

ALert AI - GenAI Application Security Product

sheet *overview of the features and services*

ALert AI - GenAI Application Security Product sheet overview of the features and services	1
Alert AI Services	2
Alert AI - Security & Posture Dashboards	3
Alert AI - Inline/Offline Gateway	4
Alert AI SEP (Security, Ethics, Performance) View	8
Alert AI LLM, Prompt/Response Detections	10
Alert AI Metric Anomaly Detections	12
Alert AI Log Anomaly Detections	13
Alert AI Integrations with GenAI & Data Platforms	15
Alert AI AI Province & LLM/Model Lineage Service	17
Alert AI AI & Data Pipeline Analytics	18
Alert AI AI & LLM/Model Tracking Analytics	19
Alert AI AI Cluster Resources & Ops Analytics	21





Alert AI Services

Alert AI

Runs

Model Versions

Models

Summary Reports

Databricks Resource

AI Provenance & Lineage

Training Data Analytics

AI Data Provenance & L

Alerts & Detections

LLM & Prompts

Anomaly detections

Metric Anomalies

Providers & Integrations

Risk Surface Analysis

AI Gateway services

AI Red Teaming

SEP Detections

AI Incidents



Risk Analysis - Taxonomy

srini.mommileti@alertai.com

Service	Status	Description
AI Dataset Provenance	Running	Maintain a detailed history of datasets used for AI applications. The history should include information about the dataset's source as well as complete record of any modifications.
AI Telemetry Logging	Running	Implement logging of inputs and outputs of deployed AI models. Monitoring logs can help to detect security threats and mitigate impacts.
AI BOM (Bill of Materials)	Running	An AI Bill of Materials (AI BOM) contains a full listing of artifacts and resources that were used in building the AI. The AI BOM can help mitigate supply chain risks and enable rapid response to reported vulnerabilities. This can include maintaining dataset provenance, i.e. a detailed history of datasets used for AI applications. The history can include information about the dataset source as well as a complete record of any modifications.
LLM/GenAI Guardrails	Running	Guardrails are safety controls that are placed between a generative AI model and the output shared with the user to prevent undesired inputs and outputs. Guardrails can take the form of validators such as filters, rule-based logic, or regular expressions, as well as AI-based approaches, such as classifiers and utilizing LLMs, or named entity recognition (NER) to evaluate the safety of the prompt or response. Domain specific methods can be employed to reduce risks in a variety of areas such as etiquette, brand damage, jailbreaking, false information, code exploits, SQL injections, and data leakage.
LLM Vulnerability Scans	Completed	Vulnerability scanning is used to find potentially exploitable software vulnerabilities to remediate them. File formats such as pickle files that are commonly used to store machine learning models can contain exploits that allow for arbitrary code execution. Both model artifacts and downstream products produced by models should be scanned for known vulnerabilities.
Model & Data Access Control	Running	Require users to verify their identities before accessing a production model. Require authentication for API endpoints and monitor production model queries to ensure compliance with usage policies and to prevent model misuse.
Inline Prompt Gateway	Running	Analyzes LLM inputs and detects adversarial user input attacks. Detects Harmful prompts, Sensitive information from responses.
Offline Prompt Gateway	Running	Security analytics, aggregation, Rules, Signals, Features. MELT data lake from LLM Apps, Agents, Platforms, Libraries, Tools, Services. Security, Ethics, Performance, LifeCycle Analytics. Detects Rogue models, Risky pipelines.




Alert AI - Security & Posture Dashboards







Alert AI - Inline/Offline Gateway

Prompt & Response (Blocked/Alerted/Allowed)







 **Dashboard**

 AI Posture

▼ AI Discovery

 **99+** AI cluster Resources

 ML Pipelines

▼ AI Tracking Analytics

 **9+** AI Traces

 **99+** Experiments

 Runs

 Model Versions

 Models



✓ Alerts & Detections

 **14+** LLM & Prompts

 **25+** Anomaly detections

 **6+** Metric Anomalies

 Providers & Integrations

✓ Risk Surface Analysis

 AI Gateway services

 AI Red Teaming

 SEP Detections

 AI Incidents

 Defenses



Alert AI SEP (Security, Ethics, Performance) View

ALERT AI

Runs

Model Versions

Models

Summary Reports

Databricks Resource

AI Provenance & Lineage

Training Data Analytics

AI Data Provenance & L

Alerts & Detections

LLM & Prompts

Anomaly detections

Metric Anomalies

Providers & Integrations

Risk Surface Analysis

AI Gateway services

AI Red Teaming

SEP Detections

AI Incidents

Defenses

Detections

srini.mommileti@alertai.com

Type	Category	
Security	LLM Jailbreak	DETECTIONS
Security	LLM Meta Prompt Extraction	DETECTIONS
Security	LLM Plugin Compromise	DETECTIONS
Security	LLM Prompt Injection	DETECTIONS
Security	LLM Data Leakage	DETECTIONS
Security	ML Supply Chain Compromise	DETECTIONS
Security	Model Probing	DETECTIONS
Security	Exfiltration Via Inference	DETECTIONS
Security	Denial of service	DETECTIONS
Security	Rogue/Poison Training Data	DETECTIONS
Security	Rogue/Poison Datasets	DETECTIONS
Security	Rogue/Poison Models	DETECTIONS
Security	Unauthorized or Excessive ML Public Facing Application Access	DETECTIONS



