

IBM Security QRadar
7.3.3

*Troubleshooting and System
Notifications Guide*



Note

Before you use this information and the product that it supports, read the information in [“Notices” on page 57](#).

Product information

This document applies to IBM® QRadar® Security Intelligence Platform V7.3.3 and subsequent releases unless superseded by an updated version of this document.

© **Copyright International Business Machines Corporation 2012, 2019.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

- About This Guide..... vii**
- Chapter 1. Techniques for troubleshooting a problem 1**
- Chapter 2. Running health checks..... 3**
- Chapter 3. Common problems..... 5**
 - Troubleshooting DSMs..... 5
 - Disk storage not accessible error..... 6
 - Verifying partition storage problem..... 6
 - Resolving log source error after protocol update..... 7
 - Verifying disk usage levels..... 7
 - Resolving disk usage issues..... 8
 - Events FAQ..... 8
 - Event processing performance..... 10
 - Identifying DSM and optimized custom property issues..... 11
 - Incomplete report results..... 12
 - Resolving limited disk space for backup partitions..... 12
 - License system notifications..... 13
 - Removing a license to prevent recurring system notifications..... 14
 - Resolving login errors with Active Directory accounts..... 14
 - Troubleshooting automatic update failure on networks that use IP-based firewall rules..... 15
 - Verifying that QRadar receives syslog events..... 16
 - Resolving unreceived syslog events..... 17
 - Fixing the certificate security browser warning..... 18
 - Installing and updating a Certificate Authority after a software update..... 18
- Chapter 4. Increased DNS requests..... 19**
- Chapter 5. QRadar system notifications..... 21**
 - Disk usage system notifications..... 21
 - Asset notifications for QRadar appliances..... 21
 - Asset change discarded..... 21
 - Asset growth deviations detected..... 22
 - Blacklist notification..... 22
 - External scan of an unauthorized IP address or range..... 23
 - Automatic update notifications for QRadar appliances..... 23
 - Auto update installed with errors..... 23
 - Automatic update error 23
 - Automatic update successful..... 24
 - Automatic updates successfully downloaded..... 24
 - Deployment of an automatic update 24
 - Custom rules notifications for QRadar appliances..... 24
 - CRE failed to read rules..... 25
 - Cyclic custom rule dependency chain detected..... 25
 - Expensive custom rule found..... 25
 - Disk notifications for QRadar appliances..... 26
 - Asset persistence queue disk full..... 27
 - Asset update resolver queue disk full..... 27
 - Disk failure..... 27

Disk full for the asset change queue	27
Disk replication falling behind.....	28
Disk storage available.....	28
Disk storage unavailable.....	29
Disk usage exceeded max threshold.....	29
Disk usage exceeded warning threshold.....	29
Disk usage returned to normal.....	30
Insufficient disk space to export data	30
Predictive disk failure.....	30
Process monitor must lower disk usage.....	31
Event and flow notifications for QRadar appliances.....	31
Event or flow data not indexed.....	31
Event pipeline dropped connections.....	31
Event pipeline dropped events.....	32
Events routed directly to storage.....	32
Expensive custom properties found.....	33
Flow collector cannot establish initial time synchronization.....	33
Maximum events or flows reached.....	34
Failure notifications for QRadar appliances.....	34
Accumulator cannot read the view definition for aggregate data	34
Accumulator is falling behind.....	34
Filter initialization failed.....	35
Infrastructure component is corrupted or did not start.....	36
Process monitor application failed to start multiple times.....	36
Store and forward schedule did not forward all events.....	36
Time synchronization failed.....	36
User authentication failed for automatic updates.....	37
User does not exist or is undefined.....	37
Certificate expires soon.....	37
Certificate is expired.....	38
High Availability notifications for QRadar appliances.....	38
Active high-availability (HA) system failure.....	38
Failed to uninstall a high-availability (HA) appliance.....	39
Failed to install high availability	39
Standby high-availability (HA) system failure.....	39
License notifications for QRadar appliances.....	39
License expired.....	40
License near expiration.....	40
Process monitor license expired or invalid.....	40
Limit notifications for QRadar appliances.....	41
Aggregated data limit was reached.....	41
Found an unmanaged process that is causing long transaction.....	41
Long running reports stopped.....	42
Long transactions for a managed process.....	42
Maximum sensor devices monitored.....	43
Process exceeds allowed run time.....	43
SAR sentinel operation restore.....	43
SAR sentinel threshold crossed.....	44
Threshold reached for response actions.....	44
Log and log source notifications for QRadar appliances.....	44
An error occurred when the log files were collected.....	44
Expensive DSM extensions were found.....	45
Log files were successfully collected	46
Log source created in a disabled state.....	46
Unable to determine associated log source.....	46
Memory and backup notifications for QRadar appliances.....	47
Backup unable to complete a request.....	47
Backup unable to run a request.....	47

Device backup failure.....	48
Last backup exceeded the allowed time limit.....	48
Backup unable to find storage directory error.....	48
Out of memory error.....	49
Out of memory error and erroneous application restarted.....	49
Offense notifications for QRadar appliances.....	50
Magistrate is unable to persist offense updates.....	50
Maximum active offenses reached.....	50
Maximum total offenses reached.....	51
Repair notifications for QRadar appliances.....	51
Accumulation is disabled for the anomaly detection engine.....	51
An infrastructure component was repaired.....	51
Custom property disabled.....	52
Data replication difficulty.....	52
Replication cleanup skipped for host.....	52
MPC: Process not shutdown cleanly.....	53
Protocol source configuration incorrect.....	53
Raid controller misconfiguration.....	53
Restored system health by canceling hung transactions.....	54
Vulnerability scan notifications for QRadar appliances.....	54
External scan gateway failure.....	54
Scan failure error.....	54
Scan tool failure.....	55
Scanner initialization error.....	55
Notices.....	57
Trademarks.....	58
Terms and conditions for product documentation.....	58
IBM Online Privacy Statement.....	59
General Data Protection Regulation.....	59

About This Guide

This information is intended for use with IBM QRadar and provides diagnostic and resolution information for common system notifications and errors that can be displayed when using QRadar SIEM.

IBM QRadar Troubleshooting and System Notifications Guide provides information on how to troubleshoot and resolve system notifications that display on the QRadar console. System notifications that display on the console can apply to any appliance or QRadar product in your deployment.

Unless otherwise noted, all references to QRadar can refer to the following products:

- IBM QRadar SIEM
- IBM QRadar Log Manager

Intended audience

System administrators responsible for troubleshooting must have administrative access to IBM QRadar and your network devices and firewalls. The system administrator must have knowledge of your corporate network and networking technologies.

Network administrators who are responsible for installing and configuring QRadar systems must be familiar with network security concepts and the Linux operating system.

Technical documentation

To find IBM QRadar product documentation on the web, including all translated documentation, access the [IBM Knowledge Center](http://www.ibm.com/support/knowledgecenter/SS42VS/welcome) (<http://www.ibm.com/support/knowledgecenter/SS42VS/welcome>).

For information about how to access more technical documentation in the QRadar products library, see [QRadar Support – Assistance 101](https://ibm.biz/qradarsupport) (<https://ibm.biz/qradarsupport>).

Contacting customer support

For information about contacting customer support, see [QRadar Support – Assistance 101](https://ibm.biz/qradarsupport) (<https://ibm.biz/qradarsupport>).

Statement of good security practices

IT system security involves protecting systems and information through prevention, detection and response to improper access from within and outside your enterprise. Improper access can result in information being altered, destroyed, misappropriated or misused or can result in damage to or misuse of your systems, including for use in attacks on others. No IT system or product should be considered completely secure and no single product, service or security measure can be completely effective in preventing improper use or access. IBM systems, products and services are designed to be part of a lawful comprehensive security approach, which will necessarily involve additional operational procedures, and may require other systems, products or services to be most effective. IBM DOES NOT WARRANT THAT ANY SYSTEMS, PRODUCTS OR SERVICES ARE IMMUNE FROM, OR WILL MAKE YOUR ENTERPRISE IMMUNE FROM, THE MALICIOUS OR ILLEGAL CONDUCT OF ANY PARTY.

Please Note:

Use of this Program may implicate various laws or regulations, including those related to privacy, data protection, employment, and electronic communications and storage. IBM QRadar may be used only for lawful purposes and in a lawful manner. Customer agrees to use this Program pursuant to, and assumes all responsibility for complying with, applicable laws, regulations and policies. Licensee represents that it will obtain or has obtained any consents, permissions, or licenses required to enable its lawful use of IBM QRadar.

Chapter 1. Techniques for troubleshooting a problem

Troubleshooting is a systematic approach to solving a problem. The goal of troubleshooting is to determine why something does not work as expected and how to resolve the problem. Certain common techniques can help with the task of troubleshooting.

The first step in the troubleshooting process is to describe the problem completely. Problem descriptions help you and the technical-support representative know where to start to find the cause of the problem. This step includes asking yourself basic questions:

- What are the symptoms of the problem?
- Where does the problem occur?
- When does the problem occur?
- Under which conditions does the problem occur?
- Can the problem be reproduced?

The answers to these questions typically lead to a good description of the problem, which can then lead to a resolution of the problem.

What are the symptoms of the problem?

When you start to describe a problem, the most obvious question is "What is the problem?" This question might seem straightforward; however, you can break it down into several focused questions that create a more descriptive picture of the problem. These questions can include:

- Who, or what, is reporting the problem?
- What are the error codes and messages?
- How does the system fail? For example, is the problem a loop, hang, crash, performance degradation, or incorrect result?

Where does the problem occur?

Determining where the problem originates is not always easy, but it is one of the most important steps in resolving a problem. Many layers of technology can exist between the reporting and failing components. Networks, disks, and drivers are only a few of the components to consider when you are investigating problems.

The following questions help you to isolate the problem layer:

- Is the problem specific to one appliance?
- Is the current environment and configuration supported?

If one layer reports the problem, the problem does not necessarily originate in that layer. Part of identifying where a problem originates is understanding the environment in which it exists. Take some time to completely describe the problem environment, including the operating system and version, all corresponding software and versions, and the hardware. Confirm that you are running within an environment that is supported; many problems can be traced back to incompatible levels of software that are not intended to run together or are not fully tested together.

When does the problem occur?

Develop a detailed timeline of events that lead up to a failure, especially for cases that are one-time occurrences. You can most easily develop a timeline by working backward: Start at the time an error was reported (as precisely as possible, even down to the millisecond), and work backward through the available logs and information. Typically, you need to look only as far as the first suspicious event that you find in a diagnostic log.

To develop a detailed timeline of events, answer these questions:

- Does the problem happen only at a certain time of day or night?
- How often does the problem happen?
- What sequence of events leads up to the time that the problem is reported?
- Does the problem happen after an environment change, such as an upgrade or an installation of software or hardware?

Under which conditions does the problem occur?

Knowing which systems and applications are running at the time that a problem occurs is an important part of troubleshooting. These questions about your environment can help you to identify the cause of the problem:

- Does the problem always occur when the same task is being performed?
- Does a certain sequence of events need to occur for the problem to occur?
- Do any other applications fail at the same time?

Answering these types of questions can help you explain the environment in which the problem occurs and correlate any dependencies. Remember when multiple problems occur around the same time, the problems are not necessarily related.

Can the problem be reproduced?

Problems that you can reproduce are often easier to solve. However, problems that you can reproduce can have a disadvantage. If the problem has a significant business impact, you do not want it to recur. If possible, re-create the problem in a test or development environment, which typically offers you more flexibility and control during your investigation. Answer the following questions:

- Can the problem be re-created on a test system?
- Are multiple users encountering the same type of problem?
- Can the problem be re-created by running a single command or a set of commands?

Chapter 2. Running health checks

DrQ is an extensible health check framework for QRadar. Run DrQ health checks before major events, such as upgrades, to determine whether there are any issues that need to be addressed first. You can also run DrQ routinely to monitor the health of your system. You can run all health checks at once, an individual check, or a group of checks.

Procedure

To run health checks, type the following command.

```
drq
```

This command runs all available checks in `/opt/ibm/si/diagnostiq` with the `checkup` mode, and with the `summary` output mode.

The following table shows the general parameters for DrQ.

Parameter	Description
-h	Shows the help information for DrQ.
-l	Lists all tests and shows which are valid and which are not applicable to the system.

The following table shows the filtering parameters for DrQ.

Parameter	Description
-d <directory>	Run all checks in a directory. You can include this flag more than once, to specify multiple directories. <pre>drq -d <path/to/directory1> -d <path/to/directory2></pre>
-f <filename>	Run a check by filename. You can include this flag more than once, to specify multiple checks. <pre>drq -f <path/to/filename1> -f <path/to/filename2></pre>
-m	Run all checks in a mode. The default mode is <code>checkup</code> .
-r	Recursively run all checks in a directory. You can include this flag more than once, to specify multiple directories. <pre>drq -r <path/to/directory1> -r <path/to/directory2></pre>
-t <tag>	Run a group of checks by tag. You can include this flag more than once, to specify multiple tags. <pre>drq -t <tag1> -t <tag2></pre>

The following table shows the output parameters for DrQ. These parameters are mutually exclusive.

