The information in this document is governed by ITQ's information classification policy @ ITQ 2026
For more information regarding security management, please contact security@itq.eu



Red Hat Summit Connect Utrecht

11 November 2026

The Pit Wall on OpenShift AI

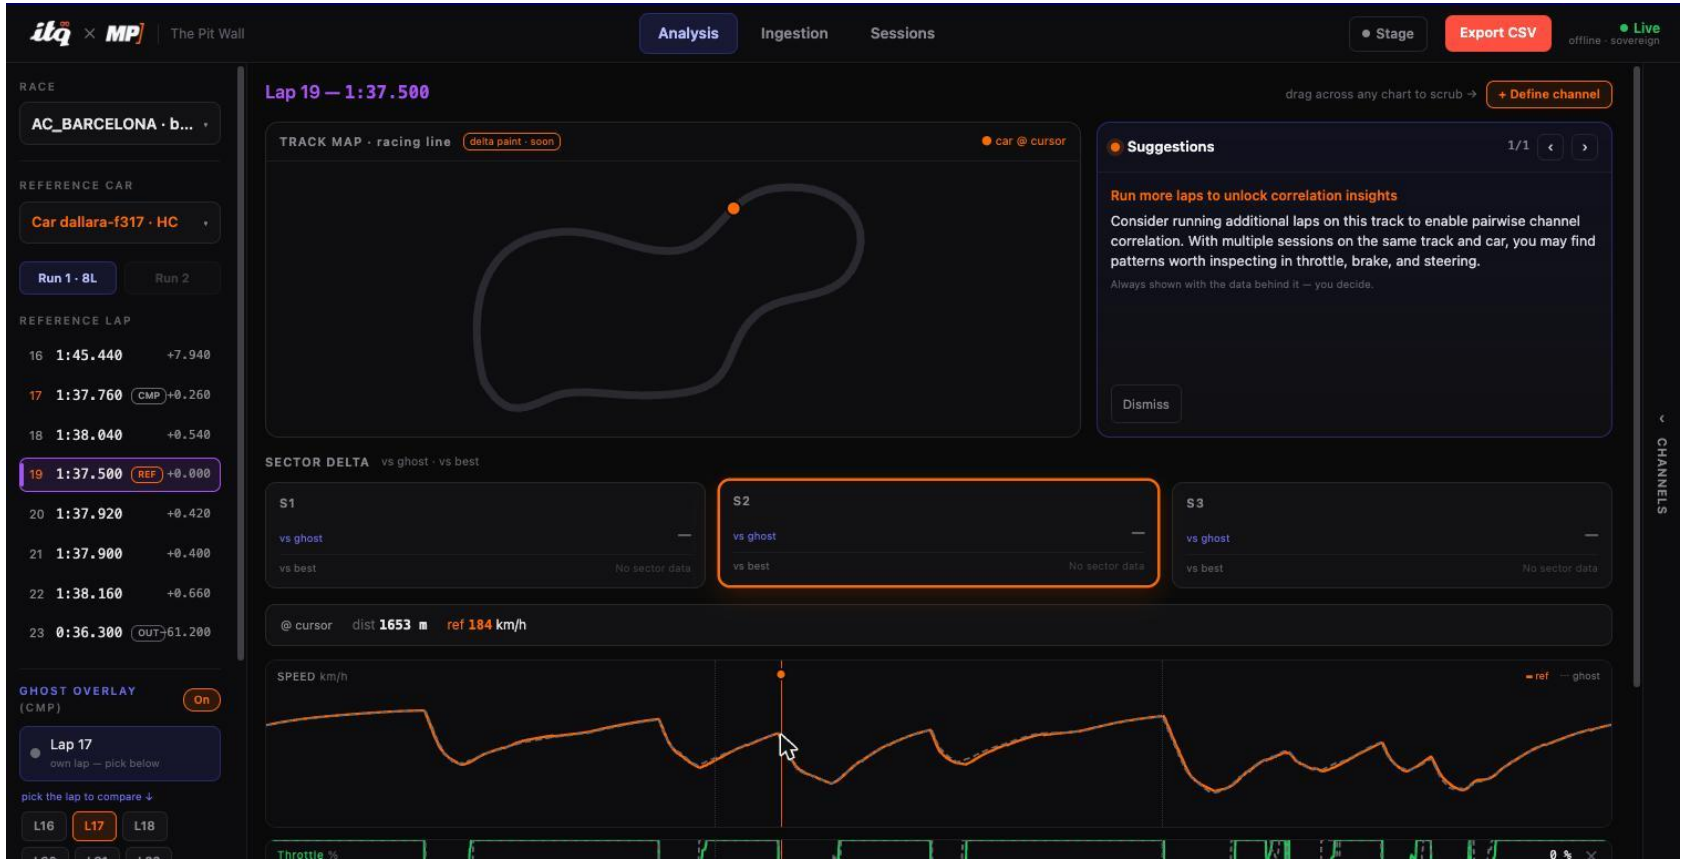
The information in this document is governed by ITQ's information classification policy @ ITQ 2026
For more information regarding security management, please contact security@itq.eu

Decisions, at the speed of racing

- Race engineering app for MP Motorsport
  - Built by the ITQ team
- One view instead of 25 spreadsheets
- Drop the race files: telemetry, laps, features
- Lap times and telemetry in a dashboard
- A coaching bar with AI hints
- Ran on a Dell GB10 with a local LLM

The information in this document is governed by ITQ's information classification policy @ ITQ 2026
For more information regarding security management, please contact security@itq.eu

Demo: the Pit Wall



The information in this document is governed by ITQ's information classification policy @ ITQ 2026
For more information regarding security management, please contact security@itq.eu

How a lap becomes a hint

Hints, never causes. The engineers know things the data does not, so the app suggests and never instructs.



From file to finding

Drop the session file. Telemetry becomes laps and features. The session is scored with the model for this circuit and this car.



From finding to sentence

The strongest link becomes one sentence, written by an LLM in about a second. A phrase gate blocks "causes" and orders; otherwise a template.

The information in this document is governed by ITQ's information classification policy @ ITQ 2026
For more information regarding security management, please contact security@itq.eu

The correlation engine

- Barcelona, Dallara F317: rear-left tyre pressure is the strongest link
  - r = −0.87, 95% CI −0.96 to −0.63, 58 laps, 5 sessions
- A channel is one column of data: speed, braking, air in each tyre
- Tyres wear, every lap gets slower
- Per setting: when it goes up, does the wear speed up?
- A confidence interval says how sure we are
- One model per circuit and per car
- Linked to, never caused by

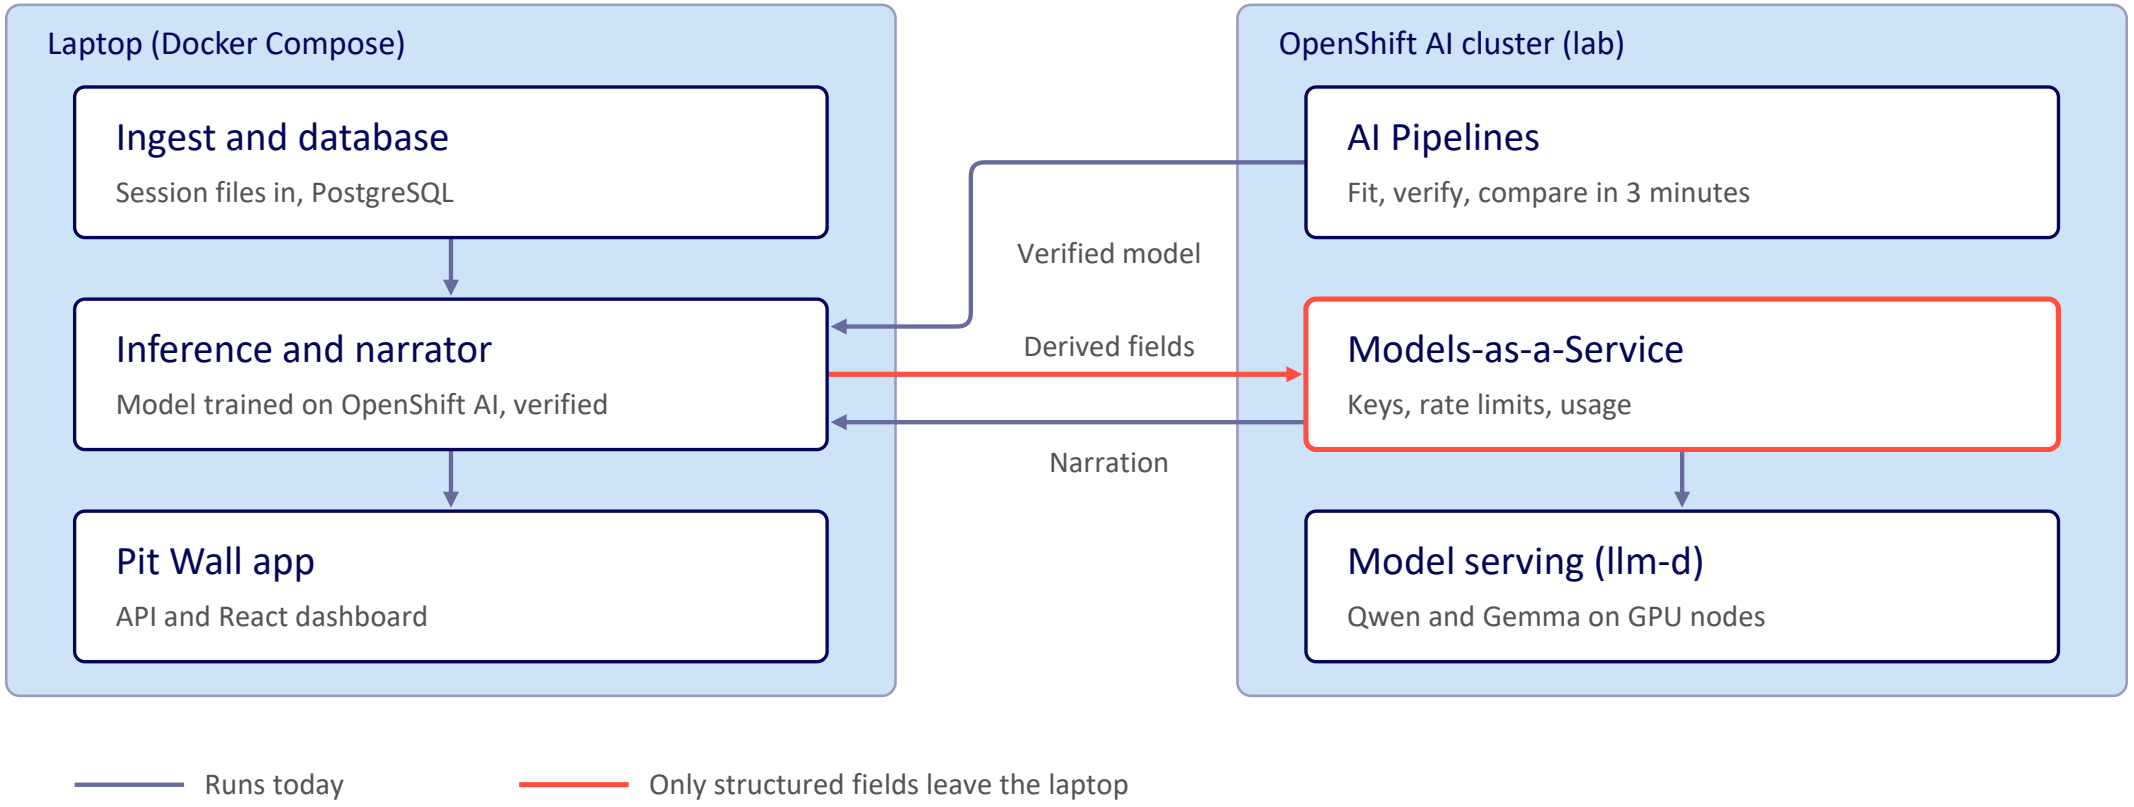
The information in this document is governed by ITQ's information classification policy @ ITQ 2026
For more information regarding security management, please contact security@itq.eu

