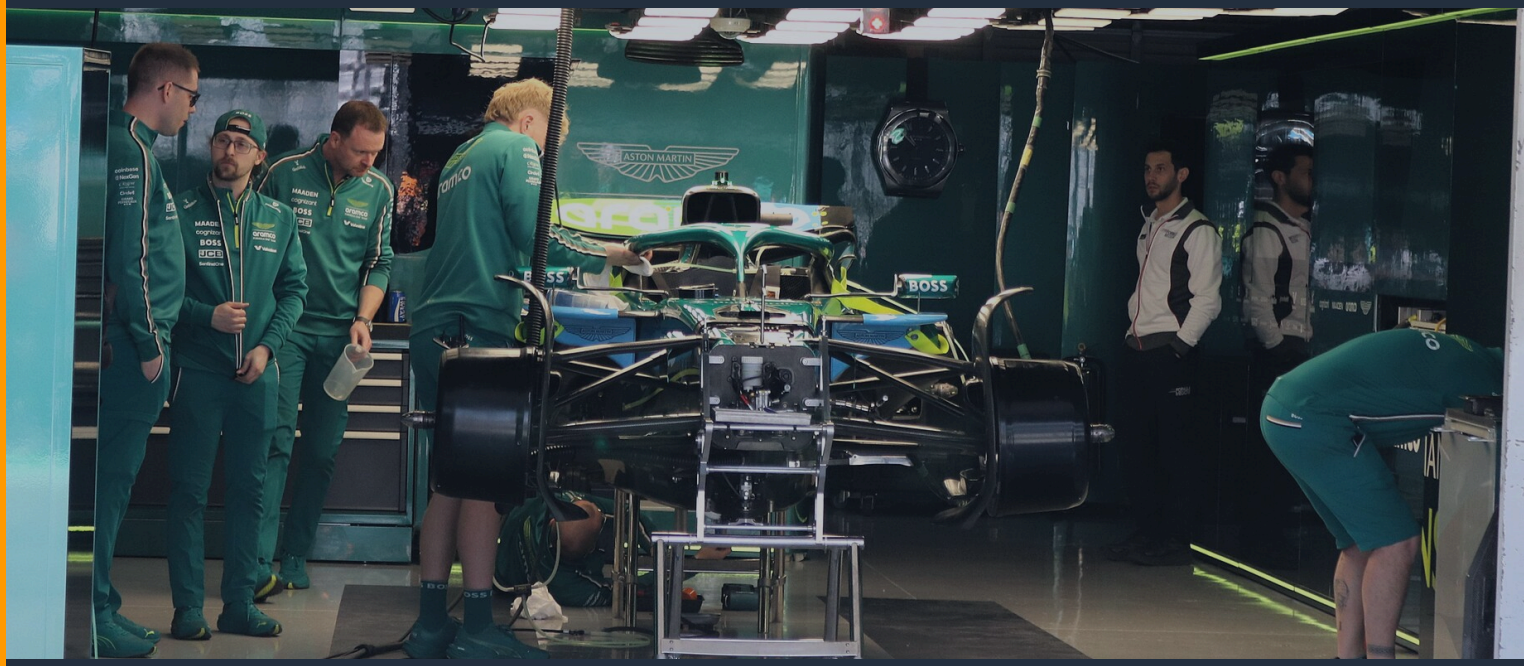
CLOUD INNOVATION CHALLENGE · LMU MUNICH · 1ST EDITION 2026




PROPOSAL TO THE VELOCITY RACING EXECUTIVE TEAM

# The fastest decision loop on the grid

A cloud strategy that turns 300+ sensors per car into race-day decisions, protects the design IP that wins the development race, and converts the same data into fan and commercial growth.

## Faster

Track-edge decisions in milliseconds, regional strategy in seconds, design iteration in hours.

## Safer

Design IP isolated by account, key and enclave. Fans never touch engineering data.

## Elastic

Pay for the ~20% of the year that is race weekends, not for peak capacity all year.

AUTHOR AND CONTACT

Duc-Anh Nguyen from Team 2 · anh.nguyen1@campus.lmu.de

INTERACTIVE DEMO

ducanhvalentinonguyen.com/velocity-racing-demo

DATE

October 2026

**Credits.** Cover photo: "2025 Japan GP – Aston Martin – AMR25 – Thursday" by Liauzh, via Wikimedia Commons, licensed CC BY-SA 4.0 (creativecommons.org/licenses/by-sa/4.0); cropped and darkened. Illustrative only; it does not show Velocity Racing. AWS logo: Amazon Web Services, Inc. via Wikimedia Commons.
Amazon Web Services, AWS and the AWS logo are trademarks of Amazon.com, Inc. or its affiliates. Formula 1 and team names belong to their owners. This document was written by LMU students for the AWS Cloud Innovation Challenge and is not an official AWS publication. Velocity Racing is a fictional team; figures marked "assumption" are planning hypotheses, not AWS quotes or commitments.

EXECUTIVE SUMMARY

# Win the decision, protect the design, grow the brand.

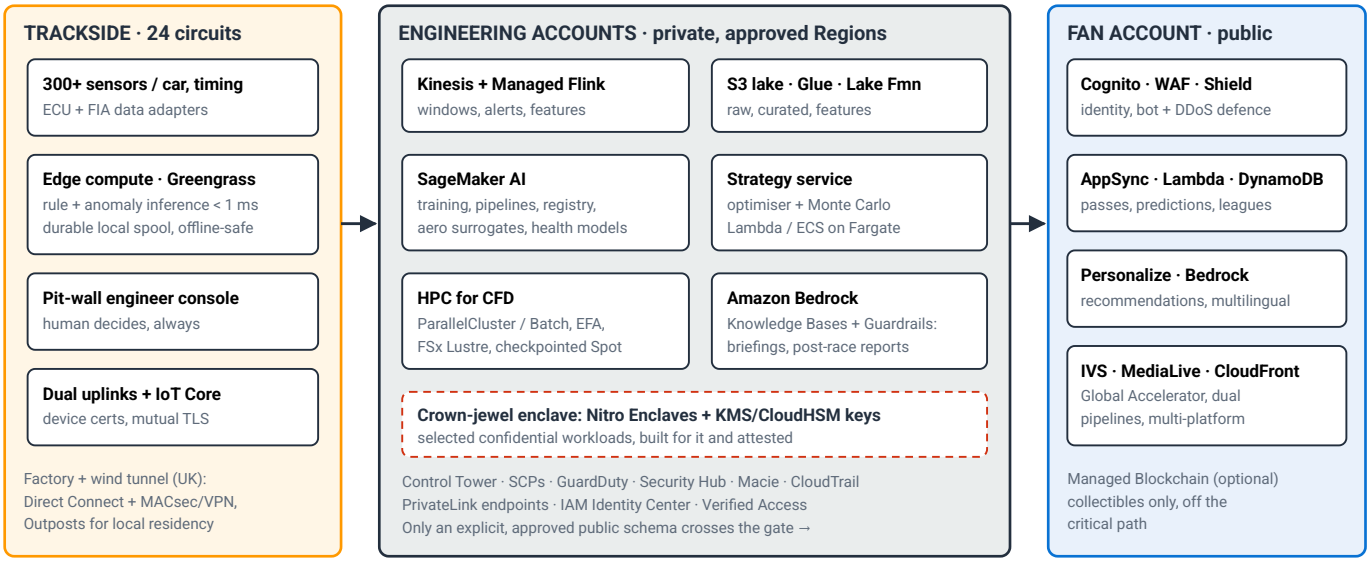
A 2026 championship push is decided by three loops: the **race-day loop** (sensor to pit call), the **development loop** (idea to validated aero upgrade under tight testing limits) and the **audience loop** (fan attention to sponsor value). Each loop is limited today by data plumbing, compute peaks and trust in the data. We propose one AWS platform with three hard boundaries, so each loop can run at full speed without endangering the others. **The innovation:** F1 already runs AWS-powered strategy insight for broadcasts [17][18]; we bring that pattern inside the garage, on Velocity's private telemetry, with an evidence-checked AI layer and a break-even business case.

**< 1 ms**
target for a defined local sensor-to-rule path, trackside on AWS IoT Greengrass. Not a cloud round trip.

**< 2 s**
proposed p95 for a regional strategy re-plan (pit window, tyre, energy), engineer-approved.

**~20%**
of the year is race weekend (24 × 72 h = 1,728 h of 8,760 h). Elastic AWS capacity matches that shape.

## Reference architecture: three boundaries, one data foundation




Design rule: a failure or traffic surge in one boundary cannot degrade another. The pit wall keeps working with the WAN down; fans can be a hundred times busier on race day without touching engineering capacity.

