IBM® Storage

Proactive Early Threat Detection and Securing SQL Database With IBM QRadar and IBM Spectrum Copy Data Management Using IBM FlashSystem Safeguarded Copy

Contents

About this document

This IBM® blueprint publication focuses on early threat detection within a database environment by using IBM QRadar®. It also highlights how to proactively start a cyber resilience workflow in response to a cyberattack or potential malicious user actions.

The workflow that is presented here uses IBM Spectrum® Copy Data Management as orchestration software to start IBM FlashSystem® Safeguarded Copy functions. The Safeguarded Copy creates an immutable copy of the data in an air-gapped form on the same IBM FlashSystem for isolation and eventual quick recovery.

This document describes how to enable and forward SQL database user activities to IBM QRadar.

This document also describes how to create various rules to determine a threat, and configure and start a suitable response to the detected threat in IBM QRadar.

Finally, this document outlines the steps that are involved to create a Scheduled Job by using IBM Spectrum® Copy Data Management with various actions.

Executive summary

The financial effects of cyberattacks continue to rise. Cyberattacks can occur in various ways. They can take the form of malware or ransomware, which is targeted at stealing confidential data or holding valuable information for ransom.

Sometimes, these attacks are designed to destroy confidential data to cripple organizations. In many cases, the data breaches involve internal threat actors.

IBM QRadar® offers solutions to protect sensitive and regulated data. It provides this protection by continuously monitoring the data activity and accelerating compliance reporting that supports a zero trust approach to data management across environments, lifecycles, and platforms.

Detecting a threat before it starts can help speed recovery even more.

IBM QRadar is a security information and event management (SIEM) and threat management system that monitors activities while looking for signs that might indicate the start of an attack, such as logins from unusual IP addresses or outside business hours.

Now, IBM QRadar can proactively start the Safeguarded Copy function to create a protected backup at the first sign of a threat.

The IBM FlashSystem Safeguarded Copy function helps businesses recover quickly and safely from a cyberattack, which helps reduce recovery to minutes or hours. It also creates multiple recovery points for a production volume. These recovery points are called *Safeguarded Copy backups*.

The recovery data is not stored in separate regular volumes, but in a storage space that is called *Safeguarded Copy backup capacity*, which creates a logical air gap. The backups are not directly accessible by a host. The data can be used only after a backup is recovered to a separate recovery volume.

If an attack occurs, the orchestration software (IBM Spectrum Copy Data Management) helps create and identify the best Safeguarded backup to use. It also automates the process to restore data to online volumes. Because a restore action uses the same snapshot technology, it is almost instant; that is, it is much faster than the use of offline copies or copies that are stored in the cloud.

Scope

The focus of this document is to showcase the early threat detection by using IBM QRadar for threats, such as brute force and ransomware attacks on an SQL database. The database host uses storage that is mapped from IBM FlashSystem.

SQL database audit events are forwarded to IBM QRadar. By using the preconfigured rules, audit data is analyzed to detect a potential threat and proactively start Safeguarded Copy restore to a clean space environment for data validation and investigation.

As part of the discussion of early threat detection, several rules are described in this publication. Also, a sample Python script is used to start the Safeguarded Copy restore action that is provided.

This document also provides several sample control path and data path use cases.

Customers are encouraged create control path and data path use cases, IBM QRadar rules, and custom response scripts that are best suited to their environment.

Note: Although the use cases, rules, and Python script from this publication can be seen as templates, they cannot be used in a real-world environment.

The solution that is featured in this document is created by using the following products:

- IBM QRadar release 7.4.3
- IBM FlashSystem 8.4.2.0
- IBM Spectrum Copy Data Management 2.2.16

Note: All components that are described in this blueprint, such as IBM QRadar, IBM Spectrum Copy Data Management, and IBM FlashSystem are in same network segment. Suitable network planning is required if these systems are in different networks.

For more information about IBM QRadar, IBM FlashSystem, Safeguarded Copy, and IBM Spectrum Copy Data Management, see “Resources” on page 33.

Introduction

Combining the capabilities of IBM FlashSystem Safeguarded Copy, IBM QRadar enables enterprises to build comprehensive cyber resilience solutions that address the protect and recover functions of the NIST framework and the detect and respond function.

IBM QRadar can provide full protection to the enterprise data by combining IBM FlashSystem administration access logs, application logs, network or server logs, flow, packet data, and database events.

IBM FlashSystem Safeguarded Copy function

The IBM FlashSystem Safeguarded Copy feature creates safeguarded backups that are not accessible by the host system. It also protects these backups from corruption that can occur in the production environment. A Safeguarded Copy schedule can be defined to create multiple backups regularly (such as hourly or daily).

Safeguarded Copy can create backups with more frequency and capacity compared to IBM FlashCopy® volumes. Creating Safeguarded backups also affects performance less than the multiple target volumes that are created by IBM FlashCopy.

> **Note:** The Safeguarded source volume cannot be removed before the Safeguarded backups are deleted.

The Safeguarded Copy function provides backup copies to recover data if a logical corruption occurs or primary data is destroyed.

Safeguarded Copy uses a backup capacity, production volume, and recovery volume. Consider the following points:

- Backup capacity can be created for any production volume. The size of the backup capacity depends on the frequency of the backups and the duration that backups must be retained.
- The Safeguarded Copy session creates a consistency group across the source volumes to create a safeguarded backup, which stores the required data in the backup capacity.
- A recovery volume is used to restore a backup copy for host access while production continues to run on the production volume. The recovery volume is the target volume for a Safeguarded Copy recovery, which enables a previous backup copy to be accessed by a host that is attached to this volume. The recovery volume typically is thin-provisioned; however, it does not have to be thin-provisioned.

  Managing Safeguarded Copy is supported by IBM Spectrum Copy Data Management 2.2.16 or later. The management software helps to create and recover backups and define policies for expiration.

IBM Spectrum Copy Data Management

IBM Spectrum Copy Data Management is a Copy Data Management (CDM) platform that can bring modernization to an environment without disruption.

Organizations of all sizes must modernize their IT processes to enable critical new use cases, such as operational automation, DevOps, and integration of system-of-record data with Cloud compute. They are equally challenged with improving management efficiencies for long established IT processes, such as data protection, disaster recovery, reporting, and test and development.

IBM Spectrum Copy Data Management delivers “in-place” copy data management to enterprise storage arrays from IBM, NetApp, and Pure Storage, which allows the IT team to easily use its existing infrastructure and data in a manner that is efficient, automated, and scalable.

IBM Spectrum Copy Data Management also modernizes IT processes, enables key use cases; all without extra hardware.

Consider the following related concepts:

- Sites and providers

  A *site* is a user-defined grouping of providers that is generally based on location to help quickly identify and interact with data that is created through Copy Data Management jobs. Sites are assigned when registering providers.

  When creating Backup and Restore jobs, sites clearly identify where your data is replicated by location.

  Providers are physical servers that host objects and attributes. After a provider is registered in IBM Spectrum Copy Data Management, cataloging, searching, and reporting can be performed.

- Inventory job

  An inventory job definition provides the framework to collect and catalog information about objects on a registered provider. A job that is based on an Inventory job definition discovers object information, catalogs it, and populates the IBM Spectrum Copy Data Management database.

- Backup

  Create copies of your data. The RPO and copy data parameters are defined in an SLA Policy, which is then applied to the Backup job definition along with a specified activation time to meet your copy data criteria.

- Restore

  IBM Spectrum Copy Data Management technology can be used for testing, cloning, and recovering copy data.

For more information, see *IBM Spectrum Copy Data Management 2.2.16 User’s Guide*.

IBM QRadar Security Intelligence Platform

IBM QRadar Security Intelligence Platform products provide a unified architecture for integrating SIEM, log management, anomaly detection, incident forensics, and configuration and vulnerability management.

It is one of the most popular SIEM solutions on the market today. It provides powerful cyber resilience and threat detection features, such as centralized visibility, flexible deployment, automated intelligence, machine learning, and proactive threat hunting.

IBM QRadar can detect malicious patterns by using various data sources and analysis tools and techniques, including access logs, heuristics, correlation with logs from other systems (such as network logs or server logs), network flow, and packet data. Its open architecture enables third-party interoperability so that many solutions can be integrated, which makes it even more scalable and robust.

To apply the security and compliance policies, IBM QRadar administrators can perform following tasks:

- Search event data by using specific criteria and display events that match the search criteria in a results list. The columns of event data can be selected and grouped.
- Visually monitor and investigate flow data in real time or perform advanced searches to filter the displayed flows. The flow information can be viewed to determine how and what network traffic is communicated.
- View all the learned assets or search for specific assets in the environment.
- Investigate offenses, source and destination IP addresses, network behaviors, and anomalies in the network.
- Edit, create, schedule, and distribute default or custom reports.

Prerequisites

The following prerequisites must be met for the solution:

- The firewall rules between IBM QRadar and IBM FlashSystem storage are adjusted to allow traffic on 514/TCP or 514/UDP.

  > **Note:** IBM QRadar accepts incoming events on TCP/UDP protocol on port 514. The choice of protocol that is used for communication depends on an organization’s guidelines.

- A policy configuration is available that consists of various rules that are aimed to capture any specific database action.
- A running SQL database instance is available. For this solution, Microsoft SQL Server 2016 single instance database was used.
- SCDM 2.2.16 or later is available and the IBM FlashSystem storage is registered in IBM Spectrum Copy Data Management by using administrator privileges. For more information, see “Resources” on page 33.
- A scheduled job is defined in SCDM that consists of various operations, depending on the functions that are used in storage system.

- The safeguarded virtual capacity is provisioned. For more information about configuring safeguarded virtual capacity, see “Resources” on page 33.
- Understand IBM FlashSystem storage for working with volumes and safeguarded virtual capacity allotment.
- SCDM was tested with 8.4.2, but it is recommended to run 8.5.1 as specified in the SCDM user guide.

Solution overview

Data is the most important asset of any company. It empowers businesses to make decisions. These decisions determine the future of the business and eventually the organization.

Organizations can face various threats, including the following examples:

- A rogue user within the organization
- Cyberattacks that result in compromised user credentials by using spear phishing attacks
- Brute force attempts
- Ransomware

All of these threats pose grave risks to storage systems that are used for storing the data.

To track the administrative actions, the solution describes various control path use cases. To track the changes from application data, data path use cases are implemented here.

An SQL database audit is created at the database level that allows forwarding of SQL audit events to IBM QRadar. IBM QRadar understands the authorization events that are forwarded by SQL database and categorizes them correctly.

After the event classification is completed, the IBM QRadar administrator can define several rules to detect threats that are categorized under control and data path.

Upon threat detection, a cyber resiliency response is started in the form of a Python script that uses API commands to run a predefined SCDM scheduled job. The scheduled job feature of SCDM is chosen to restore the last known good configuration to clean space.

An overview of solution is shown in Figure 1.

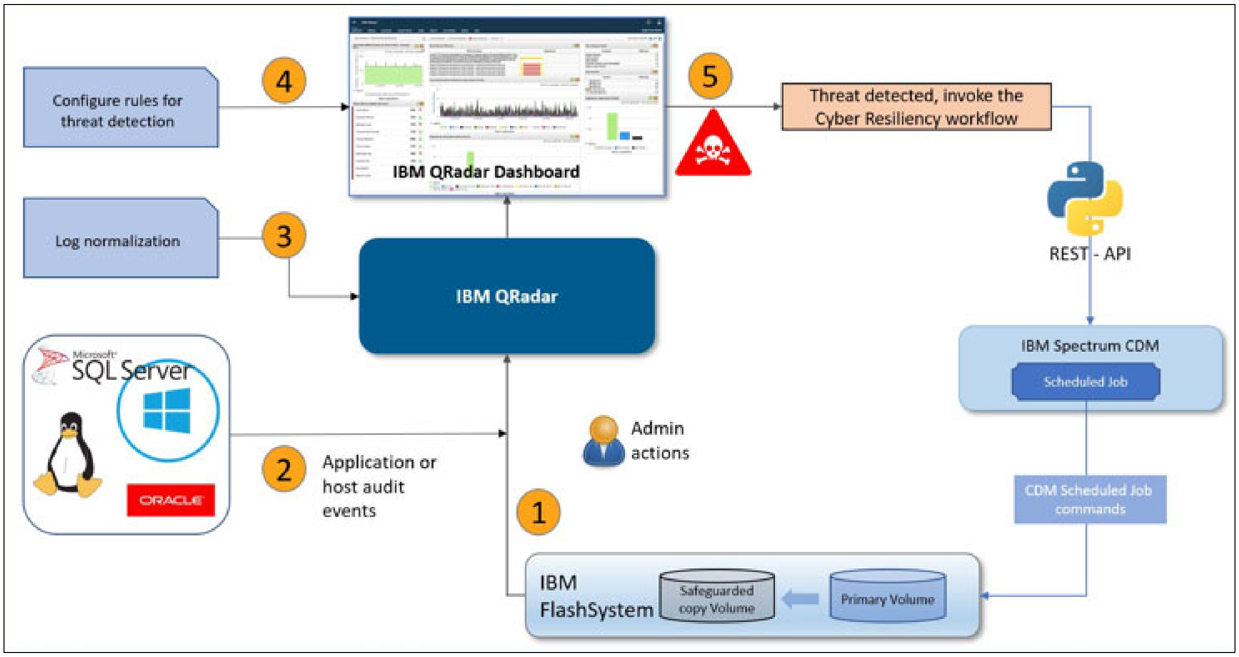

Figure 1 Solution overview

Use cases

This section describes the control and data path use cases that are considered for building the solution.

Control path

The sample control path use cases include the following examples. This list is *not* an exhaustive collection; rather, it generalizes the idea of a threat. Ultimately, the security policy of the organization defines a threat:

- A database SYS user is logged on from multiple locations or IP addresses at the same time. This use case is a classic example of compromised or shared credentials.

  A legitimate user might be oblivious of the second sessions activity under same login. Moreover, how can the remote SYS logins can be justified? What if one of the sessions is malicious?

