FOR SECURITY AND PLATFORM TEAMS

Posit and Amazon Bedrock

A guide for security and platform teams

November 2026

Your data scientists want AI help in their daily tools. With Posit and Amazon Bedrock you can say yes and keep control. You decide who can use AI and which models. Every call is recorded. Spend is visible and capped. Guardrails your organization defines check every prompt. These controls are standard AWS services that your admins already run, so there are no new vendor contracts and no API keys to issue, rotate, or leak.

How it fits together

Posit tools call Amazon Bedrock with AWS credentials that your platform already provides:

- **Posit Assistant in Positron on Amazon SageMaker** uses the Space's execution role.
- **Positron in Posit Workbench** uses Workbench-managed AWS credentials, which assume an IAM role for the signed-in user.
- **Positron Desktop** uses the user's AWS IAM Identity Center (SSO) profile.
- **Apps and agents on Posit Connect**, built with ellmer, chatlas, or Commons, use the AWS credentials configured for that content.

Bedrock checks IAM, applies your guardrails and logging, and then sends the request to the model. Bedrock runs on AWS infrastructure: AWS hosts the models, and prompts, responses, and context stay within AWS.

Common security questions

Question	Answer	Where you control it
Who can use AI, and which models?	Only principals you allow, and only the models you list. Bedrock denies a request for any other model, whatever the user picks in the tool.	IAM policies on `bedrock:InvokeModel`, scoped to approved models and inference profiles. Optionally an SCP.
What is recorded?	Every call: who made it, which model, and when. With invocation logging on, also the prompt, the response, and token counts.	AWS CloudTrail, plus Bedrock model invocation logging to CloudWatch Logs or Amazon S3.
What does it cost, and who is spending it?	Spend by model, by month, and by IAM principal, with alerts before a budget is exceeded.	AWS Cost Explorer, the CUR 2.0 export with IAM principal data, and AWS Budgets.

Question	Answer	Where you control it
What stops sensitive data or prohibited use?	Guardrails you define, such as sensitive data filters and denied topics, check what is sent to the model and what comes back. Posit Assistant needs no change.	Amazon Bedrock Guardrails, enforced at the account level.
Where does our data go?	Bedrock does not use prompts and responses to train models and does not share them with model providers. US-only inference profiles keep processing in US Regions.	Model and inference profile choice in IAM. AWS PrivateLink for private network access.
Who holds the keys?	No one. Credentials are short-lived and come from IAM roles.	SageMaker execution roles and IAM Identity Center.

Rollout checklist

1. **Allow approved models only.** Grant `bedrock:InvokeModel` and `bedrock:InvokeModelWithResponseStream` on the inference profiles and models you approve.
2. **Record each person, not a shared role.** CloudTrail must show the user, not only the role. On SageMaker, give each user their own execution role, or set the domain's `ExecutionRoleIdentityConfig` to `USER_PROFILE_NAME` so CloudTrail shows the user in `sourceIdentity`. With Workbench-managed credentials and IAM Identity Center (Positron Desktop), each user gets their own role session. Make sure the session name identifies the user.
3. **Turn on model invocation logging.** CloudTrail shows who called which model, but not the prompt, the response, or token counts.
4. **Grant `bedrock:ApplyGuardrail` to every caller of a guarded model.** Without it, every call to that model fails, including harmless ones.
5. **Accept third-party model agreements once per account.** This gives you an approval step for each new model provider.
6. **Set up cost reporting.** Create a CUR 2.0 export with IAM principal data. Per-user cost allocation tags must be activated in the payer account. Third-party models bill as separate "(Amazon Bedrock Edition)" services, so include them in reports and budgets. Count cached tokens: most AI assistant input is read from the prompt cache.
7. **Optional: keep traffic private.** Add a `bedrock-runtime` VPC endpoint (AWS PrivateLink). This covers SageMaker Spaces and Workbench servers in your VPC. Positron Desktop traffic uses the endpoint only when the device connects through your network.

Good to know

- **Guardrails work at the model level.** In our tests they checked the user's prompt and the model's output, but not tool results sent back to the model, such as rows returned by a query. For controls on agent tool calls, see Amazon Bedrock AgentCore Policy.
- **The model list in a tool is not the control. IAM is.** A user who picks a model you have not approved gets an access-denied error.

Learn more

- Positron on Amazon SageMaker: 5-minute walkthrough
- Positron on SageMaker administration guide

- Prompt to production: governed AI workflows with Posit and AWS
- Demo repository
- Meet with a Posit and AWS expert