Table 3. DrQ output parameters

Parameter	Description
-j	Outputs json output of the check results. You can pipe the content to the jq tool to parse and format the json output. <pre>drq -j jq</pre>
-q	Runs in quiet mode. Outputs one of the following exit codes: <ul style="list-style-type: none">• 0 for all success• 1 for checkups with failures• 2 for invalid lua files
-s	Runs in summary mode. Outputs the number of successes and failures. This is the default output mode for DrQ.
-v	Runs in verbose mode. Outputs success and failure messages for each check.

Chapter 3. Common problems

The following information can help you identify and resolve common problems in your IBM QRadar deployment.

Troubleshooting DSMs

Device Support Modules (DSMs) parse the events in IBM QRadar. You can think of DSMs as software plug-ins that are responsible for understanding and parsing events that are provided by an event source. An event source can be a security appliance, server, operating system, firewall, or database. DSMs can be any type of system that generates an event when an action occurs.

What is the difference between an unknown event and a stored event?

When events aren't parsed correctly, they appear on the **Log Activity** tab as one of the following event types:

Unknown events

The event is collected and parsed, but cannot be mapped or categorized to a specific log source. Log sources that aren't automatically discovered are typically identified as an unknown event log until a log source is manually created in the system. When an event cannot be associated to a log source, the event is assigned to a generic log source. You can identify these events by searching for events that are associated with the SIM `Generic` log source or by using the `Event is Unparsed` filter.

Stored events

The event cannot be understood or parsed by QRadar. When QRadar cannot parse an event, it writes the event to disk and categorizes the event as stored.

How can you find these unknown or stored events in the Log Activity tab?

To find events specific to your device, you search in QRadar for the source IP address of your device. You can also select a unique value from the event payload and search for `Payload Contains`. One of these searches might locate your event, and it is likely either categorized as unknown or stored.

You can also add a search filter for `Event in Unparsed`. This search locates all events that either cannot be parsed (stored) or events that might not be associated with a log source or auto discovered (unknown).

What do you do if the product version you have is not listed in the DSM Configuration Guide?

The *DSM Configuration Guide* contains a list of product manufacturers and the DSMs that are officially tested and validated against specific products. If the DSM is for a product that is officially supported by QRadar, but the version is out-of-date, you might need a DSM update to resolve any parsing issues. The product versions in the DSM guide were officially tested in-house, but software updates by vendors might add or change the event format for a specific DSM. In these cases, open a support ticket in [IBM Support](https://www.ibm.com/support/home/) for a review of the log source. (<https://www.ibm.com/support/home/>)

What do you do if the product device you have is not listed in the DSM Configuration Guide?

If your product device is not listed in the *DSM Configuration Guide*, it is not officially supported. For example, DSMs that appear on the IBM Security App Exchange are supplied by vendors and aren't officially supported by IBM. Not having an official DSM doesn't mean that the events are not collected. It indicates that the event that is received by QRadar might be identified as unknown on the **Log Activity** tab. You have these options:

- Open a request for enhancement (RFE) to have your device become officially supported.
 1. Go to the [QRadar Security RFE Community](https://ibm.biz/BdRPx5). (https://ibm.biz/BdRPx5)
 2. Log in to the Security RFE Community.
 3. Click the **Submit** tab and type the necessary information.

Note: If you have event logs from a device, attach the event information and include the product version of the device that generated the event log.
- Write a [log source extension to parse events for your device](https://www.ibm.com/support/knowledgecenter/SS42VS_DSM/com.ibm.dsm.doc/c_LogSourceGuide_ExtDocs_about.html). (https://www.ibm.com/support/knowledgecenter/SS42VS_DSM/com.ibm.dsm.doc/c_LogSourceGuide_ExtDocs_about.html)

Related information

[DSM Configuration Guide](#)

[QRadar: Log Activity tab \(Event viewer\) reports "unknown" events](#)

Disk storage not accessible error

Each host in your IBM QRadar deployment monitors the availability of partitions. Disk availability is tested every minute by opening a file, writing to it, and deleting it.

If the disk availability test takes longer than the default 5 seconds, then the host context process reports an error in the QRadar logs. An error might also occur when the QRadar system experiences high load and large volumes of data are written, searched, purged, or copied to another system.

The error might resemble the following output:

```
Jun 24 07:22:41 127.0.0.1 [hostcontext.hostcontext]
[5b3acf9a-aa8a-437a-b059-01da87333f43/SequentialEventDispatcher]
com.q11abs.hostcontext.ds.DiskSpaceSentinel: [ERROR]
[NOT:0150062100][172.16.77.116/- -]
[-/- -]The storage partition(s) /store/backup on qradarfc (172.16.77.116)
are not currently accessible. Manual intervention may be required to
restore normal operation.
```

If the message is displayed repeatedly, then verify the problem. For more information, see [“Verifying partition storage problem”](#) on page 6.

Verifying partition storage problem

You verify a partition storage problem by creating a temporary file on the IBM QRadar Console or a managed host.

Before you begin

Verify that the partition storage problem is not caused by external storage that is slow or unavailable.

Procedure

1. Use SSH to log in QRadar Console.
2. Create a test by typing the following commands:


```
touch /store/backup/testfile
ls -la /store/backup/testfile
```
3. If one of the following two messages is displayed, increase the partition test timeout period.
 - touch: cannot touch `/store/backup/testfile': Read-only file system
 - nfs server time out
 - a) Click the **Admin** tab.
 - b) On the **System Configuration** menu, click **System Settings > Advanced**.
 - c) In the **Partition Tester Timeout (seconds)** list box, select or type 20.

- d) Click **Save**.
4. Choose one of the following options:
- If you use a network file system, such as iSCSI, Fibre Channel, or Network File System (NFS), then contact your storage administrator to verify that the file servers are accessible and operational.
 - If you use a local file system, then you might have a file system issue or a failed disk.

Resolving log source error after protocol update

An error message might appear when you attempt to edit a log source after you upgrade IBM QRadar, a Device Support Module (DSM), a protocol, or Vulnerability Information Services (VIS) components. To remove cached files, restart the QRadar web service and clear the QRadar files from your browser cache.

Before you begin

You must have SSH access and root account credentials.

About this task

The following message indicates that the web server didn't restart after QRadar was updated:

An error has occurred. Refresh your browser (press F5) and attempt the action again. If the problem persists, please contact customer support for assistance.

A file might be cached by QRadar web service or your desktop browser. You must restart QRadar web service and remove the cached files on your desktop.

Procedure

1. Use SSH to log in QRadar.
2. Stop the QRadar web service by typing the following command:

```
service tomcat stop
```
3. Keep one web browser window open.
4. To clear your browser cache, go to your web browser's preference settings.
5. Restart the browser.
6. Restart the QRadar web service by typing the following command:

```
service tomcat start
```

Verifying disk usage levels

The /var/log partition continues to operate when disk usage reaches 100%. However, log data might not be written to the disk, which might affect IBM QRadar startup processes and components.

Procedure

1. Use SSH to log in QRadar or a managed host.
2. To review the disk partition usage, type the following command:

```
df -h
```
3. Review the partitions to check their disk usage levels.

What to do next

If any of the monitored partitions reach 95%, see [“Resolving disk usage issues”](#) on page 8.

Resolving disk usage issues

File system partitions reach 95% when the data retention period settings are too high or the available storage is insufficient for the rate at which IBM QRadar receives data. If you reconfigure your retention bucket storage settings, the storage across your entire QRadar deployment is affected.

Procedure

1. Identify and remove older debug or patch files in the / file system.
2. Reduce disk usage on the /store file system.
3. Choose one of the following options:
 - Remove the oldest data from the /store/ariel/events file system.
 - Reduce your data retention period by adjusting the default retention bucket storage settings. For more information, see the *IBM QRadar Administration Guide*.
 - If the /store file is full, identify which log sources you can retain for shorter periods. Use the retention buckets to manage the log sources. For more information, see the *IBM QRadar Administration Guide*.
 - Consider an offboard storage solution such as iSCSI or Fibre Channel. For more information, see the *Offboard Storage Guide*.
 - If the /var/log file system reaches 100% capacity, QRadar does not shut down. Other issues might cause your log files to grow faster than expected.

Events FAQ

Use these frequently asked questions and answers about events to help you understand how QRadar correlates user activities in log files to generate offenses.

What is an Event?

In QRadar, an *event* is a message that is received and processed from a device on your network, and is a log of a particular action on that device. For example, an SSH login on a UNIX server, a VPN connection to a VPN device, or a firewall deny logged by your perimeter firewall are all events. These actions occur at an instance of time and are recorded in log files.

What is a unique event?

QRadar identifies a unique event based on a number of properties: source IP, destination IP, destination port, protocol, username, and log source ID or event ID. In some circumstances, source port is also used. When four events come in with the same key properties, they are coalesced into a single record for 10 seconds. When this period passes, the cycle repeats.

What is coalescing?

Coalescing is used to reduce data that is processed by the event pipeline. As data comes in and is coalesced, a large burst of events can convert hundreds of thousands of events into only a few dozen records. This action is done while QRadar maintains the count of the number of actual events. Coalescing gives QRadar the ability to detect, enumerate, and track an attack on a huge scale. It also protects the performance of the pipeline by reducing the workload of the system, including storage requirements for those events.

One limitation of coalescing occurs when data is being normalized. The first event in the coalesced record, which is used as the base record, is the only one that is kept in its entirety, including the payload. You can disable coalescing for devices and log sources that are used to track audit and compliance requirements in your environment. Examples of these kinds of devices might be custom applications, any customer-facing services, critical assets, or other important devices.

How do different event and log sources compare?

Many different log sources, and kinds of log sources, are supported by QRadar, such as firewalls, authentication devices, scanners, file servers, application platforms, and more. Each of these log sources types, as they are referred to in QRadar, provide a different perspective and type of information about your network. For example, a firewall reports the number of remote systems that are trying to get into your network. Simultaneously, a Windows or LDAP authentication server provides you with information about local staff members that are logging in to network resources. Your monitoring, audit, and security needs influence the kinds of log sources that you send to QRadar.

If a load balancer is used, do events get parsed by any event collector? Do multiple log sources get created?

Any Syslog-based source that is sending data to a load balancer in front of QRadar can be parsed on any of the event collectors. All auto-detected log sources in QRadar can be processed by any event collector in the deployment. When auto-detection is triggered and a request to create a log source is sent to the QRadar Console, the log source is created. Within a minute, all event processors and collectors are aware of this new log source, and any data that is sent to any event processor is automatically associated with that log source. Therefore, you can enable a load balancer in front of multiple event collectors and processors.

One log source is created in this scenario. Multiple create commands can be sent from multiple processors during the first few minutes that a log source is being detected, but it is only created once. When the log source manager on the QRadar Console receives the create command, it creates the log source if the log source does not exist. The log source manager ignores the create request if the log source exists.

What do the time stamps in event details mean?

These time stamps can have different values, depending on where the data originated, when data arrived, and when it was written to QRadar. The following list describes each time stamp:

- **Start Time**

An event record that represents when the event is received by a QRadar Event Collector. When events arrive in the pipeline, an object is created in memory, and the **Start Time** time stamp is set to that time.

- **Storage Time**

The time when data is written to disk by the Ariel component at the end of processing by the event pipeline. This time stamp is useful for determining whether events are being queued in the event pipeline for performance or licensing reasons.

- **Log Source Time**

The time from the event payload, usually the time in the Syslog header. However, some log sources include the time stamps in the payload such as Windows logs that have a `MessageTime` field in the body of the payload. If no time is available in the payload, the **Log Source Time** field is populated with the same value as **Start Time**.

How does QRadar assign a source and destination IP address to events?

QRadar events require both a source and destination IP. QRadar uses these locations to locate an IP address:

- **From the event payload (First method)**

Supported log sources (and universal log sources, if you create your own parsing regex patterns) search the payload of received events for a source and destination IP address. When located, these addresses are placed into the associated `Source` and `Destination` fields of the event records. If no IP addresses are in the payload, the other methods described next are used.

- **From the Hostname field in Syslog header (Second method)**

If no IP addresses are in the event payload, the **Hostname** field of the Syslog header is used. The **Hostname** field is common for events from log sources that include only a **Source address** in the field, such as web service logs, which include only the remote host IP address. In these types of events, the destination IP address is populated by the Syslog header or the **Hostname** field of the event. If the Syslog header or **Hostname** field is a host name, and not an IP address, no DNS look-up is done and instead, the third method is used.

- **Source IP address of the network packet (Third method)**

If no IP addresses are available in the event payload, or in the **Hostname** field in the Syslog header, then the network packet's source IP address is used for the IP address. If the source IP address is from the payload, then this packet IP address is used only as the **Destination IP address** field. In this instance, the device that sends QRadar the event message was the destination address of the event. Sometimes, both source and destination IP addresses are not found in either the payload, or in the **Hostname** field in the Syslog header. In this case, both the source and destination IP addresses of the event are assigned the network packet IP address.

