

# Cloudvisor AI Agent - One-Click Deployment Guide

Deploy a complete Bedrock RAG agent in your AWS account in under 10 minutes.

## Quick Deploy (Recommended)

Choose **one** of the following options. Do not deploy both – they create conflicting Bedrock Knowledge Bases and will fail:

- **Option A: Amazon S3 Vectors (lower cost)**
- **Option B: OpenSearch Serverless (higher cost)**

 **Important:** Pick Option A OR Option B, not both.

## Estimated Monthly Costs (eu-central-1)

Component	Option A (S3 Vectors)	Option B (OpenSearch)
<b>Daily Cost</b>	~\$0.6/day	~\$16.33/day
<b>Monthly Cost</b>	~\$19/month	~\$506/month (Standard)
		~\$262/month (Dev-Test)

**Includes:** Bedrock Nova Lite (inference), Titan Embeddings (ingestion), vector storage, Lambda, API Gateway, and CloudFront.

### Cost Notes:

- S3 Vectors is ~26× cheaper than OpenSearch Standard

- Costs scale with document ingestion volume and chat frequency
- Use OpenSearch Dev-Test collection if exploring/prototyping only
-  **Important:** Verify current AWS pricing at [aws.amazon.com](https://aws.amazon.com) – prices change frequently and vary by region.

## Option A: S3 Vectors

This option uses Amazon S3 Vectors as the Bedrock Knowledge Base vector store.

### Step 1: Deploy S3 Vectors Stack

Click this link to deploy the full S3 Vectors stack:

#### [Deploy S3 Vectors Stack](#) →

- Check "I acknowledge that AWS CloudFormation might create IAM resources"
- Click **Create stack**
- Wait 5-10 minutes for completion

Learn more: [S3 Vectors + Bedrock KB integration guide](#)

## Option B: OpenSearch

### Step 1: Deploy OpenSearch Infrastructure

Click this link to deploy the OpenSearch Serverless collection and index:

#### [Deploy OpenSearch Stack](#) →

- Check "I acknowledge that AWS CloudFormation might create IAM resources"
- Click **Create stack**
- Wait 3-5 minutes for completion

## Step 2: Deploy Main Application

Click this link to deploy the agent application:

### [Deploy Agent Stack](#) →

- Check "I acknowledge that AWS CloudFormation might create IAM resources"
- Click **Create stack**
- Wait 5-10 minutes for completion

## What Gets Deployed

### Vector Store (Choose One)

- **S3 Vectors:** Lower-cost vector storage integrated with Bedrock KB
- **OpenSearch Serverless:** Vector collection + index, IAM roles, and security policies

### Main Application Stack (Common)

- **Frontend:** CloudFront + S3 hosted UI
- **Backend:** API Gateway + Lambda functions
- **AI:** Bedrock Knowledge Base + Amazon Nova Lite model
- **Storage:** S3 buckets for documents and UI assets
- **Security:** IAM roles with least privilege access

# Getting Started

## 1. Get your agent URL:

- Go to CloudFormation → `cloudvisor-chatbot` stack → Outputs tab
- Copy the `CloudFrontURL` value

## 2. Upload documents:

- Option A: Use the upload button in the agent UI
- Option B: Upload directly to the S3 documents bucket

## 3. Sync Knowledge Base:

- Click "Rebuild Knowledge Base" in the UI
- Wait 2-3 minutes for processing

## 4. Start chatting:

- Ask questions about your documents
- The AI will only answer based on your uploaded content

# Troubleshooting

## Upload Issues

If file uploads fail:

1. Check S3 bucket CORS configuration
2. Try uploading directly to S3 console
3. Use "Rebuild Knowledge Base" after manual upload

## Chat Not Working

1. Verify Knowledge Base has documents synced

2. Check Lambda logs in CloudWatch
3. Ensure Nova Lite model access is granted

## Cleanup

### Option 1: AWS Console (Recommended)

#### Step 1: Delete Application Stack

1. Go to CloudFormation Console
2. Select the stack you deployed:
  - S3 Vectors option: `cloudvisor-chatbot`
  - OpenSearch option: `cloudvisor-chatbot`
3. Click **Delete**
4. Click **Delete stack** to confirm
5. Wait 5-10 minutes for deletion to complete

#### Step 2: Delete OpenSearch Stack (OpenSearch option only)

If you deployed Option B (OpenSearch), delete its infrastructure:

1. In CloudFormation console, select `cloudvisor-opensearch` stack
2. Click **Delete**
3. Click **Delete stack** to confirm
4. Wait 3-5 minutes for deletion to complete

#### Step 3: Verify Cleanup (Optional)

1. Check S3 console - document and UI buckets should be deleted
2. Check Bedrock console - Knowledge Base should be removed

3. Check S3 Vectors - vector bucket and index should be deleted (Option A)
4. Check OpenSearch Serverless - collection should be deleted (Option B)

## Option 2: AWS CLI

```
# Delete main stack first
aws cloudformation delete-stack --stack-name clo
udvisor-chatbot --region eu-central-1

# Wait for completion, then delete OpenSearch st
ack (Option B only)
aws cloudformation delete-stack --stack-name clo
udvisor-opensearch --region eu-central-1
```

### Important Notes

- **For OpenSearch option:** Delete OpenSearch stack after the main stack – the main stack depends on OpenSearch resources
- **Document and UI buckets** are automatically emptied and deleted by CloudFormation
- **S3 vector buckets are not auto-emptied** - delete vectors or the index first if Option A stack deletion fails
- **All data will be permanently lost** - download any important documents first
- **Costs stop immediately** once resources are deleted

### Support

For issues or questions:

- Check AWS CloudFormation console for deployment errors
  - Review CloudWatch logs for runtime issues
  - Contact Cloudvisor support for assistance
- 

 **Ready to deploy? Click the links above to get started!**