Part two: the platform

From Dell GB10 to OpenShift AI production setup

The information in this document is governed by ITQ's information classification policy @ ITQ 2026
For more information regarding security management, please contact security@itq.eu

What runs where



The information in this document is governed by ITQ's information classification policy @ ITQ 2026
For more information regarding security management, please contact security@itq.eu

Demo: the pipeline on OpenShift AI

- OpenShift AI dashboard: create a run of pitwall-correlation-fit
  - The run graph: fit, verify, compare turning green
  - Fit log: the winning channel and r
  - Verify log: verify passed
  - Compare report, rendered in the dashboard
- Pit Wall, Train tab: a replay of that run, about 15 seconds
- Pit Wall, sessions: a Dallara session scored by the cluster-trained model
  - The coaching sentence, written live by the LLM on MaaS

The information in this document is governed by ITQ's information classification policy @ ITQ 2026
For more information regarding security management, please contact security@itq.eu

The pipeline: fit, verify, compare

Step	What it does	Time	Result
Fit	Runs the app's own fitting code on public sim laps	60 s	Rear-left tyre pressure, r = −0.87
Verify	The app's startup gate: numbers must reproduce	20 s	Verify passed, else the run stops
Compare	New model against the one in the app	20 s	Same winner, difference in r 0.01
Whole run	7 pods, 1 CPU and 2 GiB per step, no GPU	3 min	4 runs, identical model every time

The information in this document is governed by ITQ's information classification policy @ ITQ 2026
For more information regarding security management, please contact security@itq.eu

Models-as-a-Service: one platform trains and serves

- Qwen 3.5 35B, MoE, about 3B active, served on the cluster
- OpenAI-compatible endpoint
- API key per app, rate limits, usage ledger
- About 1 to 1.5 s per coaching sentence
- Endpoint down? The app falls back to template sentences

The information in this document is governed by ITQ's information classification policy @ ITQ 2026
For more information regarding security management, please contact security@itq.eu

What the platform gives

The hackathon build was custom code. On OpenShift AI, governance comes with the platform.

Governed and sovereign

API keys, rate limits and usage per team. A namespace boundary and offline installs. Only public sim data leaves the laptop.

Repeatable and verified

Same inputs, same model, every run. Verified by the app's own gate before use. Compared with the model already in production.

The information in this document is governed by ITQ's information classification policy @ ITQ 2026
For more information regarding security management, please contact security@itq.eu

Thank you

Questions?



ITQ in Belgium
Tiensestraat 111
3000 Leuven
Belgium
T: +31 251 82 88 03

ITQ in Luxembourg
Rue Pafebruch 89E
8308 Mamer
Luxembourg
T: +31 251 82 88 03

ITQ in France
La Défense
92042 Paris
France
T: +31 251 82 88 03

ITQ in Beverwijk (NL)
Parallelweg 94
1948 NM Beverwijk
The Netherlands
T: +31 251 82 88 03

ITQ in Amersfoort (NL)
Databankweg 26 - 32
3821 AL Amersfoort
The Netherlands
T: +31 251 82 88 03

ITQ in Germany
Wanheimer Straße 43
40472 Düsseldorf
Germany
T: +49 211 514 321 98

ITQ in Sweden
Vasagatan 38
111 20 Stockholm
Sweden
T: +31 251 82 88 03

ITQ in Norway
Hagaløkkveien 26
1383 Asker
Norway
T: +31 251 82 88 03

ITQ in Denmark
H. C. Andersens Boulevard 38
1553 København
Denmark
T: +31 251 82 88 03

Feel free to contact me for more information!

Tom Heyerick

theyerick@itq.eu