How does QRadar assign IP addresses from central Syslog servers and NAT devices?

If you have an existing central Syslog service infrastructure, or you want to add a forwarding rule to this device that copies a stream of all events to the QRadar system. The IP address that QRadar uses is the packet IP address. If you use a central Syslog server, you might see the server's IP address in many events, and in the log source names.

To avoid this situation, configure the central Syslog server to add a prefix to a new Syslog header. This new header includes the original source IP address of the packet that was received. In this practice, common when forwarding events, QRadar provides this option as part of the forwarding destinations configuration. When you add the prefix, the IP address of the original event source device is always in the Syslog header **Hostname** field, and QRadar uses that IP address in the events. With NAT devices, you might need to go back to the log source devices and reconfigure them to use the IP address of the host in the Syslog header **Hostname** field, rather than a string-based host name. For example, syslog-ng services refer to this option as `chain_hostname`.

Event processing performance

Your IBM QRadar configuration might impact the event processing pipeline.

Event processing can be affected by DSM extensions, custom properties, rule tests, and global views. Event parsing and the custom rules engine automatically detect dropped events, run self-monitoring diagnostics, and report which DSM extensions, rules, and properties are slow.

Non-optimized custom properties

Custom properties are marked as optimized when they are regularly used for QRadar rules or for searching and filtering.

Non-optimized custom properties are parsed by the system, which affects search speeds and the loading rate of the web browser.

Rule tests that impact performance

Rules that test for regular expressions in an event payload affect QRadar performance because they search the entire payload.

Before you add a payload test to a rule, use rule filters to reduce the number of events. For example, when you search for a specific message in the active directory logs, apply the following filters to the rule:

- Log source type filter
- Log source group or specific log source filter
- An optional source IP address filter

The **Host with port open** test can impact performance because it compares passive and active ports with the events and flows that are received by QRadar. Before you use the test, do a bidirectional check to ensure that the host responds to the communication request.

Global views

A saved search that is grouped by multiple fields generates a global view that has many unique entries. As the volume of data increases, disk usage, processing times, and search performance can be impacted.

To prevent increasing the volume of data, only aggregate searches on necessary fields. You can reduce the impact on the accumulator by adding a filter to your search criteria.

Identifying DSM and optimized custom property issues

To help you troubleshoot performance degradation, identify issues with any DSM extensions that were recently installed or custom property that was recently enabled.

About this task

A DSM extension creates custom parsing methods by using regex pattern matching to extract event data from unsupported or incomplete log sources. Optimized custom properties use regular expression patterns to extract data from events as they are parsed.

The regex patterns that are used in your DSM extension or optimized custom property can impact event processing in IBM QRadar. Inefficient regular expressions can incorrectly route data directly to storage, degrade QRadar performance, and affect event processing.

DSM and optimized custom property issues can cause the following system notification:

```
Performance degradation has been detected in the event pipeline.  
Events were routed directly to storage.
```

Procedure

1. Disable any DSM extension or custom property that is recently installed or enabled.
2. Choose one of the following options:
 - If QRadar stops dropping events and you receive a system notification, then review your DSM extensions or custom properties to identify and improve the inefficient regex patterns.
 - If QRadar continues dropping events, then multiple DSM extensions or custom properties might be causing a problem with the event pipeline.
3. Use SSH to log in to the QRadar Event Processor that is dropping events and type the following command:

```
/opt/qradar/support/threadTop.sh -p 7777
```

The command displays the data processing engine activity. The following table describes the columns in the output:

Columns	Description
Server	Port or process.
ID	Process ID.
MSecs	CPU time.
Name	Process name.

4. If parser threads run longer than 1500 milliseconds, review the Java thread stacks by typing the following command:

```
/opt/qradar/support/threadTop.sh -p 7777 -s -e ".*Event Parser.*" | less
```

What to do next

If the Java thread stack contains `java.util.regex.Pattern$Curly.match`, then the performance degradation might be caused by your expensive DSM extensions or custom properties. For more information, see [“Expensive DSM extensions were found” on page 45](#) or [“Expensive custom properties found” on page 33](#).

If the Java thread stack doesn't have expensive regular expressions, then your DSM extensions or custom properties might have parsing issues. For more information, see the parsing issues topic in the *IBM QRadar Log Sources User Guide*.

Incomplete report results

After you configure and run IBM QRadar reports, you might see unexpected results. A report might seem like it does not display all the data that you require.

Data accumulation for a search starts only when the search is added to a scheduled report. For example, a report that is created on Wednesday and is scheduled to run every Monday does not display data for a full week. However, the next report contains data for the full week.

Try one of these solutions:

Run the search again.

Use the **Network, Activity, or Log Activity** tab to run the search again. You can compare the results with the generated report.

Review the notification message on the Reports tab.

The **Reports** tab displays a notification message when your data is incomplete.

Run your report against raw data from the initial time period.

Ensure that you capture all the report data by running your report against the raw data from the initial time period. For more information, see *IBM QRadar User Guide*.

Incomplete reports also occur when the system is unable to accumulate data aggregations within a 60-second interval. Every minute, QRadar creates data aggregations for each aggregated search. If the search counts and unique values in the searches are too large, the time that is required to process the aggregations might exceed 60 seconds. When the accumulation is unable to finish within 60 seconds, the accumulation interval is dropped. Time-series graphs and reports might be missing columns for the time period when the problem occurred. For more information, see [“Accumulator cannot read the view definition for aggregate data ” on page 34](#).

Related concepts

[Accumulator is falling behind](#)

Resolving limited disk space for backup partitions

A system notification appears because the destination file system has limited disk space. IBM QRadar cannot complete a backup with insufficient disk space.

You might receive the following system notification:

```
Backup: Not enough free disk space to perform backup.
```

System notifications about limited disk space occur when free space in the `/store/backup/` partition is less than double the last backup file size. Limited disk space results from the volume of data and your backup retention period settings. For more information, see the *IBM QRadar Administration Guide*.

When you configure the retention bucket storage settings, a global impact occurs on the storage across your QRadar deployment.

Disk usage warnings can occur on the QRadar Console or any managed host in your QRadar deployment. To check disk usage levels, review the monitored partitions on your QRadar Console or managed hosts.

Procedure

1. Verify backup partition disk levels.
 - a) Use SSH to log in to QRadar Console or the managed host.
 - b) Type the following command:

```
df -PTh /store/backup
```
2. Review the backup partition to check the disk utilization levels.
 - a) If the backup partition is greater than double the last backup file size, identify the location of your backup.
 - If your backup is on the same file system as the /store/ariel directory, move it to another storage system.
 - If your backup is external, check your usage and ensure that your backup retention period does not require more space than you have available.
 - b) Reduce disk usage on the /store file system.
 - Consider increasing your external storage size by using an offboard storage solution such as iSCSI or Fibre Channel. For more information, see the *Offboard Storage Guide*.
 - If your QRadar backup partition is mounted on an NFS share, try lowering the retention period for the backup. The default backup retention period is two days. For more information on configuring backup retention periods, see the *IBM QRadar Administration Guide*.

License system notifications

IBM QRadar Console manages all the licenses in the deployment. Daily system notifications generate before and after a license expires.

The following table displays the QRadar components that depend on active licenses.

Type of expired license	Results
Console	When the QRadar Console license expires, event sources or flow sources that are forwarded directly to the Console are not processed or stored. Existing data is not affected. Until the Console license is renewed, direct event sources to a managed host with a valid license in the deployment.
Managed host	When a managed host license expires in your deployment, host context is disabled on that managed host. The expired appliance does not process event or flow data.
Feature	When a feature license expires, such as X-Force®, the feature stops receiving data updates. The feature relies on the latest data stream provided by X-Force.

Removing a license to prevent recurring system notifications

When an IBM QRadar appliance expires, the appliance cannot be used until its license is updated.

About this task

If a feature, such as X-Force, has an expired license, delete the license to prevent recurring system notifications.

If you have high-availability (HA) appliances, delete the secondary HA host license.

Procedure

1. Log in to QRadar Console.
2. Click the **Admin** tab.
3. On the navigation menu, click **System Configuration**.
4. Click the **System and License Management** icon.
5. From the **Display** list box, select **Licenses**.
6. Select the expired license.
License Information Messages lists the expired licenses.
7. Click **Actions > Delete License**.
8. Click **Confirm**.

What to do next

Update your expired license. For more information, see [Uploading a license key \(http://www.ibm.com/support/knowledgecenter/SS42VS_7.2.7/com.ibm.qradar.doc/t_qradar_adm_upload_license_key.html\)](http://www.ibm.com/support/knowledgecenter/SS42VS_7.2.7/com.ibm.qradar.doc/t_qradar_adm_upload_license_key.html).

Resolving login errors with Active Directory accounts

If you get an error when you log in to IBM QRadar with a valid Active Directory account, verify whether you have time synchronization issues.

About this task

When a valid Active Directory account is not synchronized with your QRadar Console, a login error similar to the following might occur:

The username and password you supplied are not valid. Please try again.

You can manually synchronize data between the QRadar server and the LDAP authentication server.

If you use authorization that is based on user attributes or groups, user information is automatically imported from the LDAP server to the QRadar console.

Each group that is configured on the LDAP server must have a matching user role or security profile that is configured on the QRadar console. For each group that matches, the users are imported and assigned permissions that are based on that user role or security profile.

By default, synchronization happens every 24 hours. The timing for synchronization is based on the last run time. For example, if you manually run the synchronization at 11:45 pm, and set the synchronization interval to 8 hours, the next synchronization will happen at 7:45 am. If the access permissions change for a user that is logged in when the synchronization occurs, the session becomes invalid. The user is redirected back to the login screen with the next request.

do these steps.

Procedure

1. If your Active Directory was not recently configured, use SSH to log in to QRadar as the root user.

2. Type the following command:

```
cat /opt/qradar/conf/login.conf
```

3. Verify that the server is configured for Active Directory authentication.
For example, an authenticated server might resemble the following output:

```
LDAPServerURL=ldaps://<server>:<port>
```

The `<server>` option is the Active Directory domain controller that receives the QRadar authentication. 389 is the default Active Directory LDAP port.

4. Copy the Active Directory domain controller IP address.
5. Type the following command and use the Active Directory domain controller IP address for the `<server>` option:

```
ntpdate -q <server>
```

6. Verify that the offset time is more than +/- 300 seconds.
The output might resemble the following example:

```
server 192.0.2.0, stratum 3, offset -10774.586000, delay 0.04221 19 Nov 13:59:16 ntpdate[22011]: step time server 192.0.2.0 offset -10774.586000 sec
```

If the offset time is more than +/- 300 seconds, then the time interval between the QRadar Console and the Active Directory server causes the authentication issues.

7. Restart the QRadar web service by typing the following command:

```
service tomcat restart
```

Restarting the QRadar web service logs off all users, stops exporting events, and stops generating reports. You might need to manually restart some reports or wait for a maintenance window to complete this procedure.

8. If the QRadar Console system time and the Active Directory server system time differ by at least 5 minutes, follow these steps:
 - a) Click the **Admin** tab.
 - b) On the navigation menu, click **System Configuration**.
 - c) Click **Authentication**.
 - d) In the **Authentication Module** list, select **LDAP**.
 - e) Click **Manage Synchronization > Run Synchronization Now**.

Troubleshooting automatic update failure on networks that use IP-based firewall rules

A change in the IP address for the auto update server can cause errors if you use IP-based firewall rules. An error message might display if IBM QRadar does not download the daily or weekly updates from the QRadar automatic update server. If you block communication based on the IP address that is associated with `auto-update.qradar.ibmcloud.com/`, you must update the firewall rules to allow the new IP address before you can receive automatic updates.

About this task

If you have IP-based firewall rules to allow automatic updates between the QRadar Console and the internet, you must update your firewall configuration with the following host name or IP address:

Description	Host name	IP address	Location	Status
Automatic Update Server	auto-update.qradar.ibmcloud.com/	69.20.113.167	United States	Active host name and IP address for US-based administrators.

Important: Configure your firewall rules to allow host names to ensure that automatic updates are not interrupted if a server IP address is changed. QRadar uses a speed test that can fail over to another auto update server. At run time for an auto update, the QRadar Console starts the speed test to download a test file from both server locations. The fastest connection to the QRadar Console is used to download daily and weekly automatic updates.

After you update the firewall rules, you can test your auto update connection by manually retrieving a QRadar automatic update.

Procedure

1. Log in to the QRadar Console as an administrative user.
2. On the **Admin** tab, click the **Auto Update** icon.
3. Click **Get New Updates**. Wait for the connection and updates to complete. A dashboard system notification is generated when updates are successfully downloaded or when errors occur.
4. Click **View Log** to view a detailed summary.
 - If the update fails, a connection error message is displayed.
 - If the update is successful, the log provides a success message and displays the most current updates as "already installed."
5. If the test fails, try the test again or verify that any corporate firewall and proxy settings are enabled to allow external connections.

Related information

[QRadar: Important auto update server changes for administrators](#)

Verifying that QRadar receives syslog events

To verify that IBM QRadar receives events, review the full syslog header for remote syslog source events. QRadar might not receive syslog events because a fire wall blocked communication or the device did not send the events.

