IBM Security QRadar
Version 7.1.0 (MR1)

Partition Splitting Technical Note

Note: Before using this information and the product that it supports, read the information in Notices and Trademarks on page 9.

CONTENTS

1 PARTITION SPLITTING

A NOTICES AND TRADEMARKS

1 PARTITION SPLITTING

This document provides information on how to use a QRadar SIEM script to create a partition and move the `/store/ariel/persistent_data` location and contents into the new partition for systems running High Availability (HA). This technical note only applies to HA systems.

The partition splitting process affects both the primary and secondary HA hosts. Before running the script, you must remove the HA secondary from the HA cluster configuration. This script takes several hours to complete. During this time, the secondary host is offline, however, the primary host continues to collect data and is still available to access using the user interface. The script performs the required actions and preserves the data integrity of the contents of the `/store` location. After the script is complete, you can reconfigure your HA cluster.

Unless otherwise noted, all references to QRadar SIEM refer to QRadar SIEM, IBM Security QRadar Log Manager, and IBM Security QRadar Network Anomaly Detection.

Before You Begin

The procedures in this document assume you have:

- Advanced knowledge of the Linux operating system.
- Administrative privileges for the QRadar SIEM software.
- Administrative privileges for the systems running QRadar SIEM and HA.

Be aware that there are potential risks involved with running the partition script.

- Determine the disc capacity of the system. You must give the new partition an appropriate size. Typically, the new partition should be approximately 25% the size of the `/store` location. The script does not have safeguards in place to prevent the introduction of values that are incorrect or too large.
- Investigate and find the root cause of your performance issues before you run the script. Partitioning and migrating the`/store` location can resolve throttling issues where HA data replication is the reason for the slowdown.
- There is a low risk of data loss. Make sure that the host has sufficient space for a new partition. For example, if you have 100 GB of free space, you should not allocate a 400 GB partition.

For assistance, contact Customer Support.

Partition Splitting Script

If you have experienced performance issues caused by HA data replication that partition splitting can resolve, you can use the partition splitting script to modify the boundaries of the `/store/` partition and move the associated temporary results to the newly created partition. This document provides information on preparing, configuring, and running the partition splitting script available with your QRadar SIEM installation.

To prepare and run the partitioning script, you need to log on to QRadar SIEM as an administrator, and then SSH to both the primary and secondary HA host.

The script is stored in the bin directory of QRadar SIEM: `/opt/qradar/bin`. The script takes two commands:

- **`size`** - Sets the disc space required for the new partition.
- **`continue`** - Resumes the processing after a reboot.

The partitioning script contains the complete set of instructions required; running the script may take several hours. You may be prompted to restart the host, if so, you can resume the script with the **`continue`** command.

Partitioning the HA Cluster Hosts

Before running the partition splitting script, the HA cluster must be disconnected. After disconnecting the HA cluster, run the partitioning script on each of the two HA systems. The script can take several hours to complete, however, you can run the partition splitting script on both hosts at the same time. After the script is complete on both hosts, you must reconnect the HA cluster.

To partition the HA cluster hosts, perform the following procedures:

1. Disconnect the HA Cluster.

   For more information see, **Disconnecting the HA Cluster**
2. Run the partition splitting script on the primary and secondary HA hosts:
   - **Partitioning the Primary HA Host**
   - **Partitioning the Secondary HA Host**
3. Reconnect the HA cluster. For information on reconnecting the HA cluster, see the *Adding an HA Cluster* section in your *IBM Security QRadar SIEM Administration Guide*.

Disconnecting the HA Cluster

To disconnect the HA cluster:

Step 1 Click the **Admin** tab.

Step 2 On the navigation menu, click **System Configuration**.

The **System Configuration** panel is displayed.

Step 3 Click the **System and License Management** icon.

The System and License Management window is displayed.

Step 4 Select the HA host you want to remove.

