

RMACC 2026

Building boisestate.ai

Lessons Learned from Developing a Cost-Effective
Internal AI Platform for Higher Education on **AWS**
Bedrock AgentCore.

<https://github.com/Boise-State-Development/agentcore-public-stack>



Why an Internal AI Platform?

- **Avoid Vendor Lock-in** – freedom to choose, swap best-fit models, avoid long term contracts and CapEx expenditure
- **Data Sovereignty** – keep institutional data inside our AWS account.
- **Model Training** – Avoid model training with our data
- **Cost Reality** – 30k subs × \$20 = \$7.2 M/yr → unsustainable at enterprise scale

Why an Internal AI Platform? (continued)

- **Equitable Access** – every student gets access to the same models. No “haves and have-nots”
- **Campus Systems Integration** – Safely leverage institutional data via MCP servers (Canvas, Google Workspace, PeopleSoft, Library, etc.)
- **Risk/Reward** – exploring an internal platform is low risk, high reward. If it works, we’ve built lasting infrastructure

“\$300 Billion Evaporated. The SaaS - Apocalypse Has Begun.”



Have agentic coding tools shifted the landscape of buy vs. build?

<https://www.forbes.com/sites/donmuir/2026/02/04/300-billion-evaporated-the-saaspocalypse-has-begun/>

The Challenges of Developing an Internal Platform

Spoiler alert: It's cost-effectiveness and context engineering



The Trade-off: The Consumption-Based Paradigm Shift

We must answer these three questions

<i>Question</i>	<i>Challenge</i>
How do we match user expectations?	Users compare us to \$20/month products with unlimited usage
How do we control costs?	Every conversation, every tool call, every token has a price
How do we make AI useful for the enterprise?	Generic AI is interesting; contextual AI is transformational

boisestate.ai: Lessons Learned



AgentCore Public Stack



An open-source, full-stack Generative AI platform for public institutions.

Version

1.0.0-beta.24

Commits

1,329

across all branches

Test Cases

3,601

Contributors

13

B AgentCore Public Stack

An open-source, production-ready Generative AI platform for institutions *Built by Boise State University, designed for everyone.*

🔗 Release **v1.0.0-beta.25** 🔄 Nightly Build & Test **passing**

📦 Python **3.13+** 📦 Angular **v21** 📦 AWS **Bedrock AgentCore** 📦 Tailwind **CSS v4.2** 📦 License **PolyForm Noncommercial**

[Features](#) · [Architecture](#) · [Admin Dashboard](#) · [Deployment](#) · [Release Notes](#) · [Contributing](#)

🚀 Ready to deploy? Fork the repo and follow the [GitHub Actions Quick Start](#) to deploy to your AWS environment in under an hour.

🤔 What Is This?