Before you begin

Review the event source that sends the syslog events and verify its IP address.

Procedure

1. Use SSH to log in to QRadar as the root user.
2. If the syslog destination is on another appliance, such as an event collector, use SSH to log in to the event collector.
3. Choose one of the following options.
 - For a TCP syslog, type the following command:

```
tcpdump -s 0 -A host Device_Address and port 514
```
 - For a UDP syslog, type the following command:

```
tcpdump -s 0 -A host Device_Address and udp port 514
```

The *Device_Address* must be an IPv4 address or a host name. The **tcpdump** command must run on the QRadar appliance that receives the events from your device. By default, QRadar appliances are configured to receive syslog events by using TCP or UDP and port 514. Do not configure the QRadar firewall.

4. If the **tcpdump** command do not display events, then the syslog events are not sent to the QRadar Console.
 - a) Ask your firewall administrator or operations group to check for firewalls that block communication between the QRadar appliance and the device.

- b) Verify that a TCP port is open to Telnet by typing the following command on QRadar:

```
telnet Device_IPAddress 514
```
- c) Review your remote device's syslog configuration to ensure that events are sent to the proper appliance.

Resolving unreceived syslog events

If the **tcpdump** command lists events, but no events are shown on the log activity, then the syslog events are not received by the IBM QRadar Console.

Procedure

1. Review your system notifications.
2. If the system notifications display the incorrect source address for the log source, choose one of the following options:
 - Manually re-create the log source.
 - Update the **Log Source Identifier** field with the correct host name or IP address.
3. Verify that the device supports QRadar automatic discovery.

The *IBM QRadar DSM Configuration Guide* appendix lists which Device Support Modules (DSMs) support automatic log source creation.

4. Verify that the log sources in QRadar match the tcpdump results.
 - a) Search for the log source host name or packet IP address in the tcpdump results.
 - b) Click the **Admin** tab.
 - c) On the navigation menu, click **Data Sources**.
 - d) In the Events pane, click **Log Sources**.
 - e) Search for the log source host name or packet IP address.

If the QRadar host name or packet IP address does not match the tcpdump results, then the log source might be created with an incorrect address. For some devices, unexpected values occur in the syslog header when the event source handles events from multiple devices. Your device might be able to preserve the original event IP address before the syslog event is sent.

5. Search for a unique payload value in QRadar.
 - a) Review the tcpdump raw payloads.
 - b) Select an identifier that is unique to your event source.
 - c) Click the **Log Activity** tab.
 - d) On the toolbar, click **Add Filter**.
 - e) From the **Parameter** menu, select **Payload Contains**.
 - f) In the **Value** field, type your unique identifier.
 - g) Review the search results.

What to do next

If the results return a different log source, then an auto-detection false positive occurred. Delete the wrongly detected log source.

If the log source is discovered incorrectly, verify that your QRadar Console is installed with the latest DSM version. Rediscover the log source.

Fixing the certificate security browser warning

To fix a browser warning that the QRadar security certificate is not valid or not secure, you can download a certificate from the issuing URL and import it into your browser.

Procedure

1. Download the Certificate Authority (CA) content from the QRadar server:
 - a) Download the root CA from `http://<host_ip>:9381/vault-qrd_ca.pem`
 - b) Download the intermediate CA from `http://<host_ip>:9381/vault-qrd_ca_int.pem`

Tip: If you need the CA bundle, you can concatenate the intermediate CA with the root CA.
2. Copy the CA files to your local computer, and then log out of QRadar.
3. Import the CA into your browser by using the appropriate method for your browser:
 - [Mozilla Firefox](https://www.cyberciti.biz/faq/firefoxadding-trusted-ca/) [https://www.cyberciti.biz/faq/firefoxadding-trusted-ca/]
 - [Google Chrome](https://wiki.wmtransfer.com/projects/webmoney/wiki/Installing_root_certificate_in_Google_Chrome)[https://wiki.wmtransfer.com/projects/webmoney/wiki/Installing_root_certificate_in_Google_Chrome]
 - [Microsoft Internet Explorer](https://msdn.microsoft.com/en-us/library/cc750534.aspx) [https://msdn.microsoft.com/en-us/library/cc750534.aspx]
 - [Apple Safari](https://portal.threatpulse.com/docs/sol/Solutions/ManagePolicy/SSL/ssl_safari_cert_ta.htm) [https://portal.threatpulse.com/docs/sol/Solutions/ManagePolicy/SSL/ssl_safari_cert_ta.htm]
 - [Opera](http://wiki.wmtransfer.com/projects/webmoney/wiki/Installing_root_certificate_in_Opera) [http://wiki.wmtransfer.com/projects/webmoney/wiki/Installing_root_certificate_in_Opera]
4. Restart the browser to ensure that the certificate is loaded.
5. Log in to QRadar and verify that the browser no longer displays the security warning.

Installing and updating a Certificate Authority after a software update

If you choose not to use the local Certificate Authority during the software update, you can select to use it after the installation is complete.

About this task

Some software updates include a new local Certificate Authority (CA). During the update, you can choose to add this CA or skip this process. If you didn't add the new CA during the update, you can use the following steps to create the new CA afterwards.

Procedure

1. On the QRadar Console, type the following command: `/opt/qradar/vault/bin/install-qradar-ca.sh`
2. Restart Tomcat on the Console by typing the following command: `service tomcat restart`
3. Log in to the Console, navigate to the **Admin** tab, and then click **Deploy**.

Chapter 4. Increased DNS requests

As part of reworking the App Framework technology stack for QRadar V7.3.2, a service that cached DNS requests has been removed.

If the increased DNS requests are flooding the DNS servers, you can optimize system performance by setting up your own DNS caching apart from the QRadar appliance.

Chapter 5. QRadar system notifications

Use the system notifications that are generated by IBM QRadar to monitor the status and health of your system. Software and hardware tools and processes continually monitor the QRadar appliances and deliver information, warning, and error messages to users and administrators.

System notifications are displayed on the QRadar dashboard or in the notification window when unexpected system behavior occurs. You can troubleshoot the most common QRadar notifications.

Disk usage system notifications

IBM QRadar disk sentry monitors the /, /store, /storetmp, /transient, and /var/log partitions before the partitions reach a pre-defined usage threshold.

The following topics can help you identify and resolve common problems in your IBM QRadar deployment. The following table displays the host context system notifications that depend on the disk usage of each monitored partition.

Notification	Description	Suggested action
Disk Sentry: Disk Usage exceeded warning threshold.	Disk usage is at 90% for a monitored partition. QRadar is not affected when the partition reaches this threshold. Continue to monitor your partition levels.	See “Disk usage exceeded warning threshold” on page 29.
Disk Sentry: Disk Usage exceeded max threshold.	Disk usage is at 95% for a monitored partition. QRadar data collection and search processes are shut down to protect the file system from reaching 100%.	See “Disk usage exceeded max threshold” on page 29.
Disk sentry: System disk usage back to normal levels.	After disk usage reaches a threshold of 95%, it must return to 92% before QRadar automatically restarts data collection and search processes.	To lower the disk usage threshold, manually remove data from the affected partitions. See “Disk usage returned to normal” on page 30.

Asset notifications for QRadar appliances

Asset change discarded

38750106 - Asset Changes Aborted.

Explanation

An asset change exceeded the change threshold and the asset profile manager ignores the asset change request.

The asset profile manager includes an asset persistence process that updates the profile information for assets. The process collects new asset data and then queues the information before the asset model is updated. When a user attempts to add or edit an asset, the data is stored in temporary storage and added to the end of the change queue. If the change queue is large, the asset change can time out and the temporary storage is deleted.

User response

Select one of the following options:

- Add or edit the asset a second time.
- Adjust or stagger the start time for your vulnerability scans or reduce the size of your scans.

Asset growth deviations detected

38750137 - The system detected asset profiles that exceed the normal size threshold.

Explanation

The system detected one or more asset profiles in the asset database that show deviating or abnormal growth. Deviating growth occurs when a single asset accumulates more IP addresses, DNS host names, NetBIOS names, or MAC addresses than the system thresholds allow. When growth deviations are detected, the system suspends all subsequent incoming updates to these asset profiles.

User response

Determine the cause of the asset growth deviations:

- Hover your mouse over the notification description to review the notification payload. The payload shows a list of the top five most frequently deviating assets. It also provides information about why the system marked each asset as a growth deviation and the number of times that the asset attempted to grow beyond the asset size threshold.
- In the notification description, click **Review a report of these assets** to see a complete report of asset growth deviations over the last 24 hours.
- Review [Updates to asset data](http://www.ibm.com/support/knowledgecenter/SS42VS_7.3.1/com.ibm.qradar.doc/c_qradar_ug_asset_reconciliation.html) (http://www.ibm.com/support/knowledgecenter/SS42VS_7.3.1/com.ibm.qradar.doc/c_qradar_ug_asset_reconciliation.html).

Blacklist notification

38750136 - The Asset Reconciliation Exclusion rules added new asset data to the asset blacklists.

Explanation

A piece of asset data, such as an IP address, host name, or MAC address, shows behavior that is consistent with asset growth deviations.

An *asset blacklist* is a collection of asset data that is considered untrustworthy by the asset reconciliation exclusion custom engine rules. The rules monitor asset data for consistency and integrity. If a piece of asset data shows suspicious behavior twice or more within 2 hours, that piece of data is added to the asset blacklists. Subsequent updates that contain blacklisted asset data are not applied to the asset database.

User response

- In the notification description, click **Asset Reconciliation Exclusion rules** to see the rules that are used to monitor asset data.
- In the notification description, click **Asset deviations by log source** to view the asset deviation reports that occurred in the last 24 hours.
- If your blacklists are populating too aggressively, you can tune the asset reconciliation exclusion rules that populate them.

- If you want the asset data to be added to the asset database, remove the asset data from the blacklist and add it to the corresponding asset whitelist. Adding asset data to the whitelist prevents it from inadvertently reappearing on the blacklist.
- Review [Updates to asset data](http://www.ibm.com/support/knowledgecenter/SS42VS_7.3.1/com.ibm.qradar.doc/c_qradar_ug_asset_reconciliation.html) (http://www.ibm.com/support/knowledgecenter/SS42VS_7.3.1/com.ibm.qradar.doc/c_qradar_ug_asset_reconciliation.html).

External scan of an unauthorized IP address or range

38750126 - An external scan execution tried to scan an unauthorized IP address or address range.

Explanation

When a scan profile includes a CIDR range or IP address outside of the defined asset list, the scan continues. However, any CIDR ranges or IP addresses for assets that are not within your external scanner list are ignored.

User response

Update the list of authorized CIDR ranges or IP addresses for assets that are scanned by your external scanner. Review your scan profiles to ensure that the scan is configured for assets that are included in the external network list.

Automatic update notifications for QRadar appliances

Auto update installed with errors

38750067 - Automatic updates installed with errors. See the Auto Update Log for details.

Explanation

The most common reason for automatic update errors is a missing software dependency for a DSM, protocol, or scanner update.

User response

Select one of the following options:

- In the **Admin** tab, click the **Auto Update** icon and select **View Update History** to determine the cause of the installation error. You can view, select, and then reinstall a failed RPM.
- If an auto update is unable to reinstall through the user interface, manually download and install the missing dependency on your console. The console replicates the installed file to all managed hosts.

Automatic update error

38750066 - Automatic updates could not complete installation. See the Auto Update Log for details.

Explanation

The update process encountered an error or cannot connect to an update server. The system is not updated.

User response

Select one of the following options:

- Verify the automatic update history to determine the cause of the installation error.

In the **Admin** tab, click the **Auto Update** icon and select **View Log**.

- Verify that your console can connect to the update server.

In the **Updates** window, select **Change Settings**, then click the **Advanced** tab to view your automatic update configuration. Verify the address in the **Web Server** field to ensure that the automatic update server is accessible.

Automatic update successful

38750070 - Automatic updates completed successfully.

Explanation

Automatic software updates were successfully downloaded and installed.

User response

No action is required.

Automatic updates successfully downloaded

38750068 - Automatic updates successfully downloaded. See the Auto Updates log for details.

Explanation

Software updates were automatically downloaded.

User response

Click the link in the notification to determine whether any downloaded updates require installation.

Deployment of an automatic update

38750069 - Automatic updates installed successfully. In the Admin tab, click Deploy Changes.

Explanation

An automatic update, such as an RPM update, was downloaded and requires that you deploy the change to finish the installation process.

User response

In the **Admin** tab, click **Deploy Changes**.

Custom rules notifications for QRadar appliances

CRE failed to read rules

38750107 - The last attempt to read in rules (usually due to a rule change) has failed. Please see the message details and error log for information on how to resolve this.

Explanation

The custom rules engine (CRE) on an event processor is unable to read a rule to correlate an incoming event. The notification might contain one of the following messages:

- If the CRE was unable to read a single rule, in most cases, a recent rule change is the cause. The payload of the notification message displays the rule or rule of the rule chain that is responsible.
- In rare circumstances, data corruption can cause a complete failure of the rule set. An application error is displayed and the rule editor interface might become unresponsive or generate more errors.

