity, WebSocket support and the polling fallback
- ✓ Final performance report with baseline, root cause, change and after-evidence for every module fix
- ✓ Monitoring plan naming the counters, thresholds and alerts that would catch each regression

Why it is shaped this way

Six modules of profiling have produced a faster WisejPerfLab and a folder of traces. None of that answers the question the operations team will ask on the day of the rollout: how many users fit on a server, and what happens when one is full? Profiling on a developer machine tells you the cost of a scenario; production needs that cost converted into a capacity model, thresholds that refuse new sessions before the process is in trouble, a deployment shape that respects server-side session state, and monitoring that

will notice if any of the six fixes regresses. That translation is this module, and the capstone report is what makes the work reviewable by someone who never opened Visual Studio.

These are the terms the assistant has to get right. If what comes back uses them loosely, it has not understood the module:

Term	What it means here
capacity model	Memory retained per idle session and CPU per common scenario, multiplied by expected peak concurrency and divided into the instance's resources, with headroom kept deliberately.
HealthCheck.json	Configuration through which a Wisej.NET server can report itself unavailable based on session count, memory and CPU, so a load balancer stops sending it new users.
503 with retry-after	The response an unhealthy instance returns to a first request; existing sessions keep running on their node while new users are routed elsewhere.
session affinity	Consistent routing of a client to the instance that holds its session state, usually implemented with sticky cookies; without it a Wisej.NET session cannot survive a second request.
WebSocket-aware balancing	Load-balancer configuration that allows the upgraded connection through; when WebSocket is unavailable the app falls back to polling and real-time push is lost.
production counters	Lightweight ongoing metrics — CPU, GC heap, allocation rate, exception rate, thread-pool health — that show a regression without a profiler attached.
final performance report	Scenarios, environment, baseline evidence, root cause, change, after evidence, remaining risks, thresholds and the monitoring plan, in one reviewable document.

Build it — pick your assistant

Both buttons open a new tab with the **same ready-made prompt** already filled in. The prompt points the model at this module's specification — a single Markdown file holding the lesson, the lab brief and the deliverables — so it works from the real material instead of guessing what Module 7 is about. You can open [module-07.md](#) yourself to see exactly what it will read.



Use ChatGPT

Opens [chatgpt.com](#) with this module's prompt loaded. Works with GPT-5 or Codex; any plan that can browse the web will fetch the spec.



Use Claude

Opens [claude.ai](#) with the same prompt loaded. Claude Code users can paste it into a terminal in an empty folder instead.



Send it as-is and you get a project, not a snippet

The prompt ends by asking for the whole solution as a single `.zip` with build-and-run steps. Download it, open the solution in Visual Studio, press F5 and OpsMonitor comes up on `localhost`. You need the .NET SDK and the `Wisej-4` NuGet package; everything else the prompt asks the assistant to write itself.

The prompt, in full — read it, then edit it if you want a different sample

I am working through Module 7 – "Scale, Health Checks, Load Balancing and Final Tuning" – of the LearnWisej Performance & Profiling course, and I want to build the module sample with you.

Start by reading the module specification. It is a single Markdown file holding the full lesson, the lab brief and the deliverables:

<https://learnwisej.com/assets/courses/performance-and-profiling/specs/module-07.md>

Then build a complete Wisej.NET 4 project I can open in Visual Studio:

1. Project "OpsMonitor", C#, net8.0 or later, NuGet package Wisej-4 4.1.x, Kestrel host calling `app.UseWisej()`.
2. Build what the specification's lab asks for: Deliver the WisejPerfLab capstone. Derive a capacity model from your own measurements: memory retained per idle session, CPU per common scenario, expected peak concurrency and the headroom you are keeping. Write a `HealthCheck.json` whose `maxSessions`, `maxMemory` and `maxCPU` come from those numbers and explain each threshold, returning 503 with a `retry-after` so a load balancer routes new users elsewhere while existing sessions keep running. Write the deployment notes for two instances behind a load balancer with session affinity and WebSocket support, and assemble the final report: scenarios, environment, baseline evidence, root causes, changes, after evidence, remaining risks and the production counters that will show if a regression returns.
3. Produce these deliverables, and tell me where each one lives in the solution:
 - Capacity model deriving sessions per server from measured per-session memory and per-scenario CPU
 - `HealthCheck.json` with `maxSessions`, `maxMemory`, `maxCPU`, return code and `retry-after`, each threshold justified
 - Deployment notes covering session affinity, WebSocket support and the polling fallback
 - Final performance report with baseline, root cause, change and after-evidence for every module fix
 - Monitoring plan naming the counters, thresholds and alerts that would catch each regression
4. Follow the specification wherever it names a Wisej.NET class, property or event – use exactly that one. Where it leaves a choice open, tell me which API you picked and why.
5. Keep the server authoritative: state and validation live in .NET, and the browser renders and reports. Never move a decision into JavaScript because it is easier there.
6. Make the result demonstrable: I should be able to run it and see the behaviour this module is about, including what happens when something goes wrong.