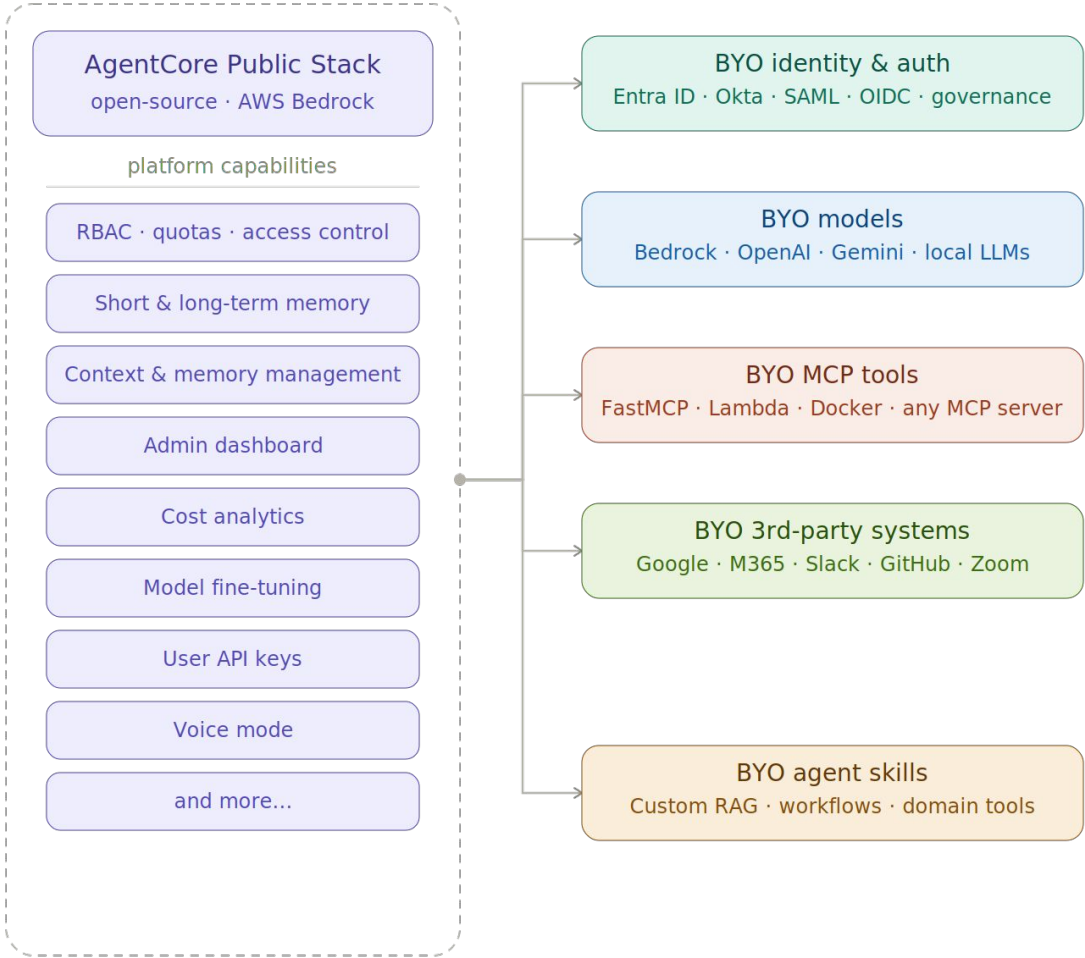
AgentCore Public Stack is the working codebase behind [boisestate.ai](#), a full-stack AI assistant platform serving students, faculty, and staff at Boise State University. It combines **AWS Bedrock AgentCore** and **Strands Agents** into a turnkey system that any institution can fork, deploy, and make their own — without writing a single line of agent code.

The platform ships with a modern chat interface, a full admin dashboard, multi-model support, tool and MCP server management, quota enforcement, cost tracking, and role-based access control — all configurable through the UI.

🌐 Why Boise State Built This

Most institutions face the same AI dilemma: vendor-hosted per-seat subscriptions don't scale, and building from scratch is too expensive.





The Key Insight: Context is everything. An AI that knows nothing about your institution is just another chatbot. An AI that understands your policies, systems, and data becomes an institutional asset.

Model Flexibility

Run inference via Bedrock, API keys to other providers, or even run local models.

← Back to Admin

Manage Models

View and manage AI models available to users. [+ Add Model](#)

Search & Filters

Search: Provider: Status:

Showing 2 models [Browse Bedrock Models](#) → [Browse Gemini Models](#) → [Browse OpenAI Models](#) →

Claude Sonnet 4.6 Enabled bedrock Anthropic

global.anthropic.claude-sonnet-4-6

Allowed Roles	Pricing (per 1M tokens)	Modalities
Default User	Input: \$3.00 Output: \$15.00	Input: text, image, TEXT, IMAGE Output: text, TEXT

[Edit](#) [Delete](#)

Claude Haiku 4.5 Enabled bedrock Anthropic

global.anthropic.claude-haiku-4-5-20251001-v1:0

Allowed Roles	Pricing (per 1M tokens)	Modalities
System Administrator Student Default User	Input: \$1.00 Output: \$5.00	Input: text, image, TEXT, IMAGE Output: text, TEXT

[Edit](#) [Delete](#)

New Feature

BYO Tools

Add MCP Tools and A2A Agents relevant to your environment

← Back to Admin

Tool Catalog

Manage tool metadata and role assignments.