User response

For a single rule read error, review the following options:

- To locate the rule that is causing the notification, temporarily disable the rule.
- Edit the rule to revert any recent changes.
- Delete and re-create the rule that is causing the error.

For application errors where the CRE failed to read rules, contact Customer Support.

Cyclic custom rule dependency chain detected

38750131 - Found custom rules cyclic dependency chain.

Explanation

A single rule referred to itself directly or to itself through a series of other rules or building blocks. The error occurs when you deploy a full configuration. The rule set is not loaded.

User response

Edit the rules that created the cyclic dependency. The rule chain must be broken to prevent a recurring system notification. After the rule chain is corrected, a save automatically reloads the rules and resolves the issue.

Expensive custom rule found

38750120 - Expensive Custom Rules Found in CRE. Performance degradation was detected in the event pipeline. Found expensive custom rules in CRE.

Explanation

The custom rules engine (CRE) is a process that validates if an event matches a rule set and then trigger alerts, offenses, or notifications.

A user can create a custom rule that has a large scope, uses a regex pattern that is not efficient, includes **Payload contains** tests, or combines the rule with regular expressions. When this custom rule is used, it negatively impacts performance, which can cause events to be incorrectly routed directly to storage. Events are indexed and normalized but they don't trigger alerts or offenses.

When multiple, expensive, or inefficient rule tests are used, the maximum event throughput rate can be reduced, causing backlogs of events to go through the rules engine. Events might be routed directly to storage, and this warning is displayed.

User response

Review the following options:

- Review the payload of the notification to determine which expensive rule in the pipeline affects performance.

For example, the following payload reports the test: "Payload Verification" rule in the pipeline and the EPS rate reported is 787 events per second, potentially reducing the maximum throughput of the rules engine.

```
Feb 23 15:56:58 ::ffff:10.1.2.4 [ecs-ep]
[Timer-27]com.q1labs.semsources.cre.CRE:
[WARN] [NOT:0040004101][10.1.2.4/- -]
[-/--]Expensive Custom Rules Based On Average Throughput in the last 60 seconds:
Test: Payload Verification=787.98045190917eps,
Monitoring: Suspect IPs seen with successful logins=899.02679830748eps
```

- On the **Offenses** tab, click **Rules** and use the search window to find and either edit or disable the expensive rule. By editing the rule, you can reduce the amount of data that goes through the rule, by applying a log source or IP address range filter. Expensive tests, such as payload contains, can also be reduced or removed if they are not required. Reference set tests are to be reviewed to ensure that they are not querying a large reference set.
- Use SSH to log in to the Event Processor and verify that parser threads are running for longer than 1500 milliseconds for EPS loads by using the following command:

```
/opt/qradar/support/threadTop.sh
```

Search the Java thread stack for `regex.Pattern.Curly`, `referenceSet`, `assets`, `host profile`, and `port profile` by using the following command:

```
/opt/qradar/support/threadTop.sh -p 7799 -s -e ".*CRE Processor.*"
```

- If the output contains `regex.Pattern.Curly`, issues with **Payload contains** tests are possible.
- If the output contains `referenceSet`, issues might occur with tests against large reference sets.
- If the output contains `assets`, `host profile`, and `port profile`, issues might occur with **Host with port open** tests or asset tests.

Rules Might Not Be the Issue

This notification can trigger when events are routed to storage around the rules engine. If, when you investigate this notice, the "EPS" rate in the notification is higher than ~20,000 EPS, it can indicate that the issue might be elsewhere. A rule that can process events upwards of 20,000 EPS is fairly optimized. The situation that triggered the 'events routed to storage' might not be a rule, but might be something else. Other items to consider are listed as follows.

- Is the system under higher load for other reasons, such as long-term data searching?
- Is the disk utilization at 85% or higher "on/store", and potentially data compression is affecting storage performance?
- If HA is in use, and event rates are higher than 10,000 EPS, ensure that sufficient bandwidth is between the two HA nodes. For example, a single 1Gbps connection, even in a dedicated crossover, can limit storage performance.
- Is there a separate `/transient/` partition. If not, then temporary data decompression might also use storage resources and contribute to the high storage demands.

Disk notifications for QRadar appliances

Asset persistence queue disk full

38750113 - Asset Persistence Queue Disk Full.

Explanation

The system detected the spillover disk space that is assigned to the asset persistence queue is full. Asset persistence updates are blocked until disk space is available. Information is not dropped.

User response

Reduce the size of your scan. A reduction in the size of your scan can prevent the asset persistence queues from overflowing.

Asset update resolver queue disk full

38750115 - Asset Update Resolver Queue Disk Full.

Explanation

The system detected that the spillover disk space that is assigned to the asset resolver queue is full.

The system continually writes the data to disk to prevent any data loss. However, if the system has no disk space, it drops scan data. The system cannot handle incoming asset scan data until disk space is available.

User response

Review the following options:

- Ensure that your system has free disk space. The notification can accompany SAR Sentinel notifications to notify you of potential disk space issues.
- Reduce the size of your scans.
- Decrease the scan frequency.

Disk failure

38750110 - Disk Failure: Hardware Monitoring has determined that a disk is in failed state.

Explanation

On-board system tools detected that a disk failed. The notification message provides information about the failed disk and the slot or bay location of the failure.

User response

If the notification persists, contact Customer Support or replace the parts.

Disk full for the asset change queue

38750117 - Asset Change Listener Queue Disk Full.

Explanation

The asset profile manager includes a process, change listener, that calculates statistics to update the CVSS score of an asset. The system writes the data to disk, which prevents data loss of pending asset statistics. However, if the disk space is full, the system drops scan data.

The system cannot process incoming asset scan data until disk space is available.

User response

Select one of the following options:

- Ensure that your system has sufficient free disk space.
- Reduce the size of your scans.
- Decrease the scan frequency.

Disk replication falling behind

38750103 - DRBD Sentinel: Disk replication is falling behind. See log for details.

Explanation

If the replication queue fills on the primary appliance, system load on the primary might increase. Replication issues are commonly caused by performance issues on the primary system, or storage issues on the secondary system, or bandwidth problems between the appliances.

User response

Select one of the following options:

- Review bandwidth activity by loading a saved search **MGMT: Bandwidth Manager** from the **Log Activity** tab. This search displays bandwidth usage between the console and hosts.
- If SAR sentinel notifications are recurring on the primary appliance, Distributed Replicated Block Device queues might be full on the primary system.
- Use SSH and the `cat /proc/drbd` command to monitor the Distributed Replicated Block Device status of the primary or secondary hosts.

Disk storage available

38750093 - One or more storage partitions that were previously inaccessible are now accessible.

Explanation

The disk sentry detected that the storage partition is available after the notification from [“Disk storage unavailable”](#) on page 29 appeared. Disk unavailability was resolved.

User response

No action is required.

Related concepts

[Disk storage unavailable](#)

38750092 - Disk Sentry has detected that one or more storage partitions are not accessible.

Disk storage unavailable

38750092 - Disk Sentry has detected that one or more storage partitions are not accessible.

Explanation

The disk sentry did not receive a response within 30 seconds. A storage partition issue might exist, or the system might be under heavy load and not able to respond within the 30-second threshold.

User response

Select one of the following options:

- Verify the status of your `/store` partition by using the **touch** command.

If the system responds to the **touch** command, the unavailability of the disk storage is likely due to system load.

- Determine whether the notification corresponds to dropped events.

The system drops events when it cannot write events to disk. Investigate the status of storage partitions.

Related concepts

[Disk storage available](#)

38750093 - One or more storage partitions that were previously inaccessible are now accessible.

Disk usage exceeded max threshold

38750038 - Disk Sentry: Disk Usage Exceeded Max Threshold.

Explanation

At least one disk on your system is 95% full.

To prevent data corruption, some processes shut down. Event collection is suspended until the disk usage falls below 92%.

User response

Identify which partition is full, such as the `/` and `/store` file systems. Free disk space by deleting files that are not needed. For example, remove debug output and patch files from the `/` file system. If the `/store` file system is at 95% capacity, look to the subdirectories to determine whether you can move the files to a temporary location or you can delete any files.

Note: If the files are deleted, you cannot search these events.

You can also manually delete older data in the `/store/ariel/` directories. The system automatically restarts processes after you free enough disk space to fall below a threshold of 92% capacity.

Disk usage exceeded warning threshold

38750076 - Disk Sentry: Disk Usage Exceeded warning Threshold.

Explanation

The disk sentry detected that the disk usage on your system is greater than 90%.

To prevent data corruption, the system disables processes when the disk space on your system reaches 95% full. This includes the event collection processes.

User response

You must free some disk space by deleting files or by changing your data retention policies. The system can automatically restart processes after the disk space usage falls below a threshold of 92% capacity.

Disk usage returned to normal

38750077 - Disk Sentry: System Disk Usage Back To Normal Levels.

Explanation

The disk sentry detected that the disk usage is below 90% of the overall capacity.

User response

No action is required.

Insufficient disk space to export data

38750096 - Insufficient disk space to complete data export request.

Explanation

If the export directory does not contain enough space, the export of event, flow, and offense data is canceled.

User response

Select one of the following options:

- Free some disk space in the `/store/exports` directory.
- Configure the **Export Directory** property in the **System Settings** window to use to a partition that has sufficient disk space.
- Configure an offboard storage device.

Predictive disk failure

38750111 - Predictive Disk Failure: Hardware Monitoring has determined that a disk is in predictive failed state.

Explanation

The system monitors the status of the hardware on an hourly basis to determine when hardware support is required on the appliance.

The on-board system tools detected that a disk is approaching failure or end of life. The slot or bay location of the failure is identified.

User response

Schedule maintenance for the disk that is in a predictive failed state.

Process monitor must lower disk usage

38750045 - Process Monitor: Disk usage must be lowered.

Explanation

The process monitor is unable to start processes because of a lack of system resources. The storage partition on the system is likely 95% full or greater.

User response

Free some disk space by manually deleting files or by changing your event or flow data retention policies. The system automatically restarts system processes when the used disk space falls below a threshold of 92% capacity.

Event and flow notifications for QRadar appliances

Event or flow data not indexed

38750101 - Event/Flow data not indexed for interval.

Explanation

If too many indexes are enabled or the system is overburdened, the system might drop the event or flow from the index portion.

User response

Select one of the following options:

- If the dropped index interval occurs with SAR sentinel notifications, the issue is likely due to system load or low disk space.
- To temporarily disable some indexes to reduce the system load, on the **Admin** tab, click the **Index Management** icon.

Event pipeline dropped connections

38750061 - Connections were dropped by the event pipeline.

Explanation

A TCP-based protocol dropped an established connection to the system.

The number of connections that can be established by TCP-based protocols is limited to ensure that connections are established and events are forwarded. The event collection service (ECS) allows a maximum of 15,000 file handles and each TCP connection uses three file handles.

TCP protocols that provide drop connection notifications include the following protocols:

- TCP syslog protocol
- TLS syslog protocol
- TCP multi-line protocol

User response

Review the following options:

- Distribute events to more appliances. Connections to other event and flow processors distribute the work load from the console.
- Configure low priority TCP log source events to use the UDP network protocol.
- Tune the system to reduce the volume of events and flows that enter the event pipeline.

Event pipeline dropped events

38750060 - Events/Flows were dropped by the event pipeline.

Explanation

If there is an issue with the event pipeline or you exceed your license limits, an event or flow might be dropped.

Dropped events and flows cannot be recovered.

User response

Review the following options:

- Verify the incoming event and flow rates on your system. If the license is exceeded and the event pipeline is dropping events, expand your license to handle more data.
- Review the recent changes to rules or custom properties. Rule or custom property changes can cause changes to your event or flow rates and might affect system performance.
- Determine whether the issue is related to SAR notifications. SAR notifications might indicate that queued events and flows are in the event pipeline. The system usually routes events to storage, instead of dropping the events.
- Tune the system to reduce the volume of events and flows that enter the event pipeline.

Events routed directly to storage

38750088 - Performance degradation has been detected in the event pipeline. Event(s) were routed directly to storage.

Explanation

To prevent queues from filling, and to prevent the system from dropping events, the event collection system (ECS) routes data to storage. Incoming events and flows are not categorized. However, raw event and flow data is collected and searchable.

User response

Review the following options:

- Verify the incoming event and flow rates. If the event pipeline is queuing events, expand your license to hold more data. To determine how close you are to your EPS/FPM license limit, monitor the **Event Rate (Events Per Second Raw)** graph on the **System Monitoring** dashboard. The graph shows you the current data rate. Compare the data rate to the per-appliance license configuration in your deployment. For more information about EPS/FPM license limits, see [QRadar: About EPS & FPM Limits](https://www.ibm.com/support/pages/qradar-about-eps-fpm-limits) (https://www.ibm.com/support/pages/qradar-about-eps-fpm-limits).
- Review recent changes to rules or custom properties. Rule or custom property changes might cause sudden changes to your event or flow rates. Changes might affect performance or cause the system to route events to storage.
- DSM parsing issues can cause the event data to route to storage. To verify whether the log source is officially supported, see the *DSM Configuration Guide*.
- SAR notifications might indicate that queued events and flows are in the event pipeline.

- Tune the system to reduce the volume of events and flows that enter the event pipeline. Events must be tuned at the source, not in the product. You can set coalescing on and configure your retention buckets to limit the number of stored events. License throttling monitors the number of incoming events to the system to manage input queues and licensing. For more information about retention buckets, see the *Administration Guide*.

Related tasks

[“Identifying DSM and optimized custom property issues” on page 11](#)

To help you troubleshoot performance degradation, identify issues with any DSM extensions that were recently installed or custom property that was recently enabled.

Expensive custom properties found

38750138 - Performance degradation was detected in the event pipeline. Expensive custom properties were found.

Explanation

During normal processing, custom event and custom flow properties that are marked as optimized are extracted in the pipeline during processing. The values are used in the custom rules engine (CRE) and search indexes.

Regex statements, which are improperly formed regular expressions, can cause events to be incorrectly routed directly to storage.

User response

Select one of the following options:

- Disable any custom property that was recently installed.
- Review the payload of the notification. If possible, improve the regex statements that are associated with the custom property.

For example, the following payload reports the regex pattern:

```
Feb 23 11:44:43 ::ffff:10.1.12.12 [ecs-ec]
[Timer-60] com.q1labs.semsources.filters.normalize.DSMFilter:
[WARN] [NOT:0080004105][10.130.126.12/- -]
[-/- -]Expensive Custom Properties Based On Average
Throughput in the last 60 seconds (most to least expensive)
- (\w+) /\S+=1136.0eps
```

- Modify the custom property definition to narrow the scope of categories that the property tries to match.
- Specify a single event name in the custom property definition to prevent unnecessary attempts to parse the event.
- Order your log source parsers from the log sources with the most sent events to the least and disable unused parsers.

Flow collector cannot establish initial time synchronization

38750009 - Flow collector could not establish initial time synchronization.

Explanation

The QFlow Collector process contains an advanced function for configuring a server IP address for time synchronization. In most cases, do not configure a value. If configured, the QFlow process attempts to synchronize the time every hour with the IP address time server.

User response

In the deployment actions, select the QFlow process. Click **Actions** > **Configure** and click **Advanced**. In the **Time Synchronization Server IP Address** field, clear the value and click **Save**.

Maximum events or flows reached

38750008 - The appliance exceeded the EPS or FPM allocation within the last hour.

Explanation

Each appliance is allocated a specific volume of event and flow data from the license pool. In the last hour, the appliance exceeded the allocated EPS or FPM.

If the appliance continues to exceed the allocated capacity, the system might queue events and flows, or possibly drop the data when the backup queue fills.

User response

- Adjust the license pool allocations to increase the EPS and FPM capacity for the appliance.
- Tune the system to reduce the volume of events and flows that enter the event pipeline.

Failure notifications for QRadar appliances

Accumulator cannot read the view definition for aggregate data

38750108 - Accumulator: Cannot read the aggregated data view definition in order to prevent an out of sync problem. Aggregated data views can no longer be created or loaded. Time series graphs will no longer work as well as reporting.

Explanation

A synchronization issue occurred. The aggregate data view configuration that is in memory wrote erroneous data to the database.

To prevent data corruption, the system disables aggregate data views. When aggregate data views are disabled, time series graphs, saved searches, and scheduled reports display empty graphs.

User response

Contact Customer Support.

Accumulator is falling behind

38750099 - The accumulator was unable to aggregate all events/flows for this interval.

Explanation

This message appears when the system is unable to accumulate data aggregations within a 60-second interval.

Every minute, the system creates data aggregations for each aggregated search. The data aggregations are used in time-series graphs and reports and must be completed within a 60-second interval. If the count of searches and unique values in the searches are too large, the time that is required to process the aggregations might exceed 60 seconds. Time-series graphs and reports might be missing columns for the time period when the problem occurred.

You do not lose data when this problem occurs. All raw data, events, and flows are still written to disk. Only the accumulations, which are data sets that are generated from stored data, are incomplete.

User response

The following factors might contribute to the increased workload that is causing the accumulator to fall behind:

Frequency of the incomplete accumulations

If the accumulation fails only once or twice a day, the drops might be caused by increased system load due to large searches, data compression cycles, or data backup.

Infrequent failures can be ignored. If the failures occur multiple times per day, during all hours, you might want to investigate further.

High system load

If other processes use many system resources, the increased system load can cause the aggregations to be slow. Review the cause of the increased system load and address the cause, if possible.

For example, if the failed accumulations occur during a large data search that takes a long time to complete, you might prevent the accumulator drops by reducing the size of the saved search.

Large accumulator demands

If the accumulator intervals are dropped regularly, you might need to reduce the workload.

The workload of the accumulator is driven by the number of aggregations and the number of unique objects in those aggregations. The number of unique objects in an aggregation depends on the group-by parameters and the filters that are applied to the search.

For example, a search that aggregates for services filters the data by using a local network hierarchy item, such as DMZ area. Grouping by IP address might result in a search that contains up to 200 unique objects. If you add destination ports to the search, and each server hosts 5 - 10 services on different ports, the new aggregate of `destination.ip + destination.port` can increase the number of unique objects to 2000. If you add the source IP address to the aggregate and you have thousands of remote IP addresses that hit each service, the aggregated view might have hundreds of thousands of unique values. This search creates a heavy demand on the accumulator.

To review the aggregated views that put the highest demand on the accumulator:

1. On the **Admin** tab, click **Aggregated Data Management**.
2. Click the **Data Written** column to sort in descending order and show the largest views.
3. Review the business case for each of the largest aggregations to see whether they are still required.

Related concepts

Incomplete report results

After you configure and run IBM QRadar reports, you might see unexpected results. A report might seem like it does not display all the data that you require.

Filter initialization failed

```
38750091 - Traffic analysis filter failed to initialize.
```

Explanation

If a configuration is not saved correctly, or if a configuration file is corrupted, the event collection service (ECS) might fail to initialize. If the traffic analysis process is not started, new log sources are not automatically discovered.

User response

Select one of the following options:

- Manually create log sources for any new appliances or event sources until traffic analysis process is working.
All new event sources are classified as SIM Generic until they are mapped to a log source.
- If you get an automatic update error, review the automatic update log to determine whether an error occurred when a DSM or a protocol was installed.

Infrastructure component is corrupted or did not start

38750083 - Infrastructure component corrupted.

Explanation

When the message service (IMQ) or PostgreSQL database cannot start or rebuild, the managed host cannot operate properly or communicate with the console.

User response

Contact Customer Support.

Process monitor application failed to start multiple times

38750043 - Process Monitor: Application has failed to start up multiple times.

Explanation

The system is unable to start an application or process on your system.

User response

Review which components are failing. For example, QFlow Collector fails to start when no flow sources are assigned. Use the deployment actions to remove that QFlow component.

Store and forward schedule did not forward all events

38750109 - A store and forward schedule finished while events were left on disk. These events will be stored on the local event collector until the next forwarding sessions begins.

Explanation

If the schedule contains a short start and end time or many events to forward, the event collector might not have sufficient time to transfer the queued events. Events are stored until the next opportunity to forward events. When the next store and forward interval occurs, the events are forwarded to the event processor.

User response

Increase the event forwarding rate from your event collector or increase the time interval that is configured for forwarding events.

Time synchronization failed

38750129 - Time synchronization to primary or Console has failed.

Explanation

The managed host cannot synchronize with the console or the secondary HA appliance cannot synchronize with the primary appliance.

Administrators must allow **ntpdate** communication on port 123. When time synchronization is incorrect, data might not be reported correctly to the console. The longer the systems go without synchronization, the higher the risk that a search for data, report, or offense might return an incorrect result. Time synchronization is critical to successful requests from managed host and appliances

User response

Contact customer support.

User authentication failed for automatic updates

38750127 - Automatic updates user authentication failed. A valid individual IBM ID is required.

Explanation

Valid credentials are required to authorize automatic downloads from the update server.

User response

Select one of the following options:

- Administrators must register for an account on the [IBM support website](http://www.ibm.com/support/) (<http://www.ibm.com/support/>).
- To view the automatic update settings, on the **Admin** tab, click the **Auto Update** icon and select **Change Settings > Advanced**. Administrators can confirm that the user name and password in the **Settings** window are correct.

User does not exist or is undefined

38750075 - User either does not exist or has an undefined role.

Explanation

The system attempted to update a user account with more permissions, but the user account or user role does not exist.

User response

On the **Admin** tab, click **Deploy Changes**. Updates to user accounts or roles require that you deploy the change.

Certificate expires soon

38750161 - The certificate named <certificate_name> will expire on <date>. Please update the certificate soon.

Explanation

Servers and clients use certificates to establish communication that uses Secure Sockets Layer (SSL) or Transport Layer Security (TLS). Certificates are issued with an expiration date that indicates how long the certificate remains valid. This message is first shown when QRadar determines that the certificate that is used for SAML authentication is set to expire within the next 14 days. The message is shown again at specific intervals that lead up to the expiration date.

User response

Select one of the following options:

- If you are using the QRadar_SAML certificate that is provided with QRadar, renew the certificate.
- If you are using a 3rd-party certificate, add a certificate.

If the certificate expires before you renew it or add a new one, QRadar cannot communicate with the SAML authentication server, and users can't log in.

For more information, see [SAML single sign-on authentication](#) in the *IBM QRadar Administration Guide*.

Certificate is expired

38750162 - The certificate named <certificate_name> has expired. Please update the certificate as soon as possible.

Explanation

Servers and clients use certificates to establish communication that uses Secure Sockets Layer (SSL) or Transport Layer Security (TLS). Certificates are issued with an expiration date that indicates how long the certificate remains valid.

This message appears when the certificate that is used for SAML authentication is expired. The message appears once a day and QRadar users cannot log in until the expired certificate is replaced or renewed.

User response

Select one of the following options:

- If you are using the QRadar_SAML certificate that is provided with QRadar, renew the certificate.
- If you are using a 3rd-party certificate, add a certificate.

For more information, see [SAML single sign-on authentication](#) in the *IBM QRadar Administration Guide*.

High Availability notifications for QRadar appliances

Active high-availability (HA) system failure

38750081 - Active HA System Failure.

Explanation

The active system cannot communicate with the standby system because the active system is unresponsive or failed. The standby system takes over operations from the failed active system.

User response

Review the following resolutions:

- Inspect the active HA appliance to determine whether it is powered down or experienced a hardware failure.
- If the active system is the primary HA, restore the active system.

Click the **Admin** tab and click **System and License Management**. From the **High Availability** menu, select the **Restore System** option.

- Review the `/var/log/qradar.log` file on the standby appliance to determine the cause of the failure.
- Use the **ping** command to check the communication between the active and standby system.
- Check the switch that connects the active and standby HA appliances.

Verify the IPtables on the active and standby appliances.

Failed to uninstall a high-availability (HA) appliance

38750087 - There was a problem while removing High Availability on the cluster.

Explanation

When you remove a HA appliance, the installation process removes connections and data replication processes between the primary and secondary appliances. If the installation process cannot remove the HA appliance from the cluster properly, the primary system continues to work normally.

User response

Try to remove the HA appliance a second time.

Failed to install high availability

38750086 - There was a problem installing High Availability on the cluster.

Explanation

When you install a high availability (HA) appliance, the installation process links the primary and secondary appliances. The configuration and installation process contains a time interval to determine when an installation requires attention. The high-availability installation exceeded the six-hour time limit.

No HA protection is available until the issue is resolved.

User response

Contact Customer Support.

Standby high-availability (HA) system failure

38750080 - Standby HA System Failure.

Explanation

The status of the secondary appliance switches to `failed` and the system has no HA protection.

User response

Review the following resolutions:

- Restore the secondary system.

Click the **Admin** tab, click **System and License Management**, and then click **Restore System**.

- Inspect the secondary HA appliance to determine whether it is powered down or experienced a hardware failure.
- Use the **ping** command to check the communication between the primary and standby system.
- Check the switch that connects the primary and secondary HA appliances.

Verify the IPtables on the primary and secondary appliances.

- Review the `/var/log/qradar.log` file on the standby appliance to determine the cause of the failure.

License notifications for QRadar appliances

License expired

38750123 - An allocated license has expired and is no longer valid.

Explanation

When a license expires on the console, a new license must be applied. When a license expires on a managed host, the appliance continues to process events and flows up to the rate that is allocated from the shared license pool.

When the license contributes EPS and FPM capacity to the shared license pool, the expiry might force the shared license pool into a deficit where it does not have enough capacity to meet the requirements of the deployment. In a deficit situation, functionality on the **Network Activity** and **Log Activity** tabs is blocked.

User response

1. Determine which appliance has the expired license.
 - a. On the **Admin** tab, click **System and License Management**.
 - b. In the **Display** box, select **Licenses**.

Expired licenses are shown in the **License Information Messages** section.
2. If the expired license is on the console, replace it.
3. If the expired license is on a managed host, review the shared license pool to ensure that the system has enough EPS and FPM capacity.
 - a. If the shared license pool is over-allocated, replace the expired license with a new license that has enough EPS and FPM to meet the system capacity requirements.
 - b. If the license pool has enough capacity, delete the expired license. In the **License** table, select the row for the expired license (shown nested beneath the managed host summary row), and select **Actions > Delete License**.

License near expiration

38750124 - A license is nearing expiration. It will need to be replaced soon.

Explanation

The system detected that a license for an appliance is within 35 days of expiration.

User response

No action is required.

Process monitor license expired or invalid

38750044 - Process Monitor: Unable to start process: license expired or invalid.

Explanation

The license is expired for a managed host. All data collection processes stop on the appliance.

User response

Contact your sales representative to renew your license.

Limit notifications for QRadar appliances

These limit notifications are related to thresholds, usage, maximums, stability, and reaching and returning to normal limits.

Aggregated data limit was reached

38750130 - The aggregated data view could not be created due to an aggregated limit.

Explanation

The accumulator process counts and prepares events and flows in data accumulations to assist with searches, displaying charts, and report performance. The accumulator process aggregates data in pre-defined time spans to create aggregate data views. An *aggregate data view* is a data set that is used to draw a time series graph, create scheduled reports, or trigger anomaly detection rules.

The Console is limited to 130 active aggregate data views.

The following user actions can create a new aggregate data view:

- New anomaly detection rules.
- New reports.
- New saved searches that use time series data.

When the aggregate data view limit is reached, the notification is generated. As users attempt to create new anomaly rules, reports, or saved searches, they are prompted in the user interface that the system is at the limit.

User response

To resolve this issue, administrators can review the active aggregate data views on the **Admin** tab in the **Aggregated Data Management** window. The aggregated data management feature provides information on the reports, searches, and anomaly detection rules in use by each aggregate data view. The administrator can review the list of aggregate data views to determine what data is most important to the users. Aggregate data views can be disabled to allow users to create a new rule, report, or saved search that requires an aggregate data view.

If the administrator decides to delete an aggregate data view, a summary provides an outline of the searches, rules, or reports affected. To re-create a deleted aggregate data view, the administrator needs only to re-enable or re-create the search, anomaly rule, or report. The system automatically creates the aggregate data view based on the data required.

Found an unmanaged process that is causing long transaction

38750048 - Transaction Sentry: Found an unmanaged process causing unusually long transaction that negatively effects system stability.

Explanation

The transaction sentry determines that an outside process, such as a database replication issue, maintenance script, auto update, or command line process, or a transaction is causing a database lock. Most processes cannot run for more than an hour. Repeated occurrences with the same process need to be investigated.

User response

Select one of the following options:

- Review the `/var/log/qradar.log` file for the word `TxSentry` to determine the process identifier that is causing your transaction issues.
- Wait to see whether the process completes the transaction and releases the database lock.
- Manually release the database lock by restarting the process identifier.

Long running reports stopped

38750054 - Terminating a report which was found executing for longer than the configured maximum threshold.

Explanation

The system cancels the report that exceeded the time limit. Reports that run longer than the following default time limits are canceled.

<i>Table 7. Default time limits by report frequency</i>	
Report frequency	Default time limits (hours)
Hourly	2
Daily	12
Manual	12
Weekly	24
Monthly	24

User response

Select one of the following options:

- Reduce the time period for your report, but schedule the report to run more frequently.
- Edit manual reports to generate on a schedule.

A manual report might rely on raw data but not have access to accumulated data. Edit your manual report and change the report to use an hourly, daily, monthly, or weekly schedule.

Long transactions for a managed process

38750056 - Transaction Sentry: Found managed process causing unusually long transaction that negatively effects system stability.

Explanation

The transaction sentry determines that a managed process, such as Tomcat or event collection service (ECS) is the cause of a database lock.

A managed process is forced to restart.

User response

To determine the process that caused the error, review the `qradar.log` for the word `TxSentry`.

Maximum sensor devices monitored

38750006 - Traffic analysis is already monitoring the maximum number of log sources.

Explanation

The system contains a limit to the number of log sources that can be queued for automatic discovery by traffic analysis. If the maximum number of log sources in the queue is reached, then new log sources cannot be added.

Events for the log source are categorized as SIM Generic and labeled as Unknown Event Log.

User response

Select one of the following options:

- Review SIM Generic log sources on the **Log Activity** tab to determine the appliance type from the event payload.
- Ensure that automatic updates can download the latest DSM updates to properly identify and parse log source events.
- Verify whether the log source is officially supported.

If your appliance is supported, manually create a log source for the events that were not automatically discovered.

- If your appliance is not officially supported, create a universal DSM to identify and categorize your events.
- Wait for the device to provide 1,000 events.

If the system cannot auto discover the log source after 1,000 events, it is removed from the traffic analysis queue. Space becomes available for another log source to be automatically discovered.

Process exceeds allowed run time

38750122 - Process takes too long to execute. The maximum default time is 3600 seconds.

Explanation

The default time limit of 1 hour for an individual process to complete a task is exceeded.

User response

Review the running process to determine whether the task is a process that can continue to run or must be stopped.

SAR sentinel operation restore

38750072 - SAR Sentinel: normal operation restored.

Explanation

The system activity reporter (SAR) utility detected that your system load returned to acceptable levels.

User response

No action is required.

SAR sentinel threshold crossed

38750073 - SAR Sentinel: threshold crossed.

Explanation

The system activity reporter (SAR) utility detected that your system load is above the threshold. Your system can experience reduced performance.

User response

Review the following options:

- In most cases, no resolution is required.

For example, when the CPU usage over 90%, the system automatically attempts to return to normal operation.

- If this notification is recurring, increase the default value of the SAR sentinel.

Click the **Admin** tab, then click **Global System Notifications**. Increase the notification threshold.

- For system load notifications, reduce the number of processes that run simultaneously.

Stagger the start time for reports, vulnerability scans, or data imports for your log sources. Schedule backups and system processes to start at different times to lessen the system load.

Threshold reached for response actions

38750102 - Response Action: Threshold reached.

Explanation

The custom rules engine (CRE) cannot respond to a rule because the response threshold is full.

Generic rules or a system that is tuned can generate a many response actions, especially systems with the **IF-MAP** option enabled. Response actions are queued. Response actions might be dropped if the queue exceeds 2000 in the event collection system (ECS) or 1000 response actions in Tomcat.

User response

- If the **IF-MAP** option is enabled, verify that the connection to the **IF-MAP** server exists and that a bandwidth problem is not causing rule response to queue in Tomcat.
- Tune your system to reduce the number of rules that are triggering.

Log and log source notifications for QRadar appliances

An error occurred when the log files were collected

38750141 - Collecting the required support logs failed with errors. See System and License Manager.

Explanation

Errors were encountered while the log files were being collected. The log file collection failed.

User response

To view information about why the collection failed, follow these steps:

1. Click **System and License Manager** in the notification message.

2. Expand **System Support Activities Messages**.
3. View additional information about why the log file collection failed.

Expensive DSM extensions were found

38750143 - Performance degradation was detected in the event pipeline. Expensive DSM extensions were found.

Explanation

A log source extension is an XML file that includes all of the regular expression patterns that are required to identify and categorize events from the event payload. Log source extensions might be referred to as *device extensions* in error logs and some system notifications.

During normal processing, log source extensions run in the event pipeline. The values are immediately available to the custom rules engine (CRE) and are stored on disk.

Improperly formed regular expressions (regex) can cause events to be routed directly to storage.

User response

Select one of the following options:

- Disable any DSM extension that was recently installed.
- Review the payload of the notification to determine which expensive DSM extension in the pipeline affects performance. If possible, improve the regex statements that are associated with the device extension.

For example, the following payload reports that the pipeline is blocked by the Checkpoint DSM:

```
Oct 23 12:32:53 ::ffff:10.1.2.4 [ecs-ec]
[Timer-57] com.q1labs.sensources.filters.normalize.DSMFilter: [WARN]
[NOT:0080014100][10.1.2.4/- -][-/ - -]Expensive Log Source or Log Source
Extensions Based On Average Throughput in the last 60 seconds
(most to least expensive) - Checkpoint=0.0eps, CatOS=86.0eps, Apache=2500.0eps,
Endpointprotection=2905.0eps
```

- Ensure that the log source extension is applied only to the correct log sources.
On the **Admin** tab, click **System Configuration > Data Sources > Log Sources**. Select each log source and click **Edit** to verify the log source details.
- If you are working with protocol-based log sources, reduce the event throttle to ensure that the events do not buffer to disk. The event throttle settings are part of the protocol configuration for the log source.
- Order your log source parsers from the log sources with the most sent events to the least and disable unused parsers.
- Verify that your Console is installed with the latest DSM versions.
- If log sources are created for devices that aren't in your environment, remove the log sources by using the following command:

```
/opt/qradar/bin/tatoggle.pl
```

If you have multiple event processors, copy the `/opt/qradar/conf/TrafficAnalysisConfig.xml` file to the `/store/configservices/staging/globalconfig/` directory. On the **Admin** tab, click **Deploy Full Configuration** for all managed hosts to obtain the configuration file.

Log files were successfully collected

38750142 - The required support logs have been successfully collected. See System and License Manager.

Explanation

The log files were successfully collected.

User response

To download the log file collection, follow these steps:

1. Click **System and License Manager** in the notification message.
2. Expand **System Support Activities Messages**.
3. Click **Click here to download file**.

Log source created in a disabled state

38750071 - A Log Source has been created in the disabled state due to license limits.

Explanation

Traffic analysis is a process that automatically discovers and creates log sources from events. If you are at your current log source license limit, the traffic analysis process might create the log source in the disabled state. Disabled log sources do not collect events and do not count in your log source limit.

User response

Review the following options:

- On the **Admin** tab, click the **Log Sources** icon and disable or delete low priority log sources. Disabled log sources do not count towards your log source license.
- Ensure that deleted log sources do not automatically rediscover. You can disable the log source to prevent automatic discovery.
- Ensure that you do not exceed your license limit when you add log sources in bulk.
- If you require an expanded license to include more log sources, contact your sales representative.

Unable to determine associated log source

38750007 - Unable to automatically detect the associated log source for IP address <IP address>. Unable to automatically detect the associated log source for IP address.

Explanation

When events are sent from an undetected or unrecognized device, the traffic analysis component needs a minimum of 25 events to identify a log source.

If the log source is not identified after 1,000 events, the system abandons the automatic discovery process and generates the system notification. The system then categorizes the log source as SIM Generic and labels the events as Unknown Event Log.

User response

Review the following options:

- Review the IP address in the system notification to identify the log source.

- Review the **Log Activity** tab to determine the appliance type from the IP address in the notification message and then manually create a log source.

Ensure that the **Log Source Identifier** field matches the host name in the original payload syslog header. Verify that the events are appearing on the device by deploying the changes and searching on the manually created log source.

- Review any log sources that forward events at a low rate. Log sources that have low event rates commonly cause this notification.
- To properly parse events for your system, ensure that automatic update downloads the latest DSMs.
- Review any log sources that provide events through a central log server. Log sources that are provided from central log servers or management consoles might require that you manually create their log sources.
- Verify whether the log source is officially supported. If your appliance is supported, manually create a log source for the events and add a log source extension.
- If your appliance is not officially supported, create a universal DSM to identify and categorize your events.

Memory and backup notifications for QRadar appliances

Backup unable to complete a request

38750033 - Backup: Not enough free disk space to perform the backup.

Explanation

Disk Sentry is responsible for monitoring system disk and storage issues. Before a backup begins, Disk Sentry checks the available disk space to determine whether the backup can complete successfully. If the free disk space is less than two times the size of the last backup, the backup is canceled. By default, backups are stored in `/store/backup`.

User response

To resolve this issue, select one of the following options:

- Free up disk space on your appliance to allow enough space for a backup to complete in `/store/backup`.
- Configure your existing backups to use a partition with free disk space.
- Configure more storage for your appliance. For more information, see the *Offboard Storage Guide*.

Backup unable to run a request

38750035 - Backup: Unable to Execute Backup Request.

Explanation

A backup cannot start or cannot complete for one of the following reasons:

- The system is unable to clean the backup replication synchronization table.
- The system is unable to run a delete request.
- The system is unable to synchronize backup with the files that are on the disk.
- The NFS-mounted backup directory is not available or has incorrect NFS export options (`no_root_squash`).
- The system cannot initialize on-demand backup.
- The system cannot retrieve configuration for the type of backup that is selected.

- Cannot initialize a scheduled backup.

User response

Manually start a backup to determine whether the failure reoccurs. If multiple backups fail to start, contact Customer Support.

Device backup failure

38750098 - Either a failure occurred while attempting to backup a device, or the backup was cancelled.

Explanation

The error is commonly caused by configuration errors in the configuration source management (CSM) or if a backup is canceled by a user.