- Operating system login activity tracking is a major task. Running commands, such as unmap, can result in dangerous consequences. It can easily go undetected and cause logical corruption on the storage volumes by overwriting the data that is in the volumes.

Data path

Data path use cases include the following scenarios:

- SQL inside attack

  Considering access to the database is already granted (whether through brute force, a DBA account that was compromised, or even a malicious insider who has access), an attacker can drop, insert, or update data and modify the data. This process can be done with a few simple SQL transactions or SQL commands.

- Encrypting the database

  This type of attack is the same as traditional ransomware attack that targets data. In this use case, critical data is encrypted by the attacker, which renders database operations useless. Then, ransom is demanded against the release of the data to bring back the database in operational mode.

A typical database infrastructure with IBM QRadar monitoring audit logs from database server is shown in Figure 2.

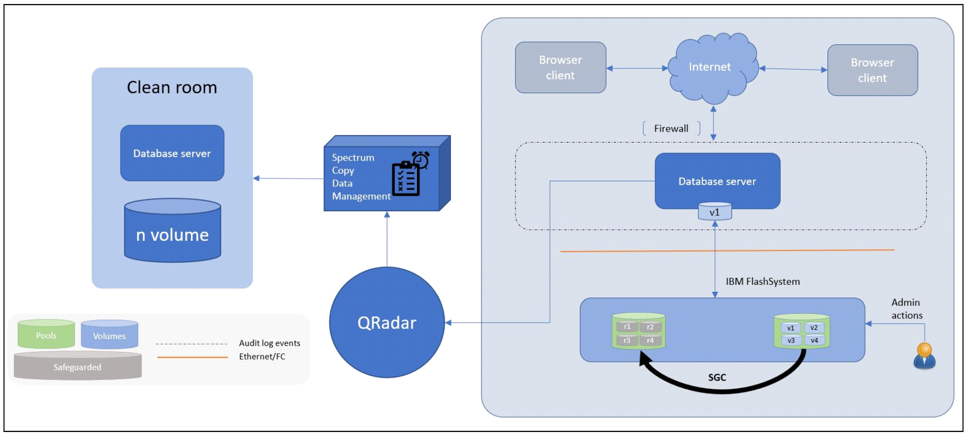

Figure 2 Sample application infrastructure

For this solution, several use cases were tracked, such as a brute force login, multiple and remote user(sa) logins, and user access to sensitive tables. The events that were captured by audit logs were sent to IBM QRadar to analyze the threat conditions.

The cyber resiliency workflow was started by IBM Spectrum Copy Data Management scheduled job to restore Safeguarded Copy restore.

Lab setup

The lab setup was created by using VMware ESXi that hosts the database server virtual machine. One storage volume was mapped to the SQL database server host by using iSCSI connectivity to store database and log files.

Audit logging was enabled on SQL database server and SQL JDBS log source was created on IBM QRadar to categorize the database audit events.

For more information about configuring SQL JDBC log source on IBM QRadar, see “Resources” on page 33.

Various components that were used in the sample setup are described next.

IBM FlashSystem

On IBM FlashSystem 9100, a Safeguarded copy child pool was created under the parent pool. The shield icon that is next to the child pool's name indicates that the child pool is the Safeguarded pool (see Figure 3).

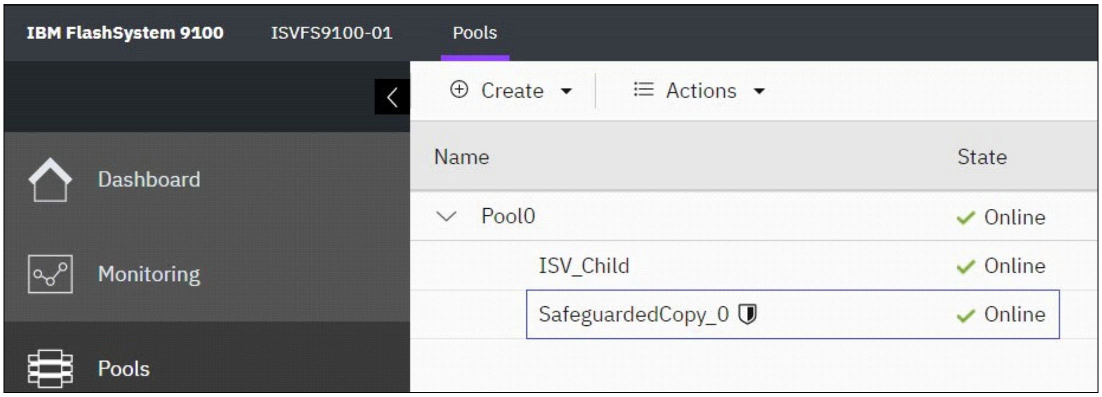

Name	State
Pool0	✓ Online
ISV_Child	✓ Online
SafeguardedCopy_0	✓ Online

Figure 3 Child pool that is configured with Safeguarded properties

To make the immutable backups of the database volumes, a volume group was created. A Safeguarded policy also was created, and the database volumes were added to the volume group (see Figure 4).

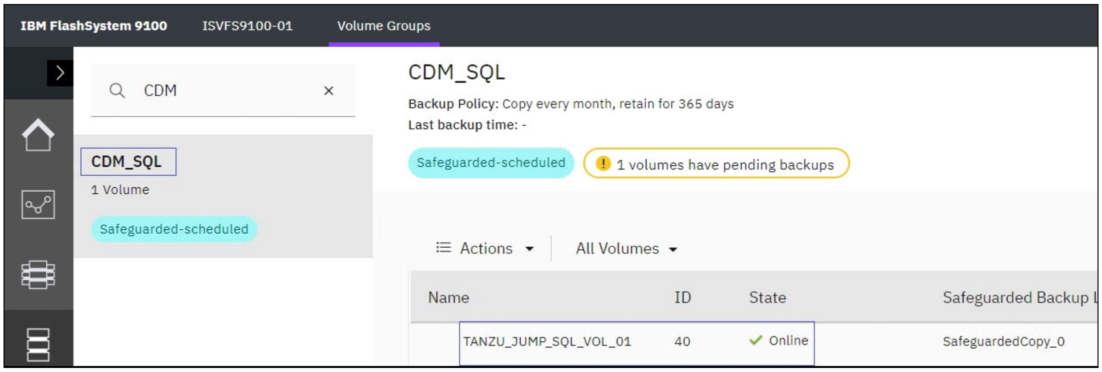

CDM_SQL

Backup Policy: Copy every month, retain for 365 days

Last backup time: -

Safeguarded-scheduled

1 volumes have pending backups

Name	ID	State	Safeguarded Backup L
TANZU_JUMP_SQL_VOL_01	40	✓ Online	SafeguardedCopy_0

Figure 4 Volume group with Safeguarded copy policy

IBM Spectrum Copy Data Management

The IBM Spectrum Copy Data Management workflow includes registering a provider, cataloging data, searching for objects, generating reports, and copying and using data.

Registering providers

Add providers, such as application servers, IBM, or NetApp storage devices, or VMware ESX resources to the Inventory by registering them. Before registering providers, create a site to assign to your provider. A *site* is a user-defined grouping of providers that is generally based on location.

Registering IBM FlashSystem 9100 storage system

Complete the following steps to register the IBM FlashSystem 9100 storage system:

1. Click the **Configure** tab. In the Views window, select **Sites & Providers** and then, select the **Providers** tab.
2. In the Provider Browser window, select **IBM Spectrum Virtualize**.
3. Right-click **IBM Spectrum Virtualize**. Then, click **Register**. The Register dialog window opens.
4. Complete the fields in the dialog window. Select **New** to add credentials if they are not yet added through identities. We used a superuser account to register IBM FlashSystem 9100 storage system (see Figure 5).

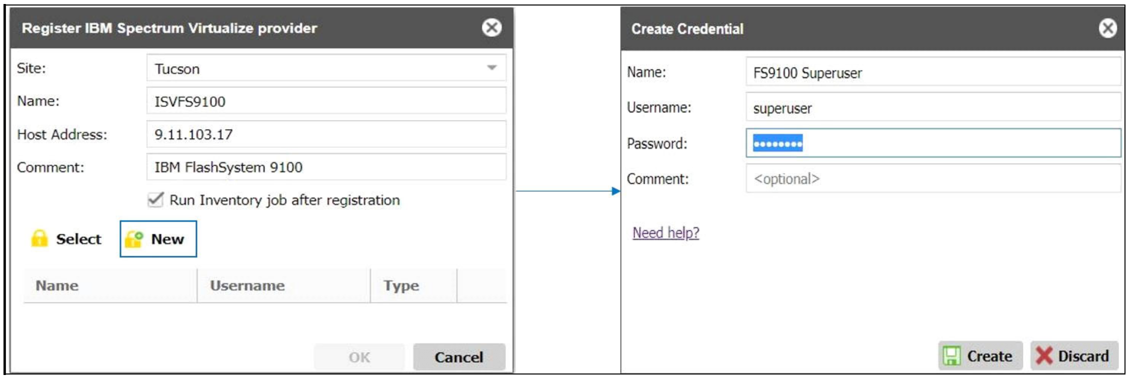

Register IBM Spectrum Virtualize provider

Site: Tucson

Name: ISVFS9100

Host Address: 9.11.103.17

Comment: IBM FlashSystem 9100

☑ Run Inventory job after registration

Name	Username	Type

Create Credential

Name: FS9100 Superuser

Username: superuser

Password: ••••••••

Comment: <optional>

Figure 5 Registering FS9100 storage system

5. Click **OK**. IBM Spectrum Copy Data Management first confirms that a network connection exists and then, adds the provider to the database.

Registering SQL database server

Complete the following steps to register the SQL database server:

1. Click the **Configure** tab. In the **Views** window, select **Sites & Providers** and then, select the **Providers** tab.
2. In the Provider Browser window, select **Application Server**.
3. Right-click **Application Server** and then, click **Register**. The Register Application Server dialog opens.
4. Select **SQL** as the Application Type.
5. Complete the fields in the dialog window. Select **New** to add credentials if they are not yet added through identities. We used a Windows server's Administrator account because mixed mode authentication is configured on SQL database (see Figure 6).

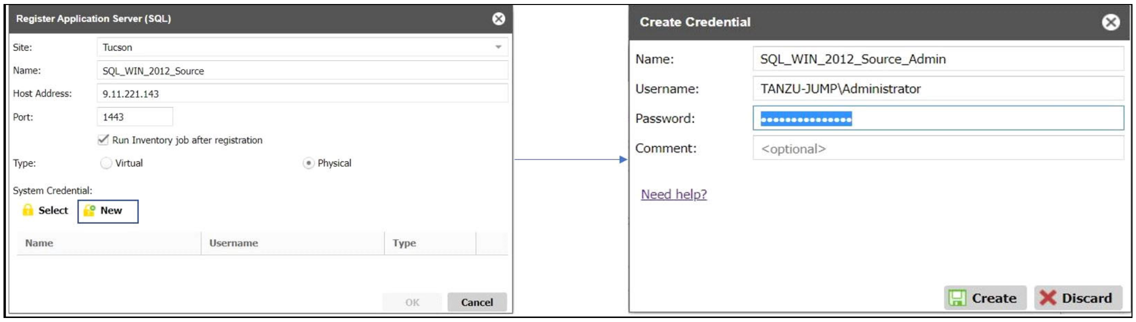

Register Application Server (SQL)

Site: Tucson

Name: SQL_WIN_2012_Source

Host Address: 9.11.221.143

Port: 1443

☑ Run Inventory job after registration

Type: ○ Virtual ● Physical

System Credential:

Name	Username	Type

Create Credential

Name: SQL_WIN_2012_Source_Admin

Username: TANZU-JUMP\Administrator

Password: •••••••••••••••

Comment: <optional>

Figure 6   Registering SQL database server

6. Click **OK**. IBM Spectrum Copy Data Management first confirms that a network connection exists and then, adds the provider to the database.

Configuring an SLA policy

Complete the following steps to configure an SLA policy:

1. Click the **Configure** tab. In the Views window, select **SLA Policies**. The All SLA Policies window opens.
2. In the All SLA Policies window, click **New**. The New SLA Policies window opens.
3. Select **IBM Spectrum Virtualize** and **Add Safeguarded Copy**.
4. Enter a name and a meaningful description for the SLA Policy.
5. Select the source icon and define the recovery point objective to determine the minimum frequency and interval with which backups must be made. In the Frequency field, select Minutes, Hourly, Daily, Weekly, or Monthly and then, set the interval in the Interval field.
6. In the Associated Safeguarded Volume Group window, expand the storage device and select the volume group that you want to back up as a Safeguarded Copy. Any volume group that you want to back up as a Safeguarded Copy must be a volume that belongs to one of these groups. If it is not a member of any of these groups, it does not back up as a Safeguarded Copy.

> **Note:** The Associated Safeguarded Volume Group lists only those volume groups that have the Safeguarded Copy policy applied on the storage array side.

7. In the Options window, set the following Safeguarded Copy subpolicy options (see Figure 7):

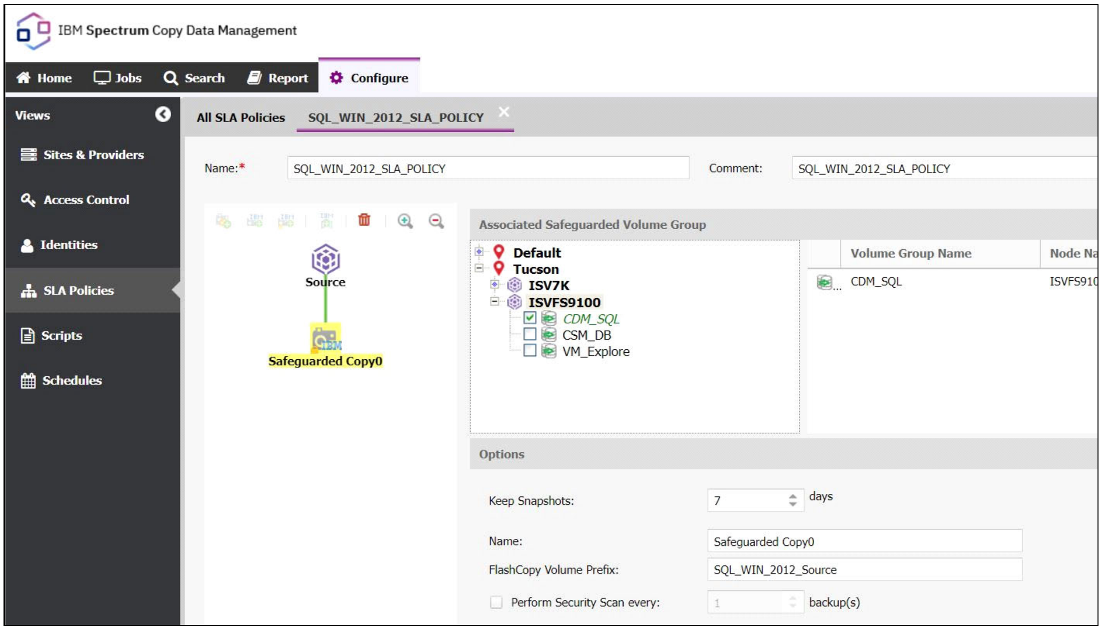

Name:* SQL_WIN_2012_SLA_POLICY

Comment: SQL_WIN_2012_SLA_POLICY

Volume Group Name	Node Na
CDM_SQL	ISVFS910

Keep Snapshots: 7 days

Name: Safeguarded Copy0

FlashCopy Volume Prefix: SQL_WIN_2012_Source

□ Perform Security Scan every: 1 backup(s)

Figure 7 Configuring SLA policy

- Keep Snapshots

  After a specific number of snapshot instances are created for a resource, older instances are purged from the storage controller. Enter the age of the snapshot instances to purge in the Days field, or the number of instances to keep in the Snapshots field.

- Name

  Enter an optional name to replace the default FlashCopy subpolicy name that is displayed in IBM Spectrum Copy Data Management. The default name is `Safeguarded Copy0`.

- FlashCopy Volume Prefix

  Enter an optional label to identify the FlashCopy. This label is added as a prefix to the FlashCopy name that is created by the job.

- Perform Security Scan

  You can enable this scan and select your security scan servers. This setting allows to scan for every backup number you specified.

8. Click **Finish**.

Configuring an SQL Backup Job

Complete the following steps to configure an SQL backup job:

1. Click the **Jobs** tab. Expand the **Database** folder and then, select **SQL**.
2. Click **New** and then, select **Backup**. The job editor opens.
3. Select a Standalone or Failover Cluster or Always On Availability Group workflow template.
4. Enter a name and meaningful description for your job definition.
5. From the list of available sites, select one or more resources to back up. Expand **Servers** to view associated application databases.
6. Select an SLA Policy that meets your backup data criteria.
7. Click the job definition's associated Schedule Time field and select **Enable Schedule** to set a time to run the SLA Policy. If a schedule is not enabled, run the job on-demand through the Jobs tab.

To create the job definition by using default options, click **Create Job**. The job runs as defined by your triggers, or can be run manually from the Jobs tab (see Figure 8).

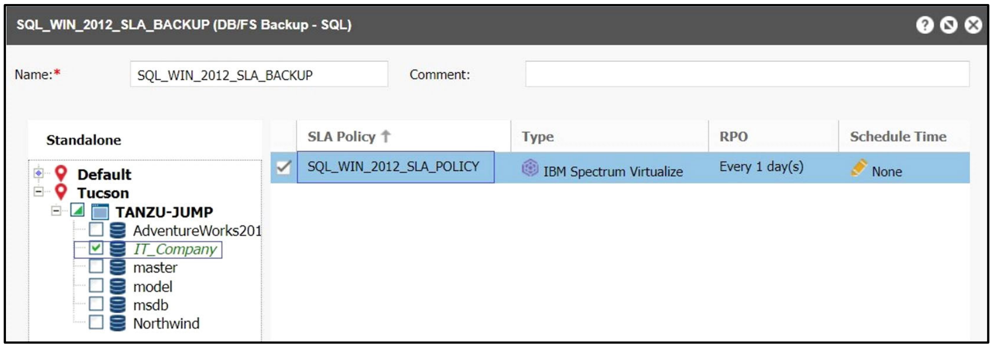

SQL_WIN_2012_SLA_BACKUP (DB/FS Backup - SQL)

Name:* SQL_WIN_2012_SLA_BACKUP Comment:

Standalone

- Default
- Tucson
  - ☑ TANZU-JUMP
    - □ AdventureWorks201
    - ☑ IT_Company
    - □ master
    - □ model
    - □ msdb
    - □ Northwind

	SLA Policy ↑	Type	RPO	Schedule Time
☑	SQL_WIN_2012_SLA_POLICY	IBM Spectrum Virtualize	Every 1 day(s)	None

Figure 8 Configuring SQL backup job

Configuring an SQL Restore Job

Complete the following steps to configure an SQL restore job:

1. Click the **Jobs** tab. Expand the **Database** folder and then, select **SQL**.
2. Click **New** and then, select **Restore**. The job editor opens.
3. Enter a name and a meaningful description for your job definition.
4. Select **Microsoft SQL (Standalone and Failover Cluster)** or **Microsoft SQL (Always On Availability Group)**.
5. Select a template. The following options are available:
   - Instant Database Restore for Microsoft SQL Standalone/Failover Cluster and Always On Availability Group jobs
   - Instant Disk Restore
   - Instant Seeding for Microsoft SQL Always On Availability Group jobs
6. Click **Source**. From the drop-down menu, select **Application Browse** to select a source site and an application server to view available database recovery points.

   Then, select **Resources** and change the order in which the resources are recovered by dragging and dropping the resources in the grid (see Figure 9).

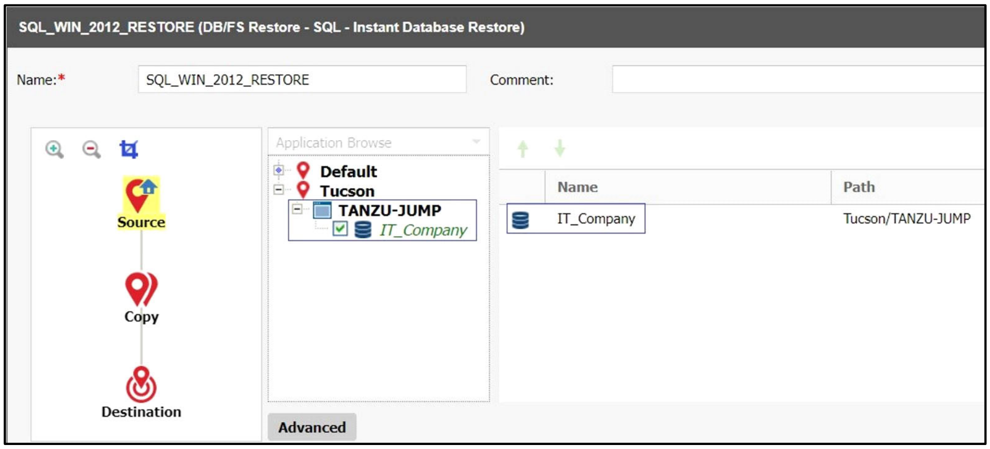

SQL_WIN_2012_RESTORE (DB/FS Restore - SQL - Instant Database Restore)

Name:* SQL_WIN_2012_RESTORE

Comment:

Application Browse

- Default
- Tucson
  - TANZU-JUMP
    - ☑ IT_Company

Name	Path
IT_Company	Tucson/TANZU-JUMP

Figure 9 Adding a source

7. Click **Copy**. Then, click **Select Backup Date Range** to provide a date range for which recovery points are displayed. Select a range of dates by selecting a date for Start Time and a date for End Time. After a range is established, click **OK**.

Sites that contain copies of the selected data are displayed. Select a site. By default, the latest copy of your data is used. To choose a specific version, select a site and then, click **Select Version** (see Figure 10).

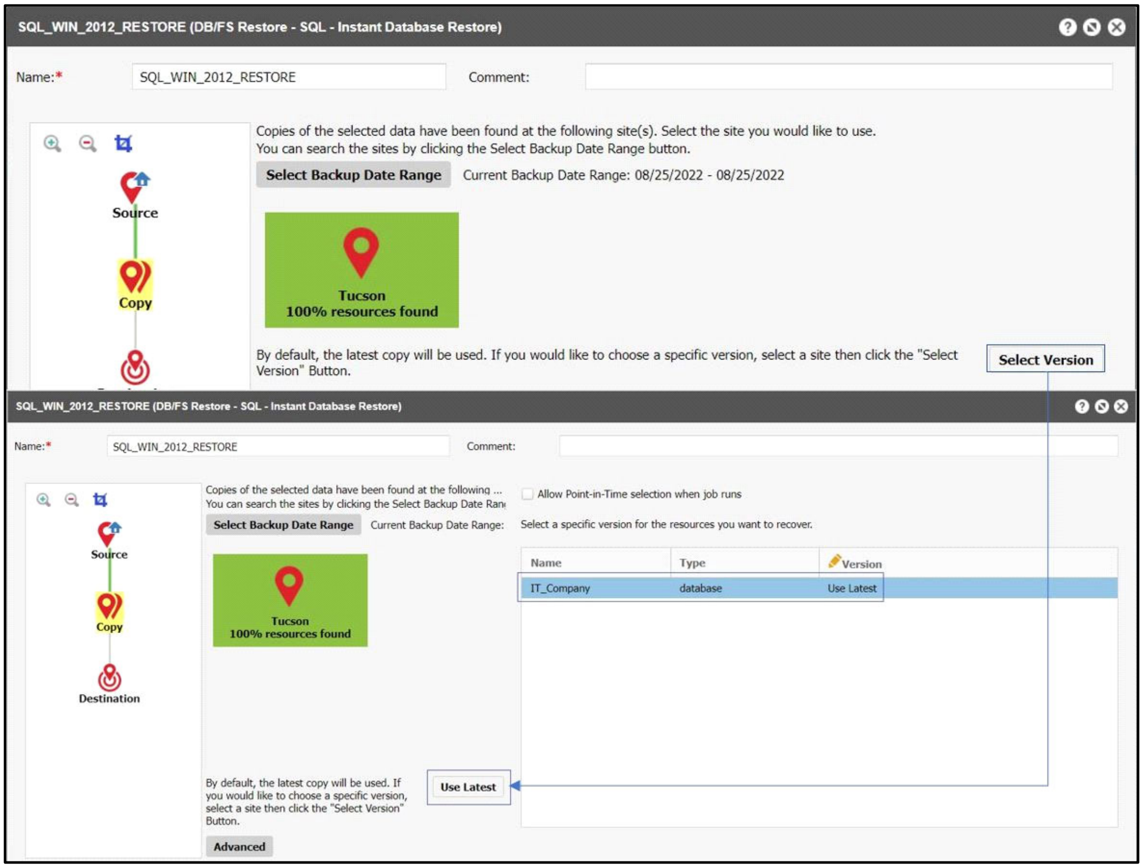

SQL_WIN_2012_RESTORE (DB/FS Restore - SQL - Instant Database Restore)

Name:* SQL_WIN_2012_RESTORE Comment:

Copies of the selected data have been found at the following site(s). Select the site you would like to use.
You can search the sites by clicking the Select Backup Date Range button.

Select Backup Date Range Current Backup Date Range: 08/25/2022 - 08/25/2022

Tucson
100% resources found

By default, the latest copy will be used. If you would like to choose a specific version, select a site then click the "Select Version" Button.

SQL_WIN_2012_RESTORE (DB/FS Restore - SQL - Instant Database Restore)

Name:* SQL_WIN_2012_RESTORE Comment:

Copies of the selected data have been found at the following ...
You can search the sites by clicking the Select Backup Date Ran

Select Backup Date Range Current Backup Date Range:

□ Allow Point-in-Time selection when job runs

Select a specific version for the resources you want to recover.

Name	Type	Version
IT_Company	database	Use Latest

Tucson
100% resources found

By default, the latest copy will be used. If you would like to choose a specific version, select a site then click the "Select Version" Button.

Figure 10 Selecting the Copy option

8. Click **Destination**. Select a target site and an associated Microsoft SQL database. Click the **New database name** field to enter an optional, alternative name for the database (see Figure 11).

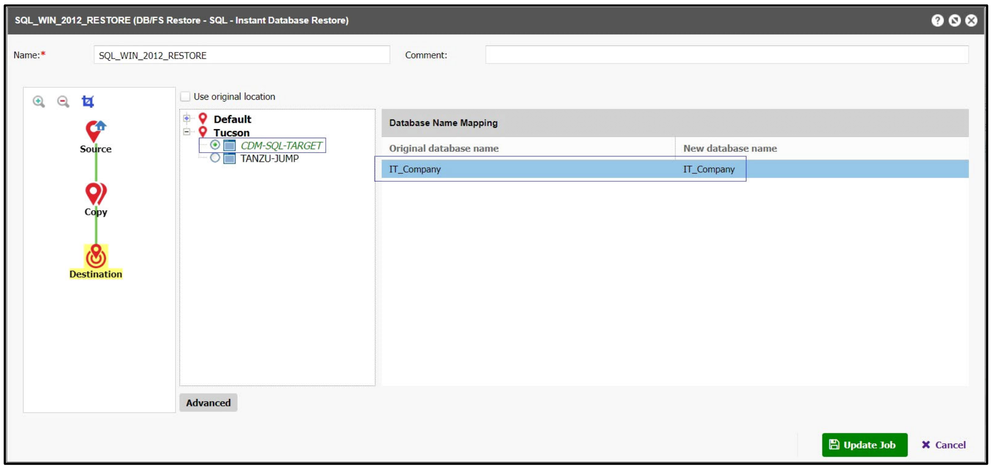

SQL_WIN_2012_RESTORE (DB/FS Restore - SQL - Instant Database Restore)

Name:* SQL_WIN_2012_RESTORE Comment:

□ Use original location

- Default
- Tucson
  - ● CDM-SQL-TARGET
  - ○ TANZU-JUMP

Database Name Mapping

Original database name	New database name
IT_Company	IT_Company

Figure 11 Selecting Target Server and Database Name

9. To create the job definition by using default options, click **Create Job**. To update a job, click **Update Job**. The job can be run manually from the Jobs tab.

Running a backup job

Complete the following steps to run a backup job:

1. Click the **Jobs** tab. Expand the **Database** folder and then, select **SQL**.
2. Select the backup job name that was created and select the **Start** option (see Figure 12).

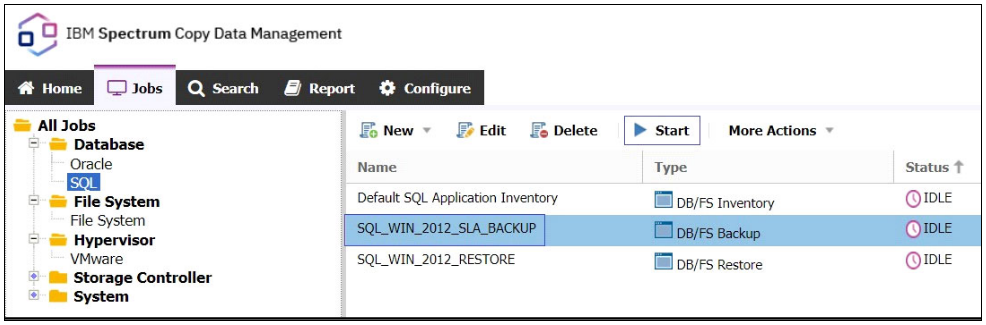

Name	Type	Status
Default SQL Application Inventory	DB/FS Inventory	IDLE
SQL_WIN_2012_SLA_BACKUP	DB/FS Backup	IDLE
SQL_WIN_2012_RESTORE	DB/FS Restore	IDLE

Figure 12   Running a backup job

After it is completed successfully, the backup operation reflects the status as `COMPLETED` (see Figure 13).

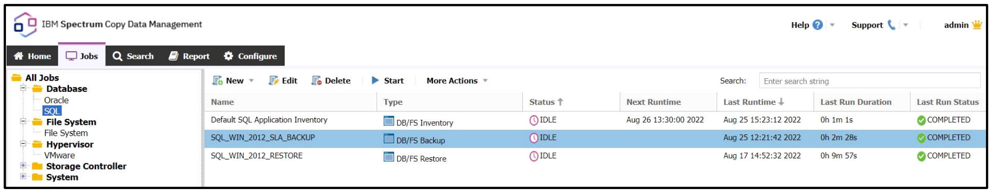

Name	Type	Status	Next Runtime	Last Runtime	Last Run Duration	Last Run Status
Default SQL Application Inventory	DB/FS Inventory	IDLE	Aug 26 13:30:00 2022	Aug 25 15:23:12 2022	0h 1m 1s	COMPLETED
SQL_WIN_2012_SLA_BACKUP	DB/FS Backup	IDLE		Aug 25 12:21:42 2022	0h 2m 28s	COMPLETED
SQL_WIN_2012_RESTORE	DB/FS Restore	IDLE		Aug 17 14:52:32 2022	0h 9m 57s	COMPLETED

Figure 13   Backup status

Threat detection in IBM QRadar

Threats are detected by the IBM QRadar rules engine. The rules engine applies various conditions on the normalized events to determine the threat.

After the threat is detected, its severity can be determined. Then, a response can be generated that is based on properties that are extracted from the events.

In addition to the response, the IBM QRadar administrator can choose to raise an offense.

Defining a custom action in IBM QRadar

Complete the following steps to define a custom action in IBM QRadar:

1. Log in to IBM QRadar by using administrator's privileges. Click the **Admin** tab and then, click **Custom Actions** → **Define Actions**. Then, click **Add** (see Figure 14).

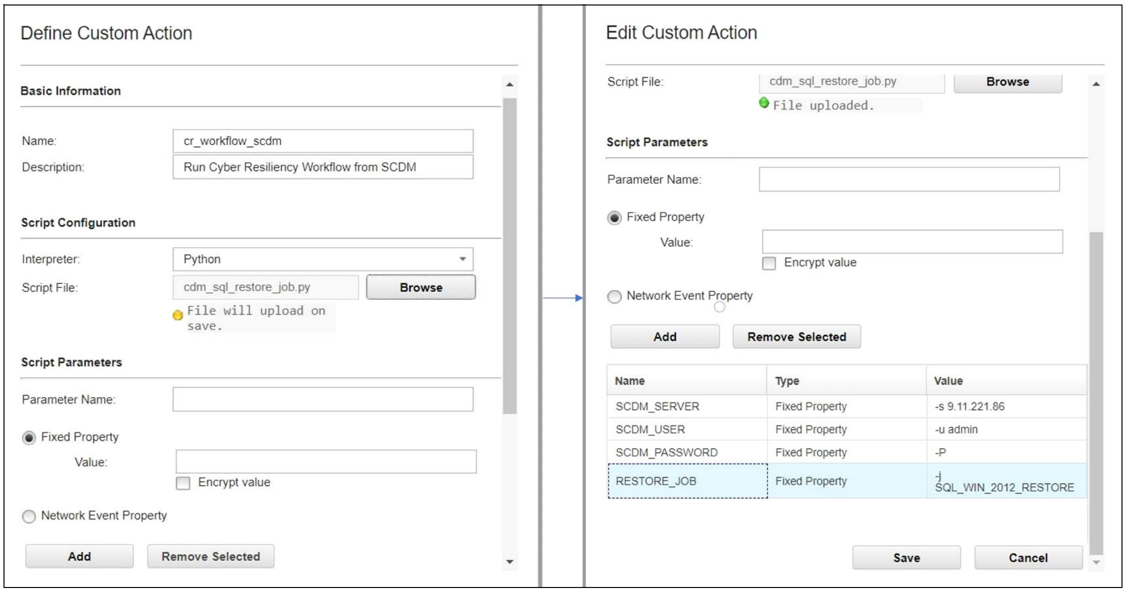

Define Custom Action

Basic Information

Name: cr_workflow_scdm
Description: Run Cyber Resiliency Workflow from SCDM

Script Configuration

Interpreter: Python
Script File: cdm_sql_restore_job.py
File will upload on save.

Script Parameters

Parameter Name:
● Fixed Property
Value:
□ Encrypt value
○ Network Event Property

Edit Custom Action

Script File: cdm_sql_restore_job.py
File uploaded.

Script Parameters

Parameter Name:
● Fixed Property
Value:
□ Encrypt value
○ Network Event Property

Name	Type	Value
SCDM_SERVER	Fixed Property	-s 9.11.221.86
SCDM_USER	Fixed Property	-u admin
SCDM_PASSWORD	Fixed Property	-P
RESTORE_JOB	Fixed Property	-j SQL_WIN_2012_RESTORE

Figure 14 Custom Action Definition windows

The property and value details that are used in this solution are listed in Table 1.

Table 1 Properties and values

Name	Property
SCDM_SERVER	`-s scdm_server_IP_address`
SCDM_USER	`-u scdm_user_name`
SCDM_PASSWORD	`-P scdm_user_password`
RESTORE_JOB	`-j restore_job_name`

2. Click **OK** to save the changes and acknowledge the dialog box to deploy the script.

3. Return to the Admin tab. Notice the messages about undeployed changes (see Figure 15). Click **Deploy Changes** to deploy the changes.

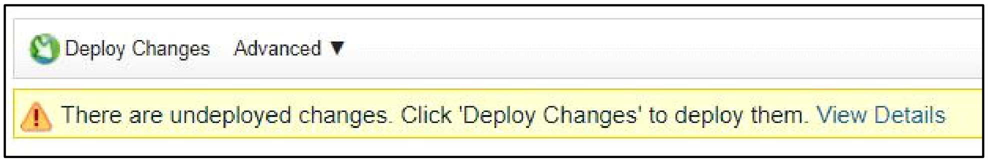

There are undeployed changes. Click 'Deploy Changes' to deploy them. View Details

Figure 15 Deploying changes after custom action is defined

4. Click the **Log Activity** tab and click the **Rules** drop-down. Select the **Rules** option. Click **Next** in the Custom Rule Wizard Welcome window. The sample rule that is displayed is based on a drop database event that was captured by SQL audit log and forwarded to IBM QRadar.

5. Select **Events** as the source to generate the rule and then, click **Next** (see Figure 16).

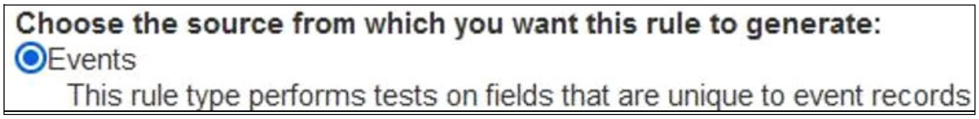

Choose the source from which you want this rule to generate:

● Events

This rule type performs tests on fields that are unique to event records

Figure 16 Rule wizard

6. The Rules Test Stack Editor window opens. Use the event matches criteria to filter the rules and click the green (**+**) icon to add the first rule. The text in bold act as hyperlinks to select suitable properties (see Figure 17).

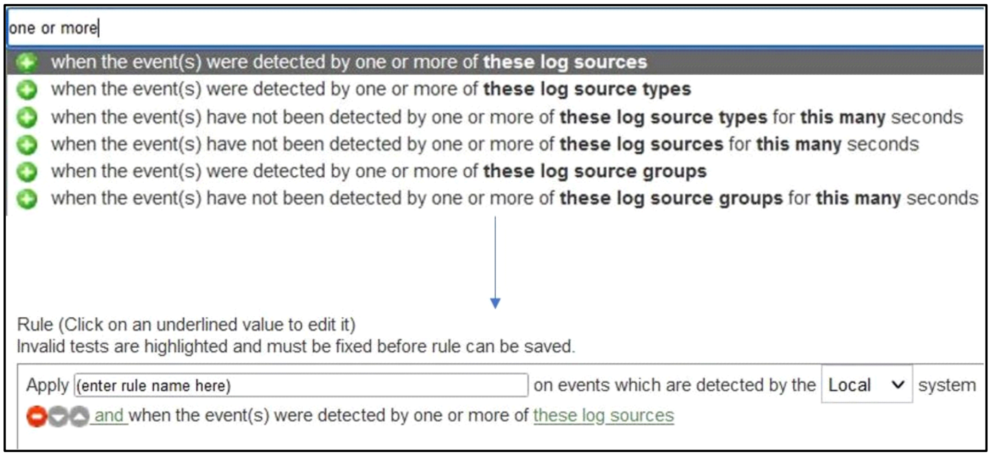

one or more

- when the event(s) were detected by one or more of **these log sources**
- when the event(s) were detected by one or more of **these log source types**
- when the event(s) have not been detected by one or more of **these log source types** for **this many** seconds
- when the event(s) have not been detected by one or more of **these log sources** for **this many** seconds
- when the event(s) were detected by one or more of **these log source groups**
- when the event(s) have not been detected by one or more of **these log source groups** for **this many** seconds

Rule (Click on an underlined value to edit it)
Invalid tests are highlighted and must be fixed before rule can be saved.

Apply (enter rule name here) on events which are detected by the Local system

and when the event(s) were detected by one or more of these log sources

Figure 17 Rule Test Stack Editor window

7. Click these log sources to select the property and search for the log source identifier (Tanzu, in our solution). Select the **TANZU-SQL_JDBC** value from the search results and then, click **Add+**. Then, click **Submit** (see Figure 18 and Figure 19).

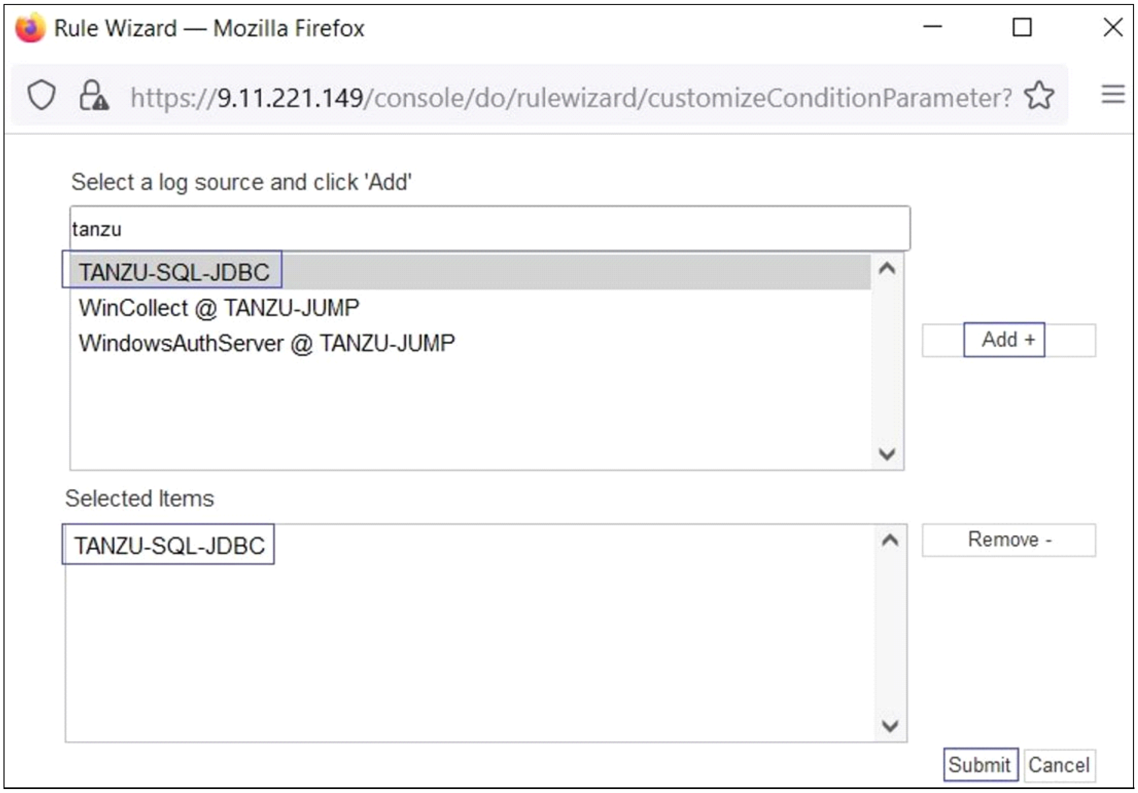

Rule Wizard — Mozilla Firefox

https://9.11.221.149/console/do/rulewizard/customizeConditionParameter?

Select a log source and click 'Add'

tanzu

TANZU-SQL-JDBC
WinCollect @ TANZU-JUMP
WindowsAuthServer @ TANZU-JUMP

Selected Items

TANZU-SQL-JDBC

Figure 18 Selecting the log source

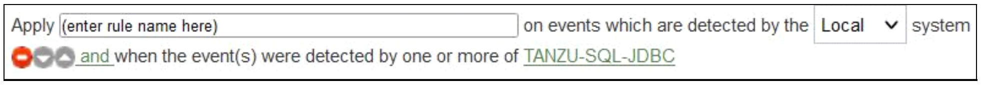

Apply (enter rule name here) on events which are detected by the Local system
and when the event(s) were detected by one or more of TANZU-SQL-JDBC

Figure 19 Event added by log source

8. Use the filter text and add the next rule (see Figure 20).

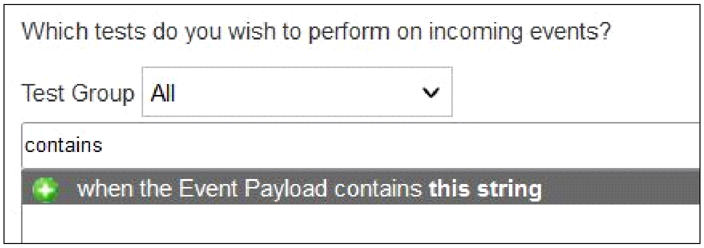

Which tests do you wish to perform on incoming events?

Test Group All

contains

when the Event Payload contains **this string**

Figure 20 Adding event by using payload string

9. Click this string and enter the payload string as the property value; in this example, `drop database IT_Company`. Here, `IT_Company` is the name of the critical database (see Figure 21, Figure 22, and Figure 23).

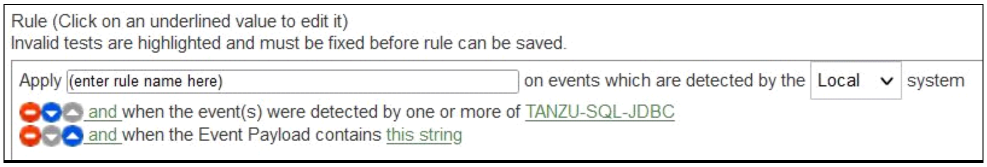

Rule (Click on an underlined value to edit it)
Invalid tests are highlighted and must be fixed before rule can be saved.

Apply (enter rule name here) on events which are detected by the Local system
and when the event(s) were detected by one or more of TANZU-SQL-JDBC
and when the Event Payload contains this string

Figure 21 Selecting payload string

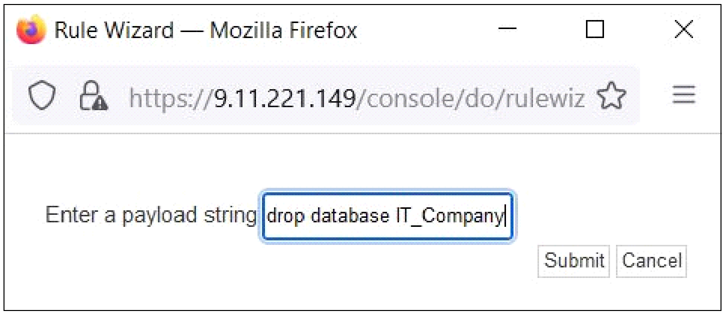

Rule Wizard — Mozilla Firefox
https://9.11.221.149/console/do/rulewiz

Enter a payload string drop database IT_Company

Figure 22 Property value drop database IT_Company

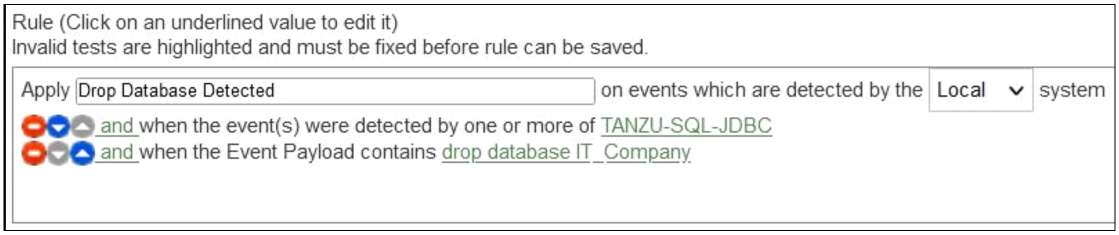

Rule (Click on an underlined value to edit it)
Invalid tests are highlighted and must be fixed before rule can be saved.

Apply Drop Database Detected on events which are detected by the Local system
and when the event(s) were detected by one or more of TANZU-SQL-JDBC
and when the Event Payload contains drop database IT_Company

Figure 23 Event added by payload string

The rule also is assigned a name for identifying its purpose and a group is chosen for which this rule can become a member. In our example, the rule was made part of the Threats group for rule categorization (see Figure 24).

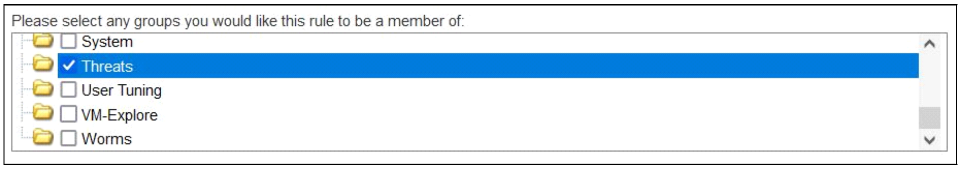

Please select any groups you would like this rule to be a member of:

- ☐ System
- ☑ Threats
- ☐ User Tuning
- ☐ VM-Explore
- ☐ Worms

Figure 24 Selecting group for a rule

10. Click **Next** to open the Rule Response window, which is divided into the following sections:

- Rule Action

  Under the Rule action section, configure the Severity, Credibility, and Relevance properties to determine the suitableness of the event. Use the Annotate even field to provide specific annotations to the event (see Figure 25).

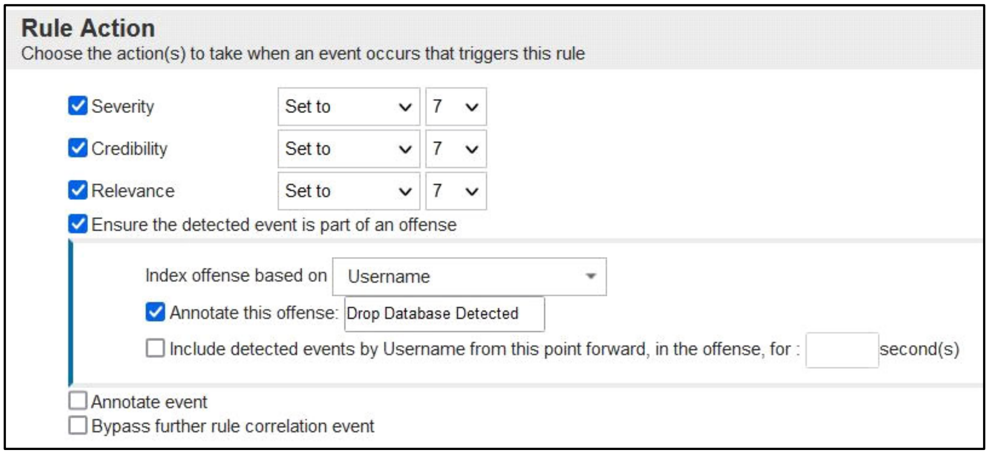

Rule Action

Choose the action(s) to take when an event occurs that triggers this rule

☑ Severity Set to 7

☑ Credibility Set to 7

☑ Relevance Set to 7

☑ Ensure the detected event is part of an offense

Index offense based on Username

☑ Annotate this offense: Drop Database Detected

□ Include detected events by Username from this point forward, in the offense, for : second(s)

□ Annotate event

□ Bypass further rule correlation event

Figure 25 Rule Action window

- Rule Response

  This section is used to configure a rule’s response, such as, the use of a property to base the offense on and the execution of custom action. It also allows configuring a new event with a user-defined name and description to indicate that the rule was triggered (see Figure 26 on page 24).

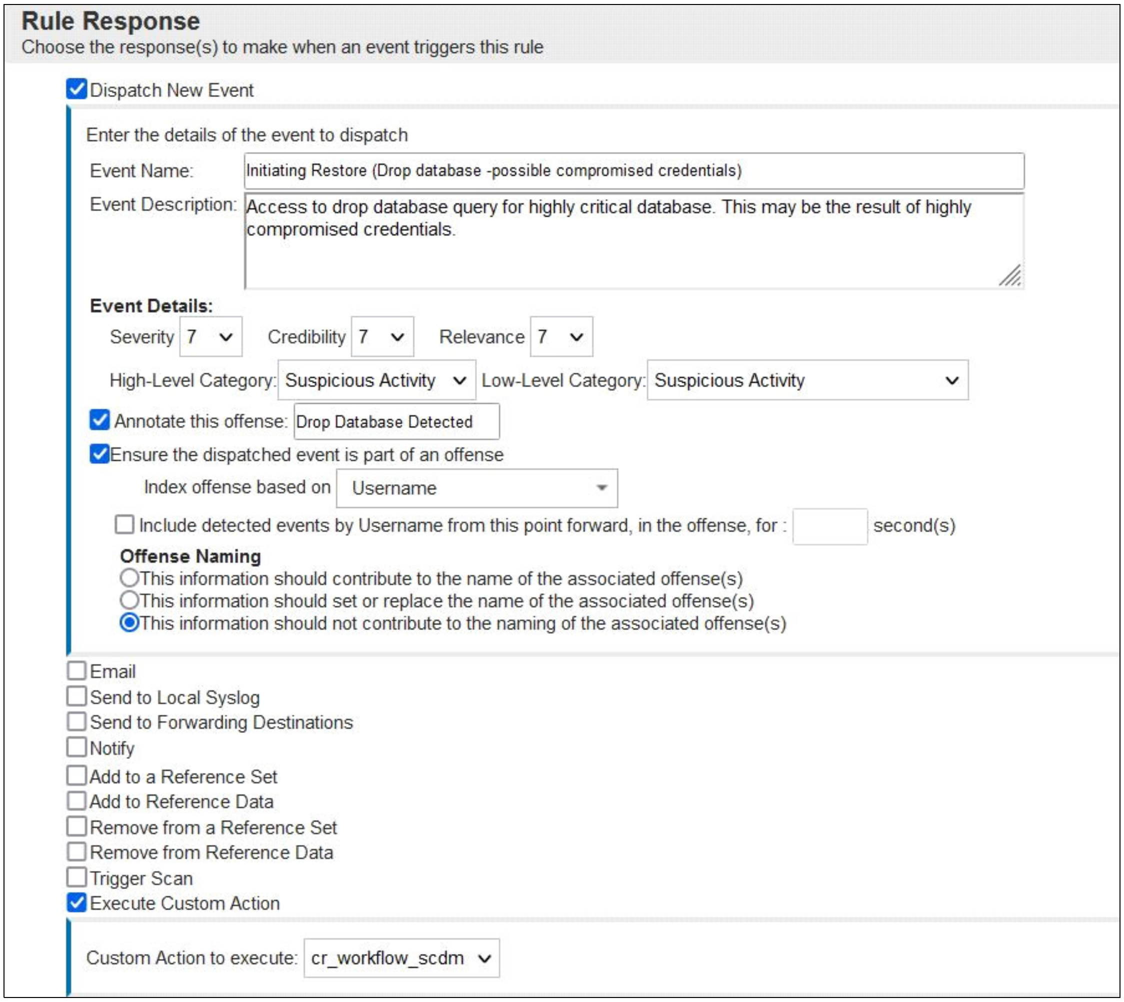

Rule Response

Choose the response(s) to make when an event triggers this rule

☑ Dispatch New Event

Enter the details of the event to dispatch

Event Name: Initiating Restore (Drop database -possible compromised credentials)

Event Description: Access to drop database query for highly critical database. This may be the result of highly compromised credentials.

Event Details:

Severity 7 | Credibility 7 | Relevance 7

High-Level Category: Suspicious Activity | Low-Level Category: Suspicious Activity

☑ Annotate this offense: Drop Database Detected

☑ Ensure the dispatched event is part of an offense

Index offense based on Username

□ Include detected events by Username from this point forward, in the offense, for : ____ second(s)

Offense Naming

○ This information should contribute to the name of the associated offense(s)
○ This information should set or replace the name of the associated offense(s)
● This information should not contribute to the naming of the associated offense(s)

□ Email
□ Send to Local Syslog
□ Send to Forwarding Destinations
□ Notify
□ Add to a Reference Set
□ Add to Reference Data
□ Remove from a Reference Set
□ Remove from Reference Data
□ Trigger Scan
☑ Execute Custom Action

Custom Action to execute: cr_workflow_scdm

Figure 26 Rule Response window

- ◦ Response Limiter

  This parameter limits the response by the rule. In this example, the rule response was set to single execution for every 30 minutes (see Figure 27).

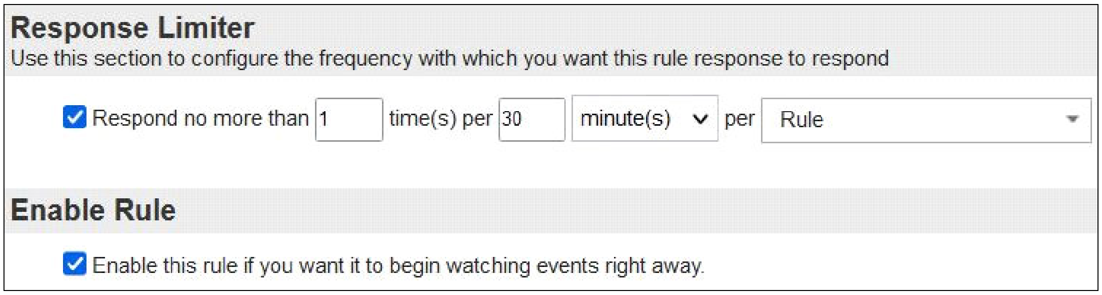

Response Limiter

Use this section to configure the frequency with which you want this rule response to respond

☑ Respond no more than 1 time(s) per 30 minute(s) per Rule

Enable Rule

☑ Enable this rule if you want it to begin watching events right away.

Figure 27 Response Limiter window

- Enable Rule

  Multiple rules can be configured for testing different conditions to detect the threat. A single rule can be enabled by using this property (see Figure 27).

The final window of the Rule Wizard shows the summary of the rule that was created (see Figure 28).

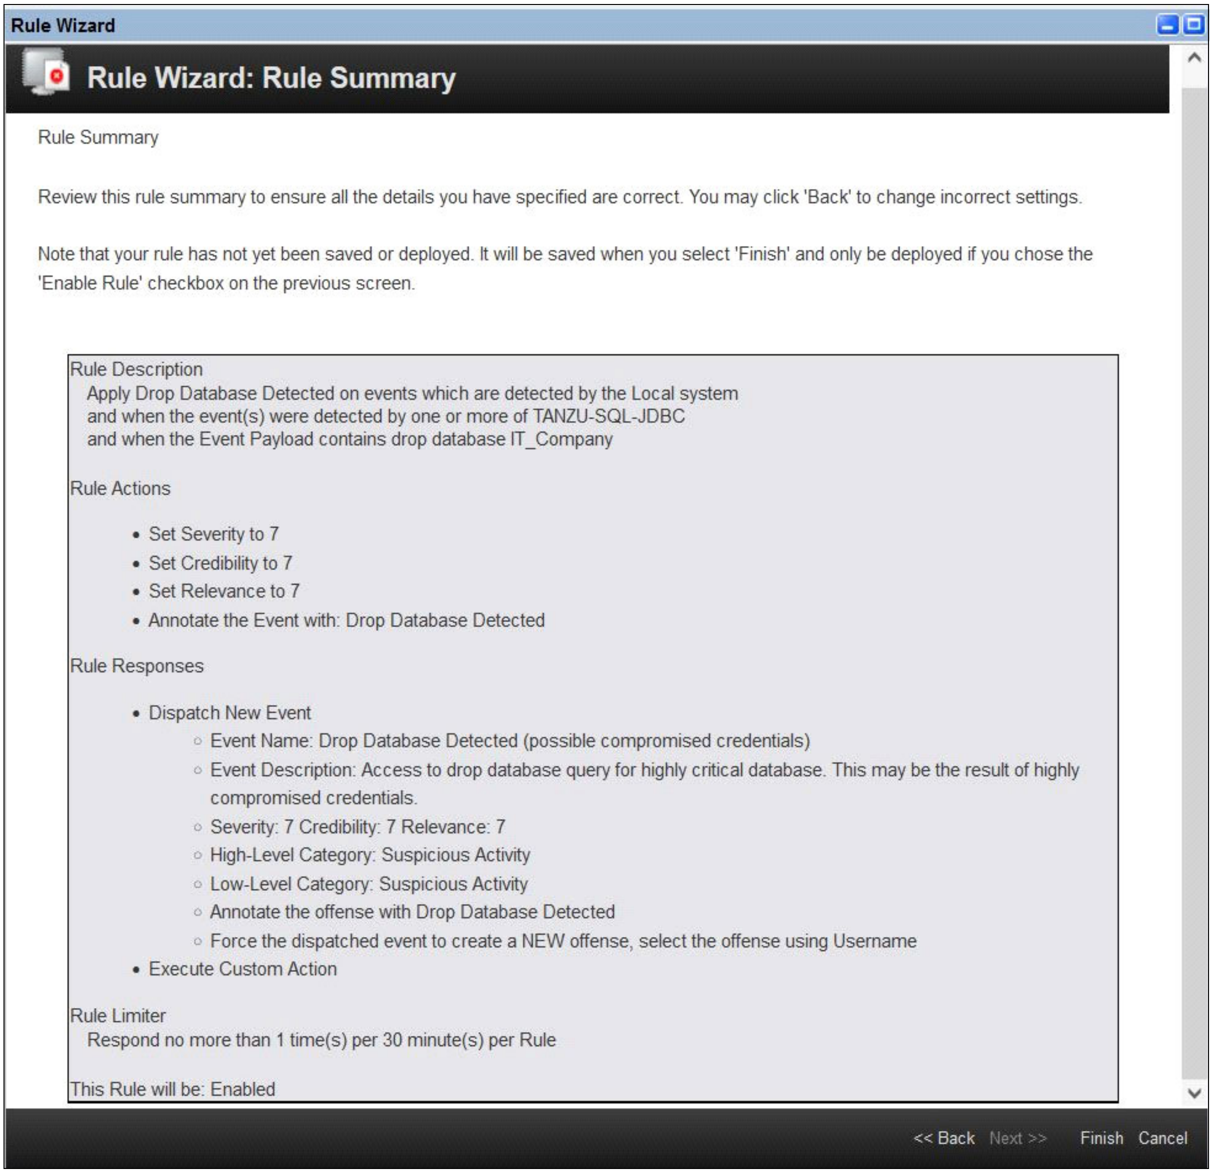

Rule Wizard: Rule Summary

Rule Summary

Review this rule summary to ensure all the details you have specified are correct. You may click 'Back' to change incorrect settings.

Note that your rule has not yet been saved or deployed. It will be saved when you select 'Finish' and only be deployed if you chose the 'Enable Rule' checkbox on the previous screen.

Rule Description
Apply Drop Database Detected on events which are detected by the Local system
and when the event(s) were detected by one or more of TANZU-SQL-JDBC
and when the Event Payload contains drop database IT_Company

Rule Actions

- Set Severity to 7
- Set Credibility to 7
- Set Relevance to 7
- Annotate the Event with: Drop Database Detected

Rule Responses

- Dispatch New Event
  - Event Name: Drop Database Detected (possible compromised credentials)
  - Event Description: Access to drop database query for highly critical database. This may be the result of highly compromised credentials.
  - Severity: 7 Credibility: 7 Relevance: 7
  - High-Level Category: Suspicious Activity
  - Low-Level Category: Suspicious Activity
  - Annotate the offense with Drop Database Detected
  - Force the dispatched event to create a NEW offense, select the offense using Username
- Execute Custom Action

Rule Limiter
Respond no more than 1 time(s) per 30 minute(s) per Rule

This Rule will be: Enabled

Figure 28 Rule Summary window

11. Validate the selection that was made and click **Finish** to save the rule and close the wizard.

Detect alter base events rule summary

The IBM QRadar rule to detect alter database events also can be defined in IBM QRadar to work with the events that are received from SQL audit logs (see Figure 29).

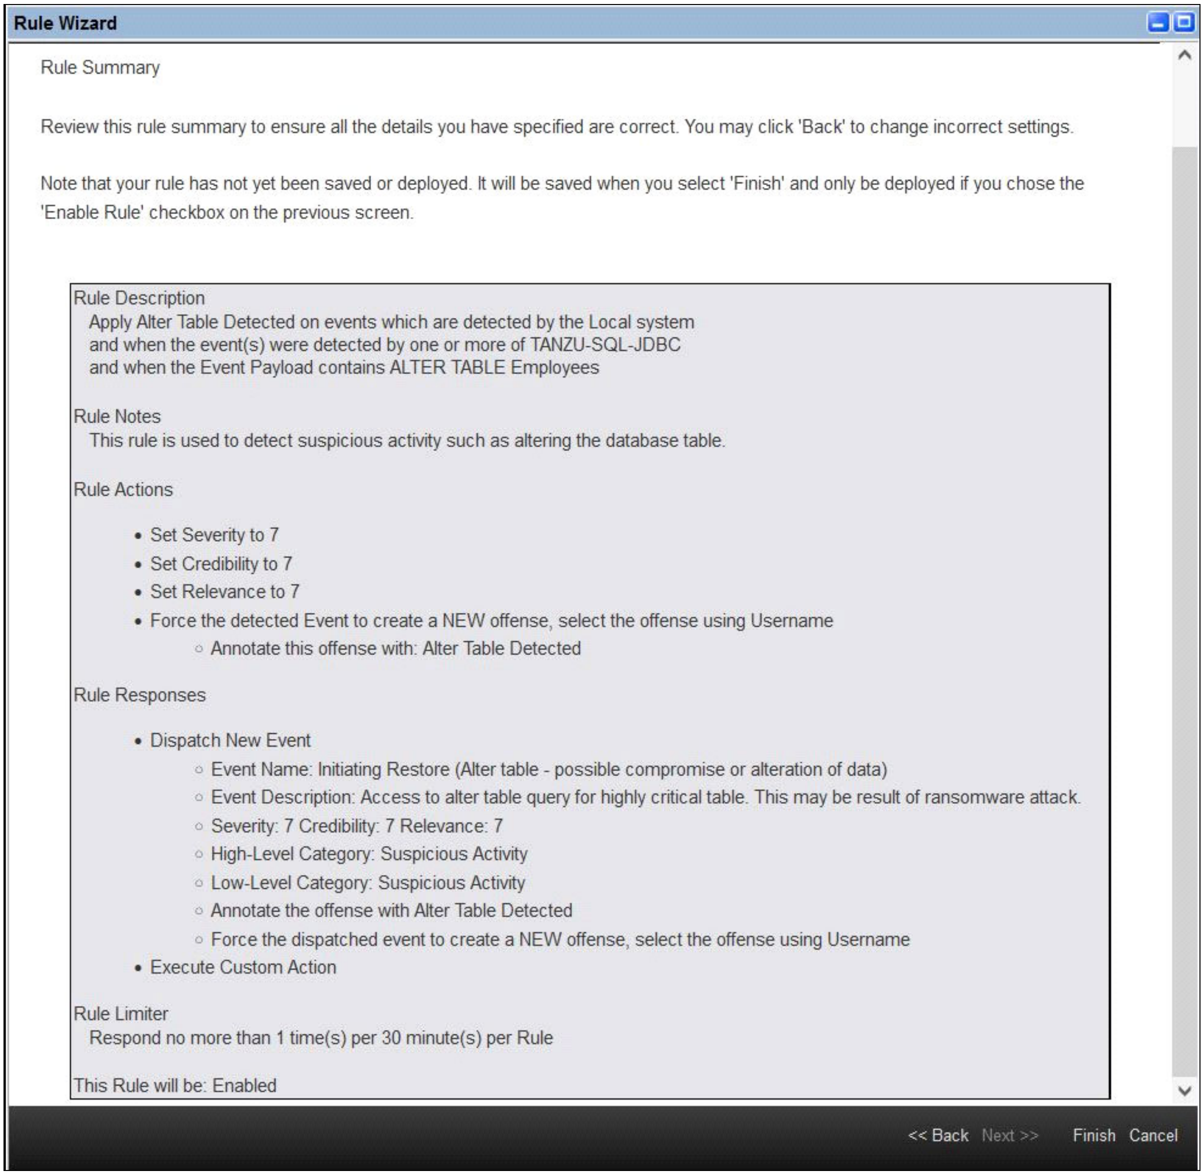

Rule Summary

Review this rule summary to ensure all the details you have specified are correct. You may click 'Back' to change incorrect settings.

Note that your rule has not yet been saved or deployed. It will be saved when you select 'Finish' and only be deployed if you chose the 'Enable Rule' checkbox on the previous screen.

Rule Description
Apply Alter Table Detected on events which are detected by the Local system
and when the event(s) were detected by one or more of TANZU-SQL-JDBC
and when the Event Payload contains ALTER TABLE Employees

Rule Notes
This rule is used to detect suspicious activity such as altering the database table.

Rule Actions

- Set Severity to 7
- Set Credibility to 7
- Set Relevance to 7
- Force the detected Event to create a NEW offense, select the offense using Username
  - Annotate this offense with: Alter Table Detected

Rule Responses

- Dispatch New Event
  - Event Name: Initiating Restore (Alter table - possible compromise or alteration of data)
  - Event Description: Access to alter table query for highly critical table. This may be result of ransomware attack.
  - Severity: 7 Credibility: 7 Relevance: 7
  - High-Level Category: Suspicious Activity
  - Low-Level Category: Suspicious Activity
  - Annotate the offense with Alter Table Detected
  - Force the dispatched event to create a NEW offense, select the offense using Username
- Execute Custom Action

Rule Limiter
Respond no more than 1 time(s) per 30 minute(s) per Rule

This Rule will be: Enabled

Figure 29 Alter Table Rule Summary window

Brute force attack on a database

In this section, we discuss the issue of compromised “sa” user credentials. Attackers can take advantage of the sa user authority to delete critical databases.

In this solution, a brute force drop database attack was generated from `mssql-cli` by using remote login.

The GitHub repository shows a sample Python script that is registered as part of the custom user action. The script makes API calls to IBM Spectrum Copy Data Management to run the predefined Scheduled Job with restore action. For more information, see “Resources” on page 33.

The brute force login case that is described here represents threat detection from the database environment by using the database audit logs and forwarding them to IBM QRadar as events. IBM QRadar’s rules engine is used to detect threats to the database and run custom user actions as response.

These events can be categorized (and threat detection rules can be defined) based on the security compliance matrix that is defined by the organization.

Consider the following points about this use case:

- In this solution, `Tanzu-jump` is the source server hosting critical SQL database `IT_Company` (see Figure 30).

