

AI DEVELOPER BOOTCAMP

Artificial intelligence has moved decisively from the laboratory into everyday operations. Governments and enterprises everywhere are adopting AI, but adoption and capability are not the same thing. A critical distinction has emerged between organisations that consume AI - buying finished tools and hoping for the best - and those that can develop AI: building tailored solutions that fit their own data, processes, languages and legal environment.

The consumption trap

Most institutions today sit firmly in the consumer category. They subscribe to generic AI products, obtain inconsistent results, and remain uncertain how to navigate the intersection of AI with privacy law and other emerging regulation. In the language of the AI Academy, they are “leaving value on the table.” They cannot adapt these tools to the specifics of Mauritian governance, they cannot integrate them with their own systems, and they cannot assure themselves that sensitive citizen data is being handled appropriately. Every new requirement becomes a new procurement.

The scarcity of builders

The reason so few organisations escape the consumption trap is a scarcity of people who can build. Fewer than five percent of developers in the region possess production-ready AI skills, and conventional computer-science curricula typically lag industry practice by several years. The tools, patterns and platforms that define modern AI development - retrieval-augmented generation, vector databases, agentic workflows, model fine-tuning, cost-aware deployment - are largely absent from formal education. The result is a workforce that can use a chatbot but cannot construct one grounded in an organisation's own knowledge.

Why this matters for government

For a public institution, the bespoke-development gap is more than an efficiency problem. Reliance on external vendors for core AI capability creates strategic dependency, recurring cost, and exposure of sensitive data to third parties. It weakens the state's ability to make informed “build versus buy” decisions and to define its own AI roadmap. In a resource-constrained environment, the inability to build cost-effective, locally-controlled solutions directly limits what government can afford to do with AI at all. The skill to develop AI in-house is, therefore, a matter of both economic sustainability and digital sovereignty.

2. The Value of the Training

The AI Developer Bootcamp exists to convert AI consumers into AI builders. It does this by concentrating on the practical, portfolio-producing skills that let a technical professional design, build and deploy working AI applications. The value to the sponsoring agency accrues on several fronts.

Strategic autonomy and reduced vendor dependency

Participants learn to build solutions in-house rather than commission them. This shifts the balance of power in every future AI decision: leadership can weigh a genuine build option against a buy option, negotiate from a position of understanding, and avoid lock-in to a single supplier.

Data sovereignty and compliance by design

Because participants build the systems themselves, they control where data lives and how it is processed. The programme embeds responsible-AI and data-protection thinking throughout, so that compliance with Mauritian and emerging privacy legislation is designed in from the outset rather than retrofitted.

Economic sustainability

A recurring theme of the Bootcamp is cost-optimisation: choosing the right model for each task, caching to avoid needless spend, and favouring open and self-hosted options where they make sense. In a resource-constrained public sector, the ability to build AI that is affordable to run is as important as the ability to build it at all.

Immediate, demonstrable productivity

The Bootcamp is built around working artifacts. Participants leave not with theory but with software they have personally constructed - a grounded chatbot and an automated workflow - that can be adapted to real departmental problems. Hours of manual knowledge work become minutes of AI-assisted process.

A foundation for a home-grown AI workforce

Finally, the five-day Bootcamp seeds capability. It gives government a cohort with a shared vocabulary and hands-on experience, and a clear onward pathway into the full nine-to-twelve-month AI Developer Bootcamp for those who wish to specialise. It is the first, deliberate step in building an enduring national capacity to develop AI rather than merely buy it.

3. The Five-Day AI Developer Bootcamp

The five-day Bootcamp is an intensive, hands-on programme that carries a mixed-skill cohort of technical staff from foundational concepts to working AI software over a single week. It is a purposeful condensation of the AI Academy's full AI Developer Bootcamp, prioritising the applied skills of building AI applications and workflows. Each day follows a consistent rhythm - concept in the morning, construction in the afternoon - and closes with a tangible, working result.

Format	Five consecutive days; instructor-led, intensive, hands-on.
Focus	Applied AI application building - from first API call to a deployed, grounded AI app and an automated AI workflow.
Audience	Technical staff of mixed ability. Some programming exposure is helpful, but the programme is scaffolded so that varied backgrounds can keep pace.
Approach	Concept in the morning, build in the afternoon. Paired and small-group work with facilitator support throughout.
Result	Each participant leaves with working software they built themselves and a personal AI action plan.
Pathway	A direct on-ramp to the full nine-to-twelve-month AI Developer Bootcamp for those who wish to specialise.