## Why AWS, for this team, this season

| Reason | What it means for Velocity |
|---|---|
| **Proven in Formula 1** | AWS is Formula 1's official cloud and machine-learning provider. F1 Insights already turns 300+ sensors per car (over 1.1 M data points per second) into Battle Forecast, Predicted Pit Stop Strategy, Alternative Strategy and, from the 2026 Hungarian GP, a Strategy Insight graphic, on SageMaker and S3 [17][18]. Track Pulse runs a serverless broadcast stack on Fargate, DynamoDB, AppSync and CloudFront [1][21]. F1's CFD work tripled throughput and halved turnaround per run [2]; Ferrari reports 4× data-processing capacity and about 50% faster insight in power-unit assembly on AWS [3]. In F1's words: "AWS outperforms all other cloud providers, in speed, scalability, reliability, global reach…" (Pete Samara, F1) [2]. |
| **Global reach, clear responsibility** | 24 races on five continents: ingest to the nearest AWS Region, Direct Connect to London, CloudFront's global edge for fans, Multi-AZ for regional services. Under the shared responsibility model AWS secures the cloud; Velocity secures what runs in it (identity, data, configuration), so every control in this report has an owner. |
| **One platform, five needs** | Edge, streaming, ML, HPC, security, media and engagement services share IAM, encryption and networking. No stitching of five vendors, and one audit trail. |
| **Elastic economics** | Peak compute for CFD and race weekends is rented, not owned. Idle capacity between events is switched off. |
| **Security and residency controls** | Region choice, per-data-class keys, private connectivity, enclave isolation and a separate AWS European Sovereign Cloud option [4], so data location is your decision. |

PILLARS 1 AND 5 · STRATEGY, TELEMETRY AND DESIGN

# From 300 sensors to a pit call the engineer trusts.

Telemetry and strategy are one pipeline, so we treat them together. The principle: **time-critical work runs at the track, learning and heavy compute run in the cloud, and every AI output arrives with its evidence.**

## 1 Three latency tiers, each with an honest promise

| Tier | Where it runs | Target | What it does |
|---|---|---|---|
| **Track edge** | Ruggedised trackside servers running AWS IoT Greengrass | < 1 ms p99 for a defined sensor-to-rule path | Threshold and small-model checks (oil pressure, coolant, brakes, battery), driver-coaching deltas, local console; works with no uplink. |
| **Stream** | IoT Core → Kinesis Data Streams → Managed Service for Apache Flink (Region) | ≈ 1 s end-to-end | Event-time windows across both cars, fuel/energy and tyre-wear features, live dashboards (Timestream for InfluxDB or Flink + Pinot [20]; Managed Grafana). |
| **Strategy** | Strategy service on Lambda / ECS Fargate, SageMaker endpoints | < 2 s p95 (proposed) | Re-plans pit window, compound and energy deployment over thousands of simulated race futures; ranks options with P10-P90 ranges. |

Sub-millisecond is claimed only for the local path on benchmarked hardware with synchronised clocks. Greengrass manages deployment and rollback but is not a real-time guarantee, so the pilot proves the number on Velocity's hardware.

## 2 Volume: sized, not assumed

The brief's 1.5 TB per weekend is ambiguous, so we size for the harder reading, **1.5 TB per car**: 3 TB over 72 h averages **11.6 MB/s**; a 10× burst is **116 MB/s**. Bulk raw data goes straight to S3 in bounded objects; only selected live features go through Kinesis (pre-scaled or on-demand mode, partitioned by session, car and sensor group to avoid hot keys [6]). A 15-minute uplink outage at burst rate needs about 105 GB of local spool.

## 3 Predictive component health: a model, not a hunch

- **Train** on historical sessions in SageMaker: gearbox, ERS/battery, brake and cooling systems. Use temporal hold-out, not random splits.
- **Output** a calibrated risk and time-to-threshold, not a bare yes/no.
- **Gate** on measured false alerts and missed events, run in **shadow mode** for several sessions before any engineer sees an alert.
- **Deploy** to the edge through SageMaker Model Registry and staged Greengrass rollout, with automatic rollback.
- **Monitor** drift with SageMaker Model Monitor; retrain between sessions.
- **Boundary**: the model advises. It never commands the car.

## 4 Generative AI that engineers can audit

- **Strategy copilot.** An engineer asks, "What if the safety car comes at lap 31?" The optimiser produces the ranking; **Amazon Bedrock** explains it in plain language, citing the telemetry and runs it used. Every number comes from the optimiser, briefings are checked against its output, and unsupported statements are rejected or sent for engineer review. The hard part is turning millions of simulations into clear Plan A/B/C calls [19]; this layer does that.
- **Post-race analysis.** Bedrock Knowledge Bases over the S3 lake and past debriefs generate a draft race report within minutes of the flag, with links to evidence; Bundesliga uses the same pattern of automated, editor-reviewed narratives [9]. Bedrock Guardrails reduce the risk of unsupported output but do not guarantee factuality, so reports stay drafts until an engineer signs them off.
- **Engineer search.** Amazon Q Business (or Bedrock) answers questions over setup notes, logs and regulations inside the private account.