ALERT AI				Detections		srimi.mommileti@alertai.com	
Runs	Security	Unauthorized or Excessive AI Model Inference API Access	DETECTIONS				
Model Versions	Security	LLM Plugin Compromise	DETECTIONS				
Models	Security	LLM Prompt Self replication: RAG-Based Attack	DETECTIONS				
Summary Reports	Security	Repeated Token Attack	DETECTIONS				
Databricks Resource	Security	Erode LLM App Integrity	DETECTIONS				
AI Provenience & Lineage	Security	Erode ML Model Integrity	DETECTIONS				
Training Data Analytics	Security	Model Supply Chain Compromise	DETECTIONS				
AI Data Provenience & L	Security	LLM Vulnerabilities	DETECTIONS				
Alerts & Detections	Security	Over-Permissive API- Information Leak	DETECTIONS				
LLM & Prompts	Security	Over-Permissive API- Excessive Queries	DETECTIONS				
Anomaly detections	Security	Model Bypass- Bad Features	DETECTIONS				
Metric Anomalies	Security	Model Bypass- Insufficient Training Data	DETECTIONS				
Providers & Integrations	Security	Model Bypass- Adversarial	DETECTIONS				
Risk Surface Analysis	Security	Exfiltration- Model inversion	DETECTIONS				
AI Gateway services	Security	Exfiltration- Model theft	DETECTIONS				
AI Red Teaming							
SEP Detections							
AI Incidents							
Defenses							

ALERT AI				Detections		srimi.mommileti@alertai.com	
Dashboard	Security	Data Poisoning Ingest Poisoning	DETECTIONS				
AI Posture	Ethics	Bias/Discrimination- Group fairness	DETECTIONS				
AI Discovery	Ethics	Bias/Discrimination- Individual fairness	DETECTIONS				
AI cluster Resources	Ethics	Explainability- Global explanations	DETECTIONS				
ML Pipelines	Ethics	Explainability- Local explanations	DETECTIONS				
AI Tracking Analytics	Ethics	User actions- Toxicity	DETECTIONS				
AI Traces	Ethics	User actions- Polarization/Exclusion	DETECTIONS				
Experiments	Ethics	Misinformation- Deliberative misinformation	DETECTIONS				
Runs	Ethics	Misinformation- Generative misinformation	DETECTIONS				
Model Versions	Performance	Data issues- Data drift	DETECTIONS				
Models	Performance	Data issues- Concept drift	DETECTIONS				
Summary Reports	Performance	Data issues- Data entanglement	DETECTIONS				
Databricks Resource	Performance	Data issues- Data quality issues	DETECTIONS				
AI Provenience & Lineage	Performance	Data issues- Feedback loops	DETECTIONS				
Training Data Analytics	Performance	Model issues- Resilience/stability	DETECTIONS				
AI Data Provenience & L							
Alerts & Detections							
LLM & Prompts							
Anomaly detections							



Alert AI LLM, Prompt/Response Detections

ALERT AI

Dashboard

AI Posture

AI Discovery

AI cluster Resources

ML Pipelines

AI Tracking Analytics

AI Traces

Experiments

Runs

Model Versions

Models

Summary Reports

Databricks Resource

AI Provenance & Lineage

Training Data Analytics

AI Data Provenance & L

Alerts & Detections

LLM & Prompts

Anomaly detections

Model & LLM Risks

2/21/2025, 2:00:02 AM

LLM application

alertai-mitre

AlertAI-Mitre-Atlas-Protection

Malware

Indirect Document Attack

execution via malicious link

2/21/2025, 2:00:02 AM

RAG application

staging-alertai-mitre

staging-AlertAI-Mitre-Atlas-Protection

RAG-Based Attack

LLM Data leakage

LLM Prompt Self replication

2/21/2025, 2:00:02 AM

LLM application

staging-alertai-mitre

staging-AlertAI-Mitre-Atlas-Protection

Attempt to change system rules

LLM Prompt Indirect Document Attack

Attempt to instruct to ignore previous turns

2/21/2025, 2:00:02 AM

LLM application

staging-alertai-mitre

staging-AlertAI-Mitre-Atlas-Protection

LLM Prompt Injection: Indirect

LLM Meta Prompt Extraction

Direct LLM to access malicious website

2/21/2025, 2:00:02 AM

LLM application

staging-alertai-mitre

staging-AlertAI-Mitre-Atlas-Protection

Fraud

LLM Prompt Indirect Document Attack

Attempt to acting on behalf of the user without authorization

2/21/2025, 2:00:02 AM

LLM application

staging-alertai-mitre

staging-AlertAI-Mitre-Atlas-Protection

Malware

LLM Prompt Indirect Document Attack

Attempt to Code execution via Email

2/21/2025, 2:00:02 AM

LLM application

staging-alertai-mitre

staging-AlertAI-Mitre-Atlas-Protection

Information Gathering

LLM Prompt Indirect Document Attack

Attempt to Unauthorized data exfiltration

2/21/2025, 2:00:02 AM

LLM application

staging-alertai-mitre

staging-AlertAI-Mitre-Atlas-Protection

Intrusion

LLM Prompt Indirect Document

Attempt to embed grounded data by

Dashboard

AI Posture

AI Discovery

AI cluster Resources

ML Pipelines

AI Tracking Analytics

AI Traces

Experiments

Runs

Model Versions

Models

Summary Reports

Databricks Resource

AI Provenance & Lineage

Training Data Analytics

AI Data Provenance & L

Alerts & Detections

LLM & Prompts

Anomaly detections

Last seen

Resource

Provider

Source

Category

Type

Cause

2/20/2025, 4:00:02 PM

AI Assistant

staging-azureai-prompt-shield

staging-AlertAI-azureai-prompt-shield

Role-Play Attack

Jailbreak User Prompt attacks

Instructing the AI assistant to act as another system persona

2/19/2025, 4:00:02 PM

llm

staging-azureai-cfs

staging-azureai-prompt-shield

Attempt to change system rules

Jailbreak User Prompt attacks

Attempt to change system rules

4/23/2024, 5:00:00 PM

model

staging-miflow-eks

staging-miflow-eks

Performance Bias

Vulnerabilities

Biased Training Data

4/23/2024, 5:00:00 PM

model

staging-miflow-eks

staging-miflow-eks

Unethical behaviour

Vulnerabilities

Ethical Guidelines and Standards

4/23/2024, 5:00:00 PM

model

staging-miflow-eks

staging-miflow-eks

Performance Bias

Vulnerabilities

Sampling Bias

4/23/2024, 5:00:00 PM

model

staging-miflow-eks

staging-miflow-eks

Performance Bias

Vulnerabilities

Inadequate Feature Representation

4/23/2024, 5:00:00 PM

model

staging-miflow-eks

staging-miflow-eks

Data Leakage

Vulnerabilities

Inclusion of Future Info

4/23/2024, 5:00:00 PM

model

staging-miflow-eks

staging-miflow-eks

Unethical behaviour

Vulnerabilities

Interdisciplinary Collaboration

Rows per page: 25

1-25 of 54

1 2 3



ALERT AI

Dashboard

AI Posture

AI Discovery

99+

AI cluster Resources

ML Pipelines

AI Tracking Analytics

9+

AI Traces

99+

Experiments

Runs

Model Versions

Models

Summary Reports

Databricks Resource

AI Provenance & Lineage

Training Data Analytics

AI Data Provenance & L

Alerts & Detections

14+

LLM & Prompts

25+

Anomaly detections

Model & LLM Risks

Search

ADD FILTER

EXPORT

☐	Last seen	Resource	Provider	Source	Category	Type	Cause
☐	2/21/2025, 2:00:02 AM	LLM application	staging-alertai-mitre	staging-AlertAI-Mitre-Atlas-Protection	Attempt to change system rules	LLM Prompt Indirect Document Attack	Attempt to instruct to disregard rules
☐	2/21/2025, 2:00:02 AM	LLM application	staging-alertai-mitre	staging-AlertAI-Mitre-Atlas-Protection	Information Gathering	LLM Prompt Indirect Document Attack	Attempt to embed grounded data with instruction to information gathering
☐	2/21/2025, 2:00:02 AM	LLM application	staging-alertai-mitre	staging-AlertAI-Mitre-Atlas-Protection	Information Gathering	LLM Prompt Indirect Document Attack	Attempt to Data removal from a system
☐	2/21/2025, 2:00:02 AM	RAG application	staging-alertai-mitre	staging-AlertAI-Mitre-Atlas-Protection	Protected material text Detection	Detect protected material	Copyrighted content detected
☐	2/21/2025, 2:00:02 AM	LLM application	staging-alertai-mitre	staging-AlertAI-Mitre-Atlas-Protection	Manipulated Content	LLM Prompt Indirect Document Attack	Attempt to embed hidden instructions
☐	2/21/2025, 2:00:02 AM	RAG application	staging-alertai-mitre	staging-AlertAI-Mitre-Atlas-Protection	Ungroundedness Detection	LLM Prompt Indirect Document Attack	Fabricated or misleading information in Summarization task



Alert AI Metric Anomaly Detections

ALERT AI

99+

AI cluster Resources

ML Pipelines

AI Tracking Analytics

9+

AI Traces

99+

Experiments

Runs

Model Versions

Models

Summary Reports

Databricks Resource

AI Province & Lineage

Training Data Analytics

AI Data Province & L

Alerts & Detections

14+

LLM & Prompts

25+

Anomaly detections

6+

Metric Anomalies

Providers & Integrations

Metric Anomalies

Search

ADD FILTEREXPORT

Timestamp	Metric	Current value	Forecasted value	Anomaly
5/28/2024, 5:00:00 PM	kube_pod_status_phase	0	0	No
5/28/2024, 5:00:00 PM	kube_node_status_allocatable	0	0	No
5/28/2024, 5:00:00 PM	kube_persistentvolume_status_phase	0	0	No
5/28/2024, 5:00:00 PM	kube_pod_status_reason	0	0	No
5/28/2024, 5:00:00 PM	instance.node.network.transmit_bytes.rate.sum	252,071.57	240,125.104	Yes
5/28/2024, 5:00:00 PM	kube_pod_status_qos_class	0	0	No
5/28/2024, 5:00:00 PM	kube_persistentvolume_status_phase	0	0	No
5/28/2024, 5:00:00 PM	kube_pod_status_phase	0	0	No
5/28/2024, 5:00:00 PM	kube_endpoint_ports	1	1	No
5/28/2024, 5:00:00 PM	kube_lease_owner	1	1	No
5/28/2024, 5:00:00 PM	kube_pod_status_reason	0	0	No
5/28/2024, 5:00:00 PM	instance.node.cpu.ratio	0.02	0.047	Yes
5/28/2024, 5:00:00 PM	kube_service_info	1	1	No
5/28/2024, 5:00:00 PM	kube_pod_status_qos_class	0	0	No
5/28/2024, 5:00:00 PM	kube_service_info	1	1	No
5/28/2024, 5:00:00 PM	kube_persistentvolumeclaim_created	1,708,118,363	1,708,118,363	No
5/28/2024, 5:00:00 PM	kube_persistentvolumeclaim_created	1,708,118,363	1,708,118,363	No
5/28/2024, 5:00:00 PM	kube_node_status_allocatable	0	0	No

ALERT AI

kube_pod_status_reason

srini.mommileti@alertai.com

99+

AI cluster Resources

ML Pipelines

AI Tracking Analytics

9+

AI Traces

99+

Experiments

Runs

Model Versions

Models

Summary Reports

Databricks Resource

AI Province & Lineage

Training Data Analytics

AI Data Province & L

Alerts & Detections

14+

LLM & Prompts

25+

Anomaly detections

6+

Metric Anomalies

Providers & Integrations

Last seen

5/28/2024

Labels

container = kube-state-metrics

prometheus_replica = prometheus-kube-prometheus-stack-prometheus-0reason = NodeLost

uid = 2064d1cf-8781-4ce9-a231-f7c708746e9eendpoint = httpinstance = 10.1.24.45:8080

pod = cluster-proportional-autoscaler-kube-dns-autoscaler-5d8dcctlqw

__name__ = kube_pod_status_reasonservice = kube-prometheus-stack-kube-state-metrics

namespace = kube-systemprometheus = kube-prometheus-stack/kube-prometheus-stack-prometheus

job = kube-state-metrics

Current value

0

Forecasted value

0

Forecasted value Lower

-1.223105849624286e-9

Forecasted value Upper

1.25067517905238e-9

Error

0

Uncertainty

0

Details

First Seen5/28/2024, 5:00 PM

Tuesday, May 28, 2024, 5:00 PM

Predicted value: 0

Tuesday, May 28, 2024, 5:00 PM

Current value: 0

Anomaly: No



Alert AI Log Anomaly Detections

ALERT AI

99+

99+

14+

25+

6+

SAVED QUERIES

LAST VISITED

Today

This week

Last week

This month

Last month

Earlier

IS ANOMALY

Yes

No

Anomaly detections

srini.mommileti@alertai.com

ADD FILTEREXPORT

☐	Last seen	Component	Resource	resources.detections.fields.anomaly_score	Anomaly
☐	5/16/2024, 5:00:00 PM	aws.eks.control-plane-logs	s3://staging-spark-spark-logs-20240131220343017000000001/spark-event-logs/eventlog_v2_spark-20c87193f3f44696aa615eb8cfa8feb5/events_3180_spark-20c87193f3f44696aa615eb8cfa8feb5	0	
☐	5/16/2024, 5:00:00 PM	aws.eks.control-plane-logs	s3://staging-spark-spark-logs-20240131220343017000000001/spark-event-logs/eventlog_v2_spark-20c87193f3f44696aa615eb8cfa8feb5/events_3180_spark-20c87193f3f44696aa615eb8cfa8feb5	0	
☐	5/16/2024, 5:00:00 PM	aws.eks.control-plane-logs	s3://staging-spark-spark-logs-20240131220343017000000001/spark-event-logs/eventlog_v2_spark-20c87193f3f44696aa615eb8cfa8feb5/events_3180_spark-20c87193f3f44696aa615eb8cfa8feb5	0	
☐	5/16/2024, 5:00:00 PM	aws.eks.control-plane-logs	s3://staging-spark-spark-logs-20240131220343017000000001/spark-event-logs/eventlog_v2_spark-20c87193f3f44696aa615eb8cfa8feb5/events_3180_spark-20c87193f3f44696aa615eb8cfa8feb5	0	
☐	5/16/2024, 5:00:00 PM	aws.eks.control-plane-logs	s3://staging-spark-spark-logs-20240131220343017000000001/spark-event-logs/eventlog_v2_spark-20c87193f3f44696aa615eb8cfa8feb5/events_3180_spark-20c87193f3f44696aa615eb8cfa8feb5	0	



AI Tracking Analytics 9+ AI Traces 99+ Experiments Runs Model Versions Models Summary Reports Databricks Resource AI Provenance & Lineage Training Data Analytics AI Data Provenance & L Alerts & Detections 14+ LLM & Prompts 25+ Anomaly detections 6+ Metric Anomalies Providers & Integrations Risk Surface Analysis AI Gateway services AI Red Teaming	☐	Last seen	Component	Resource	resources.detections.fields.anom
	☐	5:00:00 PM	aws.eks.control-plane-logs	logs/eventlog_v2_spark-20c87193f3f44696aa615eb8cfa8feb5/events_3180_spark-20c87193f3f44696aa615eb8cfa8feb5	
	☐	5/16/2024, 5:00:00 PM	aws.eks.control-plane-logs	s3://staging-spark-spark-logs-20240131220343017000000001/spark-event-logs/eventlog_v2_spark-20c87193f3f44696aa615eb8cfa8feb5/events_3180_spark-20c87193f3f44696aa615eb8cfa8feb5	
	☐	5/16/2024, 5:00:00 PM	aws.eks.control-plane-logs	s3://staging-spark-spark-logs-20240131220343017000000001/spark-event-logs/eventlog_v2_spark-20c87193f3f44696aa615eb8cfa8feb5/events_3180_spark-20c87193f3f44696aa615eb8cfa8feb5	
	☐	5/16/2024, 5:00:00 PM	aws.eks.control-plane-logs	s3://staging-spark-spark-logs-20240131220343017000000001/spark-event-logs/eventlog_v2_spark-20c87193f3f44696aa615eb8cfa8feb5/events_3180_spark-20c87193f3f44696aa615eb8cfa8feb5	
	☐	5/16/2024, 5:00:00 PM	aws.eks.control-plane-logs	s3://staging-spark-spark-logs-20240131220343017000000001/spark-event-logs/eventlog_v2_spark-20c87193f3f44696aa615eb8cfa8feb5/events_3180_spark-20c87193f3f44696aa615eb8cfa8feb5	

Rows per page: 25 1-25 of 4065159 < 1 2 3

ALERT AI

s3://staging-spark-spark-logs-202401312203430

sriini.mommileti@alertai.com

AI Tracking Analytics

9+ AI Traces

99+ Experiments

Runs

Model Versions

Models

Summary Reports

Databricks Resource

AI Provenance & Lineage

Training Data Analytics

AI Data Provenance & L

Alerts & Detections

14+ LLM & Prompts

25+ Anomaly detections

6+ Metric Anomalies

Providers & Integrations

Risk Surface Analysis

AI Gateway services

AI Red Teaming

Component

aws.eks.control-plane-logs

Resource

s3://staging-spark-spark-logs-20240131220343017000000001/spark-event-logs/eventlog_v2_spark-20c87193f3f44696aa615eb8cfa8feb5/events_3180_spark-20c87193f3f44696aa615eb8cfa8feb5

Last seen

5/16/2024

resources.detections.fields.anomaly_score

Anomaly

0

Log message

{\"Event\":\"SparkListenerTaskEnd\",\"Stage ID\":\"137125\",\"Stage Attempt ID\":\"0\",\"Task Type\":\"ResultTask\",\"Task End Reason\":{\"Reason\":\"Success\"},\"Task Info\":{\"Task ID\":\"34349821\",\"Index\":\"245\",\"Attempt\":\"0\",\"Partition ID\":\"245\",\"Launch Time\":\"1715902712255\",\"Executor ID\":\"30\",\"Host\":\"100.64.173.49\",\"Locality\":\"PROCESS_LC Result Time\":\"0\",\"Finish Time\":\"1715902712260\",\"Failed\":false,\"Killed\":false,\"Accot\":{\"ID\":\"9941731\",\"Name\":\"duration\",\"Update\":\"1\",\"Value\":\"2 Failed Values\":true,\"Metadata\":\"sql\"}, {\"ID\":\"9941732\",\"Name\":\"sort time\",\"Update\":\"0\",\"Value\":\"0\",\"Internal\":true,\"Count

Details

Last seen

5/16/2024

Last seen

5/16/2024

Thursday, May 16, 2024, 5:00 PM

aws.eks.control-plane-logs

Thursday, May 16, 2024, 5:00 PM

s3://staging-spark-spark-logs-20240131220343017000000001/spark-event-logs/eventlog_v2_spark-20c87193f3f44696aa615eb8cfa8feb5/events_3180_s20c87193f3f44696aa615eb8cfa8feb5(Event\":\"SparkLID\":\"137125\",\"Stage Attempt ID\":\"0\",\"Task Type\":\"ResultTask\",\"Task End Reason\":{\"Reason\":\"Success\"},\"Task Info\":{\"Task ID\":\"34349821\",\"Index\":\"245\",\"Attempt\":\"0\",\"Partition ID\":\"245\",\"Launch Time\":\"1715902712255\",\"Executor ID\":\"30\",\"Host\":\"100.64.173.49\",\"Locality\":\"PROCESS_LC Result Time\":\"0\",\"Finish Time\":\"1715902712260\",\"Failed\":false,\"Killed\":false,\"Accot\":{\"ID\":\"9941731\",\"Name\":\"duration\",\"Update\":\"1\",\"Value\":\"2 Failed Values\":true,\"Metadata\":\"sql\"}, {\"ID\":\"9941732\",\"Name\":\"sort time\",\"Update\":\"0\",\"Value\":\"0\",\"Internal\":true,\"Count



Alert AI Integrations with GenAI & Data Platforms

ALERT AI Providers srini.mommileti@alertai.com

Providers	Integration	Edit
Aws	INTEGRATION	EDIT
Airflow	INTEGRATION	EDIT
Anthropic	INTEGRATION	EDIT
AutoGen	INTEGRATION	EDIT
Azure	INTEGRATION	EDIT
Azureml	INTEGRATION	EDIT
Beam	INTEGRATION	EDIT
Bedrock	INTEGRATION	EDIT
Cohere	INTEGRATION	EDIT
Crewai	INTEGRATION	EDIT
Databricks	INTEGRATION	EDIT
Dataflow	INTEGRATION	EDIT

DSPy	Gcp	Googlegemini	Groq
INTEGRATION	INTEGRATION	INTEGRATION	INTEGRATION
Huggingface	Instructor	Jupyter	Kafka
INTEGRATION	INTEGRATION	INTEGRATION	INTEGRATION
Kubeflow	Langchain	Langgraph	LiteLLM
INTEGRATION	INTEGRATION	INTEGRATION	INTEGRATION



AI Posture

AI Discovery

99+ AI cluster Resources

ML Pipelines

AI Tracking Analytics

AI Provenance & Lineage

Training Data Analytics

AI Data Provenance & L

Alerts & Detections

14+ LLM & Prompts

25+ Anomaly detections

6+ Metric Anomalies

Providers & Integrations

Risk Surface Analysis

AI Gateway services

AI Red Teaming

SEP Detections

AI Incidents

Defenses

LlamaIndex

Llamaindex

INTEGRATION EDIT

MISTRAL AI_

Mistralai

INTEGRATION EDIT

mlflow

Mlflow

INTEGRATION EDIT

mongoDB.

Mongo

INTEGRATION EDIT

thumbnail

Mosaicml

INTEGRATION EDIT

thumbnail

Nvidia

INTEGRATION EDIT

thumbnail

Ollama

INTEGRATION EDIT

OpenAI

Openai

INTEGRATION EDIT

Swarm

Openaiswarm

INTEGRATION EDIT

PaLM 2

Palm

INTEGRATION EDIT

RAY

Ray

INTEGRATION EDIT

Amazon Redshift

Redshift

INTEGRATION EDIT

AI Posture

AI Discovery

99+ AI cluster Resources

ML Pipelines

AI Tracking Analytics

AI Provenance & Lineage

Training Data Analytics

AI Data Provenance & L

Alerts & Detections

14+ LLM & Prompts

25+ Anomaly detections

6+ Metric Anomalies

Providers & Integrations

Risk Surface Analysis

AI Gateway services

AI Red Teaming

SEP Detections

AI Incidents

Defenses

Swarm

Openaiswarm

INTEGRATION EDIT

PaLM 2

Palm

INTEGRATION EDIT

RAY

Ray

INTEGRATION EDIT

Amazon Redshift

Redshift

INTEGRATION EDIT

S3

S3

INTEGRATION EDIT

Amazon SageMaker

Sagemaker

INTEGRATION EDIT

APACHE Spark

Spark

INTEGRATION EDIT

AI21labs

Ai21labs

INTEGRATION EDIT

HUGGING FACE

Huggingface



ALERT AI
AI Data Province & Lineage anal
🔄 🧑 srini.mommileti@alertai.com

- Runs
- Model Versions
- Models
- Summary Reports
- Databricks Resource
- AI Province & Lineage
- Training Data Analytics
- AI Data Province & Lineage**
- Alerts & Detections
 - LLM & Prompts
 - Anomaly detections
 - Metric Anomalies
- Providers & Integrations
- Risk Surface Analysis
 - AI Gateway services
 - AI Red Teaming
 - SEP Detections
 - AI Incidents

☐	Timestamp ↓	Spark Application	Provider	Data source	Data sink
☐	5/28/2024, 5:00:00 PM	test-sparkapp	staging-spark-cluster	KafkaV2[Subscribe(metrics-kube-prometheus-stack)]	org.apache.spark.sql.kafka010.KafkaSourceProviderSK
☐	5/28/2024, 5:00:00 PM	test-sparkapp	staging-spark-cluster	KafkaV2[Subscribe(metrics-kube-prometheus-stack)]	org.apache.spark.sql.kafka010.KafkaSourceProviderSK
☐	5/28/2024, 5:00:00 PM	test-sparkapp	staging-spark-cluster	KafkaV2[Subscribe(metrics-kube-prometheus-stack)]	org.apache.spark.sql.kafka010.KafkaSourceProviderSK
☐	5/28/2024, 5:00:00 PM	test-sparkapp	staging-spark-cluster	KafkaV2[Subscribe(metrics-kube-prometheus-stack)]	org.apache.spark.sql.kafka010.KafkaSourceProviderSK
☐	5/28/2024, 5:00:00 PM	test-sparkapp	staging-spark-cluster	KafkaV2[Subscribe(metrics-kube-prometheus-stack)]	org.apache.spark.sql.kafka010.KafkaSourceProviderSK
☐	5/28/2024, 5:00:00 PM	test-sparkapp	staging-spark-cluster	KafkaV2[Subscribe(metrics-kube-prometheus-stack)]	org.apache.spark.sql.kafka010.KafkaSourceProviderSK
☐	5/28/2024, 5:00:00 PM	test-sparkapp	staging-spark-cluster	KafkaV2[Subscribe(metrics-kube-prometheus-stack)]	org.apache.spark.sql.kafka010.KafkaSourceProviderSK

ALERT AI
test-sparkapp
🔄 🧑 srini.mommileti@alertai.com

- Runs
- Model Versions
- Models
- Summary Reports
- Databricks Resource
- AI Province & Lineage
- Training Data Analytics
- AI Data Province & Lineage**
- Alerts & Detections
 - LLM & Prompts
 - Anomaly detections
 - Metric Anomalies
- Providers & Integrations
- Risk Surface Analysis
 - AI Gateway services
 - AI Red Teaming
 - SEP Detections
 - AI Incidents

Last seen

5/28/2024

Spark Application

test-sparkapp

Provider

staging-spark-cluster

Details

⌚ First seen
5/28/2024
⌚ Last seen
5/28/2024

🔍 Tuesday, May 28, 2024, 5:00 PM

Dataset Source

KafkaV2[Subscribe(metrics-kube-prometheus-stack)]

Topic and Offset

[object Object]

📄 Tuesday, May 28, 2024, 5:00 PM

Dataset Sink

org.apache.spark.sql.kafka010.KafkaSourceProviderSK



Alert AI AI & Data Pipeline Analytics

ALERT AI

Dashboard

AI Posture

AI Discovery

AI cluster Resources

ML Pipelines

AI Tracking Analytics

AI Traces

Experiments

Runs

Model Versions

Models

Summary Reports

Databricks Resource

AI Province & Lineage

Training Data Analytics

AI Data Province & L

Alerts & Detections

LLM & Prompts

Anomaly detections

Pipeline analysis

Search

ADD FILTEREXPORT

Timestamp	Spark Event log	Provider
4/23/2024, 5:00:00 PM	s3://staging-spark-spark-logs-20240131220343017000000001/spark-event-logs/spark-649cc8a3dfc34421990626e09e559abf	aws.eks.control-plane-logs
4/23/2024, 5:00:00 PM	s3://staging-spark-spark-logs-20240131220343017000000001/spark-event-logs/spark-aac6fe6bfc494aacbd034eda263faa62	aws.eks.control-plane-logs
4/23/2024, 5:00:00 PM	s3://staging-spark-spark-logs-20240131220343017000000001/spark-event-logs/spark-aac6fe6bfc494aacbd034eda263faa62	aws.eks.control-plane-logs
4/23/2024, 5:00:00 PM	s3://staging-spark-spark-logs-20240131220343017000000001/spark-event-logs/spark-aac6fe6bfc494aacbd034eda263faa62	aws.eks.control-plane-logs
4/23/2024, 5:00:00 PM	s3://staging-spark-spark-logs-20240131220343017000000001/spark-event-logs/spark-1132829058194434bc20824a95a0d476	aws.eks.control-plane-logs
4/23/2024, 5:00:00 PM	s3://staging-spark-spark-logs-20240131220343017000000001/spark-event-logs/spark-1132829058194434bc20824a95a0d476	aws.eks.control-plane-logs
4/23/2024, 5:00:00 PM	s3://staging-spark-spark-logs-20240131220343017000000001/spark-event-logs/spark-aac6fe6bfc494aacbd034eda263faa62	aws.eks.control-plane-logs
4/23/2024, 5:00:00 PM	s3://staging-spark-spark-logs-20240131220343017000000001/spark-event-logs/spark-649cc8a3dfc34421990626e09e559abf	aws.eks.control-plane-logs
4/23/2024, 5:00:00 PM	s3://staging-spark-spark-logs-20240131220343017000000001/spark-event-logs/spark-1132829058194434bc20824a95a0d476	aws.eks.control-plane-logs
4/23/2024, 5:00:00 PM	s3://staging-spark-spark-logs-20240131220343017000000001/spark-event-logs/spark-649cc8a3dfc34421990626e09e559abf	aws.eks.control-plane-logs

ALERT AI

Dashboard

AI Posture

AI Discovery

AI cluster Resources

ML Pipelines

AI Tracking Analytics

AI Traces

Experiments

Runs

Model Versions

Models

Summary Reports

Databricks Resource

AI Province & Lineage

Training Data Analytics

AI Data Province & L

Alerts & Detections

LLM & Prompts

Anomaly detections

s3://staging-spark-spark-logs-202401312203430

srini.mommileti@alertai.cor

Last seen

4/23/2024

Spark Event log

s3://staging-spark-spark-logs-20240131220343017000000001/spark-event-logs/spark-aac6fe6bfc494aacbd034eda263faa62

Provider

aws.eks.control-plane-logs

Analysis

TaskEnd	80
ExecutorRemoved	0
JobStart	0
ExecutorAdded	0
QueryEnded	0
LogStart	0
QueryProgress	0



Alert AI AI & LLM/Model Tracking Analytics

Alert AI

ML Pipelines

AI Tracking Analytics

AI Traces

Experiments

Runs

Model Versions

Models

Summary Reports

Databricks Resource

AI Provenance & Lineage

Training Data Analytics

AI Data Provenance & L

Alerts & Detections

LLM & Prompts

Anomaly detections

Metric Anomalies

Providers & Integrations

Risk Surface Analysis

Experiments

mm/dd/yyyy

mm/dd/yyyy

EXPORT

	Last seen ↓	Provider	Resource Type
>	3/9/2025	mlflow-tracking-client.resource.experiment	mlflow.cluster1.host
>	3/9/2025	mlflow-tracking-client.resource.experiment	mlflow.cluster1.host

Experiment

mlflow-tracking-client.resource.experiment

Last seen

3/9/2025

Provider

mlflow.cluster1.host

Details

experiment_name	Default
lifecycle_stage	active
experiment_id	0
artifact_location	mlflow-artifacts:/0

Alert AI

ML Pipelines

AI Tracking Analytics

AI Traces

Experiments

Runs

Model Versions

Models

Summary Reports

Databricks Resource

AI Provenance & Lineage

Training Data Analytics

AI Data Provenance & L

Alerts & Detections

LLM & Prompts

Anomaly detections

Metric Anomalies

Providers & Integrations

Risk Surface Analysis

AI Gateway services

Tracking - Databricks

Alerts

Jobs

Instance Pools

Job Runs Report

Experiments

Node Pools

Catalogs

Models

Data Sources

DBC Resources

Clusters

Compute submit runs

Pipelines

Runs



▼ AI Tracking Analytics

-  9+ AI Traces
-  99+ Experiments
-  Runs
-  Model Versions
-  Models
-  Summary Reports

 **Databricks Resource**

▼ AI Provenance & Lineage

-  Training Data Analytics
-  AI Data Provenance & L



Alert AI AI Cluster Resources & Ops Analytics

ALERT AI

Dashboard

AI Posture

AI Discovery

AI cluster Resources

ML Pipelines

AI Tracking Analytics

AI Traces

Experiments

Runs

Model Versions

Models

Summary Reports

Databricks Resource

AI Province & Lineage

Training Data Analytics

AI Data Province & L

Alerts & Detections

LLM & Prompts

Anomaly detections

Cluster Resources

Lastseen gte

mm/dd/yyyy

Lastseen lte

mm/dd/yyyy

EXPORT

		Last seen ↓		Provider	
>	☐	4/23/2024	node	ip-10-1-23-107.us-east-2.compute.internal	kubernetes
>	☐	4/23/2024	pod	kube-prometheus-stack-kube-state-metrics-65cfbc8b9-9jp7k	kubernetes
>	☐	4/23/2024	endpoint	cluster-kafka-brokers	kubernetes
>	☐	4/23/2024	daemonset	ebs-csi-node	kubernetes
>	☐	4/23/2024	endpoint	cluster-cruise-control	kubernetes
>	☐	4/23/2024	endpoint	cluster-kafka-external4-1	kubernetes
>	☐	4/23/2024	pod	aws-node-wqd67	kubernetes
>	☐	4/23/2024	pod	coredns-688c955fbc-b2f5h	kubernetes
>	☐	4/23/2024	networkpolicy	staging-mongodb	kubernetes
>	☐	4/23/2024	pod	strimzi-cluster-operator-79b6bc479b-7lm9q	kubernetes
>	☐	4/23/2024	endpoint	kube-prometheus-stack-operator	kubernetes
>	☐	4/23/2024	persistentvolumeclaim	data-prometheus-kube-prometheus-stack-prometheus-0	kubernetes
>	☐	4/23/2024	networkpolicy	cluster-network-policy-kafka	kubernetes
>	☐	4/23/2024	configmap	kube-prometheus-stack-grafana	kubernetes
>	☐	4/23/2024	configmap	kube-root-ca.crt	kubernetes
>	☐	4/23/2024	persistentvolume	pvc-377bd36f-d021-487e-9944-1033998a7484	kubernetes
>	☐	4/23/2024	endpoint	kube-prometheus-stack-coredns	kubernetes

https://x1-staging.alertai.com/#/kubernetes

ALERT AI

Dashboard

AI Posture

AI Discovery

99% AI cluster Resources

ML Pipelines

AI Tracking Analytics

9% AI Traces

99% Experiments

Runs

Model Versions

Models

Summary Reports

Databricks Resource

AI Provenance & Lineage

Training Data Analytics

AI Data Provenance & L

Alerts & Detections

14% LLM & Prompts

25% Anomaly detections

Pipeline Operations

Lastseen gte

mm/dd/yyyy

Lastseen lte



mm/dd/yyyy

EXPORT

		Last seen ↓	Provider	Current value	
>	☐	4/23/2024	spark_app_start_latency_microseconds_count	spark-operator	10
>	☐	4/23/2024	spark_app_failure_count	spark-operator	10
>	☐	4/23/2024	spark_app_start_latency_seconds_bucket	spark-operator	3
>	☐	4/23/2024	spark_app_start_latency_seconds_sum	spark-operator	46,545
>	☐	4/23/2024	spark_app_executor_failure_count	spark-operator	42
>	☐	4/23/2024	spark_app_failure_execution_time_microseconds_count	spark-operator	10
>	☐	4/23/2024	spark_app_start_latency_microseconds_sum	spark-operator	46,552,507,258
>	☐	4/23/2024	spark_app_submit_count	spark-operator	10
>	☐	4/23/2024	spark_app_failure_execution_time_microseconds_sum	spark-operator	6,684,888,000,000
>	☐	4/23/2024	spark_app_start_latency_seconds_bucket	spark-operator	3
>	☐	4/23/2024	spark_app_running_count	spark-operator	4
>	☐	4/23/2024	spark_app_executor_success_count	spark-operator	12
>	☐	4/23/2024	spark_app_start_latency_seconds_bucket	spark-operator	3
>	☐	4/23/2024	spark_app_start_latency_seconds_bucket	spark-operator	4
>	☐	4/23/2024	spark_app_start_latency_seconds_bucket	spark-operator	1
>	☐	4/23/2024	spark_app_count	spark-operator	8
>	☐	4/23/2024	spark_app_start_latency_seconds_bucket	spark-operator	9



Alert AI - Security & Posture Dashboards





Alert AI - Inline/Offline Gateway

Prompt & Response (Blocked/Alerted/Allowed)



