

Case Study – Sharyat Solutions LLP

Generative AI on AWS for Live Race Engagement and AI-Powered Advertising

1. About Client

Sharyat Solutions LLP, headquartered in Pune, Maharashtra, operates the Sharyat digital platform for bailgada (bullock cart racing), a traditional rural sport in Maharashtra. The platform supports race organizers, cattle owners, fans and advertisers through event creation, race registration and token allocation, live token status, results and scoreboards, owner profiles, race history and community content.

Sharyat also generates revenue through in-app advertising, including home-screen and location-based ad placements. Its users are predominantly Marathi-speaking and mobile-first, with race days generating high levels of activity.

2. Business Need

On race days, users had to repeatedly check the application or contact organizers to know the current token, their upcoming turn and the latest results. Results were primarily static, with limited real-time engagement and no AI-generated commentary or proactive notifications.

At the same time, local advertisers often lacked design and Marathi copywriting capabilities, limiting Sharyat's ability to scale its advertising business. Advertising content also needed screening for inappropriate content and material that could indicate animal mistreatment.

- Provide real-time race scores and token status.
- Generate short AI-powered race commentary in Marathi and English.
- Generate event recaps and notify owners when their turn is near.
- Enable advertisers to create Marathi and English ad copy and banner creatives.
- Apply automated text and image moderation before advertising content is published.
- Ensure AI-generated commentary is grounded in verified race data.
- Maintain responsible AI, security, monitoring and cost-governance controls.
- Keep the GenAI workload within an estimated monthly AWS consumption envelope.

3. Proposed Solution

Converg-IN designed and implemented two GenAI capabilities on AWS using Amazon Bedrock and Amazon Nova models. The solution uses serverless and managed AWS services to provide scalable, event-driven processing while keeping the workload cost-efficient.

Use Case 1 – Live Score & AI Race Commentary

- Race results and token information are captured through the race-desk application.
- Amazon API Gateway provides REST and WebSocket interfaces, with AWS WAF and Amazon Cognito for protection and authentication.
- AWS Lambda validates incoming results and stores the verified race data in Amazon DynamoDB.
- DynamoDB Streams triggers the live update and commentary workflow when new results arrive.
- Amazon Bedrock with Amazon Nova Micro generates short Marathi/English commentary using verified race data and a defined style guide.

- Bedrock Guardrails checks generated commentary and recaps for harmful or inappropriate content.
- Generated content is validated against the source race record; a template message is used if validation fails. API Gateway WebSocket pushes live scores, token status and commentary to connected application users.
- Amazon SNS sends mobile notifications when an owner's turn is approaching.

Use Case 2 – AI Ad Studio

- Advertisers submit a business, offer, location and language brief through the application.
- AWS Step Functions orchestrates the advertising workflow.
- Amazon Nova Lite generates Marathi and English headlines, body copy and calls to action.
- Amazon Nova Canvas generates advertising banner creatives.
- Amazon Bedrock Guardrails moderates generated text and Amazon Rekognition checks generated images.
- Creative assets are stored in Amazon S3 as pending content.
- A Sharyat administrator reviews and approves or rejects the creative before publication.
- Approved advertisements are delivered through Amazon CloudFront to application ad slots.

Responsible AI, Security and Cost Optimization


- AI commentary is generated only from verified race information.
- Generated names, times and positions are checked against the source data.
- Bedrock Guardrails applies content filters, denied topics and word filters.
- Human approval is required before AI-generated advertisements are published.
- AWS KMS provides encryption at rest and TLS protects data in transit.
- IAM follows least-privilege access, with Amazon Cognito and AWS Secrets Manager supporting access and credential management.
- CloudTrail and Bedrock model invocation logging provide auditability and GenAI traceability.
- CloudWatch provides logs, metrics, dashboards and alarms.
- AWS Budgets, Cost Explorer and Cost Anomaly Detection support cost governance.
- Model selection considers quality, Marathi language capability, latency, token consumption, request volume and cost.

4. AWS Services Used

Category	AWS Services
Generative AI	Amazon Bedrock – Amazon Nova Micro, Nova Lite and Nova Canvas; Amazon Bedrock Guardrails
AI Services	Amazon Rekognition
Compute	AWS Lambda, Amazon EC2
Application Integration	Amazon API Gateway, AWS Step Functions, Amazon DynamoDB Streams, Amazon SNS
Database	Amazon DynamoDB, Amazon RDS for PostgreSQL
Storage & Delivery	Amazon S3, Amazon CloudFront
Networking	Amazon VPC, Application Load Balancer, NAT Gateway, Amazon Route 53
Security & Identity	AWS IAM, Amazon Cognito, AWS KMS, AWS Secrets Manager, AWS WAF, Amazon GuardDuty
Monitoring & Governance	Amazon CloudWatch, AWS CloudTrail, AWS Budgets, AWS Cost Explorer, AWS Cost Anomaly Detection

Backup & Recovery	AWS Backup, DynamoDB point-in-time recovery, Amazon RDS automated backups
Infrastructure as Code	AWS CloudFormation

5. AWS Deployment Architecture



6. Business Benefits

- **50–70% reduction** in repetitive race-status enquiries through real-time scores, token updates and notifications.
- **20–30% increase** in race-day user engagement through live AI commentary, updates and event recaps.
- Improved accessibility for Marathi-speaking users through Marathi AI-generated commentary and advertising content.
- **100% automated screening** of AI-generated advertisements, with **100% human approval** before publishing.
- Improved factual reliability through grounding AI commentary in verified race data and validating generated output.
- Cost-efficient GenAI adoption through serverless services, model selection and token controls.
- Improved operational visibility through CloudWatch, CloudTrail and Bedrock model invocation logging.
- Enhanced security through least-privilege IAM, encryption, secrets management, WAF and threat detection.

7. About Converg-IN

Converg-IN Private Limited is an AWS Advanced Consulting Partner that delivers cost-effective cloud and Generative AI solutions built on AWS Well-Architected principles. Our team of 15+ certified cloud architects covers the full cloud lifecycle: consulting and cost optimisation, migration (including SAP and VMware), security, modernisation, DevOps, and 24x7 managed services.

Our Modernisation & GenAI practice helps customers move Generative AI from experimentation to production using Amazon Bedrock and the wider AWS AI/ML stack. We have delivered GenAI workloads for customers in sports technology, education, life sciences analytics, and AI-driven contact centres. We design every solution to be simple to deploy, secure by default, and cost-efficient to operate.

We deliver with an automation-first approach, using infrastructure-as-code to stand up environments quickly and repeatably. Security and governance are central to every engagement. After go-live, our proactive operations team monitors performance, health, and cost, so customers' GenAI workloads stay reliable and optimised. With a proven track record across industry verticals, Converg-IN is a trusted partner for organisations adopting Generative AI on AWS.