Private Document Evidence Analyst for AWS by Code Creator

Customer Instructions and Operations Guide

From sensitive PDF to defensible page-cited report privately on AWS.

START SIMPLE

Verified Evidence works immediately after launch and does not require an AI provider, API key, IAM role, S3 bucket, or GPU. Configure advanced AWS or AI features only when you need them.

Built for AWS Marketplace subscribers, auditors, risk teams, due-diligence users, and administrators.

Ubuntu LTS | x86_64 | 80 GiB root EBS | HTTPS web interface

Version 1.0 | September 2026

Contents

IMPORTANT ABOUT COSTS

The AWS Marketplace software charge is separate from AWS infrastructure and service charges. EC2, EBS, public IPv4/Elastic IP, S3, Amazon Bedrock, and third-party AI providers can each create additional charges depending on what you enable.

1. What the product does and what is optional

Private Document Evidence Analyst runs inside your AWS account and reviews long-form PDFs to produce a professional Word report with exact page-cited findings. The standard workflow is deliberately CPU-first and can run without an LLM.

Capability	What it does	Required?
Verified Evidence	Deterministic complete-document evidence review with exact page citations and a Word report.	Yes - core workflow
Verified Evidence + AI Analysis	Adds optional AI synthesis after verified evidence is established.	Optional
AWS Bedrock	AWS-native optional AI provider using the EC2 IAM role.	Optional
OpenAI / Anthropic	External optional AI providers using buyer-supplied API keys.	Optional
Amazon S3 Report Delivery	Copies finished DOCX reports to your private S3 bucket.	Optional
Stop EC2 when report is ready	Automatically stops EC2 after the review and any required S3 delivery.	Optional
Local Muse	Private local GPU inference using the model included on EBS.	Optional
MCP Servers	Connects trusted Streamable HTTP MCP tools for advanced workflows.	Optional

PRIVACY MODEL

Source documents and review state remain on your EC2/EBS environment by default. The product does not upload source PDFs to S3. If you select OpenAI or Anthropic, selected review/document content may be sent to that external provider. Bedrock uses your AWS account and IAM permissions.

2. Launch the EC2 instance

2.1 Launch requirements

Item	Recommended / required setting	Why
Network	Public subnet with a public IPv4 address for the validated Marketplace configuration.	The appliance uses the public IP for browser access and per-instance HTTPS certificate setup.
TCP 22	Allow from your administrator IP or trusted CIDR only.	SSH administration.
TCP 80	Allow inbound internet access.	HTTP redirect and certificate/bootstrap traffic.
TCP 443	Allow from users who need the web interface.	HTTPS application access.
Internal TCP 8090	Do not expose publicly.	The application container is intentionally bound to localhost.
EC2 key pair	Select an Amazon EC2 key pair you control.	Cloud-init injects this key for the ubuntu SSH user.
IAM role	Not required for first launch or Verified Evidence.	Attach later only for Bedrock and/or S3.
Root EBS	Use the Marketplace default 80 GiB root volume or larger.	Includes application, Docker assets, reports, and optional local model files.
GPU	Not required for normal use.	Only Local Muse needs a compatible NVIDIA GPU instance.
Outbound TCP 443	Allow standard outbound HTTPS, or provide approved equivalent connectivity.	Used by optional AWS/AI providers, S3 delivery, remote MCP servers, and supporting platform operations.

Launch steps

1. Launch the product from AWS Marketplace into the AWS Region and VPC you want to use.
2. Select a supported x86_64 EC2 instance type. Normal Verified Evidence is CPU-first and does not require a GPU.
3. Select your EC2 key pair.
4. Assign the security group described above.
5. Enable a public IPv4 address. This is required by the validated Marketplace browser-access and HTTPS setup.
6. Launch the instance and wait for both EC2 status checks to pass.

DO NOT COPY SELLER CREDENTIALS

The AMI contains no Code Creator AWS access keys, no seller IAM role, and no fixed administrator password. Each fresh instance generates customer-specific access information on first boot.

3. First login and credentials

Connect by SSH using the Amazon private key selected during launch and the SSH username **ubuntu** Password-based SSH login and root SSH login are disabled.

```
ssh -i YOUR_PRIVATE_KEY.pem ubuntu@PUBLIC_IP
```

After connecting, display the generated login file:

```
cat ~/FIRST_LOGIN.txt
```

The file shows the current HTTPS URL, username, generated administrator password, SSH username, quick-start guidance, and helper commands. Protect this file because it contains the web administrator password.

You can also display only the application credentials with:

```
codecreator-agent-credentials
```

Primary browser URL: https://PUBLIC_IP/

Advanced configuration URL: https://PUBLIC_IP/advanced-workflows

This release uses one generated web administrator credential. It does not provide individual web accounts, role-based access control, or single sign-on. Limit inbound HTTPS access to approved networks whenever practical and protect the generated password as an administrative credential.

401 IS NORMAL WITHOUT CREDENTIALS

The HTTPS application requires authentication. An unauthenticated request returning HTTP 401 is expected and confirms that the web authentication gate is active.

4. The two GUI pages and when to use each

4.1 Evidence Review - the main customer workflow

Use the main page for routine document work. It intentionally reduces the workflow to four decisions.

Main page area	What it does
Choose Document	Upload a PDF or select an existing PDF already stored on this instance.
Choose Review	Select Supplier Risk, Internal Audit, Executive Due Diligence, or Custom Review. The review assignment can be edited.
Choose Output	Choose Verified Evidence or Verified Evidence + AI Analysis.
Run Review	Start the review and optionally request Stop EC2 when report is ready.
Active Reviews	Shows whether work is queued or running.
Cost Control	Shows SAFE TO STOP only when the platform has no active work or required S3-delivery block.
Review History	Shows completed reviews, Word report downloads, and S3 delivery state when S3 is enabled.
Advanced Workflows link	Moves to provider, S3, MCP, and advanced configuration.

4.2 Advanced Workflows - configuration and integrations

Use Advanced Workflows when you need optional providers, AWS integrations, external tools, or more operational control. Standard PDF evidence review does not require this page.

Advanced area	Purpose
AI Provider	Select Local Muse, OpenAI, Anthropic, or AWS Bedrock and save provider settings.
Amazon S3 Report Delivery	Configure a private S3 destination, save settings, and test delivery.
Documents	Upload, view, and delete documents stored on the instance.
MCP Servers	Add trusted Streamable HTTP MCP servers for advanced tool use.
New Review / Completion control	Run advanced work and choose Keep Running or Stop EC2 when report is ready.
Review History	View advanced job results, report links, S3 status, retry failed S3 delivery, and delete history.
Status cards	Show active provider, MCP state, active reviews, and cost-control status.

RULE OF THUMB

If your goal is simply “upload a sensitive PDF and get a page-cited Word report,” stay on Evidence Review. Go to Advanced Workflows only when you need AI configuration, S3 delivery, MCP tools, or advanced job controls.

5. Run a standard document review

5.1 Review types

Review type	Best suited for
Supplier Risk	Supplier concentration, dependencies, obligations, deadlines, operating risk, and third-party evidence.
Internal Audit	Controls, exceptions, remediation commitments, policy evidence, dates, amounts, and management actions.
Executive Due Diligence	Business, financial, customer, supplier, location, liquidity, strategic decision, and executive evidence.
Custom Review	A buyer-defined assignment when the three built-in review templates do not match the objective.

5.2 Output modes

Output mode	Behavior	Dependencies
Verified Evidence	Deterministic whole-document review. Produces page-cited findings without an AI call.	No IAM role, API key, or GPU required.
Verified Evidence + AI Analysis	Builds the same verified evidence first, then asks the selected provider for optional synthesis. AI is fail-closed: if the synthesis does not pass the output gate, the verified evidence remains the authoritative result.	Requires a configured/available AI provider.