4. Five-Day Curriculum

The daily plan below shows the arc of the week: a shared foundation on Day 1, a deep and practical grasp of how AI models are directed on Day 2, a complete working application on Day 3, automation and agents on Day 4, and the discipline to make it safe, affordable and deployable on Day 5.

Day 1 - Foundations - A Working Baseline for Everyone		
Objective: Level the field across a mixed-skill cohort and get every participant making live calls to an AI model by the end of the day.		
Session	Topic	Focus
Morning	The Modern AI Landscape	What generative AI is and is not; narrow vs. general vs. generative AI; the AI development lifecycle; the Mauritian and African AI landscape, opportunities and constraints.
Morning	Development Environment Set-up	Guided set-up of Python, Visual Studio Code, Git and GitHub, and API access. Every participant leaves with a working, reproducible environment.
Afternoon	Essential Python for AI	Just enough Python to build: variables, functions, working with libraries, JSON and calling web APIs. Pitched for mixed backgrounds with paired support.
Afternoon	The AI Tools & Model Ecosystem	Surveying the major providers (OpenAI, Anthropic, Google, Meta) and open models; hosted vs. self-hosted; selecting the right model on cost, speed, quality and data-residency grounds.
Outcome Each participant makes their first successful call to a large language model and interprets the response.		

Day 2 - How AI Models Work & Prompt Engineering

Objective: Build an accurate mental model of how LLMs behave, then master the prompting techniques that turn inconsistent outputs into reliable ones.

Session	Topic	Focus
Morning	Inside a Large Language Model	A practical, non-mathematical view of tokens, embeddings, context windows and temperature - enough to predict and control model behaviour.
Morning	Prompt Engineering Fundamentals	Zero-shot, few-shot and chain-of-thought prompting; system prompts and persona design; structuring prompts for consistent, high-quality results.
Afternoon	Context, Reliability & Prompt Security	Managing context across a conversation; detecting and reducing hallucinations; prompt-injection risks and safe handling of sensitive inputs.
Afternoon	Introduction to Retrieval-Augmented Generation (RAG)	Why grounding a model in your own data matters; the RAG pattern at a conceptual level, setting up Day 3's build.

Outcome Participants build a reusable prompt library and a working question-and-answer assistant.

Day 3 - Building a Real AI Application

Objective: Move from single calls to a functioning, grounded AI application built on the participant's own documents.

Session	Topic	Focus
Morning	APIs & Integration	Wrapping AI behind a simple backend service (FastAPI); managing keys, rate limits and errors; the anatomy of a production-ready AI endpoint.
Morning	Embeddings & Vector Databases	Turning documents into searchable meaning; semantic search; storing and querying embeddings in a vector database.
Afternoon	Build a RAG Chatbot	Assembling the pieces into a retrieval-augmented chatbot that answers questions from a custom knowledge base - e.g. a policy or operations document.
Afternoon	A Simple User Interface	Exposing the application through a lightweight interface so a non-technical colleague could use it.

Outcome A working RAG chatbot answering questions over the participant's own content.

Day 4 - No-Code, Automation & Agentic AI

Objective: Extend from single applications to automated workflows and autonomous agents - the skills behind bespoke, high-value AI solutions.

Session	Topic	Focus
Morning	No-Code & Low-Code AI Builders	Rapidly delivering value with platforms such as Microsoft Copilot Studio and Make.com; knowing when no-code is the right tool and when custom code is needed.
Morning	Agentic AI Fundamentals	The shift from chatbots to agents that reason, plan and use tools; function calling and the ReAct pattern; safety and human-in-the-loop control.
Afternoon	Workflow Automation with n8n	The open-source automation platform: triggers, nodes and connecting AI to everyday business applications; self-hosting for data control.
Afternoon	Build an Automated Workflow	Constructing an end-to-end automation: monitor an inbox or folder, extract information with AI, store it, and generate a response.

Outcome A working automation workflow and a simple tool-using agent.

Day 5 - Responsible AI, Cost, Deployment & Showcase

Objective: Make what participants have built safe, affordable and deployable - then consolidate the week into a demonstrable result.

Session	Topic	Focus
Morning	Responsible & Compliant AI	Bias, transparency and accountability; data privacy and consent; aligning AI systems with Mauritian and emerging data-protection legislation.
Morning	Cost Optimisation for Resource-Constrained Settings	Token economics; caching and model-selection strategies; maximising value while controlling spend - a core competency for the public sector.
Afternoon	Deployment Basics	Getting an AI application online; introductory monitoring and logging; what changes between a prototype and something colleagues rely on.
Afternoon	Capstone Consolidation & Showcase	Participants finalise and present the application or automation they built during the week to peers and facilitators, and leave with a personal AI action plan.

Outcome A presented, working AI artifact plus a personal roadmap and a clear pathway into the full AI Developer Bootcamp.

5. Learning Outcomes

On completing the five-day Bootcamp, participants will be able to:

- Explain, in practical terms, how modern AI and large language models work, and select an appropriate model for a given task on cost, quality and data-residency grounds.
- Set up a professional AI development environment and make reliable, well-structured calls to AI models through code.
- Apply prompt-engineering techniques to obtain consistent, high-quality outputs, and recognise and mitigate hallucinations and prompt-injection risks.
- Build a retrieval-augmented (RAG) application that answers questions grounded in an organisation's own documents.
- Construct an automated AI workflow that connects AI to everyday business systems and build a simple tool-using agent.
- Apply no-code and low-code platforms appropriately, and judge when custom code is warranted.
- Design AI systems that respect data privacy and align with Mauritian and emerging data-protection legislation.
- Optimise AI applications for cost and deploy a working application beyond the local machine.

6. Who Should Attend

The Bootcamp is designed for a mixed-ability technical cohort. It suits IT professionals, developers, data staff and technically-inclined officers who will be responsible for building or overseeing AI solutions. Prior programming exposure in any language is helpful, but the programme is deliberately scaffolded - with paired work and facilitator support - so that participants of varying backgrounds can progress together. It is not aimed at complete beginners with no technical exposure, for whom the AI Academy offers separate awareness-level programmes.

7. Delivery Approach & Requirements

The Bootcamp is delivered in person over five consecutive days, combining short concept sessions with extended, supported build time. The following arrangements support effective delivery.

Facilities and equipment

- A training room configured for hands-on work, with reliable power and stable, adequate internet connectivity.
- Data projection and a writing surface (whiteboard or equivalent) for live demonstration.
- One workstation per participant. Participants may use their own laptops meeting the minimum specification; loan machines can be arranged where required.
- Access to a shared GPU server (for example, an Nvidia RTX A4500-class device) for exercises that benefit from local acceleration and for demonstrating data-sovereign, self-hosted deployment.

Software and accounts

- Python 3.10 or later with virtual-environment support; Node.js and npm.
- Visual Studio Code as the development environment.
- Git installed locally and a GitHub account per participant.
- Provisioned access to AI provider APIs (for example OpenAI and Anthropic) and a vector-database service, arranged ahead of the programme.

Participant prerequisites

- Basic programming familiarity in any language is recommended.
- Comfort with fundamental web concepts (what an API is, basic HTML) is helpful but not essential.
- Sound logical reasoning and comfort working with straightforward mathematical notation.
- Full-time availability across all five consecutive days, as each day builds directly on the last.

Cohort size

To preserve the hands-on, supported character of the Bootcamp, a cohort of up to twenty participants is recommended, with facilitator support scaled to the group. Larger groups can be accommodated by arrangement with additional facilitation.

AI resources budget

A modest per-participant allowance for AI API credits, subscriptions and cloud resources ensures every participant can complete the build exercises. In the full programme this is set at approximately US\$100 per participant; the five-day intensive requires only a fraction of this, and the exact figure can be confirmed with the agency during planning.

8. Facilitators

The Bootcamp is designed, developed and delivered by two senior practitioners whose combined experience spans public-sector digital leadership and world-class technical research.

Dr Kiru Pillay

Dr Kiru Pillay is a highly accomplished ICT executive, academic and researcher with more than 38 years of experience spanning the public sector, academia, telecommunications, cybersecurity and digital transformation. From 2016 to 2026 he served as Chief Director responsible for South Africa's national cybersecurity mandate, leading the development of the national CSIRT, chairing the SADC Regional CIRT Task Force, and representing South Africa in international forums including the United Nations, the African Union, BRICS and the ITU. He has developed postgraduate programmes in cybersecurity, digital diplomacy and artificial intelligence, and advises governments and international organisations on digital transformation and capacity building across Africa.

Professor Bruce W. Watson

Professor Bruce W. Watson is an internationally recognised computer scientist and technology leader with more than three decades of experience across academia and advanced industrial research in artificial intelligence, cybersecurity, high-performance computing and algorithms. Full Professor of Information Science at Stellenbosch University, he is Co-Director of the Centre for Artificial Intelligence Research and holds a Visiting Professorship in Algorithmics at King's College London. He has held senior technical roles at ASML, Cisco Systems Security, Microsoft and other leading technology organisations, and holds two doctoral degrees in Computer Science and Computer Engineering.