Before you write any code, restate in your own words what you are going to build and where each piece will live, so I can correct the design while it is still cheap.

When the code is done, package the whole solution as a single .zip I can download, and give me the exact steps to build and run it.

Finally, answer these questions from the specification, using your own implementation as the evidence:

- Your health check refuses new sessions at 150. A colleague asks why not 300, since the process has not crashed at 300 in testing. Answer with your own numbers.
- A deployment goes out behind a load balancer without session affinity. Describe exactly what a user experiences, and why this is a correctness failure rather than a slowdown.
- Six months from now the dashboard refresh is slow again. Which single production signal would tell you first, and what would you collect next?

Ask me about anything that is ambiguous before you start.

Then review what came back

Code that compiles is not yet code that is right. The interesting questions are design questions: is this the Wisej.NET control the job actually calls for, does it follow the lifecycle the framework gives you, is the authoritative state on the server rather than parked in the browser because that was convenient? Read the result as a reviewer, not as a customer.

This module gives you the checklist. Answer each one against the code you just received:

- ✓ Your health check refuses new sessions at 150. A colleague asks why not 300, since the process has not crashed at 300 in testing. Answer with your own numbers.
- ✓ A deployment goes out behind a load balancer without session affinity. Describe exactly what a user experiences, and why this is a correctness failure rather than a slowdown.
- ✓ Six months from now the dashboard refresh is slow again. Which single production signal would tell you first, and what would you collect next?

Anything you had to correct is worth writing down. A short note of the conventions and decisions you settled on, pasted into the first message of your next session, is worth more than any single clever prompt.



We already built this one — compare yours with ours

The Module 7 reference implementation lives in the public samples repository. Build yours first, then open ours: the differences are the interesting part.



[Open the Module 7 sample on GitHub](#)

Or clone the whole set with `git clone https://github.com/iceteagroup/LearnWisej-Samples` and open `Performance and Profiling Course/Module 7`.



Our sample is a starting point, not a finished product

It was built the same way you are building yours — by prompt, then corrected — and it deliberately stops where the module stops. Expect stand-ins where production code would go, thin error handling and no tests. Treat that as an invitation: use what you learned in this module, plus the assistant, to push it further. Improving code you did not write is the skill this course is actually teaching.

Three prompts to take it further

Point your assistant at the project you just built — or at ours and try these. Each one is a review-and-extend exercise, not a rewrite.

- ✓ **Prove it actually works.** *"Go through this project deliverable by deliverable against the module specification at <https://learnwisej.com/assets/courses/performance-and-profiling/specs/module-07.md> and show me the evidence that each one is genuinely done — the code path, what to click, and what I should see. Where a deliverable is only partly met, say so plainly instead of arguing it through. Then add whatever is missing."*
- ✓ **Review it like a pull request.** *"Walk me through this OpsMonitor code as if you were reviewing a colleague's pull request and did not write it. Where does state live, what happens end to end on one user action, what would you send back for changes, and which parts would you not be comfortable putting in front of a customer? Be specific and quote the lines you mean — do not summarise the design back to me."*
- ✓ **Swap a piece without moving the seams.** *"Pick the part of this OpsMonitor project that is most obviously a stand-in for something real — a fake data source, an in-memory store, a simplified service — and replace it with a production-shaped implementation. The contract either side of it must not change: same types, same calls, same behaviour under failure. Tell me what the real implementation forced you to handle that the stand-in did not."*



Carry it into the lab

The hands-on lab that closes this module asks for the same deliverables you just had the assistant produce. Use what you have here as the shape of your answer and the evidence behind it — not as the answer itself. The lab is graded on your reasoning, and that part has to be yours.