## 5 Accelerating aero development, inside the rules

CFD is where the 2026 regulations will be won. The loop: **generate** candidate geometries → **screen** with a SageMaker surrogate that predicts downforce, drag and uncertainty → **run the most informative** designs through licensed CFD on **AWS ParallelCluster or Batch** (EFA networking, FSx for Lustre, checkpointed Spot with On-Demand fallback) → **feed results back** to retrain the surrogate (active learning). AWS-published F1 work reports that a scalable CFD workflow on AWS tripled throughput and halved turnaround per run, with SageMaker-based optimisation as the next step [2]. Every job is metered against the FIA aerodynamic testing allowance before it is scheduled, so speed never becomes a compliance breach.

> **Why this beats a bigger on-premises cluster:** Velocity pays for thousands of cores only during a design sprint; the same S3 lake feeds the strategy models.

PILLARS 2 AND 4 · SOVEREIGN SECURITY AND GLOBAL DELIVERY

# Assume a competitor is inside the wire. Design so it doesn't matter.

Design IP and race strategy are the team's most valuable assets, so we use **zero trust and isolation by default**: separate accounts and keys, no standing access, and enclaves for selected crown-jewel workloads.

## 1 Layered protection of IP and strategy

| Layer | AWS services | Threat addressed |
|---|---|---|
| **Account isolation** | AWS Organizations, Control Tower, service control policies. Separate engineering, fan, security and log-archive accounts | Blast radius; a fan-side breach cannot reach the design lake |
| **Identity** | IAM Identity Center with phishing-resistant MFA, short-lived roles, Verified Access, Verified Permissions | Stolen credentials, insider misuse, contractor access |
| **Data** | KMS keys per data class (CloudHSM or external key store for crown jewels), S3 Object Lock, Macie | Exfiltration, tampering, silent copies |
| **Compute** | **AWS Nitro Enclaves** for selected confidential workloads (e.g. strategy-model inference on encrypted setup data), with attestation | Even a privileged operator cannot read data in use there |
| **Network** | PrivateLink, Direct Connect with MACsec or VPN, Network Firewall; Shield Advanced and WAF on public surfaces | Interception, DDoS, application attacks on race day |
| **Detection** | GuardDuty, Security Hub, Detective, Config; an organisation CloudTrail trail in S3 Object Lock, queried with Athena [16]; on-call runbooks | Slow discovery; missing forensic trail |

## 2 Secure, resilient links: factory, tunnel, 24 circuits

The UK factory and wind tunnel connect over diverse **Direct Connect** paths with MACsec or VPN encryption (not encrypted by default [5]). Each circuit has two independently tested carriers, IoT Core device certificates and a local spool that reconciles after an outage. Coverage is verified *per venue*; Outposts and Wavelength are not assumed. FIA timing and data sharing use a read-only adapter into the trackside process and a logged outbound export; no FIA feed passes through the fan account.

## 3 Sovereignty as an explicit decision

Data location follows written policy: UK engineering data in **London (eu-west-2)**, with Outposts at the factory if residency demands it; EU fan data in an EU Region. Replication, support access and Bedrock routing are limited to approved Regions. The **AWS European Sovereign Cloud** (generally available since January 2026 [4]) is an option where its catalogue and contract fit. Hosting on AWS is a strong control set, not a compliance certificate: GDPR, FIA data-sharing and retention rules are confirmed with counsel and the FIA in the pilot.

PILLAR 4 · GLOBAL BROADCASTING

## 4 One media plane, two latency classes

**Interactive and mass audience: Amazon IVS**

- **Low-latency channels** (typically under 3 s [7]) for millions of concurrent viewers: onboard cameras, team-radio clips, pit-wall views.
- **Real-time stages** (under 300 ms [7]) for interactive watch-alongs and driver Q&A.
- Player SDKs for iOS, Android and web; TV and headset playback validated per platform.

**Packaged and redundant: Elemental + CloudFront**

- **MediaLive** with dual pipelines, **MediaPackage** and **CloudFront** for multi-bitrate, DRM-protected delivery.
- **Global Accelerator** and redundant contribution paths from the circuit; hot backup ingest, failover rehearsed before each season.

**Real-time multilingual commentary:** commentary audio flows through **Transcribe** → **Translate** → captions (optionally **Polly** or Bedrock-drafted summaries) beside the video path, so language generation never delays the picture. Synthetic voices need rights and disclosure.

> **Sizing reality check:** 2 million concurrent viewers at an assumed 5 Mb/s is 10 Tb/s, roughly 9 PB over two hours. That is a CDN capacity and commitment question, not a compute one, so it is planned first with AWS's edge and quotas teams. **Rights:** the world feed belongs to the commercial rights holder; Velocity's own content is published under agreed delay and approval rules.