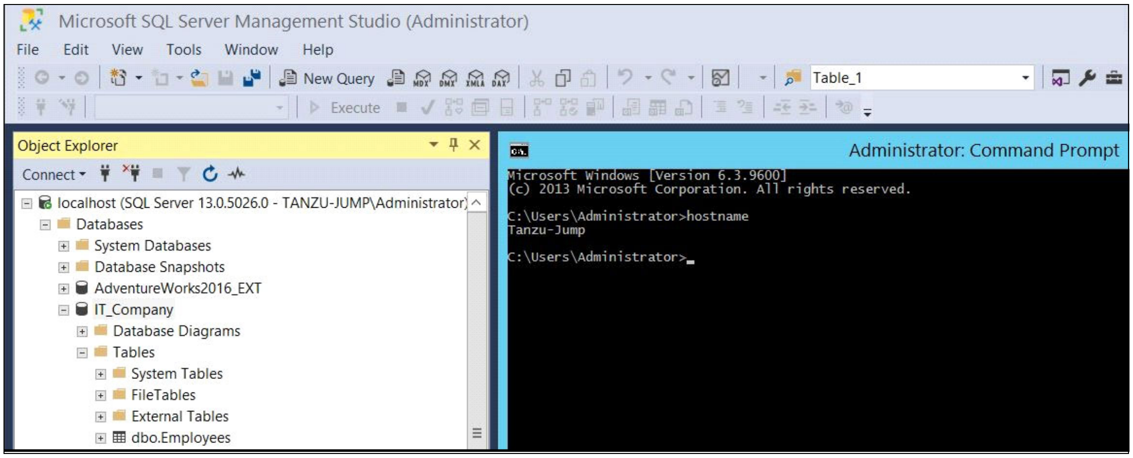

```
Microsoft Windows [Version 6.3.9600]
(c) 2013 Microsoft Corporation. All rights reserved.

C:\Users\Administrator>hostname
Tanzu-Jump

C:\Users\Administrator>_
```

Figure 30 Source server that hosts IT_Company SQL database

- The target server is `cdm-sql-target` with SQL server software installed on it (see Figure 31).

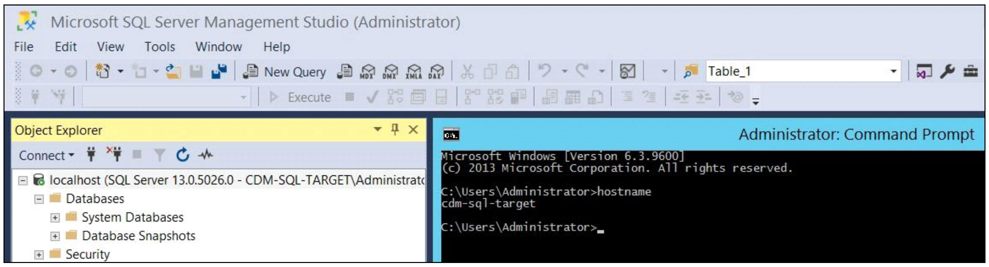
Microsoft SQL Server Management Studio (Administrator)
File
Edit
View
Tools
Window
Help
New Query
Table_1
Execute
Object Explorer
Connect
localhost (SQL Server 13.0.5026.0 - CDM-SQL-TARGET\Administrato
Databases
System Databases
Database Snapshots
Security
Administrator: Command Prompt

```
Microsoft Windows [Version 6.3.9600]
(c) 2013 Microsoft Corporation. All rights reserved.

C:\Users\Administrator>hostname
cdm-sql-target

C:\Users\Administrator>_
```

Figure 31 Target server with SQL server software installed

- Drop database custom script is fired from remote Linux server through `mssql-cli` to drop critical database `IT_Company` (see Figure 32).

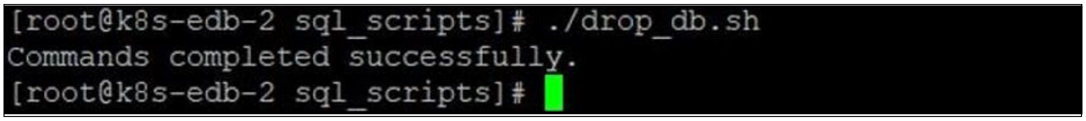

```
[root@k8s-edb-2 sql_scripts]# ./drop_db.sh
Commands completed successfully.
[root@k8s-edb-2 sql_scripts]#
```

Figure 32 Drop database custom script from mmsql-cli

- SQL audit captures the audit logs and forwards them to IBM QRadar in the form of events. Drop database event is captured on IBM QRadar (see Figure 33).

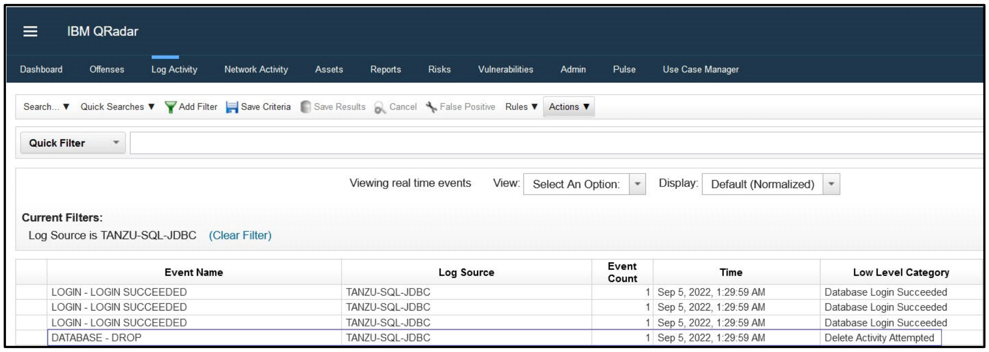

Viewing real time events View: Select An Option: Display: Default (Normalized)

Current Filters:
Log Source is TANZU-SQL-JDBC (Clear Filter)

Event Name	Log Source	Event Count	Time	Low Level Category
LOGIN - LOGIN SUCCEEDED	TANZU-SQL-JDBC	1	Sep 5, 2022, 1:29:59 AM	Database Login Succeeded
LOGIN - LOGIN SUCCEEDED	TANZU-SQL-JDBC	1	Sep 5, 2022, 1:29:59 AM	Database Login Succeeded
LOGIN - LOGIN SUCCEEDED	TANZU-SQL-JDBC	1	Sep 5, 2022, 1:29:59 AM	Database Login Succeeded
DATABASE - DROP	TANZU-SQL-JDBC	1	Sep 5, 2022, 1:29:59 AM	Delete Activity Attempted

Figure 33 Captured IBM QRadar drop database event

- The event response is generated according to the custom event rule, which is defined in the Defining Custom Action in IBM QRadar section. This response triggers the restore job from SCDM through the custom Python script (see Figure 34).

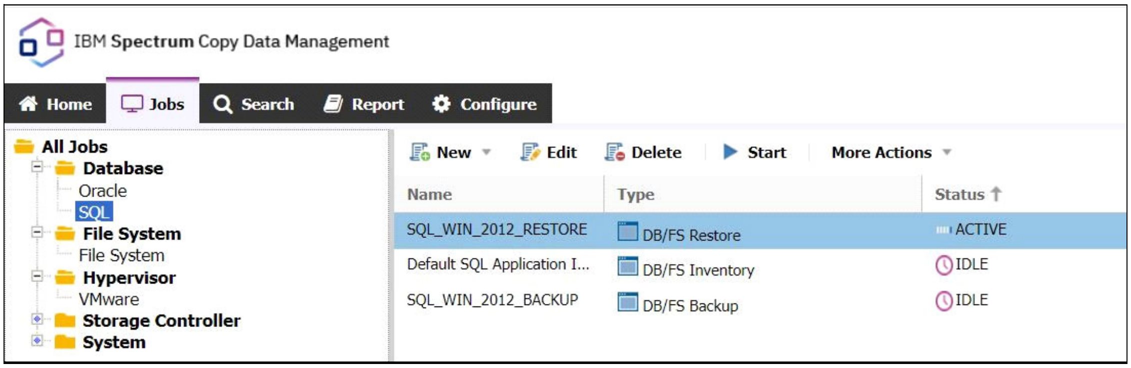

Name	Type	Status
SQL_WIN_2012_RESTORE	DB/FS Restore	ACTIVE
Default SQL Application I...	DB/FS Inventory	IDLE
SQL_WIN_2012_BACKUP	DB/FS Backup	IDLE

Figure 34 Triggered SCDM restore job

- After the restore job is completed, the `IT_Company` critical database is mounted from latest SGC backup on the clean space environment (in our case, the target server), `cdm-sql-target`, as shown Figure 35.

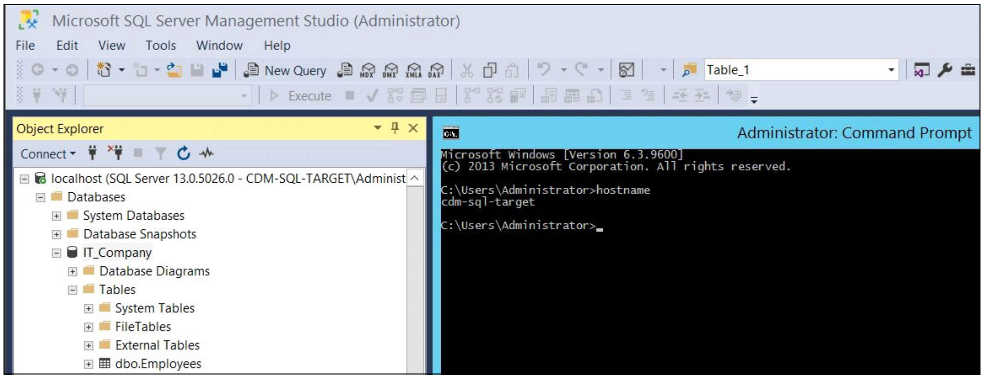

```
Microsoft Windows [Version 6.3.9600]
(c) 2013 Microsoft Corporation. All rights reserved.

C:\Users\Administrator>hostname
cdm-sql-target

C:\Users\Administrator>_
```

Figure 35 Critical database mounted on clean space environment

Now, a good copy of the critical database is available in a clean space environment for user data validation and further investigation.

Ransomware attack simulation

In this section, we discuss the issue of ransomware attacks. Attackers can take advantage of the sa user authority to encrypt the critical data from database table.

In this solution, a table from critical database was encrypted from `mssql-cli` to simulate a ransomware attack.

The GitHub repository shows a sample Python script that is registered as part of the custom user action. The script makes API calls to IBM Spectrum Copy Data Management to run the predefined Scheduled Job with restore action, as described in “Configuring an SQL Restore Job” on page 16.

The ransomware attack simulation case that is described here represents threat detection from the database environment by using the database audit logs and forwarding them to IBM QRadar as events. IBM QRadar’s rules engine is used to detect threats to the database and run custom user actions as a response.

These events can be categorized, and threat detection rules can be defined, based on the security compliance matrix that is defined by the organization.

Consider the following points about this use case:

1. In this solution, `Employees` is the table available on critical SQL database `IT_Company` (see Figure 36).

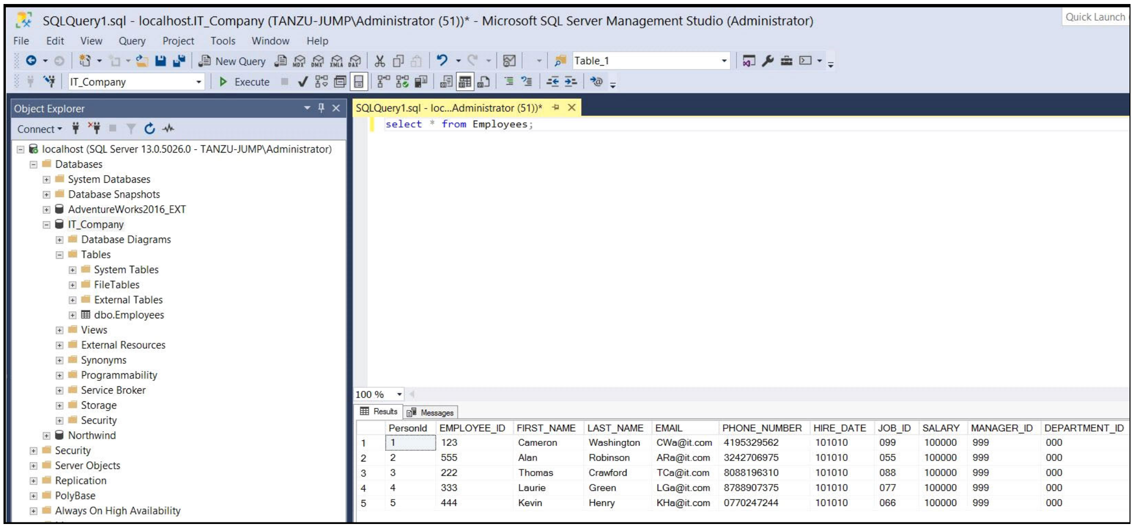

```
select * from Employees;
```

	PersonId	EMPLOYEE_ID	FIRST_NAME	LAST_NAME	EMAIL	PHONE_NUMBER	HIRE_DATE	JOB_ID	SALARY	MANAGER_ID	DEPARTMENT_ID
1	1	123	Cameron	Washington	CWa@it.com	4195329562	101010	099	100000	999	000
2	2	555	Alan	Robinson	ARa@it.com	3242706975	101010	055	100000	999	000
3	3	222	Thomas	Crawford	TCa@it.com	8088196310	101010	088	100000	999	000
4	4	333	Laurie	Green	LGa@it.com	8788907375	101010	077	100000	999	000
5	5	444	Kevin	Henry	KHa@it.com	0770247244	101010	066	100000	999	000

SQLQuery1.sql - localhost.IT_Company (TANZU-JUMP\Administrator (51))* - Microsoft SQL Server Management Studio (Administrator)
Object Explorer
localhost (SQL Server 13.0.5026.0 - TANZU-JUMP\Administrator)
IT_Company
dbo.Employees

Figure 36 Table Employees that is available on the IT_Company critical database

- Encrypt Column custom script is fired from remote Linux server through `mmsql-cli` to encrypt table `Employees` (see Figure 37).