+ Add Tool

Q Search by name or ID... All Statuses All Categories

TOOL	CATEGORY	ACCESS	DEFAULT	STATUS	ACTI
Analyze Spreadsheet analyze_spreadsheet	Data	Public	✓	active	🗑️ ✎️
Charts & Graphs create_visualization	Data	3 roles	⊙	active	🗑️ ✎️
List Spreadsheets list_spreadsheets	Data	Public	✓	active	🗑️ ✎️
Grants.gov grants_gov	Research	0 roles	⊙	active	🗑️ ✎️
Semantic Scholar semantic_scholar	Research	0 roles	⊙	active	🗑️ ✎️
Class Search class_search	Search	Public	✓	active	🗑️ ✎️
Policy Search policy_search	Search	Public	⊙	active	🗑️ ✎️
Search Boise State search_boise_state	Search	Public	✓	active	🗑️ ✎️
URL Fetcher fetch_url_content	Search	Public	✓	active	🗑️ ✎️
Web Search web_search	Search	Public	⊙	active	🗑️ ✎️

New Feature

Role-based Quotas

Create quota tiers and assign to different users

Quota Assignments

Assign quota tiers to users, roles, or email domains

[+ Create Assignment](#)

Search & Filters

Search:

Tier:

Type:

Status:

Showing 4 assignments [View Tiers](#) [Quota Inspector](#)

Staff Enabled Priority: 400 [App Role](#)

a2a908ac-5d17-4679-9513-aa4762c52848

Assignment Value: Staff (staff) Created: 3/8/26, 1:15 PM by 111653897

[Edit](#) [Delete](#)

Faculty Enabled Priority: 400 [App Role](#)

69049b25-0a68-4886-a041-38ff6ee8feaf

Assignment Value: Faculty (faculty) Created: 3/8/26, 1:16 PM by 111653897

[Edit](#) [Delete](#)

Student Enabled Priority: 300 [App Role](#)

e053e3af-27fb-4e1e-8c01-e788163ef7f3

Assignment Value: Student (student) Created: 3/8/26, 1:16 PM by 111653897

[Edit](#) [Delete](#)

Model Fine-Tuning


*Commercial LLMs today;
fine-tuned open-weight for
routine work tomorrow*

New Training Job

Upload your dataset, select a base model, configure hyperparameters, and start training.

1. Upload Dataset

Upload your training data as a JSONL, CSV, or TXT file. Each record should include "text" and "label" fields.



Train / Test Split



2. Select Base Model

Choose a pre-trained model to fine-tune, or enter any HuggingFace model ID.



New Feature


Api Keys

*Integrate your system
into other systems*

API Keys

Programmatic access to the chat API

[+ Create New API Key](#)

 No API key yet
Create your first key to get started with the API.

Available Models
4 models you can use via API


Important Information
Key policies and usage guidelines

INTERACTIVE EXAMPLES WITH YOUR AVAILABLE MODELS

Usage Examples

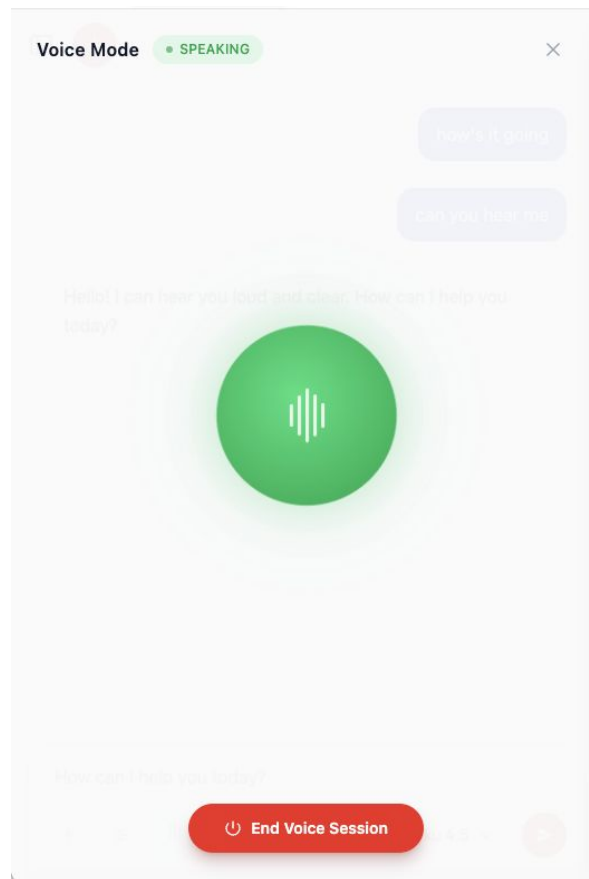
Select Model

Nova 2 Sonic (amazon.nova-2-sonic-v1:0)

Response Type 

Voice Mode

Real-time, bidirectional voice chat feature, which transcribes, reasons with tools, and streams synthesized speech back



Connectors + MCP Tools

Demo:

Canvas for Faculty Connector

List Courses:

"Pull up my courses — I want to find DEMO-101."

"Show me everything I'm teaching, including concluded sections."

List Assignments:

"What's on the schedule for Faculty Demo: Intro to MCP? I want to sanity-check the points add up — Syllabus Ack, two homeworks, a quiz, and the final project."

"Which assignments in DEMO-101 are still upcoming?"

List Submissions:

"For Homework 2: Control Flow — who's marked missing? I want to nudge them before grades post."

"Pull the late submissions on Quiz 1: Functions so I can decide whether to accept them at full credit."

"On Homework 1: Variables and Types, who got an excused entry vs. a real zero?"

"Show me submissions for the Syllabus Acknowledgment — anyone still hasn't done the 5-point gimme?"

Demo:

Canvas for Faculty Connector

Get Assignment Analytics:

"Was Quiz 1: Functions too hard? Give me the score distribution — I'm seeing complaints."

"On Homework 2: Control Flow, how does the late-submission cohort compare to the on-time scores?"

"Pull analytics for the Syllabus Acknowledgment — I want to confirm nearly everyone aced it before I use it as a baseline."

Get Student Analytics:

"Three weeks into DEMO-101 — who's at risk? Sort by lowest engagement."

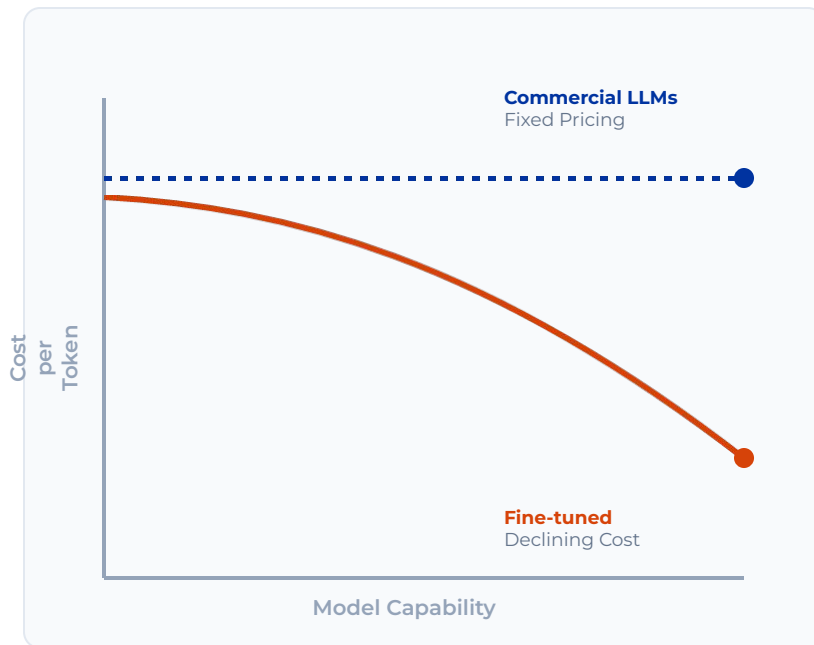
"Has Maya Morales been keeping up? Pull just her row." (uses student_id for one student)

"Give me the bottom 5 on engagement so I can flag them for early-alert."





Open-weight, tuned, cost-controlled


Strategic Implementation

- **Future Outlook:** Commercial LLMs today; fine-tuned open-weight for routine work tomorrow
- **Key Partner:** Casey Kennington — NSF NAIRR AI Education Fellow · Foundation Board Endowed Professor, CS
- **Efficiency Lever:** Cost lever for agentic automation (ServiceNow-style deflection)



GRAPHIC ASSETS







BOISE STATE UNIVERSITY