PILLAR 3 · FAN ENGAGEMENT, THEN THE BUSINESS CASE

# Turn approved insight into fans, sponsors and a measured payback.

## 1 A fan platform that scales on its own

- **Digital paddock pass.** Cognito identity; signed, short-lived QR tokens; scanners perform an **atomic conditional write in DynamoDB**, so a second scan of the same pass is rejected. One authoritative redemption store avoids duplicate entry from replication lag.
- **Anti-fraud.** Velocity signals (scan velocity, device and geography) in Kinesis and SageMaker anomaly detection, with WAF Bot Control at the edge.
- **Collectibles and membership.** Optional, on **Amazon Managed Blockchain** [8], storing only public metadata on-chain. Personal data, access rights and IP stay off-chain. A collectible is a collectible, not an identity proof or a ticket.
- **Personalised experience.** Amazon Personalize on consented preference and behaviour data; driver, language and content picks; AppSync subscriptions push live updates.
- **Gamification.** Prediction games and fantasy leagues on Lambda + DynamoDB + ElastiCache, with server-side scoring, idempotent submissions and versioned rules. A leaderboard spike is absorbed by autoscaling, not pre-purchased hardware.
- **AR overlays.** Licensed telemetry synchronised to video timecode, delivered through the approved public event schema only.

The gate between engineering and fans is a strict public event schema plus editorial approval. Nothing reaches fans by direct database access. Track Pulse already uses this serverless pattern in F1 [1][21].

## 2 Business case: pays back on a small fraction of one position

Racing value is asymmetric: one championship position is worth far more than the platform costs. We therefore frame the case as a **break-even** rather than a savings claim. All values are planning **assumptions**; the pilot replaces them with Velocity's own data.

| Year-one view (EUR) | No sporting effect | Conservative | Base | Upside |
|---|---|---|---|---|
| Programme cost: AWS usage 0.47 M + team 0.72 M (Appendix D) | 1.2 M | 1.2 M | 1.2 M | 1.2 M |
| Expected championship positions gained (≈ 9 M each [10][11]) | 0 | 0.15 → 1.35 M | 0.3 → 2.7 M | 0.6 → 5.4 M |
| Commercial: fan/sponsor uplift (≈ 0.7% / 2% of ≈ 72 M mid-field sponsorship [14]) | 0 | 0 | 0.5 M | 1.5 M |
| Engineering hours redeployed (20 h × 24 races × EUR 90) | 0.04 M | 0.04 M | 0.04 M | 0.04 M |
| **Net value / benefit-cost ratio** | **≈ −1.1 M / 0.04×** | **≈ +0.2 M / 1.2×** | **≈ +2.1 M / 2.7×** | **≈ +5.8 M / 5.8×** |

Fractional positions are expected values: 0.3 means, for example, a 30% chance of finishing one place higher. Break-even needs about 0.13 of a position (1.2 M ÷ 9 M). Adjacent constructors' positions are publicly estimated at USD 8-12 M a year [10][11]; we use EUR 9 M (assumed 0.90 EUR/USD). With no sporting effect the full programme would lose about 1.1 M, which is why it is only funded after the pilot gate: stopping there caps the loss at the pilot budget. No avoided DNFs are assumed. Every assumption has a pilot check (Appendix D).

## 3 A gated plan, not a leap of faith

| Phase | When | Scope | Gate to proceed |
|---|---|---|---|
| **0 · Pilot** | Weeks 1-6 | Landing zone and security baseline; replay one season of historic telemetry into the lake; shadow strategy advice and component-health model; measure edge latency on real hardware. | Acceptance criteria in Appendix F met; spend within the EUR 100k ceiling |
| **1 · Race ops** | Months 2-4 | Live trackside deployment at 3 circuits; Bedrock briefings and post-race reports; DR rehearsal. | RTO / RPO targets met in a failure drill |
| **2 · Development** | Months 3-6 | CFD on AWS with surrogate screening, FIA allowance metering. | Design-iteration time cut vs baseline |
| **3 · Fans and media** | Months 5-9 | Paddock pass, personalisation, games, IVS streaming; blockchain collectibles last. | Load test at 3× forecast peak |

**Our ask:** approve a six-week pilot with a **budget ceiling of EUR 100k** (team ≈ EUR 83k, AWS usage capped at EUR 15k with budgets and alerts), a named engineering owner and a security owner. You fund the full programme only if the acceptance criteria in Appendix F are met.

## APPENDIX A · REQUIREMENT TRACEABILITY

| Brief requirement | AWS answer | Page |
|---|---|---|
| AI race strategy, pit, tyre, fuel | Optimiser + Monte Carlo on Lambda/Fargate; SageMaker models; Bedrock explanations | 2 |
| Generative AI for aero and CFD | SageMaker surrogates, active learning, ParallelCluster/Batch, FSx Lustre, EFA | 2 |
| Post-race analysis, natural-language reports | Bedrock Knowledge Bases + Guardrails; Amazon Q Business | 2 |
| Scalable AI/ML training | SageMaker (training jobs, Pipelines, Spot); GPU quotas agreed ahead of the season | 2 |
| Protect IP; espionage; FIA/GDPR | Control Tower, KMS/CloudHSM, Nitro Enclaves, PrivateLink, Macie, GuardDuty; Region policy | 3 |
| FIA data sharing and timing integration | Read-only timing adapter; logged, controlled outbound export; counsel sign-off | 3 |
| Secure factory, tunnel, 24 circuits | Direct Connect + MACsec/VPN, IoT Core, dual carriers, local spool | 3 |
| Digital ticketing and anti-fraud | Cognito, DynamoDB atomic redemption, WAF Bot Control | 4 |
| Blockchain collectibles / NFTs | Amazon Managed Blockchain, metadata only on-chain | 4 |
| Personalisation and gamification | Personalize, AppSync, Lambda, ElastiCache | 4 |
| Low-latency multi-platform streaming | IVS, MediaLive, MediaPackage, CloudFront, Global Accelerator | 3 |
| Multilingual commentary | Transcribe, Translate, Polly, Bedrock | 3 |
| 300+ sensors, 1.5 TB/weekend, failure prediction | Greengrass, IoT Core, Kinesis, Flink, S3, SageMaker | 2 |
| Trackside edge, sub-ms processing | Greengrass on benchmarked trackside hardware; Outposts at factory | 1, 2 |
| Race-day computational spikes | Autoscaling, Spot for checkpointable HPC, on-demand streams; separate fan account | 1, 4 |

## APPENDIX B · KEY ASSUMPTIONS

- 1.5 TB read as *per car*, 2 cars, 72 h weekend, decimal units.
- 10× burst = 116 MB/s; 1 KB records ≈ 116,000 records/s before batching.
- Video: 2 M viewers × 5 Mb/s = 10 Tb/s; ≈ 9 PB per 2 h.
- Position value: USD 8-12 M per adjacent position (public estimates), used as EUR 9 M. Sponsorship base: mid-field teams about USD 50-110 M a year [14]. Cost model and the plan to firm up each number: Appendix D.
- 24 races = 1,728 race-weekend hours of 8,760 per year (≈ 20%).
- Latency targets are proposed acceptance criteria, not delivered service levels.

## APPENDIX C · RISKS AND MITIGATIONS

- **Model accuracy.** Shadow mode, temporal validation, agreed false-alert budget.
- **Connectivity.** Local-first design, dual carriers, spool and replay.
- **Regulation.** FIA testing allowance metering; counsel sign-off on GDPR and data sharing.
- **Media rights.** Confirm with rights holder before any live-stream work.
- **Quotas and capacity.** Agree SageMaker GPU, Kinesis and media quotas early; capacity reservations for peak weekends.
- **Cost drift.** Budgets and alerts, tagging, weekly review; credits are not a hard spending cap.

**Interactive demo:** [ducanhvalentinonguyen.com/velocity-racing-demo](ducanhvalentinonguyen.com/velocity-racing-demo)
A browser-based demo on synthetic data with three pages. **Race control**: replayed telemetry, a thermal-anomaly rule, a 500-trial pit-strategy comparison with P10-P90 ranges, an evidence-cited briefing and engineer acknowledgement (page 2). **Fan experience**: single-use paddock-pass redemption (page 4). **ML predictions**: a component early-warning classifier and a tyre-degradation model, trained in the browser and scored on held-out races against the Appendix F criteria (pages 2 and Appendix F). The data are generated, so results are optimistic. The demo is an illustrative model, not a validated F1 simulation or failure predictor, and not a live AWS deployment.

## APPENDIX D · YEAR-ONE COST MODEL (LIST PRICES, BEFORE DISCOUNTS)

| Line | Basis | USD |
|---|---|---|
| SageMaker training | 1,200 h of ml.p4d.24xlarge at $23.42/h [12] | 28,104 |
| SageMaker inference | 2 × ml.g5.12xlarge, 1,728 race-weekend h at $10.18/h; 5 dev notebooks × 1,000 h (ml.m5.xlarge, $0.23/h) | 36,332 |
| Kinesis | 10% of 3 TB/weekend live (7.2 TB in, 3 consumers), on-demand $0.08/GB in, $0.04/GB out, 10 streams × $0.04/h [12] | 4,944 |
| Streaming (phase 3) | 20k avg viewers × 2 h × 12 races = 480k viewer-hours at IVS $0.072/h HD ($34,560), plus input: 4 channels × 4 h × 12 races at $2/h ($384) [12] | 34,944 |
| CFD compute | *Planning assumption:* 5 M core-hours at $0.05. FIA CFD usage is calculated separately from Velocity's declared hardware, because Allocation Unit hours scale with core count, solver time and clock frequency [15] | 250,000 |
| Bedrock | Briefings and reports; token volume is small [12] | 10,000 |
| Lake, Flink, network, security, support | *Allowance*, not priced line by line (S3, Flink, Direct Connect, GuardDuty, support) | 160,000 |
| **AWS subtotal** | ≈ EUR 0.47 M at 0.90 EUR/USD. Team: 6 FTE × EUR 120k loaded = EUR 0.72 M (Munich ML median EUR 97k [13] × 1.25 on-costs) | **524,324** |

Mass streaming is priced per viewer-hour, not as a lump sum: about USD 0.05-0.07 on IVS low-latency, or about USD 0.05-0.19 through CloudFront at 5 Mb/s (2.25 GB per hour at $0.02-0.085/GB [12]). Two million viewers for two hours across 24 races would cost several million dollars a year, so audience scale must be paid for by sponsorship, subscription or rights revenue. Prices are us-east-1 list prices from AWS pages; Regions, discounts and commitments change them.

## APPENDIX D (CONT.) · HOW WE WILL FIRM UP EACH NUMBER

| Number | Today (basis) | How we firm it up | When |
|---|---|---|---|
| AWS usage, EUR 0.47 M | Published list prices [12] | AWS Pricing Calculator on the pilot workload; quota check; account-team quote with commitment discounts | Pilot wk 1-2 |
| CFD compute, USD 250k | Planning assumption, 5 M core-hours | Translate Velocity's FIA CFD allocation into core-hours using its declared hardware and the FIA formula; benchmark its solver on HPC instances; price a Spot/On-Demand mix | Pilot wk 1-4 |
| Team, EUR 0.72 M | Salary benchmarks [13] × 1.25 | Use Velocity's internal rates and partner day rates; agree the FTE plan with engineering leads | Pilot wk 1-2 |
| Position value, EUR 9 M | Public estimates [10][11] | Velocity finance supplies its own prize-money model and points-to-position history | Pilot wk 1-2 |
| Positions gained | Scenario labels, 0.15 / 0.3 / 0.6 | Back-test strategy advice on past races; measure modelled seconds gained and decision-preparation time; convert to expected points with the team | Pilot wk 3-4 |
| Fan and sponsor uplift | 0.7% / 2% of ≈ EUR 72 M [14] | Agree sponsor KPIs with the commercial team; measure engagement in a phase-3 test against a control group | Phase 3 |
| Latency targets | Proposed acceptance criteria | Measure on trackside hardware with synchronised clocks; report p99 against the threshold | Pilot wk 3-4 |