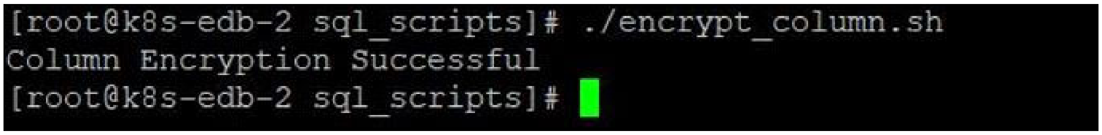

```
[root@k8s-edb-2 sql_scripts]# ./encrypt_column.sh
Column Encryption Successful
[root@k8s-edb-2 sql_scripts]#
```

Figure 37 Alter Table custom script

- Table `Employees` is now encrypted and the critical data that is on the table is in indecipherable format (see Figure 38).

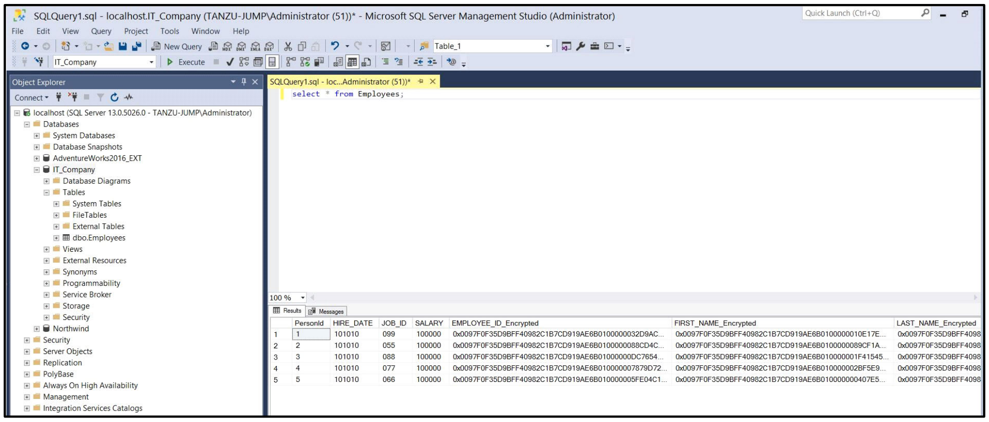

```
select * from Employees;
```

	PersonId	HIRE_DATE	JOB_ID	SALARY	EMPLOYEE_ID_Encrypted	FIRST_NAME_Encrypted	LAST_NAME_Encrypted
1	1	101010	099	100000	0x0097F0F35D9BFF40982C1B7CD919AE6B0100000032D9AC...	0x0097F0F35D9BFF40982C1B7CD919AE6B0100000010E17E...	0x0097F0F35D9BFF4098
2	2	101010	055	100000	0x0097F0F35D9BFF40982C1B7CD919AE6B0100000088CD4C...	0x0097F0F35D9BFF40982C1B7CD919AE6B0100000089CF1A...	0x0097F0F35D9BFF4098
3	3	101010	088	100000	0x0097F0F35D9BFF40982C1B7CD919AE6B01000000DC7654...	0x0097F0F35D9BFF40982C1B7CD919AE6B010000001F41545...	0x0097F0F35D9BFF4098
4	4	101010	077	100000	0x0097F0F35D9BFF40982C1B7CD919AE6B010000007879D72...	0x0097F0F35D9BFF40982C1B7CD919AE6B010000002BF5E9...	0x0097F0F35D9BFF4098
5	5	101010	066	100000	0x0097F0F35D9BFF40982C1B7CD919AE6B010000005FE04C1...	0x0097F0F35D9BFF40982C1B7CD919AE6B01000000407E5...	0x0097F0F35D9BFF4098

SQLQuery1.sql - localhost.IT_Company (TANZU-JUMP\Administrator (51))* - Microsoft SQL Server Management Studio (Administrator)
Object Explorer
localhost (SQL Server 13.0.5026.0 - TANZU-JUMP\Administrator)
IT_Company
dbo.Employees

Figure 38 Alter table custom script

- SQL audit captures the audit logs and forwards them to IBM QRadar in the form of events. The Alter Table event is captured on IBM QRadar (see Figure 39).

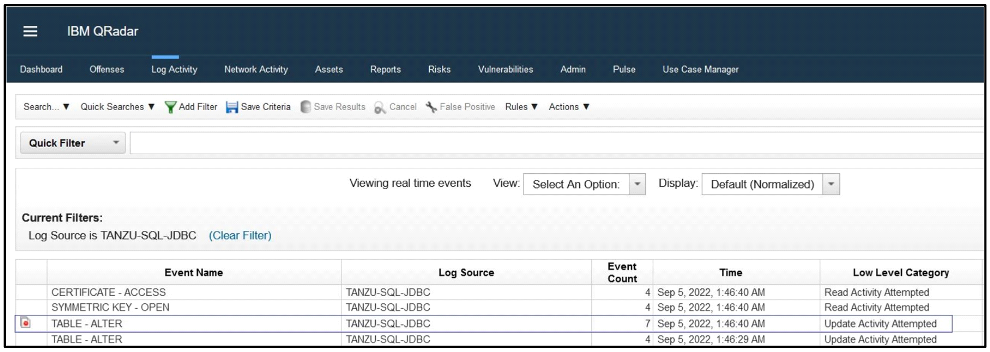

Viewing real time events View: Select An Option: Display: Default (Normalized)

Current Filters:
Log Source is TANZU-SQL-JDBC (Clear Filter)

Event Name	Log Source	Event Count	Time	Low Level Category
CERTIFICATE - ACCESS	TANZU-SQL-JDBC	4	Sep 5, 2022, 1:46:40 AM	Read Activity Attempted
SYMMETRIC KEY - OPEN	TANZU-SQL-JDBC	4	Sep 5, 2022, 1:46:40 AM	Read Activity Attempted
TABLE - ALTER	TANZU-SQL-JDBC	7	Sep 5, 2022, 1:46:40 AM	Update Activity Attempted
TABLE - ALTER	TANZU-SQL-JDBC	4	Sep 5, 2022, 1:46:29 AM	Update Activity Attempted

Figure 39   Captured IBM QRadar alter table event

- An event response is generated according to the custom event rule, which is defined in the Defining Custom Action in IBM QRadar section. This response triggers the restore job from SCDM through the custom Python script (see Figure 40).

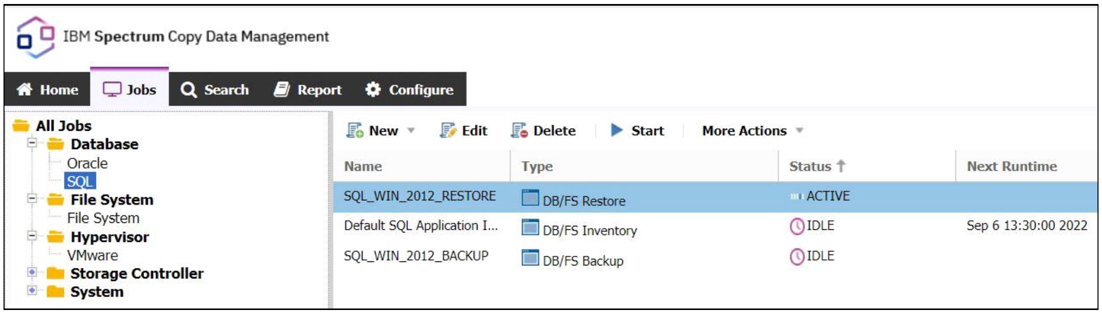

Name	Type	Status	Next Runtime
SQL_WIN_2012_RESTORE	DB/FS Restore	ACTIVE	
Default SQL Application I...	DB/FS Inventory	IDLE	Sep 6 13:30:00 2022
SQL_WIN_2012_BACKUP	DB/FS Backup	IDLE	

Figure 40   Triggered SCDM restore job

- After the restore job is completed, the critical database `IT_Company` is mounted from the latest SGC backup on the clean space environment; in our case, the target server `cdm-sql-target` (see Figure 41).

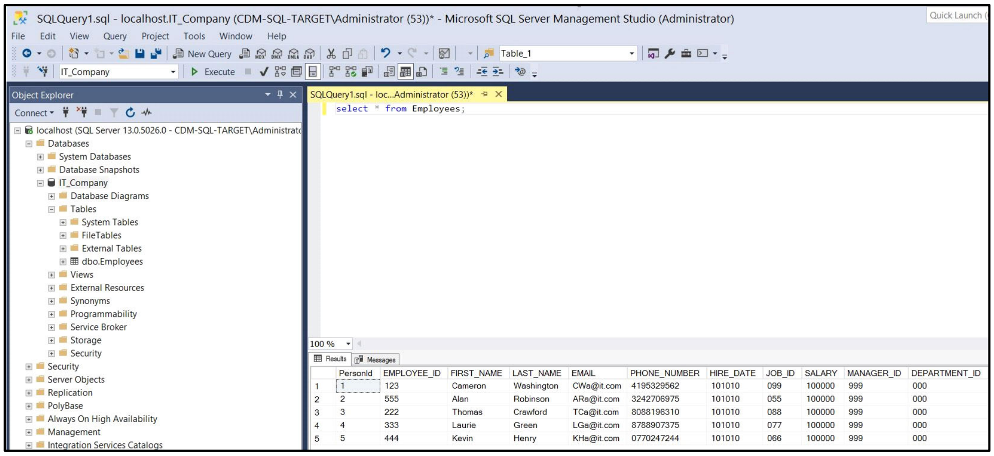

```
select * from Employees;
```

	PersonId	EMPLOYEE_ID	FIRST_NAME	LAST_NAME	EMAIL	PHONE_NUMBER	HIRE_DATE	JOB_ID	SALARY	MANAGER_ID	DEPARTMENT_ID
1	1	123	Cameron	Washington	CWa@it.com	4195329562	101010	099	100000	999	000
2	2	555	Alan	Robinson	ARa@it.com	3242706975	101010	055	100000	999	000
3	3	222	Thomas	Crawford	TCa@it.com	8088196310	101010	088	100000	999	000
4	4	333	Laurie	Green	LGa@it.com	8788907375	101010	077	100000	999	000
5	5	444	Kevin	Henry	KHa@it.com	0770247244	101010	066	100000	999	000

Figure 41   Critical database mounted on clean space environment

Now, a good copy of the critical database is available in a clean space environment. The table `Employees` is in an unencrypted format for user data validation and further investigation.

Summary

The solution that is described in this IBM blueprint publication shows the integration of IBM FlashSystem, IBM Spectrum Copy Data Management, and IBM QRadar to perform early threat detection on databases.

When a threat is detected, a cyber resiliency workflow is triggered. This workflow is used to run a predefined scheduled job in IBM Spectrum Copy Data Management to perform the required actions. Then, these actions start the restore job to mount the latest database backup on clean space environment.

The solution that is presented here can be used as template to categorize the events that are received from the database host.

Based on the events that are received, threat detection rules can be defined that confirm to security standards that are defined by organization’s compliance matrix.

The sample Python script shows how to use the API interface of IBM Spectrum Copy Data Management to perform specific tasks.

Authors

This blueprint guide was produced by a team of specialists from around the world working at IBM Redbooks, Poughkeepsie Center.

Tejas Sapkar a Storage Solutions Architect with IBM Systems, ISDL Lab Pune, India. He has over 12 years of experience with IBM storage products, such as IBM Spectrum Virtualize, IBM FlashSystems, IBM Tape Libraries, and IBM (TSM/Spectrum Protect). Currently, he is working on demonstrating cyber resilience solutions with IBM QRadar and IBM Storage Systems. Before joining IBM, Tejas worked with many storage vendors, such as EMC, Pure, and NetApp to deliver key projects to valued customers.

Shashank Shingornikar is a Storage Solutions Architect with IBM Systems, ISDL Lab Pune, India, for over 12 years. He has worked extensively with IBM Storage products, such as IBM Spectrum Virtualize, IBM FlashSystems, and IBM Spectrum Scale, building solutions that combine Oracle and Red Hat OpenShift features. Currently, he is working on demonstrating cyber resilience solutions with IBM QRadar and IBM Storage Systems. Before joining IBM, Shashank worked in The Netherlands on various high availability, disaster recovery, cluster, and replication solutions for database technologies, such as Oracle, MSSQL, and MySQL.

The authors want to thank **Pepe Lam**, of Spectrum Copy Data Management, for his valuable contributions and support of this project.

Resources

For more information about the topic that is discussed in this blueprint, see the following resources:

- Enhanced Cyber Resilience Threat Detection IBM Redpaper:

  https://www.redbooks.ibm.com/abstracts/redp5655.html?Open

- Proactive Early Threat Detection IBM Redpaper:

  https://www.redbooks.ibm.com/abstracts/redp5686.html

- IBM QRadar:

  https://www.ibm.com/docs/en/qsip

- IBM Spectrum Copy Data Management:

  https://www.ibm.com/in-en/products/spectrum-copy-data-management

- IBM Spectrum Copy Data Management User’s Guide:

  https://www.ibm.com/docs/en/SS57AN_2.2.16/pdf/b_cdm_guide.pdf

- Downloadable GitHub script:

  https://github.com/IBM/ibm-qradar-cdm-fs-safeguarded-copy

- Configuration for SQL IBM QRadar audit specification:

  https://www.ibm.com/docs/en/dsm?topic=server-microsoft-sql-preparation-communication-qradar

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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, MD-NC119, Armonk, NY 10504-1785, US

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.

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.

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.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, 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 http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

Redbooks (logo) ®
FlashCopy®
IBM®
IBM FlashSystem®
IBM Spectrum®
QRadar®

The following terms are trademarks of other companies:

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

OpenShift, Red Hat, are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

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.

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.

Privacy policy considerations

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.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

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 in 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.

October 2022

ISBN 0738460850

REDP-5691-00