User response

Select one of the following options:

- Review the credentials and address sets in CSM to ensure that the appliance can log in.
- Verify the protocol that is configured to connect to your network device is valid.
- Ensure that your network device and version is supported.
- Verify that your network device connects to the appliance.
- Verify that the most current adapters are installed.

Last backup exceeded the allowed time limit

38750059 - Backup: The last scheduled backup exceeded execution threshold.

Explanation

The time limit is determined by the backup priority that you assign during configuration.

User response

Select one of the following options:

- Edit the backup configuration to extend the time limit that is configured to complete the backup. Do not extend over 24 hours.
- Edit the failed backup and change the priority level to a higher priority. Higher priority levels allocate more system resources to completing the backup.

Backup unable to find storage directory error

38750164 - Backup: Unable to find storage directory error.

Explanation

The **Backup Repository Path** determines where backups are stored. By default, backups are stored in /store/backup. Administrators can configure the **Backup Repository Path** parameter on the **Backup Recovery Configuration** page on the **Admin** tab. If the system can't detect the **Backup Repository Path**, the backup can't complete successfully. For example, the path might not be found if there is an external storage mount failure.

User response

To resolve this issue, select one of the following options:

- On the **Backup Recovery Configuration** page, change the value for the **Backup Repository Path** to a file path that exists and has enough storage space for a backup. Deploy your changes, and then run the backup again.
- If the **Backup Repository Path** is pointing to external storage, verify that the mounted storage failed. If so, resolve the storage mount issue, and then run the backup again.

Out of memory error

38750004 - Application ran out of memory

Explanation

When the system attempts to use more than the allocated amount of memory, the application or service can stop working. Out of memory issues are often caused by software, or user-defined queries and operations that exhaust the available memory.

User response

Review the following resolutions:

- Review the error message that is written to the `/var/log/qradar.log` file to determine which component failed.
- If the Ariel proxy server is searching through large amounts of data or is using a grouping option that generates unique values in the search results, reduce the number of unique values or reduce the time frame of the search.
- If the accumulator is generating a time series graph with many aggregated unique values, reduce the size of the query.
- If a protocol-based log source is recently enabled, decrease the polling period to reduce the data queried. If multiple protocol-based log sources are running at the same time, stagger the start times.
- If a rule recently changed to track unique properties over long periods of time, reduce the time frame by half or reduce the number of matching events by adding another filter.

Out of memory error and erroneous application restarted

38750055 - Out of Memory: system restored, erroneous application has been restarted.

Explanation

An application or service ran out of memory and was restarted. Out of memory issues are commonly caused by software issues or user-defined queries.

User response

Review the following resolutions:

- Review the error message that is written to the `/var/log/qradar.log` file to determine which component failed.
- If the Ariel proxy server is searching through large amounts of data or is using a grouping option that generates unique values in the search results, reduce the number of unique values or reduce the time frame of the search.
- If the accumulator is generating a time series graph with many aggregated unique values, reduce the size of the query.

- If a protocol-based log source is recently enabled, decrease the polling period to reduce the data queried. If multiple protocol-based log sources are running at the same time, stagger the start times.
- If a rule recently changed to track unique properties over long periods of time, reduce the time frame by half or reduce the number of matching events by adding another filter.

Offense notifications for QRadar appliances

Magistrate is unable to persist offense updates

38750147 - Magistrate encountered serious errors that may prevent offenses from being updated.

Explanation

The system detected an exception when it wrote offense updates to the database. Events are processed and stored, but they might not contribute to offenses.

User response

Conduct a soft clean of the SIM data model with **Deactivate offenses** unchecked.

1. Click the **Admin** tab.
2. On the toolbar, click **Advanced > Clean SIM Model**.
3. Click **Soft Clean** to set the offenses to inactive.
4. Ensure that **Deactivate offenses** is not checked.
5. Click the **Are you sure you want to reset the data model?** check box and click **Proceed**.

When you clean the SIM model, all existing offenses are closed. Cleaning the SIM model does not affect existing events and flows.

Maximum active offenses reached

38750050 - MPC: Unable to create new offense. The maximum number of active offenses has been reached.

Explanation

The system is unable to create offenses or change a dormant offense to an active offense. The default number of active offenses that can be open on your system is limited to 2500. An active offense is any offense that continues to receive updated event counts in the past five days or less.

User response

Select one of the following options:

- Change low security offenses from open or active to closed, or to closed and protected.
- Tune your system to reduce the number of events that generate offenses.

To prevent a closed offense from being removed by your data retention policy, protect the closed offense.

Maximum total offenses reached

38750051 - MPC: Unable to process offense. The maximum number of offenses has been reached.

Explanation

By default, the process limit is 2500 active offenses and 100,000 overall offenses.

If an active offense does not receive an event update within 30 minutes, the offense status changes to dormant. If an event update occurs, a dormant offense can change to active. After five days, dormant offenses that do not have event updates change to inactive.

User response

Select one of the following options:

- Tune your system to reduce the number of events that generate offenses.
- Adjust the offense retention policy to an interval at which data retention can remove inactive offenses.

To prevent a closed offense from being removed by your data retention policy, protect the closed offense.

- To free disk space for important active offenses, change offenses from active to dormant.

Repair notifications for QRadar appliances

Accumulation is disabled for the anomaly detection engine

38750121 - Accumulation disabled for the Anomaly Detection Engine.

Explanation

Aggregate data view is disabled or unavailable or a new rule requires data that is unavailable.

A dropped accumulation does not indicate lost anomaly data. The original anomaly data is maintained because accumulations are data sets generated from stored data. The notification provides more details about the dropped accumulation interval.

The anomaly detection engine cannot review that interval of the anomaly data for the accumulation.

User response

Update anomaly rules to use a smaller data set.

If the notification is a recurring SAR sentinel error, system performance might be the cause of the issue.

An infrastructure component was repaired

38750084 - Corrupted infrastructure component repaired.

Explanation

A corrupted component that is responsible for host services on a managed host was repaired.

User response

No action is required.

Custom property disabled

38750097 - A custom property has been disabled.

Explanation

A custom property expression is disabled because the custom property expression has performance problems. Rules, reports, or searches that use this property, and which rely on the disabled expression to populate it, stop working properly.

The RegexMonitor feature monitors custom properties and disables any expressions that take longer than two seconds to parse. If inefficient custom property expressions are not disabled, the parsing queue overflows, and some events bypass parsing and do not normalize. Any searches, rules, or reports that rely on the non-normalized events do not function properly. When inefficient custom property expressions are disabled, parsing functions properly, and all events normalize. Only those searches, rules, and reports that rely on the custom property that is populated by the disabled expression do not function properly.

User response

Select one of the following options:

- Review the disabled custom property to correct your regex patterns. Do not re-enable disabled custom properties without first reviewing and optimizing the regex pattern or calculation.
- If the custom property is used for custom rules or reports, ensure that the **Optimize parsing for rules, reports, and searches** check box is selected.

Data replication difficulty

38750085 - Data replication experiencing difficulty.

Explanation

Data replication ensures that managed hosts can continue to collect data if the console is unavailable.

A managed host had difficulty downloading data. If a managed host repeatedly fails to download data, the system might experience performance or communication issues.

User response

If a managed host does not resolve the replication issue on its own, contact customer support.

Replication cleanup skipped for host

38750172 - Database replication cleanup skipped for host as it has been too long since it received an update.

Explanation

Data replication ensures that managed hosts can continue to collect data when the console is not available.

A managed host was skipped during cleanup because it was too long since it received an update. If a managed host fails to receive replication updates from the console, it isn't connecting properly to the console.

User response

To resolve this issue, select one of the following options:

- Click **Admin > System and License Management**, and then check the status of your managed host. Ensure that the **Host Status** is **Active**. If the **Host Status** is **unknown**, there are issues with the managed host that you need to investigate.
- If a managed host doesn't resolve the replication issue on its own, contact customer support.

MPC: Process not shutdown cleanly

38750058 - MPC: Server was not shutdown cleanly. Offenses are being closed in order to re-synchronize and ensure system stability.

Explanation

The magistrate process encountered an error. Active offenses close, services restarts, and the database tables are verified and rebuilt if necessary.

The system synchronizes to prevent data corruption. If the magistrate component detects a corrupted state, then the database tables and files are rebuilt.

User response

The magistrate component self-repairs. If the error continues, contact Customer Support.

Protocol source configuration incorrect

38750057 - A protocol source configuration may be stopping events from being collected.

Explanation

The system detected an incorrect protocol configuration for a log source. Log sources that use protocols to retrieve events from remote sources can generate an initialization error when a configuration problem in the protocol is detected.

User response

Resolve the protocol configuration issues by following these steps:

- Review the log source to ensure that the protocol configuration is correct.
Verify authentication fields, file paths, database names for JDBC, and ensure that the system can communicate with remote servers. Hover your mouse pointer over a log source to view more error information.
- Review the `/var/log/qradar.log` file for more information about the protocol configuration error.

Raid controller misconfiguration

38750140 - Raid Controller misconfiguration: Hardware Monitoring determined that a virtual drive is configured incorrectly.

Explanation

For maximum performance, raid controllers cache and battery backup unit (BBU) must be configured to use write-back cache policy. When write-through cache policy is used, storage performance degrades and might cause system instability.

User response

Review the health of the battery backup unit. If the battery backup unit is working correctly, change the cache policy to write-back.

Restored system health by canceling hung transactions

38750049 - Transaction Sentry: Restored system health by canceling hung transactions or deadlocks.

Explanation

The transaction sentry restored the system to normal system health by canceling suspended database transactions or removing database locks. To determine the process that caused the error, review the `qradar.log` file for the word `TxSentry`.

User response

No action is required.

Vulnerability scan notifications for QRadar appliances

External scan gateway failure

38750119 - An invalid/unknown gateway IP address has been supplied to the external hosted scanner, the scan has been stopped.

Explanation

When an external scanner is added, a gateway IP address is required. If the address that is configured for the scanner is incorrect, the scanner cannot access your external network.

User response

Select one of the following options:

- Review the configuration for any external scanners to ensure that the gateway IP address is correct.
- Ensure that the external scanner can communicate through the configured IP address.
- Ensure that the firewall rules for your DMZ are not blocking communication between your appliance and the assets you want to scan.

Scan failure error

38750090 - A scanner has failed.

Explanation

A scheduled vulnerability scan failed to import vulnerability data. Scan failures are typically caused by configuration or performance issues that result from a large volume of data to import. Scan failures can also occur when a scan report that is downloaded by the system is in an unreadable format.

User response

Follow these steps:

1. Click the **Admin** tab.
2. On the navigation menu, click **Data Sources**.
3. Click **Schedule VA Scanners**.
4. From the scanner list, hover the cursor in the **Status** column of any scanner to display a detailed success or failure message.

Scan tool failure

38750118 - A scan has been stopped unexpectedly, in some cases this may cause the scan to be stopped.

Explanation

The system cannot initialize a vulnerability scan and asset scan results cannot be imported from external scanners. If the scan tools stop unexpectedly, the system cannot communicate with an external scanner. The system tries the connection to the external scanner five times in 30-second intervals.

In rare cases, the discovery tools encounter an untested host or network configuration.

User response

Select one of the following options:

- Use the **System and Licence Management** window to review the configuration for external scanners to ensure that the gateway IP address is correct.
- Ensure that the external scanner can communicate through the configured IP address.
- Ensure that the firewall rules for your DMZ are not blocking communication between your appliance and the assets you want to scan.

Scanner initialization error

38750089 - A scanner failed to initialize.

Explanation

A scheduled vulnerability scan is unable to connect to an external scanner to begin the scan import process.

Scan initialization issues are typically caused by credential problems or connectivity issues to the remote scanner. Scanners that fail to initialize display detailed error messages in the hover text of a scheduled scan with a status of failed.

User response

Follow these steps:

1. Click the **Admin** tab.
2. On the navigation menu, click **Data Sources**.
3. Click **Schedule VA Scanners** icon.
4. From the scanner list, hover the cursor in the **Status** column of any scanner to display a detailed success or failure message.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions..

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

Trademarks

IBM, the IBM logo, and ibm.com[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Linux[®] is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

IBM Online Privacy Statement

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session cookies that collect each user’s session id for purposes of session management and authentication. These cookies can be disabled, but disabling them will also eliminate the functionality they enable.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, See IBM’s Privacy Policy at <http://www.ibm.com/privacy> and IBM’s Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled “Cookies, Web Beacons and Other Technologies” and the “IBM Software Products and Software-as-a-Service Privacy Statement” at <http://www.ibm.com/software/info/product-privacy>.

General Data Protection Regulation

Clients are responsible for ensuring their own compliance with various laws and regulations, including the European Union General Data Protection Regulation. Clients are solely responsible for obtaining advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulations that may affect the clients’ business and any actions the clients may need to take to comply with such laws and regulations. The products, services, and other capabilities described herein are not suitable for all client situations and may have restricted availability. IBM does not provide legal, accounting or auditing advice or represent or warrant that its services or products will ensure that clients are in compliance with any law or regulation.

Learn more about the IBM GDPR readiness journey and our GDPR capabilities and Offerings here: <https://ibm.com/gdpr>