## APPENDIX E · SOURCES


1. AWS, “Real-time storytelling: the AWS architecture behind Formula 1 Track Pulse” · aws.amazon.com/blogs/media/real-time-storytelling-the-aws-architecture-behind-formula-1-track-pulse/
2. AWS, “Optimize F1 aerodynamic geometries via design of experiments and machine learning” · aws.amazon.com/blogs/machine-learning/optimize-f1-aerodynamic-geometries-via-design-of-experiments-and-machine-learning/ ; and Formula 1 case study · aws.amazon.com/solutions/case-studies/formula-one-case-study/
3. AWS, “How AWS supports Scuderia Ferrari HP to optimize Formula 1 power unit assembly” · aws.amazon.com/blogs/industries/how-aws-supports-scuderia-ferrari-hp-to-optimize-formula-1-power-unit-assembly-process/
4. AWS, “Opening the AWS European Sovereign Cloud” · aws.amazon.com/blogs/aws/opening-the-aws-european-sovereign-cloud/
5. AWS docs, Direct Connect, encryption in transit · docs.aws.amazon.com/directconnect/latest/UserGuide/encryption-in-transit.html
6. AWS docs, Kinesis Data Streams capacity · docs.aws.amazon.com/streams/latest/dev/working-with-streams.html
7. AWS docs, Amazon IVS glossary (latency classes) · docs.aws.amazon.com/ivs/latest/ChatUserGuide/ivs-glossary.html
8. AWS docs, Amazon Managed Blockchain endpoints and quotas · docs.aws.amazon.com/general/latest/gr/managedblockchain.html
9. AWS, Bundesliga case study (editorial workflow efficiency) · aws.amazon.com/solutions/case-studies/bundesliga-case-study/
10. PaddockIntel, “F1 Prize Money 2026: how much each team actually earns” (pool and per-position estimates) · paddockintel.com/melbourne-2026-f1-prize-money-constructors-championship/
11. PlanetF1, “F1 2025 prize money payouts” (champion share about 14%, USD 147 M) · planetf1.com/news/f1-2025-prize-money-payouts-revealed-mclaren-crowned-champions-as-ferrari-steal-the-spoils
12. AWS pricing pages (accessed 30 Sep 2026): aws.amazon.com/sagemaker/ai/pricing/ · aws.amazon.com/kinesis/data-streams/pricing/ · aws.amazon.com/ivs/pricing/ · aws.amazon.com/cloudfront/pricing/pay-as-you-go/ · aws.amazon.com/bedrock/pricing/
13. Salary benchmarks, Munich: Levels.fyi machine learning engineer median EUR 97,211; ERI SalaryExpert (ML EUR 104.6k, cloud EUR 97.6k); Glassdoor cloud engineer median EUR 75k · levels.fyi/t/software-engineer/title/machine-learning-engineer/locations/munich-metro-region (accessed 1 Oct 2026)
14. Mid-field sponsorship estimates: Sportico, “F1 teams combine for $2B in sponsorship revenue in 2024”; Sportcal partner pages (Alpine over USD 87 M, Haas about USD 50 M, 2023); Williams estimated USD 70-110 M · sportico.com/business/sponsorship/2025/f1-team-sponsor-deals-nfl-report-1234854316/ (estimates, accessed 1 Oct 2026)
15. FIA 2026 Formula 1 Regulations, Section F (Operational): CFD limit in MAUh per Aerodynamic Testing Period, Allocation Unit hours from cores, solver time and clock frequency (F4.4.4), scaled by championship position (about 70% of baseline for the champion, 115% for the last team) · fia.com/system/files/documents/fia_2026_f1_regulations_-_section_f_operational_-_iss_05_-_2025-12-10_1.pdf
16. AWS docs, CloudTrail Lake availability change (closed to new customers from 31 May 2026) · docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-lake-service-availability-change.html
17. About Amazon, “F1 AI insights” (Battle Forecast, Predicted Pit Stop Strategy, Alternative Strategy; 300+ sensors, 1.1 M data points per second; SageMaker, S3, Direct Connect) · aboutamazon.com/news/aws/f1-ai-insights-grand-prix
18. Formula 1, “F1 to launch new Strategy Insight graphic powered by AWS at the Hungarian Grand Prix” · formula1.com/en/latest/article/f1-to-launch-new-strategy-insight-graphic-powered-by-aws-at-hungarian-grand-prix.3h211Yy2IaW04biHGRY4EM
19. AWS Executive Insights podcast, “AI and data strategy in F1 with Ruth Buscombe” · aws.amazon.com/executive-insights/podcast/ai-and-data-strategy-in-f1-with-ruth-buscombe/
20. AWS Builder Center, “Building a real-time F1 telemetry dashboard with Apache Pinot and Apache Flink on AWS” · builder.aws.com/content/2xUz2tAg0rAdAzJN8RZCetkydgD/
21. AWS, “F1 revs up race-day broadcasts with real-time data storytelling” (Track Pulse) · aws.amazon.com/blogs/media/f1-revs-up-race-day-broadcasts-with-real-time-data-storytelling/


## APPENDIX F · PILOT ACCEPTANCE CRITERIA (PROPOSED; THRESHOLDS AGREED WITH THE CTO IN WEEK 1)

| Area | Criterion |
|---|---|
| Edge latency | p99 < 1 ms on the defined sensor-to-rule path during a one-hour replay at burst rate, on trackside hardware with synchronised clocks |
| Strategy service | p95 < 2 s for bounded re-plan requests at 10× expected concurrency; numerical ranking always delivered before any explanation |
| Component health (shadow) | On a held-out season: flags at least 80% of recorded failures before the failure lap, with at most one false critical alert per race weekend |
| Briefings | 100% of numbers in sampled briefings match the optimiser output; unsupported statements are rejected or routed to engineer review |
| Resilience | No telemetry loss across a simulated 15-minute uplink outage (spool and replay); duplicate events removed |
| Security | Landing zone live; no open critical or high findings in Security Hub at week 6; data stays in approved Regions |
| Cost | Pilot within EUR 100k (AWS usage ≤ EUR 15k); Velocity-specific cost model replaces the Appendix D assumptions |