5.3 Supported PDF input

For reliable page-cited findings, use a PDF that contains a readable text layer. Image-only scans must be processed with OCR before upload. Password-protected, damaged, or text-inaccessible PDFs may fail or produce incomplete findings.

The product does not claim built-in OCR. Remove document passwords only when permitted by your organization, and validate the extracted findings before relying on a report for consequential decisions. Consult the AWS Marketplace listing and the product interface for any supported-format or enforced upload limits.

Standard run sequence

1. Upload or select a PDF.
2. Select a review type and edit the Review assignment if needed.
3. Choose Verified Evidence for the simplest deterministic workflow, or Verified Evidence + AI Analysis if a provider is configured.
4. Leave Stop EC2 when report is ready unchecked if you want to download the report immediately from the GUI.
5. Choose Start Review.
6. Wait until Review History shows completed.
7. Choose Download Report to retrieve the DOCX.

IMPORTANT FOR AUTO-STOP USERS

If you choose Stop EC2 when report is ready but do not use S3 Report Delivery, the browser becomes unavailable after EC2 stops. The report remains on EBS, but you must start the instance again before downloading it from the GUI. S3 delivery is therefore strongly recommended when using unattended auto-stop.

6. Stop, restart, changing public IPs, and Elastic IP

6.1 What happens when EC2 stops

Stopping is different from terminating. Stopping preserves the EBS root volume and the product data stored on it. Terminating the instance may delete the root EBS volume if Delete on termination is enabled.

With an automatically assigned public IPv4 address, AWS releases that address when the instance is stopped and normally assigns a new public IPv4 address when the instance starts again. The product is designed to detect the new public IP, refresh HTTPS for the new address, and update FIRST_LOGIN.txt.

After a restart without an Elastic IP:

1. Open the EC2 console and note the new Public IPv4 address.
2. Connect with SSH to the new address using the same EC2 private key.
3. Run cat ~/FIRST_LOGIN.txt or codecreator-agent-url to obtain the current application URL.
4. Use the updated HTTPS URL in your browser. Your documents, reports, history, and saved configuration remain on EBS.

AWS reference: How EC2 stop/start affects public IPv4 addresses

6.2 Optional: use an Elastic IP for a stable address

An Elastic IP is an AWS-owned static public IPv4 address that remains associated with the instance across stop/start cycles. This avoids changing bookmarks and user instructions each time the server restarts. AWS charges for Elastic IP/public IPv4 addresses, including Elastic IPs whether associated or idle; check current AWS pricing.

1. Open EC2 > Network & Security > Elastic IPs.
2. Choose Allocate Elastic IP address and allocate an address in the same network border group/Region as the instance.
3. Select the new Elastic IP and choose Actions > Associate Elastic IP address.
4. Choose Resource type: Instance and select the Evidence Analyst instance.
5. Choose Associate.
6. Stop and start the EC2 instance so the validated boot-time public-IP and HTTPS synchronization runs for the Elastic IP. Then use codecreator-agent-url and verify the displayed HTTPS address before distributing a bookmark.
7. Allow the product time to synchronize HTTPS to the new stable address. Confirm the current application URL with codecreator-agent-url.

AWS reference: Allocate and associate an Elastic IP address

7. IAM role setup for Bedrock and S3

WHEN AN IAM ROLE IS NEEDED

No IAM role is required for Verified Evidence. Attach an EC2 IAM role only if you want AWS Bedrock AI and/or Amazon S3 Report Delivery. The application uses temporary role credentials; do not place long-lived AWS access keys on the instance.

An EC2 instance can have one IAM role attached at a time. If you use both Bedrock and S3, put both permission sets on the same role.

7.1 Create the EC2 role

1. Open IAM > Roles > Create role.
2. Choose AWS service as the trusted entity and EC2 as the use case.

3. Create or attach a permissions policy that grants only the Bedrock and/or S3 actions you need.
4. Give the role a clear name, for example EvidenceAnalystRuntimeRole.
5. Create the role.

7.2 Attach the role to a running instance

1. Open EC2 > Instances and select the Evidence Analyst instance.
2. Choose Actions > Security > Modify IAM role.
3. Select your IAM role / instance profile.
4. Choose Update IAM role.

AWS reference: Attach an IAM role to an EC2 instance

7.3 Example combined runtime policy

The following is a practical starting policy. Replace YOUR_BUCKET and YOUR_PREFIX. If you do not use S3 or Bedrock, remove the unneeded statement. Your organization may require a narrower policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "BedrockInference",
      "Effect": "Allow",
      "Action": [
        "bedrock:InvokeModel",
        "bedrock:InvokeModelWithResponseStream",
        "bedrock:GetInferenceProfile",
        "bedrock:ListInferenceProfiles"
      ],
      "Resource": "*"
    },
    {
      "Sid": "S3ReportDelivery",
      "Effect": "Allow",
      "Action": ["s3:PutObject"],
      "Resource": "arn:aws:s3:::YOUR_BUCKET/YOUR_PREFIX/*"
    }
  ]
}
```

If you want reports directly in the bucket root rather than under a prefix, use arn:aws:s3:::YOUR_BUCKET/* for the S3 resource. For production, restrict the policy to the exact bucket/prefix and Bedrock resources your organization approves.

BEDROCK CROSS-REGION NOTE

The validated default model is us.amazon.nova-2-lite-v1:0 in us-east-1. This is a US cross-Region inference profile. AWS Organizations SCPs or Region-deny policies can block one of the destination Regions and cause inference to fail even when the role contains Bedrock permissions. If Bedrock returns AccessDenied, check both IAM and your organization/SCP Region restrictions.

8. Amazon S3 Report Delivery

S3 Report Delivery is designed for unattended reviews and auto-stop. It sends only the finished DOCX report to your private S3 bucket before EC2 is allowed to stop. The source PDF is not copied to S3.

8.1 Create a private S3 bucket

1. Open Amazon S3 and choose Create bucket.
2. Create a General purpose bucket with a globally unique bucket name.
3. Choose the AWS Region you want to use. Using the same Region as the EC2 instance is usually simplest.
4. Keep Object Ownership at the default Bucket owner enforced setting unless your organization requires otherwise.
5. Keep all S3 Block Public Access settings enabled. Reports do not need public access.
6. Keep default encryption enabled. New S3 objects are encrypted with SSE-S3 by default unless you choose another encryption option.
7. Choose Create bucket.

AWS references: Create an S3 bucket | S3 default encryption

8.2 Grant the EC2 role permission to write reports

The tested minimum permission for ordinary SSE-S3 delivery is s3:PutObject to the report destination. Restrict it to the bucket and prefix you plan to use.

```
{
  "Effect": "Allow",
  "Action": "s3:PutObject",
  "Resource": "arn:aws:s3:::YOUR_BUCKET/evidence-reports/*"
}
```

CUSTOMER-MANAGED KMS

If the bucket requires a customer-managed AWS KMS key (SSE-KMS), the EC2 role can also need KMS permissions such as kms:Encrypt and kms:GenerateDataKey on that key, and the KMS key policy must allow the role.

8.3 Configure S3 in Advanced Workflows

1. Open Advanced Workflows.
2. Under Amazon S3 Report Delivery, select Enable S3 Report Delivery.
3. Enter the bucket name only, for example company-evidence-reports. Do not include s3:// in the bucket field.
4. Enter the AWS Region, for example us-east-1.
5. Optionally enter a prefix such as evidence-reports. A leading slash is not needed.
6. Choose Save S3 Settings.
7. Choose Test S3 Delivery and confirm the page reports a successful test.

Finished reports use a customer-friendly object name similar to:

```
evidence-reports/2026-09-08_1923_My_Document_Executive_Due_Diligence_a0254faa.docx
```

S3 delivery status appears in Review History. A delivered job shows the full s3:// URI.

The leading date and time in the report object name are UTC so completed reports remain useful when sorted lexically in Amazon S3.

8.4 Failure and retry behavior

If S3 is required for a job and delivery fails, the application intentionally prevents automatic EC2 shutdown. This protects you from losing browser access before the report reaches the configured S3 destination.

1. Leave the instance running.
2. Correct the bucket, Region, prefix, IAM role, bucket policy, or KMS permissions.
3. Use Retry S3 Delivery on the failed job in Review History.
4. After the retry succeeds, the original stop request is preserved and automatic shutdown can resume.

9. Optional AI providers and API credentials

AI IS OPTIONAL

Verified Evidence does not call any AI provider. Configure an AI provider only when you choose Verified Evidence + AI Analysis or an advanced workflow that needs AI.

Provider	Authentication	Where content is processed	Separate charges
AWS Bedrock	EC2 IAM role - no AWS access key stored by Code Creator.	AWS Bedrock using your AWS account and configured inference profile.	AWS Bedrock usage charges may apply.
OpenAI	Buyer-supplied OpenAI API key.	OpenAI API outside this EC2 instance.	OpenAI API charges may apply.
Anthropic	Buyer-supplied Anthropic API key.	Anthropic API outside this EC2 instance.	Anthropic API charges may apply.
Local Muse	No external API key. Uses local GPU inference.	On the EC2 instance.	Requires a compatible GPU EC2 instance; EC2 costs are higher.

9.1 AWS Bedrock

Bedrock is the default optional provider. The validated configuration is:

- AWS Region: us-east-1
- Model / inference profile: us.amazon.nova-2-lite-v1:0
- Authentication: EC2 IAM role

1. Attach an EC2 IAM role with Bedrock inference permissions as described in Section 7.
2. Open Advanced Workflows > AI Provider.
3. Select AWS Bedrock.
4. Confirm or enter the AWS Region and model/inference profile.
5. Choose Save AI Provider.
6. Run Verified Evidence + AI Analysis on a suitable document.

AWS references: Bedrock inference prerequisites | Amazon Nova 2 getting started

9.2 OpenAI

Use OpenAI only if your organization approves sending selected review/document content to the OpenAI API. You need an OpenAI API account with API billing and a secret API key.

1. Create or retrieve an API key from your OpenAI API project.
2. Open Advanced Workflows > AI Provider and select OpenAI API.

3. Enter the model and API key requested by the page.
4. Choose Save AI Provider.
5. Protect and rotate the key according to your organization policy.

The product treats provider secrets as write-only configuration: the saved secret is not returned in normal API responses. A saved-key field may therefore appear blank or masked on later visits. Enter a new key when you intentionally want to replace it.

OpenAI references: Find and manage OpenAI API keys | API key safety

9.3 Anthropic

Use Anthropic only if your organization approves sending selected review/document content to the Anthropic API. You need an Anthropic Console/API account, API billing, and an API key.

1. Create an API key in the Anthropic Console.
2. Open Advanced Workflows > AI Provider and select Anthropic API.
3. Enter the model and API key requested by the page.
4. Choose Save AI Provider.
5. Rotate or replace the key according to your organization policy.

10. Optional Local Muse GPU inference

Local Muse is optional private local inference. It is not started during normal CPU-first boot. Do not start it with raw Docker Compose commands; use the Code Creator helper so hardware and model checks run first.

```
sudo codecreator-local-ai-start
```

Local Muse item	Guidance
Validated GPU size	g6.4xlarge was validated successfully with NVIDIA L4 and approximately 60 GiB host RAM.
Not supported	g6.xlarge is not supported for this release because its host RAM was insufficient for acceptable loading behavior.
Other GPU sizes	Other compatible NVIDIA GPU instances may work, but Code Creator should not be assumed to have validated every size.
First local start	The helper warms the bundled model files from EBS, starts the local-gpu profile, and waits for health. This can take several minutes.
Output safety	If local AI output does not pass the product output gate, Verified Evidence remains valid and the optional AI synthesis may be omitted rather than returned.

EBS INITIALIZATION OPTION

Fresh volumes created from AMI snapshots can have higher first-read latency while snapshot blocks initialize. For predictable initialization, AWS offers Amazon EBS Provisioned Rate for Volume Initialization (100-300 MiB/s) at additional cost. This can be useful for buyers who plan to use the large Local Muse model immediately after launch.

To stop Local Muse without stopping the EC2 instance:

```
cd /opt/codecreator/muse-glimmer
sudo docker compose --profile local-gpu stop glimmer
```

AWS reference: Initialize Amazon EBS volumes

11. Optional MCP server integrations

MCP integrations are for advanced tool-enabled workflows. No MCP server is configured in the Marketplace AMI by default.

1. Open Advanced Workflows > MCP Servers.
2. Enter a clear server name.
3. Enter the Streamable HTTP MCP URL for a server you control or trust.
4. Choose Add Server.
5. Confirm the server shows the expected status before relying on tools in a review.

MCP SECURITY

An MCP tool can cause data or tool arguments to leave this EC2 instance depending on the server and tool being called. Connect only trusted MCP servers, review their authentication and data-handling policies, and restrict outbound/network access according to your organization policy.

12. Data persistence, security, and cost control

12.1 What persists across EC2 stop/start

- Uploaded documents on the EBS root volume.
- Generated Word reports.
- Review history and job state.
- Saved provider configuration and buyer-supplied provider secrets.
- S3 delivery configuration.
- Bundled Local Muse model files.
- The generated web administrator password.

The public IPv4 address normally does not persist unless you use an Elastic IP. The product refreshes the public-IP HTTPS configuration after restart.

12.2 Security practices

- Restrict SSH TCP 22 to your administrator IP or trusted CIDR whenever practical.
- Keep internal TCP 8090 closed to the public internet.
- Use HTTPS on TCP 443 for the web interface.
- Protect FIRST_LOGIN.txt and the administrator password.
- Use EC2 IAM roles instead of long-lived AWS access keys.
- Keep S3 Block Public Access enabled for report buckets.
- Use least-privilege IAM permissions for the exact S3 prefix and Bedrock functions required.
- Review organizational data-residency requirements before using cross-Region Bedrock inference or external AI providers.
- Take EBS snapshots or other backups if the reports/documents are important. Stopping preserves EBS; terminating may not.

Use encrypted EBS volumes and restrict snapshot sharing. Uploaded documents, generated reports, configuration, and job history can contain sensitive customer data and persist on EBS after the instance is stopped.

Reports delivered to Amazon S3 are managed in the customer bucket. Apply bucket encryption, least-privilege access, retention rules, and S3 Lifecycle policies that match your organization's requirements; manage deletion of delivered copies in S3 separately from local Review History.

12.3 Cost-control states

State	Meaning	Recommended action
ACTIVE REVIEWS > 0	Work is queued or running.	Do not stop EC2.
SAFE TO STOP	No active review and no required S3 delivery block.	You may stop EC2 from AWS or allow an auto-stop job to do so.

State	Meaning	Recommended action
S3 DELIVERY ATTENTION	A required report upload is pending or failed.	Correct S3/IAM/KMS settings and retry delivery before stopping.
EC2 Stopped	Compute charge is stopped, but EBS and applicable IP/storage charges can continue.	Start again when you need the GUI or local reports.

13. Server commands

Command	Purpose
cat ~/FIRST_LOGIN.txt	Show the full first-login summary including current URL and credentials.
codecreator-agent-url	Show the current application URL.
codecreator-agent-credentials	Show the application username and generated password.
codecreator-agent-status	Show application/platform status.
codecreator-agent-restart	Restart the application platform.
codecreator-agent-logs	View application logs.
codecreator-agent-help	Show product help and commands.
sudo codecreator-local-ai-start	Validate hardware, warm EBS model files, and start optional Local Muse.
cd /opt/codecreator/muse-glimmer && sudo docker compose --profile local-gpu stop glimmer	Stop optional Local Muse without stopping the EC2 instance. Use sudo codecreator-local-ai-start to start it safely.

14. Troubleshooting

Problem	What to check / do
Browser URL stopped working after EC2 restart	Check the new Public IPv4 address in EC2. Run codecreator-agent-url or cat ~/FIRST_LOGIN.txt. Auto-assigned public IPv4 addresses normally change after stop/start.
Want the address never to change	Allocate and associate an Elastic IP. AWS charges for public IPv4/Elastic IP resources.
SSH connection fails	Confirm the instance is running, TCP 22 allows your current source IP, you are using the private key selected at launch, and the username is ubuntu.
FIRST_LOGIN.txt is missing immediately after launch	Wait until both EC2 status checks pass, then run sudo systemctl status codecreator-agent-firstboot.service. First boot should finish with status 0/SUCCESS.
Browser returns 401	Authentication is working. Enter the generated admin credentials from FIRST_LOGIN.txt.
Verified Evidence works but AI Analysis fails	Verified Evidence does not need AI. For Bedrock, check IAM role/Region/model/SCPs. For OpenAI or Anthropic, check API key, billing, model, and outbound internet access.
Bedrock says AccessDenied	Confirm the EC2 IAM role is attached and allows model invocation. For cross-Region inference, also check AWS Organizations SCP/Region restrictions.
S3 Test Delivery fails with AccessDenied	Confirm s3:PutObject is allowed for the exact bucket/prefix. Check bucket policy, Region, and customer-managed KMS permissions if used.
Auto-stop did not occur after an S3 failure	This is intentional. Correct the S3 problem and choose Retry S3 Delivery. Automatic shutdown can resume after delivery succeeds.
EC2 stopped before report was downloaded	Start EC2 again. If no Elastic IP is used, obtain the new public IP and new HTTPS URL. The report remains on EBS.
Local Muse helper rejects the instance	Use a larger compatible NVIDIA GPU instance. g6.xlarge is not supported. g6.4xlarge was validated.
Local Muse takes time on first start	The bundled model is large and first reads may initialize EBS snapshot blocks. Wait for the helper to complete; consider EBS Provisioned Rate for Volume Initialization for predictable initialization.
SAFE TO STOP is not shown	Check for active jobs or an S3 delivery block. Do not force-stop until the page reports a safe state unless you accept interrupting work.
PDF uploads but findings are empty or incomplete	Confirm the PDF has a selectable text layer. Run OCR on image-only scans before upload, remove an authorized document password, and retry with an undamaged copy.

Problem	What to check / do
HTTPS is unavailable or the certificate is not ready after launch or an IP change	Wait for both EC2 status checks, run codecreator-agent-url, and use the displayed address. If needed, inspect sudo systemctl status codecreator-ip-cert-sync.service and confirm inbound TCP 80 and 443 plus outbound connectivity are allowed.

SUPPORT

For product support, use the Code Creator support contact shown on the AWS Marketplace product page. Include the instance Region, approximate time of the issue, the relevant Review History job ID, and output from codecreator-agent-status when possible. Do not send API keys, private SSH keys, or sensitive source documents in a support request unless explicitly required through an approved secure channel.

15. AWS and provider references

- Amazon EC2 - Attach an IAM role to an instance
- Amazon EC2 - Stop/start behavior and changing public IPv4 addresses
- Amazon EC2 - Allocate and associate Elastic IP addresses
- Amazon S3 - Create a general purpose bucket
- Amazon S3 - Default bucket encryption
- Amazon Bedrock - Prerequisites for model inference
- Amazon Nova - Getting started with Nova 2
- Amazon EBS - Initialize volumes and Provisioned Rate for Volume Initialization
- OpenAI - Find and manage API keys
- OpenAI - API key safety
- Anthropic - Claude Platform documentation

Private Document Evidence Analyst for AWS by Code Creator

This software assists evidence review and does not replace professional judgment, legal advice, audit opinions, or organizational approval processes.