Step 5 From the **High Availability** menu, select **Remove HA Host**.

A confirmation message is displayed, indicating that removing an HA host reboots the user interface.

Step 6 Click **OK**.

When you remove an HA host, the host restarts.

Partitioning the Primary HA Host

To partition the primary HA host:

Step 1 Using SSH, log into the primary HA host as the **`root`** user:

Username: **root**

Password: **<password>**

Step 2 Change to the **`/opt/qradar/bin`** directory.

Step 3 Type **`./create_cursor_partition.sh size=<size>.`**

`<size>` should be approximately one quarter the `/store` capacity. **`<size>`** is written as a numeric value and the measurement specification. The partition size on the primary and secondary HA host must be the same. Type the measurement using one of the following:

- **`M`** for Megabyte
- **`G`** for Gigabyte
- **`T`** for Terabyte

If the script prompts you to restart the host, do the following steps:

a Restart the primary host and log in as the **`root`** user.

b Change to the **`/opt/qradar/bin`** directory.

c Type the following command to restart the script:

```
./create_cursor_partition.sh --continue.
```

Step 4 To check the partition when the script has finished, type **`df -h`**.

The results resemble the following output:

```
[root@xxxx-primary ~]# df -h
Filesystem  Size  Used  Avail  Use%  Mounted on
/dev/sda2   20G   7.0G  12G    38%   /
/dev/sda3   9.7G  581M  8.7G   7%    /var/log
/dev/sda1   97M   25M   68M    27%   /boot
tmpfs       12G   0     12G    0%    /dev/shm
```

```
/dev/drbd0  2.0T  58G   2.0T   3%    /store
/dev/sda5   9.7G  187M  9.0G   2%    /store/tmp
/dev/sda9   626G  2.0G  593G   1%    /store/ariel/persistent_data
```

Partitioning the Secondary HA Host

To partition the secondary HA host:

Step 1 Using SSH, log into the secondary HA host as the **`root`** user:

Username: **root**

Password: **<password>**

Step 2 Change to the **`/opt/qradar/bin`** directory.

Step 3 Type **`./create_cursor_partition.sh size=<size>.`**

`<size>` should be approximately one quarter the `/store` capacity. **`<size>`** is written as a numeric value and the measurement specification. The partition size on the primary and secondary HA host must be the same. Type the measurement using one of the following:

- **`M`** for Megabyte
- **`G`** for Gigabyte
- **`T`** for Terabyte

If the script prompts you to restart the host, perform the following steps:

a Restart the secondary host and log in as the **`root`** user.

b Change to the **`/opt/qradar/bin`** directory.

c Type the following command to restart the script:

```
./create_cursor_partition.sh --continue.
```

Step 4 To check the partition when the script has finished, type **`df -h`**.

The results resemble the following output:

```
[root@xxxx-secondary ~]# df -h
Filesystem  Size  Used  Avail  Use%  Mounted on
/dev/sda2   20G   7.0G  12G    38%   /
/dev/sda3   9.7G  581M  8.7G   7%    /var/log
/dev/sda1   97M   25M   68M    27%   /boot
tmpfs       12G   0     12G    0%    /dev/shm
/dev/drbd0  2.0T  58G   2.0T   3%    /store
/dev/sda5   9.7G  187M  9.0G   2%    /store/tmp
/dev/sda9   626G  2.0G  593G   1%    /store/ariel/persistent_data
```

5 Reconnect the HA cluster.

For more information on reconnecting an HA cluster, see the *IBM Security QRadar SIEM High Availability (HA) Guide*.

A NOTICES AND TRADEMARKS

What's in this appendix:

- **Notices**
- **Trademarks**

This section describes some important notices, trademarks, and compliance information.

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

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

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 states 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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply 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 Corporation
170 Tracer Lane,
Waltham MA 02451, USA

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.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

All 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 the names and addresses used by an actual business enterprise is entirely coincidental.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks